METHOD ARTICLE

REVISED Effective immunity and second waves: a dynamic causal modelling study [version 2; peer review: 2 approved]

Karl J. Friston ¹, Thomas Parr ¹, Peter Zeidman¹, Adeel Razi ^{1,2}, Guillaume Flandin¹, Jean Daunizeau³, Oliver J. Hulme^{4,5}, Alexander J. Billig ⁶, Vladimir Litvak¹, Cathy J. Price¹, Rosalyn J. Moran ⁷, Anthony Costello⁸, Deenan Pillay⁹, Christian Lambert ¹

¹The Wellcome Centre for Human Neuroimaging, University College London, London, UK

²Turner Institute for Brain and Mental Health, Monash University, Clayton, VIC, 3800, Australia

³Institut du Cerveau et de la Moelle épinière, INSERM UMRS 1127, Paris, France

⁴Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital, Hvidovre, Denmark

⁵London Mathematical Laboratory, Hammersmith, London, UK

⁶UCL Ear Institute, University College London, London, UK

⁷Centre for Neuroimaging Science, Department of Neuroimaging, IoPPN, King's College London, London, UK

⁸UCL Institute for Global Health, Institute of Child Health, University College London, London, UK

⁹UCL Division of Infection and Immunity, University College London, London, UK

V2 First published: 01 Sep 2020, 5:204
<https://doi.org/10.12688/wellcomeopenres.16253.1>
 Latest published: 30 Sep 2020, 5:204
<https://doi.org/10.12688/wellcomeopenres.16253.2>

Abstract

This technical report addresses a pressing issue in the trajectory of the coronavirus outbreak; namely, the rate at which effective immunity is lost following the first wave of the pandemic. This is a crucial epidemiological parameter that speaks to both the consequences of relaxing lockdown and the propensity for a second wave of infections. Using a dynamic causal model of reported cases and deaths from multiple countries, we evaluated the evidence models of progressively longer periods of immunity. The results speak to an effective population immunity of about three months that, under the model, defers any second wave for approximately six months in most countries. This may have implications for the window of opportunity for tracking and tracing, as well as for developing vaccination programmes, and other therapeutic interventions.

Keywords

coronavirus, epidemiology, compartmental models, dynamic causal modelling, variational, Bayesian

Open Peer Review

Approval Status

	1	2
version 2 (revision) 30 Sep 2020		 view
version 1 01 Sep 2020	 view	 view

1. **Michael Stern**, National Institutes of Health, Baltimore, USA

2. **Anotida Madzvamuse** , University of Sussex, Brighton, UK

Any reports and responses or comments on the article can be found at the end of the article.

This article is included in the [Coronavirus \(COVID-19\)](#) collection.

Corresponding author: Adeel Razi (adeel.razi@monash.edu)

Author roles: **Friston KJ:** Conceptualization, Data Curation, Formal Analysis, Funding Acquisition, Methodology, Software, Validation, Visualization, Writing – Original Draft Preparation; **Parr T:** Methodology, Writing – Review & Editing; **Zeidman P:** Methodology, Writing – Review & Editing; **Razi A:** Methodology, Validation, Writing – Review & Editing; **Flandin G:** Methodology, Software, Validation, Writing – Review & Editing; **Daunizeau J:** Methodology, Writing – Review & Editing; **Hulme OJ:** Methodology, Writing – Review & Editing; **Billig AJ:** Methodology, Writing – Review & Editing; **Litvak V:** Methodology, Writing – Review & Editing; **Price CJ:** Writing – Review & Editing; **Moran RJ:** Methodology, Writing – Review & Editing; **Costello A:** Writing – Review & Editing; **Pillay D:** Writing – Review & Editing; **Lambert C:** Methodology, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: The Wellcome Centre for Human Neuroimaging is supported by core funding from Wellcome [203147]. A.R. is funded by the Australian Research Council (Refs: DE170100128 and DP200100757). A.J.B. is supported by a Wellcome Trust grant WT091681MA. CL is supported by an MRC Clinician Scientist award (MR/R006504/1).

The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2020 Friston KJ *et al.* This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Friston KJ, Parr T, Zeidman P *et al.* **Effective immunity and second waves: a dynamic causal modelling study [version 2; peer review: 2 approved]** Wellcome Open Research 2020, 5:204 <https://doi.org/10.12688/wellcomeopenres.16253.2>

First published: 01 Sep 2020, 5:204 <https://doi.org/10.12688/wellcomeopenres.16253.1>

REVISED Amendments from Version 1

In the revised version we have made following changes.

1. We have added a new Discussion section (which also absorbed Conclusions in the previous draft). Further description of the generative model (Friston *et al.*, 2020a) and Bayesian model comparison procedures is provided in this section.
2. We have included text in the Discussion section to delimit the scope and target of this report and also added study limitations and future work.
3. We have added mathematical expressions and description of how effective reproduction number is calculated in the Results section.

Any further responses from the reviewers can be found at the end of the article

Background

Over the past months, an alternative to standard epidemiological modelling has been considered in the form of dynamic causal modelling (Friston *et al.*, 2020a). This approach inherits from statistical physics and variational procedures in Bayesian modelling and machine learning (Dauwels, 2007; Feynman, 1972; Friston *et al.*, 2007; MacKay, 1995; MacKay, 2003; Winn & Bishop, 2005). The validity of this approach has been partly established in a series of reports looking at the role of population immunity within an outbreak in a single region (Friston *et al.*, 2020a), the effect of population fluxes between multiple regions (in the United States of America) (Friston *et al.*, 2020b) and the genesis of rebounds following lockdown, in relation to a second wave of infections (Friston *et al.*, 2020c). In brief, the conclusions of this kind of modelling are: (i) population immunity—inherited from the initial phases of the pandemic—plays a key role in nuancing its subsequent progression (ii) in the context of population exchange between regional outbreaks, social distancing and lockdown strategies based upon the local prevalence of infection reduce morbidity and mortality. Finally (iii), the mechanism that underwrites a second wave depends sensitively on the rate at which population immunity is lost following the first wave. This affords a window of opportunity within which track and trace protocols may delay or defer any second wave until it can be rendered innocuous through vaccination or clinical advances (Aleta *et al.*, 2020; Chinazzi *et al.*, 2020; Hellewell *et al.*, 2020; Keeling *et al.*, 2020; Kissler *et al.*, 2020; Kucharski *et al.*, 2020; Simonsen *et al.*, 2018; Streeck *et al.*, 2020; Wu *et al.*, 2020). In this report, we consider a key question: how long is this window—or, equivalently, what is the period of effective immunity inherited at the population level from the first wave (Kissler *et al.*, 2020).

Dynamic causal modelling¹ can be characterised as a generalisation of state-space modelling based upon differential equations. This contrasts with advanced descriptive approaches that fit curves to timeseries data, without any explicit reference to the underlying dynamics: e.g., (Tsallis & Tirnakli, 2020). Dynamic causal modelling differs from conventional epidemiological modelling

in that it uses mean field approximations and variational procedures to model the evolution of probability densities—in a way that is similar to quantum mechanics and statistical physics (Greenland, 2006). This contrasts with epidemiological modelling that uses stochastic realisations of epidemiological dynamics to approximate probability densities with sample densities (Kermack & McKendrick, 1997; Rhodes & Hollingsworth, 2009; Vineis & Kriebel, 2006; White *et al.*, 2007). One advantage of variational procedures is that they are orders of magnitude more efficient; enabling model inversion or fitting within minutes (on a laptop) as opposed to hours or days on a supercomputer (Rhodes & Hollingsworth, 2009). More importantly, variational procedures provide an efficient way of assessing the quality of one model relative to another, in terms of model evidence (a.k.a., marginal likelihood) (Penny, 2012). This enables one to compare different models using Bayesian model comparison (a.k.a. structure learning) and use the best model for nowcasting, forecasting or, indeed, test competing hypotheses about viral transmission.

We have used this technology to build epidemiological models of how data are generated—in terms of latent causes like the prevalence of infection—that embed conventional epidemiological models (e.g., susceptible, exposed, infected, recovered (SEIR) models) in an extended state space. For example, dynamic causal modelling allows certain probability densities to be factorized. A key example of this is to model a joint distribution over states of infection and clinical manifestation. In other words, instead of assuming that there is a difference between being infected (I) and having recovered (R), one can accommodate the fact that it is possible to express symptoms without being infected: e.g., a secondary bacterial infection following interstitial pneumonia (Huang *et al.*, 2020). Conversely, one can be infected without showing symptoms. Crucially, dynamic causal models can be extended to generate any kind of data at hand: for example, the number of positive tests. This requires careful consideration of how positive tests are generated, by modelling latent variables such as the bias towards testing people with or without infection or, indeed, the capacity for testing, which may itself be time-dependent. In short, everything that might matter—in terms of the latent (hidden) causes of the data—can be installed in the model, including social distancing, self-isolation and other processes that underwrite transmission. When all such latent causes are included, model comparison can then be performed to assess whether they are needed to explain the data. Here, we leverage the efficiency of dynamic causal modelling to evaluate the evidence for a series of models that are distinguished by the rate at which effective immunity to SARS-CoV-2 is lost. This provides a probability distribution over the rate of loss that determines when, or if, a second wave will ensue (Friston *et al.*, 2020c; Kissler *et al.*, 2020). In what follows, effective population immunity refers to the proportion of people who cannot contract or transmit the virus. This means that the loss of effective immunity can be mediated in several ways. For example, a decline in antibody levels, viral mutation, or a dilution of population immunity due to population fluxes. These are all important mechanistic hypotheses that can, in principle, be addressed using Bayesian model comparison. To do this, it would be necessary to parameterise the model in a way that allowed one to withdraw one or other mechanism and

¹ http://www.scholarpedia.org/article/Dynamic_causal_modeling

evaluate whether the model evidence increased or decreased. An example of this can be found in (Friston *et al.*, 2020a), where the relative contribution of lockdown and population immunity to prevalence and mortality was evaluated. Interestingly, both lockdown and herd immunity were necessary to explain the data: in the sense that removing either mechanism substantially reduced model evidence. Please see (Friston *et al.*, 2020a) for details.

Details about the dynamic causal model can be found in the above technical reports (Friston *et al.*, 2020a; Friston *et al.*, 2020b;

Friston *et al.*, 2020c). Please see Figure 1 and Table 1 for a summary of its form and parameters. The model was fitted to new cases and deaths using data available from Johns Hopkins University². The inversion and subsequent model comparison used standard variational (Laplace) procedures (Friston *et al.*, 2007; Marreiros *et al.*, 2009), as implemented in academic (open source) software³. The particular model used here has a degree

² <https://github.com/CSSEGISandData/COVID-19>.

³ <https://www.fil.ion.ucl.ac.uk/spm/covid-19/>

Figure 1. The LIST model. This schematic summarises the LIST (*location, infection, symptom, and testing*) generative model used for the following analyses. This model is formally identical to that described in (Friston *et al.*, 2020c). It includes a state (*isolation*) to model people who are self-isolating because they think they may be infectious (within their home or elsewhere). It also includes another (*resistant*) state to model individuals who are shielded or have pre-existing immunity, e.g., via cross-reactivity (Grifoni *et al.*, 2020; Ng *et al.*, 2020) or other protective host factors (Bunyavanich *et al.*, 2020; Zheng *et al.*, 2020). This absorbing state also plays the role of the *recovered* or *removed* states of SEIR models, namely, once entered, people stay in the state for the duration of the outbreak. One can leave any of the remaining states. For example, one only occupies the state of testing *positive* or *negative* for a day, and then moves to the *untested* state the following day. This ensures that the total population is conserved; e.g., deaths are offset by births into the susceptible state. Furthermore, it enables the occupancy of various states to be interpreted in terms of the rate of daily expression. The blue discs represent the four factors of the model, and the segments of these discs correspond to their states (i.e., compartments). The states within any factor are mutually exclusive, whereas the factors embody the factorial form of the compartmental model. In other words, every individual has to be in one of the states associated with the four factors or attributes. The orange boxes represent the observable outputs that are generated by this dynamic causal model, in this instance, daily reports of positive tests and deaths. The rate of transition between states—or the dwell time within any state—rests upon the model parameters that, in many instances, can be specified with fairly precise prior densities. These are listed in Table 1.

Table 1. Parameters of the epidemic (LIST) model and priors, $N(\eta, C)$. NB: prior means are for scale parameters $\theta = \exp(\vartheta)$.

Number	Parameter	Mean	Variance	Description
1	θ_n	4	1	Number of initial cases
2	θ_r	1/2	1/256	Proportion of non-susceptible cases
3	θ_N	8	1	Effective population size (millions)
Location				
4	θ_{out}	1/3	1/256	Probability of going out
5	θ_{sde}	1/32	1/256	Social distancing threshold
6	θ_{cap}	16/100000	1/256	CCU capacity threshold (per capita)
Infection				
7	θ_{res}	1/2	1/256	Proportion of non-contagious cases
8	θ_{Rin}	4	1/16	Effective number of contacts: home
9	θ_{Rou}	48	1/16	Effective number of contacts: work
10	θ_{trn}	1/3	1/16	Transmission strength
11	$\theta_{inf} = \exp(-\frac{1}{\tau_{inf}})$	$\tau_{inf} = 4$	1/256	Infected period (days)
12	$\theta_{con} = \exp(-\frac{1}{\tau_{con}})$	$\tau_{con} = 4$	1/256	Infectious period (days)
13	$\theta_{imm} = \exp(-\frac{1}{\tau_{imm}})$	$\tau_{imm} = 1:32$	1/512	Period of immunity (months)
Symptoms				
14	$1 - \theta_{dev} = \exp(-\frac{1}{\tau_{inc}})$	$\tau_{inc} = 16$	1/256	Incubation period (days)
15	θ_{sev}	1/32	1/256	Probability of ARDS
16	$\theta_{sym} = \exp(-\frac{1}{\tau_{sym}})$	$\tau_{sym} = 8$	1/256	Symptomatic period (days)
17	$\theta_{rds} = \exp(-\frac{1}{\tau_{rds}})$	$\tau_{rds} = 10$	1/256	ARDS period (days)
18	θ_{fat}	1/2	1/256	ARDS fatality rate: CCU
19	θ_{sur}	1/8	1/256	ARDS fatality rate: home
Testing				
20	θ_{ttt}	1/10000	1	Efficacy of tracking and tracing
21	θ_{lat}	2	1	Latency of sustained testing (months)
22	θ_{sus}	4/10000	1/256	Sustained testing
23	θ_{bas}	4/10000	1/256	Baseline testing
24	θ_{tes}	8	1/16	Selectivity of testing infected people
25	$\theta_{del} = \exp(-\frac{1}{\tau_{del}})$	$\tau_{del} = 2$	1/256	Delay in reporting test results (days)

Secondary sources (Huang *et al.*, 2020; Kissler *et al.*, 2020; Mizumoto & Chowell, 2020; Russell *et al.*, 2020; Verity *et al.*, 2020; Wang *et al.*, 2020). These prior expectations should be read as the effective rates and time constants as they manifest in a real-world setting. For example, a six-day period of contagion is shorter than the period that someone might be infectious (Wölfel *et al.*, 2020)⁴, on the (prior) assumption that they will self-isolate, when they realise they could be contagious. The priors for the non-susceptible and non-contagious proportion of the population are based upon clinical and serological studies reported over the past few weeks; e.g., (Ing *et al.*, 2020; Stringhini *et al.*, 2020). Please see the code base for a detailed explanation of the role of these parameters in transition probabilities among states. Although the (scale) parameters are implemented as probabilities or rates, they are estimated as log parameters, denoted by $\vartheta = \ln \theta$.

⁴ Shedding of COVID-19 viral RNA from sputum can outlast the end of symptoms. Seroconversion occurs after 6–12 days but is not necessarily followed by a rapid decline of viral load. However, RNA shedding usually lasts longer than the shedding of infectious virus: many viruses produced are defective in some way, but still present RNA. As a rule of thumb 1/100 to 1/1000 virions produced are infective: van Kampen *et al.*, 2020. Shedding of infectious virus in hospitalized patients with coronavirus disease-2019 (COVID-19): duration and key determinants. medRxiv, 2020.2006.2008.20125310.

of face validity. Formal Bayesian model comparison—with the closest conventional epidemiological models—speak to a higher model evidence (Moran *et al.*, 2020), i.e., it provides a more accurate and parsimonious account of the data via optimising a (variational) bound on model evidence. Its predictive validity has been partly established. For example, it predicted death rates would peak on 10 April in the United Kingdom, with an initial relaxation of lockdown on 8 May 2020. In what follows, we use dynamic causal modelling to ask a simple but crucial question: how quickly will immunity to SARS-CoV-2 be lost at a population level?

Before addressing this question, we reiterate that this paper is a technical report illustrating how questions of this sort can be answered using variational Bayes and dynamic causal modelling. It explicitly does not purport to provide definitive answers. In other words, as the models are improved through Bayesian model comparison—or as more data become available—the inferences and posterior predictions below will change. Although these inferences are described definitively, they are entirely conditional upon the model used in this analysis, and the data available at the time of writing (8 June 2020).

Results

The dynamic causal model above was fitted (i.e., inverted) using timeseries data from Johns Hopkins University⁵ covering reported new cases and deaths from countries showing the highest cumulative number of deaths. The priors over the (25) model parameters can be found in Table 1. Crucially, this model inversion was repeated with different rates at which effective immunity is lost (i.e., the expected period of immunity following infection). These ranged from one month through to 32 months. This range was chosen to cover worst to best case scenarios. The worst-case scenario would correspond to a short-term period of immunity, less than that associated with the betacoronaviruses that cause the common cold: SARS-CoV-2 belongs to the betacoronavirus genus, which includes the SARS, MERS, and two other human coronaviruses, HCoV-OC43 and HCoV-HKU1 that cause the common cold (Kissler *et al.*, 2020; Su *et al.*, 2016). Immunity to HCoV-OC43 and HCoV-HKU1 appears to be lost over a few months. However, betacoronaviruses might induce immune responses against each another. For example, SARS can generate neutralizing antibodies against HCoV-OC43 that can endure for years, while HCoV-OC43 infection can generate cross-reactive antibodies against SARS (Chan *et al.*, 2013). A period of 32 months corresponds to a level of effective immunity for close to three years, comparable to SARS-CoV-1.

The dynamic causal model used in this analysis accommodates heterogeneity of susceptibility and transmission at three levels, including a non-contagious proportion of the population that stands in for people who cannot transmit the virus. This inclusion speaks to the increasing appreciation of how heterogeneity in

the population can have a fundamental effect on the epidemiological dynamics. This is variously described in terms of overdispersion, super spreaders, and amplification events (Endo *et al.*, 2020; Lloyd-Smith *et al.*, 2005; Paynter, 2016). In the current model, such heterogeneity was modelled in terms of three successive bipartitions (see Figure 2):

Heterogeneity in exposure: This was implicitly modelled in terms of an effective population size that is a subset of the total (census) population. The effective population is constituted by individuals who are in contact with contagious individuals. The remainder of the population are assumed to be geographically sequestered from a regional outbreak or are shielded from it. For example, if the population of the UK was 68 million, and the effective population was 39 million, then only 57% are considered to participate in the outbreak⁶. Of this effective population, a certain proportion are susceptible to infection:

Heterogeneity in susceptibility: This was modelled in terms of a portion of the effective population that are not susceptible to infection. For example, they may have pre-existing immunity, e.g., via cross-reactivity (Grifoni *et al.*, 2020; Ng *et al.*, 2020) or particular host factors (Bunyavanich *et al.*, 2020; Zheng *et al.*, 2020) such as mucosal immunity (Seo *et al.*, 2020). This non-susceptible proportion is assigned to the state of *resistance* at the start of the outbreak. Of the susceptible proportion of the effective population, a certain proportion can transmit the virus to others:

Heterogeneity in transmission: We modelled heterogeneity in transmission with a free parameter (with a prior of one half and a prior standard deviation of 1/16). This parameter corresponds to the proportion of susceptible people that cannot transmit the virus, i.e., those who move directly from a state of being *exposed* to a state of *resistance*, as opposed to moving from a state of being *infectious* to subsequent *immunity*. We associated this with a potentially mild illness—e.g., (Chau *et al.*, 2020)—that does not entail seroconversion, e.g., recovery in terms of T-cell mediated responses (Grifoni *et al.*, 2020; Zheng *et al.*, 2020). Note that this construction conflates transmission with the probability of developing symptoms, in that being infectious means you can transmit the virus but also increases the period during which you could move from a *healthy* state to a *symptomatic* state.

The *resistant* state therefore plays the role of an immune state for people who never become contagious, either because they are not susceptible to infection or become resistant after a mild illness. This model reconciles the apparent disparity between the relatively high morbidity/mortality rates and the relatively

⁵ Available from <https://github.com/CSSEGISandData/COVID-19>. These timeseries were smoothed with a Gaussian kernel to suppress spurious fluctuations at the weekends.

⁶ And of those 57%, some will be more exposed than others, conferring a further degree of heterogeneity, e.g., people working in care homes and hospitals, whose staff show a high seroprevalence: Houlihan *et al.*, 2020. SARS-CoV-2 virus and antibodies in front-line Health Care Workers in an acute hospital in London: preliminary results from a longitudinal study. medRxiv, 2020.2006.2008.20120584.

Figure 2. Heterogeneity of exposure, susceptibility, and transmission. Schematic illustrating the composition of a population in terms of people who are not exposed to contagious contact, not susceptible to contagion, susceptible but not contagious and, finally susceptible and contagious.

low seroprevalence observed empirically e.g., (Stringhini *et al.*, 2020)⁷. Bayesian model comparison confirmed that there was very strong evidence (Kass & Raftery, 1995) for all three types of heterogeneity (portrayed as ‘immunological dark matter’ in the media); namely, an effective population that is a subset of the census population, a susceptible population that is a subset of the effective population and a contagious population that is a subset of the susceptible population (c.f., a super spreaders). In this model, only susceptible individuals who become contagious develop antibodies to SARS-CoV-2, typically around 8% of the total population.

Crucially, we did not impose any prior constraints on the effective population size⁸. In other words, we treated the data from each country as reflecting an outbreak in a population of unknown size that comprised a mixture of susceptible and non-susceptible individuals, where susceptible individuals comprised a mixture contagious and non-contagious individuals. In this way,

we were able to model the self-evident dissociation between the total size of a population and the number of people affected in each country.

We specified a total of 32 models, each differing in their assumption about how long immunity would last, from 1 month to 32 months, in monthly increments. The log evidence for each of these 32 models was pooled over the 10 countries with the highest reported deaths (listed in Table 2). This evidence accumulation furnishes the marginal likelihood of each period of immunity (i.e., model) that—under uninformative priors over the period of immunity—corresponds to a posterior distribution, having marginalised out conditional uncertainty about all other parameters. Model inversion itself maximises the marginal likelihood that implicitly penalises overfitting, with respect to model complexity⁹. The resulting accuracy of the data fits are shown in Figure 3, in terms of cumulative death rates and new cases for the countries considered.

The accompanying distribution over the period of immunity is shown in Figure 4, suggesting that the expected period of immunity is about three months, with fairly precise 90%

⁷ See also: <https://www.gov.uk/government/publications/national-covid-19-surveillance-reports/sero-surveillance-of-covid-19>

⁸ The susceptible population can be defined operationally as that proportion of the population that will eventually succumb to infection, with consequent immunity or resistance that may or may not be lost over time.

⁹ Technically, the Kullback-Leibler divergence between posterior and the prior.

Bayesian credible intervals (less than the one month resolution of the model search). This does not mean that individuals will suddenly lose immunity after three months, rather that the effective population immunity will decline exponentially with a time constant of about three months. The ‘effective’ immunity refers to the fact that this characterisation of resilience is conditioned upon the model of aggregated or population dynamics. In other words, the effective population is behaving ‘as if’ its immunity is lost at this rate. There are many mechanisms that could contribute to this loss; for example, population fluxes could slowly increase the proportion of susceptible individuals (e.g., by relaxing lockdown); thereby diluting immunity acquired by the contagious proportion. Other viral and host factors (Beutler *et al.*, 2007; Su *et al.*, 2016) may clearly play a role (e.g., viral mutation or loss of antibodies)

The rate at which immunity is lost is important because it constrains the onset of any putative second wave. Figure 5 illustrates this in terms of three scenarios for the effective population in the United Kingdom: first, a worst-case scenario with rapid loss of immunity (over a period of one month), a most likely scenario based upon the posterior expectation from Figure 4 (left panel) and, finally, a best-case scenario with a period of immunity lasting for years (32 months)¹⁰. We see that a very short period of immunity effectively merges the second wave into the first to produce a protracted time course of fatality rates. In effect, (a quasi) endemic equilibrium is obtained quickly as people lose immunity and become re-infected. With an immune period of three months, a second wave can be anticipated shortly after Christmas, in the New Year. With enduring immunity

(here of 32 months) any second wave will be deferred for a year or more.

Interestingly, cumulative death rates appear to be higher with a three-month period of immunity, relative to a one-month period. This is because the predictions are posterior predictive densities, which are the most likely outcomes under the two periods of immunity. In other words, the best explanation for the data—under a one-month period of immunity—rests upon other model parameters that attenuate fatality rates, relative to a three-month period. Anecdotally, this kind of result suggests that we should be fairly confident about the loss of effective

Table 2. List of countries and the dates of predicted fatality peaks of first and second waves.

Country	First wave	Second wave
US	24-Apr-2020	02-Nov-2020
Brazil	12-Jun-2020	27-Aug-2021
UK	18-Apr-2020	09-Jan-2021
France	16-Apr-2020	28-Sep-2020
Spain	10-Apr-2020	03-Oct-2020
Italy	06-Apr-2020	13-Jan-2021
Mexico	12-Jun-2020	29-Aug-2021
Belgium	19-Apr-2020	26-Oct-2020
Germany	24-Apr-2020	14-Sep-2020
Canada	14-May-2020	14-Feb-2021

¹⁰ Note that the best and worst scenarios are not determined by the credible intervals of the posterior distribution, but are the limits of the scenarios considered *a priori*.

Figure 3. Model accuracy. This figure illustrates the accuracy of model inversion by plotting the empirical data for cumulative deaths (left panel) and cumulative new cases (right panel). The empirical data are shown as black dots overlaid on country specific predictions (coloured lines) based upon the latent states summarised in the subsequent figure. The trajectories have been shifted in time such that zero weeks corresponds to the time point at which the prevalence of the infection was estimated to be 0.1%.

Figure 4. Loss of immunity. The left-hand panel reports the posterior distribution over the period of immunity based upon the marginal likelihood of a series of models that assume a particular prior expectation (with a precise prior covariance of 1/512). Here, the log posterior has been plotted on log scale. The form of the posterior over this key parameter reflects the fact that the trajectory of new cases and deaths contains sufficient information to make fairly precise inferences about the rate at which immunity is lost. The right-hand panel shows the same results in terms of a cumulative distribution. The broken lines correspond to 90% Bayesian credible intervals.

Figure 5. Second waves. This figure reports the expected death rates as a function of time for the effective United Kingdom population. The three trajectories (blue lines) and accompanying 90% Bayesian credible intervals (shaded areas) correspond to posterior predictions with a loss of immunity over 1, 3 and 32 months. These represent the smallest, most likely and longest period of immunity considered in the Bayesian model comparison (summarised in the previous figure). The black dots correspond to empirical data, after smoothing with a four-day Gaussian kernel. The right panel reproduces the data in the left panel (now using different colours for the three trajectories) with a focus on the second wave peaking in January of next year.

immunity in a month when the predictions under short (one and three months) and long term (32 months) scenarios diverge. One would be hoping to see death rates fall to negligible levels by October. If they persist at 20 deaths per day, then one might anticipate a second wave in January.

A common metric of viral spread is the effective reproduction ratio (R_t).

The effective reproduction rate provides a useful statistic that reflects the exponential growth of the prevalence of infection. There are several ways in which it can be estimated. For our

purposes, we can evaluate an instantaneous reproduction rate directly from the time varying prevalence of infection as follows:

$$R_t = \exp(K_t \cdot \tau_{con})$$

$$K_t = \ln \frac{P(\text{infected}_{t+1})}{P(\text{infected}_t)} = \frac{\ln(2)}{T_d}$$

These expressions show that the reproduction rate reflects the growth of the (logarithm of the) proportion of people infected—and the period of being infectious. This is related to the doubling

time T_d . Note that the reproduction rate is not a parameter of the model: it is an outcome that is generated by the latent states inferred by inverting (i.e., fitting) the model to empirical timeseries.

Figure 6 uses the same format as the previous figure to show the effective reproduction ratio for the United Kingdom. The initial fall in the effective reproduction ratio is subtended by lockdown in the first instance, followed by an acquisition of population immunity in the effective population. After reaching a minimum of about 0.7, the effective reproduction ratio slowly increases with loss of population immunity to peak in the late autumn, portending a second wave infections in January. Following this, the reproduction ratio remains largely below one and slowly drifts back to one after a year.

Figure 7 illustrates the underlying or latent causes of the predicted fatality rates over 18 months for the most likely (three-month) loss of immunity. These are the hidden states that we can infer from the modelling. In this model, the latent states are factorized into various *locations*, different states of *infection*, *symptom* expression and the states that underwrite the generation of *test* results. Please see figure legend for details. It is evident from these posterior predictions that the UK might expect a second wave in about eight months (around January 2021). This is important because there is a window of opportunity in the next few months during which nonpharmacological interventions—especially, tracking and tracing—will, in principle, be in a position to defer or delay the second wave indefinitely (or until an effective treatment or vaccination programme is in place).

Please see (Friston *et al.*, 2020c) for a more detailed treatment. Note that this model includes a latent state of *immunity* that peaks around 11% and then falls gently as immunity is lost (yellow line in the infection panel). In contrast, the *resistant* proportion (purple line) slowly accumulates people who *recover* from a mild illness and are *removed* from the susceptible proportion of the effective population.

The predictions above are generated from the parameters of a single country. However, these predictions conceal a large amount of between-country variability due to the non-linear relationship between the model parameters and trajectories of latent causes and states. Figure 8 shows the equivalent predictions of fatality rates for all 10 countries, under the most likely period of immunity (three months). Note that there is considerable variation in the onset of the second wave due to country-specific differences in the underlying epidemiological parameters. Table 2 summarises these differences in terms of the predicted dates of the first and second waves, respectively. For countries like the United Kingdom, this analysis suggests that one can anticipate a second wave in early 2021, which is later than the prediction for Germany, which might experience a second wave in October of this year (2020).

This variation from country to country reflects differences in their epidemiological parameters. Table 3 summarises a few of these parameters and their variation. The first column lists the proportion of the effective population that are immune at the peak of population immunity. These range from 7 to 17% (4.5% to 5% of the total population), in line with current serological

Figure 6. Reproduction ratios. This figure uses the same format as the previous figure; here, showing the predictions of the effective reproduction ratio for the United Kingdom. The initial fall in the reproduction rate is subtended by lockdown in the first instance, followed by an acquisition of population immunity in the effective population.

Figure 7. Latent causes. This figure reproduces the predictions of the second wave for the United Kingdom in [Figure 4](#). Here, these outcomes are supplemented with the underlying latent causes or expected states in the lower four panels (the first state in each factor has been omitted for clarity: i.e., *home*, *susceptible*, *healthy*, and *untested*). These latent or expected states generate the observable outcomes in the upper two panels. The solid lines are colour-coded and correspond to the states of the four factors in [Figure 1](#). For example, under the *location* factor, the probability of being found at work declines steeply from about 25% to 3% at the onset of the outbreak. At this time, the probability of isolating oneself rises to about 3% during the peak of the pandemic. After about six weeks, the implicit lockdown starts to relax and slowly tails off, with accompanying falls in morbidity (in terms of symptoms) and mortality (in terms of death rate). As population immunity (yellow line in the *infection* panel) declines, the prevalence of infection accelerates to generate a second wave that peaks at about 50 weeks. Note that the amplitude of the second wave is much smaller than the first.

Figure 8. Second waves over countries. This figure illustrates the posterior predictive expectations of daily death rates for the countries considered at the time of writing (8 June 2020).

Table 3. List of countries and posterior estimates and population size. The non-contagious proportion is a percentage of the susceptible population, while the non-susceptible proportion is a percentage of the effective population. The numbers in brackets express the effective population as a percentage of the total (census) population.

Country	Population immunity (percent)	First peak deaths (per day)	Second peak deaths (per day)	Non-contagious (percent)	Non-susceptible (percent)	Effective population (millions)	Total population (millions)
US	15%	2254	244	48%	53%	127 (38%)	331
Brazil	7%	1136	6	54%	61%	138 (65%)	213
UK	11%	988	60	57%	58%	33 (49%)	68
France	17%	964	74	49%	53%	19 (29%)	65
Spain	11%	863	73	54%	66%	33 (70%)	47
Italy	10%	819	44	61%	59%	37 (62%)	60
Mexico	8%	648	5	53%	57%	27 (21%)	129
Belgium	16%	331	20	52%	52%	5 (42%)	12
Germany	15%	266	24	47%	59%	13 (15%)	84
Canada	13%	174	8	60%	47%	4 (11%)	38

data¹¹. The subsequent two columns make the point that the peak fatality rates at the second wave (based upon posterior predictions) are considerably smaller than the corresponding peak fatality rates at the first wave. For most countries, this second peak is in the order of tens of deaths per day, as opposed to hundreds.

¹¹ See also: <https://www.gov.uk/government/publications/national-covid-19-surveillance-reports/sero-surveillance-of-covid-19>

The proportion of the susceptible population who cannot transmit the virus ranges from 47% to 61% (non-contagious column). This reflects heterogeneity in transmission. The corresponding heterogeneity of susceptibility is reflected in the proportion of the effective population that are not susceptible to infection (non-susceptible column). Finally, the difference between the effective and total population size reflects heterogeneity of exposure. In most instances, the effective population constitutes a large proportion of the total population (largest in Brazil, Spain,

and Italy), with the exception of Canada, where the effective population is only four out of 38 million.

Discussion

As noted in the introduction, this technical report should be read as part of a series demonstrating the application of variational Bayesian inference to quantitative epidemiological modelling. The first, foundational report (Friston *et al.*, 2020d) established the nature of the model and attending inversion scheme. The second dealt with coupling models of a single region or population to illustrate how exchange between populations can be handled (Friston *et al.*, 2020b). The third report (Friston *et al.*, 2020c) focused on posterior predictions and projections under different scenarios (e.g., testing and tracing). This report illustrates a particular application of Bayesian model comparison (strictly speaking, Bayesian model reduction) known in some fields as structure learning (Friston *et al.*, 2016). As such, the estimates and inferences reported here should not be taken as definitive. Rather, we have described the procedures that could be used to furnish these kinds of estimates during the course of the current epidemic or in the future.

One may ask why we chose to use Bayesian model comparison to form posterior beliefs about a particular parameter of the model, as opposed to simply evaluating its posterior under the Laplace (i.e., Gaussian form) assumption? Our motivation was twofold. First, it shows how one can eschew the Laplace assumption and use Bayesian model reduction to build a non-Gaussian posterior belief over a parameter of interest. For example, the posterior could have been bimodal. This application of Bayesian model comparison shows how it is possible to leverage the computational efficiency of variational Bayes, without committing to a fixed-form posterior over one or a small number of interesting parameters. The second reason—for illustrating Bayesian model comparison in this way—was to show how to accumulate evidence from multiple datasets (here, different countries). This pooling reduces to adding the (logarithms of) evidence for the same model from independent data¹². Note that a model is defined here in terms of prior beliefs. This means that one can use Bayesian model reduction to score the quality of any prior beliefs empirically. In one sense, this is an example of empirical Bayes (Efron & Morris, 1973; Friston *et al.*, 2016; Kass & Steffey, 1989).

Although our focus is on Bayesian model comparison, it may be useful to rehearse the distinction between the variational approaches used in dynamic causal modelling and the usual approaches found in the epidemiological literature. Perhaps the most important difference is the way that model evidence or

marginal likelihood is handled or evaluated. In variational Bayes, this is computed explicitly in terms of an evidence bound afforded by the variational free energy (Beal, 2003; Fox & Roberts, 2011; Winn & Bishop, 2005). This uses the entire time series and a computationally efficient scheme afforded by assumptions about the shape and factorisation of an approximate posterior. The alternative would be to eschew any assumptions about the form of the posterior and approximate the marginal likelihood of a model in terms of its crossvalidation accuracy. Technically, the log of model evidence is accuracy minus complexity. In the setting of crossvalidation, one can ignore the complexity term and approximate model evidence with the accuracy with which some new (i.e. test) data are explained (MacKay, 2003). On this view, crossvalidation accuracy becomes another approximation to log evidence. Models with a greater evidence are those that, on average, generalise to new data and therefore have the greatest predictive validity. This means that variational approaches evaluate model evidence or marginal likelihood explicitly, while other approaches (e.g. stochastic or sampling approaches) use crossvalidation or predictive validity. Although it is possible to demonstrate the predictive validity of dynamic causal models by withholding test data—see Figure 13 in (Friston *et al.*, 2020d)—this is not necessary because we already know that model selected has a greater evidence than another model.

In a similar vein, the conventional epidemiological modelling literature often features sensitivity analyses of the parameters. For a sensitivity analysis of the dynamic causal model used above, please see Figure 9 in (Friston *et al.*, 2020d). These kinds of analyses allow one to eyeball which parameters make a difference or, from the point of view of model inversion, which parameters are informed by the data. Although a useful adjunct to dynamic causal modelling, sensitivity analyses of this sort are not necessary to understand the relationship between model parameters and data. This is because the sensitivity (i.e., the derivative of the data with respect to model parameters) is an integral part of the posterior uncertainty under the Laplace assumption: see Equation 23 in (Friston *et al.*, 2007). In other words, the sensitivity is handled implicitly in terms of uncertainty quantification or the posterior credible intervals ascribed to various parameters. In brief, a parameter that has tight credible interval implies that small changes in this parameter produce large changes in data space. The ensuing estimates of posterior uncertainty are, effectively, then used to marginalise out uncertainty about the parameters to furnish the marginal likelihood required for model comparison.

Note that the posterior maximises model evidence, which is the same as maximising accuracy while minimising complexity. Complexity in variational inference corresponds to the Kullback-Leibler divergence between the prior and posterior (Penny *et al.*, 2004). Heuristically, this can be regarded as the degrees of freedom used up to explain the data. Crucially, this means that the data fit or accuracy is only half the story. One has to provide an accurate account of the data as simply as possible. Procedures based upon the Akaike and Bayesian information criteria do not evaluate the complexity explicitly and can be

¹² By pooling the evidence over countries in this way, we have allowed for country-specific differences in the parameters shaping their epidemics. However, we then assume that all countries share the same loss of immunity. A consideration of between country variations and conditional dependencies between parameters would require a different approach. Usually, this would be addressed using hierarchical Bayesian models (a.k.a., parametric empirical Bayes). An example of this can be found in the first report examining within and between country differences: see Figure 6 in (Friston *et al.*, 2020d).

dangerously misleading when used for model comparison (Cornish & Littenberg, 2007; Penny, 2012). This speaks again to the potential utility of variational procedures in epidemiological modelling.

There are clearly many limitations to the modelling here. These include modelling each outbreak as a point process and ignoring geospatial aspects and waves of infection (Chinazzi *et al.*, 2020). Furthermore, we have assumed idealised dynamics that do not consider interactions with seasonal influenza or any other annual fluctuations (Kissler *et al.*, 2020). As with all dynamic causal modelling studies, the conclusions based upon Bayesian model comparison and posterior inferences are limited to the models considered. Additionally, the posterior predictions will change as more data becomes available. Having said this, it is interesting to note, irrespective of the modelling, that there is sufficient information—in the current epidemiological trajectories—to support fairly precise posterior beliefs about how quickly we will lose immunity. In our treatment of heterogeneity, we have not explicitly modelled things like age, ethnicity, population density etc. Instead, we have simply modelled the implicit heterogeneity—at a coarse-grained level—by using a series of bipartitions of the latent states. More refined models could consider stratification by age or ethnicity, with appropriate contact matrices. Whether this fine graining of heterogeneity is justified by the data can be cast as a question of model comparison. If the model is too expressive or complex for the data at hand, including age stratification will reduce model evidence. However, with disaggregated and sufficiently long time series, Bayesian model comparison could, in principle, identify whether the attributes above play an important role. And, if they do, one could assess their quantitative contribution in terms of the posterior parameter estimates.

Death rates in the United Kingdom over the next few weeks will be telling: if they can be suppressed to zero, then it is possible that the effective (population) immunity will be enduring, and we may elude a second wave. If, on the other hand, fatality rates continue above 20 a day, then according to the model presented here, it is likely we will see a slow increase in the reproduction rate and a second wave after Christmas. Note that the analyses in this report are predicated on a track and trace process whose efficacy is estimated based upon the data to date. As discussed in (Friston *et al.*, 2020c) and elsewhere, any second wave could be deferred by introducing a more efficacious tracking and tracing protocol, even in the context of a relatively rapid loss of population immunity, such as the three

month period estimated here. This deferment rests upon finding a substantial proportion of infected individuals before they can transmit the virus by identifying local outbreaks and clusters. On one view, this takes us out of the arena of ensemble dynamics and epidemiological modelling into the pragmatic considerations of an effective local surveillance and public health response.

Methods

Software note

The annotated (MATLAB/Octave) code is available as part of the free and open source academic software SPM (<https://www.fil.ion.ucl.ac.uk/spm/>), released under the terms of the GNU General Public License version 2 or later. The routines are called by a demonstration script that can be invoked by typing DEM_COVID_I at the MATLAB prompt. For this technical report, we used MATLAB R2019b and SPM12 r7872 (archived at <https://doi.org/10.6084/m9.figshare.12174006.v5>; Friston *et al.*, 2020d).

We recommend anyone interested in applying this model should use the latest version of the software available. Details about future developments of the software will be available from <https://www.fil.ion.ucl.ac.uk/spm/covid-19/>.

Data availability

Source data

The data used in this technical report are available for academic research purposes from the COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University, hosted on GitHub at <https://github.com/CSSEGISandData/COVID-19>.

Software availability

Software is available from: <https://www.fil.ion.ucl.ac.uk/spm/covid-19/>.

Archived source code at time of publication: <https://doi.org/10.6084/m9.figshare.12174006.v5> (Friston *et al.*, 2020d).

License: GPL 2.0+.

Acknowledgements

An earlier version of this article can be found on arXiv (<https://arxiv.org/abs/2006.09429>).

References

Aleta A, Martin-Corral C, Pastore Y, *et al.*: **Modeling the impact of social distancing, testing, contact tracing and household quarantine on second-wave scenarios of the Covid-19 epidemic.** *medRxiv*. 2020.
[Publisher Full Text](#)

Beal MJ: **Variational Algorithms for Approximate Bayesian Inference.** PhD Thesis, University College London, 2003.
[Reference Source](#)

Beutler B, Eidschinken C, Crozat K, *et al.*: **Genetic analysis of resistance to**

- viral infection. *Nat Rev Immunol.* 2007; **7**(10): 753–766.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Bunyanich S, Do A, Vicencio A: **Nasal Gene Expression of Angiotensin-Converting Enzyme 2 in Children and Adults.** *JAMA.* 2020; **323**(23): 2427–2429.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Chan KH, Chan JFW, Tse H, *et al.*: **Cross-reactive antibodies in convalescent SARS patients' sera against the emerging novel human coronavirus EMC (2012) by both immunofluorescent and neutralizing antibody tests.** *J Infect.* 2013; **67**(2): 130–140.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Chau NVV, Thanh Lam V, Thanh Dung N, *et al.*: **The natural history and transmission potential of asymptomatic SARS-CoV-2 infection.** *Clin Infect Dis.* 2020; ciaa711.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Chinazzi M, Davis JT, Ajelli M, *et al.*: **The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak.** *Science.* 2020; **368**(6489): 395–400.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Cornish NJ, Littenberg TB: **Tests of Bayesian model selection techniques for gravitational wave astronomy.** *Phys Rev D.* 2007; **76**(8).
[Publisher Full Text](#)
- Dauwels J: **On Variational Message Passing on Factor Graphs.** 2007 *IEEE International Symposium on Information Theory.* 2007; 2546–2550.
[Publisher Full Text](#)
- Efron B, Morris C: **Stein's estimation rule and its competitors – an empirical Bayes approach.** *J Am Stat Assoc.* 1973; **68**(341): 117–30.
[Publisher Full Text](#)
- Endo A, Abbott S, Kucharski AJ, *et al.*: **Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China [version 3; peer review: 2 approved].** *Wellcome Open Res.* 2020; **5**: 67.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Feynman RP: **Statistical mechanics.** Benjamin, Reading MA. 1972.
- Fox C, Roberts S: **A tutorial on variational Bayes.** *Artificial Intelligence Review.* 2011; **38**(2): 1–11.
[Publisher Full Text](#)
- Friston K, Mattout J, Trujillo-Barreto N, *et al.*: **Variational free energy and the Laplace approximation.** *NeuroImage.* 2007; **34**(1): 220–234.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Friston K, Parr T, Zeidman P, *et al.*: **Dynamic causal modelling of COVID-19 [version 2; peer review: 2 approved].** *Wellcome Open Res.* 2020a; **5**: 89.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Friston K, Parr T, Zeidman P, *et al.*: **Second waves, social distancing, and the spread of COVID-19 across America [version 1; peer review: awaiting peer review].** *Wellcome Open Res.* 2020b; **5**: 103.
[Publisher Full Text](#)
- Friston KJ, Parr T, Zeidman P, *et al.*: **Tracking and tracing in the UK: a dynamic causal modelling study.** *arXiv e-prints, arXiv: 2005.07994.* 2020c.
[Reference Source](#)
- Friston K, Parr T, Zeidman P, *et al.*: **Dynamic Causal Modelling of COVID-19.** *figshare.* Software. 2020d.
<http://www.doi.org/10.6084/m9.figshare.12174006.v5>
- Friston KJ, Litvak V, Oswal A, *et al.*: **Bayesian model reduction and empirical Bayes for group (DCM) studies.** *NeuroImage.* 2016; **128**: 413–31.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Greenland S: **Bayesian perspectives for epidemiological research: I. Foundations and basic methods.** *Int J Epidemiol.* 2006; **35**(3): 765–775.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Grifoni A, Weiskopf D, Ramirez SI, *et al.*: **Targets of T cell responses to SARS-CoV-2 coronavirus in humans with COVID-19 disease and unexposed individuals.** *Cell.* 2020; **181**(7): 1489–1501.e15.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Hellewell J, Abbott S, Gimma A, *et al.*: **Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts.** *Lancet Glob Health.* 2020; **8**(4): e488–e496.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Houlihan C, Vora N, Byrne T, *et al.*: **SARS-CoV-2 virus and antibodies in front-line Health Care Workers in an acute hospital in London: preliminary results from a longitudinal study.** *medRxiv.* 2020; 2020.2006.2008.20120584.
[Publisher Full Text](#)
- Huang C, Wang Y, Li X, *et al.*: **Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China.** *Lancet.* 2020; **395**(10223): 497–506.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Ing AJ, Cocks C, Green JP: **COVID-19: in the footsteps of Ernest Shackleton.** *Thorax.* 2020; **75**(8): 693–694.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Kass RE, Raftery AE: **Bayes Factors.** *J Am Stat Assoc.* 1995; **90**(430): 773–795.
[Publisher Full Text](#)
- Kass RE, Steffey D: **Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical Bayes models).** *J Am Stat Assoc.* 1989; **407**: 717–26.
[Publisher Full Text](#)
- Keeling MJ, Hollingsworth TD, Read JM: **Efficacy of Contact Tracing for**
- the Containment of the 2019 Novel Coronavirus (COVID-19).** *J Epidemiol Community Health.* 2020; **74**(10): 861–866.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Kermack WO, McKendrick AG: **A contribution to the mathematical theory of epidemics.** *P R Soc Lond A-Conta.* 1997; **115**(772).
[Publisher Full Text](#)
- Kissler SM, Tedijanto C, Goldstein E, *et al.*: **Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period.** *Science.* 2020; **368**(6493): 860–868.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Kucharski AJ, Russell TW, Diamond C, *et al.*: **Early dynamics of transmission and control of COVID-19: a mathematical modelling study.** *Lancet Infect Dis.* 2020; **20**(5): 553–558.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Lloyd-Smith JO, Schreiber SJ, Kopp PE, *et al.*: **Superspreading and the effect of individual variation on disease emergence.** *Nature.* 2005; **438**(7066): 355–359.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- MacKay DJ: **Free energy minimisation algorithm for decoding and Cryptanalysis.** *Electron Lett.* 1995; **31**(6): 446.
[Publisher Full Text](#)
- MacKay DJC: **Information Theory, Inference and Learning Algorithms.** Cambridge University Press, Cambridge. 2003.
[Reference Source](#)
- Marreiros AC, Kiebel SJ, Daunizeau J, *et al.*: **Population dynamics under the Laplace assumption.** *Neuroimage.* 2009; **44**(3): 701–714.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Mizumoto K, Chowell G: **Estimating Risk from Death from 2019 Novel Coronavirus Disease, China, January-February 2020.** *Emerg Infect Dis.* 2020; **26**(6): 1251–1256.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Moran R, Fagerholm E, Cullen M, *et al.*: **Estimating required 'lockdown' cycles before immunity to SARS-CoV-2: model-based analyses of susceptible population sizes, 'S0', in seven European countries, including the UK and Ireland [version 1; peer review: 1 approved with reservations, 1 not approved].** *Wellcome Open Res.* 2020; **5**.
[Publisher Full Text](#)
- Ng K, Faulkner N, Cornish G, *et al.*: **Pre-existing and de novo humoral immunity to SARS-CoV-2 in humans.** *bioRxiv.* 2020; 2020.2005.2014.095414.
[Publisher Full Text](#)
- Paynter S: **Incorporating Transmission Into Causal Models of Infectious Diseases for Improved Understanding of the Effect and Impact of Risk Factors.** *Am J Epidemiol.* 2016; **183**(6): 574–582.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Penny WD: **Comparing dynamic causal models using AIC, BIC and free energy.** *Neuroimage.* 2012; **59**(1): 319–330.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Penny WD, Stephan KE, Mechelli A: **Comparing dynamic causal models.** *Neuroimage.* 2004; **22**(3): 1157–72.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Rhodes CJ, Hollingsworth TD: **Variational data assimilation with epidemic models.** *J Theor Biol.* 2009; **258**(4): 591–602.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Russell TW, Hellewell J, Jarvis CI, *et al.*: **Estimating the infection and case fatality ratio for coronavirus disease (COVID-19) using age-adjusted data from the outbreak on the Diamond Princess cruise ship, February 2020.** *Euro surveillance: bulletin Européen sur les maladies transmissibles = European communicable disease bulletin.* 2020; **25**.
[Reference Source](#)
- Seo GY, Giles DA, Kronenberg M: **The role of innate lymphoid cells in response to microbes at mucosal surfaces.** *Mucosal Immunol.* 2020; **13**(3): 399–412.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Simonsen L, Chowell G, Andreasen V, *et al.*: **A review of the 1918 herald pandemic wave: importance for contemporary pandemic response strategies.** *Ann Epidemiol.* 2018; **28**(5): 281–288.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Streeck H, Schulte B, Kuemmerer B, *et al.*: **Infection fatality rate of SARS-CoV-2 infection in a German community with a super-spreading event.** *medRxiv.* 2020; 2020.2005.2004.20090076.
[Publisher Full Text](#)
- Stringhini S, Wisniak A, Piumatti G, *et al.*: **Seroprevalence of anti-SARS-CoV-2 IgG antibodies in Geneva, Switzerland (SEROCoV-POP): a population-based study.** *Lancet.* 2020; **396**(10247): 313–319.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Su S, Wong G, Shi W, *et al.*: **Epidemiology, Genetic Recombination, and Pathogenesis of Coronaviruses.** *Trends Microbiol.* 2016; **24**(6): 490–502.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Tsallis C, Tirnakli U: **Predicting COVID-19 Peaks Around the World.** *Frontiers in Physics.* 2020; **8**: 217.
[Publisher Full Text](#)
- van Kampen JJA, van de Vijver DA, Fraaij PL, *et al.*: **Shedding of infectious virus in hospitalized patients with coronavirus disease-2019 (COVID-19):**

duration and key determinants. *medRxiv*. 2020; 2020.2006.2008.20125310.
[Publisher Full Text](#)

Verity R, Okell LC, Dorigatti I, *et al.*: **Estimates of the severity of coronavirus disease 2019: a model-based analysis.** *Lancet Infect Dis*. 2020; **20**(6): 669–677.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)

Vineis P, Kriebel D: **Causal models in epidemiology: past inheritance and genetic future.** *Environ Health*. 2006; **5**: 21–21.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)

Wang D, Hu B, Hu C, *et al.*: **Clinical Characteristics of 138 Hospitalized Patients With 2019 Novel Coronavirus-Infected Pneumonia in Wuhan, China.** *JAMA*. 2020; **323**(11): 1061–1069.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)

White LJ, Mandl JN, Gomes MG, *et al.*: **Understanding the transmission dynamics of respiratory syncytial virus using multiple time series and**

nested models. *Math Biosci*. 2007; **209**(1): 222–239.

[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)

Winn J, Bishop CM: **Variational message passing.** *J Mach Learn Res*. 2005; **6**: 661–694.
[Reference Source](#)

Wölfel R, Corman VM, Guggemos W, *et al.*: **Virological assessment of hospitalized patients with COVID-2019.** *Nature*. 2020; **581**(7809): 465–469.
[PubMed Abstract](#) | [Publisher Full Text](#)

Wu JT, Leung K, Leung GM: **Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study.** *Lancet*. 2020; **395**(10225): 689–697.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)

Zheng M, Gao Y, Wang G, *et al.*: **Functional exhaustion of antiviral lymphocytes in COVID-19 patients.** *Cell Mol Immunol*. 2020; **17**(5): 533–535.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)

Open Peer Review

Current Peer Review Status:

Version 2

Reviewer Report 09 October 2020

<https://doi.org/10.21956/wellcomeopenres.17961.r40671>

© 2020 Madzvamuse A. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Anotida Madzvamuse

School of Mathematical and Physical Sciences, Department of Mathematics, University of Sussex, Brighton, UK

The authors have addressed fully my comments and happy to approve the revised version.

I would recommend the authors to consider referencing this archived report (entirely as optional by the authors as they see fit) since it also sets similar approaches for data-driven parameter values for models for COVID-19 with forecasting capabilities. Predicting and forecasting the impact of local resurgence and outbreaks of COVID-19: Use of SEIR-D quantitative epidemiological modelling for healthcare demand and capacity (2020)¹.

References

1. Campillo-Funollet E, Van Yperen J, Allman P, Bell M, et al.: Predicting and forecasting the impact of local resurgence and outbreaks of COVID-19: Use of SEIR-D quantitative epidemiological modelling for healthcare demand and capacity. *medRxiv*. 2020. [Publisher Full Text](#)

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Mathematical and Computational Biology, Analysis of epidemiological models, Inverse problems, Bayesian Approach.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Version 1

Reviewer Report 17 September 2020

<https://doi.org/10.21956/wellcomeopenres.17854.r40237>

© 2020 Madzvamuse A. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Anotida Madzvamuse

School of Mathematical and Physical Sciences, Department of Mathematics, University of Sussex, Brighton, UK

The authors present a dynamic causal model to study the quantify the effects and impact of the rate at which effective immunity to COVID-19 is lost. The study follows several works by the authors where some of the details involving the model and the code are described. I found the study lacking key aspects associated with rigorous scientific validation of the model output and its predictions.

1. The dynamical causal model is not fully described but that the reader is referred to previous work by the authors. This is a crucial part of this study, instead the model in its full glory should have been included as supplementary material to aid readability of the work.
2. The study fails to convince the reviewer that a rigorous mathematical/statistical validation has been carried out to understand predictive power of the dynamical causal model.
3. The study also does not present sensitivity of the parameters to understand which parameters are more sensitive to small changes in data.
4. Given that the authors employ the Bayesian approach with detailed datasets, the 25 parameters driving the model should be inferred based on how best the model fits to data. This is not apparent in the study.
5. To really demonstrate the predicting and forecasting power of the methodology, the authors should have taken a subset of the data sets, say, in Figure 3 for example, take data for 20, 40, 60 and 80 days to infer model parameters and use these in the model to demonstrate how accurate they can predict data-points which are known but are in the future. At the moment, predictions are made where it is impossible to compare with data since there is no such data (after say, 100 days). This is a major weakness of the forecasting approach adopted in this work.
6. The authors state that there are many mechanisms that could contribute to the loss immunity and give some examples. However, they are not able to differentiate which ones are key mechanisms. Given the use of Bayesian approach, they should be able to use data to select the best mechanism such the model with such a mechanism best-fits data.
7. In Figure 5, it is extremely difficult to differentiate between model solutions given that the same color is used throughout.
8. The timescale for loss of immunity is prescribed a priori, the authors really could have exploited the inverse approach associated with Bayes theorem to infer this timescale.
9. It is not clear how the effective reproduction number R_t is computed under this proposed

approach.

10. It is also not clear if the same graphs are obtained for all the countries included in the study. Does it mean that country-specific characteristics are not important in determining the rate at which immunity is lost?
11. Is R_t the same for all countries?
12. The comparison between the UK and Germany does not seem to make much sense since Germany has had some of the fewest infections due to COVID.
13. In this study, the authors ignore patient specifics, such as age, ethnicity, location, etc. How important are these factors?

Is the rationale for developing the new method (or application) clearly explained?

Yes

Is the description of the method technically sound?

Partly

Are sufficient details provided to allow replication of the method development and its use by others?

Partly

If any results are presented, are all the source data underlying the results available to ensure full reproducibility?

Partly

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?

Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Mathematical and Computational Biology, Analysis of epidemiological models, Inverse problems, Bayesian Approach.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to state that I do not consider it to be of an acceptable scientific standard, for reasons outlined above.

Author Response 21 Sep 2020

Adeel Razi, University College London, London, UK

Dear Prof. Madzvamuse,

We would like to thank you for the time and effort you have spent reviewing our manuscript. Below are the replies to the comments. We hope these revisions are what you had in mind.

The authors present a dynamic causal model to study the quantify the effects and impact of the rate at which effective immunity to COVID-19 is lost. The study follows several works by the authors where some of the details involving the model and the code are described. I found the study lacking key aspects associated with rigorous scientific validation of the model output and its predictions.

1. *The dynamical causal model is not fully described but that the reader is referred to previous work by the authors. This is a crucial part of this study, instead the model in its full glory should have been included as supplementary material to aid readability of the work.*
2. *The study fails to convince the reviewer that a rigorous mathematical/statistical validation has been carried out to understand predictive power of the dynamical causal model.*
3. *The study also does not present sensitivity of the parameters to understand which parameters are more sensitive to small changes in data.*
4. *Given that the authors employ the Bayesian approach with detailed datasets, the 25 parameters driving the model should be inferred based on how best the model fits to data. This is not apparent in the study.*
5. *To really demonstrate the predicting and forecasting power of the methodology, the authors should have taken a subset of the data sets, say, in Figure 3 for example, take data for 20, 40, 60 and 80 days to infer model parameters and use these in the model to demonstrate how accurate they can predict data-points which are known but are in the future. At the moment, predictions are made where it is impossible to compare with data since there is no such data (after say, 100 days). This is a major weakness of the forecasting approach adopted in this work.*

Many thanks for these comments and critiques. These are understandable points that we hope to have addressed by delimiting the scope and target of this report. For example, this is not an example of forecasting – it was meant as a worked example of Bayesian model comparison to get at a key model parameter. We have included the following qualifications in the new Discussion section and also updated the Background section in the revised manuscript – but would be happy to revisit or extend this discussion if the reviewer thought appropriate.

- "As noted in the introduction, this technical report should be read as part of a series demonstrating the application of variational Bayesian inference to quantitative epidemiological modelling. The first, foundational report (Friston *et al.*, 2020d) established the nature of the model and attending inversion scheme. The second dealt with coupling models of a single region or population to illustrate how exchange between populations can be handled (Friston *et al.*, 2020b). The third report (Friston *et al.*, 2020c) focused on posterior predictions and projections under different scenarios (e.g., testing and tracing). This report illustrates a particular application of Bayesian model comparison (strictly speaking, Bayesian model reduction) known in

some fields as structure learning (Friston *et al.*, 2016). As such, the estimates and inferences reported here should not be taken as definitive. Rather, we have described the procedures that could be used to furnish these kinds of estimates during the course of the current epidemic or in the future."

- "One may ask why we chose to use Bayesian model comparison to form posterior beliefs about a particular parameter of the model, as opposed to simply evaluating its posterior under the Laplace (i.e., Gaussian form) assumption? Our motivation was twofold. First, it shows how one can eschew the Laplace assumption and use Bayesian model reduction to build a non-Gaussian posterior belief over a parameter of interest. For example, the posterior could have been bimodal. This application of Bayesian model comparison shows how it is possible to leverage the computational efficiency of variational Bayes, without committing to a fixed-form posterior over one or a small number of interesting parameters. The second reason—for illustrating Bayesian model comparison in this way—was to show how to accumulate evidence from multiple datasets (here, different countries). This pooling reduces to adding the (logarithms of) evidence for the same model from independent data. Note that a model is defined here in terms of prior beliefs. This means that one can use Bayesian model reduction to score the quality of any prior beliefs empirically. In one sense, this is an example of empirical Bayes (Efron and Morris, 1973; Kass and Steffey, 1989; Friston *et al.*, 2016)."
- "Although our focus is on Bayesian model comparison, it may be useful to rehearse the distinction between the variational approaches used in dynamic causal modelling and the usual approaches found in the epidemiological literature. Perhaps the most important difference is the way that model evidence or marginal likelihood is handled or evaluated. In variational Bayes, this is computed explicitly in terms of an evidence bound afforded by the variational free energy (Beal, 2003; Winn and Bishop, 2005; Fox and Roberts, 2011). This uses the entire time series and a computationally efficient scheme afforded by assumptions about the shape and factorisation of an approximate posterior. The alternative would be to eschew any assumptions about the form of the posterior and approximate the marginal likelihood of a model in terms of its crossvalidation accuracy. Technically, the log of model evidence is accuracy minus complexity. In the setting of crossvalidation, one can ignore the complexity term and approximate model evidence with the accuracy with which some new (i.e. test) data are explained (MacKay, 2003). On this view, crossvalidation accuracy becomes another approximation to log evidence. Models with a greater evidence are those that, on average, generalise to new data and therefore have the greatest predictive validity. This means that variational approaches evaluate model evidence or marginal likelihood explicitly, while other approaches (e.g. stochastic or sampling approaches) use crossvalidation or predictive validity. Although it is possible to demonstrate the predictive validity of dynamic causal models by withholding test data—see Figure 13 in (Friston *et al.*, 2020d)—this is not necessary because we already know that model selected has a greater evidence than another model."
- "In a similar vein, the conventional epidemiological modelling literature often features sensitivity analyses of the parameters. For a sensitivity analysis of the dynamic causal model used above, please see Figure 9 in (Friston *et al.*, 2020d). These kinds of analyses allow one to eyeball which parameters make a difference or, from the point of view of model inversion, which parameters are informed by the data."

Although a useful adjunct to dynamic causal modelling, sensitivity analyses of this sort are not necessary to understand the relationship between model parameters and data. This is because the sensitivity (i.e., the derivative of the data with respect to model parameters) is an integral part of the posterior uncertainty under the Laplace assumption: see Equation 23 in (Friston *et al.*, 2007). In other words, the sensitivity is handled implicitly in terms of uncertainty quantification or the posterior credible intervals ascribed to various parameters. In brief, a parameter that has tight credible interval implies that small changes in this parameter produce large changes in data space. The ensuing estimates of posterior uncertainty are, effectively, then used to marginalise out uncertainty about the parameters to furnish the marginal likelihood required for model comparison."

- "Note that the posterior maximises model evidence, which is the same as maximising accuracy while minimising complexity. Complexity in variational inference corresponds to the Kullback-Leibler divergence between the prior and posterior (Penny *et al.*, 2004). Heuristically, this can be regarded as the degrees of freedom used up to explain the data. Crucially, this means that the data fit or accuracy is only half the story. One has to provide an accurate account of the data as simply as possible. Procedures based upon the Akaike and Bayesian information criteria do not evaluate the complexity explicitly and can be dangerously misleading when used for model comparison (Cornish and Littenberg, 2007; Penny, 2012). This speaks again to the potential utility of variational procedures in epidemiological modelling."

It was useful to have your perspective on what constitutes a rigorous statistical validation and we hope to have addressed your concerns above. Highlighting what variational (Bayes) procedures bring to the table (including a much simpler and easier way of doing things) is the *raison d'être* of these reports.

6. *The authors state that there are many mechanisms that could contribute to the loss immunity and give some examples. However, they are not able to differentiate which ones are key mechanisms. Given the use of Bayesian approach, they should be able to use data to select the best mechanism such the model with such a mechanism best-fits data.*

This is a very good point. We have included the following in the Background section to address it:

- "In talking about a loss of effective immunity, we assume that there could be many contributions to this loss. For example, a decline in antibody levels, viral mutation, or a dilution of population immunity due to population fluxes. These are all important mechanistic hypotheses that can, in principle, be addressed using Bayesian model comparison. To do this, it would be necessary to parameterise the model in a way that allowed one to withdraw one or other mechanism and evaluate whether the model evidence increased or decreased. An example of this can be found in (Friston *et al.*, 2020a), where the relative contribution of lockdown and population immunity to prevalence and mortality was evaluated. Interestingly, both lockdown and herd immunity were necessary to explain the data: in the sense that removing either mechanism substantially reduced model evidence. Please see (Friston *et al.*, 2020a) for details."

7. *In Figure 5, it is extremely difficult to differentiate between model solutions given that the same color is used throughout.*

Thank you – we have adjusted the colour scheme.

8. *The timescale for loss of immunity is prescribed a priori, the authors really could have exploited the inverse approach associated with Bayes theorem to infer this timescale.*

It would certainly have been possible to report the Gaussian posterior following a model inversion under uninformative priors. However, as noted above, this would commit to a Gaussian form for the posterior. Furthermore, this would involve a slightly more complicated approach to assimilating the data from multiple countries (using Bayesian parameter averaging). We could go into this in more detail if you thought it necessary.

9. *It is not clear how the effective reproduction number R_t is computed under this proposed approach.*

The effective reproduction number is computed as described in (Friston *et al.*, 2020c). We have now included this in the revised manuscript (in the Results section) as follows:

- "The effective reproduction rate provides a useful statistic that reflects the exponential growth of the prevalence of infection. There are several ways in which it can be estimated. For our purposes, we can evaluate an instantaneous reproduction rate directly from the time varying prevalence of infection as follows: (Mathematical typesetting feature is not available in this Editor. Please refer to the Results section in the revised draft.) These expressions show that the reproduction rate reflects the growth of the (logarithm of the) proportion of people infected—and the period of being infectious. This is related to the doubling time T_d . Note that the reproduction rate is not a parameter of the model: it is an outcome that is generated by the latent states inferred by inverting (i.e., fitting) the model to empirical timeseries."

10. *It is also not clear if the same graphs are obtained for all the countries included in the study. Does it mean that country-specific characteristics are not important in determining the rate at which immunity is lost?*

This is a very good question. We have tried to answer it – as a footnote 1 in the revised draft – with the following:

- "By pooling the evidence over countries in this way, we have allowed for country-specific differences in the parameters shaping their epidemics. However, we then assume that all countries share the same loss of immunity. A consideration of between country variations and conditional dependencies between parameters would require a different approach. Usually, this would be addressed using hierarchical Bayesian models (a.k.a., parametric empirical Bayes). An example of this can be found in the first report examining within and between country differences: see Figure 6 in (Friston *et al.*, 2020d)."

11. *Is R_t the same for all countries?*

No. It is evaluated at each point in time on the basis of latent states that are generated from the model parameters (please see above). These parameters were estimated independently

for each country. The model evidence under different priors about the immunity parameter is then pooled over countries. This can be regarded as a form of empirical Bayes (please see above).

12. *The comparison between the UK and Germany does not seem to make much sense since Germany has had some of the fewest infections due to COVID.*

This question reflects prior assumptions about the relationship between the prevalence of infection at the onset of the outbreak and the occurrence of a second wave. It is precisely these prior assumptions that the current modelling allows one to address. The nonlinearity in these models can sometimes confound prior assumptions of the sort.

13. *In this study, the authors ignore patient specifics, such as age, ethnicity, location, etc. How important are these factors?*

It is likely that they are very important. Again, questions of this sort can only be answered definitively by including models with the patient specifics in them. We have added the following in the new Discussion section where we discussed limitations and future work:

- "In our treatment of heterogeneity, we have not explicitly modelled things like age, ethnicity, population density etc. Instead, we have simply modelled the implicit heterogeneity—at a coarse-grained level—by using a series of bipartitions of the latent states. More refined models could consider stratification by age or ethnicity, with appropriate contact matrices. Whether this fine graining of heterogeneity is justified by the data can be cast as a question of model comparison. If the model is too expressive or complex for the data at hand, including age stratification will reduce model evidence. However, with disaggregated and sufficiently long time series, Bayesian model comparison could, in principle, identify whether the attributes above play an important role. And, if they do, one could assess their quantitative contribution in terms of the posterior parameter estimates."

References:

- Beal MJ. Variational Algorithms for Approximate Bayesian Inference. PhD Thesis, University College London 2003.
- Cornish NJ, Littenberg TB. Tests of Bayesian model selection techniques for gravitational wave astronomy. *Physical Review D* 2007; 76(8).
- Efron B, Morris C. Stein's estimation rule and its competitors – an empirical Bayes approach. *J Am Stats Assoc* 1973; 68: 117-30.
- Fox C, Roberts S. A tutorial on variational Bayes. *Artificial Intelligence Review* 2011; 38(2): 1-11.
- Friston K, Costello A, Pillay D. Dark matter, second waves and epidemiological modelling. *medRxiv* 2020a: 2020.09.01.20185876.
- Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W. Variational free energy and the Laplace approximation. *NeuroImage* 2007; 34(1): 220-34.
- Friston K, Parr T, Zeidman P, Razi A, Flandin G, Daunizeau J, *et al.* Second waves, social distancing, and the spread of COVID-19 across America [version 1; peer review: awaiting peer review]. *Wellcome Open Research* 2020b; 5(103).

Friston K, Parr T, Zeidman P, Razi A, Flandin G, Daunizeau J, *et al.* Testing and tracking in the UK: A dynamic causal modelling study [version 1; peer review: awaiting peer review]. Wellcome Open Research 2020c; 5(144).

Friston KJ, Litvak V, Oswal A, Razi A, Stephan KE, van Wijk BCM, *et al.* Bayesian model reduction and empirical Bayes for group (DCM) studies. *Neuroimage* 2016; 128: 413-31.

Friston KJ, Parr T, Zeidman P, Razi A, Flandin G, Daunizeau J, *et al.* Dynamic causal modelling of COVID-19. Wellcome Open Research 2020d; 5(89): 89.

Kass RE, Steffey D. Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical Bayes models). *J Am Stat Assoc* 1989; 407: 717-26.

MacKay DJC. *Information Theory, Inference and Learning Algorithms*. Cambridge: Cambridge University Press; 2003.

Penny WD. Comparing dynamic causal models using AIC, BIC and free energy. *Neuroimage* 2012; 59(1): 319-30.

Penny WD, Stephan KE, Mechelli A, Friston KJ. Comparing dynamic causal models. *Neuroimage* 2004; 22(3): 1157-72.

Winn J, Bishop CM. Variational message passing. *Journal of Machine Learning Research* 2005; 6: 661-94.

Competing Interests: No competing interests were disclosed.

Reviewer Report 15 September 2020

<https://doi.org/10.21956/wellcomeopenres.17854.r40238>

© 2020 Stern M. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The author(s) is/are employees of the US Government and therefore domestic copyright protection in USA does not apply to this work. The work may be protected under the copyright laws of other jurisdictions when used in those jurisdictions.

Michael Stern

Laboratory of Cardiovascular Science, Intramural Research Program, National Institute on Aging, National Institutes of Health, Baltimore, MD, USA

This article uses a version of the established LIST model to attempt to estimate the rate of decay of immunity of COVID-19 from published mortality data from a number of countries. It finds that the most probable immunity lifetime is 3 months. While perfectly consistent with present evidence about other coronaviruses, this is an alarming finding. Several topics require discussion, particularly possible limitations that might allow escape from this dire prophecy.

1. This was done by a novel method, running multiple models for a range prior values of the immune lifetime, and comparing pooled model evidence among models. Why was this done? If the signal for early decay of immunity is as strong as indicated, I would have thought it would come out by simply floating the decay rate as one of the many parameters and estimating its posterior in the fit.

2. Obviously, any signal about the decay of immunity must derive from the slope of the decaying phase of the infection wave. There are a variety of ways this could be misleading. While a simple form of heterogeneity is included in the form of sub-populations that are resistant, unable to transmit or effectively isolated, events and antibody data since this paper was submitted have shown, at least in the US which unfortunately dominates mortality data, that there is extreme heterogeneity of the infection on almost every spatial scale, down to the single neighborhood in New York City. These local waves of infection are no longer synchronous, if they ever were, and would be expected to put a "tail" on the first wave that might be hard to distinguish from the decay of immunity. Experience in nursing homes and prisons has shown that it is extremely difficult to have any subgroup that is truly isolated from the epidemic, and those areas again may respond late but still make a major contribution (e.g. ~45% of US mortality). In the absence of age stratification, mortality data from this virus is likely to be highly biased towards groups whose connection to the rest of society may be weaker (and therefore slower) than average.
3. The model uses an internal feedback loop to effectively estimate how social distancing interventions respond to epidemic numbers. This may be applicable in some societies, but the libertarian ethos (and then politicization) in the US has resulting in "re-opening" that is both chaotic in time and heterogeneous geographically. As the authors acknowledge, premature relaxation of mitigations can be conflated with loss of immunity. Experience since the end of the authors' data sets has shown new waves of infection related to relaxation of controls well in advance of the predicted time of the "second wave" estimated using the simple, exponential decay of effective immunity. Indeed, in many areas, the (apparent) number of infections has formed plateaus that are hard to account for by any homogeneous, compartmental model.
4. A naive, technical question for my own understanding about the underlying model structure: From the earlier papers, I gather that the likelihood of outcome variables is computed as a product of independent binomials (approximated by gaussians) over time, i.e. the only uncertainty is from sampling error in determining the individual counts at each time point. Direct stochastic simulation would predict that outcome points would be correlated, e.g. more cases now predicts more deaths after a lag. Granted that this likelihood is supposed to be marginal over the (unknown) distributions of parameters and latent compartments, is that sufficient to justify the independence assumption? The structure of the transition matrix itself would seem to imply temporal correlations.

Is the rationale for developing the new method (or application) clearly explained?

Yes

Is the description of the method technically sound?

Yes

Are sufficient details provided to allow replication of the method development and its use by others?

Yes

If any results are presented, are all the source data underlying the results available to

ensure full reproducibility?

Yes

Are the conclusions about the method and its performance adequately supported by the findings presented in the article?

Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Stochastic simulation and propagation, usually applied cardiac calcium signaling, but lately to simulation of the ongoing coronavirus pandemic.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.