

# Was $R < 1$ before the English lockdowns? On modelling mechanistic detail, causality and inference about Covid-19

S. N. Wood<sup>1</sup> and E. C. Wit<sup>2</sup>

<sup>1</sup>School of Mathematics, University of Edinburgh, UK.

<sup>2</sup>Institute of Computing, Università della Svizzera italiana, Switzerland.

simon.wood@ed.ac.uk. wite@usi.ch

March 15, 2021

## Abstract

Detail is a double edged sword in epidemiological modelling. The inclusion of mechanistic detail in models of highly complex systems has the potential to increase realism, but it also increases the number of modelling assumptions, which become harder to check as their possible interactions multiply. In a major study of the Covid-19 epidemic in England, Knock et al. (2020) fit an age structured SEIR model with added health service compartments to data on deaths, hospitalization and test results from Covid-19 in seven English regions for the period March to December 2020. The simplest version of the model has 684 states per region. One main conclusion is that only full lockdowns brought the pathogen reproduction number,  $R$ , below one, with  $R \gg 1$  in all regions on the eve of March 2020 lockdown. We critically evaluate the Knock et al. epidemiological model, and the semi-causal conclusions made using it, based on an independent reimplemention of the model designed to allow relaxation of some of its strong assumptions. In particular, Knock et al. model the effect on transmission of both non-pharmaceutical interventions and other effects, such as weather, using a piecewise linear function,  $b(t)$ , with 12 breakpoints at selected government announcement or intervention dates. We replace this representation by a smoothing spline with time varying smoothness, thereby allowing the form of  $b(t)$  to be substantially more data driven. We conclude that there is no sound basis for using the Knock et al. model and their analysis to make counterfactual statements about the number of deaths that would have occurred with different lockdown timings. However, if fits of this epidemiological model structure are viewed as a reasonable basis for inference about the time course of incidence and  $R$ , then without very strong modelling assumptions, the pathogen reproduction number was probably below one, and incidence in substantial decline, some days before either of the first two English national lockdowns. Of course this does not imply that lockdowns had no effect, but it does suggest that other non-pharmaceutical interventions (NPIs) were much more effective than Knock et al. imply.

**Keywords:** SARS-CoV-2, statistics, epidemiology, epidemic model, spline.

<https://www.maths.ed.ac.uk/~swood34/albus.zip> contains the replication code and data for this paper.

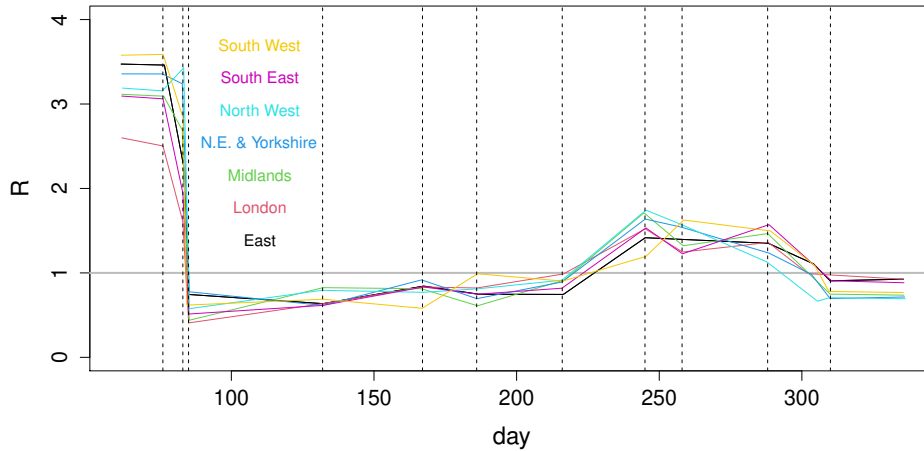


Figure 1: Estimates of  $R$  by English region against day of year, as reported in Knock et al. (2020). The plot is based on data digitized from figure 1 of Knock et al. Uncertainties were not reported. The vertical lines mark: 16th March movement restrictions, 23rd March lockdown announcement, 25 March ‘Lockdown in full effect’, May 11th initial easing, June 15th shops re-open, July 4th restaurants re-open, August 3rd eat-out-to-help-out scheme, September 1st schools open, September 14th rule of 6, October 14th Tier system, November 5th Lockdown. The kinks preceding November 5th are at a further model breakpoint.

## 1 Introduction

In principle the inclusion of known mechanisms into models used for statistical inference should improve inference by reducing the bias caused by model misspecification. But there is a catch. What happens if the mechanisms are themselves described only in an approximate manner by ad hoc sub-models? It is then possible for the assumptions built into the sub-models to introduce substantial misspecification bias. The real world consequences of such bias could be substantial if the model is used to determine major public policies. This paper examines and re-implements the model of Knock et al. (2020) to investigate the robustness of the inferences about Covid-19 lockdowns made using it. We show that key results are entirely dependent on strong but incidental assumptions introduced in the model formulation, and that relaxation of those assumptions effectively reverses the conclusions. This may matter in assessing the effectiveness of lockdowns, which have other consequences in addition to reducing viral spread. For example, they substantially modify the evolutionary landscape for the pathogen in ways that seem unlikely to offer a selective advantage for milder strains. Lockdowns also cause substantial economic hardship and exacerbate inequality. In England economic hardship and inequality are associated with very substantial loss of life, as reviewed in detail in Marmot et al. (2020).

Knock et al. (2020) is the 41st report of the COVID-19 response team from Imperial College London, whose reports have played a profound role in the shaping of UK government policy on Covid-19. Report 9 in the series provided a major component of the official justification for the first UK lockdown from March 24th 2020, and Knock et al. (2020) was accompanied by substantial press release material. A major message from Knock et al. is that the pathogen reproductive number  $R$  was only reduced below one by full lockdowns in England in March and November (see figure 1), with incidence apparently increasing until the eve of the March lockdown. We show that this result does not survive relaxation of some strong modelling assumptions. Knock et al. also present ‘counterfactual’ simulations from the calibrated model from which they draw conclusions about the deaths that could have been avoided by an earlier first lockdown. We show that these simulations can not be viewed as ‘counterfactuals’ in the usual inferential sense (e.g. Pearl et al., 2016). The avoidable death figures are simple model extrapolations.

The Knock et al. model is an age-structured SEIR model with age-structured hospital compartments.

The population is divided into 5-year age classes with a final 80+ class and two unstructured classes for care home residents and staff. There are 36 states in each of 19 classes (see figure 2) . The model was specified as a set of ODEs and converted to a discrete time stochastic model for fitting by the  $\tau$ -leap method (Gillespie, 2001). The model was fitted to daily data on hospital deaths, care home deaths, hospital admissions, general ward occupancy, ICU occupancy, antibody test results and PCR test results from surveys, which Knock et al. supply as supplementary material. Knock et al also attempt to fit data on test results from the health system. However their model does not attempt to deal with the non-random, opportunistic nature of the sampling in this data stream, despite the continual changes in test capacity, criteria for testing, and operation of the contact tracing system over the course of the data. We therefore believe that there is substantial danger of these data simply undermining the analysis and they should not be included in data to be fitted<sup>1</sup>. Data were available for seven English regions, which were fitted separately. The model has 26 free parameters.

Knock et al. based model inference (fitting) on particle filtering methods, with full fit to all regions reported to take over 100 CPU days, despite using only 96 particles per fit. This computational cost makes model checking difficult, particularly if a more usual number of particles is used and the stronger model assumptions are relaxed: the latter involves allowing substantially more free parameters plus hyper-parameters. Additionally Knock et al. specify massive overdispersion in all but the test data streams. Decreasing this over-dispersion to levels consistent with the data would likely increase particle depletion problems in filtering, leading to yet longer computing times. Given these issues, we will work directly with the ODE based model. The neglect of stochasticity in the state equations seems likely to be a minor issue here, relative to the other approximations made in the model. In particular, the only non-linearity in the model dynamics is in the transmission between infectious and susceptible sub-populations, which contain large numbers except right at the epidemic start. Other model components are controlled by simple linear flows and are also aggregated over multiple age classes for fitting. Additionally the data sampling interval and total data duration are fairly short relative to the model’s dynamic timescales. In any case, any results *dependent* on stochasticity would then require a much stronger justification for the stochastic formulation than that it was produced by discretisation of an underlying ODE model.

Furthermore, a generic strength and weakness of the particle filtering methods used by Knock et al. is that they necessarily filter the state variables as well as model parameters. This is advantageous for state forecasting, but can be more problematic for inferential tasks. For an ill-specified dynamic model the filter is often forced to repeatedly select state transitions that are improbable under the model, in order to be sufficiently close to the data. This can result in the filtered states being in an extreme tail of the posterior predictive distribution of the model: that is, of the distribution implied by simulating unfiltered states from the model given the posterior distribution of parameters. Hence model adequacy needs to be checked by comparison of the data with simulations from the posterior predictive distribution. Knock et al. do not report such checks, showing only the outputs of filtering. This is problematic when reality is then contrasted to ‘counterfactual’ simulations, necessarily from the posterior predictive distribution. The simple ODE approach used here does not filter. Instead the states are determined entirely by the model equations and the parameter values. This approach is unforgiving of model misspecification: adequacy is directly assessable from the model fit. It also reduces fit time by four orders of magnitude.

## 2 Evaluation of original Knock et al. age-structured SEIR model

In this section we review the Knock et al. (2020) model, before presenting some corrections and assumption relaxations in section 3. Figure 2 is a schematic showing the compartments in each 5-year age or care home class. The exposed, but pre-symptomatic,  $E$  stage is modelled by two sequential compartments. It

---

<sup>1</sup>We made this decision at the outset, having concluded that we would strongly advise against use of these data if acting as statistical consultants, and have never attempted to fit these data.

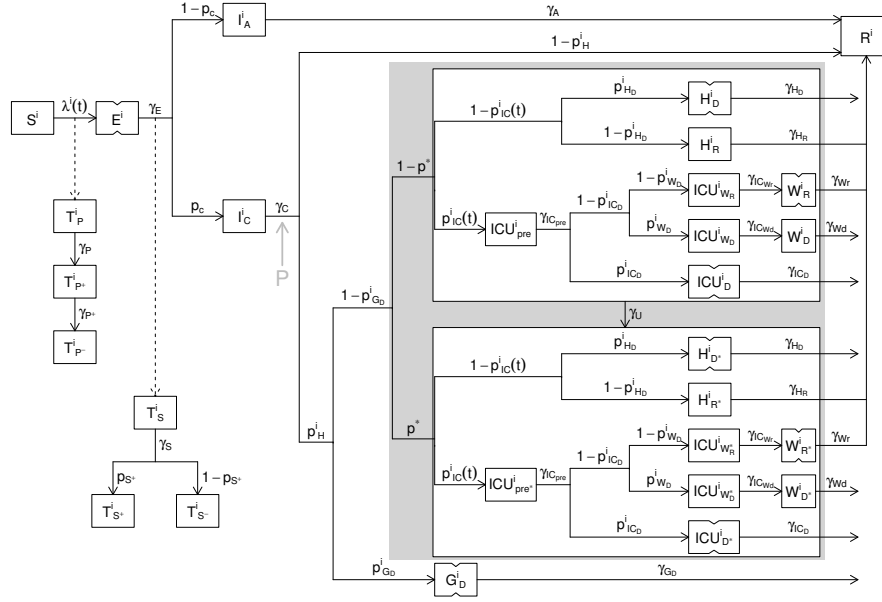


Figure 2: Schematic diagram of the model compartments (boxes) and flows (arrows) for a single model age class, following supplementary figure 2 of Knock et al. (2020), but with the notational modifications used here, stages represented as two sequential compartments indicated with notched boxes, and the location of the extra stage  $P$  that we insert to relax the generation time assumptions shown by the grey arrow and ‘ $P$ ’. To obtain the rate of flow from one compartment to another, follow the path joining them in the direction of the arrow, multiplying the source state variable by the rate parameters labelling the segments of the path. Rates with a superscript  $i$  vary with age class. The relative rates in different classes was obtained from a separate analysis reported in Knock et al., with only a common multiplier of the class specific rates left as a free parameter. For example  $p_H^i = p_H^{max} \psi_H^i$ , where  $\psi_H^i$  is fixed, but  $p_H^{max}$  is free. See section 2 and supplementary appendix A for full definitions.

is assumed that no infections are caused by this class. Symptomatic and asymptomatic stages  $I_C$  and  $I_A$  follow and cause infections, both are single compartment. The duration of the  $I_C$  stage is set from data on time from onset of symptoms to hospital admission. The absence of pre-symptomatic infection will lead to longer generation times than are reported in the literature (e.g. Flaxman et al., 2020; Anderson et al., 2020, table 1), elevating the  $R$  estimates required to achieve observed epidemic growth rates. Care home residents are not hospitalised, and the  $G_D^i$  class shown actually only receives patients for the care home resident class.

Model compartments for PCR and antibody test positivity are fed by the infection rate and the progression rate from the  $E$  state, respectively. The infection rate is driven by an age-structured mixing model with contact matrix,  $C$  based on the POLYMOD survey data for the UK (Mossong et al., 2017). Most elements of  $C$  are multiplied by a function  $b(t)$  modelling the impact of NPIs and effects such as weather on contact rates. In Knock et al.  $b(t)$  is piecewise linear with 12 breakpoints (and 12 free parameters) at policy change points. A major aim here is to relax the very strong assumptions built in to such a restrictive model. Care home contact rates are separately parameterized.

Hospitalized patients follow an ICU or general ward route. There are separate compartments for those eventually recovering or dying on the general ward. The ICU route has a pre-ICU compartment, from which patients enter compartments for those dying in ICU, entering ICU but dying later on the general ward, or entering ICU and recovering on the general ward. All compartments are duplicated for confirmed Covid (starred) and not yet confirmed (not starred), with a parameter,  $\gamma_U$ , controlling the rate of testing based transfer from unconfirmed to confirmed. It is assumed that, from the start, 25% of patients arrive at hospital with confirmed Covid. This is improbable given initial testing capacity.

The model captures many features in impressive detail, but several aspects are not modelled:

1. Separation into locked down and key worker sub-populations at lockdown is not modelled, despite

the very different values of  $R$  that must apply in these sub-populations, if lockdown is effective.

2. The assumed linearity of  $b(t)$  during lockdown precludes compensation for point 1 in fitting.
3. Seasonality or other non-NPI temporal effects on transmission are not modelled explicitly and are therefore confounded with the NPI effects, invalidating counterfactual manipulations of the latter.
4. Region-to-region transmission at the epidemic start is not represented, compromising early model fit and  $R$  estimates, as imported cases are modelled as local.
5. The assumption of no pre-symptomatic infectivity is inconsistent with empirical estimates of the serial interval and generation time, reviewed in Anderson et al. (2020), for example.
6. Within hospital transmission is not modelled, although hospital-acquired infections has been reported to account for a quarter of cases at times in both waves<sup>2</sup>, and there is good evidence that the actual figure was higher (McKeigue et al., 2021). This will compromise the hospital data fit.
7. No interaction between NPIs and age is allowed, which is unlikely given the risk-by-age profiles.
8. Differential transmission rates between symptomatics and asymptomatics are not modelled.
9. The reported differences in disease progression between men and women are not modelled.
10. Changes in testing rates with capacity changes are not modelled.

## 2.1 The basic SEI(R) model

For concreteness we describe the core of the SEIR model, giving the equations for other compartments in supplementary A. Denoting the time derivative of a variable  $x$  by  $\dot{x}$ , then for the  $i$ th class,

$$\dot{S}^i = -\lambda_i(t)S^i \quad (1)$$

$$\dot{E}^{i,1} = \lambda_i(t)S^i - \gamma_E E^{i,1} \quad (2)$$

$$\dot{E}^{i,2} = \gamma_E E^{i,1} - \gamma_E E^{i,2} \quad (3)$$

$$\dot{I}_A^i = (1 - p_c)\gamma_E E^{i,2} - \gamma_A I_A^i + \mathbb{I}(2 < i < 13)\phi_{t_0, \sigma_t}(t) \quad (4)$$

$$\dot{I}_C^i = p_c\gamma_E E^{i,2} - \gamma_C I_C^i. \quad (5)$$

$\lambda_i(t)$  is the force of infection defined below, and is the only interesting interaction between age classes.  $p_c$  is the proportion of the infected showing symptoms, and the  $\gamma$  parameters determine between compartment flow rates, given in Knock et al. (2020).  $\mathbb{I}(\cdot)$  is an indicator function and  $\phi_{t_0, \sigma_t}$  is an  $N(t_0, \sigma_t^2)$  p.d.f. where  $t_0$  is a free parameter. This initialization differs slightly from Knock et al. who put 10 individuals in the age 15-20 asymptomatics at  $t_0$ . It is unclear why this is sensible, although it may slightly delay the first wave model care home epidemic. Susceptibles,  $S^i$ , are initialized from regional demography supplied in the Knock et al. supplementary material. Care home sizes are supplied in the `sircovid` package by the `carehomes_parameters()` function (Baguelin et al., 2021).

$R$  for this system can be computed according to Diekmann et al. (1990). See supplementary A.3. Our fitting also requires the derivatives of the model states with respect to the parameters: the *sensitivities*. These follow directly from the model specification. For example if  $S_{\theta_j}^i$  is the differential of  $S^i$  w.r.t.  $\theta_j$ ,

$$\dot{S}_{\theta_j}^i = -\frac{\partial \lambda_i}{\partial \theta_j} S^i - \lambda_i S_{\theta_j}^i.$$

<sup>2</sup><https://www.hsj.co.uk/patient-safety/covid-infections-caught-in-hospital-rise-by-a-third-in-one-week/7029211.article>

Generically each term in the model equation involving a state gets replaced by that state’s derivative w.r.t the parameter of interest, and to this are added any terms relating to direct dependence on the parameter of interest. For example, if  $\gamma_C$  was a free parameter then  $\dot{I}_{C\gamma_C}^i = p_c\gamma_E E_{\gamma_C}^{i,2} - \gamma_C I_{C\gamma_C}^i - I_C^i$ .

## 2.2 Force of Infection

Writing  $\mathbf{I}$  for the vector of infectious individuals in each class, then the model for the force of infection in each class is  $\lambda = \mathbf{M}\mathbf{I}$  where

$$\mathbf{M} = \begin{pmatrix} b(t)\mathbf{C} & b(t)\mathbf{c}^{\text{chw}} & \epsilon b(t)\mathbf{C}_{\cdot,16} \\ b(t)\mathbf{c}^{\text{chw}} & m_{\text{chw}} & m_{\text{chw}} \\ \epsilon b(t)\mathbf{C}_{16,\cdot} & m_{\text{chw}} & m_{\text{chr}} \end{pmatrix}.$$

$\epsilon$ ,  $m_{\text{chw}}$  and  $m_{\text{chr}}$  are free parameters.  $b(t)$  is a parameterized function of time controlling the variation of infection causing contact over time.  $\mathbf{C}$  is a symmetric matrix of contact rates and  $\mathbf{c}^{\text{chw}}$  a vector (derived from it for carehome workers).  $I_j$  is the sum of asymptomatic ( $I_A^j$ ) and symptomatic infectious ( $I_C^j$ ) in class  $j$ . Supplementary A.1 has the force of infection expressed so that sensitivities follow by inspection.

$\mathbf{C}$  is based on the POLYMOD survey (Mossong et al., 2017) accessed through the `socialmixr` R package (Funk, 2020). This had 1011 UK participants, who each recorded their contacts on one day. There were 7 participants in the 75-80 age group and none over 80. Supplementary A.2 gives details.

## 2.3 The likelihood

The likelihood is constructed from binomial components for the PCR and antibody test data (see supplementary A.9), and negative binomial components for the hospital death, care home death, hospital admissions, general ward occupancy and ICU occupancy data. For the negative binomial components Knock et al. (2020) set  $\kappa = \mu^2/(\sigma^2 - \mu)$  equal to 2 *in all cases* without justification offered. This is a huge level of overdispersion, heavily down-weighting the data relative to the priors. For example, for an expected death rate of 200 it raises the standard deviation from 14, for a Poisson deviate, to 140. It is hard to understand this choice, unless it was made to avoid particle depletion problems in filtering. Hospital deaths, for example, show no evidence of over-dispersion relative to Poisson. Still more problematic is the assumption that observed daily bed occupancy is given by a negative binomial deviate with expectation given by the model, *with these deviates independent between days*. We are at a loss to understand what mechanism could give rise to such a model. A reasonable model might have daily arrivals and discharges as independent random variables with means given by the model, but occupancy obviously integrates these arrival and discharge rates over days, leading to strong dependence between days. The stochastic version of the model might model some of this dependence, but leaves even less justification for additional independent negative binomial variability.

## 3 Modification of the Knock et al. model

In this section we present modifications of the Knock et al. model in order to deal with some of the deficiencies identified above. They consist of a number of corrections and minor modifications and, more fundamentally, relaxing some of the stronger modelling assumptions made by Knock et al.

### 3.1 Corrections and minor modifications

**Rates.** The  $\gamma$  parameters controlling rates of progression between model compartments are either taken from the literature, or are estimated from CHES<sup>3</sup> data that are not available for checking. There are at

<sup>3</sup>COVID-19 Hospitalisations in England Surveillance System

least two identifiable problems with the durations used in Knock et al. (2020). Firstly they set the mean duration of the  $E$  stage to 4.6 days citing Lauer et al. (2020). That paper actually reports a mean of 5.5 days, with 4.6 days lying just above the lower 95% confidence limit for the median. Here we used the mean of 5.8 days from the meta-analysis of McAloon et al. (2020), which includes Lauer et al. as one of the studies. In fact the most statistically careful analysis we found (Deng et al., 2020) gives an estimated mean incubation period of 9.1 days ( $n=1211$ ), and generation time of 5-6 days. Secondly Knock et al. assume that the *mean* time from symptoms to hospitalization is 4 days based on Docherty et al. (2020), but that paper gives 4 days as the *median*. An exponential distribution is used for time from symptoms to hospitalization (a model which the figures reported in Docherty et al. does seem to support), so the median is  $\log 2$  of the mean. Based on the male and female medians of 5 and 4 days reported in Docherty et al., we therefore used a mean time to hospitalization of 6.5 days.

Another issue is the assumption that 25% of patients were arriving at hospital with a test confirming their status from the start of the epidemic, despite the initial lack of infra-structure for this to happen. The assumption actually appears to make rather little difference to inference, but leaving it in place is clearly not quite right. Having no data to set up an evidence based alternative, we implemented the ad hoc modification of allowing the pre-tested rate,  $p^*$  to increase linearly from 0 to 0.25 between days 90 and 200, staying at 0.25 thereafter. While not ideal, this is arguably a less wrong assumption.

**Priors.** The priors used were not exactly those in Knock et al., rather priors were set to be vague on a working parameter scale. Any limits on parameter were set by the prior intervals reported in Knock et al.. Parameters were optimized on a working scale – either untransformed, log transformed or scaled logit transformed. Gaussian priors on the working scale were also applied, but except for  $t_0$  these were vague, and their only purpose was to allow ready detection of any parameters that were not identifiable. See supplementary A.8 for details.

### 3.2 The negative binomial likelihood

While our basic conclusions are in fact unchanged if we use the Knock et al. likelihood for the hospital occupancy data, we can see no valid justification for this part of the model formulation, and therefore replaced it with a defensible likelihood, based on the daily change in occupancy. In particular we model the ward (or ICU) arrivals and departures as independent overdispersed Poisson deviates, the difference in which gives the daily change in occupancy. A difficulty with applying this model directly is that hospital arrivals and discharges tend to have weekly pattern. This pattern shows up strongly in the ACFs and PACFs of occupancy first differences for some regions, especially east of England, but is absent from the model. We therefore base the likelihood on weekly changes. Since the changes in occupancy carry no information on the level of occupancy, we also add the sum of daily bed occupancies as a final datum to be fitted, treating this as close to Poisson (by setting  $\kappa$  to a very high constant). See supplementary A.9 for details.

For the total daily hospital admissions data and the care home deaths data we retain the negative binomial model, with the respective  $\kappa$  parameters free to be estimated. Some overdispersion here is a pragmatic way to deal with likely model mismatches in these components. For example, in addition to the mismatches expected from not modelling hospital acquired infections (e.g. McKeigue et al., 2021), it seems likely that there was some on the ground variability in the severity of disease sufficient for hospitalization, and in rates of discharge, particularly early in the epidemic and when loads were high. For the hospital deaths we set  $\kappa = 2000$ , which gives a likelihood very close to Poisson. There seems no legitimate reason to expect overdispersion here, if the model is at all fit for purpose.

### 3.3 Relaxing the model assumptions

The largest change made here is to relax the strong assumption that  $b(t)$ , which represents the effects of NPIs *and the weather*, is a piecewise linear function with slope changes only at 12 selected NPI change points. Here,  $b(t)$  is instead represented semi-parametrically by a logistic transform (see section A.8) of an adaptive smoothing spline, with 80 coefficients and 5 smoothing parameters, in which the degree of smoothness is allowed to vary smoothly with  $t$ . See section 5.3.5 of Wood (2017) for details. This relaxation substantially increases the role of data, relative to model assumptions, in inferences about the size and shape of  $b(t)$ . Of course it does nothing to remove the confounding of weather and NPI effects, but does avoid the implication that the weather changes in response to government announcements.

We also relaxed the assumption that all the  $\gamma$  parameters are fixed and known. Firstly, the reference used to justify the choice of  $\gamma_G$ , controlling the rate of progression of fatal disease in care homes, Bernabeu-Wittel et al. (2020), appears to contain no information on this parameter, so we allowed it to be a free parameter, which slightly reduces care home death mistiming. Secondly, the model also has difficulty matching the general ward and ICU occupancy data, tending to over-estimate both in the Midlands and two northern regions. To reduce this problem it seemed reasonable to relax the assumption that all the rate parameters controlling progression through the health system were fixed and known. In particular we relaxed the parameters for which there seemed likely to be most scope for some latitude in clinical judgement, perhaps driven by local circumstances, to make substantial differences. So we relaxed the assumptions on the rates related to movement of recovering patients through the system. That is  $\gamma_{IC_{W_r}}$ ,  $\gamma_{W_r}$  and  $\gamma_{H_r}$  were treated as free parameters.

A final rigidity in the model structure is that there is assumed to be no infection before individuals could at least potentially become symptomatic on leaving the  $E$  stage. At the same time the mean duration of the symptomatic infective stage is set equal to the mean time from symptom onset to hospitalisation. This makes for a very long generation time, much longer than the 5-7 days reported in the literature for the serial interval or generation time (see table 1 of Anderson et al., 2020, for a review). One consequence of this is that  $R$  estimates need to be higher than those usually quoted to meet the initial rate of increase in the disease (Knock et al., 2020, actually limit  $R$  in a way that avoids estimates being too high). To relax this link between clinical disease progression rates and the generation interval, we introduced an extra compartment between  $I_c$  and hospitalization (see the grey ‘P’ on figure 2).

$$\dot{P} = \gamma_C I_c - \gamma_{ph} P$$

where  $P$  replaces  $I_c$  in all flows into hospital compartments and the  $R$  state. By appropriate choice of  $\gamma_{ph}$ , this state allows us to shorten the  $E$  state and  $I_c$  state, hence reducing the generation time, without changing the literature based mean time from infection to hospitalisation. Specifically, we shortened the  $E$  state to have an average of 3 days to infectivity, and the  $I_c$  state to be 4 days, yielding a generation time of 6.2 days (accounting for the duration of  $I_A$ , which was unchanged).

### 3.4 Estimation and inference

The sensitivities of the model states with respect to the parameters were obtained for all 703 model state variables, yielding a system of 65379 sensitivity ODEs. Model and sensitivities were solved by fourth order Runge-Kutta integration (see e.g. Press et al., 2007) with a one day time step (having confirmed that halving the step made negligible difference to the evaluated likelihood). Hence the log likelihood and its derivatives w.r.t. the free parameters could be readily evaluated. Due to sparsity and cache efficiency, the sensitivity system less than doubles computing time for the model. Computing the likelihood, likelihood derivatives and  $R$  series for the full model takes less than a second on a single core of a low specification laptop — it is considerably faster for the original Knock et al. model with fewer free parameters.



Given the log likelihood and derivatives, the penalized log likelihood and derivatives are also readily evaluated, so the posterior modes of the free parameters can be obtained by quasi-Newton optimization. The smoothness of  $b(t)$  was controlled by a Gaussian smoothing prior, with 5 free smoothing parameters, which were estimated by the approximate marginal likelihood optimization method of Wood and Fasiolo (2017). Uncertainty was assessed using the large sample approximate posterior covariance matrix of the parameters, and the delta method. See supplementary A.10

## 4 Results

Figure 3 shows the fit of the model with the various assumption relaxations applied. The model fits imperfectly, with some systematic errors in the fit to hospital occupancy and arrival data as expected: without modelling the hospital acquired infection process and possible time variability in on-the-ground admission criteria it seems unlikely that better fits could be achieved. Given the ambitious nature of the fitting task, it seems reasonable to view the results as useful in the ‘all models are wrong, but some are useful’ sense.

Figures 4 and 5 shows the corresponding inferences about incidence and  $R$ . All regions have peak incidence prior to the first lockdown with total incidence for England in decline well before lockdown. The regional incidence picture is more mixed at the second lockdown, although the total is again falling well before lockdown. Furthermore all regions have  $R \lesssim 1$  by either lockdown, with average  $R < 1$  some days before either lockdown. Several regions relatively distant from London have the inferred  $R$  initially increasing. This is probably an artefact caused by the independent initialisation of each region, which cannot capture the initial region-to-region spread. As in Knock et al. (2020) the plotted uncertainties would be over-optimistic, even if we assumed a correct model structure, as they do not account for the uncertainty in most of the rate constants.

Although it could also be weather driven, the systematic pattern of  $R$  continuing to fall after the first lockdown and then increasing again well before restrictions were lifted, is to be expected.  $R$  is the average number of new infections *per existing infection*. Immediately after lockdown most infections are in the locked down population, with a low  $R$ , and only a minority are in the key worker population with higher  $R$  (assuming lockdown has an effect), so the average is low. After the locked down population runs out of household members to infect, the proportion of infections among key workers must increase, due to their higher  $R$ . So the average  $R$  must increase too. The piecewise linear  $b(t)$  of Knock et al. can not accommodate this effect.

Fits with the generation time relaxation (not shown), but not correcting the mean/median confusion in the time to hospitalization, shift the incidence peak and  $R = 1$  point a little later, but still before lockdown. The equivalent results without relaxation of the long generation time assumption gave similar incidence patterns, but with higher initial  $R$ , as expected. With this long generation time, the average  $R$  for England dropped below one at about lockdown: before or after can not be determined with statistical confidence. Not correcting the occupancy data likelihood again gives results very similar to figures 3 - 5.

If fits of this model to data are viewed as a reasonable basis for inference about the timing of incidence and  $R$  levels, then the implication is that  $R < 1$  probably occurred some time before both the first two English lockdowns, and that incidence was already in sharp decline before either. The contrary result of Knock et al. relies on a very restrictive model for  $b(t)$ , a much longer generation time than is reported in the literature, and the inappropriate use of an exponential median as an exponential mean.

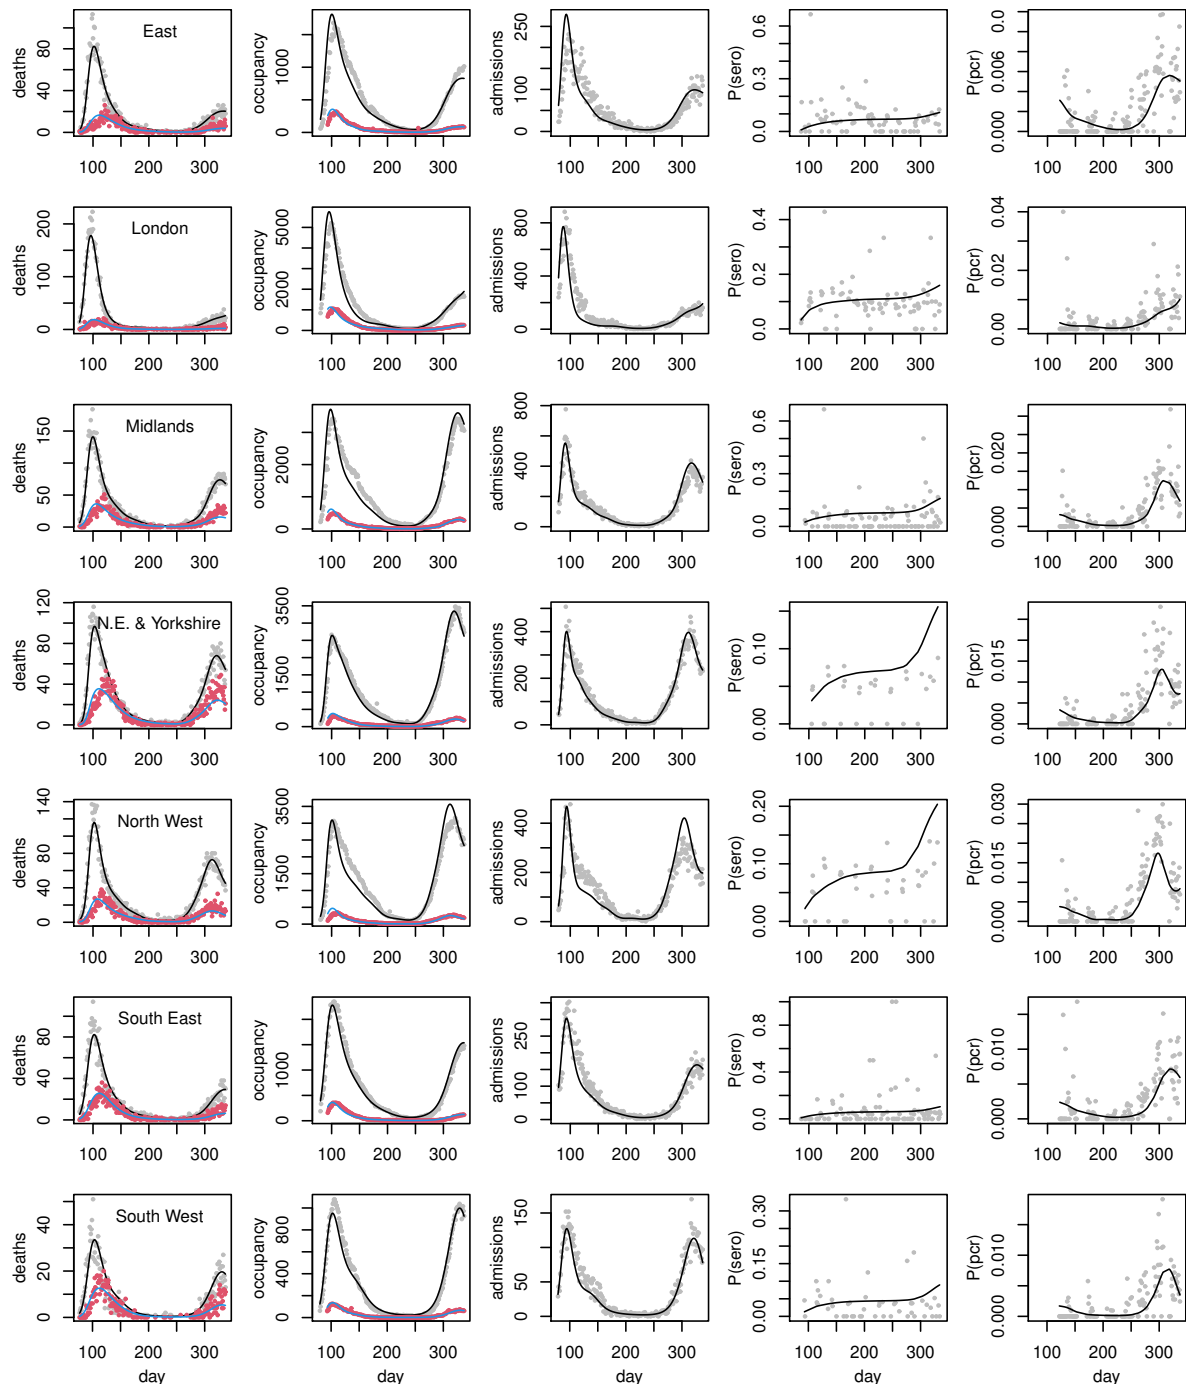


Figure 3: Model fits (posterior modes), to the death, hospital and testing data, with one region per row, against day of year. In the leftmost ‘deaths’ column, grey points are hospital deaths and red points care home deaths. In the second ‘occupancy’ column, grey is general ward occupancy and red ICU occupancy. Note some substantial discrepancies in the two northern regions.




Figure 4: Inferred incidence, for all regions (coloured) and whole of England (black). Notional 95% credible bands are shown: these do not reflect all the uncertainty in rate parameters and assume a correct model structure. As such, they provide a lower bound on uncertainty. Vertical dashed lines show times of some policy changes. ‘Eat out to help out’ was a scheme encouraging people to use the restaurants and pubs. The re-opening of schools after the first lockdown is also shown. Subsequent policies introduce increasing levels of restriction.

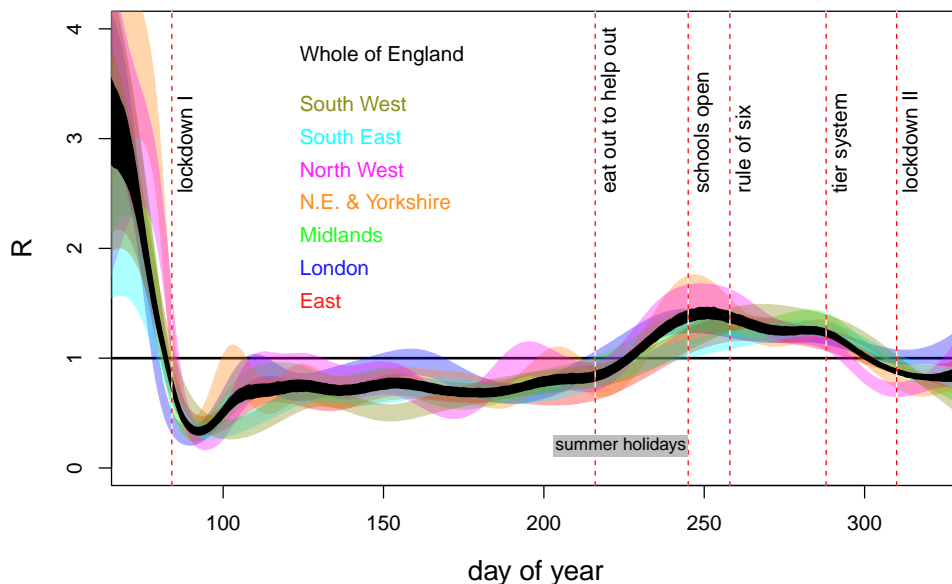


Figure 5: Inferred  $R$ , for all regions (coloured) and whole of England (black). Notional 95% credible bands are shown: these do not reflect all the uncertainty in rate parameters and assume a correct model structure. As such, they provide a lower bound on uncertainty. Vertical dashed lines as Figure 4.

## 5 Discussion

Knock et al. (2020) make three major claims in their paper. Whereas the first is of a descriptive nature, namely that the two English Covid-19 lockdowns in March and November 2020 coincide with a major drop in the reproduction rate of Covid-19 in the UK, the other two are of a so-called “counterfactual” nature: (i) if England had not gone into lockdown, then there would have not been an associated drop in reproduction rate and (ii) if England had gone down into lockdown earlier (or later) then a lot of lives would have been saved (or lost, respectively).

The key challenge is that a counterfactual cannot be directly observed and must be approximated with reference to a comparison group. There are various accepted approaches to determining an appropriate comparison group for counterfactual analysis, ideally using a prospective design. When this is not available, such as in this case, a retrospective approach is necessary. But there are stringent conditions on a retrospective design in order for it to have counterfactual validity, such as avoiding confounding, contamination, and impact heterogeneity (see Pearl et al., 2016, for an introductory treatment). Confounding occurs where certain factors, for example the various social distancing measures in place prior to the lockdowns, are correlated with exposure to the intervention and, independent of exposure, are causally related to the outcome of interest. Confounding factors are therefore alternate explanations for an observed, but possibly spurious, relationship between intervention and the outcome; in this case between lockdown and the reduction in  $R$ . The pre-lockdown social distancing measures are also an example of contamination, which may also invalidate any counter-factual statements. Contamination occurs when members of treatment group (i.e. the actual population) and/or comparison groups (i.e. the counterfactual populations) have access to another intervention which also affects the outcome of interest. Additionally, there is the issue of impact heterogeneity: the impact of the lockdown will be very different in the locked down subset of the population, compared to key workers, who are less restricted. Finally, Knock et al. explicitly state that  $b(t)$  is modelling both the effects of NPIs and the weather. There is therefore no basis on which the model can identify the effect of lockdown independent of the weather, enabling the counterfactual manipulation of one while appropriately controlling the other. But such control is absolutely fundamental to causal reasoning with counterfactuals. We conclude that the model and inference of Knock et al. do not form any sort of reasonable basis for making counterfactual statements about how many people would have died if lockdown had occurred at a different time.

Our results on the timing of  $R < 1$  and peak incidence obviously do not imply that the lockdowns had no effect. Indeed the dip and recovery seen in  $R$  after the first lockdown is *only* expected if lockdown reduces spread in the locked down population, relative to those not locked down. The point is rather that the additional effect, on top of the cumulative effects of other behavioural changes pre-dating lockdown, seems likely to have been greatly overstated. In our view, determining definitively what caused  $R$  to drop below one is not possible. In March especially, policy and behavioural changes were so rapid (public information campaign, March 4th; symptomatic self isolation 13th; work from home advice, 16th; school and hospitality closures 20th; full lockdown, 24th) that there would simply have been insufficient time to determine what had worked, even if adequate data had been gathered to answer this question. In fact, there was no surveillance testing at that point. However, it seems difficult to make the case that full lockdowns were necessary to bring  $R$  below one, whether region-by-region or in aggregate for England.

## References

Anderson, R., C. Donnelly, D. Hollingsworth, M. Keeling, C. Vegvari, R. Baggaley, and R. Madsen (2020). Reproduction number ( $R$ ) and growth rate ( $r$ ) of the COVID-19 epidemic in the UK: methods of estimation, data sources, causes of heterogeneity, and use as a guide in policy formulation. Technical report, Royal Society SET-C.

- Baguelin, M., E. Knock, L. Whittles, R. FitzJohn, J. Lees, and A. Cori (2021). *sircovid: SIR Model for COVID-19*. <https://github.com/mrc-ide/sircovid>, <https://mrc-ide.github.io>.
- Bernabeu-Wittel, M., J. Ternero-Vega, P. Díaz-Jiménez, C. Conde-Guzmán, M. Nieto-Martín, L. Moreno-Gaviño, J. Delgado-Cuesta, M. Rincón-Gómez, L. Giménez-Miranda, M. Navarro-Amuedo, et al. (2020). Death risk stratification in elderly patients with covid-19. a comparative cohort study in nursing homes outbreaks. *Archives of gerontology and geriatrics* 91, 104240.
- Deng, Y., C. You, Y. Liu, J. Qin, and X.-H. Zhou (2020). Estimation of incubation period and generation time based on observed length-biased epidemic cohort with censoring for covid-19 outbreak in china. *Biometrics*.
- Dennis, J. M., A. P. McGovern, S. J. Vollmer, and B. A. Mateen (2021). Improving survival of critical care patients with coronavirus disease 2019 in england: a national cohort study, March to June 2020. *Critical care medicine* 49(2), 209.
- Diekmann, O., J. A. P. Heesterbeek, and J. A. Metz (1990). On the definition and the computation of the basic reproduction ratio  $R_0$  in models for infectious diseases in heterogeneous populations. *Journal of mathematical biology* 28(4), 365–382.
- Docherty, A. B., E. M. Harrison, C. A. Green, H. E. Hardwick, R. Pius, L. Norman, K. A. Holden, J. M. Read, F. Dondelinger, G. Carson, et al. (2020). Features of 20 133 UK patients in hospital with covid-19 using the ISARIC WHO Clinical Characterisation Protocol: prospective observational cohort study. *BMJ* 369.
- Flaxman, S., S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, et al. (2020). Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* 584(7820), 257–261.
- Funk, S. (2020). *socialmixr: Social Mixing Matrices for Infectious Disease Modelling*. R package version 0.1.8.
- Gillespie, D. T. (2001). Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of chemical physics* 115(4), 1716–1733.
- Knock, E. S., L. K. Whittles, J. A. Lees, P. N. Perez Guzman, R. Verity, R. G. Fitzjohn, K. A. M. Gaythorpe, N. Imai, W. Hinsley, L. C. Okell, A. Rosello, N. Kantas, C. E. Walters, S. Bhatia, O. J. Watson, C. Whittaker, L. Cattarino, A. Boonyasiri, B. A. Djaafara, K. Fraser, H. Fu, H. Wang, X. Xi, C. A. Donnelly, E. Jauneijaite, D. J. Laydon, P. J. White, A. C. Ghani, N. M. Ferguson, A. Cori, and M. Baguelin (2020). Report 41: The 2020 SARS-CoV-2 epidemic in England: key epidemiological drivers and impact of interventions. *Imperial College London*.
- Lauer, S. A., K. H. Grantz, Q. Bi, F. K. Jones, Q. Zheng, H. R. Meredith, A. S. Azman, N. G. Reich, and J. Lessler (2020). The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. *Annals of internal medicine* 172(9), 577–582.
- Marmot, M., J. Allen, T. Boyce, P. Goldblatt, and J. Morrison (2020). *Health Equity in England: The Marmot Review 10 Years On*. The Health Foundation.
- McAloon, C., Á. Collins, K. Hunt, A. Barber, A. W. Byrne, F. Butler, M. Casey, J. Griffin, E. Lane, D. McEvoy, et al. (2020). Incubation period of COVID-19: a rapid systematic review and meta-analysis of observational research. *BMJ open* 10(8), e039652.

- McKeigue, P. M., D. McAllister, D. Caldwell, C. Gribben, J. Bishop, S. J. McGurnaghan, M. Armstrong, J. Delvaux, S. Colville, S. Hutchinson, et al. (2021). Relation of severe covid-19 in scotland to transmission-related factors and risk conditions eligible for shielding support: React-scot case-control study. *medRxiv*.
- Mossong, J., N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, M. Massari, S. Salmaso, G. S. Tomba, J. Wallinga, J. Heijne, M. Sadkowska-Todys, M. Rosinska, and W. J. Edmunds (2017, November). Polymod social contact data.
- Nocedal, J. and S. Wright (2006). *Numerical Optimization* (2nd ed.). New York: Springer Verlag.
- Pearl, J., M. Glymour, and N. P. Jewell (2016). *Causal inference in statistics: A primer*. John Wiley & Sons.
- Press, W., S. Teukolsky, W. Vetterling, and B. Flannery (2007). *Numerical Recipes* (3rd ed.). Cambridge: Cambridge University Press.
- RECOVERY Collaborative Group (2020). Dexamethasone in hospitalized patients with covid-19—preliminary report. *New England Journal of Medicine*.
- Wood, S. N. (2017). *Generalized Additive Models: An Introduction with R* (2 ed.). Boca Raton, FL: CRC press.
- Wood, S. N. and M. Fasiolo (2017). A generalized Fellner-Schall method for smoothing parameter optimization with application to Tweedie location, scale and shape models. *Biometrics* 73(4), 1071–1081.

## A Supplementary Appendix: full model and inference method details

This appendix gives full details of the model structure and computations not covered in the main paper, including the model compartments for care homes, hospitals and testing, as well as the output variables predicting the data, the specification of the likelihood terms, and further details of the inferential methods.

### A.1 Force of infection

Here is the force of infection term for for classes 0 to 16, written out in a form where the sensitivities follow easily.

$$\lambda_i(t) = b(t) \left\{ \sum_{j=0}^{16} C_{i,j} I_j + c_i^{\text{chw}} I_{17} + \epsilon C_{i,16} I_{18} \right\}$$

For care home workers, class 17,

$$\lambda_{17}(t) = b(t) \sum_{j=0}^{16} C_{i,j} I_j + m_{\text{chw}} I_{17} + m_{\text{chw}} I_{18}$$

and for care home residents, class 18,

$$\lambda_{18}(t) = \epsilon b(t) \sum_{j=0}^{16} C_{i,16} I_j + m_{\text{chw}} I_{17} + m_{\text{chr}} I_{18}$$

The sensitivities require the derivatives of these terms w.r.t. the parameters  $\theta$  (i.e.  $\epsilon$ ,  $m_{\text{chw}}$ ,  $m_{\text{chr}}$  and the parameters of  $b(t)$  - sensitivities to further parameters are obviously zero).

### A.2 The contact matrix, C

**C** is based on the POLYMOD survey (Mossong et al., 2017) accessed through R package `socialmixr` (Funk, 2020), which had 1011 participants in the UK, who each recorded their contacts on one day. There were 7 participants in the 75-80 age group and none over 80. Knock et al. (2020) are vague about exactly how  $C_{ij}$  is obtained, but given the statement that it is symmetric there is really only one sane option.

Let  $A_{ij}$  denote the average number of contacts of someone in age class  $i$  with someone in age class  $j$ . Because population sizes may differ between age classes,  $A_{ij} \neq A_{ji}$  in general. Let  $S_i^0$  denote the population (initially susceptible) in class  $i$ . The total number of contacts between members of class  $i$  and class  $j$  is  $T_{ij} = A_{ij} S_i^0$ . Obviously for total contacts  $T_{ij} = T_{ji}$ . In reality the polymod estimates of  $T_{ij}$  and  $T_{ji}$  will usually be different because the study population is not closed and there are likely to be some recording errors. It therefore makes sense to replace both by their average.

The number of infections generated in class  $i$  by class  $j$  is proportional to the total contacts between  $S_i$  and  $I_j$ . That is the total number of contacts between the classes, multiplied by the proportion that are between  $S_i$  and  $I_j$ ,

$$T_{ij} \frac{S_i}{S_i^0} \frac{I_j}{S_j^0} = C_{ij} S_i I_j$$

by definition of  $C_{ij}$ . The contacts for care home workers with the general population are set to the mean of the  $C_{ij}$  for the ages 20-65.

In practice there is no data on 80+ with 80+ contacts. The preceding age class has 10% of contacts within age class, so this proportion could be assumed to apply to the 80+ also, and was the assumption made here. Knock et al. (2020) do not document what they did. The `sircovid` package accompanying Knock et al. (2020) has a `carehomes_parameters()` function. Among other things this returns a

matrix  $m$ , which is *exactly* the  $\mathbf{C}$  matrix defined above, up until age 70, but thereafter something undocumented appears to have been done and it is not clear if this was the  $\mathbf{C}$  used or not. In sensitivity testing it made negligible difference whether we used the `sircovid` matrix or the version just described.

### A.3 Computing $R$

Knock et al. (2020) do not describe exactly what was done to compute  $R$ , but section 4.1 of Diekmann et al. (1990) provides what is required. For the current model we conceptually divide the population into those who will show symptoms and those who will not, so that we have  $E_A$  and  $E_C$  (exposed eventually asymptomatic and exposed eventually symptomatic). This division does not change the dynamics. Let  $\mathbf{S}$  denote the vector of susceptibles in each class. We define the matrix

$$\mathbf{K} = \begin{pmatrix} (1 - p_c)\text{diag}(\mathbf{S})\mathbf{C}/\gamma_A & (1 - p_c)\text{diag}(\mathbf{S})\mathbf{C}/\gamma_C \\ p_c\text{diag}(\mathbf{S})\mathbf{C}/\gamma_A & p_c\text{diag}(\mathbf{S})\mathbf{C}/\gamma_C \end{pmatrix}.$$

$R$  is the dominant eigenvalue of  $\mathbf{K}$ . Checking by simulation, the epidemic does indeed decline when  $R < 1$  and increase when  $R > 1$ . Note that except when  $R = 1$ , anything that lengthens the assumed generation time requires a corresponding increase(decrease) in  $R$  in order to achieve the same epidemic growth(shrinkage) rate in time. The generation times assumed by Knock et al. (2020) are fairly long, and get longer after the rate parameter corrections described in the main paper.

To average  $R$  from different regions we need to weight by the total number of infectives in each region, since conceptually  $R$  is the mean number of new infections caused by each infective.

### A.4 Care home dynamics

The first additional compartments relate to care home deaths. The equations are given for each age class, but actually only apply to class  $i = 18$ , care home residents.

$$\dot{G}_D^{i,1} = p_H^{max}\psi_H^i p_{G_D}^i \gamma_C I_C^i - \gamma_{G_D} G_D^{i,1} \quad (6)$$

$$\dot{G}_D^{i,2} = \gamma_{G_D} G_D^{i,1} - \gamma_{G_D} G_D^{i,2} \quad (7)$$

Because this only applies in a single class there is only one  $p_{G_D}$  parameter to estimate.  $\psi_H^i$  is given in the Knock et al. supplementary material as `psi_hosp_symp`.

### A.5 Hospital dynamics

Now consider the ICU flows ( $i = 0 \dots 17$ ). Define

$$r(t, \mu) = \begin{cases} 1 & t < t_0 \\ 1 + (1 - \mu)(t_0 - t)/(t_1 - t_0) & t_0 \leq t \leq t_1 \\ \mu & t > t_1 \end{cases}$$

$t_0$  is set to 1 April and  $t_1$  to 1 June. The reference given to justify this term is a report on a clinical trial that showed modest improvements, but did not finish until July (RECOVERY Collaborative Group, 2020). It is unclear how improvements developed during the trial could have wide-spread consequences before its completion. However Dennis et al. (2021), do report improvements in mortality rates of hospitalized patients in England over this period, and show that these improvements are not the result of changes in patient characteristics, although of course some component of the change may relate to on the ground changes in the severity of disease required for admission, particularly between peak hospital load and later.



In what follows starred states are for cases where Covid-19 has been confirmed by test. The ICU compartments are governed by the following ODEs - here written out in a form such that the sensitivity equations follow by inspection.  $ICU_{pre}^i$  is the compartment preceding admission to ICU,  $ICU_{W_r}^i$  is the compartment for patients in the ICU who will eventually recover, back on the ward,  $ICU_{W_D}^i$  is for ICU patients who will eventually die back on the general ward,  $ICU_D$  compartments are for patients who die in the ICU. Parameters  $\mu_{IC}$ ,  $\mu_D$ ,  $p_H^{max}$ ,  $p_{IC}^{max}$ ,  $p_{IC_D}^{max}$  and  $p_{W_d}^{max}$  are free, and we additionally treated  $\gamma_{IC_{W_r}}$  as free.

$$\dot{ICU}_{pre}^i = p_H^{max} \psi_H^i (1 - p^*) p_{IC}^{max} \psi_{IC}^i r(t, \mu_{IC}) \gamma_{IC} I_C^i - (\gamma_{IC_{pre}} + \gamma_U) ICU_{pre}^i \quad (8)$$

$$\dot{ICU}_{pre*}^i = p_H^{max} \psi_H^i p^* p_{IC}^{max} \psi_{IC}^i r(t, \mu_{IC}) \gamma_{IC} I_C^i - \gamma_{IC}^{pre} ICU_{pre*}^i + \gamma_U ICU_{pre}^i \quad (9)$$

$$\dot{ICU}_{W_r}^i = (1 - p_{IC_D}^{max} \psi_{IC_D}^i r(t, \mu_D)) (1 - p_{W_d}^{max} \psi_{W_D}^i r(t, \mu_D)) \gamma_{IC}^{pre} ICU_{pre}^i - (\gamma_{IC_{W_r}} + \gamma_U) ICU_{W_r}^i \quad (10)$$

$$\dot{ICU}_{W_r*}^i = (1 - p_{IC_D}^{max} \psi_{IC_D}^i r(t, \mu_D)) (1 - p_{W_d}^{max} \psi_{W_D}^i r(t, \mu_D)) \gamma_{IC}^{pre} ICU_{pre*}^i - \gamma_{IC_{W_r}} ICU_{W_r*}^i + \gamma_U ICU_{W_r}^i \quad (11)$$

$$\dot{ICU}_{W_D}^i = (1 - p_{IC_D}^{max} \psi_{IC_D}^i r(t, \mu_D)) p_{W_d}^{max} \psi_{W_D}^i r(t, \mu_D) \gamma_{IC}^{pre} ICU_{pre}^i - (\gamma_{IC_{W_D}} + \gamma_U) ICU_{W_D}^i \quad (12)$$

$$\dot{ICU}_{W_D*}^i = (1 - p_{IC_D}^{max} \psi_{IC_D}^i r(t, \mu_D)) p_{W_d}^{max} \psi_{W_D}^i r(t, \mu_D) \gamma_{IC}^{pre} ICU_{pre*}^i - \gamma_{IC_{W_D}} ICU_{W_D*}^i + \gamma_U ICU_{W_D}^i \quad (13)$$

$$\dot{ICU}_D^{i,1} = p_{IC_D}^{max} \psi_{IC_D}^i r(t, \mu_D) \gamma_{IC}^{pre} ICU_{pre}^i - (\gamma_{IC_D} + \gamma_U) ICU_D^{i,1} \quad (14)$$

$$\dot{ICU}_D^{i,2} = \gamma_{IC_D} ICU_D^{i,1} - (\gamma_{IC_D} + \gamma_U) ICU_D^{i,2} \quad (15)$$

$$\dot{ICU}_{D*}^{i,1} = p_{IC_D}^{max} \psi_{IC_D}^i r(t, \mu_D) \gamma_{IC}^{pre} ICU_{pre*}^i - \gamma_{IC_D} ICU_{D*}^{i,1} + \gamma_U ICU_D^{i,1} \quad (16)$$

$$\dot{ICU}_{D*}^{i,2} = \gamma_{IC_D} ICU_{D*}^{i,1} - \gamma_{IC_D} ICU_{D*}^{i,2} + \gamma_U ICU_D^{i,2} \quad (17)$$

The next 6 compartments are the step down from ICU to regular ward compartments. They originally involved no free parameters, but we treated  $\gamma_{W_r}$  as free. Subscripts  $r$  and  $D$  refer to recovery or death.

$$\dot{W}_r^{i,1} = \gamma_{IC_{W_r}} ICU_{W_r}^i - (\gamma_{W_r} + \gamma_U) W_r^{i,1} \quad (18)$$

$$\dot{W}_r^{i,2} = \gamma_{W_r} W_r^{i,1} - (\gamma_{W_r} + \gamma_U) W_r^{i,2} \quad (19)$$

$$\dot{W}_{r*}^{i,1} = \gamma_{IC_{W_r}} ICU_{W_r*}^i - \gamma_{W_r} W_{r*}^{i,1} + \gamma_U W_r^{i,1} \quad (20)$$

$$\dot{W}_{r*}^{i,2} = \gamma_{W_r} W_{r*}^{i,1} - \gamma_{W_r} W_{r*}^{i,2} + \gamma_U W_r^{i,2} \quad (21)$$

$$\dot{W}_D^i = \gamma_{IC_{W_d}} ICU_{W_d}^i - (\gamma_{W_d} + \gamma_U) W_D^i \quad (22)$$

$$\dot{W}_{D*}^i = \gamma_{IC_{W_d}} ICU_{W_d*}^i - \gamma_{W_d} W_{D*}^i + \gamma_U W_D^i \quad (23)$$

Six further compartments are for the hospital general ward, and again depend on free parameters, including new one  $p_{H_d}^{max}$ . We also treated  $\gamma_{H_r}$  as free. Again subscripts  $r$  and  $D$  refer to recovery or death.

$$\dot{H}_r^i = p_H^{max} \psi_H^i (1 - p^*) (1 - p_{IC}^{max} \psi_{IC}^i r(t, \mu_{IC})) (1 - p_{H_d}^{max} \psi_{H_d}^i r(t, \mu_D)) \gamma_{IC} I_C^i - (\gamma_{H_r} + \gamma_U) H_r^i \quad (24)$$

$$\dot{H}_{r*}^i = p_H^{max} \psi_H^i p^* (1 - p_{IC}^{max} \psi_{IC}^i r(t, \mu_{IC})) (1 - p_{H_d}^{max} \psi_{H_d}^i r(t, \mu_D)) \gamma_{IC} I_C^i + \gamma_U H_r^i - \gamma_{H_r} H_{r*}^i \quad (25)$$

$$\dot{H}_D^{i,1} = p_H^{max} \psi_H^i (1 - p^*) (1 - p_{IC}^{max} \psi_{IC}^i r(t, \mu_{IC})) p_{H_d}^{max} \psi_{H_d}^i r(t, \mu_D) \gamma_{IC} I_C^i - (\gamma_{H_D} + \gamma_U) H_D^{i,1} \quad (26)$$

$$\dot{H}_D^{i,2} = \gamma_{H_D} H_D^{i,1} - (\gamma_{H_D} + \gamma_U) H_D^{i,2} \quad (27)$$

$$\dot{H}_{D*}^{i,1} = p_H^{max} \psi_H^i p^* (1 - p_{IC}^{max} \psi_{IC}^i r(t, \mu_{IC})) p_{H_d}^{max} \psi_{H_d}^i r(t, \mu_D) \gamma_{IC} I_C^i + \gamma_U H_D^{i,1} - \gamma_{H_d} H_{D*}^{i,1} \quad (28)$$

$$\dot{H}_{D*}^{i,2} = \gamma_{H_d} H_{D*}^{i,1} - \gamma_{H_d} H_{D*}^{i,2} + \gamma_U H_D^{i,2} \quad (29)$$

The recovered compartment is governed by

$$\dot{R}^i = \gamma_A I_A^i + (1 - p_H^{max} \psi_H^i) \gamma_{IC} I_C^i + \gamma_{H_r} (H_r^i + H_{r*}^i) + \gamma_{W_r} (W_r^{i,2} + W_{r*}^{i,2}) \quad (30)$$

(this is corrected from Knock et al., 2020).

## A.6 Testing compartments

Finally there are compartments used to determine the proportions testing positive in randomized testing. They have no free parameters (S subscript for antibodies, P for PCR).

$$\dot{T}_S^i = \gamma_E E^{i,2} - \gamma_S T_S^i \quad (31)$$

$$\dot{T}_{S+} = p_{S+} \gamma_S T_S^i \quad (32)$$

$$\dot{T}_{S-} = (1 - p_{S+}) \gamma_S T_S^i \quad (33)$$

$$\dot{T}_P^i = \lambda_i S^i - \gamma_P T_P^i \quad (34)$$

$$\dot{T}_{P+}^i = \gamma_P T_P^i - \gamma_{P+} T_{P+}^i \quad (35)$$

$$\dot{T}_{P-}^i = \gamma_{P+} T_{P+}^i \quad (36)$$

(corrected from Knock et al., 2020).

## A.7 Model outputs

The outputs required from all this are as follows, again written in a form in which the sensitivities are immediate. Firstly the rate of new Covid cases in hospital, which is the sum of the flows into the  $H_{r*}$ ,  $H_{D*}$  and  $ICU_{pre*}$  states plus all the rate  $\gamma_U$  flows.

$$X_{adm} = \sum_i p_H^{max} \psi_{HP}^i \gamma_C I_C^i + \gamma_U (ICU_{pre}^i + ICU_{W_r}^i + ICU_{W_D}^i + ICU_D^{i,1} + ICU_D^{i,2} + W_r^{i,1} + W_r^{i,2} + W_D^i + H_r^i + H_D^{i,1} + H_D^{i,2}) \quad (37)$$

Note that Knock et al. (2020) modify some of the  $\gamma_U$  plumbing between the ODE statement of the model and its stochastic discretisation, but this does not change the flow totals.

The hospital regular bed occupancy is given by (correcting report notation mutation)

$$X_{hos} = \sum_i H_{r*}^i + H_{D*}^{i,1} + H_{D*}^{i,2} + ICU_{pre*}^i + W_{D*}^i + W_{r*}^{i,1} + W_{r*}^{i,2} \quad (38)$$

But a correctly specified likelihood actually requires regular bed arrivals

$$\sum_i p_H^{max} \psi_{HP}^i \gamma_C I_C^i + \gamma_U ICU_{pre}^i + \gamma_{IC_{W_r}} ICU_{W_{r*}}^i + \gamma_U W_r^{i,1} + \gamma_U W_r^{i,2} + \gamma_{IC_{W_d}} ICU_{W_{d*}}^i + \gamma_U W_D^i + \gamma_U H_r^i + \gamma_U H_D^{i,1} + \gamma_U H_D^{i,2} \quad (39)$$

and departures

$$\sum_i \gamma_{IC}^{pre} ICU_{pre*}^i + \gamma_{W_r} W_{r*}^{i,2} + \gamma_{W_d} W_{D*}^i + \gamma_{H_r} H_{r*}^i + \gamma_{H_d} H_{D*}^{i,2} \quad (40)$$

ICU occupancy is

$$X_{icu} = \sum_i ICU_{W_{r*}}^i + ICU_{W_{D*}}^i + ICU_{D*}^{i,1} + ICU_{D*}^{i,2}, \quad (41)$$

but again for a defensible likelihood we need arrivals

$$\sum_i \gamma_{IC}^{pre} ICU_{pre*}^i + \gamma_U ICU_{W_r}^i + \gamma_U ICU_{W_D}^i + \gamma_U ICU_D^{i,1} + \gamma_U ICU_D^{i,2} \quad (42)$$

and departures

$$\sum_i \gamma_{IC_{W_r}} ICU_{W_{r*}}^i + \gamma_{IC_{W_D}} ICU_{W_{D*}}^i + \gamma_{IC_D} ICU_{D*}^{i,2} \quad (43)$$

The death rate in hospital is

$$X_{HD} = \sum_i \gamma_{ICD} (ICU_D^{i,2} + ICU_{D*}^{i,2}) + \gamma_{Wd} (W_D^i + W_{D*}^i) + \gamma_{Hd} (H_D^{i,2} + H_{D*}^{i,2}) \quad (44)$$

The care home death rate is given by

$$X_{GD} = \gamma_{GD} G_D^{i,2} \quad (45)$$

The testing data require

$$X_{S+} = \sum_{i=3}^{12} T_{S+}^i \quad (46)$$

(summation over ages 15 to 65) and

$$X_{R1+} = \sum_{i=1}^{17} T_{P+}^i \quad (47)$$

(summation from age 5 upwards and care home workers, but not residents). In the Knock et al. there is one further stream for the Pillar 2 PCR data, but the model seems so crude that this stream can really only undermine inference and is better omitted.

## A.8 Priors

Parameters were optimized on a working scale, so that the dynamic model parameters were  $\theta_j = h_k(\theta'_j)$  where  $\theta'_j$  was unconstrained and  $k$  selects among 3 alternatives.  $h_1$  was the identity,  $h_2$  the exponential and  $h_3 = a_1 + (a_2 - a_1)e^{\theta'_j}/(1 + e^{\theta'_j})$ , which constrains the  $\theta_j$  to the interval  $(a_1, a_2)$ . The limits were generally set from the prior intervals given in Knock et al. (2020). Gaussian priors on the working scale were also applied, but except for  $t_0$  these were vague, and their only purpose was to allow ready detection of any parameters that were not identifiable.

	$\epsilon$	$m_{chw}$	$m_{chr}$	$t_0$	$p_H^{max}$	$p_G$	$\mu_{IC}$	$p_{IC}^{max}$	$p_{ICD}^{max}$	$\mu_D$	$p_{WD}^{max}$	$p_{HD}^{max}$
$h$	2	2	2	1	3	3	3	3	3	3	3	3
$a_1$					0	0	.66	.14	.57	.5	.25	.36
$a_2$					1	1	1	.36	.77	1	.46	.56
$\mu_p$	-.7	-12.7	-12.7	30	1.1	1.1	-1	-.2	0	-1	0	0
$\sigma_p$	2	2	2	10	2	2	2	2	2	2	1	1

$h_2$  was used for  $\gamma_G, \gamma_{IC_{W_r}}, \gamma_{W_r}$  and  $\gamma_{H_r}$ , with working scale mean vector  $\mu_p = (-1, -2.75, -1.8, -2.37)$  and  $\sigma_p = 2$  in all cases. The means were set from the Knock et al. (2020) estimates. The  $h_3$  applied to the adaptive spline to yield  $b(t)$  had  $a_1 = 0$  and  $a_2 = 0.1$ .

## A.9 Likelihood and data

The likelihood is constructed from negative binomial and binomial components. The binomial log likelihood is used for the serology and REACT PCR data.

$$l_b(p) = y \log p + (n - y) \log(1 - p) + \log n! - \log y! - \log(n - y)!$$

so

$$\frac{\partial l_b}{\partial p} = \frac{y}{p} - \frac{n - y}{1 - p}.$$

The negative binomial log likelihood is used for admissions, general ward and ICU occupancy and deaths.

$$l(\mu, \kappa) = \kappa \log \left( \frac{\kappa}{\kappa + \mu} \right) + y \log \left( \frac{\mu}{\mu + \kappa} \right) + \log \Gamma(\kappa + y) - \log \Gamma(\kappa) - \log y!$$

and so

$$\frac{\partial l}{\partial \mu} = \frac{y}{\mu} - \frac{y + \kappa}{\mu + \kappa}.$$

As discussed in the main paper, the likelihood based on independent negative binomial deviates is not justifiable for occupancy, and an alternative based on changes in occupancy, which can be modelled as independent, is more appropriate. If we model the ward (or ICU) arrivals as Poisson, and the departures as Poisson, then this daily change will follow a Skellam distribution, but that allows for no overdispersion. The difference between negative binomials (with common  $\kappa$  parameter) is a skewed generalized discrete Laplace distribution, but is computationally awkward. We therefore model ward and ICU arrivals and departures using overdispersed versions of the normal approximation to the Poisson,  $N(\mu_1, k\mu_1)$  and  $N(\mu_2, k\mu_2)$ , with difference  $N(\mu_2 - \mu_1, k(\mu_1 + \mu_2))$ . Let  $\sigma_0 = \mu_1 + \mu_2$  and  $\alpha = (y - \mu_1 + \mu_2)$ , where  $y$  is now the observed change in occupancy, then

$$l_d(\mu_1, \mu_2) = -\frac{\alpha^2}{2k\sigma_0} - \log(k\sigma_0)/2 - \log(2\pi)/2$$

and so

$$\frac{\partial l_d}{\partial \mu_1} = \frac{\alpha}{k\sigma_0} + \frac{\alpha^2}{2k\sigma_0^2} - \frac{1}{2\sigma_0}, \quad \frac{\partial l_d}{\partial \mu_2} = -\frac{\alpha}{k\sigma_0} + \frac{\alpha^2}{2k\sigma_0^2} - \frac{1}{2\sigma_0}.$$

A difficulty with applying this model directly is that hospital arrivals and discharges tend to have weekly pattern. This pattern shows up strongly in the ACFs and PACFs of occupancy first differences for some regions, especially East of England, but is absent from the model. We therefore base the likelihood on weekly changes. Since the changes in occupancy carry no information on the level of occupancy, we also add the sum of daily bed occupancies as a final datum to be fitted, treating this as close to Poisson (by setting  $\kappa$  to a very high constant).

For the total daily hospital admissions data and the care home deaths data we retain the negative binomial model, with the respective  $\kappa$  parameters free to be estimated. Some overdispersion to deal with likely model mismatches in these components seems pragmatic. For the hospital deaths we set  $\kappa = 2000$ , which gives a likelihood very close to Poisson. There seems no legitimate reason to expect overdispersion here if the model is at all fit for purpose.

## A.10 Estimation and inference

Given smoothing parameters,  $\lambda$ , and overdispersion parameters we find the posterior parameter modes

$$\hat{\theta} = \operatorname{argmax}_{\theta} l(\theta) - \frac{1}{2} \theta^T \mathbf{S}_{\lambda} \theta - \frac{1}{2} \sum_j (\theta_j - \mu_j^{\theta})^2 \nu_j^{\theta}, \quad \text{where } \mathbf{S}_{\lambda} = \sum_j \lambda_j \mathbf{S}_j, \quad (48)$$

$l$  is the log likelihood and the  $\mathbf{S}_j$  are fixed positive semi-definite matrices imposing the spline smoothing penalty.  $\mu_j^{\theta}$  and  $\nu_j^{\theta}$  are prior means and precisions. The precisions may be zero, and are for all spline coefficients. Quasi-Newton optimization was used to find  $\hat{\theta}$  (see e.g. Nocedal and Wright, 2006). Let  $\mathbf{H}_{\lambda}$  denote the Hessian of the negative of the penalized log likelihood given in (48), evaluated at  $\hat{\theta}$  (it is obtained by finite differencing the computationally exact gradients). Then we have the large sample approximation to the posterior

$$\theta \sim N(\hat{\theta}, \mathbf{H}_{\lambda}^{-1}). \quad (49)$$

The smoothing parameters were obtained using the approximate marginal likelihood maximization method of Wood and Fasiolo (2017), which alternates optimization of (48), given smoothing parameters, with updates of all smoothing parameters, given  $\hat{\boldsymbol{\theta}}$

$$\lambda_j \leftarrow \frac{\text{tr}(\mathbf{S}_\lambda^{-1} \mathbf{S}_j) - \text{tr}(\mathbf{H}_\lambda^{-1} \mathbf{S}_j)}{\hat{\boldsymbol{\theta}}^\top \mathbf{S}_\lambda \hat{\boldsymbol{\theta}}} \lambda_j.$$

At each step of this iteration the overdispersion parameters were obtained by maximum likelihood estimation, given the current model predictions. Uncertainties for incidence and  $R$  trajectories were obtained from (49) by the delta method (applied on the log scale).