Simulations of Potential Exit Strategies

University of Warwick

Matt Keeling, Mike Tildesley, Louise Dyson, Ed Hill, Erin Gorsich, Bridget Penman, Massi Tamborrino, Jane Hutton, Hector McKimm, Trystan Leng, Glen Guyver-Fletcher, Alex Holmes.

03 May 2020

Executive Summary

The Warwick model has been simulated to consider the effect of phased relaxation strategies from 11th May. The model simulates four relaxation phases starting on 11th May, 1st June, 1st July and 15th August respectively; each represents an increase in contacts at workplaces, schools and through leisure activities (specific details of the assumptions underpinning these four phases are given below). We implement differing levels of contact tracing from 18th May.

Phases 1 and 2 are predicted to have a minimal impact on the outbreak, with cases still observed to decrease, indicating that the effective reproduction number remains below 1.

Phase 3 (which sees a considerable increase in leisure activities and 60% of children in schools) is predicted to lead to a rise in cases and deaths without contact tracing; however with contact tracing the infection continues to decline.

In Phase 4, without contact tracing there is a substantial rise in cases and deaths leading to a second wave of comparable size to the first wave. With contact tracing in place the magnitude of a secondary wave is generally suppressed, but there are some parameter combinations that generate a small second wave.

Throughout, the difference between the two scenarios is minimal (compare solid and dashed lines); the difference between tracing a maximum of 30 contacts (blue) and 15 contacts (green) is relatively slight.

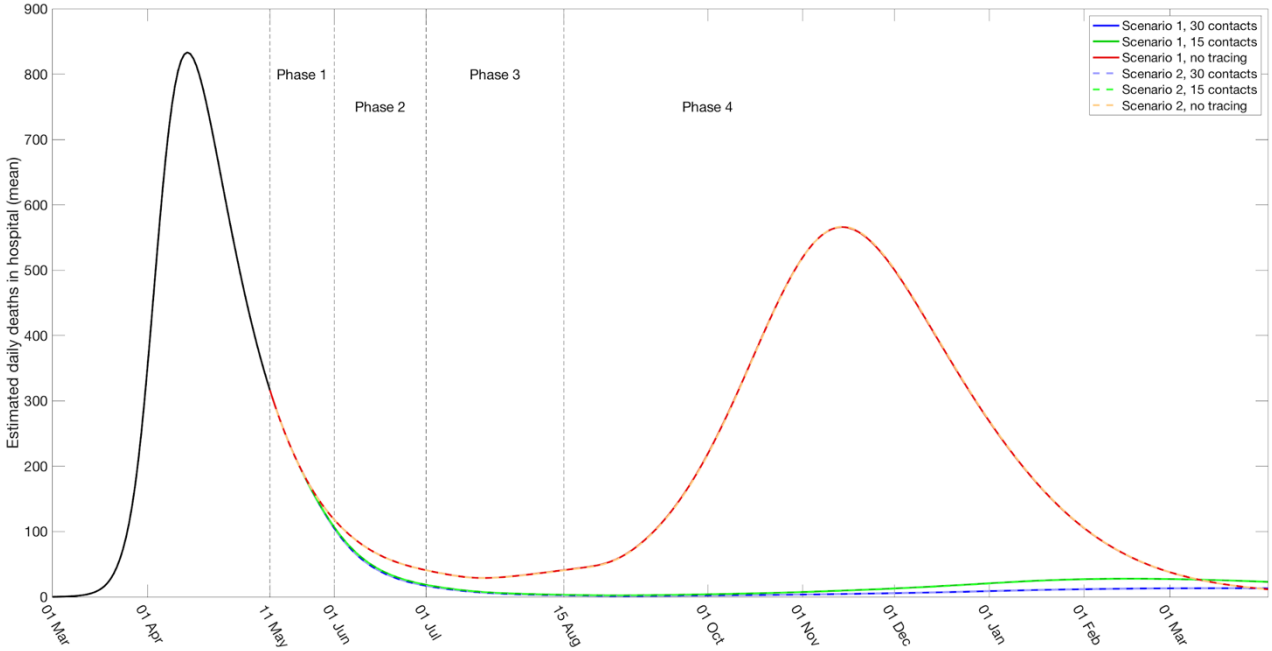

We highlight three caveats about these predictions:

1) They are sensitive to the continued compliance of the population. We have made the optimistic assumption that compliance with other measures (such as household quarantining and self-isolation) are not affected by the relaxations in Phases 1-4.
2) The damping of the outbreak during Phase 4 is highly dependent on the uptake and efficacy of contact tracing in the population. We have assumed that everyone that self-isolates will act as an index case for contact tracing, and that 80% of contacts will be traced and isolated.
3) We have assumed that school attendance is close to 100% in the age-groups that return to school. Lower levels of attendance due to worried parents may reduce cases in Phase 4 slightly.

It is therefore key that epidemiological measures are carefully monitored for signs of growth and prompt action taken.

The Warwick approach uses a deterministic SEIR-style age-structured model, matched to the early UK age-distribution of cases and then fitted to the temporal dynamics across 11 regions. There are a variety of ways in which the early matching can occur (largely reflecting assumptions about asymptomatic infections), and our results average across this uncertainty. A full description of the model formulation has been submitted to *MedRxiv* and the main details are covered in Appendix 1.

2. EXIT STRATEGY SIMULATIONS

METHODS

We performed scenarios as detailed below. Simulations were run until 1st April 2021 and performed across all thinned posterior estimates from ten separate MCMC chains. We consistently plot the median prediction as well as sets of confidence intervals (showing 1, 5, 25, 75, 95 and 99th percentile). For each of the tracing values and scenarios considered we ran the simulation with the same collection of parameter samples.

Scenario Overview

All scenarios were run over the specified four phases, starting on 11 May, 1 June, 1 July, and 15 August. In each scenario, self-isolation on symptoms, household quarantining and shielding remained in place, with adherence to these measures governed by a compliance parameter (whose values were drawn from posterior distributions acquired from inference scheme, fitting to hospitalisation and death data).

- In phase 1 starting on 11th May, we assume that school attendance remains restricted to key workers' children (11% of school-age children) and that workplace contacts have increased by 20% from the baseline level.
- In phase 2 starting on 1st June, workplace contacts are increased by 30% from the baseline level and the number of school children attending school is increased to either 25% of full attendance, representing key workers, vulnerable children and transition years (scenarios 1 and 3) or 50% of full attendance, representing key worker children, vulnerable children and all primary school children (scenarios 2 and 4). Additionally, we modify our contact matrices to increase leisure contacts by 10% from the baseline level.
- In phase 3 starting on 1st July, workplace contacts and leisure contacts are increased by 40% and 30% respectively from the baseline level, whilst key worker children, vulnerable children, transition years and all primary school children return to school (60% of full attendance). In this scenario, we assume that schools remain open until 22nd July, after which the summer vacation begins.
- In phase 4, starting on 15th August, workplace contacts and leisure contacts are increased by 50% and 75% respectively from the baseline level and all children return to school from 1st September.

In schools, we assume that class sizes remain as they were before lockdown and that children are as infectious as adults but have a higher probability of being asymptomatic or showing mild symptoms.

Contact Tracing

Contact tracing is limited to 25,000 new COVID-like symptomatic index cases per day, with tracing of contacts initiated from 18th May. We consider 3 variants on contact tracing: 1) a maximum of 30 contacts can be identified per individual of which 80% of traced and self-isolate; 2) a maximum of 15 contacts can be identified per individual of which 80% of traced and self-isolate; 3) no contact tracing. We have assumed that everyone that would self-isolate will be willing to act as an index case for contact tracing. When there is a limit on the number of contacts, we take those of the longest duration.

The action of contact tracing is determined by a separate analysis of the Warwick Contact Survey. This allows us to assess the number of contacts with age and the proportion of cases that could be traced (see Appendix 2). These tracing values then modify the dynamics of the ODE model, leading to the quarantining for 2 weeks of newly infected individuals as well as susceptible and recovered members of the population. The numbers in quarantine could be greatly reduced by testing for infection. Given that in all of the simulations the number of new index cases never approaches the limit of 25,000 per day, we ignored the potential effect of tracing from individuals not infected with COVID-19.

RESULTS

All graphs are given in Appendix 3. We focus here on the dynamics for England simulated as a single region. Given the very limited difference between the two school Scenarios and the two variants of contact tracing, we describe the results for a maximum of 30 contacts per day and Scenario 1.

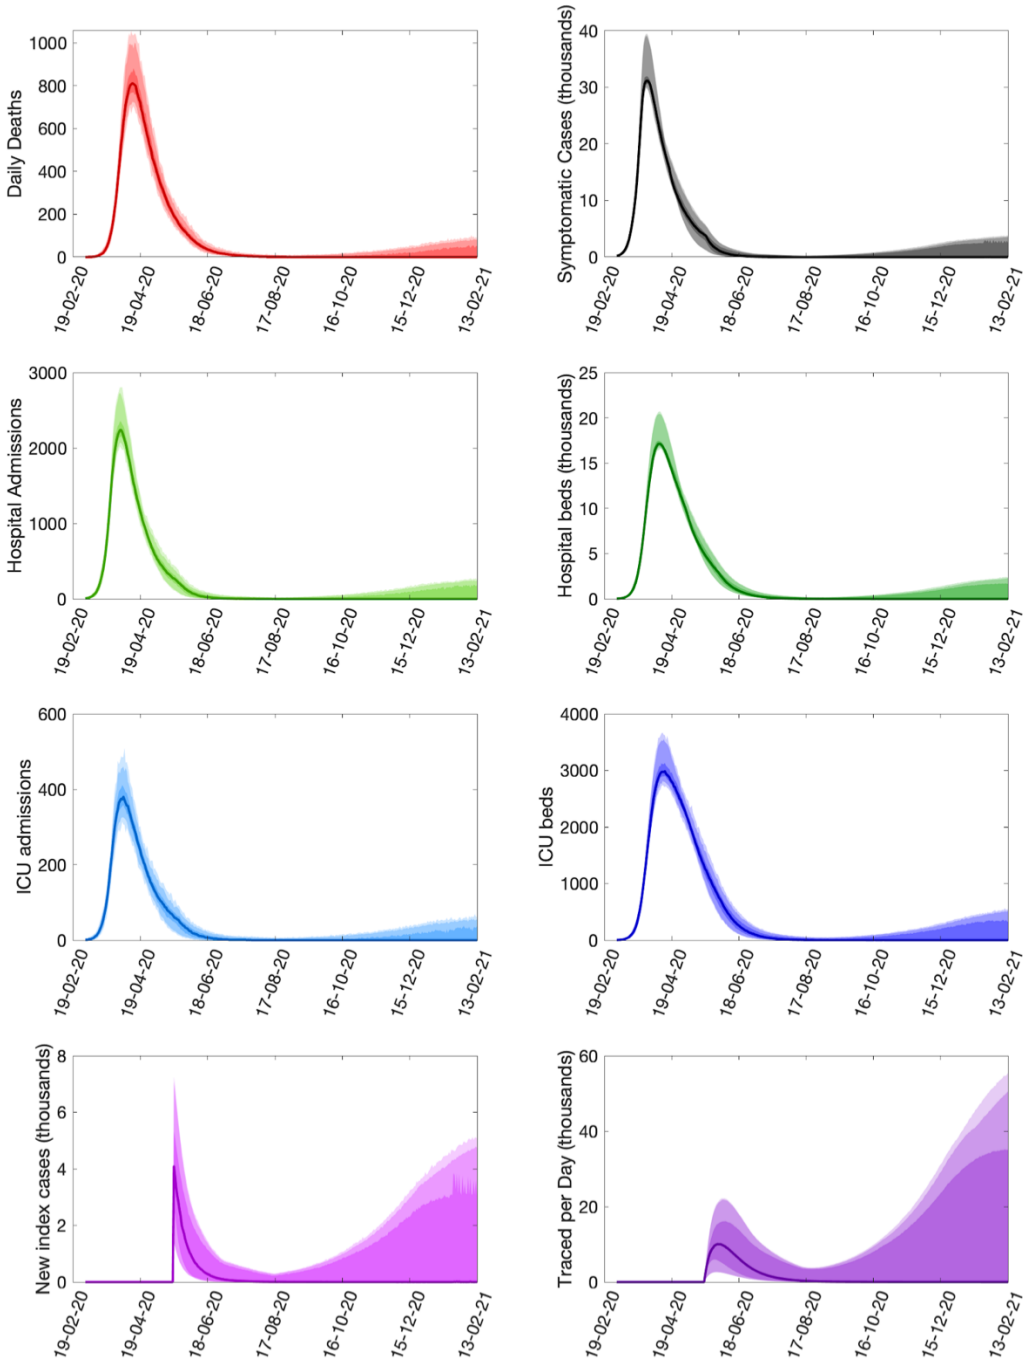

The most striking observation is that under our basic assumptions we are unlikely to experience a large second wave. The median prediction is that contact tracing is able to suppress the second wave of cases, but a long protected second wave cannot be discounted. We estimate a notable second wave (>1000 deaths) in 25% of simulations of Scenario 1 and up to 30 contacts a day (26 % for Scenario 2 and 30 contacts, 26% for Scenario 1 and 15 contacts, and 27% for Scenario 2 and 15 contacts).

In Phases 1-3 (assuming contact tracing is in place) the R_t value remains persistently below one. It is only in Phase 4, and for a limited set of posterior parameters that R_t can increase above one. Without contact tracing (see Appendix 3) both Phase 3 and 4 are consistently predicted to have R_t above one, highlighting the importance of contact tracing.

By far the greatest model uncertainty occurs in Phase 4, where some parameter combinations lead to a significant amount of tracing and quarantining of individuals (lower right graph) - the upper 95% confidence intervals peak at over 50,000 individuals traced per day, although the median is still close to zero.

From these simulations, it is therefore impossible to robustly predict the demand for contact tracing or the resources that need to be deployed.

We find very little difference between the two main relaxation Scenarios - this is largely unsurprising as previous work by the Warwick group and others suggests that school scenarios 3 and 5 have comparable impact on the dynamics).

SENITIVITY ANALYSIS

We have now undertaken a brief sensitivity analysis of the role of reducing compliance as the epidemic unfolds and we shift through different phases of relaxation. This has been undertaken in two ways: firstly choosing three different exemplars of how relaxation of controls might affect the relative degree of compliance in the population; secondly a wider parameter sweep of changing compliance in the Phase 4.

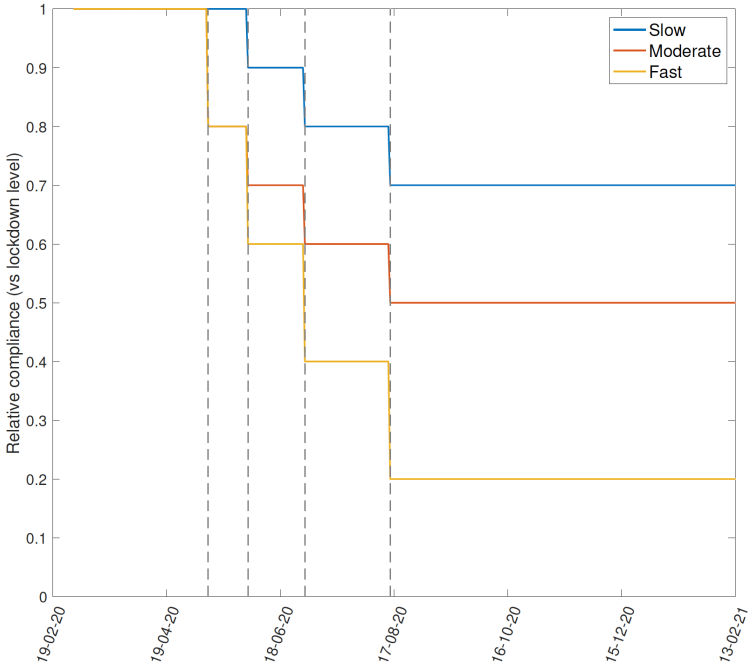

This figure shows the three examples of relative decline in compliance investigated. Slow decline drops to 70% of the current value in Phase 4, while Moderate and Fast drop to 50% and 20% respectively.

It should be noted that it is difficult to relate these aggregate changes in compliance to specific up takes of activities – it is more about a general pattern of increased mixing and less social distancing.

	Compliance relaxation speed		
	Slow	**Moderate**	**Fast**
No tracing	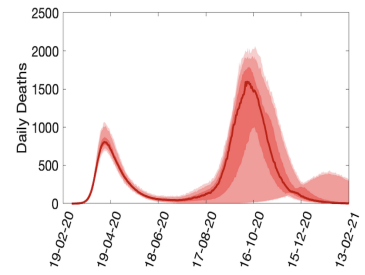	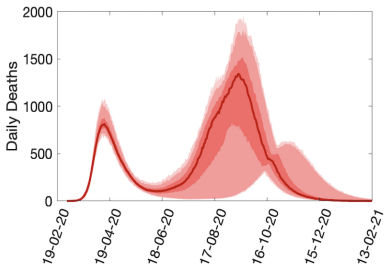	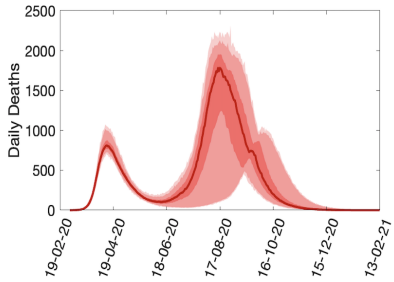
15 contacts	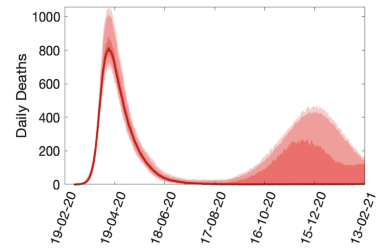	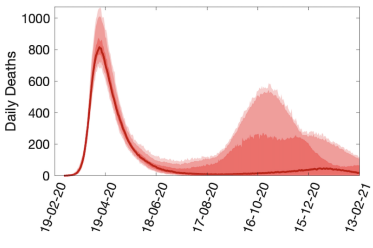	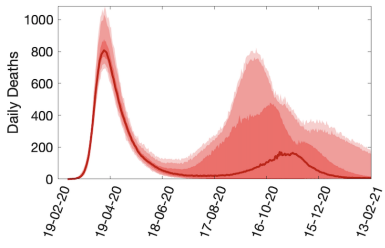
30 contacts	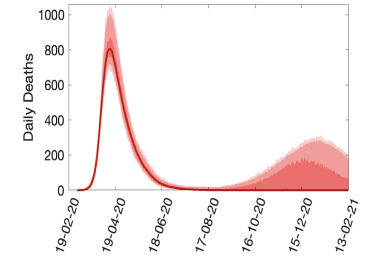	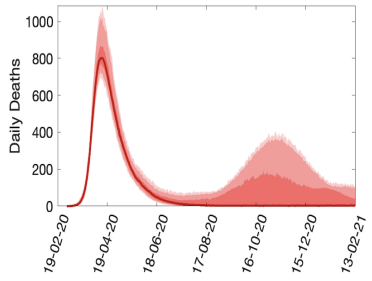	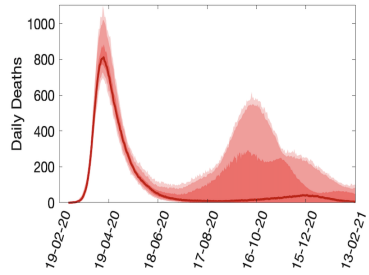

The resulting predictions of deaths in England, show that the tracing protocols can cope with a moderate change in compliance. This is far more buffered when tracing up to 30 contacts per day, but relatively large second waves are still within the confidence intervals.

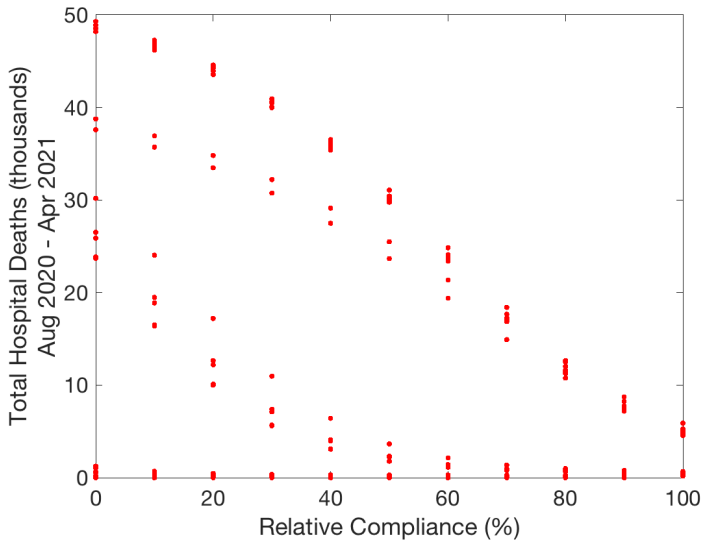

We have also looked at changes in compliance in Phase 4, assuming that we are able to trace 30 contacts per index case. The graph shows the number of hospital deaths during the second wave (from 15th Aug until the end of the simulations). In general, there is a linear increase in deaths with decreasing compliance, although for some parameter sets there is evidence of a threshold, below which much larger epidemics are predicted.

CONCLUSIONS

A simulation model, fitted to the early outbreak data, was expanded to include a more refined school structure and the impact of contact tracing. Given the time taken to develop these extensions and the short time scale, a fully comprehensive sensitivity analysis to all assumptions could not be completed. Instead we focus on the two scenarios of interest, the two variations of contact tracing (and compare to the absence of tracing) and across the posteriors of the MCMC inference. A more thorough examination of sensitivity to contact tracing assumptions (such as the 80% tracing efficacy, or delays to tracing) would be beneficial.

In general the proposal of gradually opening schools, workplaces and leisure activities, together with contact tracing limits the chance of a second wave. Where a second wave does occur it is generally low and protracted. This wave only arises for some combination of parameters, suggesting that there is currently insufficient epidemiological information to determine if a second wave is likely - some of this uncertainty may be resolved through incorporating serological data into the inference (on-going work).

For the two Scenarios investigated, there is only marginal epidemiological difference, so the choice between these needs to be determined on educational and other priorities. As expected tracing a maximum of 30 contacts per index suppresses the epidemic more than having a maximum of 15. However, given the fact that the number of new daily index cases never approaches the limit of 25,000, there should be sufficient capacity to trace 30 contacts per index case. One complication is that tracing from non-COVID index individuals has not been included, our expectation is that this effect will be minimal.

While we have focused on results for England in the main text, results for the 8 regions within England as well as Wales, Scotland and Northern Ireland are given in Appendix 3. In general the difference in age-structure and the subtle parameter differences mean that some regions (East of England, North East & Yorkshire) are expected to experience a small second waves, with greater than 50% probability. We note here that difference is data between the different home nations will affect the robustness of the MCMC inference procedure.

We have made the optimistic assumption that compliance with other measures (such as household quarantining and self-isolation) are not affected by the relaxations in Phases 1-4. We have also assumed that everyone that self-isolates will act as an index case for contact tracing, and that 80% of contacts will be traced and isolated. Both of these are largely suppositions, it is therefore key that epidemiological measures are carefully monitored for signs of growth and prompt action taken.

APPENDIX 1. THE WARWICK MODEL

Here, we describe a compartmental model that has been developed to simulate the spread of SARS-CoV-2 virus (resulting in cases of COVID-19) in the UK population. In the ongoing outbreak in the UK, cases of COVID-19 are confirmed based upon testing, with priority for testing given to patients requiring critical care in hospitals - generating biases and under-reporting. There is evidence to suggest that a significant proportion of individuals who are infected may be asymptomatic or have only mild symptoms. These asymptomatic individuals are still able to transmit infection, though it remains unclear whether they do so at a reduced level. Our modelling approach has consequently been designed to consider the interplay between symptoms (and hence detection) and transmission of COVID-19. We developed a deterministic, age-structured compartmental model, stratified into five-year age bands. Transmission was governed through age-dependent mixing matrices based on UK social mixing patterns. The population was further stratified according to current disease status, following a susceptible-exposed-infectious-recovered (SEIR) paradigm, as well as differentiating by symptoms, quarantining and household status (figure below). Susceptibles (S) infected by SARS-CoV-2 entered a latent state (E) before becoming infectious. Given that only a proportion of individuals who are infected are tested and subsequently identified, the infectious class in our model was partitioned into symptomatic (and hence potentially detectable), D, and asymptomatic (and likely to remain undetected) infections,U. We assumed both susceptibility and disease detection were dependent upon age, although the partitioning between these two components is largely indeterminable. We modelled the UK population aggregated to ten regions (Wales, Scotland, Northern Ireland, East of England, London, Midlands, North East and Yorkshire, North West England, South East England, South West England).

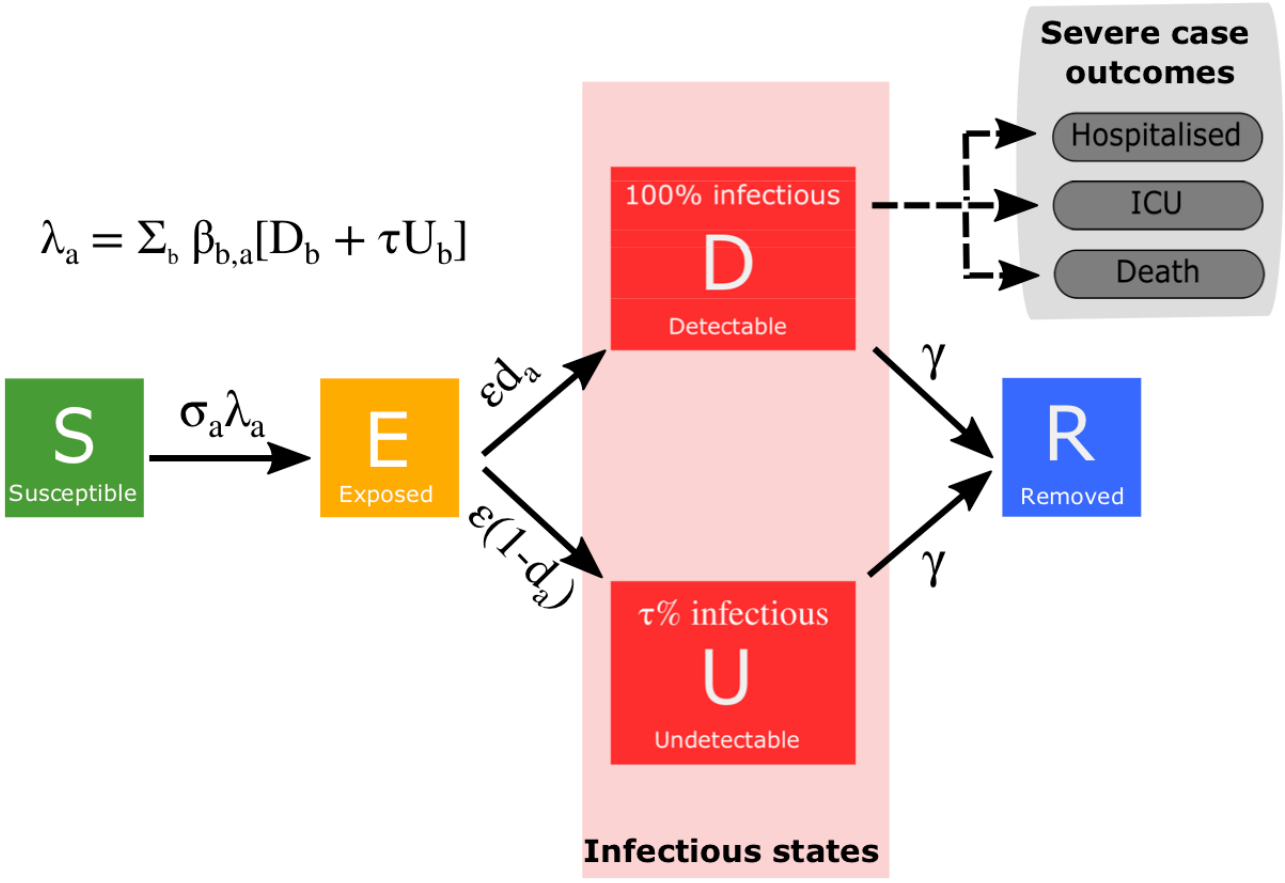

A drawback of the standard SEIR ordinary differential equation (ODE) formation in which all individuals mix randomly in the population is that it cannot readily account for the isolation of households. For example, if all transmission outside the household is set to zero in a standard ODE model, then an outbreak can still occur as within-household transmission allows infection between age-groups and does not account for local depletion of susceptibles within the household environment. We addressed this limitation by extending the standard SEIR models such that first infections without a household (EF,DF,UF) are treated differently from subsequent infections (ES,DS,US). To account for the depletion of susceptibles in the household, we assumed that all within household transmission is generated by the first infection within the household.

The model is concerned with epidemiological processes and so predicts the number of symptomatic and asymptomatic infections on each day. However, in order to provide evidence regarding the future impact of the outbreak in the UK, it is crucial to be able to predict the number of severe cases that may require hospital or critical care. We utilised two processes in order to estimate hospitalisation rates:(i) we estimated the proportion of clinical cases in each age group that would require hospitalisation by comparing the age distribution of hospital admission to the age structure of early detected cases - assuming these detected cases were an unbiased sample of symptomatic individuals. (ii) we used age independent distributions to determine the time between onset of symptoms and hospitalisation. A similar process was repeated for admission into intensive care units. Both of these distributions were drawn from the COVID-19 Hospitalisation in England Surveillance System (CHESS) data set that collects detailed data on patients infected with COVID-19 [14]. Information on the distributions of length of stay in both intensive care units (ICUs) and hospital was used to translate admissions into bed occupancy - which adds a further delay between the epidemiological dynamics and quantities of interest. In terms of matching the available data and quantities of interest, we also use the prediction of symptomatic infections to drive the estimated daily number of deaths within hospitals. The risk of death is again captured with an age-dependent probability, while the distribution of delays between hospital admission and death is assumed to be age-independent. These two quantities are determined from the Public Health England (PHE) death records.

We fit on a region-by-region basis to four timeseries: (i) new hospitalisations; (ii) hospital bed occupancy; (iii) ICU bed occupancy; (iv) daily deaths (using data on the recorded date of death, where-ever possible). The relative transmission rate from asymptomatic cases (τ) and the scaling of whether age-structure case reports are based on age-dependent susceptibility or age-dependent symptoms (α) were treated as free parameters. We performed parameter inference using the Metropolis-Hastings algorithm, computing likelihoods assuming the daily count data to be drawn from a Poisson distribution. At the outset of each iteration, based on the α and τ instances at that moment, we first set the recovery rate γ, age-dependent susceptibilites (σ_a) and age-dependent probabilities of displaying symptoms (d_a). We achieved this through calibration to early age-stratified UK data from 1st February 2020 to 1st April 2020; initial assumptions were an R0 of 2.7 and doubling time of 3.3 days, but these are allowed to vary between regions.

After a burn-in of 250,000 particles, the algorithm was run for a further 250,000 iterations. We thinned the generated parameter sets by a factor of 100, giving 2,500 parameter sets representing samples from the parameter posterior distributions.

Model equations

The full equations are given by an equation for the susceptibles

$$\frac{dS_a}{dt} = q_t Q_a^S - \left(\lambda_a^{FD} + \lambda_a^{FU} + \lambda_a^{SD} + \lambda_a^{SU} + \lambda_a^Q - T_a^c\right)\frac{S_a}{N_a},$$

the four exposed classes (first exposed in the household; second exposed, caught from a detectable infection; second exposed, caught from an undetectable infection; and exposed, but in quarantine:

$$\begin{aligned}
\frac{dE_a^F}{dt} &= (t_a^r \lambda_a^{FD} + \lambda^{FU})\frac{S_a}{N_a} - \epsilon E_a^F,\\
\frac{dE_a^{SD}}{dt} &= t_a^r \lambda^{SD}\frac{S_a}{N_a} - \epsilon E_a^{SD},\\
\frac{dE_a^{SU}}{dt} &= \lambda^{SU}\frac{S_a}{N_a} - \epsilon E_a^{SU},\\
\frac{dE_a^Q}{dt} &= \lambda^Q\frac{S_a}{N_a} - \epsilon E_a^Q,
\end{aligned}$$

the five detectable infected classes (a proportion, d_a of the exposed progress to detectable infections, D_a^F, D_a^{SD} and D_a^{SU}; and the first and second exposed can be quarantined, D_a^{QF} and D_a^{QS}:

$$\begin{aligned}
\frac{dD_a^F}{dt} &= d_a(1-H)\epsilon E_a^F - \gamma D_a^F,\\
\frac{dD_a^{SD}}{dt} &= d_a\epsilon E_a^{SD} - \gamma D_a^{SD},\\
\frac{dD_a^{SU}}{dt} &= d_a(1-H)\epsilon E_a^{SU} - \gamma D_a^{SU},\\
\frac{dD_a^{QF}}{dt} &= d_a H\epsilon E_a^F - \gamma D_a^{QF},\\
\frac{dD_a^{QS}}{dt} &= d_a H\epsilon E_a^{SU} + d_a\epsilon E_a^Q - \gamma D_a^{QS},
\end{aligned}$$

the five undetectable infected classes (the rest of the exposed progress to undetectable infections, U_a^F and U_a^S; and may be quarantined, U_a^Q:

$$\begin{aligned}
\frac{dU_a^F}{dt} &= (1-d_a)\epsilon E_a^F - \gamma U_a^F,\\
\frac{dU_a^S}{dt} &= (1-d_a)\epsilon(E_a^{SD} + E_a^{SU}) - \gamma U_a^S,\\
\frac{dU_a^Q}{dt} &= (1-d_a)\epsilon E_a^Q - \gamma U_a^Q,
\end{aligned}$$

finally, some of the contacts of detectable infections are traced and quarantined, and these usefully traced contacts may be susceptible Q_a^S or recovered Q_a^R:

$$\begin{aligned}
\frac{dQ_a^S}{dt} &= T_a^c\frac{S_a}{N_a} - q_t Q_a^S,\\
\frac{dQ_a^R}{dt} &= T_a^c\frac{R_a}{N_a} + (\lambda_a^{FD} + \lambda_a^{FD})(1-T_a^c) - q_t Q_a^R
\end{aligned}$$

where q_t is the time in quarantine and the traced contacts are given by

$$T_a^c = s_c M_a^t \sum N_a^Q - \lambda_a^{FD}(1-t_a^r) - \lambda_a^{SD}(1-t_a^r),$$

where t_r is the tracing ratio (which is affected by the total number of contacts that can be traced), s_c is a scaling ratio, and the mean number of traced contacts is given by

$$M_a^t = \frac{(D_a^F + D_a^{SD} + D_a^{SU})\beta_{ba}^N + D_a^F\beta_{ba}^H}{\gamma\sum_b(D_b^F + D_b^{SD} + D_b^{SU})}$$

with the forces of infection obeying

$$\lambda_a^{FD} = \sigma_a\sum_b\left(D_b^F + D_b^{SD} + D_b^{SU}\right)\beta_{ba}^N,$$

$$
\begin{aligned}
\lambda_a^{FU} &= \sigma_a \sum_b \tau (U_b^F + U_b^S) \beta_{ba}^N, \\
\lambda_a^{SD} &= \sigma_a \sum_b D_b^F \beta_{ba}^H, \\
\lambda_a^{SU} &= \sigma_a \tau \sum_b U_a^F \beta_{ba}^H, \\
\lambda_a^{Q} &= \sigma_a \sum_b D_b^{QF} \beta_{ba}^H,
\end{aligned}
$$

where β^H is household transmission and $\beta^N = \beta^S + \beta^W + \beta^O$ is all other transmission locations, comprising school-based transmission (β^S), work-place transmission (β^W) and transmission in all other locations (β^O). σ_a corresponds to the age-dependent susceptibility of individuals to infection, d_a the age-dependent probability of displaying symptoms (and hence being detected), and τ represents reduced transmission of infection by undetectable individuals compared to detectable infections. Each exposed class is in fact composed of multiple steps (not shown here for brevity), so that the time spent in an exposed class is Erlang distributed rather than exponential.

APPENDIX 2. RESULTS FROM THE WARWICK CONTACT SURVEY

We characterised contact patterns in the UK using a postal and online cross-sectional survey, which asked participants to report the number of social encounters with unique individuals during a given day, as well as the duration and typical frequency of those encounters (Danon et al 2012, 2013). In total, 5,802 respondents reported more than 50,000 encounters - one of the biggest studies of its kind to date. The encounter patterns of this study were in good qualitative agreement with other similar studies of social interactions (Mossong et al 2008, Isella et al 2010). In this study, the daily encounter data was first extrapolated to generate a pattern of contacts over a 14 day period (replicating random encounters and increasing the total duration associated regular contacts), to act as the basis for transmission and contact tracing simulations. Using this extrapolated data, we can classify interactions into those which satisfy the definition of a close contact for the purpose of contact tracing. From our social encounter data we can also distinguish interactions with people who could be later identified and traced, from those with unidentifiable strangers. We assume that all contact of longer than 1 hour or repeated contacts can be identified and traced, whereas shorter meetings with people for the first time are strangers who are unidentifiable. We define a contact as anyone encountered for 15 minutes or more, and within 2 meters (The full details of this method and its implications for contact tracing of COVID-19 are explored in Keeling, Hollingsworth & Read (2020) MedRxiv doi.org/10.1101/2020.02.14.20023036).

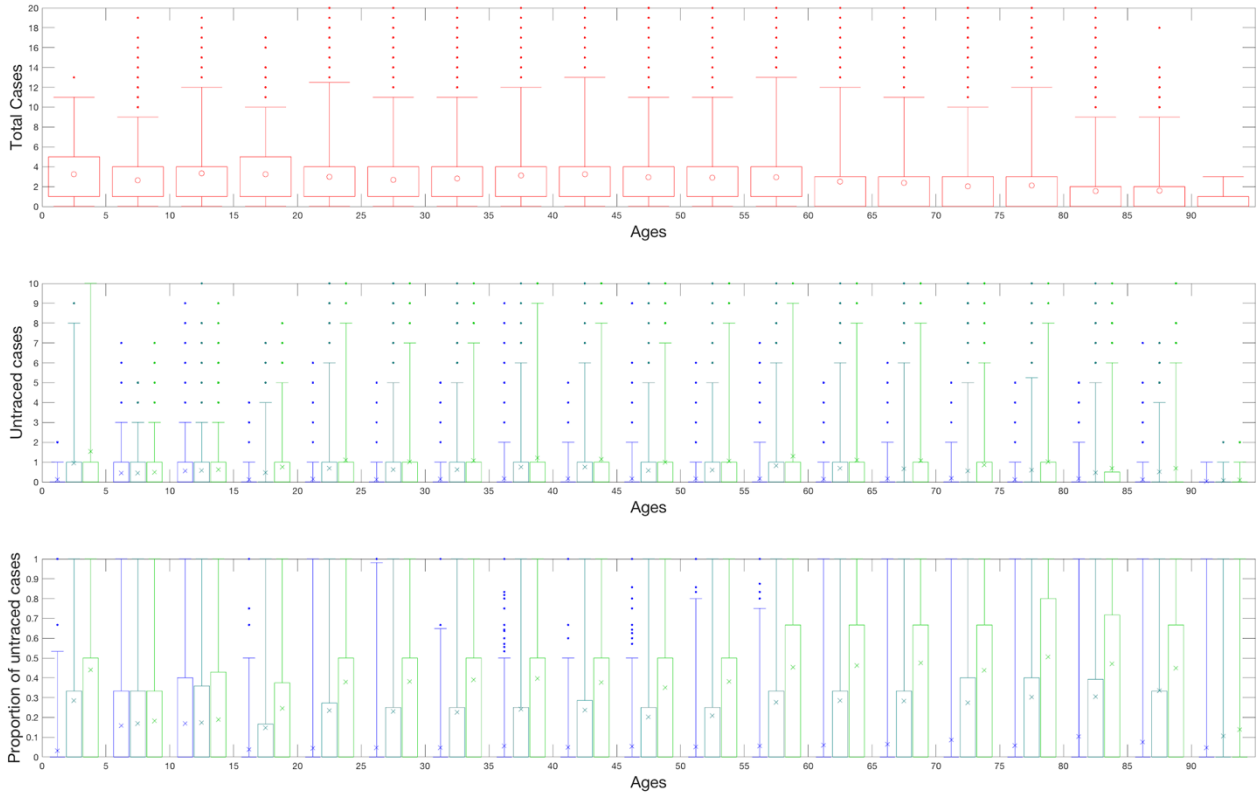

Here we show the distribution of all secondary cases from an index case of a given age (top panel). While there is some age-dependency especially for older age-groups, the pattern is dominated by individual variation. The middle panel shows how many cases would go untraced under 3 different assumptions (from left to right): unlimited tracing, tracing a maximum of 30 contacts per index case and tracing a maximum of 15 contacts per index case. As expected the impact of this cut-off to tracing effort is highly non-linear reflecting the fact that many individuals will have fewer than 15 close contacts over the period.

Finally, we estimate the reduction in cases that would be generated by different limits on the contact tracing (lower panel); while there remains considerable variability, reflecting the heterogeneity in the number of social contacts, there is generally a consistent pattern across ages and between different variants of contact tracing.

This reduction in cases is used to mimic the action of contact tracing within the ODE model; determining what fraction of secondary cases from an unidentified index case are traced and isolated. Given the assumed speed of contact tracing (80% within 48 hours) we have taken the optimistic assumption that none of the secondary cases will be actively transmitting yet. Obviously, not all contacts that are traced will be infected, so the remaining contacts will be a representative mixture of susceptible and recovered individuals; for these individuals testing would allow them to be released from quarantine much earlier than the assumpted 14 days.

APPENDIX 3. RESULTS FROM OTHER SCENARIOS AND REGIONS

Scenario 1, maximum 15 contacts traced, England

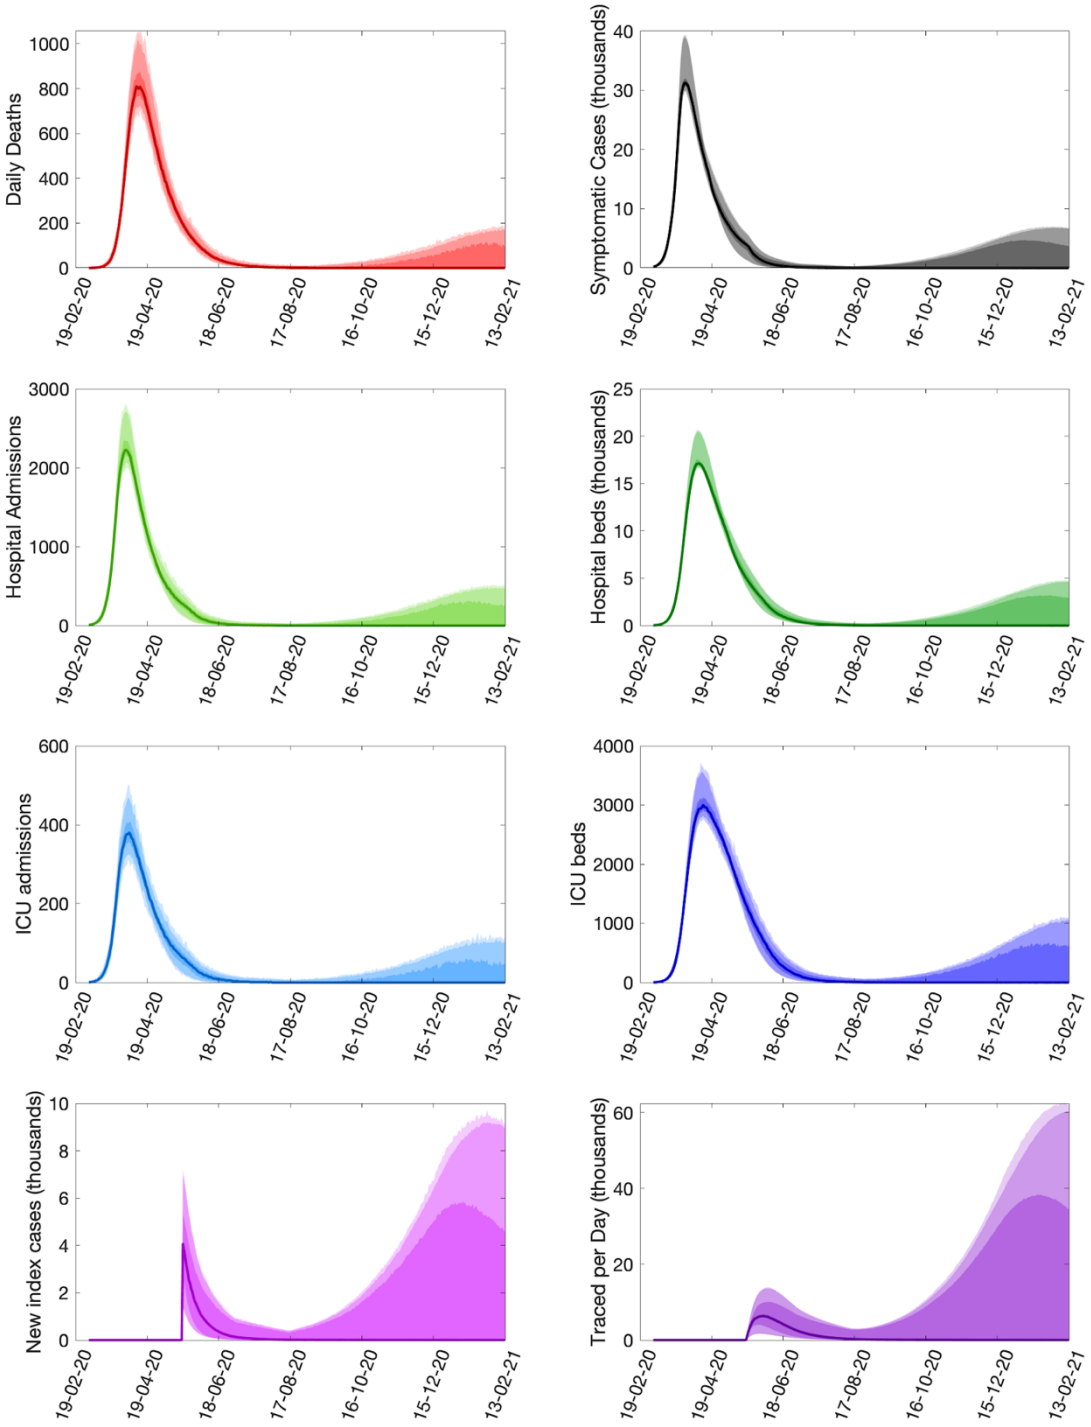

Scenario 2, maximum 30 contacts traced, England

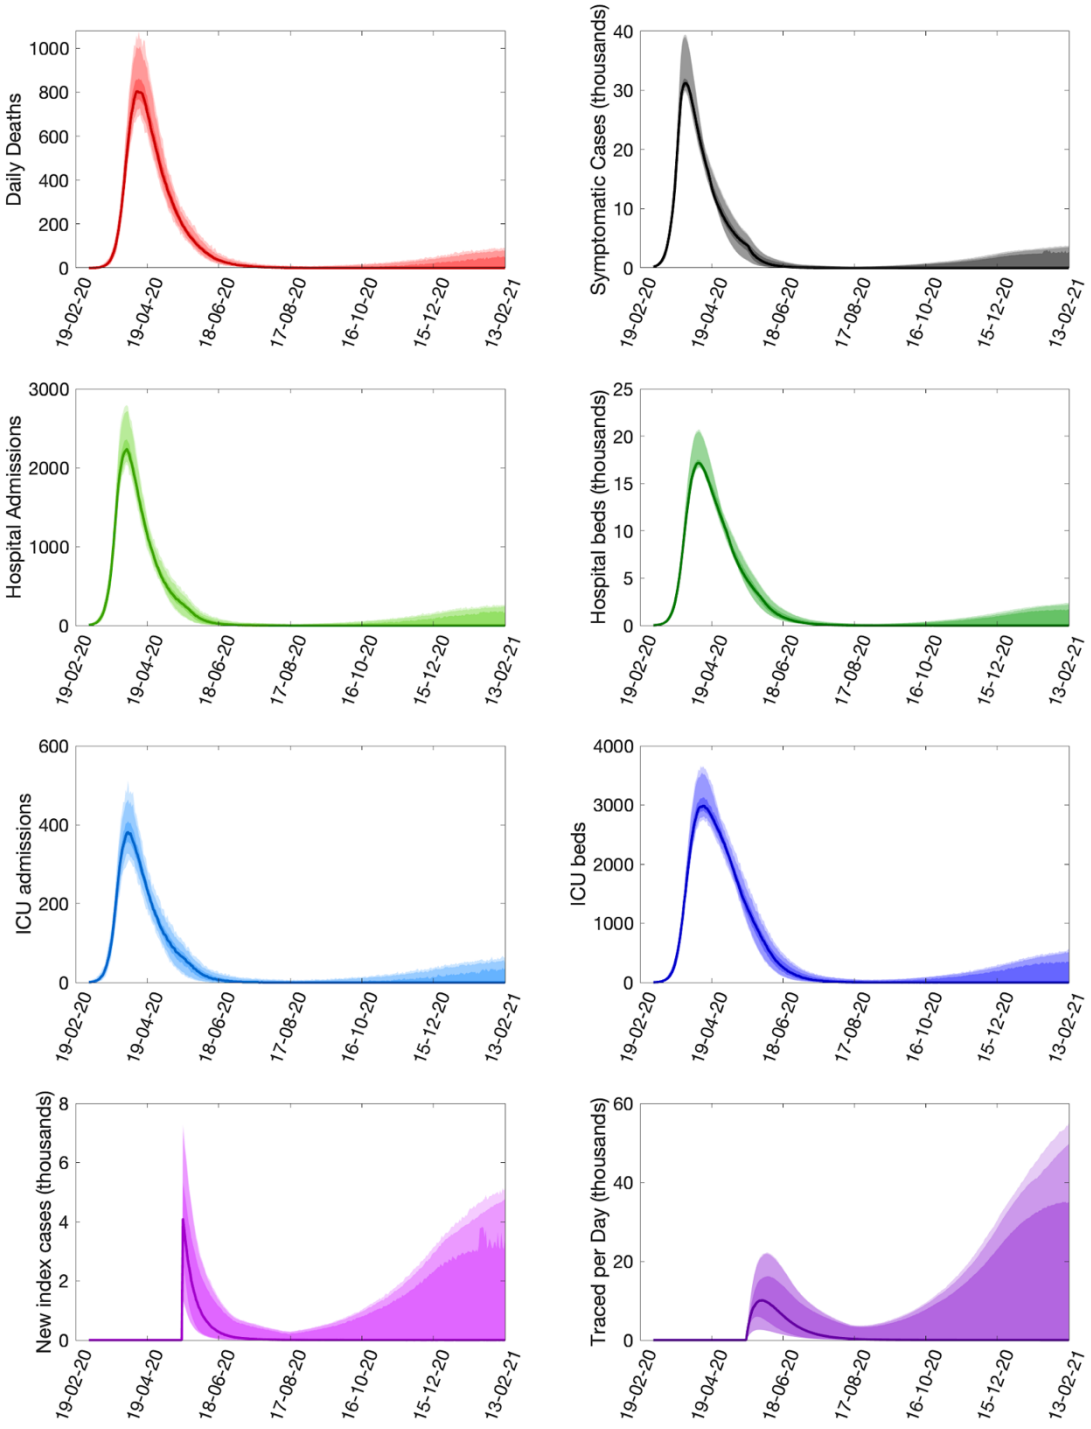

Scenario 2, maximum 15 contacts traced, England

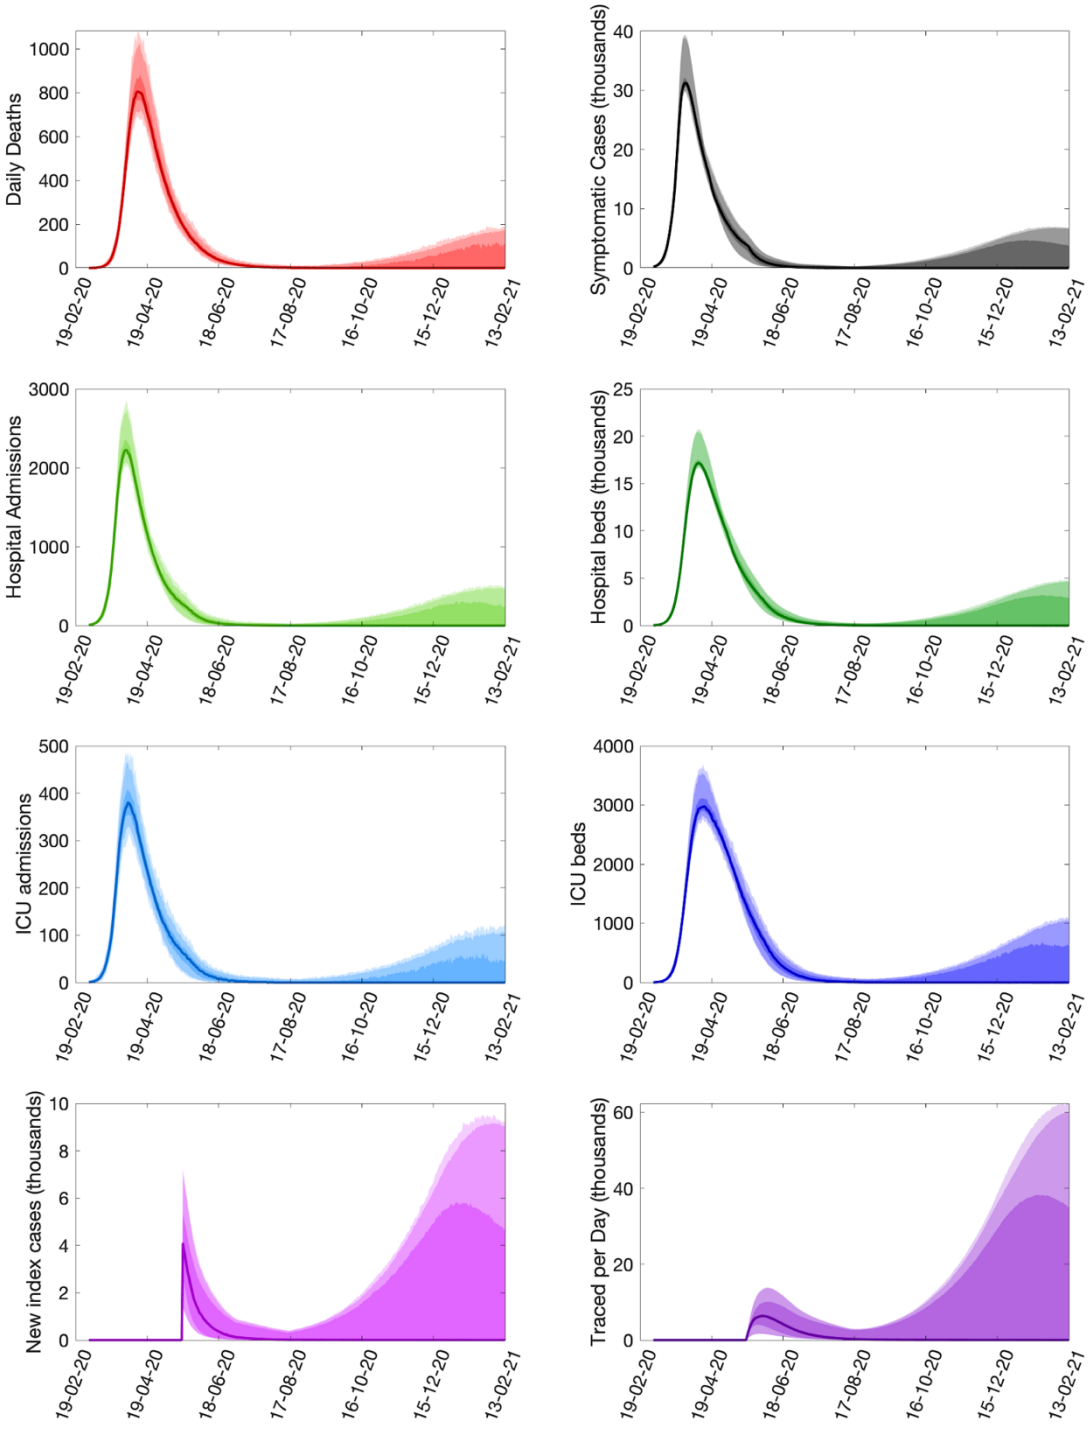

Scenario 1, no contact tracing, England

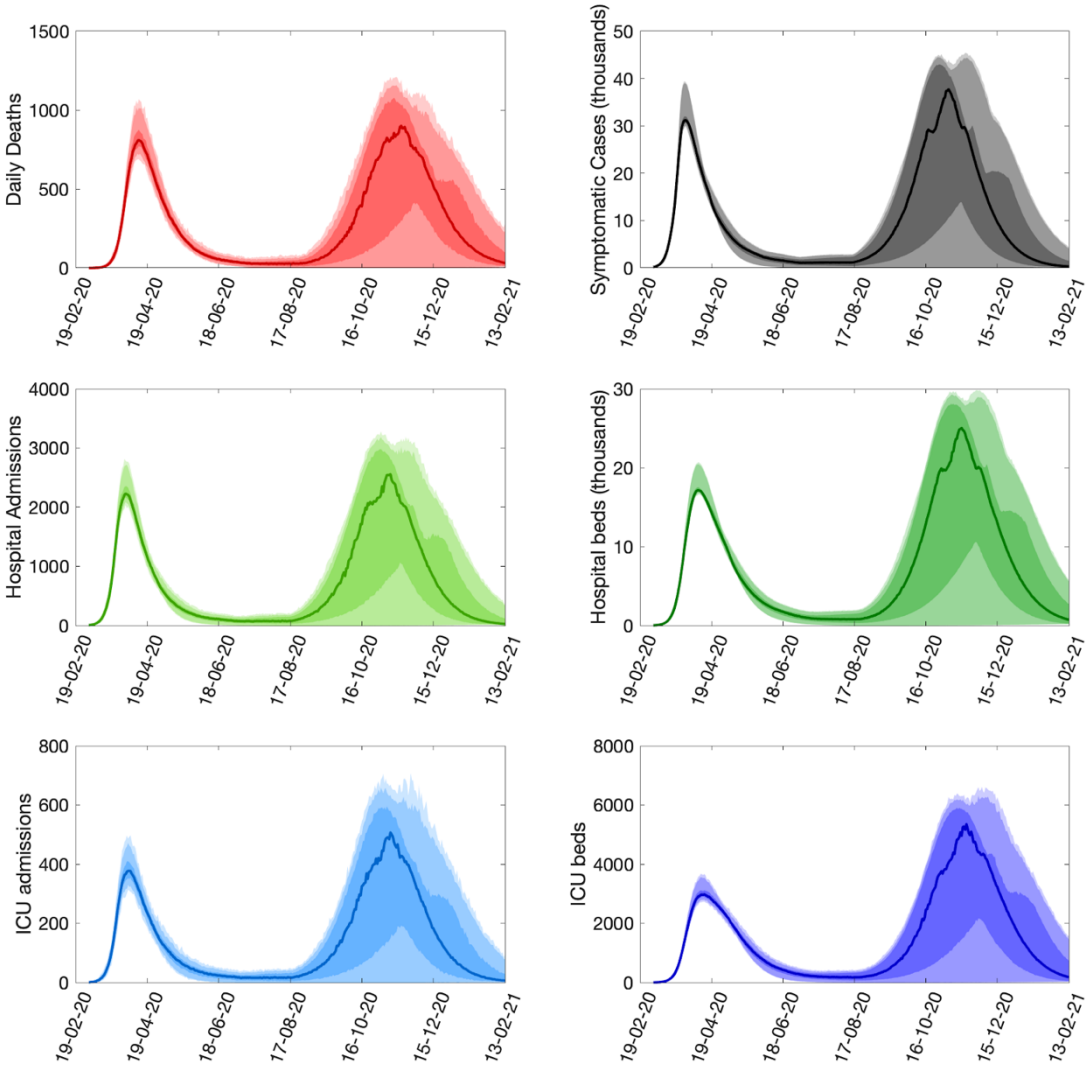

Scenario 1, maximum 30 contacts traced, Wales

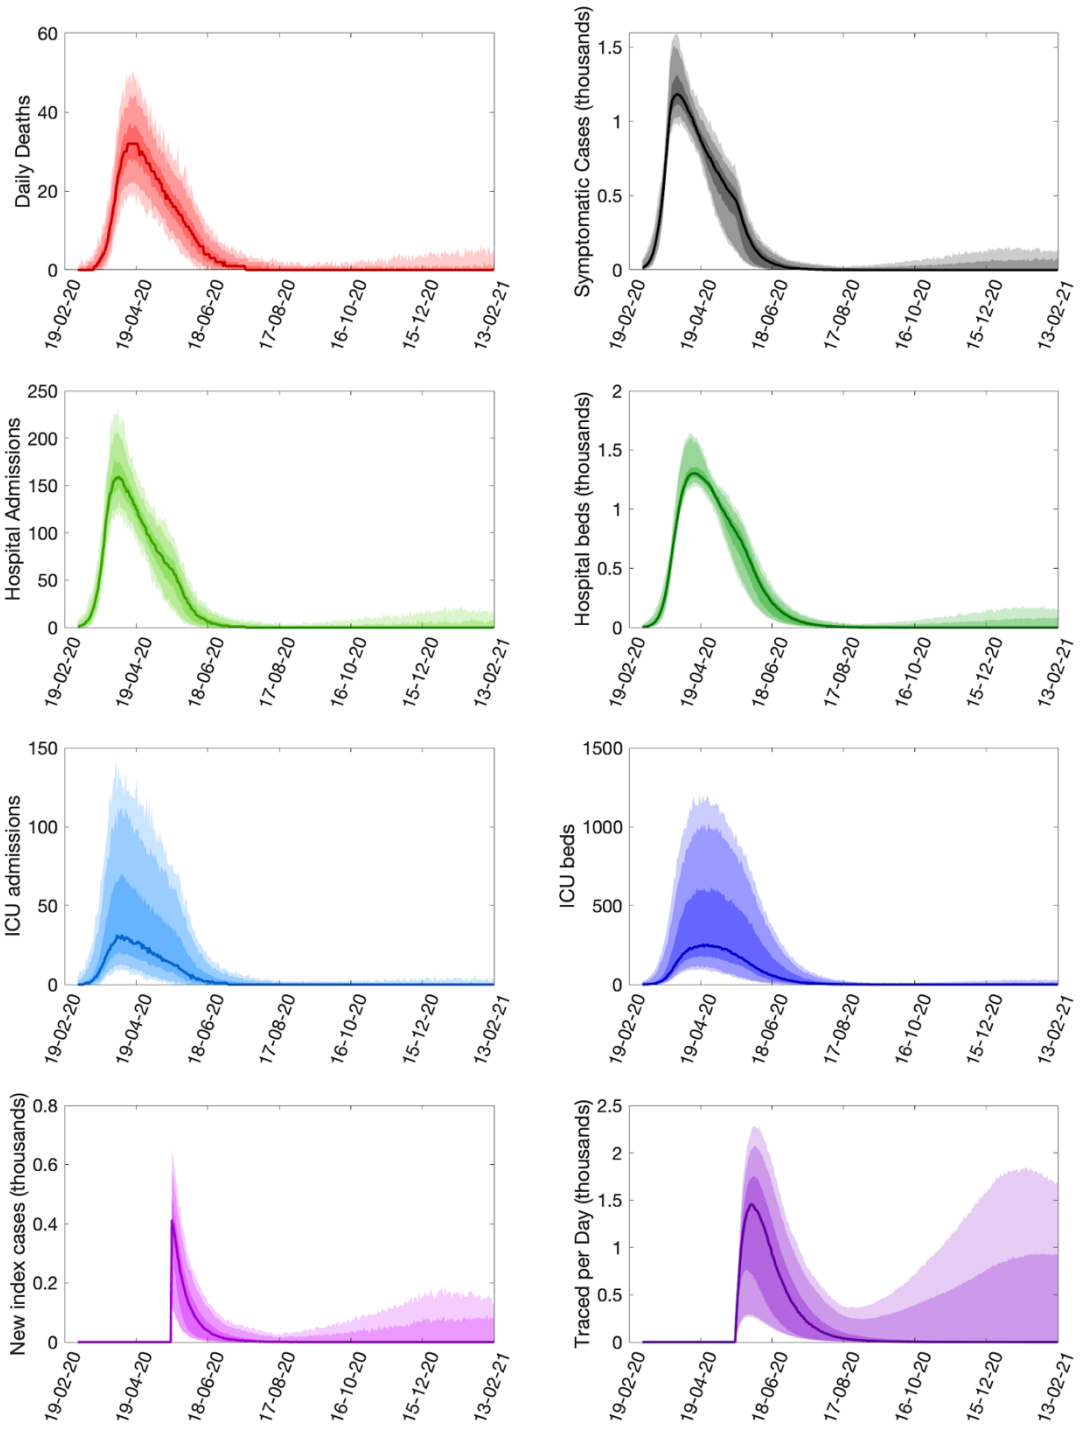

Scenario 1, maximum 30 contacts traced, Scotland

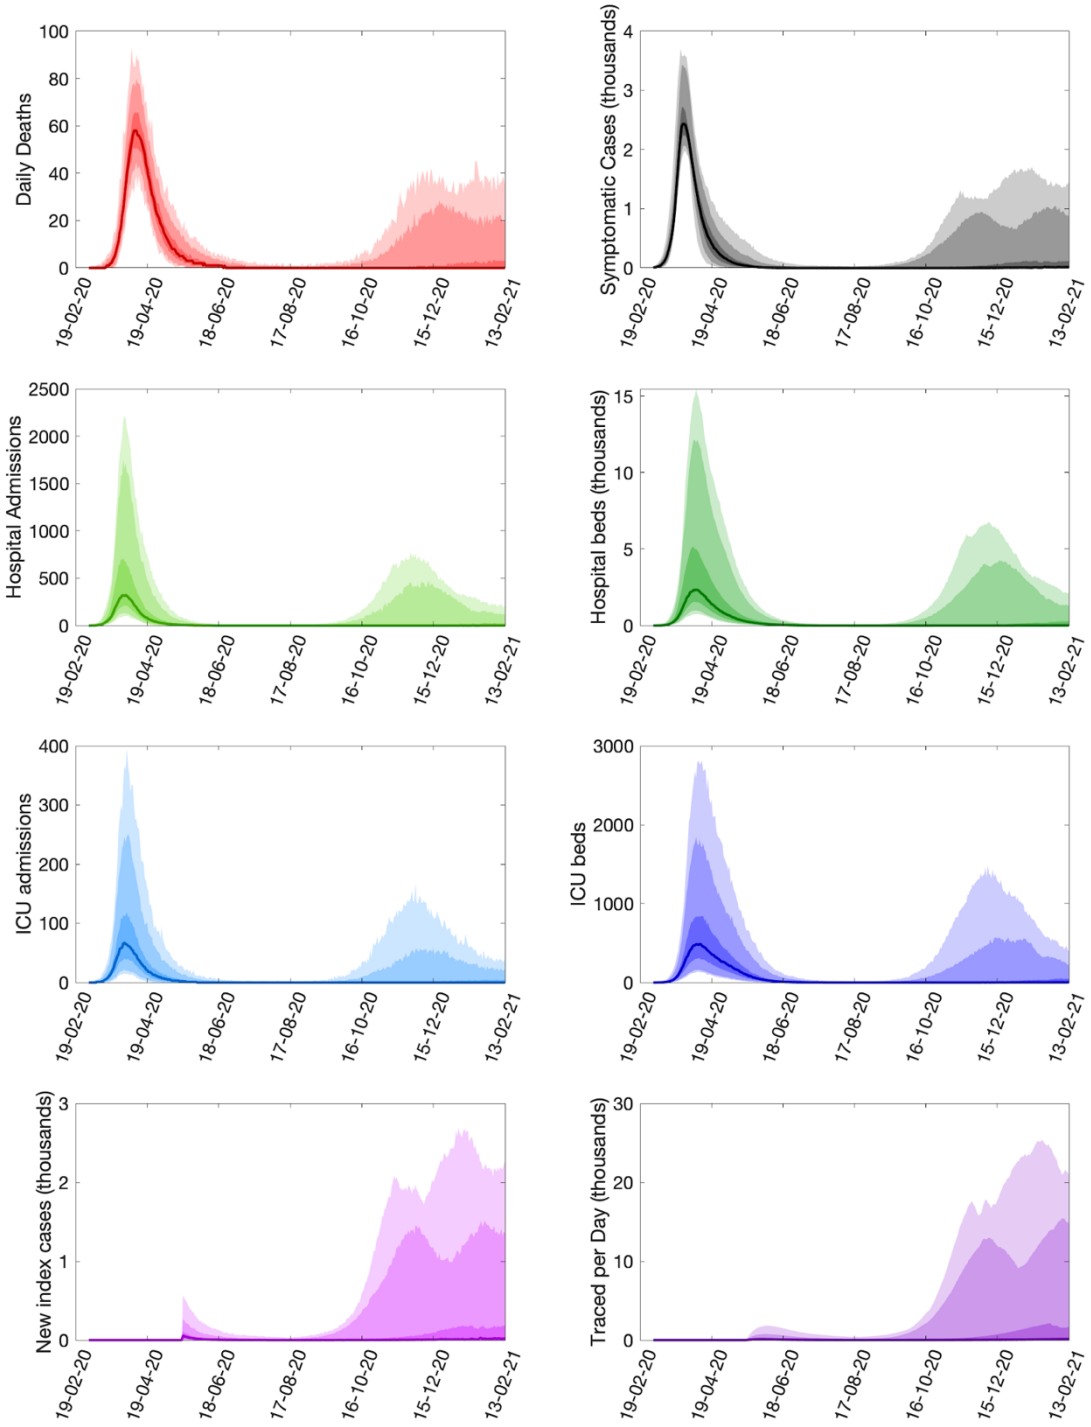

Scenario 1, maximum 30 contacts traced, Northern Ireland

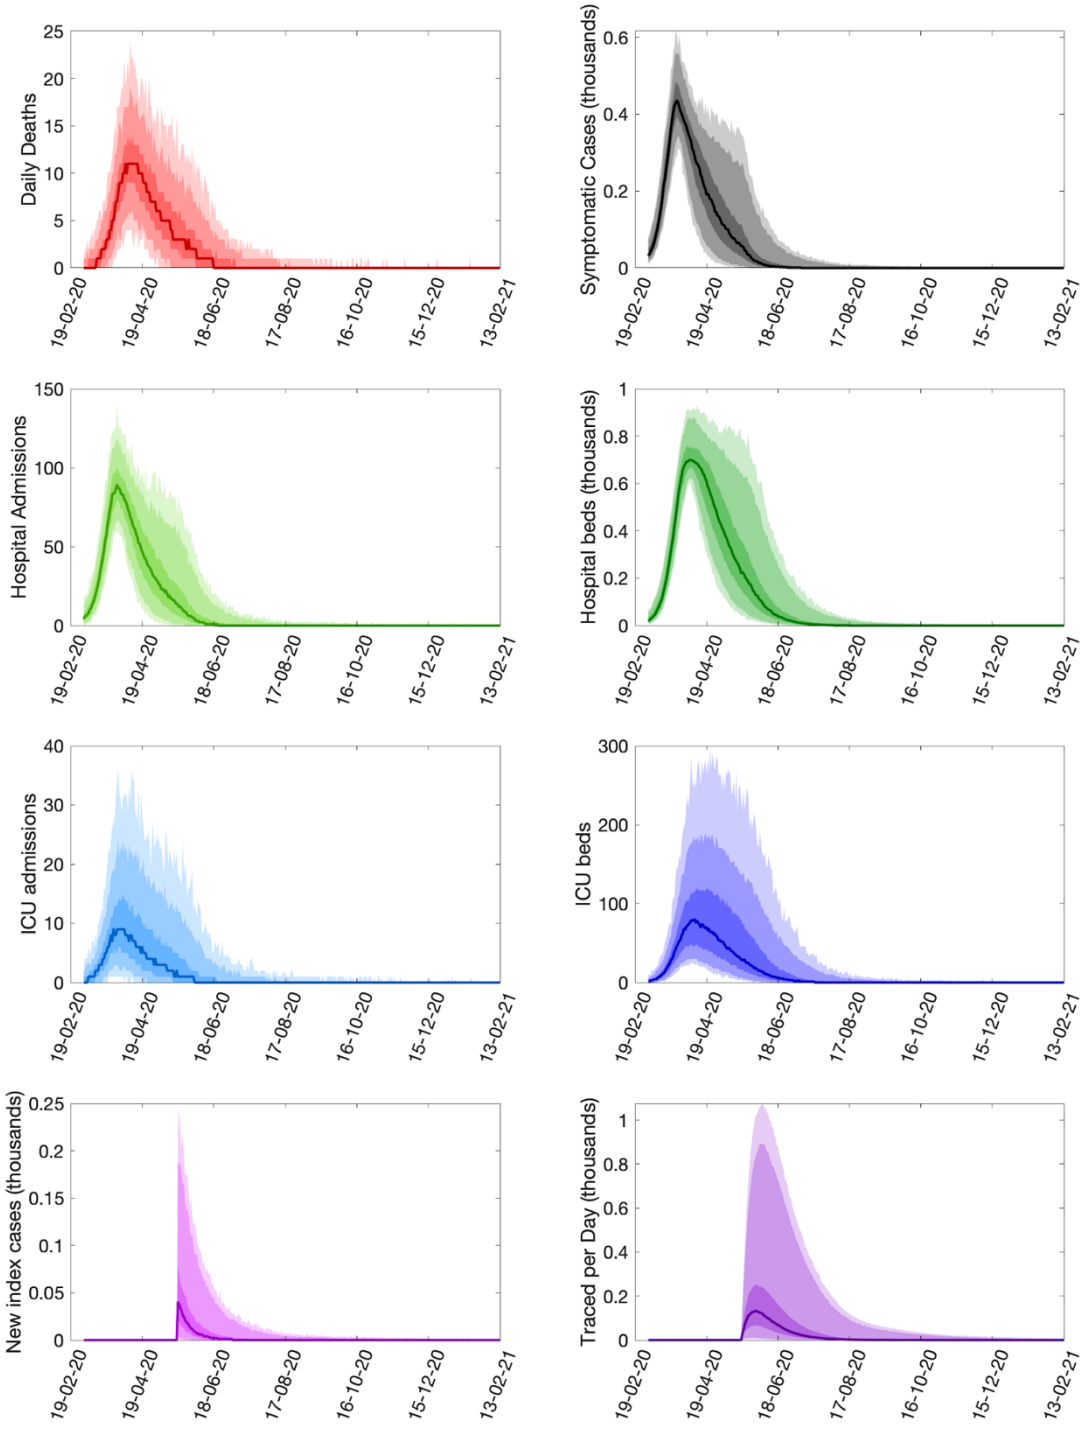

Scenario 1, maximum 30 contacts traced, London

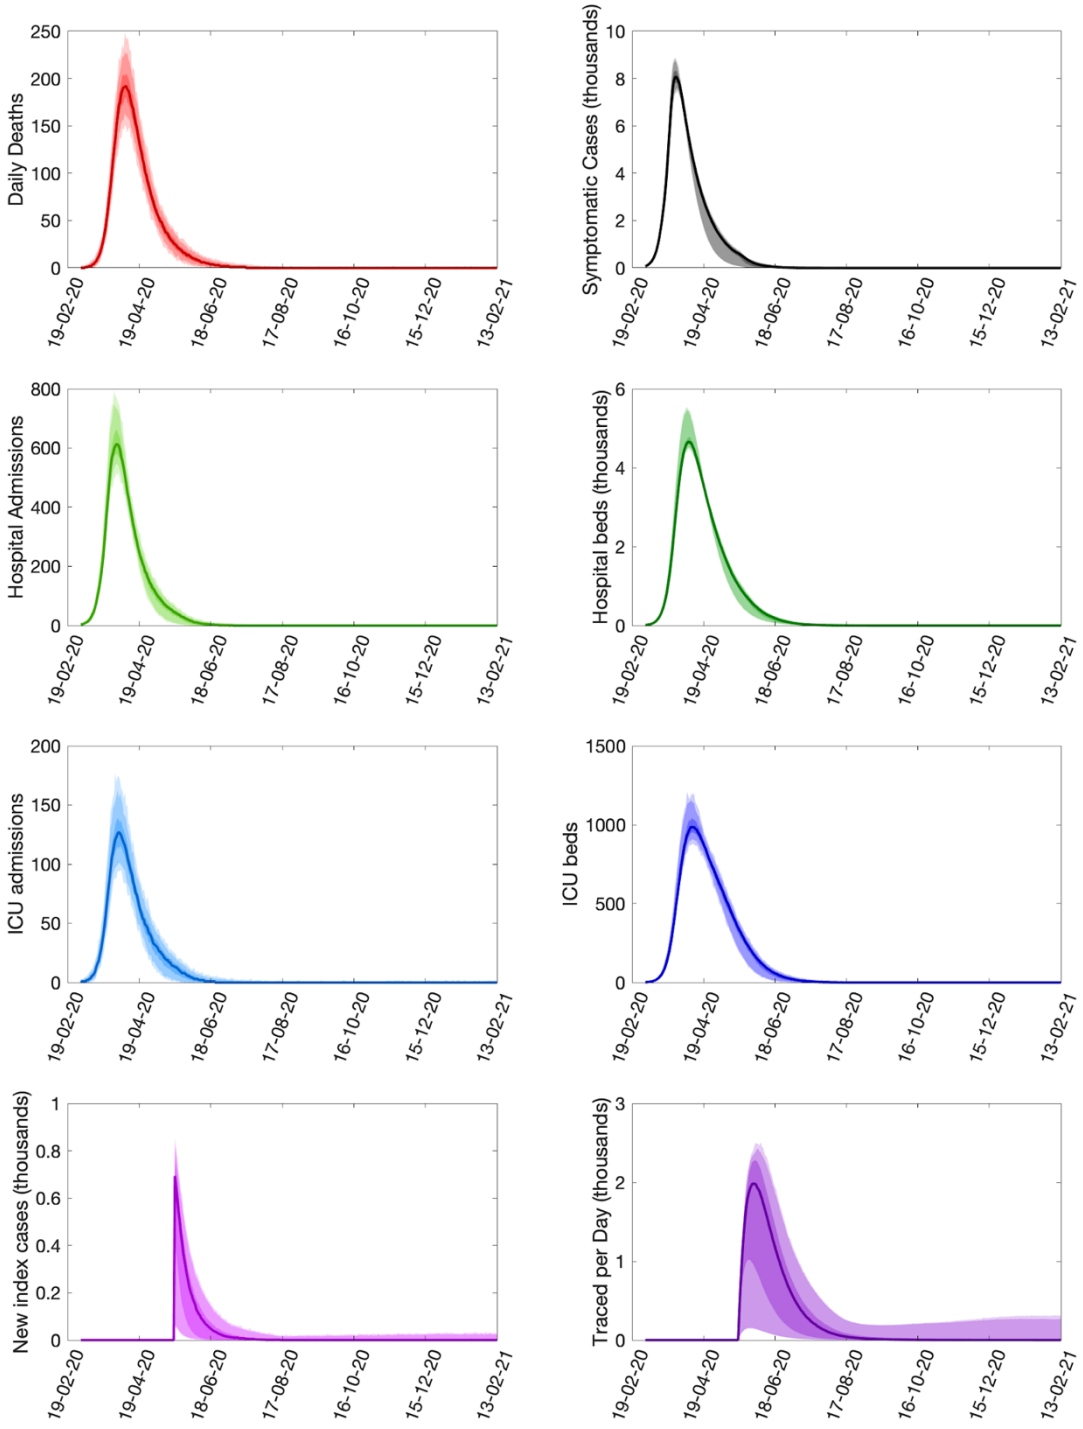

Scenario 1, maximum 30 contacts traced, Midlands

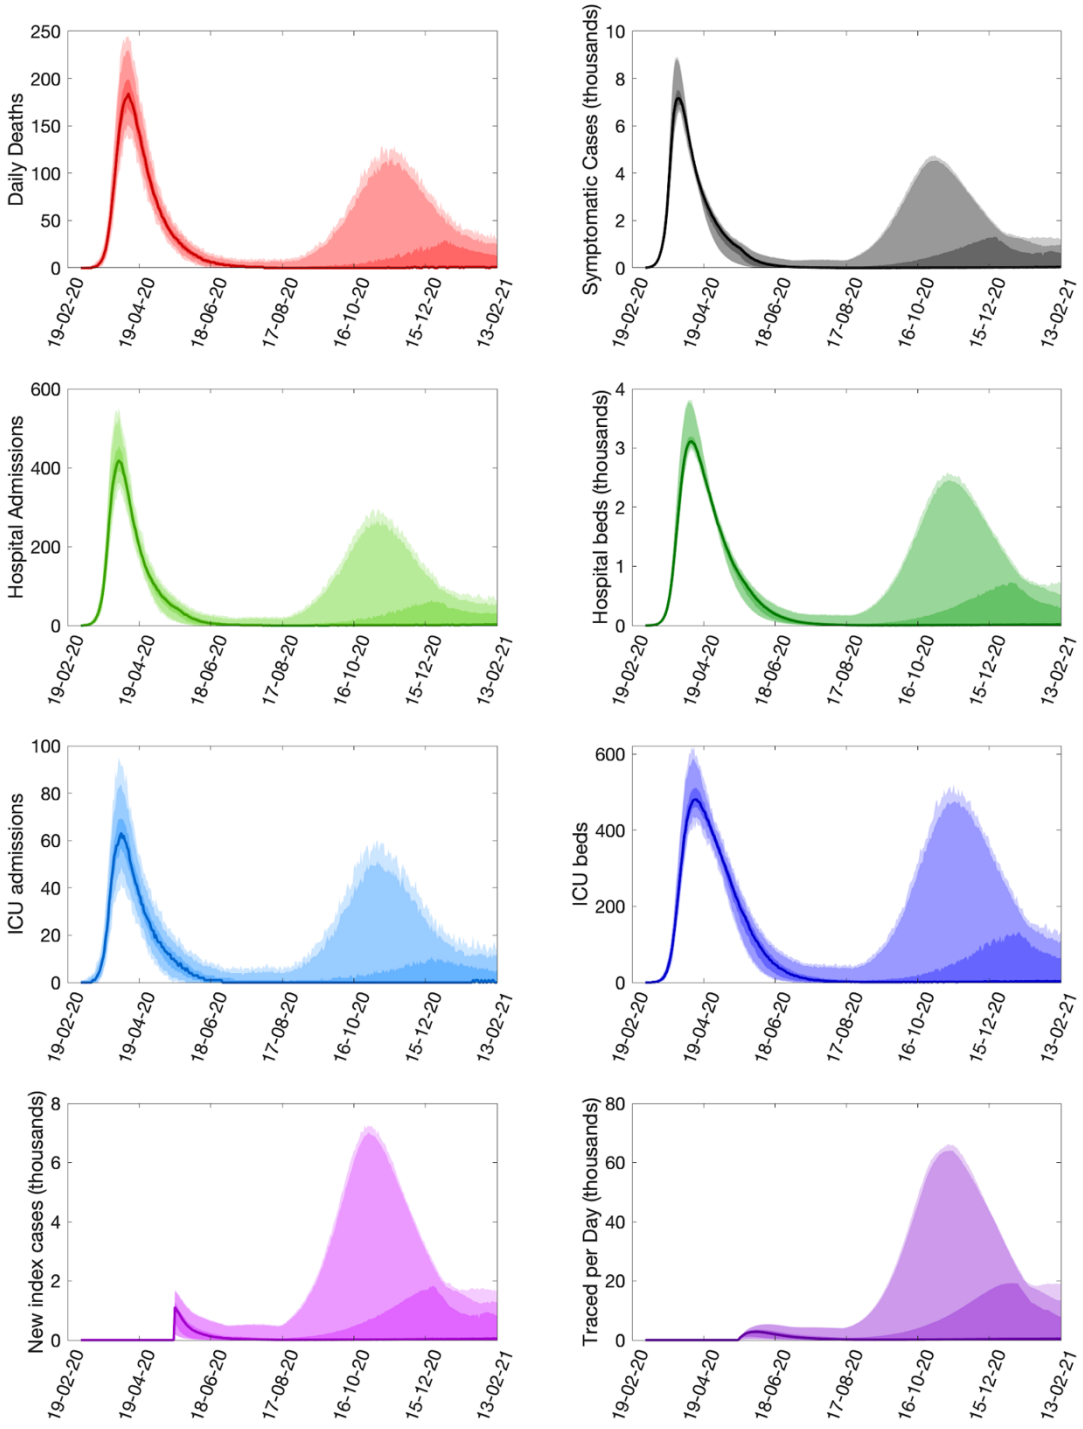

Scenario 1, maximum 30 contacts traced, East of England

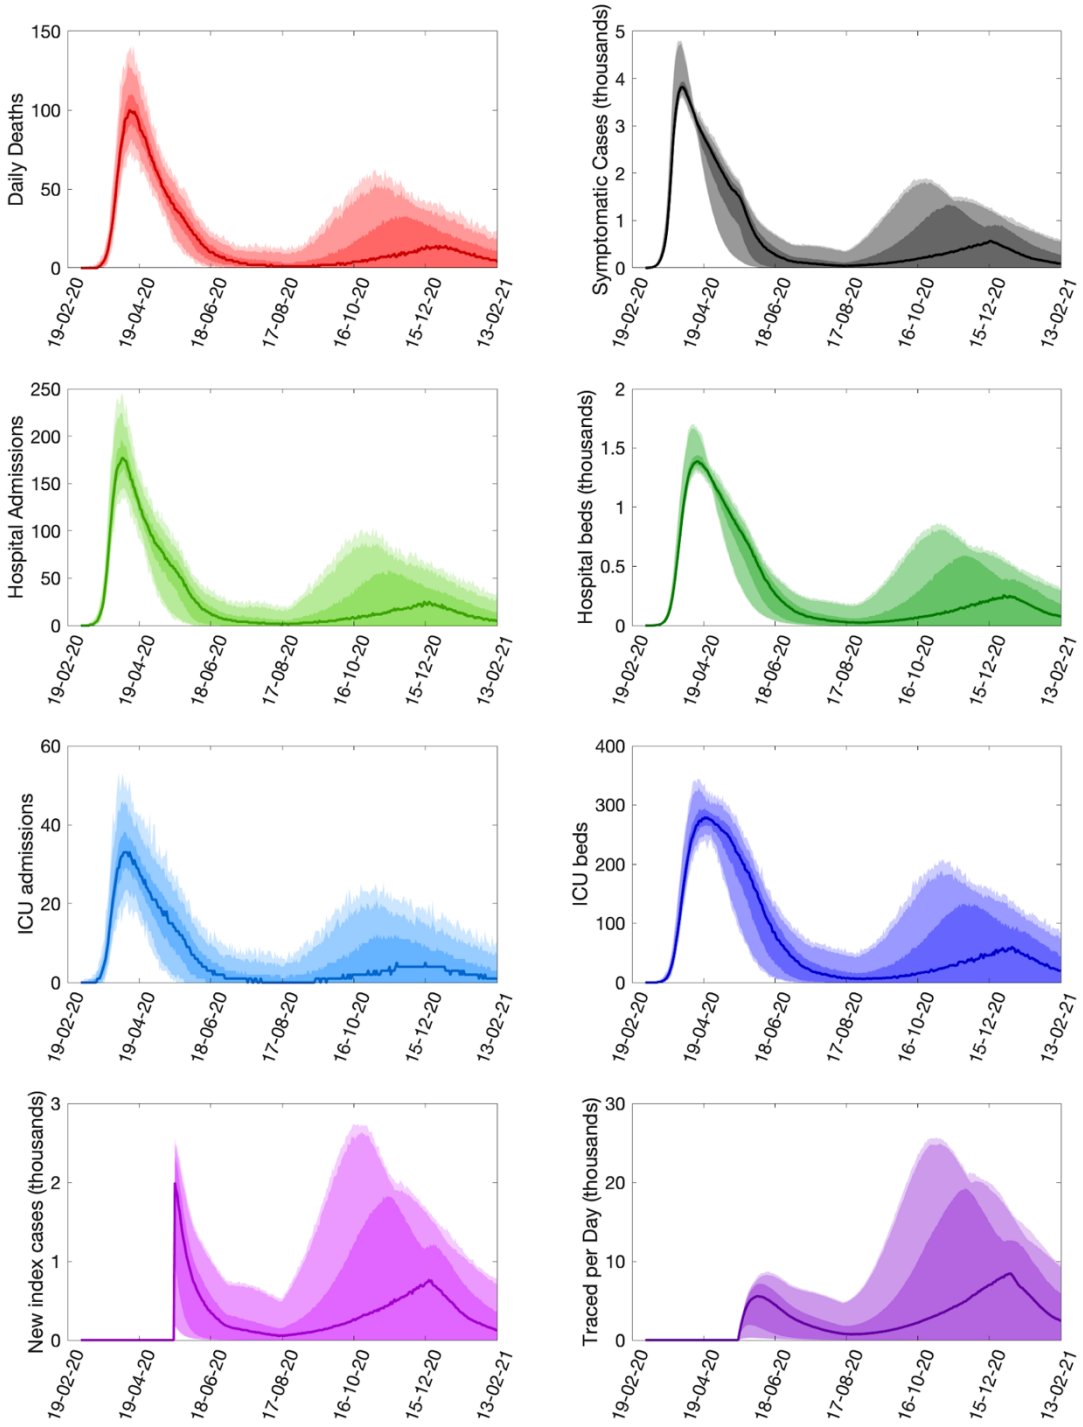

Scenario 1, maximum 30 contacts traced, North East and Yorkshire

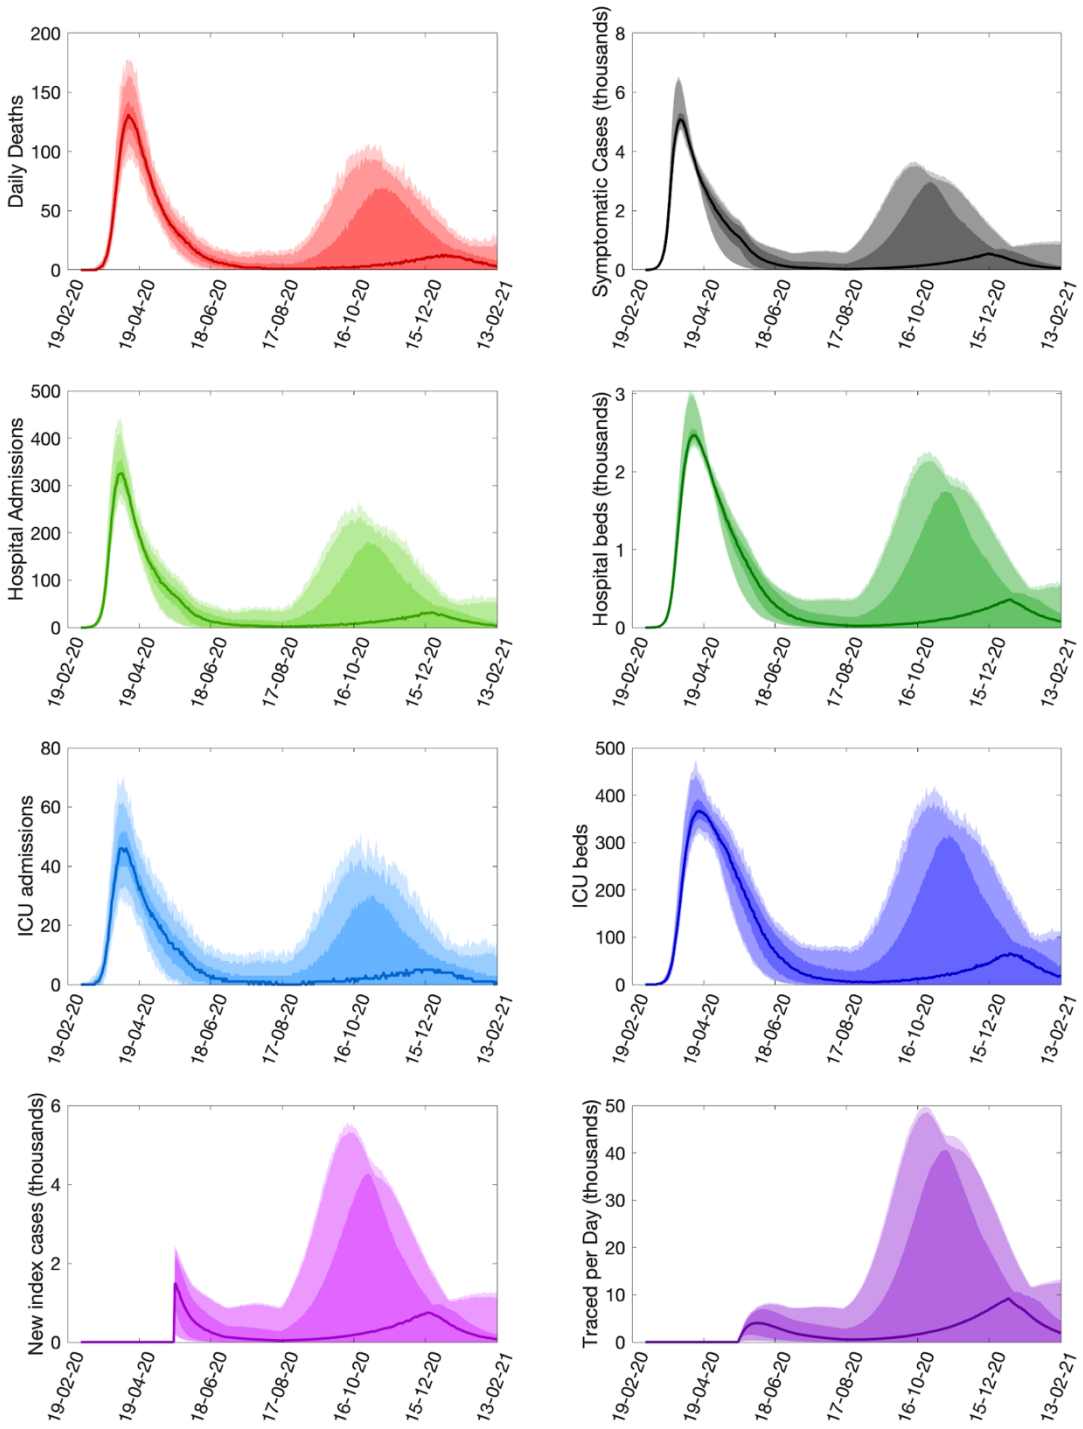

Scenario 1, maximum 30 contacts traced, North West

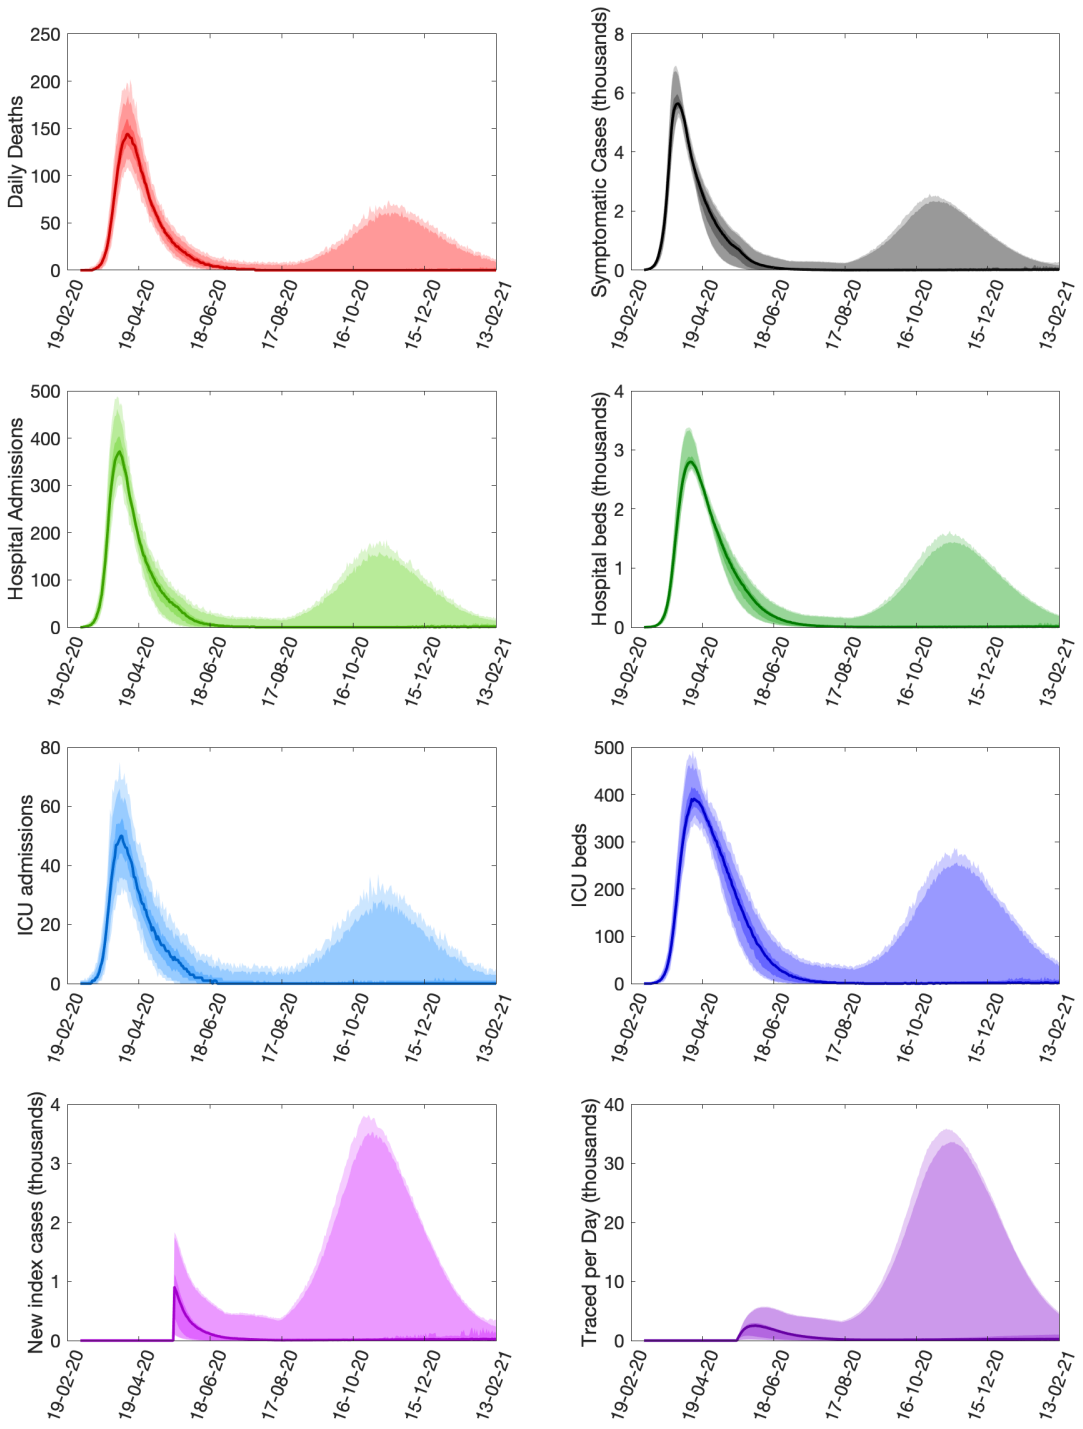

Scenario 1, maximum 30 contacts traced, South East

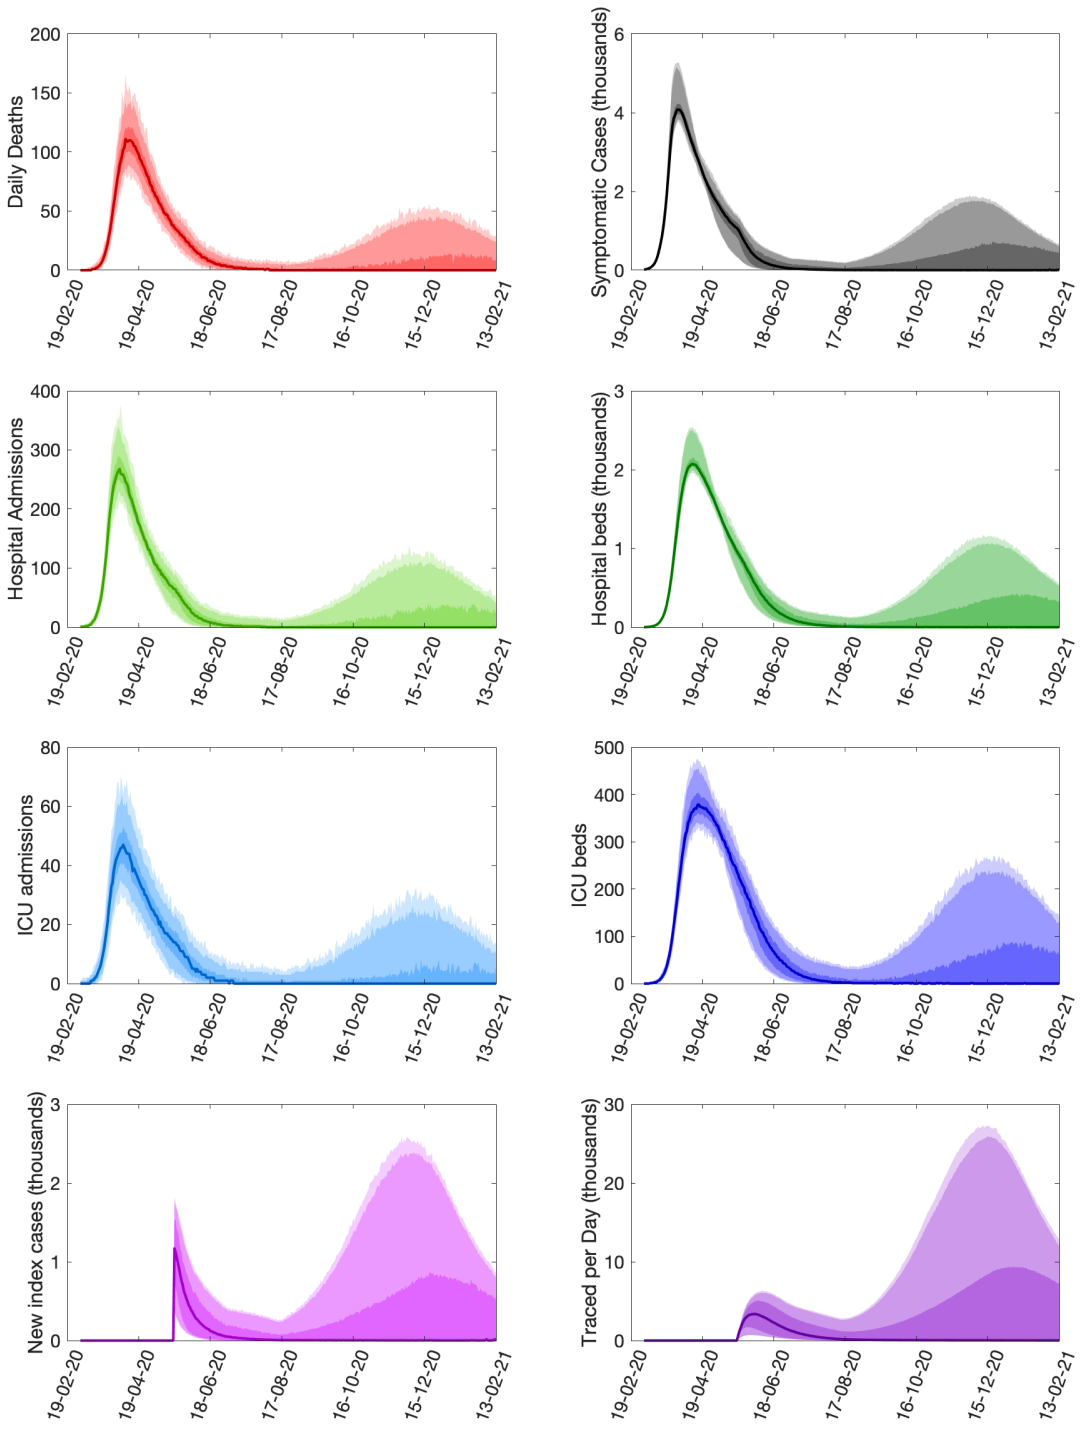

Scenario 1, maximum 30 contacts traced, South West