THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The effect of school closures on COVID mortality: old and new predictions

OPEN ACCESS

The effect of school closures on COVID mortality: old and new predictions

Ken Rice, professor of computational astrophysics, Ben Wynne, research fellow, Victoria Martin, professor of collider physics and Graeme Ackland, professor of computer simulation.

Author Affiliation

School of Physics and Astronomy, University of Edinburgh, Edinburgh EH9 3FD

Correspondence to: G.J. Ackland gjackland@ed.ac.uk

Submitted: 29th June 2020

Summary Box

What is already known on this topic:

Detailed models of individual interactions, which may take many hours of supercomputer time to run, are a reliable way to predict the course of an epidemic, and to investigate counterfactual scenarios.

The CovidSim model is the leading UK model, well-tested in influenza epidemics.

What the study adds: The predictions of the Imperial College model, including the inevitability of a second wave, were replicated and shown to be well realised in practice. Detailed investigation of the model show that it predicted that social distancing and school closures would suppress first-wave case numbers at the cost of a higher overall number of deaths, which number over 200,000 in all scenarios without a vaccine.

Print Abstract

Study Question: What information was available to government regarding school closure in March 2020 when the lockdown decision was taken.

Methods: We ran calculations using the CovidSim code which implements the Imperial College individual-based model of the COVID epidemic. This model was used to produce "Report 9", generally regarded as "The Science" behind the lockdown decision. We used only input data available in March 2020, but have adapted the code to more closely match the interventions which actually took place. By more detailed data analysis, we investigate the reason why general social distancing, school closures and isolation of younger people were predicted to increase the final number of deaths.

Study Answer and Limitations: CovidSim gives a good description of the epidemic, and predicts that isolation of less vulnerable people increases the final death toll. The excess

deaths are postponed to second and subsequent waves, and could be averted by a successful vaccination programme, which is not explicitly modelled.

What the study adds: We now know that the predictions of the Imperial College model, including the inevitability of a second wave, were well realised in practice. Furthermore, Government was already aware in March that social distancing and school closures would suppress first-wave case numbers at the cost of higher overall deaths.

1 Abstract

Objective: to establish what information was available to government when the lockdown decision was taken.

Design: Independent calculations using data known in March 2020 with the CovidSim code which implements the Imperial College individual-based model of the COVID epidemic.

Main Outcome Measures: Replication of summary data reported to SAGE. Detailed study of unpublished results, especially the effect of school closures.

Results: CovidSim would have given a good forecast of the subsequent data if initialised with a reproduction number (R0) about 3.5. We confirm the little-reported forecast that school closures and isolation of younger people were predicted to increase the final number of deaths, albeit postponed to a second and subsequent waves. We find that prompt interventions are highly effective at reducing peak ICU demand, but they also prolong the epidemic, in some cases causing more deaths long term. In the absence of an effective vaccination programme, none of the proposed mitigation strategies reduces the predicted total number of deaths below 200,000. This happens because COVID mortality is highly skewed towards older age groups.

Conclusions: It was predicted in March 2020 that a broad lockdown, as opposed to a focus on shielding the most vulnerable, would reduce immediate ICU demand at the cost of more deaths long-term. The optimal strategy for saving lives in a COVID pandemic is different from that anticipated for an influenza epidemic with different mortality age-profile.

2 Introduction

The UK national response to the coronavirus crisis has been widely reported as being primarily led by modelling based on work at Imperial College [1], although other models have been considered[1]. The key paper [2], which we will refer to as "Report 9", investigated a number of scenarios using this code with the best parameterisation available at the time. Contrary to popular perception, the lockdown which was then implemented was not specifically modelled in this work. As the pandemic has progressed, the parameterisation has been continually improved with new data as it arrives. The main

---

[1] throughout this paper, we maintain the distinction between epidemiological "model", and software implementations as "code"

conclusions of Report 9 were not especially surprising. COVID has a mortality around 1% [3], so an epidemic in a susceptible population of 70M people would cause many hundreds of thousands of deaths. In early March there may have been a case-doubling time of around 3 days in the UK [4], meaning that within a week COVID cases could go from accounting for a minority of available ICU spaces, to exceeding capacity. Furthermore, with an onset delay of over a week, and limited or delayed testing and reporting in place, there would be very little measurable warning of the explosion in ICU demand. However, in one table in Report 9 it is shown that closing schools reduces the R number, but has the unexpected effect of increasing total deaths. In this paper, we reproduce the main results from Report 9, and explain why, in the framework of the model, these counter-intuitive results were obtained. We chose not to attempt to re-parameterize the model, because we wanted to replicate the information available to policymakers at the time, specifically highlighting policies for which "suppressing the outbreak" and "saving lives" were conflicting choices.

3 Methods

The CovidSim model is developed from an influenza pandemic model [1,5,6]. The original code used for Report 9 has not been released. However, the Ferguson group has led an effort with Microsoft, GitHub and the Royal Society RAMP-initiative to recreate the model: this version has been stringently externally validated [7]. We used GitHub tagged version 0.14.0 + additional patches dated before 03-06-2020 to which we refer the reader for full technical details [8]. Input files relevant to Report 9 were supplied by Ferguson et al. [9] and were included in the GitHub release. CovidSim models the UK at the most detailed level possible without requiring personal data. The model simulates millions of individual "people" going about their daily business at home, within their community and at schools, universities, places of work, hospitals etc. The geographical representation of the UK is taken from census data, so the "people" in each area have appropriate distribution of age, health, wealth and household size. Simulated schools and workplaces have "people" with appropriate numbers, age distribution and commuting distances in line with national averages for each. The network of interactions is age dependent: people interact mainly with their own age group and with family, teachers and carers. The virus initially infects random members of this network of interacting co-workers, strangers, friends and family. Whenever an infected person meets a non-infected one, there is a probability that the virus spreads. This probability depends on the time and proximity of the interaction, and the infectiousness of the person given their stage of disease. Infected people may become hospitalised, and may die, with probability dependent on age, pre-existing conditions and stage of disease. This extremely detailed model is then parameterised using the best available expert clinical and behavioural evidence [5], with the coronavirus-specific features being updated as more data comes in from the worldwide epidemic [8]. Therefore, the model has the required complexity to consider non-pharmaceutical interventions (NPIs), which reduce the number of interactions between "people" in the model (see Table.1). To predict policymaking, it is assumed that these interventions are implemented when ICU bed occupation is observed to reach a particular "trigger" level. The model contains far more realistic detail than the data available. So results are averages over many runs with different starting conditions, all of

which are consistent with known data. The real epidemic is just one of these possibilities, so the code determines the range of scenarios which should be planned for. This is particularly important when there are low numbers of localised outbreaks: the prediction that local spikes will occur somewhere is reliable, and the most likely places can be identified, but predicting exactly when and where is not possible with the level of data available. All interventions reduce the reproduction "R" number, and slow the spread of the epidemic. However, a counter-intuitive result presented in Report 9 (their Table 3 and Table A1) is the prediction that, once all other considered interventions were in place, the additional closure of schools and universities would increase the total number of deaths. Similarly, adding general social distancing (SD) to a scenario involving case isolation and household quarantine, with appropriate estimates for compliance, was also projected to increase the total number of deaths.

3.1 Patient and public involvement

Patients or the public were not involved in the design, or conduct, or reporting, or dissemination plans of our research. All data used was retrieved from existing, public sources as referenced.

4 Results

The result tables for the scenarios presented in the original report were straightforwardly reproduced by averaging over 10 simulation runs with the same random number seeds as used in Report 9. The simulations are run for 800 days, with day 1 being 01 January 2020. The intervention period lasts for 3 months (91 days), with some interventions extended for an additional 30 days. The mitigation scenarios in Report 9 considered R_0=2.2 and R_0=2.4, but we initially only considered R_0= 2.4. As highlighted in [8] the results we obtain here are not precisely identical to those in Report 9, since they are an average over 10 stochastic realisations, the population dataset has changed to one that is open-source, and the algorithm used to generate the household-to-place network has been modified to be deterministic. The stochasticity gives a variance around 5% in total deaths and ICU demand, which explains the discrepancies with Report 9. More significant is the uncertainty of the timing of the peak of the infections, which is around ±5 days. We compare these predictions to the death rates from the actual trajectory of the disease [10,11]. We note that NHS England stopped publishing critical bed occupancy in March 2020[12], so it is not possible to compare ICU data from the model with reality.

In Table 1 we show the critical care (ICU) bed demand, while in Table 2 we show total deaths, both using the same mitigation scenarios as presented in Report 9. As in Report 9, for each mitigation scenario we consider a range of ICU triggers. In Table 1 we report the peak ICU bed demand across the full simulation for each trigger, as was presented in Report 9, but also include the peak ICU bed demand during the period of the intervention (first wave). The latter we define as the period during which general social distancing (SD) was in place, when implemented.

In Table 2 we also report the total number of deaths across the entire simulation, and also the number of deaths at the end of the first wave, again defined as the time at which general social distancing was lifted.

The full simulation numbers we present in Tables 1 and 2 are essentially the same as those presented in Table A1 in Report 9. As discussed earlier, the small difference between our numbers and those presented in Report 9 are probably because these are averaged over 10 stochastic realisations, the population dataset is slightly different, and the algorithm for generating the household-to-place network was changed to make it deterministic. Table 2 also illustrates the counter-intuitive result that adding school closures to CI_HQ_SDOL70 increases the total number of deaths across the full simulation. Moreover, it shows that social distancing of over-70s only is *more* effective than general social distancing.

It is clear from Tables 1 and 2 that in some mitigation scenarios peak ICU demand, and most deaths, occur during the period when the intervention is in place. There are, however, other scenarios where the opposite is true.

The reason for this is illustrated in Figure 2. The solid lines are the same mitigation scenarios as presented in Figure 2 of Report 9. We also show some additional scenarios (dashed lines) not shown in Figure 2 of Report 9, but which are included in Tables 1 and 2 and also in the Tables in Report 9.

In the simulations presented here, the main interventions are in place for 3 months and end on about day 200 (some interventions are extended for an additional 30 days). Figure 2 shows that some intervention scenarios lead to a single wave that occurs during the period in which the interventions are in place. Hence, the peak ICU bed demand occurs during this period, as do most deaths.

There are, however, some interventions that suppress the infection so that there is then a second wave once the interventions are lifted. For example, adding place closures to case isolation, household quarantine, and social distancing of those over 70 substantially suppresses the infection during the intervention period when compared to the same scenario without place closures. However, this suppression then leads to a second wave with a higher peak ICU bed demand than during the intervention period, and total deaths that exceed that of the same scenario without place closures.

We therefore conclude that the somewhat counter-intuitive results that school closures lead to more deaths are a consequence of the addition of some interventions suppressing the first wave, and failing to prioritise protection of the most vulnerable.

When the interventions are lifted, there is still a large population of people who are susceptible and a substantial number of people who are infected. This then leads to second wave of infections that can result in more deaths, but at a later time. Further lockdowns lead to a repeating series of waves of infection, unless herd immunity is achieved by vaccination, which is not considered in the model.

A similar result occurs in some of the scenarios involving general social distancing (SD). For example, adding general social distancing to case isolation and household quarantine

also strongly suppresses the infection during the intervention period, but then leads to a second wave that actually has a higher peak ICU demand than for the equivalent scenario without general social distancing.

Figure 3 provides an explanation for how place closure interventions affect the second wave, and why an extra intervention may result in more deaths than the equivalent scenario without this intervention. In the CI_HQ_SDOL70 scenario, without closures, a single peak of cases is seen. The data is broken-down into age groups, showing that younger people contribute most to the total cases, but that deaths come primarily from older groups. Adding the place closure intervention (and keeping all other things constant) gives the behaviour shown in the second row of plots. The initial peak is greatly suppressed, but the end of closures seems to prompt a second peak of cases amongst younger people. This then leads to a third, more deadly, peak of cases affecting the elderly when SDOL70 is removed. The postponement in the spread means there are more infectious younger people to infect the older age groups, a much larger fraction of whom then die.

One criticism of school closure is that reduced contact at school leads to increased contact at home; meaning children infect high-risk adults rather than low-risk children. We investigated this by increasing the infection rate at home to an extremely high value. Figure 1 shows that this makes an insignificant difference compared to the overall effect of adding school closures[2] to the other interventions.

4.1 CovidSim's description of a second wave

Although Report 9 does discuss the possibility that relaxing the interventions could lead to a second peak later in the year, we wanted to explore this in more detail, using the latest set of parameter files included in the GitHub repository [8].

The interventions we consider are place closures (PC), case isolation (CI), household quarantine (HQ) and general social distancing (SD) which are implemented using the PC_CI_HQ_SD parameter file. Specifically, we use the parameter file available in the data/param_files sub-directory of the GitHub respository. The only modification is to change the duration of the interventions to be 91 days.

These interventions start in late March (day 83) and last for 3 months (91 days). These simulations are also initialised so that there are about 15600 deaths by day 100 (April 9th) in all scenarios, mostly infected before the interventions were implemented.[3] This compares with Report 9 initiation which used then-reported deaths to March 14th.

The results are presented in Figure 4. The top panel shows cumulative deaths, with data from [11] and [13], while the bottom panel shows ICU bed demand per 100000 people. Although our simulations do include Northern Ireland, the available reported data does not. Therefore, the simulation results, and data, presented in Figure 4 are for England, Wales and Scotland only. We also consider a range of R_0 values and find that values higher

---

[2] Despite the description of place closure interventions in Table 2 of Report 9, university closures are not included in the (PC_)CI_HQ_SDOL70 scenario parameter files [9].

[3]This is implemented by modifying the [Number of deaths accumulated before alert] parameter in the preUK_2.0.txt parameter file.

than those considered in Report 9 best reproduce the data, with an R_0 value between 3 and 3.5 probably providing the best fit. This is consistent with the analysis presented in [14], but we acknowledge that the data could also be fitted by changes to the other scenario parameters. In both panels we also show the "Do nothing" scenario for $R_0 = 3.0$.

The ICU bed demand for the scenarios presented in Figure 4 show that the interventions are predicted to substantially reduce the ICU demand.

Random antibody tests at the time of writing suggest some 5% of the population have been exposed to coronavirus [13,15]. In the context of modelling a second wave, this is small. In the absence of interventions, predictions for the second wave are similar to those for the first. Assuming a similar response to a new wave, that exposure gives immunity, and that no vaccine become available, up to ten waves can be anticipated.

In practice, it seems that mandatory and voluntary interventions will continue, and maintain the reproduction number close to 1. This will keep ICU demand manageable, but it is worth noting that R=1 is also the value which prolongs the need for interventions for the longest time. At this level, the inhomogeneity of transmissions, particularly the unpredictability of superspreading events, becomes critical. Despite the level of detail of the model, there is insufficient data to model real people: we saw that for a major national epidemic this introduces an uncertainty of about 5 days in the predictions. At a local level, and with a lower R number, this uncertainty is greatly increased: it is impossible to predict when a particular town will suffer an outbreak (specifically, different towns are hit on different runs).

5 Conclusion

In this paper we used the recently released CovidSim code [8] to reproduce the mitigation scenarios presented in mid-March in Report 9 [2]. The motivation behind this was that some of the results presented in Report 9 suggested that the addition of extra interventions may actually increase the total number of deaths.

We find that the CovidSim code reliably reproduces the results from Report-9, and that the model underlying CovidSim can accurately track the UK death-rate data. To do so does require an adjustment to the parameters, a slightly higher R_0 than considered in Report 9, and results in an earlier start to the epidemic than suggested by Report 9. We emphasize, though, that the unavailability of these parameters in early-March is not a failure of the model.

We confirm that adding school and university closures to case isolation, household quarantine, and social distancing of those over 70 would lead to more deaths when compared to the equivalent scenario without the school and university closures. Similarly, general social distancing was also projected to reduce the number of cases but increase the total number of deaths compared with social distancing over 70s only. We note that, in assessing the impact of school closures, UK policy advice has concentrated on reducing total number of cases, not number of deaths [16].

The qualitative explanation for this is that within all mitigation scenarios in the model, the epidemic ends with widespread immunity with a large fraction of the population infected. Strategies which minimise deaths involve having the infected fraction primarily in the low-risk younger age groups, e.g. focussing stricter social distancing measures on care-homes where people are likely to die rather than schools where they are not. Optimal death reduction strategies are different from those aimed at reducing the ICU burden, and different again from those which lower the overall case rate.

We find that scenarios that are very effective when the interventions are in place, can then lead to subsequent waves during which most of the infections, and deaths, occur. Our comparison of updated model results with the published death data suggests that a similar second wave will occur later this year if interventions are fully lifted.

Since this paper was written, UK policy has moved to more local interventions. CovidSim models the geography of all towns, but the simulated people are only representative of the true population. This uncertainty means that the model cannot reliably predict which town will suffer an outbreak. Specifically, whereas the timing of the national outbreak is uncertain by days, the timing of an outbreak in a given town is uncertain by months. CovidSim is the most precise model available, but massively more personal data would be needed to obtain reliable local predictions.

Finally, we reemphasize that the results in this work are not intended to be detailed predictions for the second wave. Rather, we are re-examining the evidence available from CovidSim at the start of the epidemic. More accurate information is now available about the compliance with lockdown rules and age-dependent mortality. The difficulty in shielding care-home residents is a particularly important piece of health data that was not available to modellers at the outset.

Nevertheless, in almost all mitigation scenarios, CovidSim epidemics eventually finish with widespread immunity, and the final death toll depends primarily on the age distribution of those infected, not the total number.

References

[1] Ferguson NM, Cummings DA, Cauchemez S, Fraser C, Riley S, Meeyai A,et al. Strategies for containing an emerging influenza pandemic in Southeast Asia. Nature. 2005;437(7056):209–214.

[2] Ferguson NM. Report 9: Impact of non-pharmaceutical interventions(NPIs) to reduce COVID-19 mortality and healthcare demand. ImperialCollege London; 2020. https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/report-9-impact-of-npis-on-covid-19/.

[3] Grewelle R, De Leo G. Estimating the Global Infection Fatality Rate of COVID-19. medRxiv. 2020; Available from: https://doi.org/10.1101/2020.05.11.20098780.

[4] Pellis L, Scarabel F, Stage HB, Overton CE, Chappell LHK, Lythgoe KA, et al. Challenges in control of Covid-19: short doubling time and long delay to effect of interventions.

medRxiv. 2020; Available from: https://www.medrxiv.org/content/early/2020/06/11/2020.04.12.20059972.

[5] Ferguson NM, Cummings DA, Fraser C, Cajka JC, Cooley PC, Burke DS. Strategies for mitigating an influenza pandemic.Nature. 2006; 442(7101):448–452.

[6] Halloran ME, Ferguson NM, Eubank S, Longini IM, Cummings DA, Lewis B, et al. Modeling targeted layered containment of an influenza pandemic in the United States. Proceedings of the National Academy of Sciences.2008;105(12):4639–4644.

[7] Eglen S. CODECHECK certificate 2020-010; 2020. Available from: https://zenodo.org/record/3865491#.XuPW_y-ZPGI.

[8] Ferguson NM, Gilani GLN, Laydon DJ, Hinsley WR, et al. COVID-19CovidSim Model. GitHub. 2020; Available from: https://github.com/mrc-ide/covid-sim.

[9] Ferguson NM. github; 2020. Available from: https://github.com/mrc-ide/covid-sim/tree/master/report9.

[10] Office of National Statistics. UK Government; 2020. Available from:https://www.gov.uk/guidance/coronavirus-covid-19-information-for-the-public#number-of-cases-and-deaths.

[11] of Scotland NR. Scottish Government; 2020. Available from:https://www.nrscotland.gov.uk/covid19stats.

[12] England N.NHS England; 2020. Available from: https://www.england.nhs.uk/statistics/statistical-work-areas/critical-care-capacity/.

[13] Connors E, Fordham E. Coronavirus (COVID-19) Infection Survey pilot: England and Wales, 21 August 2020; Availablefrom:https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/bulletins/coronaviruscovid19infectionsurveypilot/englandandwales21august2020.

[14] Flaxman S, Mishra S, Gandy A, Unwin HJT, Coupland H, Mellan TA,et al. Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in European countries. Nature. 2020;584:257–261.Available from: https://doi.org/10.1038/s41586-020-2405-7.

[15] Ward H, Atchison C, Whitaker M, Ainslie K, Elliott J, Okell L, et al.Antibody prevalence for SARS-CoV-2 following the peak of the pandemic in England: REACT2 study in 100,000 adults. medarXiV. 2020; Available from: https://www.medrxiv.org/content/10.1101/2020.08.12.20173690v2.

[16] Interdisciplinary Task and Finish Group on the Role of Chil-dren in Transmission:Modelling and behavioural science responses to scenarios for relaxing school closures; Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/886994/s0257-sage-sub-group-modelling-behavioural-science-relaxing-school-closures-sage30.pdf.

What is already known on this subject

The Covid-Sim model is the most detailed individual-based model of the UK appropriate for simulation of the spread of an epidemic.

The UK-wide lockdown was implemented as a highly effective way of reducing epidemic spread.

What this study adds

The model used for "Report 9" predicts that, in the absence of a vaccine, school closures result in more overall deaths than not closing schools.

The code used, and results obtained in "Report 9" are independently verified and provided a good description of the subsequent spread of the epidemic at the national level, except that the R_0 parameter was set too low.

Mitigating a COVID epidemic requires different strategy from an influenza epidemic, with more focus on shielding the elderly and vulnerable.

While total infections are at a low level, coronavirus manifests as localised spikes. Currently available data is insufficient to reliably predict exactly where these will occur.

Contributors

KR and BW ported and validated the code across a number of computer architectures, performed the calculations and made the figures. VM supervised the testing and pre-Opensourcing test of the CovidSim code. GJA designed and supervised the project. All authors contributed to writing the paper. The corresponding author attests that all listed authors meet authorship criteria and that no others meeting the criteria have been omitted.

We would like to thank Kenji Takeda and Peter Clarke for help with the code, and Neil Ferguson for sharing data.

Competing Interests

All authors have completed the ICMJE uniform disclosure form at www.icmje.org/coi_disclosure.pdf and declare: support from UKRI for the submitted work; no financial relationships with any organisations that might have an interest in the submitted work in the previous three years; no other relationships or activities that could appear to have influenced the submitted work.

Patients and Public statement

Patients or the public were not involved in the design, or conduct, or reporting, or dissemination plans of our research.

Ethical approval

No ethical approval was required for this research.

Data sharing

The full simulation and datasets can be accessed and run from GitHub using the SHA1 hashcode 92d414769c6387a08ab65d9830f7f9775fdd3a71

Code examples and raw data sufficient to reproduce all results in this research are available at https://doi.org/10.7488/ds/2912

Transparency

The lead author affirms that the manuscript is an honest, accurate, and transparent account of the study being reported; that no important aspects of the study have been omitted; and that any discrepancies from the study as planned have been explained.

Dissemination to participants

Since this research uses public demographic data for the whole of the UK, there are no plans for dissemination of this research to specific participants, beyond publishing it.

Provenance and peer review

Not commissioned; externally peer reviewed.

Trigger	Time	PC	CI	CI_HQ	CI_HQ_SD	CI_SD	CI_HQ_S DOL70	PC_CI_HQ_SD OL70
0.1	1st wave	152	119	87	8	20	62	33
0.1	total	152	119	87	115	84	62	**51**
0.3	1st wave	153	119	87	10	22	62	34
0.3	total	153	119	87	115	73	62	**48**
1	1st wave	154	119	87	11	22	62	35
1	total	154	119	87	104	59	62	**37**
3	1st wave	159	119	87	13	22	62	37
3	total	159	119	87	82	40	62	**37**

Table 1: Table showing peak ICU bed demand (UK-wide, in thousands) for different intervention scenarios: home isolation of suspect cases (CI), home quarantine of family members (HQ), general social distancing (SD), social distancing of those over 70 (SDOL70) and "place closures" (PC), specifically the closure of schools and universities. More details of these NPIs are provided in Table 2 of Report 9, which we reproduce in Appendix Figure 5. For each trigger value of cumulative ICU cases (again in thousands) we show the peak ICU demand, and the peak during the first wave when the interventions were in place (which is sometimes the same).

Trigger	Time	PC	CI	CI_HQ	CI_HQ_SD	CI_SD	CI_HQ_S DOL70	PC_CI_HQ_SD OL70
0.1	1st wave	418	354	252	21	39	177	75
0.1	total	496	416	355	440	402	**262**	357
0.3	1st wave	456	378	281	32	58	200	104
0.3	total	495	416	355	437	390	**261**	356
1	1st wave	479	398	310	48	86	223	139
1	total	494	416	355	428	370	**261**	351
3	1st wave	490	407	325	70	114	237	172
3	total	495	416	355	411	347	**262**	342

Table 2: Table showing total deaths (UK-wide, in thousands) for different intervention scenarios and different ICU triggers. For each trigger value of cumulative ICU cases (thousands) we show the total deaths across the full simulation, and during the first wave. Bold numbers are the minimum achievable

Figure 1: Effect of place closure. The CI_HQ_SDOL70 and PC_CI_HQ_SDOL70 intervention scenarios are compared. After the trigger at 100 cumulative ICU cases, all the interventions are in place for 91 days: the general social distancing runs to day 194, and the enhanced social distancing for over 70s runs for an extra 30 days. With Place Closure (PC), we also show the effect of increasing the amount of in-household interactions by a factor (home) of up to 2. %the value of the relative household contact parameter is varied from 1.0 to 2.0 This shifts cases from first to later waves, but the additional PC intervention always leads to an increase in total cases and deaths.

Figure 2: Flattening the curve. The solid lines are the same scenarios as presented in Figure 2 of report 9. We also show three additional scenarios (dashed lines) for $R_0 = 2.4$ which are summarised in Tables 1 and 2. The PC_CI_HQ_SDOL70 scenario minimises peak ICU bed demand, but prolongs the epidemic,

resulting in more ICU cases and deaths. These illustrate why adding place closures (PC) to a scenario with case isolation (CI), household quarantine (HQ) and social distancing of those over 70 (SDOL70) can lead to more deaths than the equivalent scenario without place closures. Doing so suppresses the infection when the interventions are present, but leads to a second wave when they are lifted, which happens on around day 200. The total number of deaths in the CI_HQ_SDOL70 scenario is 260,000, while for PC_CI_HQ_SDOL70 it is 350,000. Similarly, comparing general social distancing (SD) with equivalent scenarios without SD, the second wave peak in the CI_HQ_SD scenario is actually higher than the first wave peak in the CI_HQ scenario.

Figure 3: Simulated values for daily virus cases (left) and deaths (right), for scenarios CI_HQ_SDOL70 (top) and PC_CI_HQ_SDOL70 (bottom). Interventions are triggered by reaching 100 cumulative ICU cases. After the trigger, all the interventions are in place for 91 days: the general social distancing runs to day 194, and the enhanced social distancing for over 70s runs for an extra 30 days. Results are broken down into age categories as indicated, with SDOL70 interventions affecting the three oldest groups. In the CI_HQ_SDOL70 scenario we see a single peak of cases, with greatest infection in the younger age groups but most deaths occurring in the older. In the PC_CI_HQ_SDOL70 scenario we see three peaks in the plot of daily cases, with the first peak occurring at a similar time for CI_HQ_SDOL70 above, but with reduced severity. The second peak seems to be a response to the ending of Place Closure (PC), and most affects the younger age groups, therefore having little impact on the total deaths. The third peak affects the older groups, leading to a significant increase in the total deaths.

Figure 4: Refit of the CovidSim March parameterization based on death data through to June. The top panel shows cumulative deaths, with data from [11] and [13], while the bottom panel shows ICU bed demand per 100000 people. We considered a range of R_0 values and find that values higher than that considered in Report 9 best reproduce the data. A good fit also requires us to assume that the epidemic started earlier than was previously suggested in Report 9. We see that CovidSim provides a good fit to the data with a value of R_0 between 3 and 3.5 and (inset) predicts that the ICU demand would probably be limited to around 10 per 100000.

Figure 5: Table defining the interventions considered in CovidSim copied from Report 9.