

Optimal COVID-19 lockdown strategies in an age-structured SEIR model of Northern Ireland

ABERNETHY, Gavin and GLASS, David

Available from Sheffield Hallam University Research Archive (SHURA) at:

<https://shura.shu.ac.uk/29760/>

This document is the Supplemental Material

Citation:

ABERNETHY, Gavin and GLASS, David (2022). Optimal COVID-19 lockdown strategies in an age-structured SEIR model of Northern Ireland. *Journal of the Royal Society Interface*, 19 (188). [Article]

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Optimal COVID-19 lockdown strategies in an age-structured SEIR model of Northern Ireland

Supplementary Material

Gavin M Abernethy*, David H Glass†

Journal of the Royal Society Interface

*G.M.Abernethy@shu.ac.uk

†dh.glass@ulster.ac.uk

Contents

1	Model description	3
1.1	Distinct periods of the pandemic	3
1.2	Contact matrices	3
1.3	Estimated age-dependent IFR and maximum deaths	4
1.4	Parameters chosen by fitting the model	5
2	Best fit simulation	6
2.1	Parameter searches	6
2.2	Residual re-sampling and confidence intervals	6
2.3	Simulation results	6
3	Alternate-history scenarios	11
4	Best fit simulations with alternative parameter choices	13
4.1	Modifying the incubation rate σ	14
4.2	Modifying the progression and recovery rates γ	15
4.3	Modifying the relative transmissibility from clinical and subclinical infectious cases	17
5	Role of intensity and duration of lockdowns early in the pandemic	21
5.1	Closed system	22
5.2	Constant inflow	24
5.3	Ten evenly-spaced lockdowns	26
6	Mechanistic activation of lockdowns	27
6.1	Closed system	28
6.2	Constant inflow	32
6.3	Inpatient and ICU occupant ceiling times	36
6.4	Long-term simulation of the best possible interventions	37
6.4.1	Closed system	37
6.4.2	Constant inflow	39
6.5	Understanding closed systems versus those with a constant inflow of cases	41
6.6	Comparison after 1000 days or 2000 days of lockdown	43
6.7	Vaccination	45

1 Model description

In this section we present additional details regarding the construction of the age-structured SEIIR model calibrated and used in this study to simulate the COVID-19 pandemic in Northern Ireland.

1.1 Distinct periods of the pandemic

The first 400 days beginning January 1st 2020 are divided into nine distinct time periods based on significant public policy decisions by the devolved administration in Northern Ireland (Table 1).

Period	Description	Schools
1st January - 23rd March 2020	No restrictions	Open
24th March - 11th June 2020	First lockdown	Closed
12th June - 16th August 2020	Shops re-open	Closed
17th August - 16th October 2020	Pupils begin return	Open
17th October - 1st November 2020	Additional restrictions	Closed
2nd November - 26th November 2020	Schools re-open	Open
27th November - 10th December	2-week circuit break	Open
11th December - 25th December 2020	Circuit break ends	Open
26th December - 4th February 2021	6-week lockdown begins on Boxing Day	Closed

Table 1: Distinct transmission time periods

1.2 Contact matrices

Data for the number of contacts made between individuals of each age class are derived from the major empirical study by Prem et al (see main report for reference), which estimated separate contacts for the “home”, “work”, “school” and “other” environments and presented their data for 5-year age classes only up to age 80. For use in this study, we condensed the average of the matrices for the UK and Ireland into 4×4 matrices of 20-year age classes, and extrapolated the data for the 75-80 age group to estimate contacts for the 80+ age class which was absent from the empirical data. This yields four 5×5 setting-specific matrices. However, we cannot be certain that the 80+ age group behaves in precisely same manner as those aged 75-80, with some residing in communal care homes which were uniquely affected by COVID-19, and in particular if we proceed without further modification we find that the model significantly under-predicts the amount of infections, hospitalisations and deaths occurring in those aged over 80 for whom we did not have explicit empirical data. To resolve this, we scale the amount of contact from all individuals aged 40+ (accounting for care home workers, visitors, and residents) to the 80+ age group by a parameter μ_e which will be learned by fitting to the data. Given the small proportion of the population that is aged 80+ and their already small contact with those aged 0-40, transmission from those aged 80+

to the younger age classes should be relatively small¹.

During each period of the simulation, the time-dependent final contact matrix $c = \{c_{i,j}\}$ is assembled by combining the four matrices subject to time-specific scalings. During the first period, before restrictions were put in place, we assume normal behaviour and sum the four matrices. If schools are closed (see Table 1), the corresponding matrix is omitted. During the three periods of “lockdown” or “circuit break”, the work matrix is also removed due to working-from-home and furlough. Otherwise, it is not obvious to what extent social gatherings and work practices may have changed, so the work matrix and the other contacts matrix are scaled by parameters μ_w and μ_o respectively, which will also be determined by fitting the model.

1.3 Estimated age-dependent IFR and maximum deaths

Using the parameters selected *a priori*, we may determine the probability of death for a given case of COVID-19 (symptomatic or otherwise), and apply this to the demographics of Northern Ireland to obtain upper limits for the number of deaths by age and in total in the event that the entire population is eventually infected². Note that these calculations are based entirely on the estimates from the sources in Table 1 of the main body of the paper, and not on original simulations.

Age class	Infection fatality ratio	Maximum possible deaths
0 - 20	0.0001	39
20 - 40	0.0004	215
40 - 60	0.0055	2721
60 - 80	0.0848	28727
80+	0.2112	17227

Table 2: Maximum deaths per age class

Combining the values in Table 2 yields the total possible deaths:

$$\sum_{i=1}^5 P_i \epsilon_i h_{1,i} (d_{1,i} + h_{2,i} d_{2,i}) = 48929$$

and an overall IFR of 0.0258.

¹As there do not exist specific contact matrices for the population of Northern Ireland, it is already not possible to use a true contact matrix that satisfies the reciprocity condition $P_i c_{i,j} = P_j c_{j,i}$, and this change exacerbates this disparity by around 6% only. These matrices are thus only approximations of net behaviour.

²Due to herd immunity, this would not naturally occur unless the entire population was repeatedly exposed for an infinite amount of time.

1.4 Parameters chosen by fitting the model

We select empirically-evidenced constant values for as many parameters as possible. However, some unknowns remain that must be fitted by simulation. These include the values of the transmission rate β at different periods of the pandemic, and the day and age-class of the initial infection. Backdated data from the Department of Health estimates that at least one individual may already have been in hospital with COVID-19 before the day that “patient zero” was reported in the news media to have arrived in Northern Ireland, so it is not clear which was the first day the virus was present.

Given the difficulties of accurately monitoring and reporting case numbers quickly and assigning a singular cause of death, the empirical data used to fit to the model is the 7-day rolling average of daily new hospital admissions for each of the five age classes, on each of 400 days starting January 1st 2020. Empirical data is publicly available from Northern Ireland’s Department of Health (citation in the main report), giving a total of 2000 data points. We computationally select the free parameters that minimise the total sum of the square residuals between the predicted and actual values of these points.

During an individual simulation of the model, 11 parameters for each period of the pandemic are fitted sequentially³. Three parameters are fitted during the first transmission period (the first 83 days from January 1st, as per the first row of Table 1): the transmission rate β_1 , the initial day of the first case in January 2020, and the age class of this first infection. Next one parameter, a new transmission rate β_k , is fitted during each of the subsequent eight time periods k . However, as there is a lag between the effects of restrictions or relaxations and the impact on hospitalisation rates, for an accurate fit when determining the value of β_k for time period k , we iterate from the first day of period k to the final day of period k plus an additional D days (to account for further possible delay between the change of regime and the observable effect on hospitalisations) and seek the value of β_k that minimises the error in hospital admissions over this extended interval. When the best such value of β_k has been identified, we then reset to the first day of period $k + 1$ that resulted from that choice and begin simulating the subsequent period.

The length of the delay D that yields the best overall results is one of four additional parameters that are fitted across multiple simulations. The others are the 80+ contact scaling μ_e , the work contact reduction parameter μ_w , the “other contacts” reduction parameter μ_o . Each of these are set to fixed values for an individual simulation while the best fitting values of the other 11 parameters are determined, and then the process is repeated for multiple values of these four parameters. Thus, in total there are 15 parameters to be determined by fitting: 11 within an individual simulation, and 4 across simulations.

³Whilst this yields a less accurate fit than if all possible parameter combinations were tested simultaneously, it hugely reduces the computational time required to explore a 15-dimensional parameter space and identify a reasonable fit.

2 Best fit simulation

2.1 Parameter searches

We initially tested 59,532 sets of parameters: 12 values of delay D between subsequent β -periods, from 0 to 11; 41 values of scaling μ_e for the 80+ contacts, from 1 to 5; 11 values of work contact reduction μ_w , from 0 to 1; and 11 values of other contact reduction μ_o , from 0 to 1. On identifying the set which yielded the least sum of square residuals, a finer search was conducted around those values.

Within each simulation, 31 days (January 1st - January 31st) and each of the five age classes are tested as the single initial infectious individual that yields the most accurate projections, along with 100 values of each of $\beta_1 \dots \beta_9$ over a suitable range (the order of β_1 is first determined by numerical experimentation, and subsequently $\beta_2 \in [0, 1.5\beta_1]$, and $\beta_3, \dots, \beta_9 \in [0.1\beta_2, 1.5\beta_1]$). Note that increasing the number of increments of β_k tested initially results in a worse fit overall, as the greatest divergence occurs in the later periods and a good fit earlier in the simulation does not necessarily enable a better fit during later sections.

2.2 Residual re-sampling and confidence intervals

Having located the point in 15-dimensional parameter space that gives the best fit of the rolling average of daily hospitalisations per age group that minimises the sum of the square residuals, we identify confidence intervals for the nine β parameters that are fitted within a given simulation. This is achieved by a bootstrapping method of resampling residuals:

1. For each of 30,000 resamples:
 - (a) Modify the predicted daily hospitalisations by randomly sampling from the 1850 residuals (with replacement) and adding one of these residuals to each of the 1850 predicted data points (one for each age class from February 1st 2020) from the original best fit simulation.
 - (b) This modified predicted data set is then used in place of the empirical data as the curve fitting procedure is run again, with the initial infection day and class fixed.
 - (c) This yields a new set of nine β parameters for each of these 30,000 resamples.
2. Thus we obtain a set of 30,000 estimates of each of the fitted β parameters, in addition to the original best fit value.
3. Finally, we calculate the 95% confidence interval using the percentile bootstrap method.

2.3 Simulation results

In this section, figures illustrating the best fit simulation shown in §3 of the main report are reproduced, along with the additional Figures 3(a), 5 and 6. The fitted values of the 15 parameters and their confidence intervals are shown in Table 3. Note that confidence

intervals of non-zero width can be obtained for β_1 and β_2 if the number of increments of β tested in each case is increased from 100 to 1000, but this is different from the curve fitting procedure used to obtain the principle best fit.

Across simulations		Within simulations		
Parameter	Value	Parameter	Value	95% confidence interval
D	2	Initial class	5	
μ_e	3.00	Initial day	30	
μ_w	0.00	β_1	0.1200	0.1200 - 0.1200
μ_o	0.95	β_2	0.0288	0.0288 - 0.0288
		β_3	0.0666	0.0542 - 0.0720
		β_4	0.0773	0.0702 - 0.0950
		β_5	0.0401	0.0135 - 0.0560
		β_6	0.0560	0.0418 - 0.0897
		β_7	0.0418	0.0029 - 0.0914
		β_8	0.0914	0.0153 - 0.1800
		β_9	0.0489	0.0259 - 0.0897

Table 3: Fitted parameters

Note that care must be taken when interpreting the contact matrix parameters. In particular, it seems extreme that all work contact effectively ended ($\mu_w = 0$), while all other contact continued nearly unchanged ($\mu_o = 0.95$). While work-based contacts have reduced significantly, this is likely also due to the model correcting the overall age-structure of the contact matrices, as the sources used for this are an estimation with limited accuracy, and furthermore, they do not account for other changes in patterns of social behaviour during the pandemic, which may be age-dependent (e.g. individuals of particular ages may have made more or less changes to avoid meeting friends). Further testing shows that without the effective removal of the work matrix (which may be capturing a combination of all of these effects) the model has a strong tendency to over-estimate the number of cases and hospital admissions of individuals aged 40-60.

Figure 1 shows the best fit of the model to the 7-day rolling average of the daily recorded hospital admissions over the 400 days of simulation, and the empirical data that is being fitted against.

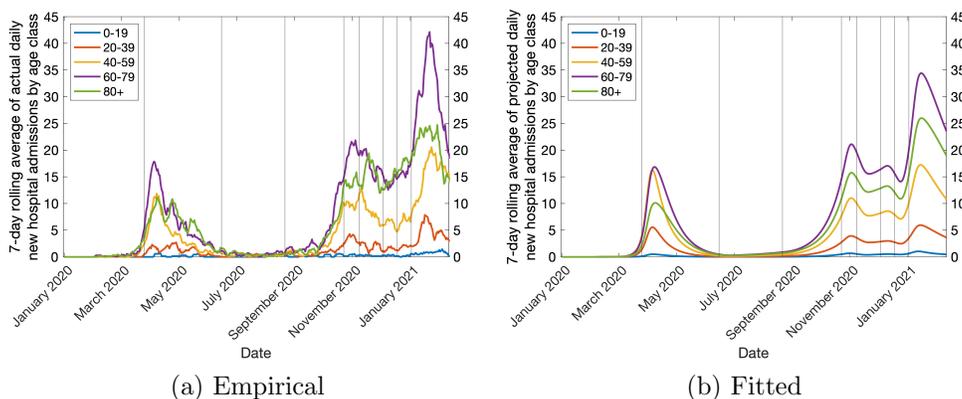


Figure 1: 7-day rolling average of daily new hospital admissions

Figure 2 illustrates the total number of ICU and hospital inpatients (across all ages classes) that were recorded by Northern Ireland’s Department of Health, and that were projected by the best fit of the model.

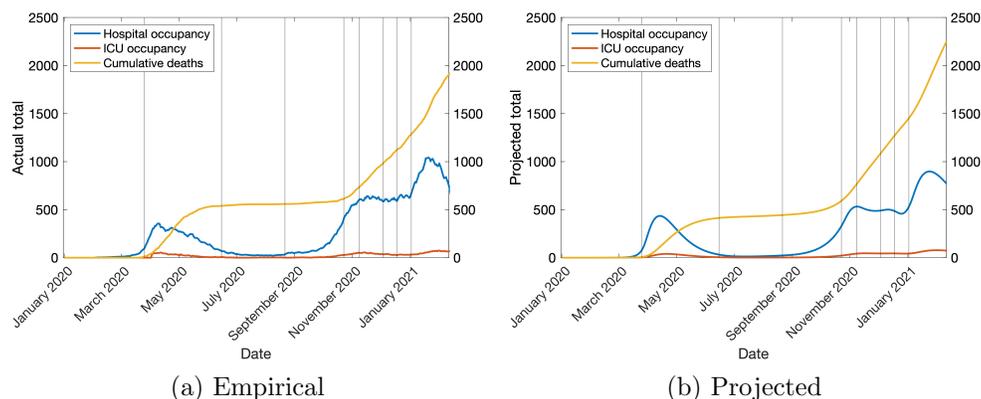


Figure 2: Hospital occupancy, ICU occupancy and cumulative deaths

Figure 3(a) shows the projected time-series of the sum of each compartment of the SEIR model across age classes, where the Infectious subclinical, Infectious clinical, and the hospital and ICU compartments are combined to a single “Infected” compartment. Figure 3(b) shows the time-series of the fraction of the total population who have been infected and left the susceptible compartment.

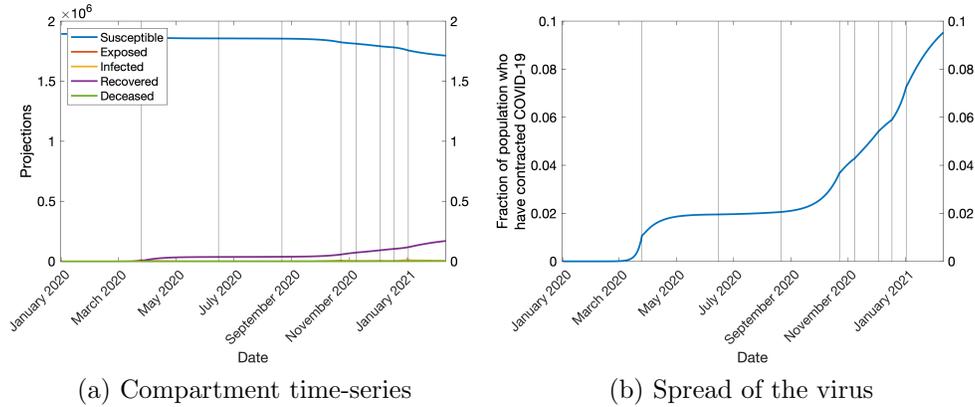


Figure 3: Time-series of the total of the compartments and the spread of the virus

Figure 4 illustrates the projected time-series of the effective reproduction ratio R_t of the virus during the best-fit simulation.

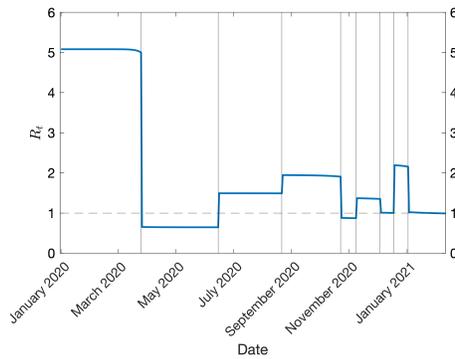


Figure 4: Time-series of the effective reproduction ratio R_t

Figure 5 shows the empirical data and the projected time-series of the number of hospital inpatients in Northern Ireland in each of the five age classes.

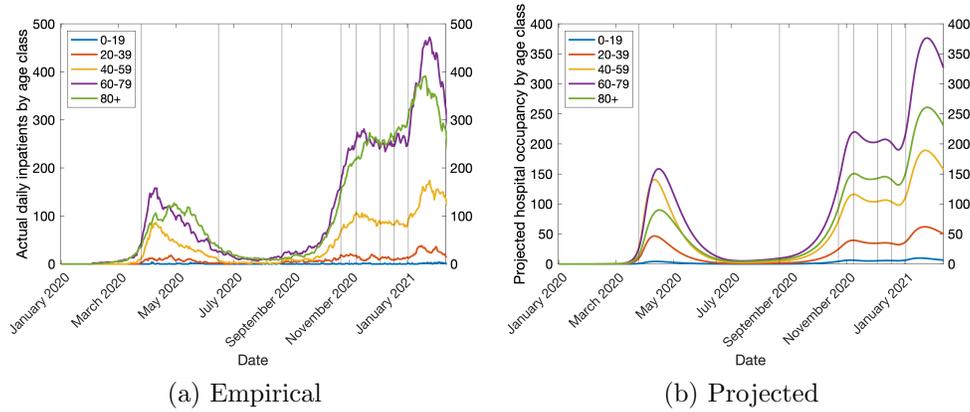


Figure 5: Hospital occupancy by age class

Figure 6 illustrates the empirical data and the projection for the cumulative number of fatal cases of infection in each of the five age classes.

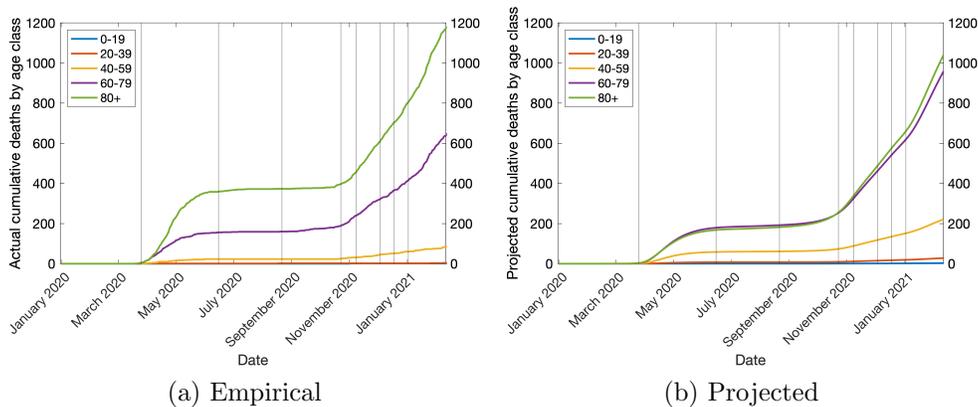


Figure 6: Cumulative deaths by age class

3 Alternate-history scenarios

Here we simulate hypothetical “alternate history” versions of the model, to determine what would have transpired in other circumstances or as a result of different policy decisions.

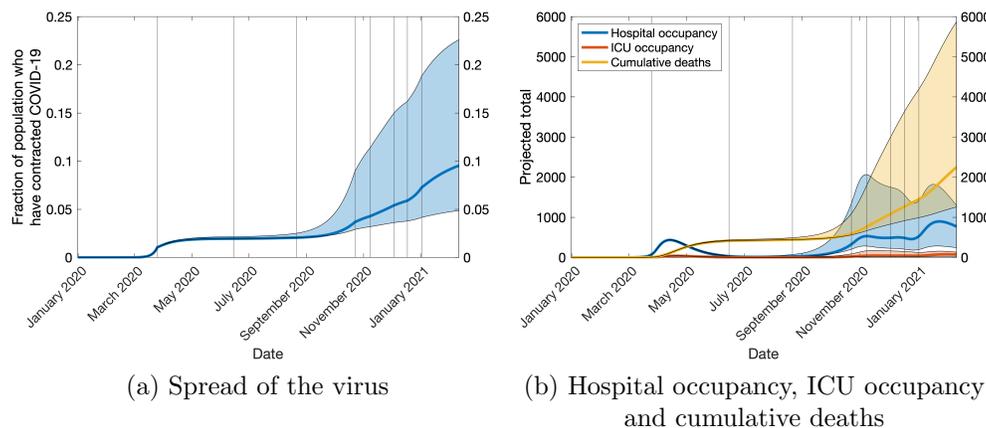


Figure 7: Effect of school closure

In Figure 7, we explore the role played by contacts made at school. For each outcome (fraction of cumulative cases, current hospital and ICU occupants, and cumulative deaths) the upper limit of the envelope shows the projection in the event that schools *never* close, even during usual school holidays, while the lower limit illustrates the projection in the event that schools had never re-opened following their closure during the initial lockdown. In each case, the central curve is the original best fit projection from Figure 3(b) and 2(b) rather than empirical data. No other factors are adjusted from the original model parameters, so we assume for example that circuit break lockdowns are not extended or introduced earlier despite the much greater hospital admissions already occurring by November 2020 in the upper limit, so although the projected number of deaths is nearly tripled by March 2021, this is more severe than a practical worst case scenario would likely be. Similarly, although this simulation indicates that a significant number of deaths could be avoided or prevented by 2021 if schools were closed indefinitely, this is only the case assuming that the actual restrictions introduced would still have been brought in nonetheless and that social behaviour is unchanged in this hypothetical scenario.

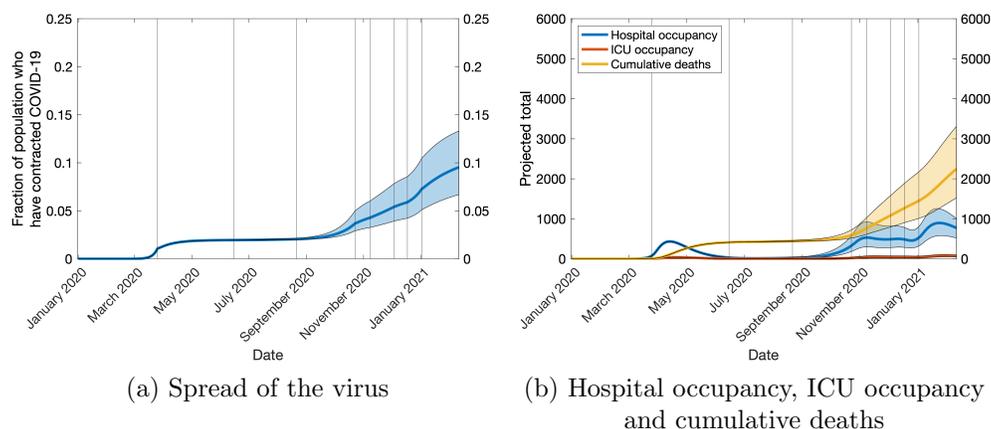


Figure 8: Effect of extending or reducing the first lockdown by one week

Next consider the sensitivity of the model to the duration of the original first lockdown. In reality, this ran with the severest restrictions from 24th March until 11th June 2020. In Figure 8, the envelopes illustrate the effect of either extending this lockdown by one week (lower limit) so that the lower transmission rate β_2 is used until 18th June, or reducing it by one week (upper limit) so that it ends on 4th June. As with the school closures, we assume that there are no other changes to the simulation and thus no policy reaction to the increased or decreased clinical outcomes, and observe that in such a hypothetical case the total number of deaths could potentially have been reduced by as much as 700 or increased by over 1000.

4 Best fit simulations with alternative parameter choices

In this section we show the best fit simulations obtained when some key parameter choices are altered from their values which are fixed throughout the investigations presented in the main body of the report, in particular the incubation rate σ , the progression rates γ_S and γ_C out of infectious compartments, and the relative transmissibility i_C of individuals in the clinical infectious compartment compared to those in the subclinical infectious compartment.

In all cases, the parameters fitted “between simulations” maintain their values determined previously ($D = 2$, $\mu_e = 3.0$, $\mu_w = 0.0$, and $\mu_o = 0.95$), and so in these simulations we allow the model procedure to learn the best initial day and age class of infection, and the value of β_1, \dots, β_9 that yields the best fit of the rolling average of hospital admissions for each age class.

4.1 Modifying the incubation rate σ

In this section we consider the effect of varying the amount of time that individuals spend in the exposed compartment before progressing to the subclinical infectious compartment.

Figure 9 shows the best fit hospital admissions and the consequent projections of hospital occupancy and total deaths in the event that exposed cases incubate and transition to subclinical infectious cases more quickly, spending an average of 4 days in the exposed compartment (rather than 5.1).

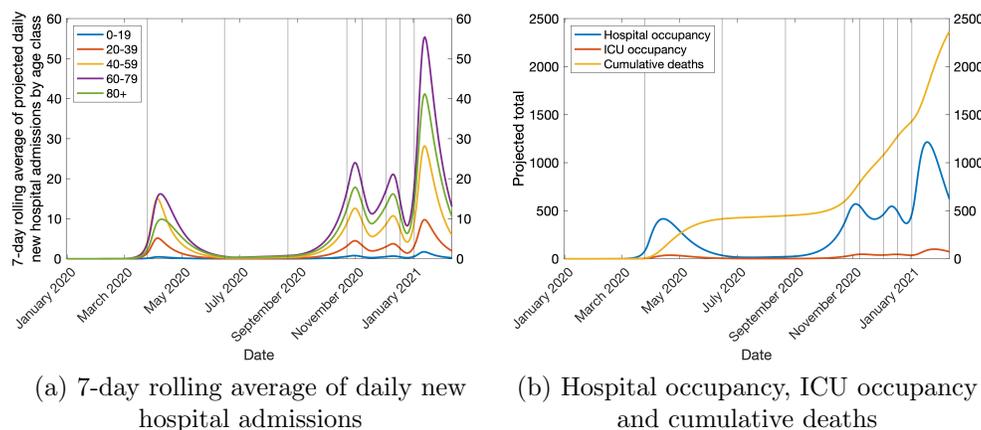


Figure 9: Faster incubation: $\sigma = 1/4$

Figure 10 shows the best fit outcomes if individuals spend slightly longer (an average of 6 days) in the exposed compartment before becoming infectious.

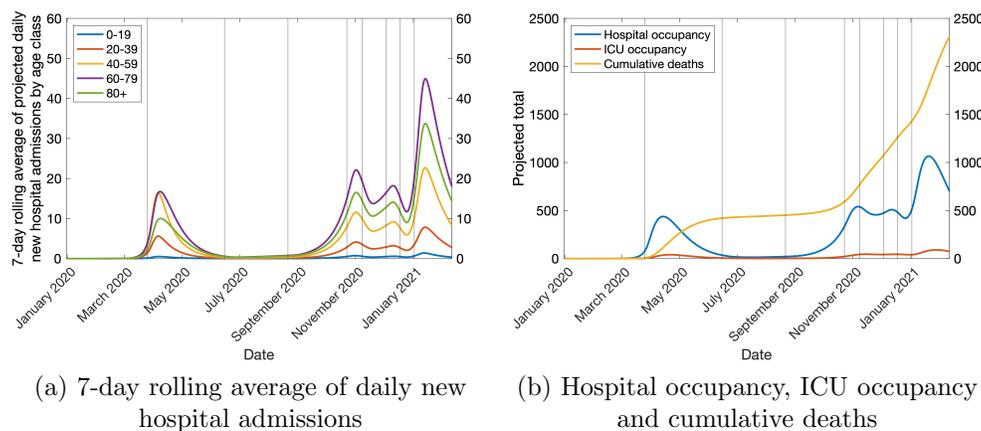


Figure 10: Slower incubation: $\sigma = 1/6$

4.2 Modifying the progression and recovery rates γ

In this section we investigate the role of the time spent in each infectious compartment in the model, and the influence of the progression rates γ_S and γ_C which control how long individuals spend within the subclinical and clinical infectious compartments respectively.

In Figure 11, we first test the effect of a variant of the SEIIR model, where rather than exposed individuals progressing to subclinical infectious cases and then some of those progressing to the clinical infectious compartment while the rest recover, individuals leaving the exposed compartment instead *either* enter the clinical infectious compartment (with a probability of ϵ_i for age class $i = 1, \dots, 5$) and spend an average of 7 days there, *or* they join the subclinical infectious compartment for 3 days - all of whom then recover.

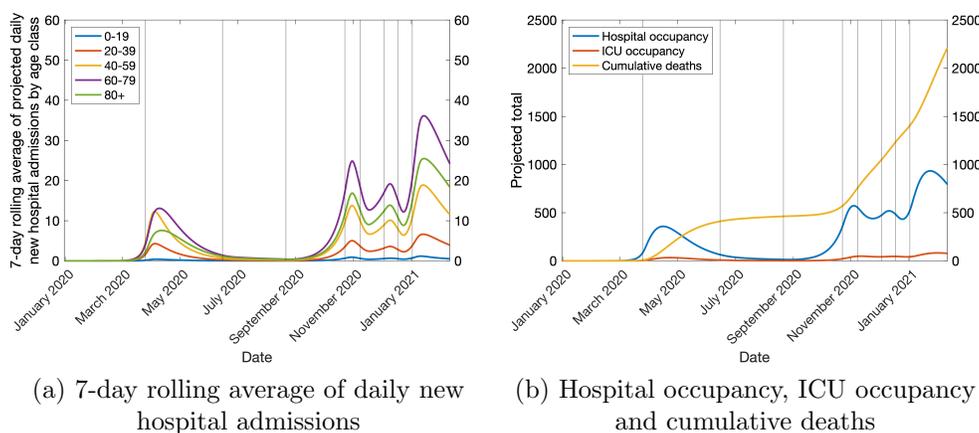


Figure 11: Alternative infection model: Exposed cases progress to either subclinical or clinical infectious

In Figures 12 and 13, we revert to the original model design of subclinical infectious cases always preceding clinical infectious cases, and consider the best-fit model projections in the event that the amount of time spent in each infectious sub-compartment is either reduced (Figure 12) or increased (Figure 13). Note that, as in all alternative scenarios in this section, the simulation remains free to modify the global transmission parameter β to compensate for greater or fewer infections and thus hospitalisations.

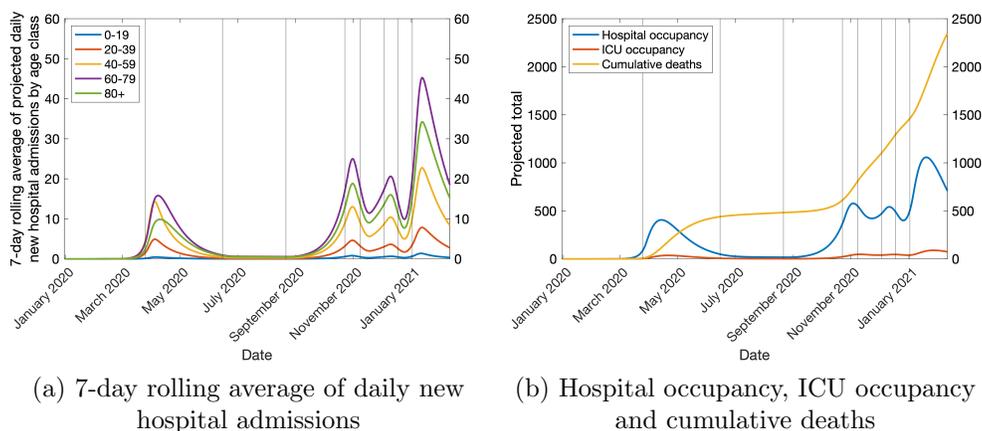


Figure 12: Faster progression: $\gamma_S = 1/2, \gamma_C = 1/3$

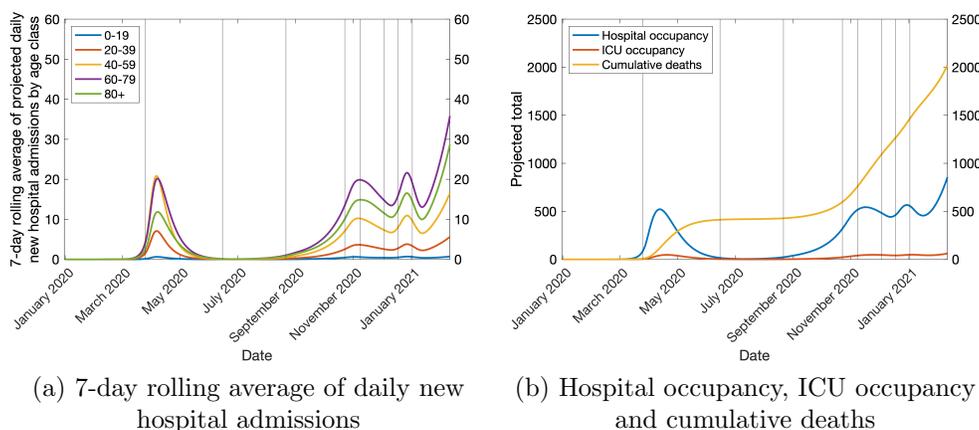


Figure 13: Slower progression: $\gamma_S = 1/4, \gamma_C = 1/5$

4.3 Modifying the relative transmissibility from clinical and sub-clinical infectious cases

In this section, we consider the effect of the relative levels of infectivity of individuals in the subclinical and clinical infectious compartments. In all simulations in this section, we assume that exposed individuals all enter the subclinical infectious class and remain there for an average of 3 days, having an age-dependent probability ϵ_i of then moving to the clinical infectious compartment for an average period of 4 days. In all cases the transmission rate of the subclinical compartment is $i_S = 1$, transmission from both infectious compartments are governed by the time-dependent fitted parameter β , and we consider the effect of making alternative choices of the relative transmission rate i_C of the clinical compartment. This encompasses the net effect of both social behaviour, such as whether or not individuals with symptoms self-isolate or not, and the biological question of how much more infectious an individual with manifesting symptoms (e.g. cough) may be.

Figure 14 shows the best fit with $i_C = 0$, and so in fact all transmission from a case occurs during the 3-day subclinical infectious period.

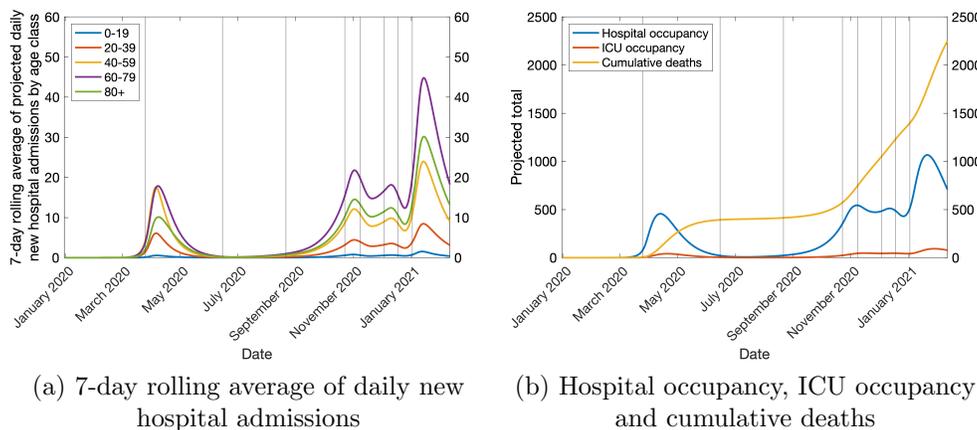


Figure 14: $i_C = 0$

Figures 15-17 show the best fit with values of $i_C < 1$, so that (as in the main report) due to self-isolation of symptomatic individuals, more transmission occurs per day when in the subclinical infectious compartment.

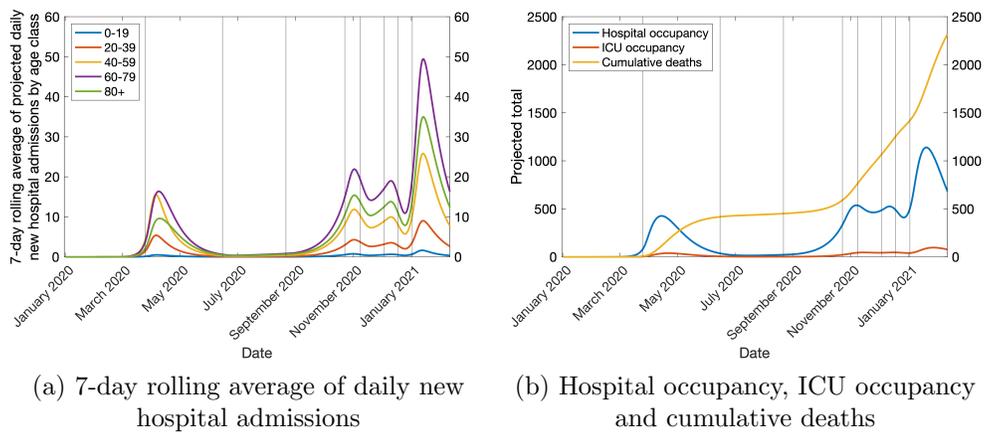


Figure 15: $i_C = 0.25$

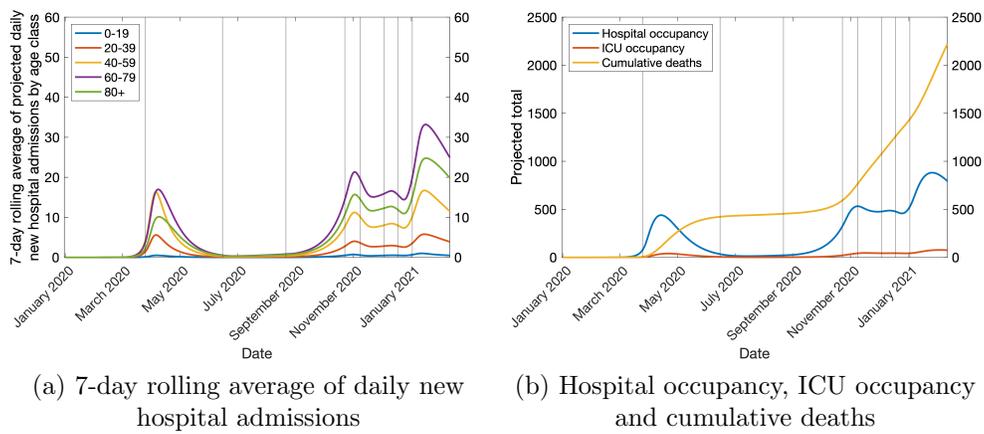


Figure 16: $i_C = 0.5$

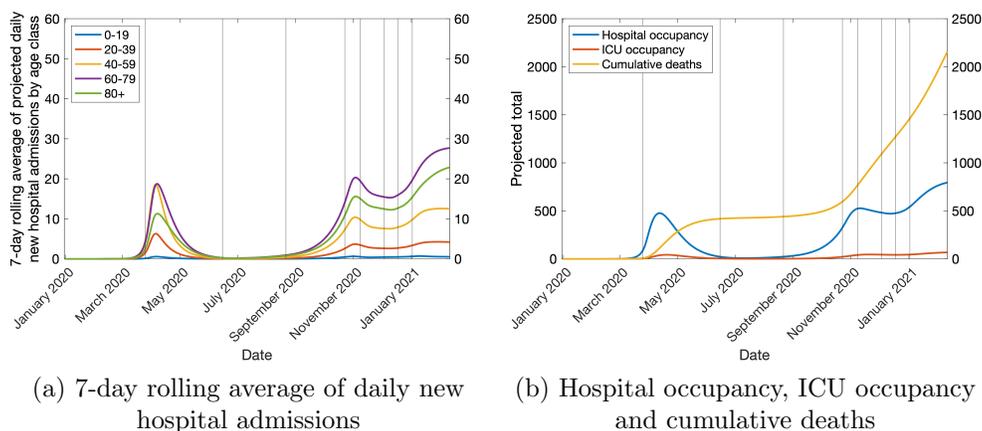
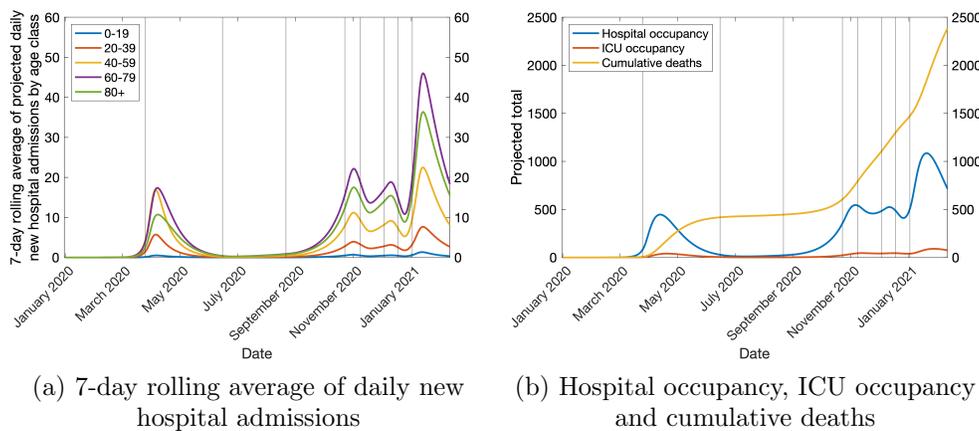
Figure 17: $i_C = 0.75$

Figure 18 illustrates the best fit with $i_C = i_S = 1$, so that an individual tends to have a net transmission effect when manifesting symptoms that is equivalent to when they are in the subclinical infectious compartment.

Figure 18: $i_C = 1$

Figures 19-20 show the best fit when $i_C > 1$ and thus individuals spread more infections per day after they have begun to show symptoms than they did beforehand.

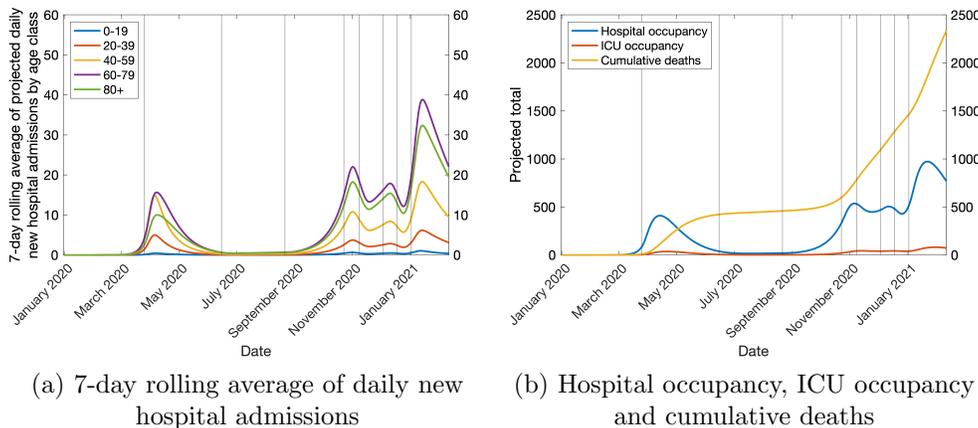


Figure 19: $i_C = 1.5$

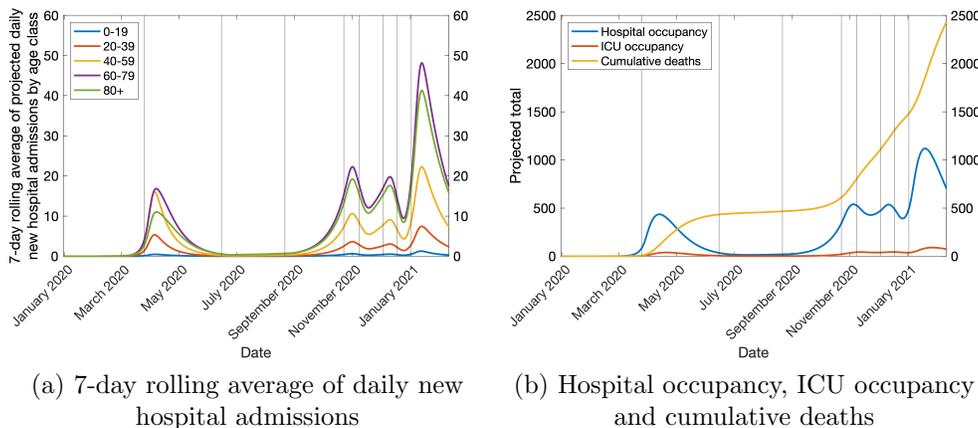


Figure 20: $i_C = 2$

5 Role of intensity and duration of lockdowns early in the pandemic

This section replicates the results shown in §4 of the main body of the report, and describes the best results of additional outcomes that can be achieved by either one or ten lockdowns in the first 100 days beginning March 23rd 2020. In particular, for the subfigure associated with each of the outcomes, the main figure shows the best result that could be achieved (minimum number of deaths, minimum peak number of hospital inpatients, etc.) over the 1000 day simulation for the particular duration of lockdown (from 1-100 days) and strength of lockdown (as a fractional reduction to the value of the transmission rate β from 1-100%) indicated by the co-ordinate, and it is the best result from implementing the lockdown on one of the 100 days following March 23rd. The (upper) inset image illustrates which of these days it is that achieves the best outcome which is shown in the main image, and in the scenarios of multiple shorter lockdowns, the lower inset image shows the best number of days (from 0-50) to separate subsequent lockdowns to yield the best result shown.

Two model scenarios for Northern Ireland are presented: whereby in Figures 21-24 the country is treated as a closed, isolated system and there is no change in the population. In particular, no new exposed or infectious individuals enter the system from an external pool, and so it is possible to entirely eliminate the disease while many in the population have not yet been exposed. In the second scenario (Figures 25-28, this is not the case, and there is instead a constant inflow of exactly one new exposed individual of a random age class added to the population each day after day 30 (when the first infection begins). For each of these scenarios, we further consider the role of the strength and duration of lockdowns that apply to the entire population (Figure 21 and 25), to only those aged over 80 (Figure 22 and 26), to only those aged 60-80 (Figure 23 and 27), and finally to only those aged over 60 (Figure 24 and 28).

For each scenario, the best of the following outcomes (by day 1000 of the simulation) that can be achieved are shown as a function of the strength, duration, and day of lockdown:

- (a) Total number of deaths.
- (b) Peak number of hospital occupants.
- (c) Peak number of ICU occupants.
- (d) Number of days during which the number of hospital inpatients exceeded 100.
- (e) Number of days during which the number of hospital inpatients exceeded 1000.
- (f) Number of days during which the number of ICU occupants exceeded 50.
- (g) Number of days during which the number of ICU occupants exceeded 200.

With the exception that only (a) is shown in Figures 29, in the case of multiple shorter lockdowns.

5.1 Closed system

In the first set of scenarios, the results of which are shown in Figures 21-24, we treat Northern Ireland as a closed system, so some hypothetical lockdowns can lead to an early end to the pandemic.

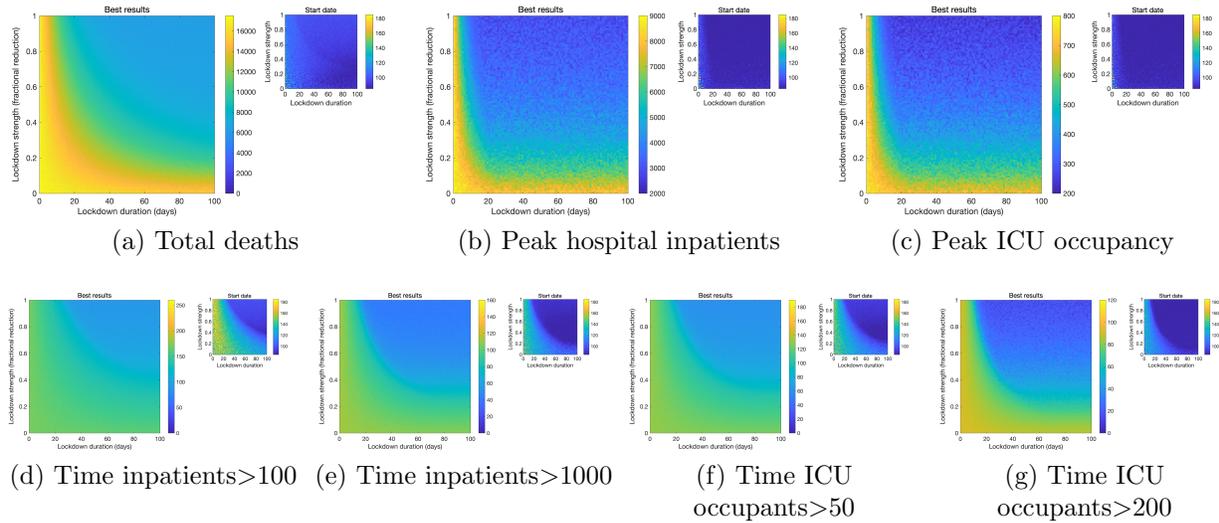


Figure 21: Optimally-timed single lockdown in a closed system. All age classes restricted.

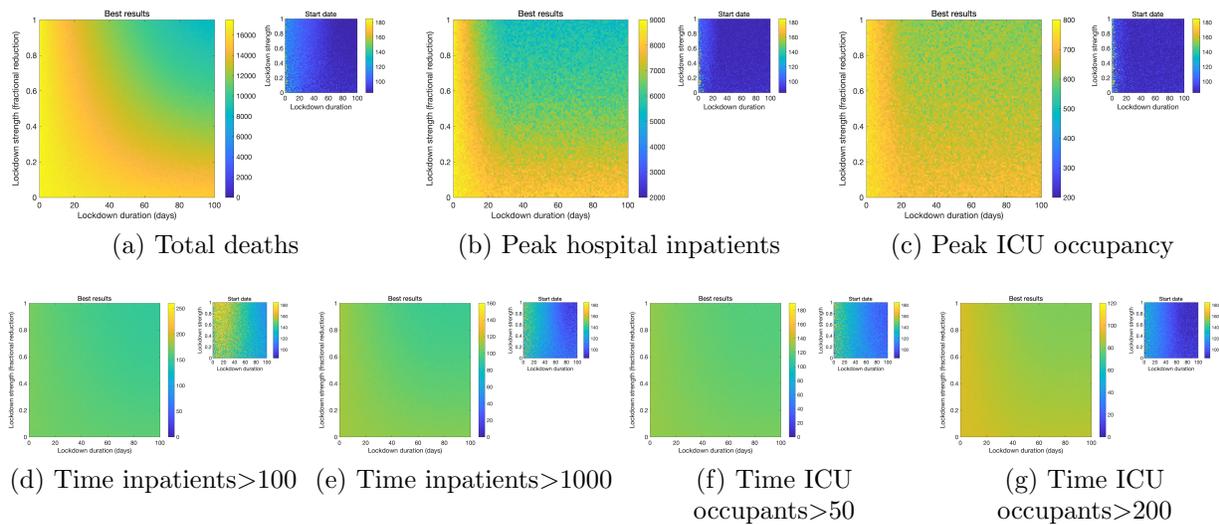


Figure 22: Optimally-timed single lockdown in a closed system. Only the 80+ age class is restricted.

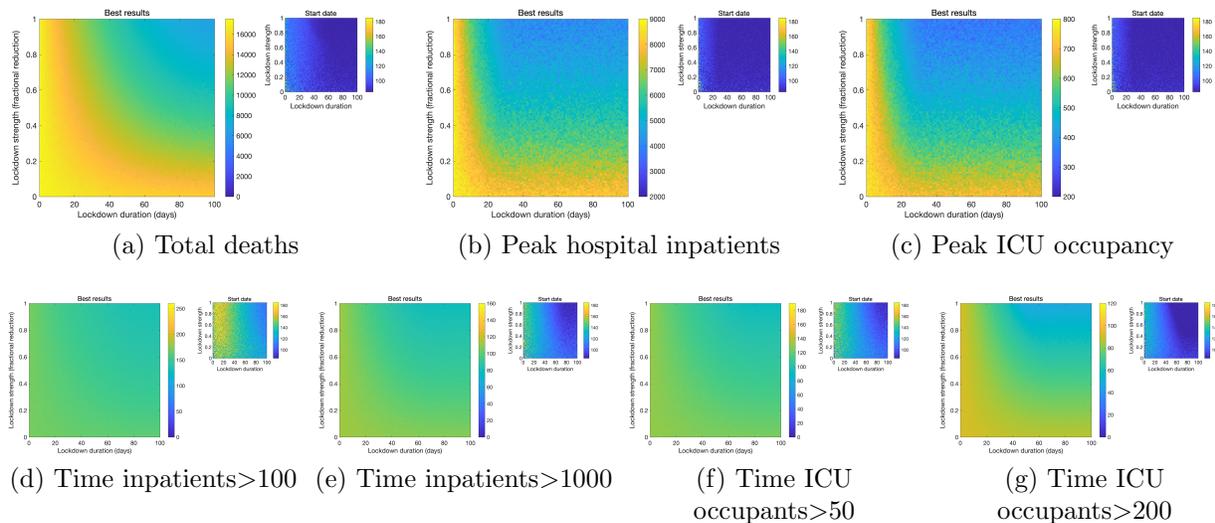


Figure 23: Optimally-timed single lockdown in a closed system. Only the 60-80 age class is restricted.

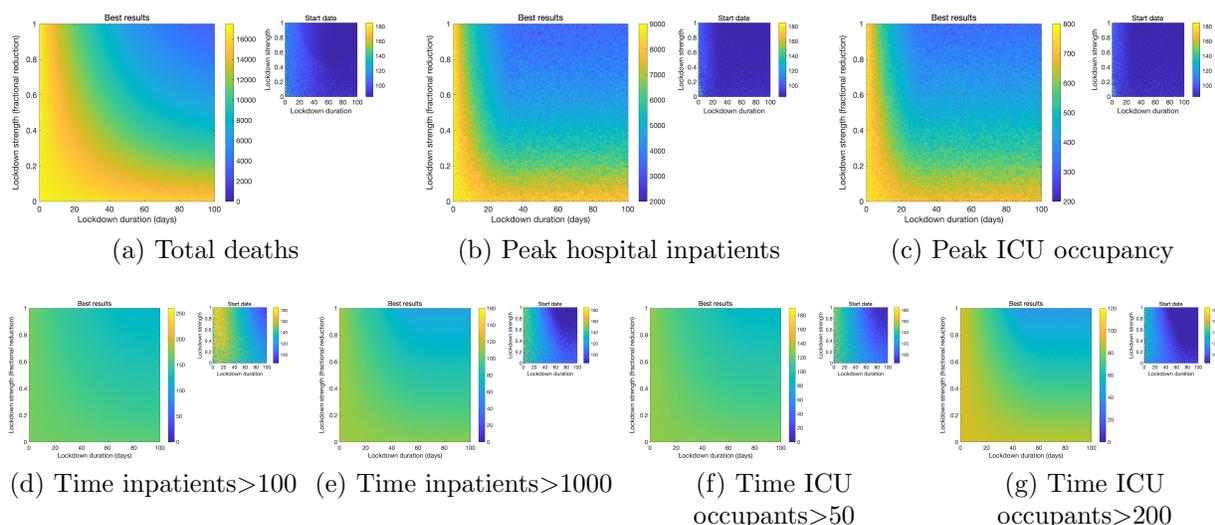


Figure 24: Optimally-timed single lockdown in a closed system. Only the 60-80 and 80+ age classes are restricted.

5.2 Constant inflow

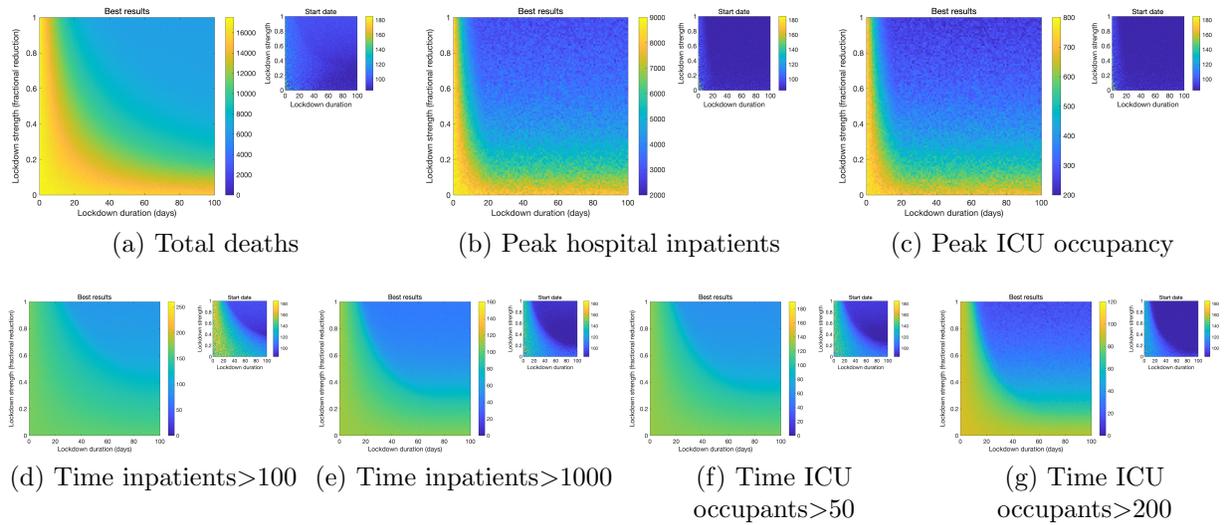


Figure 25: Optimally-timed single lockdown in a system with constant inflow. All age classes restricted.

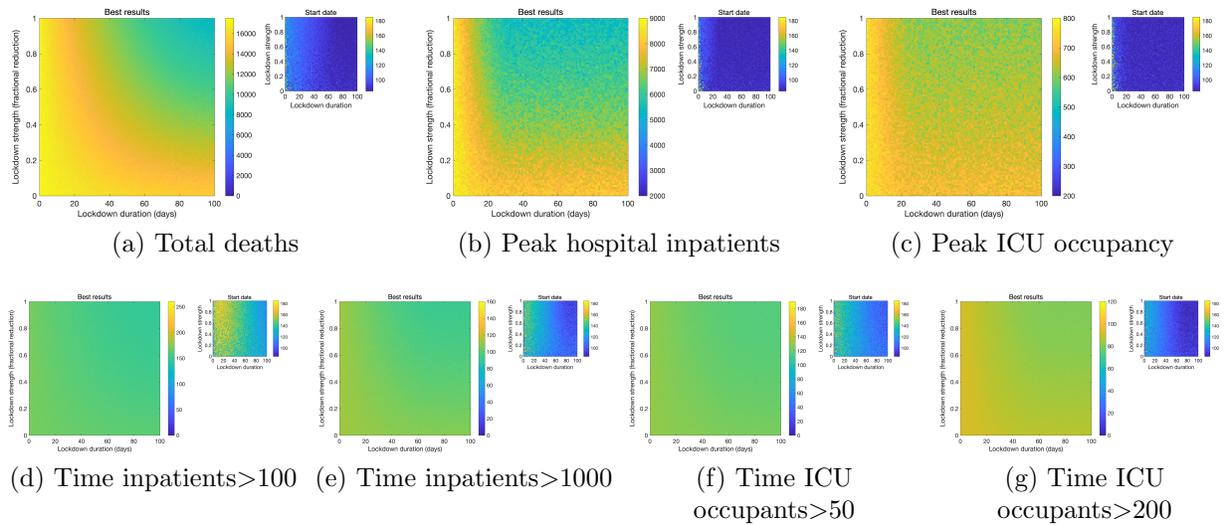


Figure 26: Optimally-timed single lockdown in a system with constant inflow. Only the 80+ age class is restricted.

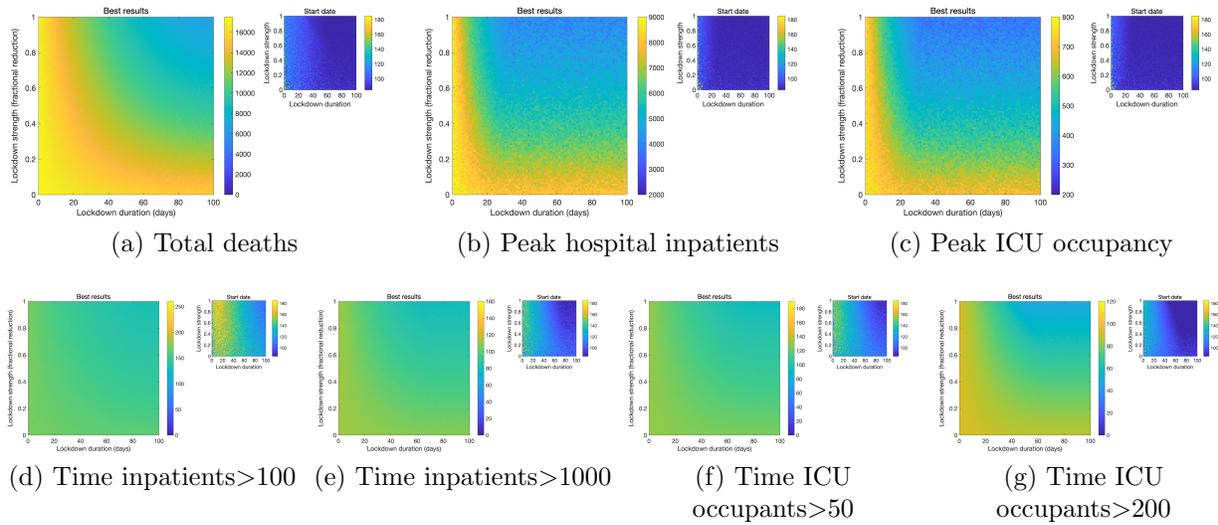


Figure 27: Optimally-timed single lockdown in a system with constant inflow. Only the 60-80 age class is restricted.

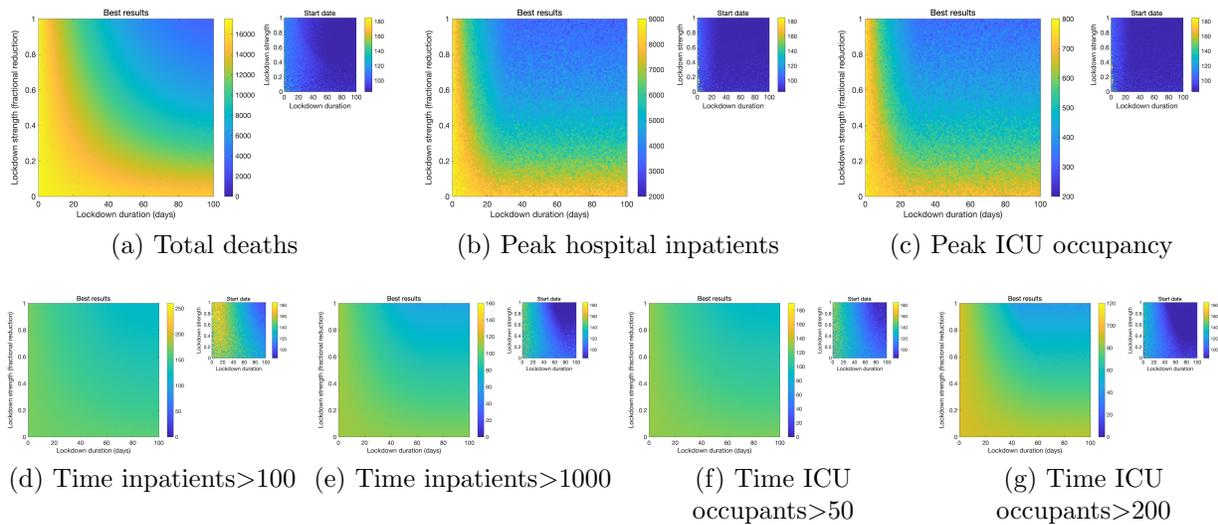


Figure 28: Optimally-timed single lockdown in a system with constant inflow. Only the 60-80 and 80+ age classes are restricted.

5.3 Ten evenly-spaced lockdowns

In place of a single lockdown of up to 100 days, in Figure 29 we illustrate the results of employing ten lockdowns of up to ten days each. In all cases, the lockdowns are of equal duration and are equally-spaced in time once the first lockdown has begun. The subfigures show both the start date (top) and the spacing (bottom) in the best case for strength and duration.

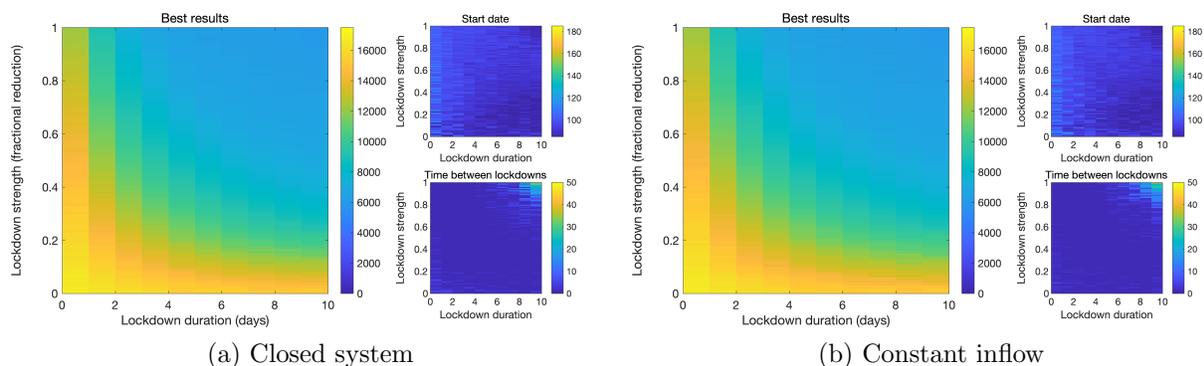


Figure 29: Total deaths using ten evenly-spaced short lockdowns with optimal initial timing. All age classes are restricted.

In the experiments with ten shorter lockdowns (Figure 29), these were not able to provide such effective suppression compared to a single continuous lockdown. However, they were more effective when weak and short compared to a single lockdown equivalent as they gradually reduced the effects of the pandemic over a longer period. In the best cases with the strongest short lockdowns, it was optimal to spread them out considerably (observe the brighter colour of the top-right of the lower subfigure in each case).

6 Mechanistic activation of lockdowns

In this section we replicate and add to the results demonstrated in §5 of the main report. These experiments test the use of mechanisms to trigger the implementation of lockdowns during a pandemic according to conditions that must be met by an observable statistic such as the current number of hospital inpatients or the total number of new daily deaths.

For each of the two model scenarios (treating Northern Ireland as a closed system with respect to the virus in Figures 30-33, or not in Figures 34-37), and each of these two possible trigger mechanisms, we consider the results on a selection of outcomes of either allowing a single or multiple lockdowns during the first 1000 days of the simulation (beginning January 1st 2020), followed by a further 1000 days during which no lockdown(s) may be implemented. Results are shown as a function of the threshold of the trigger (that is, the number of hospital inpatients or the number of daily deaths) that activates the intervention and the delay in terms of the number of days between this condition being met and the lockdown actually beginning. This is shown for three levels of intervention: low-strength (15 days, 25% reduction in β), medium strength (30 days, 50% reduction in β) and high-strength (60 days, 75% reduction).

In all cases, the outcomes shown are:

- (a) Total number of deaths.
- (b) Peak number of hospital occupants.
- (c) Peak number of ICU occupants.
- (d) Number of days during which the number of hospital inpatients exceeded 100.
- (e) Number of days during which the number of hospital inpatients exceeded 500.
- (f) Number of days during which the number of hospital inpatients exceeded 1000.
- (g) Number of days during which the number of ICU occupants exceeded 50.
- (h) Number of days during which the number of ICU occupants exceeded 100.
- (i) Number of days during which the number of ICU occupants exceeded 200.
- (j) The fraction of the population who have contracted the disease by the simulation end.
- (k) The total number of infectious individuals remaining at the end of the simulation.
- (l) The number of lockdowns which were triggered during the course of the simulation.

6.1 Closed system

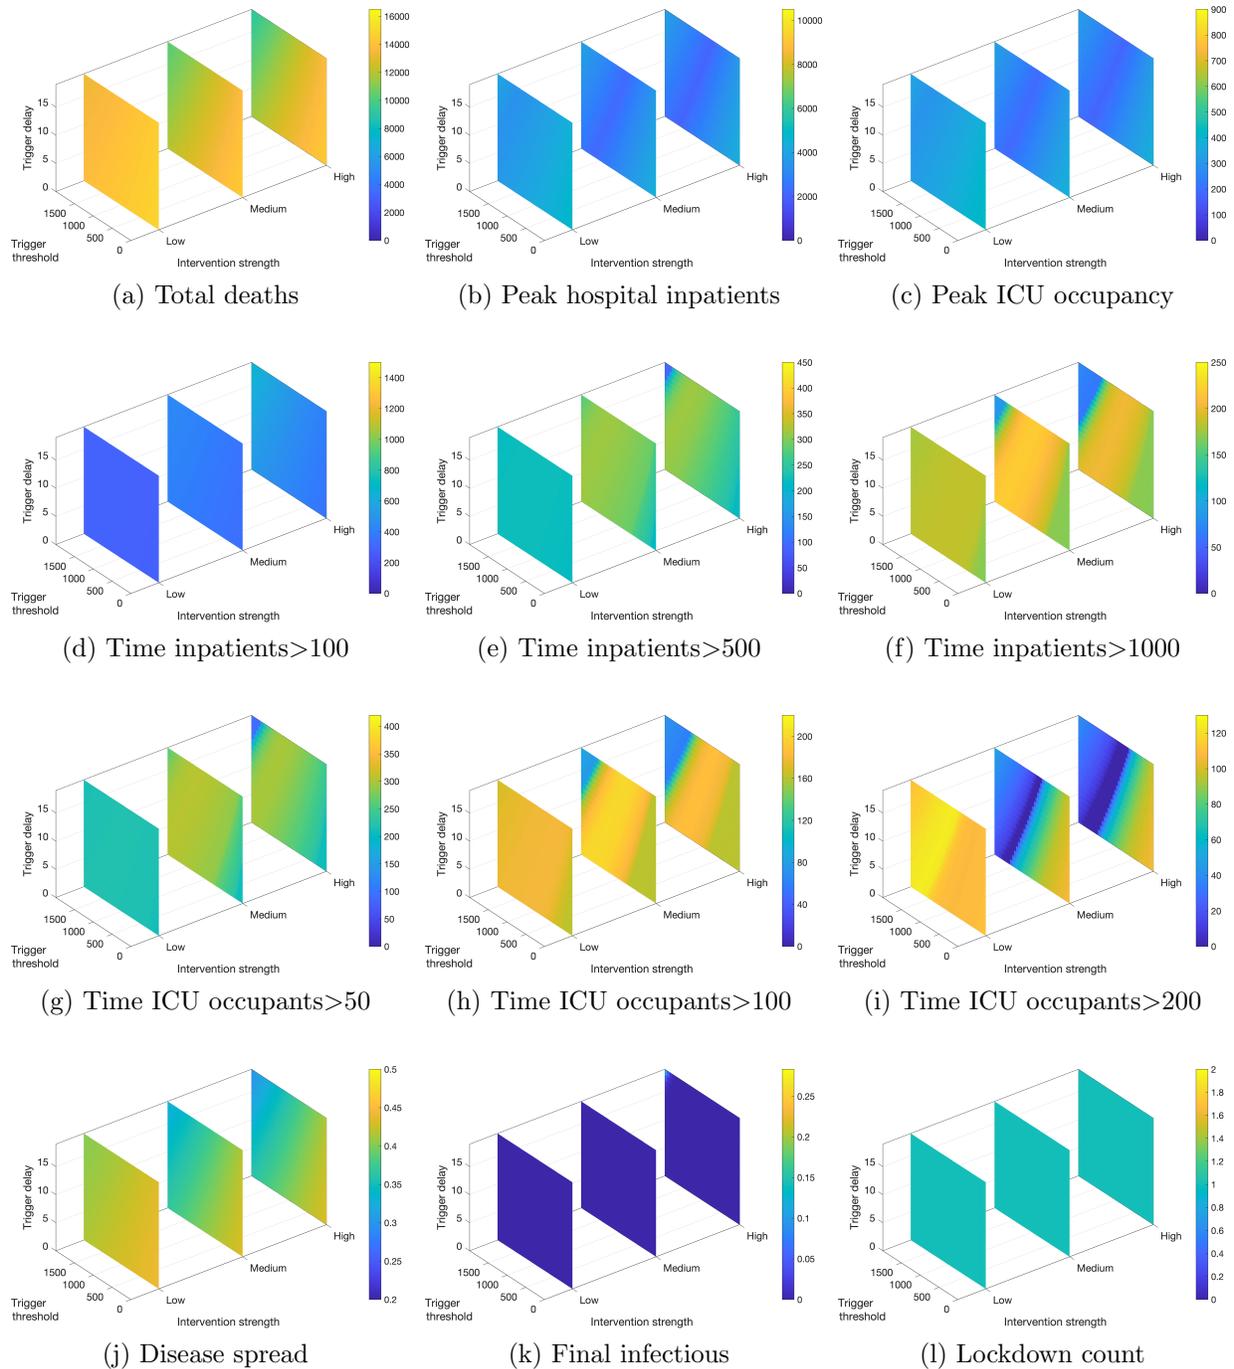


Figure 30: Dynamic lockdowns: Closed system - Hospital triggers - 1 lockdown

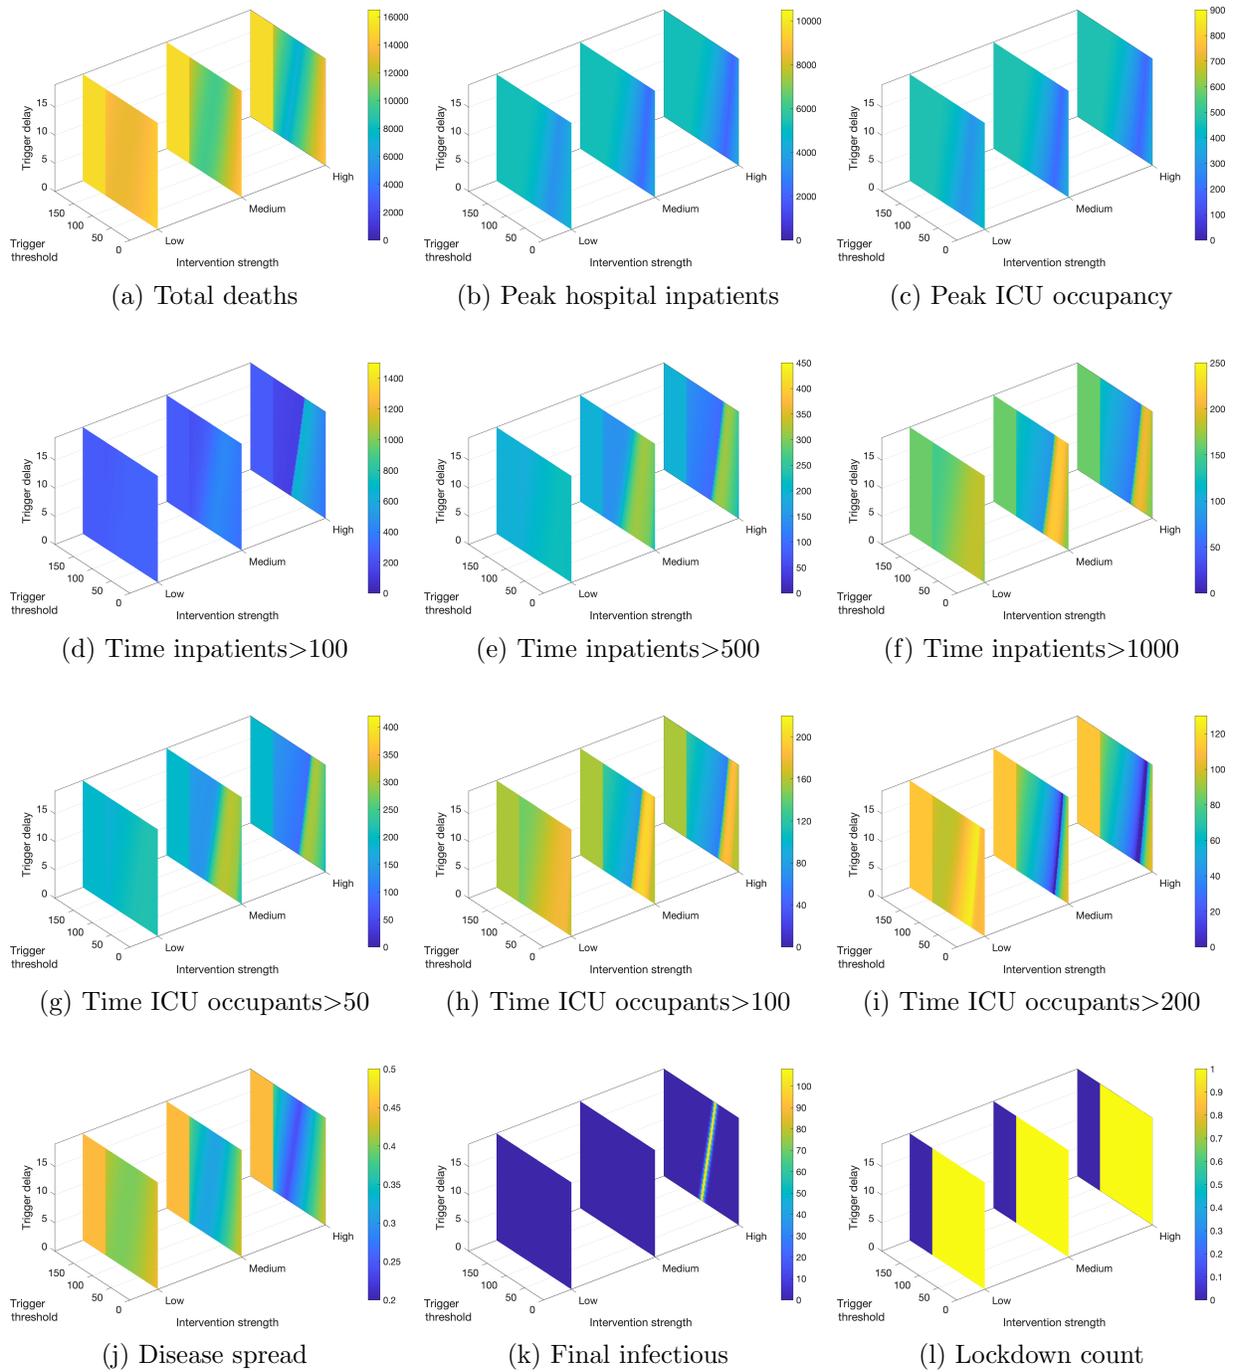


Figure 31: Dynamic lockdowns: Closed system - Death triggers - 1 lockdown

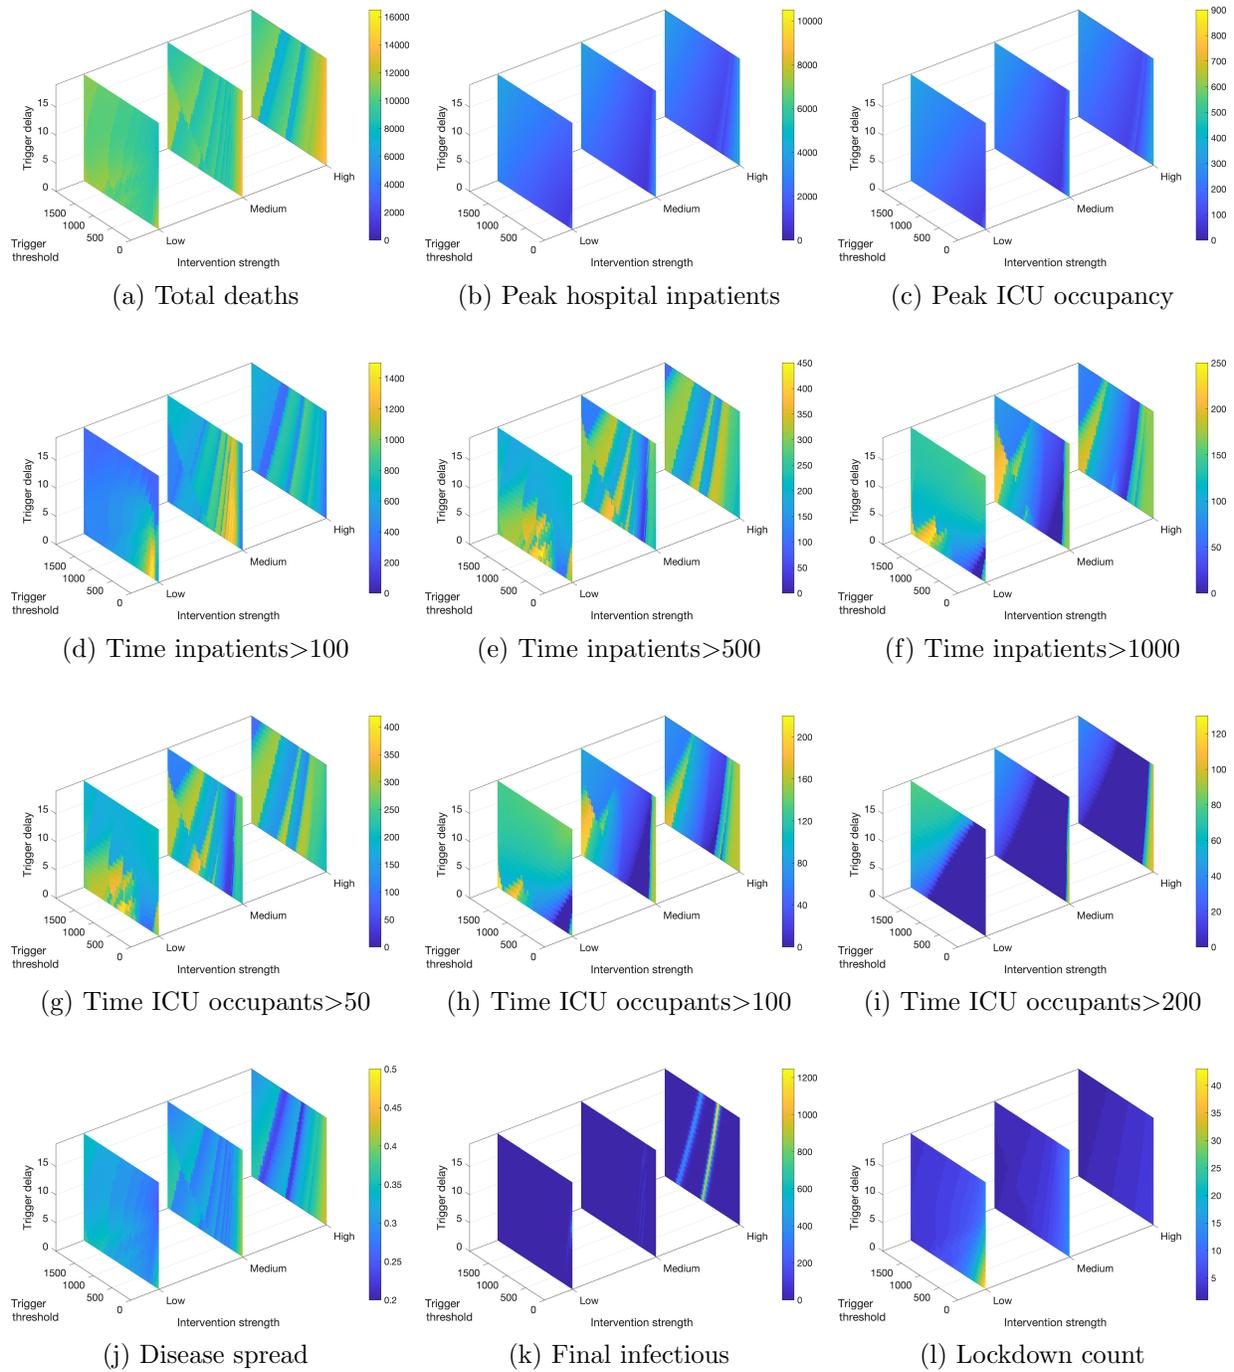


Figure 32: Dynamic lockdowns: Closed system - Hospital triggers - Multiple lockdowns

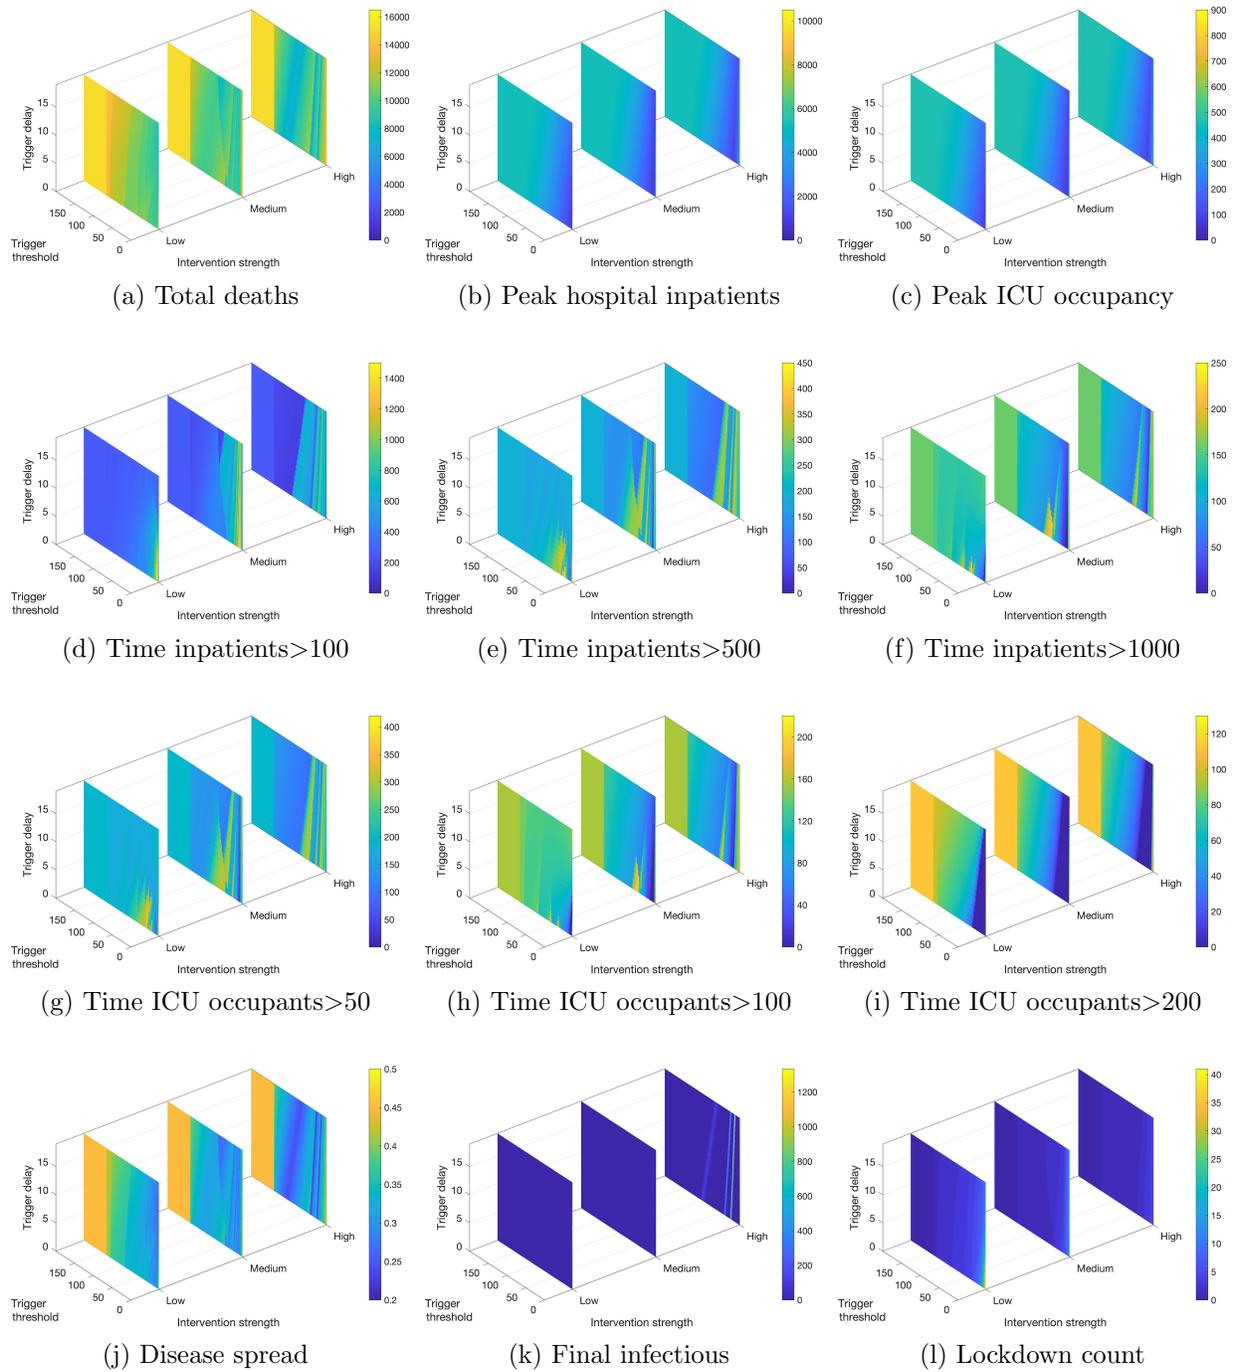


Figure 33: Dynamic lockdowns: Closed system - Death triggers - Multiple lockdowns

6.2 Constant inflow

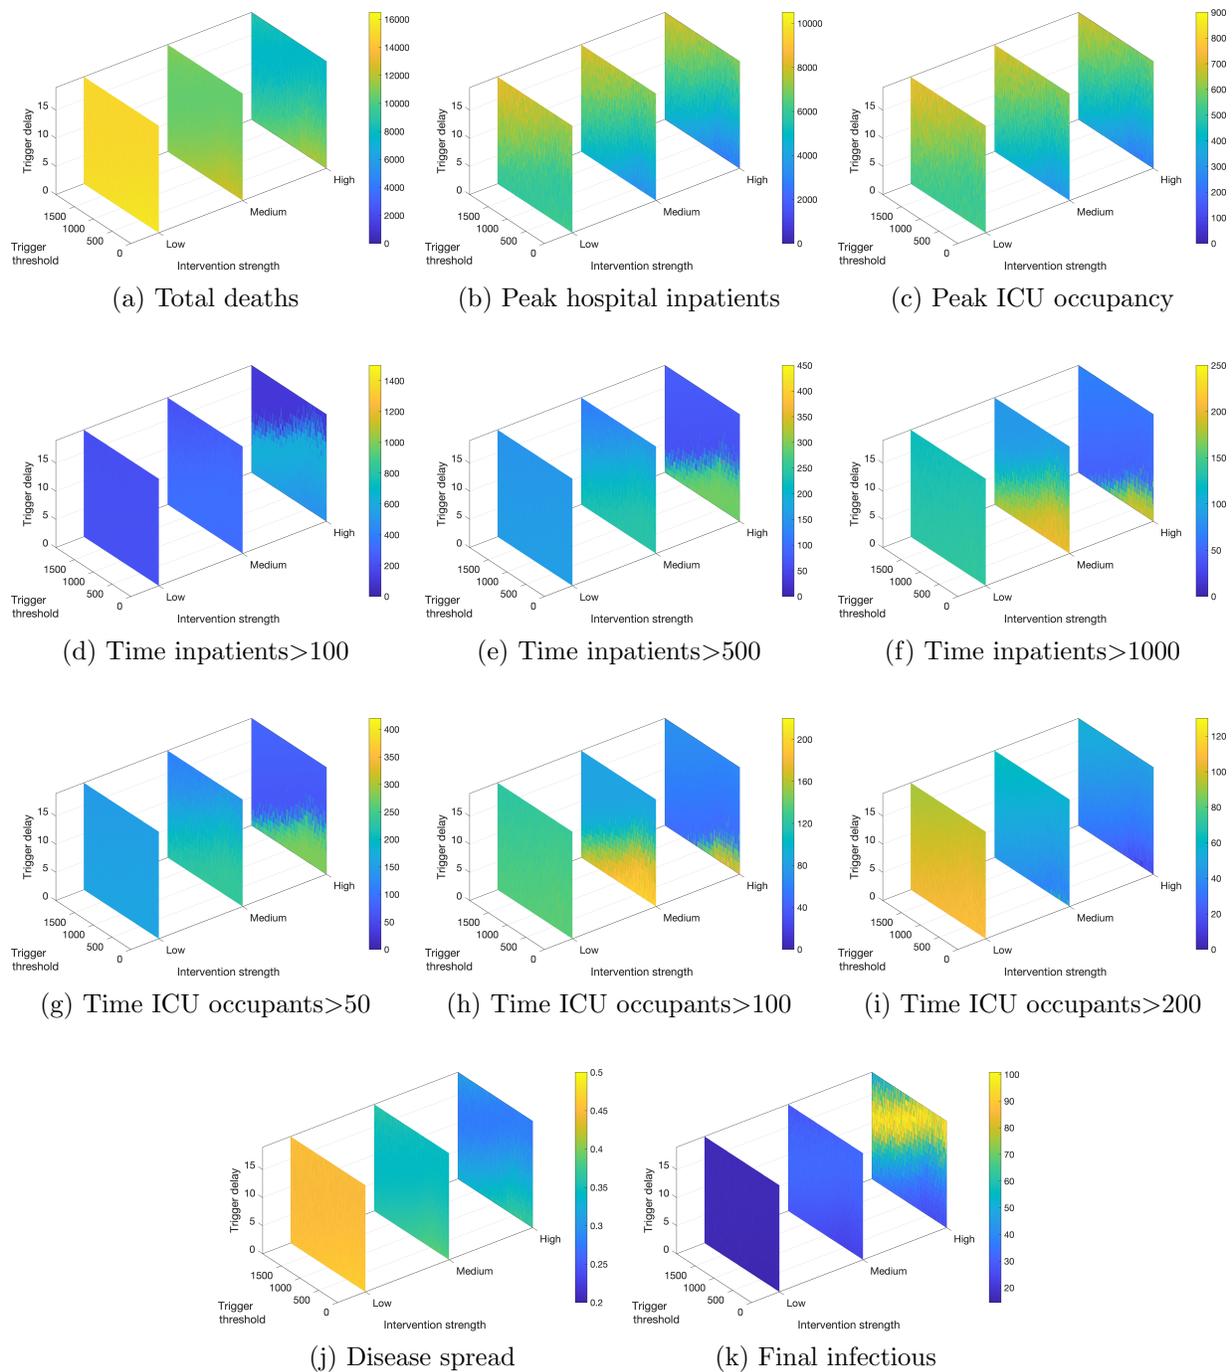


Figure 34: Dynamic lockdowns: Constant inflow - Hospital triggers - 1 lockdown

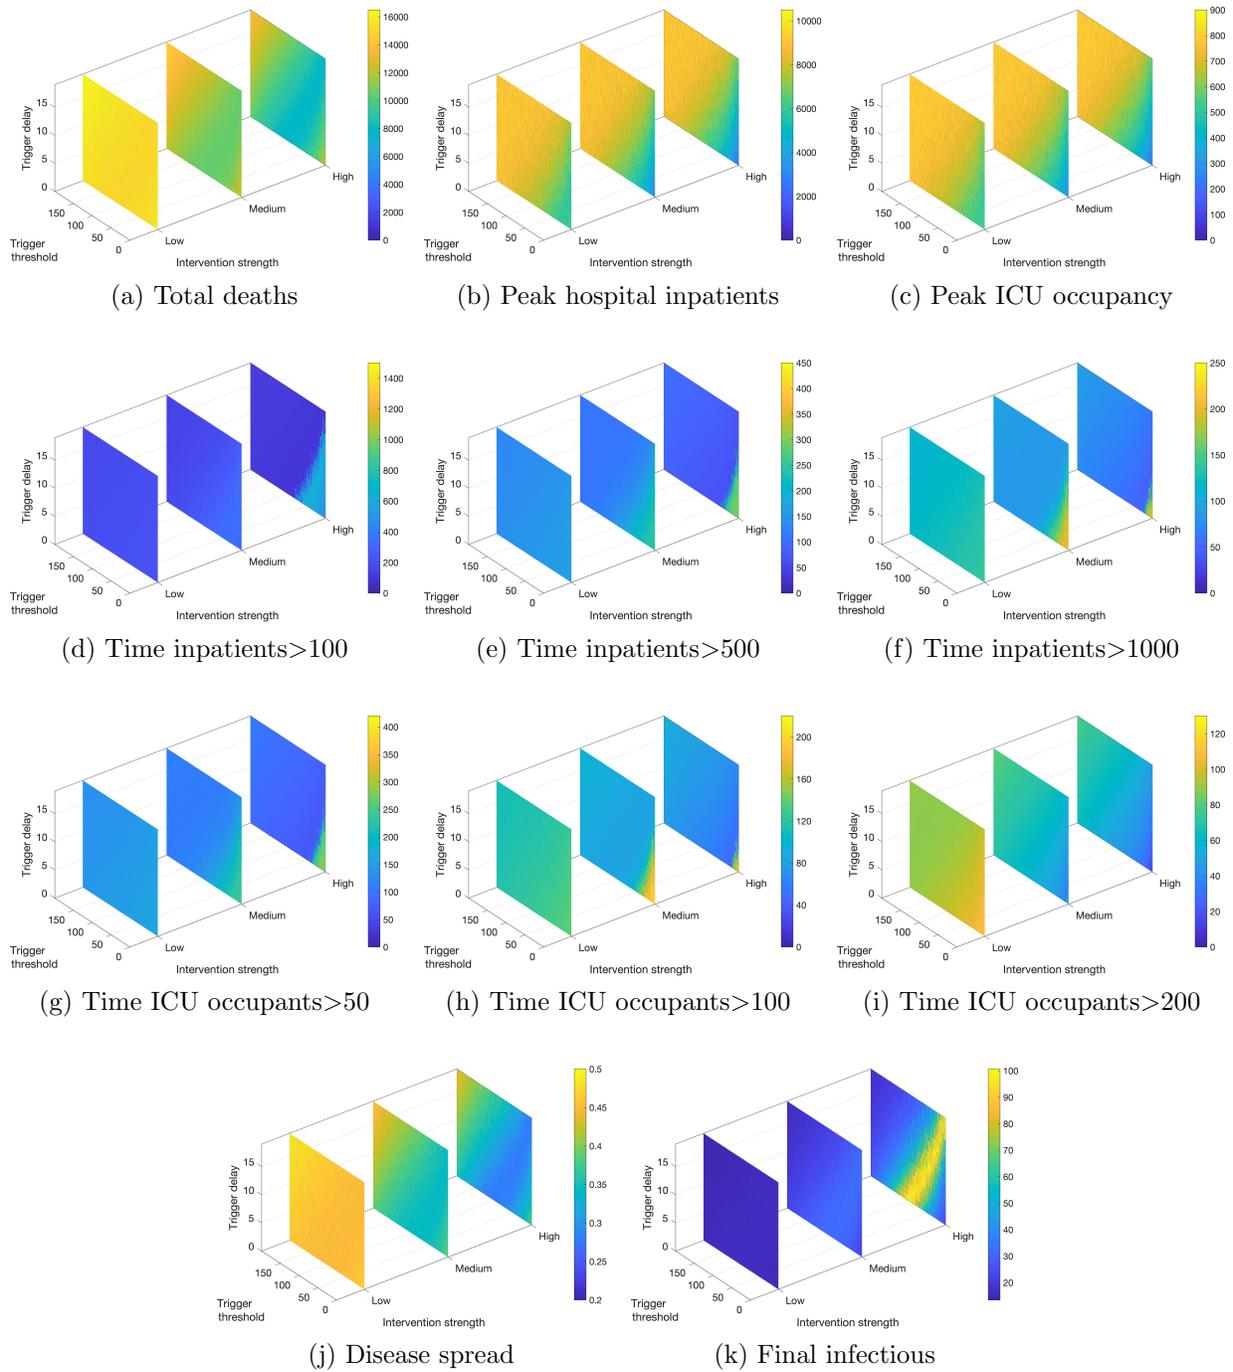


Figure 35: Dynamic lockdowns: Constant inflow - Death triggers - 1 lockdown

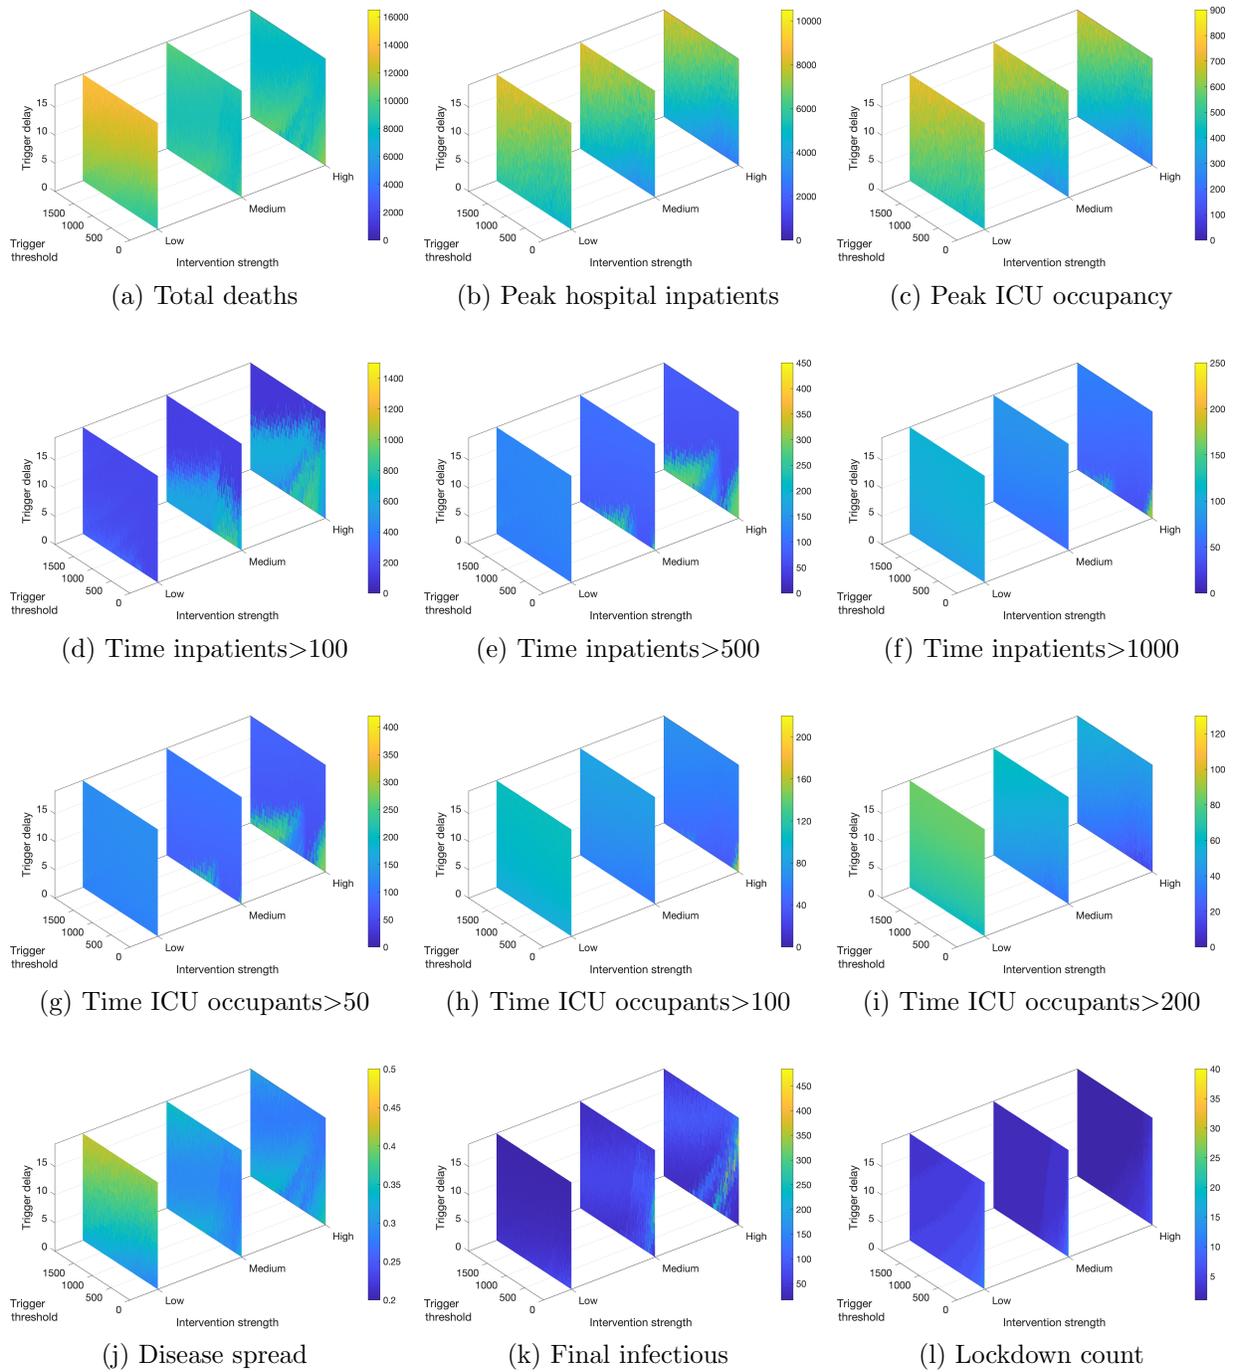


Figure 36: Dynamic lockdowns: Constant inflow - Hospital triggers - Multiple lockdowns

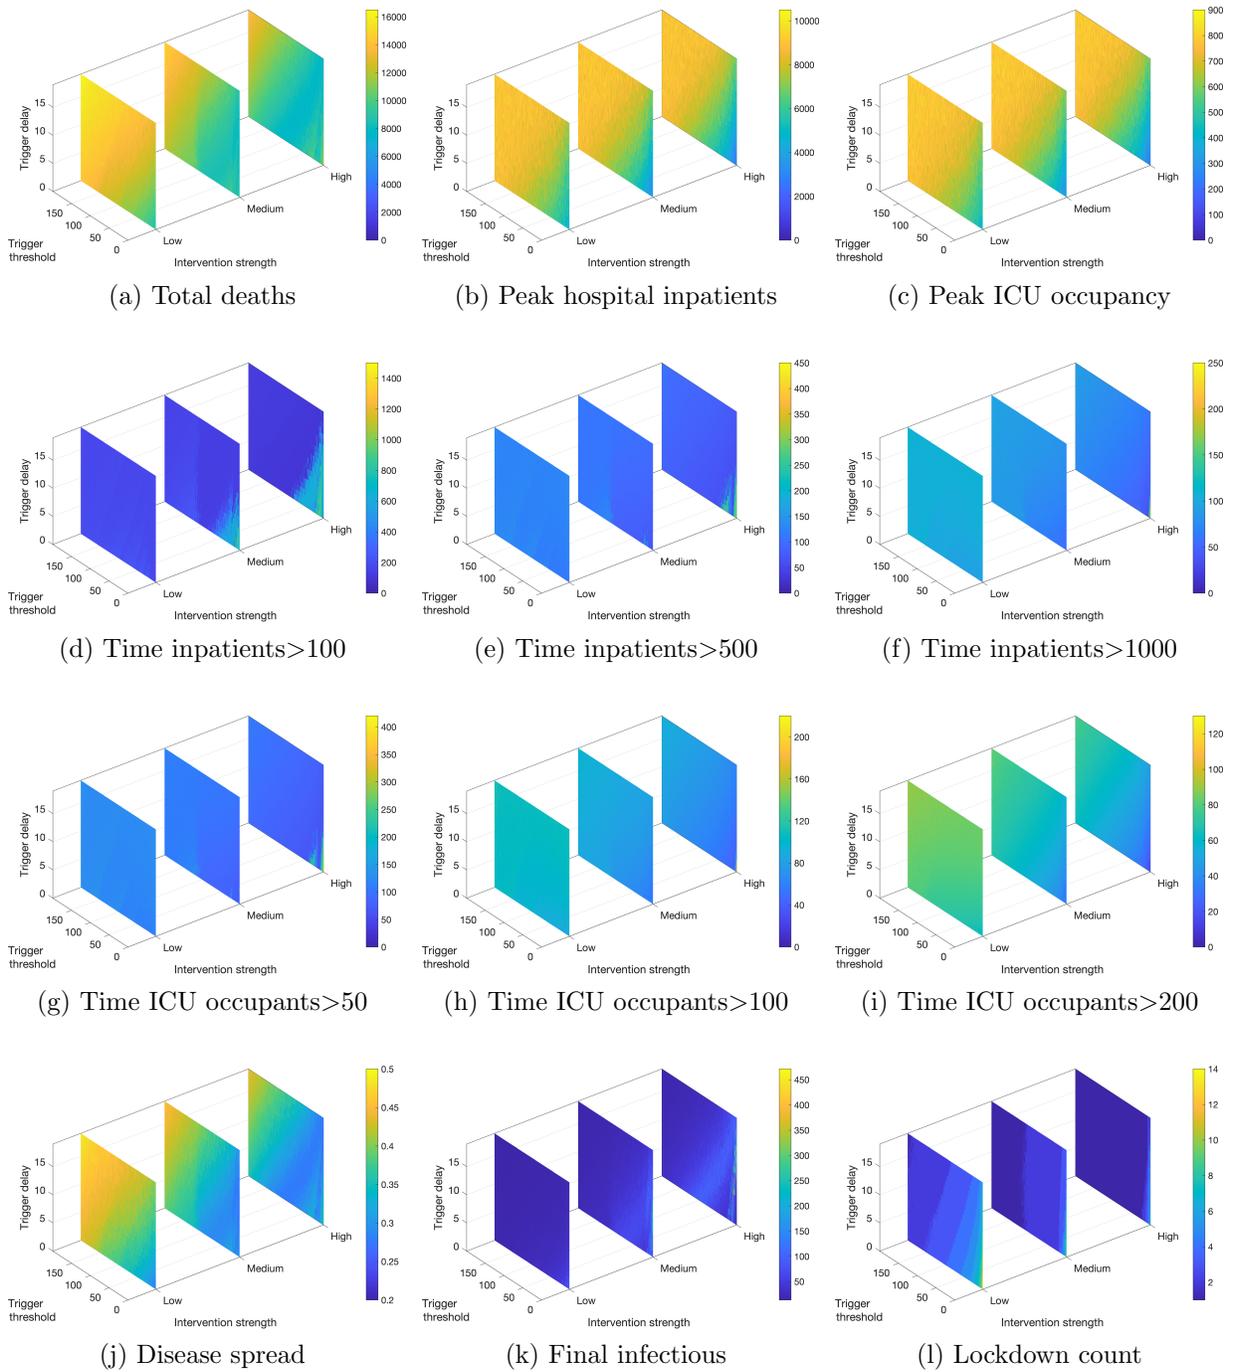


Figure 37: Dynamic lockdowns: Constant inflow - Death triggers - Multiple lockdowns

6.3 Inpatient and ICU occupant ceiling times

In a system with constant inflow, curiously both the best and worst lengths of time for exceeding specific healthcare occupancy ceilings can occur when the intervention is strongest. Regardless of whether using a single or multiple lockdowns, outcomes are poor if medium or strong interventions are activated too soon, but most successful if calibrated correctly (consider Figure 34(e), where the brightest and darkest colours both occur on the high-strength slice). We also observe that critical thresholds, that is the borders in parameter space that separate these best and worst outcomes, at each intervention strength reduce (both in delay and threshold) as the ceiling for hospitalisation and ICU occupancy are increased. In other words, the optimal trigger to minimise the amount of time with more than 500 inpatients in hospital (e.g. Figure 36(e)) occurs at a lower trigger threshold and delay than an optimally activated lockdown of the same strength to minimise the time with over 100 hospital occupants (compare with Figure 36(d)), so strong interventions can be put in place earlier to avert the most severe outcomes as long as they are not too early so as to have little positive effect. This pattern moves the threshold for minimising the amount of time with more than 200 in ICU with a strong lockdown out of the parameter window, so that this outcome alone is always improved by the strength of the intervention (Figure 34-37(i)).

In closed system scenarios, the patterns are more complex, however this same trend still emerges when limited to a single lockdown. In particular, there still exists a critical threshold for triggering interventions, just above which gives the best reduction to the amount of time each ICU or hospital occupancy ceiling is exceeded. At each intervention strength this critical value curve moves down in the parameter space as the ceiling of hospital and ICU occupants increases (for example, follow the trend in the location of the darkest band for a high-strength lockdown from Figure 31(d)-(f) and from (g)-(i)). This is understandable given the previous result in §5.3 of the main report concerning peak occupancy: as earlier interventions tend to be better for reducing the value of the greatest peak, it also makes sense that they would be better for reducing the time that the highest occupancy ceilings are exceeded.

6.4 Long-term simulation of the best possible interventions

To study the long-term behaviour of these system with constant new cases, in each scenario we select the response that yielded the best result in terms of minimising the total number of deaths after 2000 days of simulation. In all cases, the best-timed intervention(s) was always high intensity. We illustrate the time series of the spread of the virus, cumulative deaths, and the number of hospital and ICU occupants over a long-term simulation up to 10,000 days. To aid visualisation, the time axis is displayed on a logarithmic scale so that the behaviour over days 100-1000 when the lockdown(s) are permitted is clearly visible. Time periods when the lockdowns are active are shaded in grey.

However, note that over such a long timescale the assumptions of this age-structured model (principally the invariance of the age classes) would no longer hold, so such forecasts for the 10,000th day are speculative.

6.4.1 Closed system

In a closed system, the best use of single or multiple lockdowns was always using the number of daily new deaths as a trigger.

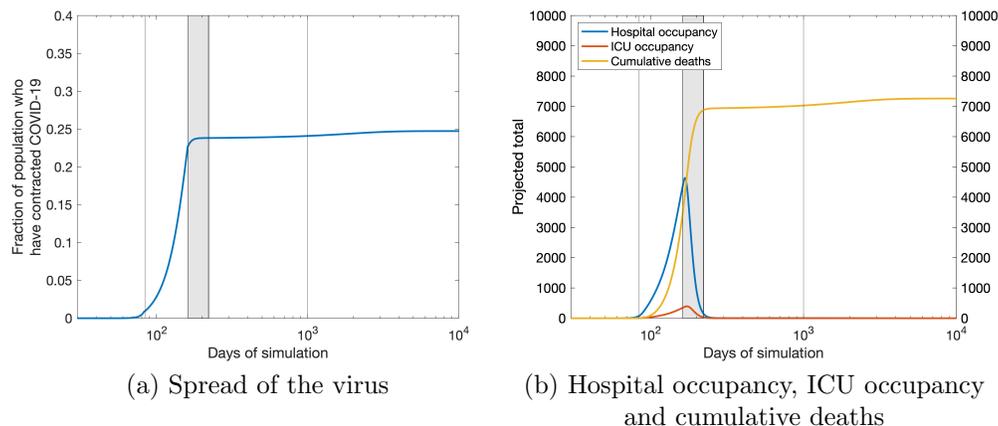


Figure 38: Long-term projection of best single lockdown triggered by daily death thresholds

Figure 38 illustrates the time series of the best result from those tested in Figure 30(a) and 31(a). That is, a single lockdown which is triggered by the number of daily deaths exceeding 67.5, followed by a 20-day implementation delay.

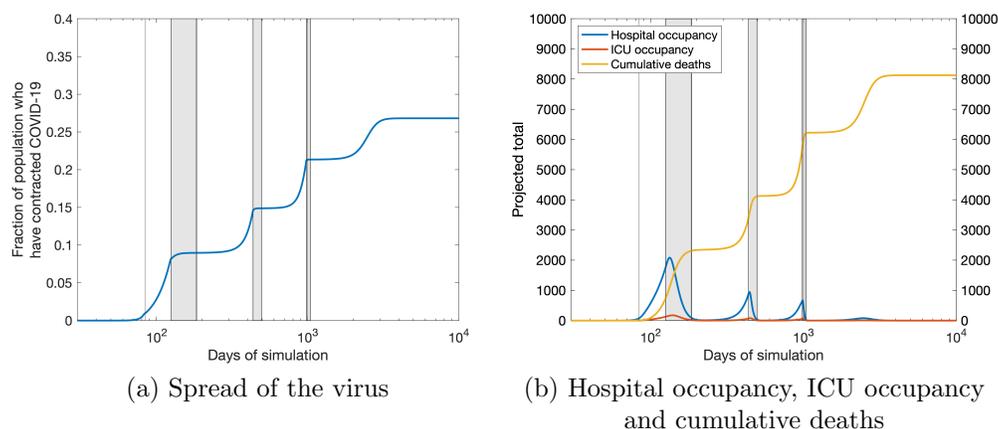


Figure 39: Long-term projection of best series of multiple lockdowns triggered by daily death thresholds

Figure 39 illustrates the time series of the best result from those tested in Figure 32(a) and 33(a). In this case, three lockdowns are triggered: one each day that the number of new deaths exceeds 14.5, following a 20-day implementation delay.

Even after 1000 days of lockdowns and 1000 days without in a closed system, it is still possible for the virus to persist and potentially cause thousands of additional deaths in the very long term. However, even in the absence of any other solution, in a closed system we can at least reduce the total number of deaths from a worst-case scenario of around 16,500 (for example, Figure 31(a), in the regions of parameter space where no intervention was triggered) to 6000-8000 without depending on an unrealistic expectation of eradicating the virus completely⁴.

The best result from a single lockdown (Figure 38) is less effective in the medium term (by day 2000) than the best reduction in overall deaths that can be achieved from multiple lockdowns (Figure 39), however it then becomes the better choice in the longer term by day 10000⁵. Regarding timing, in the closed system it is never best to lock down too early in the pandemic, and indeed if limited to a single lockdown it is better to hold off longer, resulting in one large initial peak but no further waves, whilst optimal strategies permitting multiple lockdowns require earlier interventions and these frequent restrictions result in multiple smaller waves of hospital admissions. There is therefore a trade-off of several factors to consider when selecting a public health strategy: timescale, healthcare capacity, and how many lockdowns the economy and other public services can sustain. If the population can reasonably be considered as a closed system, and if the role of lockdowns is simply to buy

⁴Furthermore, note that these best cases were from the maximum allowable delay of 20 days, so it is possible that an improved results could be obtained by expanding the parameter space and allowing even greater delays until implementation.

⁵Note that this does not guarantee that there is not some other strategy involving multiple lockdowns in the first 1000 days that gives a better result again when deaths are counted at day 10,000.

time for the medium term while a permanent vaccination exit strategy is developed, then multiple lockdowns will yield benefits both to managing peak capacity and reducing overall deaths for several years provided that such a strategy is economically and socially sustainable. However in this model, if there is essentially no hope of a vaccine within several years, then one lockdown as late as possible would be desirable, provided that the health service can survive the single peak burden.

6.4.2 Constant inflow

In a system with constant new infections, the best use of single or multiple lockdowns was always using the current number of hospital occupants as a trigger.

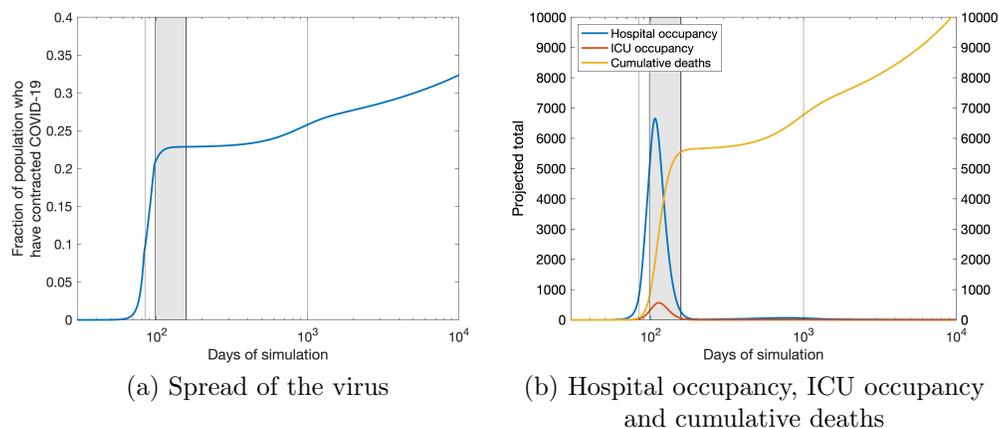


Figure 40: Long-term projection of best single lockdown triggered by hospital occupancy thresholds

Figure 40 illustrates the best result from those tested in Figure 34(a). In this case, a single lockdown is triggered by the number of hospital inpatients exceeding 985, followed by a 12-day implementation delay.

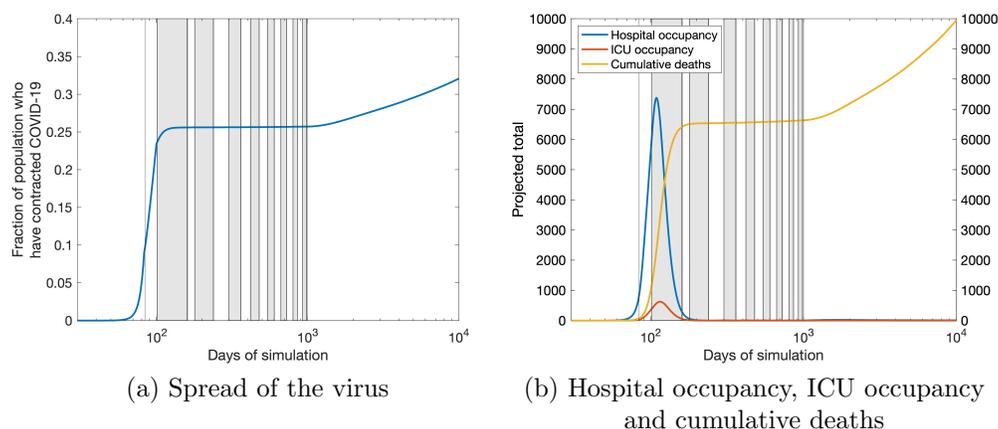


Figure 41: Long-term projection of best series of multiple lockdowns triggered by hospital occupancy thresholds

Figure 41 illustrates the best result from Figure 36(a), where a sequence of eight lockdowns are triggered by the number of hospital inpatients exceeding just 5, followed by a 17-day implementation delay. Recall that there is an element of randomness in simulations with constant inflow and hence these are only indicative examples of the optimal single- or multiple-lockdown strategies. For example, in some repeated simulations with multiple lockdowns and the same parameters, nine rather than eight lockdowns were employed.

The constant influx of infections dampens the effectiveness of any lockdown and so there is not a rapid increase in deaths following the final lockdown, in comparison to the resurgence which may be observed in closed systems (Figure 39(b)). As a result, it is possible (dependent on the randomness in these scenarios) in such a system for the total deaths after 2000 days to be fewer than the best simulations in a closed system model, and the mechanism underlying this is investigated in greater detail in §6.5. However, if there is a truly infinite importation of new cases, however slowly, then inevitably the virus will spread to the entire population with a commensurate number of deaths, assuming that the population is well-mixed. Hence, the number of deaths eventually surpasses that of equivalent closed systems, exceeding 10000 by the 10,000th day in most trialled simulations. In this case, the long-term outcomes are not greatly impacted by restrictions, and so if conditions necessitate modelling the country as a system with constant inflow of new cases, decisions should be taken based on short-term requirements.

6.5 Understanding closed systems versus those with a constant inflow of cases

In many scenarios, particularly those employing strong lockdown restrictions, it appears that enabling a steady arrival of new infections can lead to fewer overall deaths and a lower spread of the virus amongst the population after 2000 days than when using the same strategies in an identical but closed system (for example, compare the strong slices of Figure 30(a) with 34(a)). This seems counter-intuitive, as we would expect that cutting off a source of new infections would always weaken the ability of the virus to spread. To illustrate the mechanism underlying this phenomenon, we consider the corresponding simulations where either a single or multiple high-strength lockdowns are triggered by 400 hospital inpatients with a 10-day delay. The results are shown below in Figure 42.

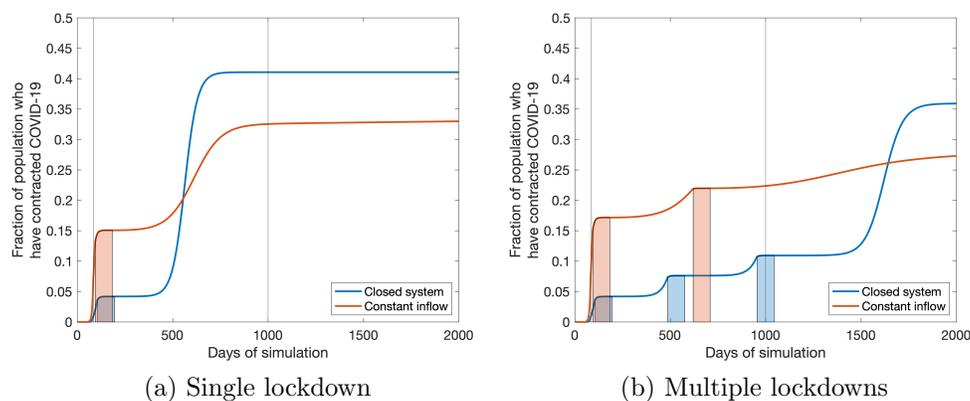


Figure 42: Demonstrating why closed systems may have greater overall spread of the virus

When only one lockdown is permitted, a constant inflow of new infections is very influential from the outset, causing the first wave to grow extremely rapidly compared to the equivalent closed system (Figure 42(a) comparing the red and blue curves), and thus it reaches higher initial peaks of hospitalisation and ICU occupancy. The rise in cases following the only lockdown is consequently smaller by comparison. However, if the system is closed, the virus initially grows more slowly, leading to a smaller first peak which is clamped down on while the rate is not yet so steep. However, the post-lockdown surge in cases is much more rapid, leading to a significantly higher second peak. The increase in spread and hence deaths during this larger second wave results in the closed system potentially overtaking in the number of deaths by day 2000 of the simulation. In either system, the total deaths immediately during the first wave account for less than half of the total, and this difference is greatly exacerbated in the closed system.

If multiple lockdowns are permitted (Figure 42(b)), a similar mechanism occurs although the effect is weaker. Again, if there is a constant feed of additional infections, the initial spread and peak is much greater. In a closed system, growth is slower and repeatedly suppressed by a series of lockdowns, but the propagation after the final lockdown is lifted sees a rapid rise

in both cases and deaths that once again are able to surpass that of the equivalent system with constant inflow.

However, we emphasise that this effect, while long-lasting, is not permanent. As we can see especially from the right-hand edge of Figure 42(b), if the constant addition of one new case per day is maintained forever, then necessarily the whole population will eventually be exposed to the disease assuming that all members of the society interact with each other. But in the scenarios illustrated above, it would take more than 110 or 50 years, in the case of single or multiple lockdowns respectively, for the proportion of the original members of the population who have contracted the disease to overtake their closed system equivalents. This age-structured model assumes both a static population and that members do not transition between age class (although this could easily be incorporated for a longer-term model), so clearly such long term trends cannot be accurately forecast. In any case, it would be reasonable to assume that the inflow of new cases could be controlled after such a time, by vaccination, testing or otherwise.

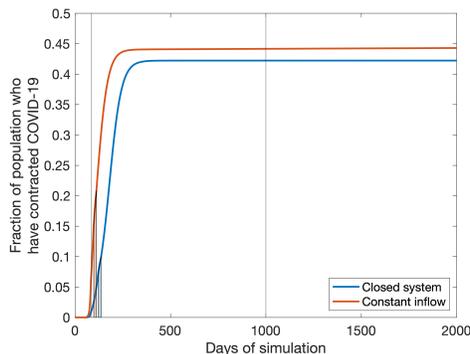


Figure 43: Comparing closed and constant systems with only weak interventions

Finally, note that this effect is not observed by day 2000 when weaker interventions are used (compare the weak slices of Figure 30(a) with 34(a)). This is because the weak interventions do not significantly slow the progression of the pandemic and so the increased rate of infection due to the constant extra cases does behave as we would intuitively expect. The subsequent resurgence of cases in a closed system is not so great, and thus it fails to overtake the case numbers of the equivalent system with constant inflow. An example comparison is shown in Figure 43 where a single weak-strength lockdown is triggered by 1000 hospital inpatients with a 10-day delay.

6.6 Comparison after 1000 days or 2000 days of lockdown

All illustrations in Figures 30-37 are the result of simulations with duration of 2000 days, where lockdowns were only permitted to occur during the first 1000 days. However, we also compared these with measuring the best outcomes immediately after the first 1000 days, and with the results of simulating 2000 days where lockdowns were permitted to potentially occur during the entire duration.

If only a single lockdown was allowed, we observe that if it occurs, this certainly will take place within the first 1000 days. Consequently, there is little change between days 1000 and 2000 as both the majority of deaths and the peak healthcare demands will have already passed. However, if there is a constant inflow of cases, then naturally we observe a slight increase in spread and number of deaths between 1000 and 2000 days, particularly for high strength interventions.

If multiple lockdowns are allowed, then in a closed system triggering these lockdowns as easily and frequently as possible seems to give excellent results if the outcomes are measured immediately after 1000 or 2000 days of lockdowns (Figure 44). However, compare this to Figure 32(a), and we see that in such circumstances the brunt of the pandemic is simply being delayed until a time when no further lockdowns are active - even if this is after a very considerable period of time has passed. The maximum number of deaths reached in this case does not quite match that of the circumstance where no lockdowns are triggered, but it is otherwise the worst possible condition of these scenarios. This post-lockdown “bounce back” by the disease explains the phenomenon identified in §5.2 of the main report whereby low-threshold interventions ultimately yield very poor benefits to reducing cumulative deaths during the full course of the pandemic.

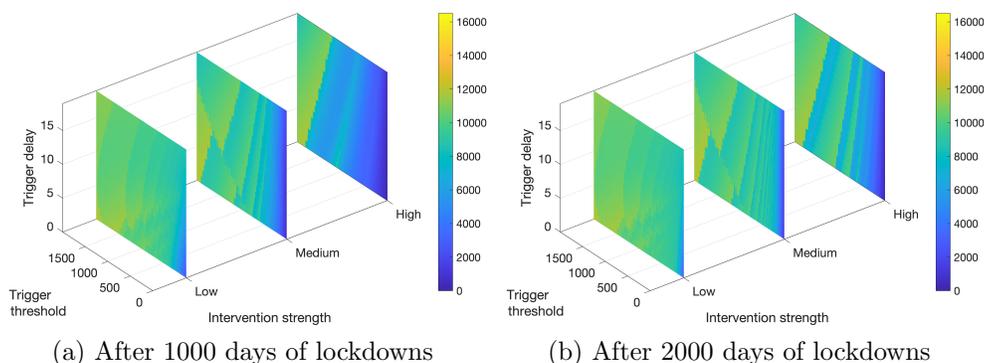


Figure 44: Total deaths: Closed system - Hospital triggers - Multiple lockdowns

In model systems with a constant inflow of cases, as previously discussed the overall suppressing effect of lockdowns are reduced and consequently so is the possible resurgence after no further interventions are permitted. Between days 1000 and 2000 of the simulations, even if further lockdowns are permitted during the second 1000 days there is an inevitable slight

increase in total deaths and spread throughout (compare Figure 45(a) with (b)). This is clearest for high strength interventions with a low-medium trigger threshold. In such circumstances, new peak hospital or ICU occupancies are not reached, although there is an increase in the number of days with over 100 patients in hospital, and these increase again once further lockdowns are not permitted.

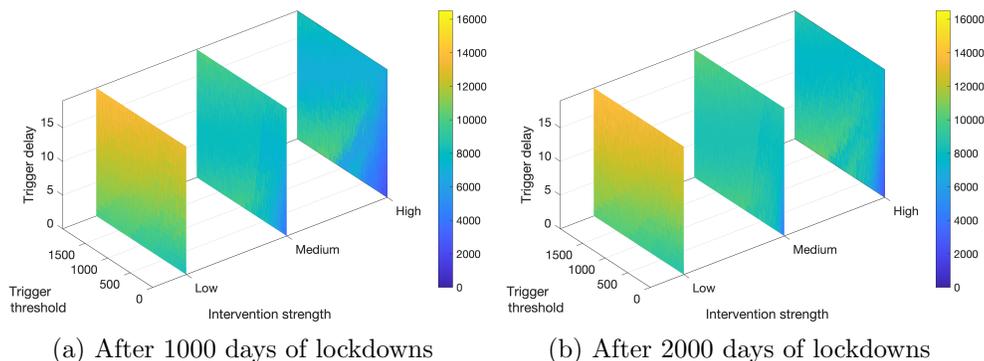


Figure 45: Total deaths: Constant inflow - Hospital triggers - Multiple lockdowns

As before, medium and high strength interventions with a very low trigger threshold are able to give the best reduction to spread and total deaths while the lockdowns can still be activated (Figure 45), but these benefits fade significantly once further lockdowns are no longer permitted (compare with Figure 36(a) where the lowest threshold is no longer optimal). Resurgence is milder than in a closed system, with a greater peak of deaths being only just achieved which is nonetheless lower than if only weak lockdowns with long trigger delays were employed. This suggests that the constant introduction of new infections from outside the system acts as a release, preventing the catastrophic rapid increase in post-lockdown cases observed when the same strategies were employed in the equivalent closed system.

6.7 Vaccination

In this section we expand on the impact of vaccination on these results discussed in §5.4 of the main report. A vaccination programme began in Northern Ireland in December 2020, and to represent this, every day from 1st January 2021 a constant number of individuals immediately transfer from the susceptible to the recovered compartments, and are removed from the pool of individuals used to calculate the number of new infections on that day. It is not unreasonable that in the future a vaccine may be available approximately one year after the significant emergence of a new virus. As we treat the transmission rate as an average of fitted values of β , we obtain generic results which may be applicable to future pandemics.

We compare several simple vaccination schemes: an age-structured distribution of two rates ($v(i)$ individuals of age class i , where either $v = \{200, 400, 650, 1250, 2500\}$ or $v = \{400, 800, 1300, 2500, 5000\}$), and three rates of uniform vaccine distribution amongst all age classes (that is, 400, 1600, or 3200 individuals per age class per day). Thus, the age-structured strategies are such that either 5000 or 10000 vaccines are administered per day, and in Northern Ireland the 7-day moving average of daily vaccinations administered by the Department of Health varied between approximately 3000 and 18,600 during January-August 2021. In each case, the total number of deaths over a 2000-day simulation with either one or multiple lockdowns in the first 1000 days is shown in Figures 46-47 along with the reproduced results in the absence of vaccination (from Figures 34(a) and 36(a)). In all cases we restrict to the more realistic scenario of constant inflow of one daily new infection of a random age class, with the number of hospital inpatients used as the lockdown trigger.

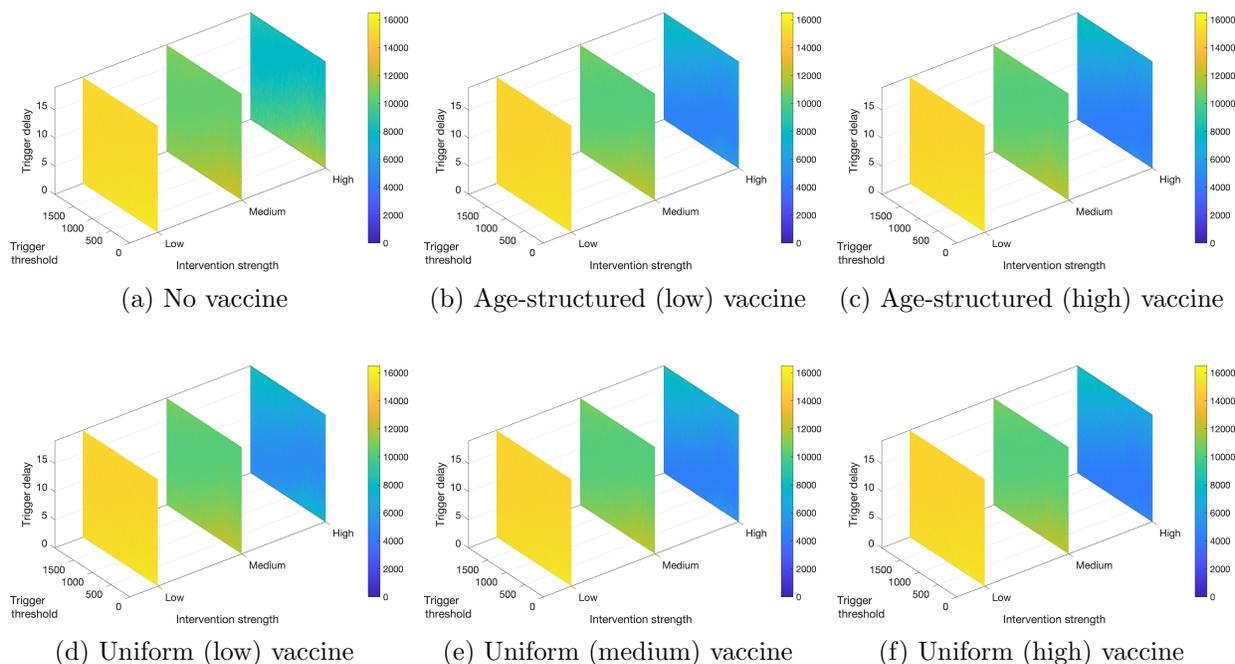


Figure 46: Total deaths using hospital triggers for a single lockdown in a system with constant inflow

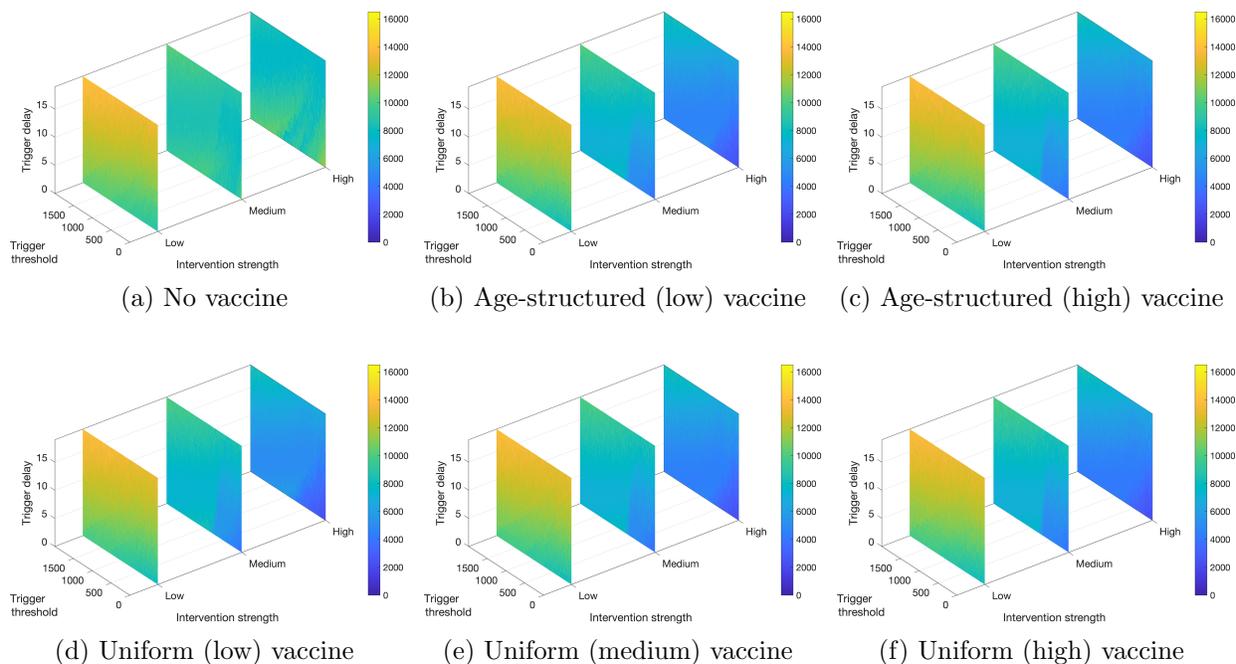


Figure 47: Total deaths using hospital triggers for multiple lockdowns in a system with constant inflow

When one (low or intermediate) lockdown is permitted, the outcomes are independent of the existence and rapidity of a vaccine (Figure 46). The virus will already have spread to most of the population (and most deaths will already have occurred) during the first year. However, a high-strength lockdown can significantly curtail transmission during this period, and then a sufficient vaccination programme will reduce the total deaths from 10,000-12,000 to 5000-6000, of which about 3200 will have occurred by the beginning of vaccination. Increasing the vaccination rate reduces the remaining deaths further, but the contrast is minor compared to the absence of a vaccine.

When a vaccine will be available and multiple lockdowns are permitted, it is desirable to trigger them at the lowest sensitivities (Figure 47) - as immunity gained from the vaccine replaces the ultimate solution from herd immunity. The rate of vaccine administration again has little impact, as fewer lockdowns are mechanically-triggered to achieve these same effects given a faster rollout. Thus if multiple but finite moderate lockdowns were permitted, they would prevent a majority of deaths in the first (pre-vaccine) year, and then further lives could be saved by a faster vaccine distribution beyond that point. Similarly, for multiple strong lockdowns, there is some difference made by the value of the trigger at low sensitivities as the faster the vaccine rollout, the less low that trigger needs to be in order to prevent what additional deaths may be possible beyond the first year.