This paper has not yet undergone external peer review

Estimates of effects on changing Alert Levels for the August 2021 outbreak

Report delivered 9 September, 2021
This version compiled: October 15, 2021 Author list alphabetical

James Gilmour[1,2], Emily Harvey[1,3,*], Joshua Looker[4], Frank Mackenzie[1,2], Oliver Maclaren[4], Dion O'Neale[1,2,*], Frankie Patten-Elliott[1,2,4], Joel Trent[4], Steven Turnbull[1,2], and David Wu[4]

[1]Te Pūnaha Matatini, University of Auckland, NZ
[2]Physics Department, University of Auckland, NZ
[3]M.E. Research, Takapuna, Auckland, NZ
[4]Engineering Science, University of Auckland, NZ
[*]Corresponding author e-mail addresses: emily@me.co.nz; d.oneale@auckland.ac.nz

EXECUTIVE SUMMARY

We use an individual-based, Aotearoa-specific Contagion Network Model (CNM) to simulate the spread of COVID-19 in the community for an outbreak comparable to that detected 17th of August 2021. The CNM explicitly simulates spreading processes and the effect of any interventions, in order to predict the spread of COVID-19. It is therefore able to avoid the need for input assumptions or external estimates of the effective reproduction, R_{eff}, number during the outbreak.

We use this model to consider two scenarios:

i) An intervention with parameters comparable to *Alert Level 4 is applied from 18th August onwards and remains in places for the duration of the simulations*; and

ii) An intervention with parameters comparable to *Alert Level 4 is applied from 18th August to 15th September, after which the intervention parameters are changed to match those expected for Alert Level 3.*

We complement the CNM simulations with results from a Branching Process Model (BPM). We use this to simulate scenarios where an Alert Level change on September 16th results in a multiplicative increase of R_{eff} of factors of 1, 1.5, 2, or 3 times that estimated during AL4.

After 60 days at AL4, it is predicted that almost all trajectories will be either eliminated or contained (i.e. all cases are either in quarantine or isolation). In contrast, only a few trajectories are eliminated or contained for the optimistic and pessimistic AL3 scenarios. In the former, many trajectories exhibit suppression like behaviour, remaining at low numbers (under 20 new cases per day) or growing only slowly.

The pessimistic AL3 scenario shows a growing number of daily new cases for the majority of trajectories. The pessimistic AL3 scenario also suggests that it would initially be difficult to distinguish between uncontrolled growth and suppression or elimination like behaviour — it takes approximately another 10 days post-deescalation before the growing trajectories start to become distinguishable.

1 Introduction

We simulate a community outbreak of the Delta variant of COVID-19 detected in Auckland on August 17th, with initial spread from an unidentified seed case prior to this date while Aotearoa was at AL1. Detection of this initial case is followed by a period at Alert Level 4. We consider two scenarios: one where there is a de-escalation of intervention measures from AL4 to AL3 in the middle of September, and another where no de-escalation occurs. We use both a simple Branching Process Model (BPM) and an individual based, Aotearoa-specific Contagion Network Model (CNM) to estimate the likely consequences of such a de-escalation, subject to a number of assumptions.

The BPM requires, as an input, estimates of the effective reproductive number, R_{eff}. We can inform estimates of R_{eff} and the effects of changing intervention levels by considering data from previous outbreaks and the present outbreak. However, past outbreaks did not involve the more contagious Delta variant of COVID-19, and hence the changes of R_{eff} when moving from AL4 to AL3 are uncertain. The CNM, on the other hand, uses an explicit representation of intervention measures and can therefore be parameterised by directly using estimates related to these interventions.

For both the BPM and the CNM we employ Gaussian Process (GP) conditioning to produce predictive bands conditional on the observed data. This conditioning works by first fitting an approximate GP to realisations of the full simulation models. Next, the approximate GP is exactly conditioned on the data to produce conditional prediction bands. In addition to predictive bands, we also retain and plot the simulations lying within these bands on future predictions. Note, however, that by only retaining predictions consistent with the future predictions, this filtering may include some simulations which give similar predictions but a less close match to the data observed so far. This issue is avoided somewhat by conditioning on cumulative case data. Still, we plan to implement an alternative simulation filtering step in future using Approximate Bayesian Computation (ABC), which will further remove those realisations that give consistent predictions but inconsistent histories.

[Note: In this update, we now also consider results from conditioning the same BPM ensemble on case data using ABC[1]. We use a sum of squares distance measure, retaining 1% of simulations that best match the observed data with this measure.]

2 Simulation settings

2.1 Contagion Network Model

The contagion network model employs a bipartite interaction network that connects around 4.7 million individuals — with attributes of age (split into four age bands), sex, and ethnicity — to group nodes representing the interaction contexts where possible disease transmission between individuals can occur.

Interaction contexts in the network have sizes and attributes determined from a mix of empirical data and from literature on social interaction networks. For example, dwellings have a geographic location (at SA2 level), a size, an age structure (number of individuals in each age band), with factors such as assortativity of ethnicity and the correlation between, e.g. ethnicity and dwelling size, determined at SA2 level, whenever possible.

The interaction network attempts to capture the heterogeneous structure of the New Zealand population, and the correlation of factors that occur within it. This helps the CNM to not only reproduce as many as possible of the factors that may influence transmission of COVID-19, but also to better address the issue of the burden of disease falling disproportionately on certain groups.

Model parameters are described in the Appendix, but important assumptions and parameters relevant to this report are:

- People can only infect people who they interact with through contexts specified in the interaction network. These contexts are broken into the four categories: Dwelling, Work, School, and Community. The different context types have different levels of infection risk, which also changes according to different Alert Level interventions. Contexts are classified as either Close or Casual, with interactions through Casual contexts having a much lower infection risk than those classified as Close.

  The Close and Casual status of interactions is also used to determine which aspects of a Test-Trace-Isolate (TTI) policy are applied to an individual. For example, Close contacts of a confirmed case are queued in the contact tracing process at a higher priority than Casual contacts.

  N.B. the definitions of Close and Casual applied in the model do not map exactly to the definitions of Close, Close Plus, and Casual that have been used in reporting of the August 2021 outbreak. However, the TTI policies applied in the model, for an individual from a particular interaction context and Close/Casual status, are designed to match as closely as possible those used in the August 2021 outbreak.

- Cases that develop symptoms have a propensity to seek testing for COVID-19 — this propensity is low during the period when cases have not yet been detected and the country is at AL1. The propensity of symptomatic individuals to seek testing is elevated after the first case of COVID-19 has been confirmed. The time from symptom onset to seeking a test also varies, being shorter once the first case is detected.

  People who seek a test themselves (i.e. not those identified via contact tracing) are not required to isolate while awaiting a test result.

- The proportion of infections that are asymptomatic varies with age, and equates to about 16% over the whole population, in line with findings from PCR based studies with inclusive symptom case definitions[2,3]. Asymptomatic cases are assumed to have zero chance of being tested for COVID-19 unless they have been identified as a casual or close contact of a confirmed case.

- All cases that are detected (confirmed cases) initiate a contact tracing process for all ‘close’ contacts, and for ‘casual’ contacts in large dwellings, schools, and workplaces. The time from positive test to isolation of contacts differs by context, with dwelling contacts traced faster, then work and school, then community. Close contacts are traced faster than casual contacts.

- Casual community contacts are not identified via contact tracing. Instead they have a higher propensity to seek a test, regardless of symptoms, once they find out that they are casual contacts — i.e. once an associated case is confirmed. this is intended to capture the effect of people changing their behaviour if they know that they have been at a Location of Interest.

- The initial Alert Level intervention is applied one day after the first case is detected. Contact tracing processes, however, begin immediately once the first case is confirmed.

- At each Alert Level, we are set a proportion of interaction contexts to be Closed (i.e. individuals do not interaction through these group nodes at the specified Alert Level). This setting differs for each interaction context type and for close and casual contacts.

  Within the interaction contexts that remain Open we specify a relative transmission reduction. This represents the impact of non-pharmaceutical interventions including physical distancing measures, mask wearing, adjusting work shifts, and other behaviour changes.

- Vaccination in the model is allocated to individuals using the number of people within each age band, sex, ethnicity, and DHB that had been fully vaccinated (two doses) by 2nd September, 2021. This allows approximately two weeks for the vaccine to become effective by the date of the Alert Level de-escalation considered in the simulations. Within each age-band/sex/DHB combination vaccination is allocated randomly.

Test trace and isolate settings

The August 2021 outbreak has seen a number changes to public health interventions in order to better counter the increased threat posed by the Delta variant. A key component of these changes has been the approaches taken to the Test-Trace-Isolate (TTI) settings. The CNM attempts to emulate the policies from the current outbreak, based on what actions are recommended (or required) for exposed and/or infected individuals. Parameters related to these policy settings are inferred from contact tracing and case data provided by Auckland Regional Public Health Service (ARPHS).

The CNM uses an ensemble of simulations to cover a range of parameters for: rates of symptomatic testing in the community; proportion of close contacts known (and traced); and proportion of casual contacts who would seek a test. Parameters related to TTI policies are unchanged between AL4 and AL3. These are[*]:

- proportion of *symptomatic* individuals who would seek a test and test positive in AL4 and AL3: [10%, 40%, 80%].
- proportion of *close contacts through school or work contexts* who would be known (contact traceable): [95%, 98%]
- proportion of *close contacts through community contexts* who would be known (contact traceable): [85%, 95%]
- proportion of *casual community contacts* who would seek a test regardless of symptoms (i.e. individuals who were not contact traced but who would seek a test after discovering that they were at a Location of Interest): [10%, 50%]

Specific differences in TTI policies that we do not include are:

- The requirement for household contacts of close contacts to isolate as well, until the day 5 test result is returned (I.e. if a dwelling member is contact traced and notified that they must isolate, on account of being a close contact of a confirmed case, the other inhabitants of the dwelling are not require to also isolate).
- The requirement to isolate while waiting for a test result.

These differences are not expected to have a significant impact on results at elevated Alert Levels. In particular, at AL4 individuals, other than essential workers, already have very low numbers of interactions beyond their dwelling (schools are closed and 95% of potential close community interactions do not occur).

AL4 intervention

In Alert Level 4 everyone is meant to stay in their ‘bubble’, all schools are closed, only essential workplaces are open, and those workplaces that are open must follow strict infection prevention practices. There is also limited allowance for ‘shared’ bubbles between households for reasons including shared custody, someone who lives alone, and care-giving responsibilities.

We implement these interventions as follows:

- We set 100% of schools to be closed.
- We represent the impact of Alert Level 4 on workplaces using estimates from the StatsNZ Household Labour Force Survey estimates for proportion of workers working on site at different Alert Levels in 2020[4]. At AL4, this corresponds to ∼30% of all workplaces, but with large differences by industry sector (at 1-digit ANZSIC06 level).
- We assume a 90% reduction in casual contact workplace transmission risk and 80% reduction in close contact workplace transmission risk due to workplace safe operating requirements at Alert Level 4.
- For all other non-work, non-school close contacts (community interaction contexts), we assume the restrictions mean that 95% of close contact community events wouldn’t occur, and that all events are limited to 10 people or fewer. This results in ∼9% of the population who have close contact with someone outside their dwelling. This in turn equates to around ∼15% of household bubbles (dwellings) having a connection to another bubble (dwelling). These figures are derived from two main pieces of work. First, a quantitative survey conducted in New Zealand during April 2020, which found than there was an average of 1.26 households in each bubble[5]. Secondly, analysis of data from a repeated survey commissioned by Ministry of Health that asked people about their contacts during different Alert Levels in New Zealand in 2020, found that ∼25% of survey participants had at least one close contact during Alert Level 4 who was not in their household or a work colleague[6].
- We assume that for these close community contacts that do occur, we assume that there is a 25% reduction in transmission risk.
- For casual contact community events (e.g. shopping, transport, etc.) we assume that 95% don’t occur and that for those that do occur, there is a 75% reduction in transmission risk, due to physical distancing measures, mask wearing, etc.

[*]Lists of values indicate that scenarios were run for each of the listed values.

- Finally, for the casual community events that do still occur, we enforce an event size limit of 10. We assume here that, if an event is larger than that, that it does not happen, as opposed to keeping the same number of events but reducing their size.

AL3 intervention

At Alert Level 3 there are a number of changes from the settings above. For close contact events, the most important changes are that many workplaces can reopen, people can 'expand' their bubble, and the children of essential workers can go back to school. There is also an increase in casual contacts through the community through takeaway food and click-and-collect retail. We do not have a clear idea of the details for how these changes will manifest at Alert Level 3 for an outbreak of the Delta variant. We therefore set parameters for both an ***optimistic AL3*** and a ***pessimistic AL3***.

- We set all schools as potentially open, but with a transmission risk reduced by 97% to represent the estimated 3% of students who attend at Alert Level 3 , as children of essential workers.
- We represent the impact of Alert Level 3 on workplaces using the StatsNZ Household Labour Force Survey estimates of proportion of workers working on site under different Alert Levels in 2020[4]. This equates to $\approx 53\%$ of total workplaces, but, as with AL4, there are large differences by 1-digit ANZSIC06 industry sector.
- We assume that transmission reduction measures in workplaces are not as strong as at AL4, with either a 50% or 70% reduction for close contacts (pessimistic and optimistic), and a 85% reduction for casual contacts.
- At Alert Level 3, survey data[6] indicates that the number of people with non-work close contacts outside their household doubled from AL4.
- We set the proportion of close community events that are closed to 70% and 80% (pessimistic and optimistic scenarios, respectively). This leads to around 45% (resp. 32%) of people having at least one contact outside their household. These interactions are still less risky than at AL1, due to restrictions on mass gatherings, bars, church services, and other community events, so we set the transmission reduction for these interactions to 25%.
- For casual contact community events (e.g. shopping, transport, etc.), there is a lot more interaction due to click-and-collect, takeaways, and other customer/business interactions. We assume that there is a 90% in the number of events from AL1, and that for those event that do occur, there is a 75% reduction in transmission risk due to physical distancing measures, mask wearing, etc.
- Finally, for the casual community events that do still occur, we enforce an event size limit of 10. Again, we assume that if an event is larger than that limit it does not happen, as opposed to keeping the same number of events but reducing their size.

2.2 Branching Process Model

In addition to the contagion network model, we also consider a simplified branching process model (BPM) using a broad ensemble of parameters across 14000 model realisations. We validate (and centre) these parameters on 2020 data, before using them (updated for Delta) on data from the August 2021 outbreak, as at 8th September. The parameters used can be seen in Table 1. Our R_{eff} is increased for the 2021 outbreak due to the higher infectivity of the COVID-19 Delta variant[7]. We use our BPM simulations as an empirical prior for a Gaussian process (GP). We then condition the GP on confirmed cumulative case numbers to get prediction bands for our forecasts, as well as retain simulation realisations lying within these bands. *This may result in a higher expected value of R_{eff} than estimated by the other BPM presented by TPM, because simulations retained are good predictors of the future, but don't necessarily predict the past as well as simulations filtered using Approximate Bayesian Computation (ABC).*

We also present results using the same ensemble when conditioned on cumulative case data using an ABC approach[1]. This is done using ABC with a sum of squares distance measure between realisations and data, retaining 1% of simulations with the smallest value of this distance. These results likely provide better estimates of R_{eff} during Alert Level 4 than those presented by conditioning using the GP, because they better represent past data. However, ABC relies more heavily on the availability of a large number of simulation realisations, and the number used here (14,000) may be too small. An ensemble with a larger number of simulations (e.g. 100,000), will provide more accurate estimates of R_{eff}. Using a coarser acceptance tolerance may also be a reasonable compromise that will also provide more a conservative range of estimates.

AL4 intervention

The parameters used in the BPM for the initial period of AL4 post-intervention measures correspond to an estimated final R_{eff} value of between 0.05 and 2.4. Conditioning of the ensemble using GP reduces R_{eff} to a median of 0.56 with an IQR between [0.50, 0.63] for retained simulations. This corresponds to an estimated proportional decrease in R_{eff}, due to the Alert Level 4 intervention, of between 0.10 and 0.12 $([0.5, 0.63] \div 5)$. This post-intervention reduction is greater than that estimated for Alert Level 4 on the March/April 2020 Outbreak of between 0.15 and 0.24 ($[0.28, 0.43] \div 1.8$)[8]. This may be due to household saturation factors or improved contact tracing since that outbreak.

In contrast, conditioning the ensemble using ABC reduces R_{eff} to a slightly lower median of 0.52 with a wider IQR between [0.32, 0.72] for retained simulations.

We note that the mean times from symptom onset to test result given in Table 1 are very short. This is to account for the lack of contact tracing built into the BPM; in reality some people will test positive when asymptomatic or pre-symptomatic (due to contact tracing). This is why a low mean symptom onset time is used within this simulation.

Based on experience with the network model, post-intervention R scaling is initially $0.8\times$ the pre-intervention R value, before linearly reducing to the values given in Table 1 over an average 6 days. This linear decrease is intended to capture the effect of a post-intervention lag in reduction of R, with infections occurring within households of infectious individuals for a period after the initial intervention. This effect decreases with time, based on the assumption that all the members of a dwelling will eventually either be already infected or will isolate away from the infected individual(s). It should be noted that the R scaling parameter decreases R_{eff}, but is not the only thing that affects R_{eff} post-intervention (for example R_{eff} is also a function of the testing rate).

AL3 intervention

To estimate the effect of a change in Alert Levels occurring on 16th September 2021 we simulate a change in the expected effectiveness of R_{eff} occurring at this date. This corresponds to an increase in R_{eff} by a factor of 1.5, 2, or 3, relative to the value for Alert Level 4. We also simulate the case of remaining in AL4 with no change in R_{eff} .

The value of R_{eff} in August/Sep 2020, when Auckland was at AL3, was estimated to be between 0.6 and 0.8[9]. This is comparable to 1–1.5 times the estimate of R_{eff} for spread during the AL4 period of the August 2021 outbreak. Hence, we would expect the value of R_{eff} due to a move to Alert Level 3 in mid-September 2021 to be around 1.5 times the value observed at AL4, or between 0.75 and 0.95, when accounting for the increased infectivity of the Delta variant (Table 4). We also include, more more pessimistic scenarios modelled by increasing R_{eff} by a factor of 2 and 3. These represent the situation of spread in an increased transmission environment, relative to that previously seen in AL3 type interventions.

Parameter	Minimum Value	Maximum Value	Average Value
Basic reproduction number R (2020)	2.5	3.8	3.15
Resultant pre-intervention R_{eff} (2020)	2.1	3.2	2.63
Post-intervention R scaling lower bound (2020)	$0.20\times R$	$0.40\times R$	$0.3\times R$
Basic reproduction number R (2021)	5.5	6.5	6
Resultant pre-intervention R_{eff} (2021)	4.6	5.4	5[7]
Post-intervention R scaling lower bound (2021)	$0.01\times R$	$0.5\times R$	$0.25\times R$
Estimated Post-intervention R_{eff} after conditioning (2021, AL4)	0.50	0.63	0.56
Speed of linear decrease to R scaling post-intervention	2 days	10 days	6 days
Pre-intervention symptomatic testing rate	0.05	0.10	0.075
Post-intervention symptomatic testing rate	0.7	1.0	0.85
Pre/post-intervention asymptomatic testing rate	0.0	0.3	0.15
Pre-intervention mean time from symptom onset to test result	1.0 days	3.0 days	2.0 days
Post-intervention mean time from symptom onset to test result	0.5 days	1.5 days	1 day
Number of seed cases	1	4	2.5
Number of cases detected before ‘detection’	1	1	1
Infectiousness after isolation	0	0	0

Table 1. Parameter settings for Branching Process Model Ensemble. Parameters used in models in the ensemble are uniformly distributed between their minimum and maximum values. Symptom onset to test result is exponentially distributed with the mean value given in this table. Other parameter values are based on those used in a similar model, described in Hendy et al.[10], unless otherwise specified.

3 Results

3.1 Contagion Network Model

We simulate a period at AL4 from August 18th to September 14th. At this point in time, the ensemble of simulations has a median of 4 [LQ=2, UQ=7] daily new cases. From September 15th onwards, we then simulate three scenarios: remaining at AL3; deescalating to an optimistic AL3; and deescalating to a pessimistic AL3. The parameterisation of each of these scenarios is described in section 2.1. We also consider a range of TTI parameter values within each of the scenarios.

We use the daily case data (midnight to midnight) from 9th September 2021 to select the simulations that are most consistent with the current outbreak. We then separate the selected (retained) simulations into the three scenarios: Alert Level 4, optimistic Alert Level 3, and pessimistic Alert Level 3.

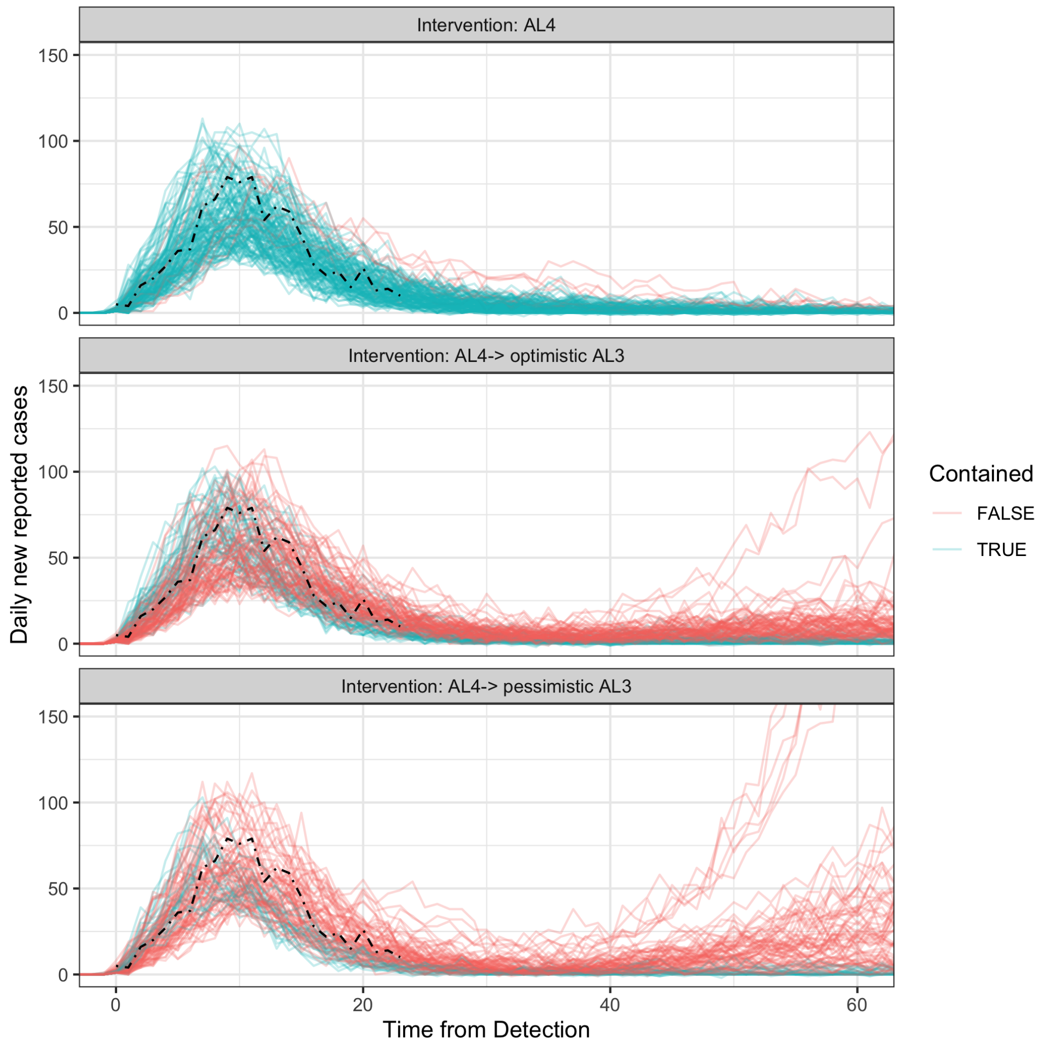

Figure 1. Simulation results coloured by whether the outbreak is contained (all active cases in isolation or quarantine before reaching 10,000 cases or 140 days), for AL4, AL3 optimistic, and AL3 pessimistic. Here we keep the 30% of simulations that are closest to observed daily case data (using least squares difference over the 23 days we have data for). The black dashed lines are the daily reported cases (midnight to midnight) until September 9th.

Figure 1 shows the daily reported cases for all filtered trajectories for the three scenarios considered. We see that after 60 days at AL4, it is predicted that almost all trajectories will be either eliminated or contained (i.e. all cases are either in quarantine or isolation). In contrast, only a few trajectories are eliminated or contained for the optimistic and pessimistic AL3 scenarios. In the former, many trajectories exhibit suppression like behaviour, remaining at low numbers (under 20 new cases per day) or growing only slowly. The pessimistic AL3 scenario shows a growing number of daily new cases for the majority of trajectories. The pessimistic AL3 scenario also suggests that it would initially be difficult to distinguish between uncontrolled growth and suppression or elimination like behaviour — it takes approximately another 10 days post-deescalation before the growing trajectories start to become distinguishable.

Rather than using a simple filtering on observed daily cases, as in Figure 1, we also use a Gaussian Process surrogate to construct conditional predictive bands, conditioned on cumulative confirmed cases in the current outbreak. Here we use case numbers as announced in the 1pm daily briefings, but since we consider the cumulative case count, this will not differ substantially from the midnight to midnight counts used in Figure 1. We further retain simulations consistent with these prediction bands to estimate the median and the credible interval of 95% quantile bands for the three Alert Level scenarios (Figures 2 and 3). Conditioning on cumulative cases and then filtering, as opposed to conditioning on daily new cases, is used here as it typically gives more reliable approximate simulation filtering for projecting forward, for reasons discussed earlier (despite looking like a 'less good' fit to daily observed data). As noted, we plan to complement this with ABC-based simulation filtering in future reports.

These figures help to reinforce the subtle, but important, distinction between the elimination behaviour seen in Figure 2 and the suppression-like and slow growth behaviour seen for the optimistic AL3, shown in Figure 3(left). The pessimistic AL3, shown in Figure 3(right) reinforces the illustration of the initial period of slow growth, post-deescalation before the uncontrolled growth in daily case numbers becomes apparent.

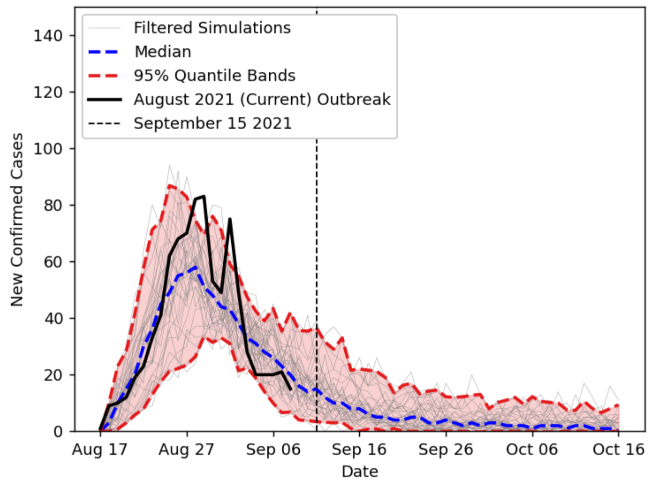

Figure 2. Daily new confirmed infections for an ensemble of contagion network simulations conditioned on cumulative cases in current outbreak. These simulations stay in Alert Level 4 until elimination.

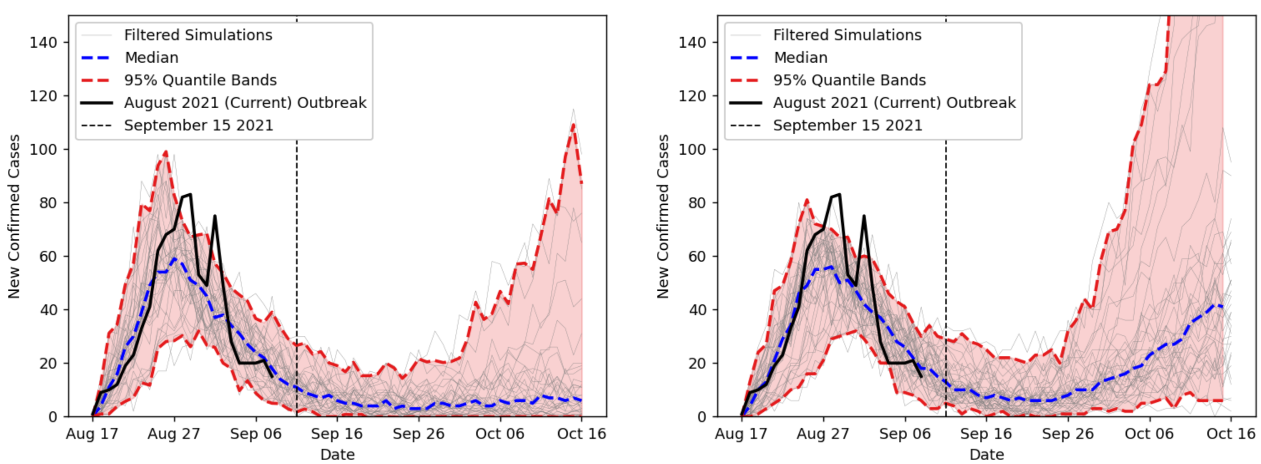

Figure 3. Daily new confirmed infections for an ensemble of contagion network simulations conditioned on cumulative cases in current outbreak. These simulations change to an optimistic AL3 (left) and pessimistic AL3 (right) on the 15th of September, 2021.

Because the network contagion model explicitly represents the different contexts of infection, we are able to look at when infections occur and through which contexts. In Figure 4 we see how there is a large amount of spread through community events (EV & EX) in the early stages of the outbreak. Post-intervention, there is a sudden decline in infections through every context, except dwellings. We do not increase the rate of transmission through dwellings at Alert Level 4, and believe the observed increase in spread is due to the large amount of spread pre-intervention that means that a significant number of dwellings already contained an infected person when AL4 was applied.

During Alert Level 4, we see the amount of spread through the remaining essential workplaces and community events increasing through time, as more people within these contexts get infected.

In the two scenarios where we shift to Alert Level 3 on September 15th, we see the largest increase in the infections through the community, through extended bubbles, reduced compliance, and customers returning to businesses. This increase is largest in the pessimistic Alert Level 3 scenario.

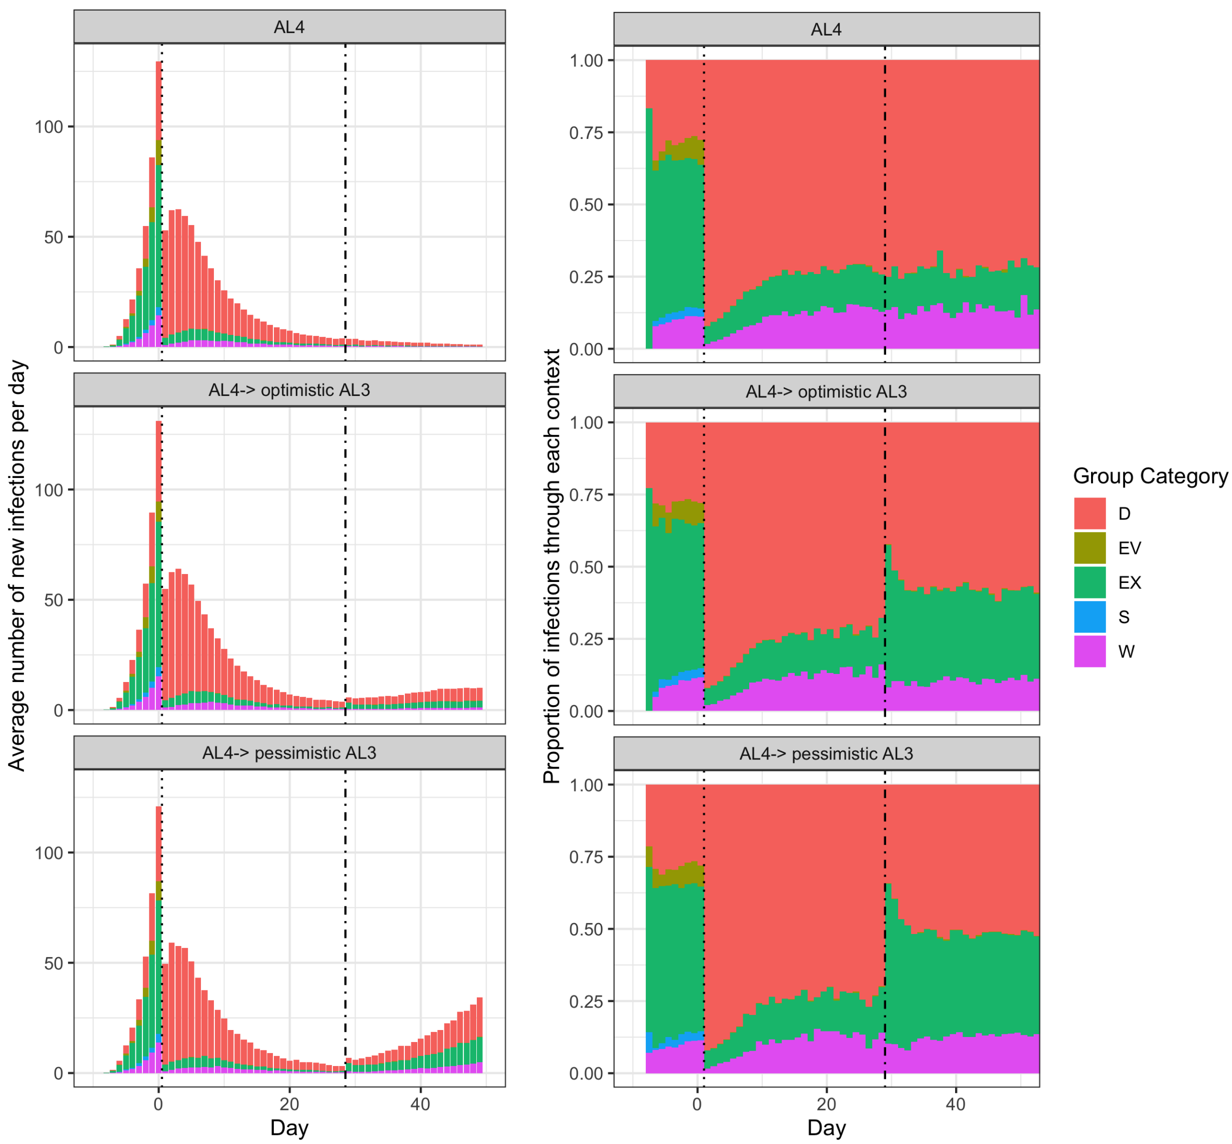

Figure 4. For the network contagion simulation results retained in Figure 2, we show the average number of daily infections (time of exposure) and the context of the spread through time (Days post detection). Colours indicate the categories D=Dwelling, W=Workplace, S=School, EX=Close community contact, EV=Casual community contact. The shift to AL4 on the 18th of August is shown with a dotted line, and the potential shift to AL3 on 15th September is shown with a dot-dashed line.

The context of infections in our model is driven by a number of key parameters which split up the relative infectiousness through different contexts, as described in the Appendix, as well as being strongly influenced by the number of connections people have in the interaction network. We have used data from ARPHS contact tracing for this outbreak to help calibrate the context of infection pre-detection and in Alert Level 4. For Alert Level 3, while the magnitude of changes in infection context is uncertain, we can confidently expect an increase in community (EX and EV) connections. This is of particular concern since these connections drive further growth in infections in other contexts by allowing infected cases to transmit to new parts of the network.

3.2 Branching Process Model

3.2.1 Conditioning using a Gaussian Process

Since the date of a possible change in Alert Level intervention, simulated in this report is several days ahead of the current date, we begin by estimating the number of cases we expect to see on 16th September, when the simulated Alert Level change occurs. **We predict a median of 8 [LQ=5, UQ=12] daily detected cases by the date of Alert Level change.** (See Table 2.) At the point of deescalating the Alert Level intervention, the estimated value of R_{eff} for the ensemble of retained simulations has a median of 0.56 [LQ=0.50, UQ=0.63]. The distribution of R_{eff} values is shown in Figure 5.

Date	25% Quartile	Median	75% Quartile
September 16 2021	5.0	8.0	12.0

Table 2. Projected detected case numbers on September 16 2021

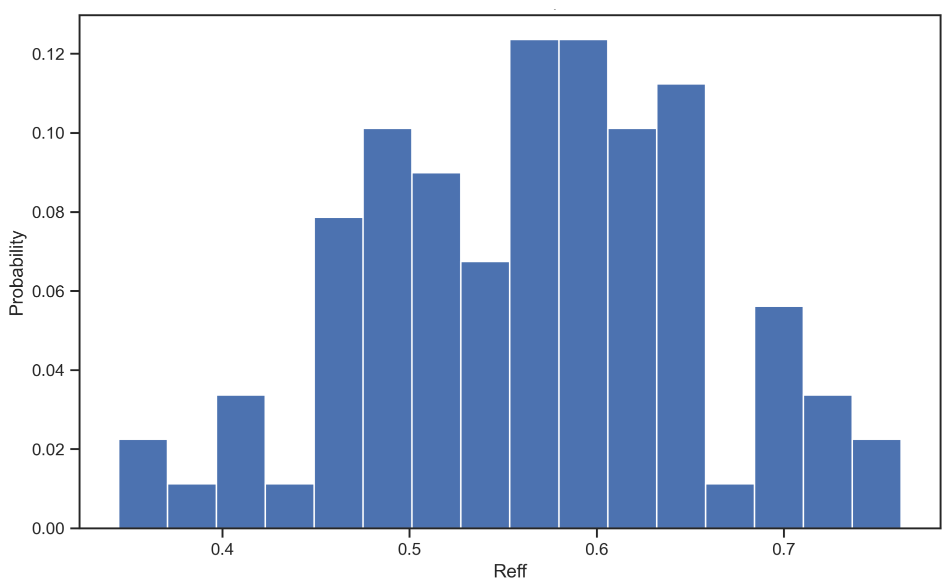

Figure 5. Conditioned distribution of R_{eff} prior to de-escalation. The median value of R_{eff} at de-escalation is 0.56.

Table 3 reports the expected number of daily new detected cases at times 3, 7, 10, and 14 days after a de-escalation of Alert Levels. Since the expected value of R_{eff} is unknown during this period, we present scenarios where the value of Reff increases by a factor of 1.5, 2, and 3, due to the Alert Level change. We also present estimates for the expected number of daily detected cases if there is no change in Alert Level. Trajectories of daily detected cases corresponding to the same scenarios are also shown in Figure 6.

These results suggest that daily case numbers will continue to decrease towards elimination if we stay at Alert Level 4. This decrease in daily case numbers does not continue for all but the most optimistic of Alert Level 3 scenarios. We would expect it to take up to 10-14 days after the alert level change for a scenario with a high increase in R_{eff} became apparent.

Figure 7 shows the distribution of estimated R_{eff} values in the period after Alert Levels are reduced. These suggest that our low increase scenario (optimistic AL3; R_{eff} 1.5 times greater than AL4) has an R_{eff} median of 0.84; our moderate increase scenario (R_{eff} 2 times greater than AL4) has a median of 1.1; and our high increase scenario (pessimistic AL3; R_{eff} 3 times greater than AL4) has a median of 1.7. These estimates, along with the lower and upper quartiles are summarised in Table 4.

R_{eff} multiplier	Data Quartiles	September 19 2021	September 23 2021	September 26 2021	September 30 2021
1.0	25% Quartile	-2.0	-3.0	-4.0	-5.0
	Median	-2.0	-4.0	-5.0	-6.0
	75% Quartile	-4.0	-8.0	-9.0	-10.0
1.5	25% Quartile	-2.0	-4.0	-4.0	-5.0
	Median	-3.0	-5.0	-6.0	-7.0
	75% Quartile	-3.0	-5.0	-4.0	-5.0
2.0	25% Quartile	-2.0	-4.0	-3.0	-4.0
	Median	-3.0	-4.0	-3.0	-4.0
	75% Quartile	-3.0	-2.0	0.0	1.0
3.0	25% Quartile	-2.0	-3.0	-3.0	-2.0
	Median	-2.0	-2.0	-1.0	2.0
	75% Quartile	-5.0	-4.0	-1.0	5.0

Table 3. Change in projected daily cases, relative to projected cases for the relevant quartile on September 16 2021 (see Table 2), for 3, 7, 10 and 14 days after a change in intervention level corresponding to an increase of R_{eff} by a factor of 1.0, 1.5, 2.0, and 3.0. If values in this table are positive (negative), it indicates that daily case numbers at the corresponding date are higher (lower) than they were at the time of Alert Level de-escalation.

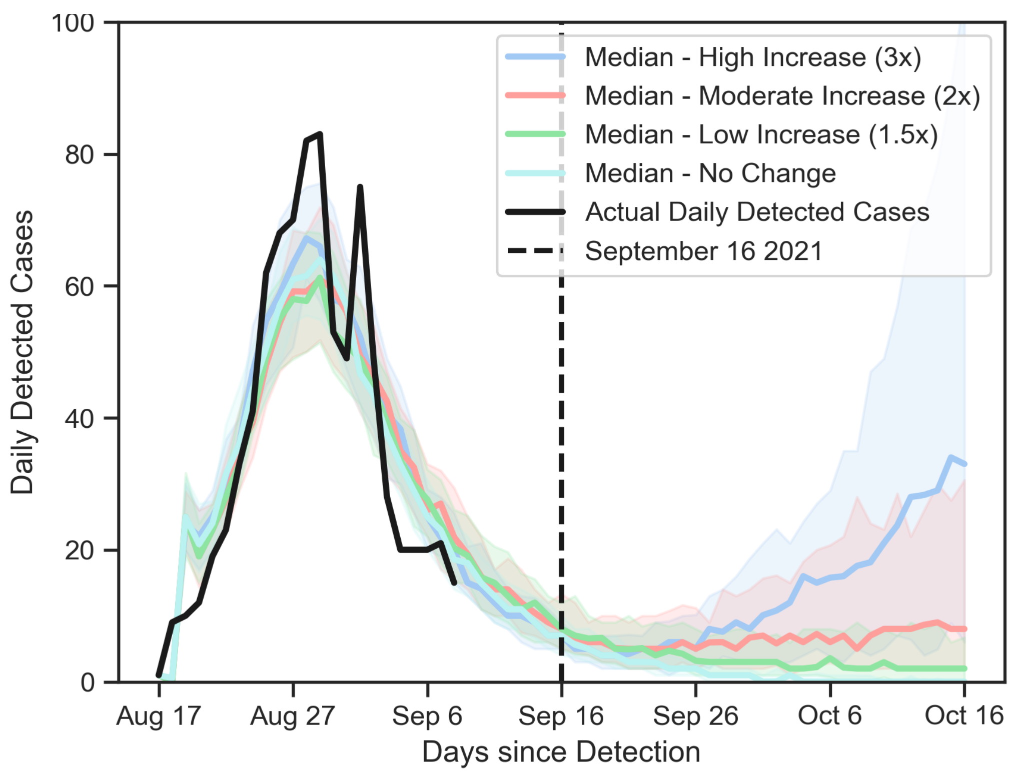

Figure 6. Daily Case Numbers After Detection, Changing Alert Level Effect on R_{eff}

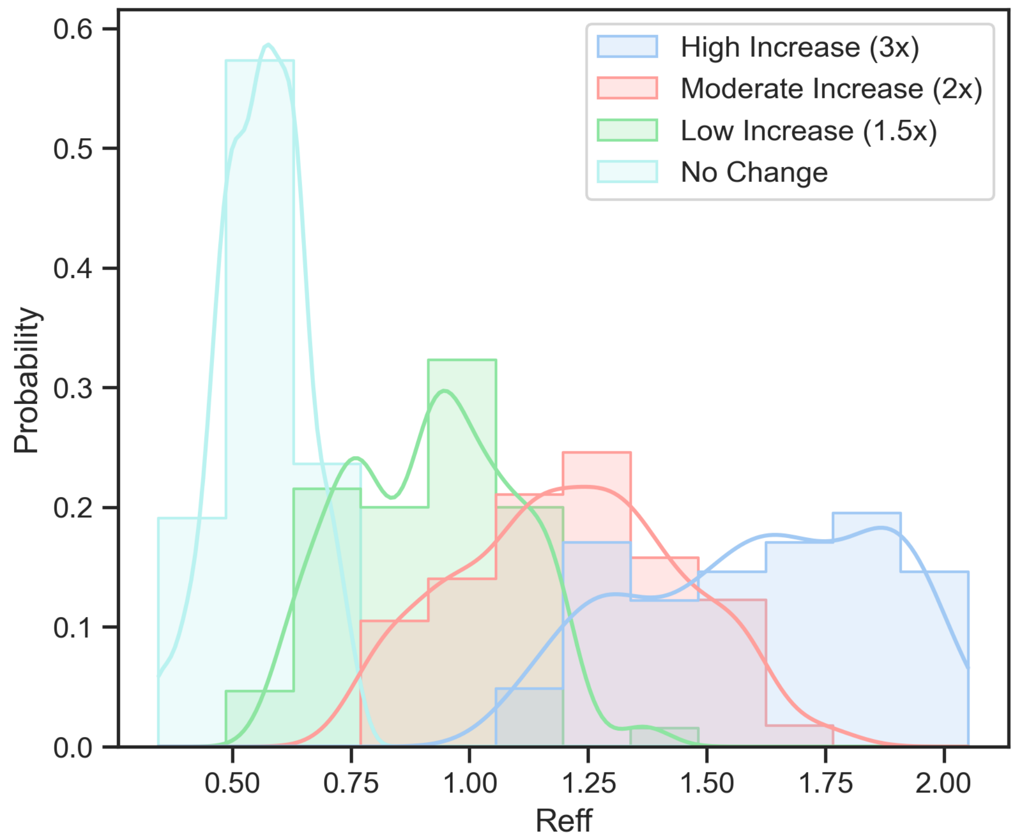

Figure 7. Distributions of R_{eff} with changing Alert Level Effect

Data Quartiles	No Change	Low Increase (1.5x)	Moderate Increase (2x)	High Increase (3x)
25% Quartile	0.50	0.75	1.0	1.5
Median	0.56	0.84	1.1	1.7
75% Quartile	0.63	0.95	1.3	1.9

Table 4. Median, 25% and 75% Quantiles for R_{eff} with changing Alert Level Effect

Since estimates of R_{eff} at 16th September indicate that staying at AL4 would put us on track towards eliminating the outbreak, it is worth calculating estimates of the first day to reach zero daily cases.[†] These estimates are shown in Table 5. Scenarios where there is no no zero case day before the end of the simulated period (16th October) are indicated as such, however it is important to note that this does not distinguish between scenarios where daily case number are increasing at AL3 and those where a zero case day may be expected to occur at a date outside the simulation range.

These results suggest that remaining at Alert Level 4 will lead to a zero case day around the beginning of October, but that a zero case day is unlikely in the near term at lower levels of intervention.

Data Quartiles	No Change	Low Increase (1.5x)	Moderate Increase (2x)	High Increase (3x)
25% Quartile	September 28 2021	October 8 2021	After October 16 2021	After October 16 2021
Median	October 2 2021	After October 16 2021	After October 16 2021	After October 16 2021
75% Quartile	October 11 2021	After October 16 2021	After October 16 2021	After October 16 2021

Table 5. Date to reach zero cases after a change in intervention level corresponding to an increase of R_{eff} by a factor of 1.0, 1.5, 2.0, and 3.0.

[†]Note - stochasticity and spread from any remaining non-quarantined cases means that this is not necessarily the same as the date at which we expect no further cases.

3.2.2 Conditioning using ABC (added 15 September)

Since the date of a possible change in Alert Level intervention, simulated in this report is several days ahead of the current date, we begin by estimating the number of cases we expect to see on 16th September, when the simulated Alert Level change occurs. **We predict a median of 6 [LQ=2, UQ=13] daily detected cases by the date of Alert Level change.** (See Table 6.) At the point of deescalating the Alert Level intervention, the estimated value of R_{eff} for the ensemble of retained simulations when using ABC has a median of 0.52 [LQ=0.32, UQ=0.72] (See Table 8). The distribution of R_{eff} values is shown in Figure 8.

Date	25% Quartile	Median	75% Quartile
September 16 2021	2.0	6.0	13.0

Table 6. Projected detected case numbers on September 16 2021

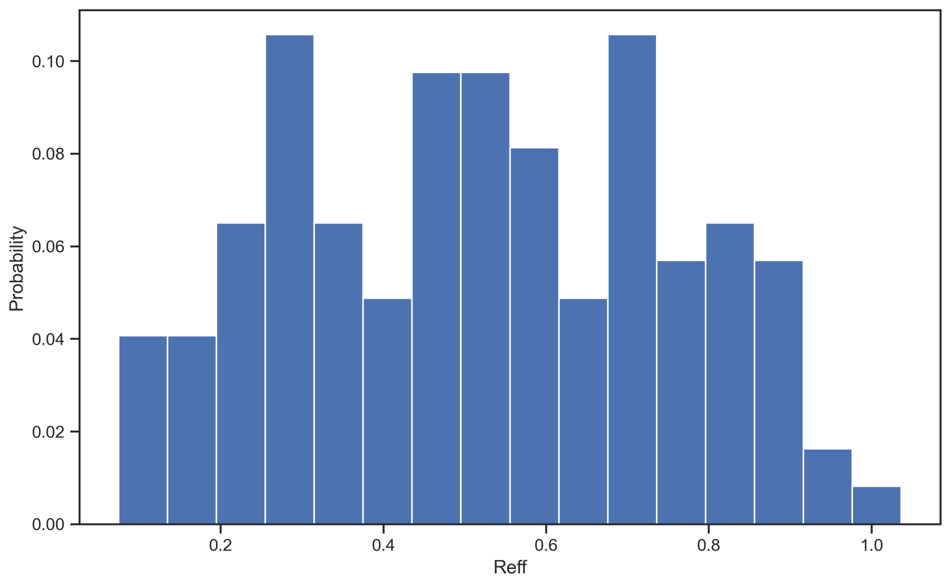

Figure 8. Conditioned distribution of R_{eff} prior to de-escalation. The median value of R_{eff} at de-escalation is 0.52.

Table 7 reports the expected number of daily new detected cases at times 3, 7, 10, and 14 days after a de-escalation of Alert Levels. Since the expected value of R_{eff} is unknown during this period, we present scenarios where the value of R_{eff} increases by a factor of 1.5, 2, and 3, due to the Alert Level change. We also present estimates for the expected number of daily detected cases if there is no change in Alert Level. Trajectories of daily detected cases corresponding to the same scenarios are also shown in Figure 9.

These results suggest that daily case numbers will continue to decrease towards elimination if we stay at Alert Level 4. This decrease in daily case numbers does not continue for all but the most optimistic of Alert Level 3 scenarios. We would expect it to take up to 10-14 days after the alert level change for a scenario with a high increase in R_{eff} became apparent.

Figure 10 shows the distribution of estimated R_{eff} values in the period after Alert Levels are reduced. These suggest that our low increase scenario (optimistic AL3; R_{eff} 1.5 times greater than AL4) has an R_{eff} median of 0.78; our moderate increase scenario (R_{eff} 2 times greater than AL4) has a median of 1.04; and our high increase scenario (pessimistic AL3; R_{eff} 3 times greater than AL4) has a median of 1.56. These estimates, along with the lower and upper quartiles are summarised in Table 8.

R_{eff} multiplier	Data Quartiles	September 19 2021	September 23 2021	September 26 2021	September 30 2021
1.0	25% Quartile	-2.0	-3.0	-4.0	-5.0
	Median	-1.0	-3.0	-4.0	-4.0
	75% Quartile	-4.0	-7.0	-9.0	-10.0
1.5	25% Quartile	0.0	-1.0	-1.0	-1.0
	Median	0.0	-3.0	-3.0	-4.0
	75% Quartile	-2.0	-3.0	-6.0	-6.0
2.0	25% Quartile	-1.0	-2.0	-2.0	-2.0
	Median	-2.0	-4.0	-5.0	-4.0
	75% Quartile	-4.0	-5.0	-5.0	-3.0
3.0	25% Quartile	-1.0	-1.0	-1.0	-2.0
	Median	-2.0	-2.0	0.0	2.0
	75% Quartile	-3.0	1.0	6.0	18.0

Table 7. Change in projected daily cases, relative to projected cases for the relevant quartile on September 16 2021 (see Table 6), for 3, 7, 10 and 14 days after a change in intervention level corresponding to an increase of R_{eff} by a factor of 1.0, 1.5, 2.0, and 3.0. If values in this table are positive (negative), it indicates that daily case numbers at the corresponding date are higher (lower) than they were at the time of Alert Level de-escalation.

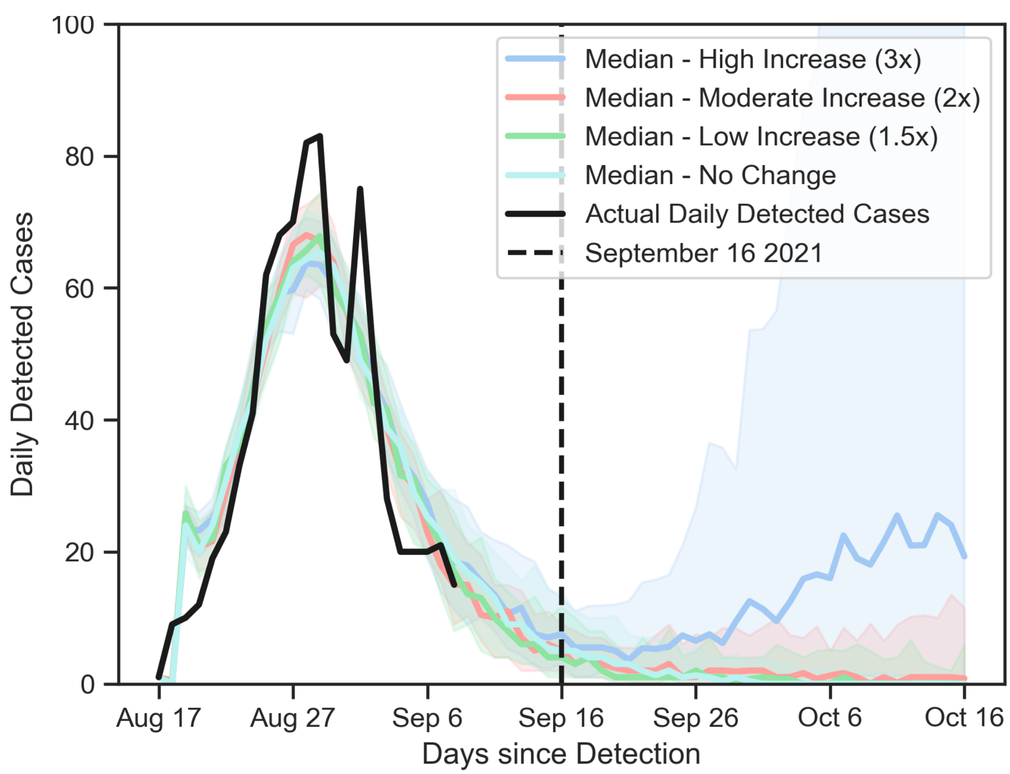

Figure 9. Daily Case Numbers After Detection, Changing Alert Level Effect on R_{eff}

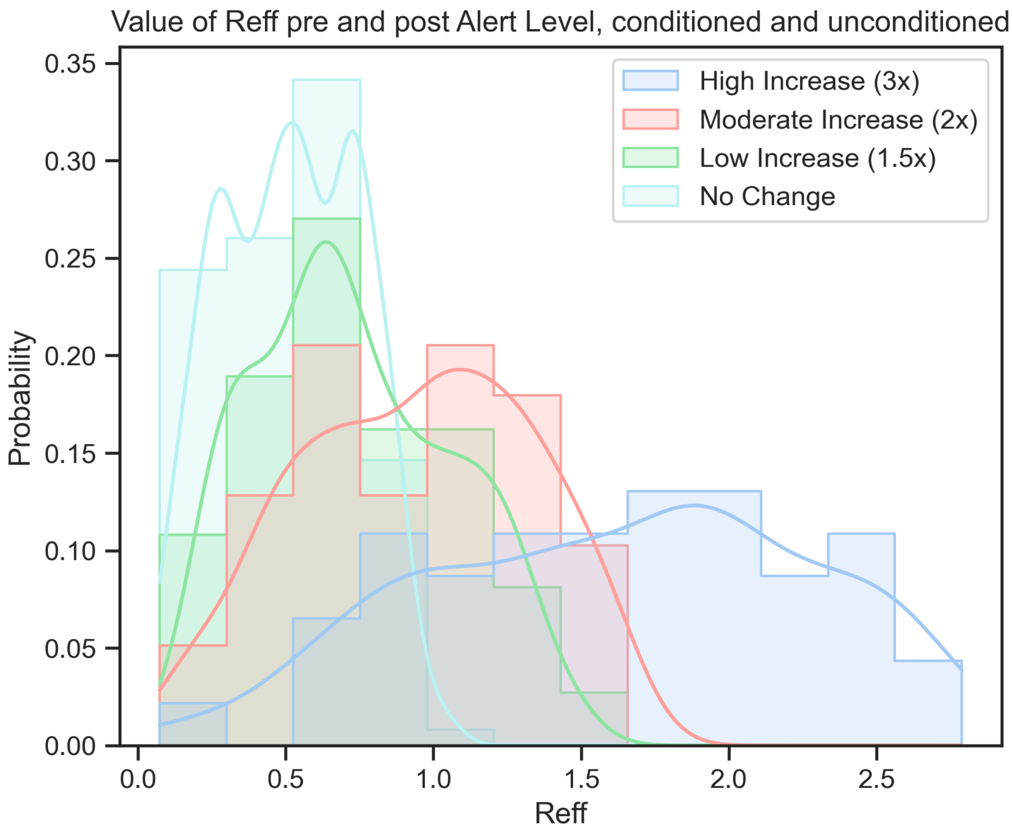

Figure 10. Distributions of R_{eff} with changing Alert Level Effect

Data Quartiles	No Change	Low Increase (1.5x)	Moderate Increase (2x)	High Increase (3x)
25% Quartile	0.32	0.48	0.64	0.96
Median	0.52	0.78	1.04	1.56
75% Quartile	0.72	1.08	1.44	2.16

Table 8. Median, 25% and 75% Quantiles for R_{eff} with changing Alert Level Effect

Since estimates of R_{eff} at 16th September indicate that staying at AL4 would put us on track towards eliminating the outbreak, it is worth calculating estimates of the first day to reach zero daily cases.[‡] These estimates are shown in Table 9. Scenarios where there is no no zero case day before the end of the simulated period (16th October) are indicated as such, however it is important to note that this does not distinguish between scenarios where daily case number are increasing at AL3 and those where a zero case day may be expected to occur at a date outside the simulation range.

These results suggest that remaining at Alert Level 4 will lead to a zero case day around the beginning of October, but that a zero case day is unlikely in the near term at lower levels of intervention.

[‡]Note - stochasticity and spread from any remaining non-quarantined cases means that this is not necessarily the same as the date at which we expect no further cases.

Data Quartiles	No Change	Low Increase (1.5x)	Moderate Increase (2x)	High Increase (3x)
25% Quartile	September 22 2021	September 21 2021	September 22 2021	September 30 2021
Median	September 29 2021	September 30 2021	October 10 2021	After October 16 2021
75% Quartile	After October 16 2021	After October 16 2021	After October 16 2021	After October 16 2021

Table 9. Date to reach zero cases after a change in intervention level corresponding to an increase of R_{eff} by a factor of 1.0, 1.5, 2.0, and 3.0.

References

1. Marin, J.-M., Pudlo, P., Robert, C. P. & Ryder, R. J. Approximate bayesian computational methods. *Stat. Comput.* **22**, 1167–1180 (2012).

2. Harvey, E., Maclaren, O. J., O'Neale, D. & Wu, D. Asymptomatic cases of COVID-19 in NZ: a study of data from the Auckland August re-emergence case data. Tech. Rep., Te Pūnaha Matatini (2020).

3. Byambasuren, O. *et al.* Estimating the extent of asymptomatic COVID-19 and its potential for community transmission: systematic review and meta-analysis. *medRxiv* DOI: 10.1101/2020.05.10.20097543 (2020). https://www.medrxiv.org/content/early/2020/09/13/2020.05.10.20097543.full.pdf.

4. Statistics NZ. Household labour force survey: June 2020 quarter. Retrieved from: https://www.stats.govt.nz/assets/Uploads/Labour-market-statistics/Labour-market-statistics-June-2020-quarter/Download-data/household-labour-force-survey-june-2020-quarter-supplementary-tables.xlsx (2020). Accessed: 9 Dec 2020.

5. Kearns, N. *et al.* How big is your bubble? characteristics of self-isolating household units ('bubbles') during the COVID-19 Alert Level 4 period in New Zealand: a cross-sectional survey. *BMJ Open* **11** (2021).

6. Analysis of contact survey data from aotearoa during 2020 (2021). Internal Te Pūnaha Matatini analysis.

7. Liu, Y. & Rocklöv, J. The reproductive number of the Delta variant of SARS-CoV-2 is far higher compared to the ancestral SARS-CoV-2 virus. *J. Travel. Medicine* DOI: 10.1093/jtm/taab124 (2021).

8. Binny, R. N. *et al.* Effective reproduction number for COVID-19 in Aotearoa New Zealand. Tech. Rep., Te Pūnaha Matatini (2020). DOI: 10.1101/2020.08.10.20172320. Company: Cold Spring Harbor Laboratory Press Distributor: Cold Spring Harbor Laboratory Press Label: Cold Spring Harbor Laboratory Press Type: article.

9. James, A., Plank, M. J., Binny, R. N., Hendy, S. C. & Lustig, A. Summary of Advice to Cabinet on Auckland's August 2020 COVID-19 Outbreak (2020).

10. Hendy, S. *et al.* Mathematical modelling to inform New Zealand's COVID-19 response. *J. Royal Soc. New Zealand* **51**, DOI: 10.1080/03036758.2021.1876111 (2021).

11. OpenABM-Covid19 model. https://github.com/BDI-pathogens/OpenABM-Covid19. Accessed: 2020-08-03.

12. Prem, K., Cook, A. R. & Jit, M. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLoS computational biology* **13**, e1005697 (2017).

13. QriousNZ. Qrious inter-regional ground travel data. https://www.transport.govt.nz/mot-resources/research-papers/inter-regional-ground-travel-data-from-qrious/ (2019). Accessed: 2020-08-10.

14. Steyn, N. *et al.* Māori and pacific people in new zealand have higher risk of hospitalisation for covid-19. *medRxiv* 2020–12 (2021).

Appendix

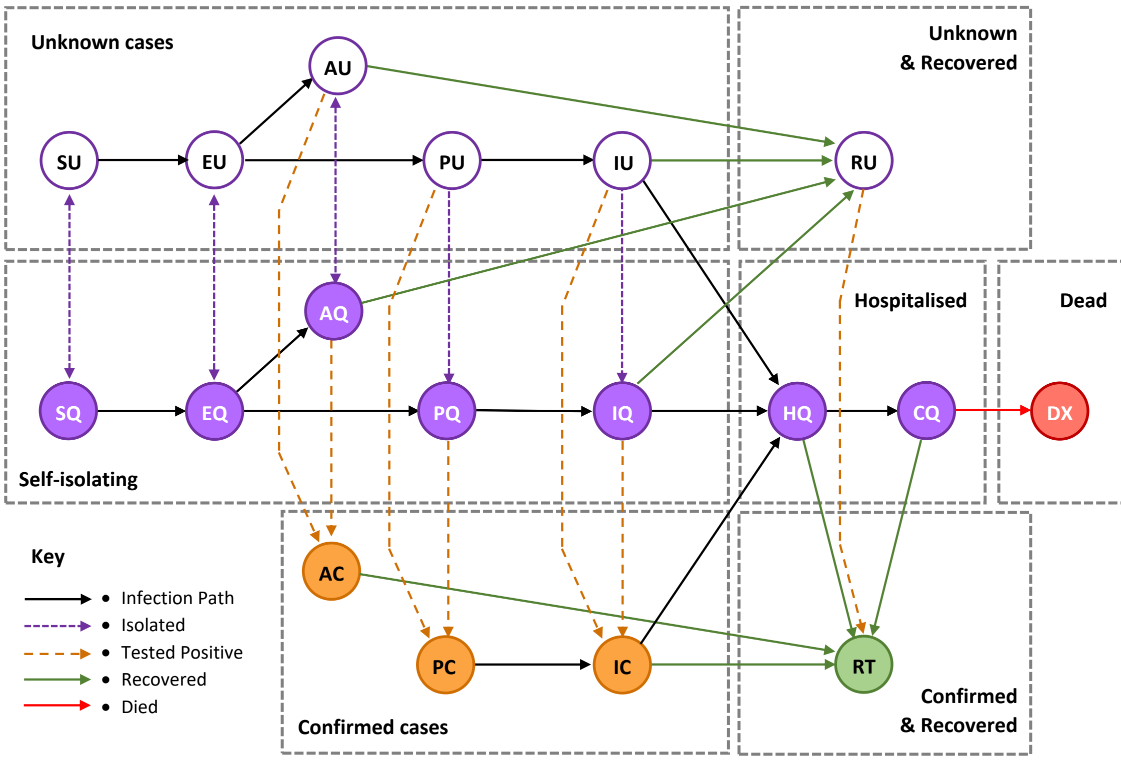

Figure 11. State and Transition Diagram for Network Contagion Model. We note that there is no transition from IU→DX, IC→DX, or HQ→DX due to the assumption that all life-threatening cases would be hospitalised first, and all would require critical care before death.

Transmission parameters

Transmissibility We model infection as a transmission process from infected individuals to susceptible individuals that is mediated through the groups, or interaction contexts, that they share. The transmission is modelled as a Markovian process, which occurs with an exponentially distributed interval between successive infection events. The overall infectiousness 'budget' that people have is represented by the parameter β. This $\beta = \beta_c + \beta_0$, where β_c, the rate of risky contacts/events occurring for 'close contacts' is 10 times β_0, the rate of risky contacts/events occurring for 'casual contacts'. Based on simulations with uncontrolled spread, and early estimates of the growth rate in the initial stages of the August 2021 outbreak, we use $\beta_c = 3.55$, $\beta_0 = 0.355$, which produces an overall $\beta = \beta_c + \beta_0 = 3.905$.

We conceptualise transmission budget as being a split up based on the 'time' that an infectious individual spends in a context with susceptible individuals, and the 'riskiness of contact'. We use parameters to split up this transmission budget between three categories of mixing context: home, work or school, and the community. When individuals have more than one mixing context within a category, we assume that individuals spend an even amount of time in each mixing context in each of the categories. For example, if an individual may spend 4 'hours' in community mixing contexts, and is linked to 4 different community mixing contexts, then we assume they spend 1 hour in each mixing context. We make a well-mixed assumption within each mixing context, where individuals spend equal amounts of time with each other individual in that mixing context. The parameters used to split the transmission budgets β_c and β_0 are key input parameters to the model, and determine the relative spread through different interaction contexts. For close contact events, we split infections up between (Dwellings):(Work or School):(Community) according to the ratios $2.5 : 0.75 : 4.5$. For casual contact events, we split infections up between (Dwellings):(Work or School):(Community) according to the ratios $1 : 2 : 1.33$.

Infection events are mediated by the attributes of the individuals involved, as well as the group that the infection event happens in. When a transmission occurs, it has a chance of being rejected, or failing, which is a function of the individuals and group involved. This is to model heterogeneities between individuals. Currently, transmission is automatically denied if both individuals are self-isolating and do not share a dwelling; down-weighted (to represent rule-breaching) if either is in a

'Confirmed' state, or is in 'Self-isolating' state and the individuals do not share a dwelling. The likelihood of transmission is also down-weighted by the age and vaccination status of the susceptible individual, as well as the current Alert Level policy's effect on transmission in the group that event is occurring in. Infection events are also rejected in groups that are 'closed', due to Alert Level policies.

Finally, infection events involve one infectious individual, but one or more susceptibles. The number of susceptibles that are infected in a single event is drawn from a log-normal distribution based on the size of the group. This currently is only applied to casual community mixing contexts.

Infectiousness of different states The states Pre-symptomatic, Symptomatic, and Asymptomatic are all infectious. Due to the state transition model being Markovian, we need to use a constant infectiousness rate for each state, rather than allowing it to vary through time. Using generation times from uncontrolled spread, and comparing to estimates in the literature, we find that setting pre-symptomatic cases and asymptomatic cases to have 70% the infectiousness of a symptomatic individual helps produce reasonable model outputs.

Susceptibility We assume age-varying susceptibility to infection, which works by increasing the chance of 'rejecting' an infection event, depending on the age of the infectee. For this we use parameters from the OpenABM model[11] for susceptibility by 10 year age band. We then use the population structure by ethnicity to aggregate these values up to the four age bands used in our network. This means that the values differ for individuals in the same age band, depending on their ethnicity.

Interaction network parameters We build the interaction network for dwellings, schools, and workplaces, from linked microdata in the StatsNZ Integrated Data Infrastructure.

In small dwellings, small schools, and small workplaces, we assume that all contacts are both 'close' and 'casual' contacts. If groups are larger than some threshold, we create smaller groups within the large group which represent the smaller number of close contacts such as a class within a school or a team within a workplace. Note: these close contact group sizes are drawn from a distribution, so some groups will have more and some fewer than the mean sizes listed below, and people can be in more than one close contact group.

For large dwellings (over size 12), the close contact groups have a mean of size 10 and 'overcover parameter' of 0.2 which means that people living in dwellings over size 12 have on average 12 close contacts. For large workplaces (over size 10), the close contact groups have a mean of size 10 and 'overcover parameter' of 0.1 which means that people in workplaces over size 10 have on average 11 close contacts. For large schools (over size 40, so almost all schools except ECE centres), the close contact groups have a mean of size 30 and 'overcover parameter' of 1 which means that people in schools over size 40 have on average 60 close contacts.

For all other close contacts ('community contexts') we use a Poisson distribution for the number of 'events' per person, with a minimum of one per person. The mean number of events per person is selected to produce a mean number of contacts per person that is double the daily contact counts in 'other' contexts from Prem et al.[12] ([10.5, 13.2, 8, 6.6] for the four age bands [0-14, 15-29, 30-59, 60+]). We select these close contact event sizes using a Power-law distribution with a minimum of 2, mean of 3.2, and maximum size of 100. We use age-weighted group allocation to produce age-contact matrices in 'other' contexts that match Prem et al.[12]. Overall this results in a mean of 9.6 close contacts per person, with median=5, and LQ=2, UQ=10.

For all other casual contacts ('community contexts') we again use a Poisson distribution for the number of 'events' per person, with a minimum of one per person and a mean of 3 per person. The casual contact event sizes are drawn from a Power-law distribution, with a mean size of 5.4. We enforce a minimum size of 2, and a maximum of 1000. Once we project this, we get a mean of almost 200 casual community contacts per person, but it is very heavy-tailed (median=50, LQ=14, UQ=207). We assume no age-structure to the casual contacts.

The majority of community interactions are well-mixed within a single Territorial Authority, but a small proportion of people are also linked to community events in other Territorial Authorities, with the density of these long-range links between different Territorial Authorities being based on cellphone movement data from 2017[13].

Case progression parameters

Latent period All infected individuals are initially in an Exposed state, where they are not infectious (and will not test positive), for an average of 2 days[11] (latent period). At that point infections will either move to a pre-symptomatic or asymptomatic state.

Asymptomatic proportion Asymptomatic cases 'recover'[§] in an average of 10 days[11]. We use NZ case data from the August Outbreak, fit with a logistic regression model depending only on age, to estimate the proportion of infections that will be asymptomatic. Despite the underlying proportion depending only on age, the proportion in our model vary by ethnicity once we aggregate it up to our four age bands due to the different age structure for different ethnic groups. At a population level, if

[§] 'Recover' only means that they are no longer infectious, and won't test positive if tested. It does not mean they are symptom-free or even that they are discharged from hospital.

we assumed all people had an equal infection risk, this comes out to ∼18%, but with a strong dependence on age (∼34% for 0–14 year olds down to ∼6% for 60 years and older).

Pre-symptomatic and symptomatic states Pre-symptomatic cases will be infectious and test positive if tested. Symptom onset will occur an average of 3.5 days[11] after progression from Exposed to Pre-symptomatic. This gives an overall incubation period of 5.5 days. Once symptomatic, cases will either 'recover'† or be hospitalised. Those that 'recover' will take an average of 10 days[11]. Note: this means that symptomatic cases are infectious for 13.5 days, compared to 10 days for asymptomatic cases. So even though we assume the same infectiousness per unit time, we would expect fewer infections from asymptomatic cases.

Pre-symptomatic and symptomatic states Following Steyn *et al.* (2021)[14], we use a logistic regression model fit to the same NZ case data but for symptomatic cases only. This gives us an estimate of the proportion of symptomatic infections that will be hospitalised by age and ethnicity. In order to account for the higher disease severity with the Delta variant, we double the risk of hospitalisation for symptomatic individuals. We then use the population counts by age and ethnicity to aggregate these hospitalisation risks by year of age up to the four age bands in our network.

Hospitalisation, critical care, and death Cases that end up in hospital move from Symptomatic to Hospitalised states in an average of 5 days[11] after symptom onset, and will then either 'recover'† or need critical care. For those that 'recover'*, it will take an average of 6 days[11], those that will need critical care will move into critical care in an average of 2 days after hospital admission[11]. We use age-based estimates for the proportion of hospitalised cases that will need critical care from[11]. We do not vary this by ethnicity or for the Delta variant. At a population level, if we assumed all people had an equal infection risk, then the overall infection hospitalisation rate for the delta variant would come out to ∼12%, but with a strong dependence on age and ethnicity.

Cases that need critical care will either 'recover' or die. We use age-based estimates for the proportion of critical cases that will die from[11]. For those that 'recover'† from critical care, it will take an average of 16 days, those that die in critical care will die in an average of 9 days[11]. At a population level, if we assumed all people had an equal infection risk, then the overall infection fatality rate for the delta variant would come out to ∼1.5%, but with a strong dependence on age and ethnicity.

Note on timings As our disease progression model is Markovian, the above average times represent the means of exponentially-distributed random durations in the corresponding states. This assumption can be relaxed by introducing delay or semi-Markov processes but is not done here.

Vaccination parameters

- We currently assume that vaccine effects are uniform over the population, and have no age-, ethnicity- or location-dependent effects. This is only a first approximation as there is e.g. documented age-dependency in vaccine efficacy (especially a significant efficacy drop off in older individuals).
- We assume that the vaccine is effective against both infection and onwards transmission given infection, and that these are affected differently. The reduction in susceptibility to infection is assumed to be 70% while the typical reduction in onward transmission, given infection, is assumed to be 50%.
- We assume that infected, vaccinated individuals are 25% less likely to be symptomatic.
- We assume that the vaccine reduces the likelihood of severe outcomes: we assume a 50% reduced chance of hospitalisation, critical care, and death of infected, symptomatic individuals.
- We assume that symptomatic, infectious, vaccinated individuals are 50% less likely to seek a test, compared to similar unvaccinated individuals.
- We assume that the duration of infectivity of vaccinated individuals is identical to unvaccinated individuals.
- We assume that interventions, testing policies, and isolation policies are applied identically to vaccinated and unvaccinated individuals.

Test/trace/isolate parameters

Symptomatic testing We have two different parameters to represent community testing. The first is the likelihood (probability) that someone with symptoms (but no known contact with a case) would seek a test, and test positive. The second is the time from symptom onset to the test result being returned. Due to the number of seeds used (20) and the known time to detection in the current outbreak, we set the likelihood of seeking a test to: 2.5%, the test positivity as 80%, which gives an overall 2% of symptomatic cases that will be detected. The time from symptom onset to return of test follows an exponential distribution with a mean of 5 days. After detection of a case, we increase the likelihood of seeking a test as listed in Section 2.1. We also decrease the time from symptom onset to the return of test to a mean of 4 days (following an exponential distribution).

Casual contact testing Casual contacts are not explicitly contact traced. We assume that a proportion of casual contacts would know they were contacts, seek a test, and test positive. This proportion is set in Section 2.1. This test-seeking is regardless of symptoms. We do not assume that casual contacts would isolate until a test result is returned. The delay between the confirmed case notification and those casual contacts testing positive is modelled as a scaled Beta distribution. ***We do not assume that people being tested because of symptoms or being casual contacts would isolate until a test result is returned.***

Close contact contact tracing and self-isolation Time from positive test to contact isolation follows a Weibull distribution (scale = 1.5, shape=1.5). These parameters match the target for P002 of 80% contacted within 48hrs. Only close contacts, as defined in our network structure, are contact traced. 'Known (traceable) proportions' of close contacts are set to in the ensemble parameters (see Section 2.1). Close contacts in dwellings are contacted first, then those in workplaces,schools,and community.As well as priority in contact order, there is a different probability for a tracing attempt to be unsuccessful, by interaction context. This is 0% for dwellings (i.e. all contacts from the same dwelling are traceable), 10% for workplaces and schools, and 10% for community contacts. We set the maximum number of trace attempts at 5 for any given contact, and set no capacity limits for contact tracing.

We assume that identified close contacts will isolate at home, so can still transmit to dwelling members as the same rate as baseline. ***We do not assume that their household members will be told to isolate.*** We also allow for a 1% 'leak' rate. Identified close contacts will isolate for 14 days from notification, and will have multiple tests and symptom follow up that will result in: 95% of symptomatic cases testing positive and 90% of asymptomatic cases. Close contacts who are infected will test positive over time following an exponential distribution with a mean of 2 days after isolation.

Confirmed cases We assume that all confirmed active cases (AC, PC, IC) will be moved to MIQ or equivalent, so there will be no transmission to any of their contacts, including those in their dwelling, but we allow for a 1% 'leak' rate. Hospitalised cases (HQ and CQ) are assumed to be confirmed on admission, if not already confirmed, but hospitals are assumed to be perfect isolation (no transmission).

Historical cases We define historical cases to be recovered individuals that were not confirmed during their course of infection, corresponding to the RU state. We assume that historical cases will not seek tests spontaneously, but will test positive if they are tested. Historical cases that are casual contacts of a confirmed individual will test, for example. Historical cases that are close contacts of a confirmed individual will be assumed to test similarly to other self-isolated individuals, if they are notified by contact tracing. Historical cases that are identified through testing will induce contact tracing processes as if they were infected, but will not be required to isolate themselves (we assume that they have no infectivity, so their isolation does not particularly matter in our model). They will also be categorised as "known" as they move to a RT state.