

---

# SOCIAL INTERVENTIONS COVID-19 ABM

## DOCUMENTATION

---

This agent-based model was developed for NHS trusts in North East England by the University of Durham COVID-19 Community Health and Social Care Modelling Team. For comments or questions, please contact [Jennifer Badham](#).

### **Contents: User Guide and Technical Reference**

1	Quick Start: Installing and Running the Model.....	1
2	Epidemic Process .....	3
3	Interventions.....	5
4	Example Scenarios .....	8
5	Advanced Settings.....	11
6	Model Design and Mechanisms.....	14
7	Calibration.....	16
8	Predefined Scenarios .....	19
9	Model History .....	19

# USER GUIDE

In the absence of a vaccine or effective treatment, the only way to manage a communicable disease epidemic is to interrupt transmission from infectious people to susceptible people. This involves controlling the ways in which people mix (such as isolating those potentially exposed) and reducing the potential for transmission where they do mix (such as promoting good hand hygiene). This model is intended to help understand the potential impact of combinations of these non-pharmaceutical interventions over time, and the uncertainties associated with estimates of the impact.

Two processes are represented in the model: transmission and disease progression. These interact through an extended person to person SEIR epidemic model (see section 2). The disease spreads directly from infectious people to susceptible people, excluding indirect real-world paths such as virus survival on hard surfaces. The interventions (see section 3) act on the transmission process, for example constraining the mixing between infectious and susceptible people. Once the disease has been successfully transmitted to a susceptible person, that person progresses through different epidemic states, which may include hospitalisation.

## 1 QUICK START: INSTALLING AND RUNNING THE MODEL

The model is built in NetLogo, specialist simulation software. You will need to install NetLogo, version 6.1 or later, which is available as a free download from <https://ccl.northwestern.edu/netlogo/>. Open the model file named *Social COVID-19.nlogo* in NetLogo to access the main user interface (shown in Figure 1).

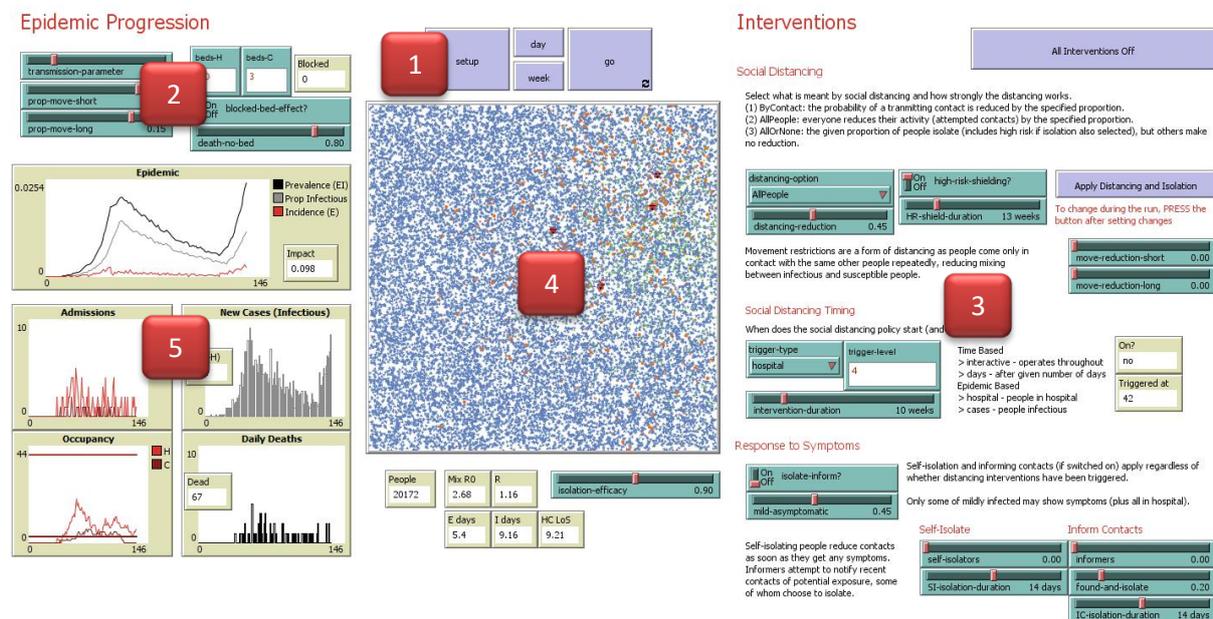
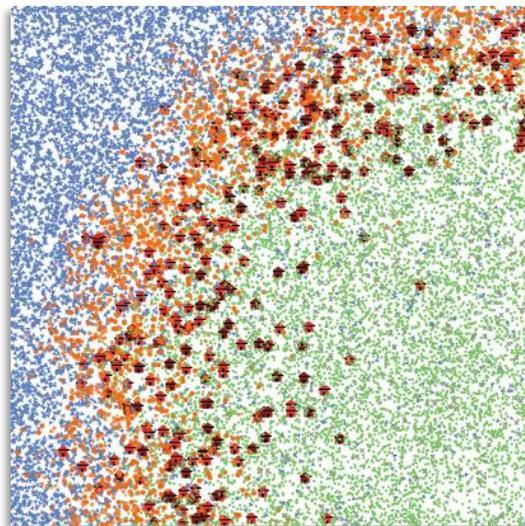


Figure 1: User interface. Area 1 has the buttons to start and stop the simulation. Area 2 controls the disease progression process. Area 3 controls the interventions in the scenario. Areas 4 and 5 are the display areas, showing simulated people and summary plots respectively.

The buttons to control the simulation are at the top centre of the interface (labelled 1 at Figure 1). The **setup** button is used to initialise the simulation, creating people throughout the world (labelled 4 at Figure 1) and setting one random person to be infected. After initialisation, the **day** and **week** buttons advance the simulation by one and seven timesteps (representing days) respectively. The **go** button runs the simulation until it either ends because there are no infectious people remaining, or the user presses the **go** button again.

Start the simulation by pressing the **setup** button and then the **go** button. If you have made no changes to the controls since downloading the model, for most runs the epidemic will start slowly but then accelerate. Part way through the simulation, the main display will be similar to the example simulation at Figure 2 (which is the section of the screen labelled as 4 in Figure 1).

In some runs, the random process that transmits the disease does not lead to enough early infections and the epidemic will not occur. If the simulation ends before the epidemic spreads, make sure there are no interventions reducing transmission (press the **All Interventions Off** button at top right of the interface) and restart the simulation with the **setup** and **go** buttons.



*Figure 2: Example run spatial display at tick 100 (with no interventions). Blue points represent susceptible people, orange and red points represent exposed and infectious people in the community, red and dark red houses represent people in hospital (general and critical care), green points represent people who have recovered from the epidemic and are now immune, and a grey X represents a person who has died.*

This main display demonstrates the spatial nature of epidemic spread. That is, the new cases spread out from the initial centres of infection over time, and the epidemic occurs at different times from the perspective of a particular location. However, the spatial structure of the model does not represent any real-world location and there is no variation in population density.

The left side of the interface (labelled 5 in Figure 1) contains plots to summarise the progress of the epidemic.

- The **Epidemic** plot summarises the active aspect of the epidemic, displaying the number of simulated people who are currently infected (prevalence, including those exposed but not yet infectious) and newly infected (incidence), both as a proportion of the population. The **Impact** box reports the proportion of the population who have ever been infected, providing a cumulative summary of the epidemic.

- **Admissions** and **Occupancy** refer to the number of new or current people in hospital respectively, separated between those requiring general (H) or critical (C) care. The horizontal lines in the occupancy plot show the hospital capacity.
- The **New Cases** and **Daily Deaths** plots display the number of people, by day, who change state to the infectious or dead status respectively. The position of these directly above each other and on the same scale makes it easier to see the delayed effect of interventions. Newly exposed is not displayed, so there is also some delay before the response will be visible in the **New Cases** plot, which better represents the information available to those making policy decisions.

Interventions to control the epidemic are set from the right side of the interface (labelled 3 in Figure 1). A simple scenario would implement social distancing measures, triggered by the epidemic reaching some level. The intervention explored in section 4 represents social distancing as applied initially in the United Kingdom. To implement it: (1) Use the **trigger-type** dropdown box to select **hospital** and set **trigger-value** to 4; and (2) Use the **distancing-option** dropdown box to select **ByContact** and set **distancing-reduction** to 0.7 with the slider. Run the scenario (with the **setup** and **run** buttons). If all other parameters are set to their defaults and no other interventions are set, this will instruct the model to introduce simulated social interventions at an approximately equivalent point in the epidemic as they were introduced in the United Kingdom, and the epidemic growth will slow to just below replacement value. That is, the new cases will flatten but not substantially reduce.

The **New Cases** and **Epidemic** plots for a typical run are at Figure 6. The epidemic clearly slows from day 25 when the interventions are introduced and is increasing after 10 weeks, when social distancing ends (default value for the **intervention-duration** slider) and the simulated people return to their pre-intervention activity levels. Note that the simulation involves random processes and different runs will have different outcomes, including the day on which the intervention is introduced.

## 2 EPIDEMIC PROCESS

---

For the disease progression, people start in the susceptible (S) state, change to the exposed state (E) if they come in adequate contact with an infectious person, become infectious (I) automatically several days after they are exposed. The infectious group includes those who are not yet showing symptoms (pre-symptomatic) or will never show symptoms (asymptomatic). Once infectious, they can recover (R – become immune) at home, die at home (R – dead), or their infection becomes severe and they are moved to hospital. From hospital, they may recover or die directly from general wards. However, if their condition worsens, they may instead move to critical care, from which they either die or recover. People in hospital remain infectious, but are treated within the model as if they are isolated and therefore have limited capacity to infect others. The potential paths are displayed at Figure 3.

The transmission from infected to susceptible people is implemented through spatial proximity. The model starts with twelve people (turtles in NetLogo terminology) on each patch (NetLogo terminology for a cell in the spatial grid). Each simulated day, any infectious person interacts with all the susceptible people on their patch and, with independent random draws, has the opportunity to transmit the infection to each of them. There is no spatial scale, where a unit of distance has some real-world equivalent number of kilometres. Such abstraction is reasonable for small regions where people mix relatively freely. It also reflects the priority of this model as a tool for policy experimentation, by making the model fast and relatively simple to use.

Some people also move each tick to facilitate mixing. Short movements are one unit of distance in a random direction, which is likely but not always moving them to a different patch (and therefore mix with different people). Long movements are three units of distance. There is no concept of regular

activities such as commuting that have a person moving between two fixed locations that represent home and work. This movement allows the epidemic to spread spatially, and also introduces variation in the number of potential contacts each simulated day.

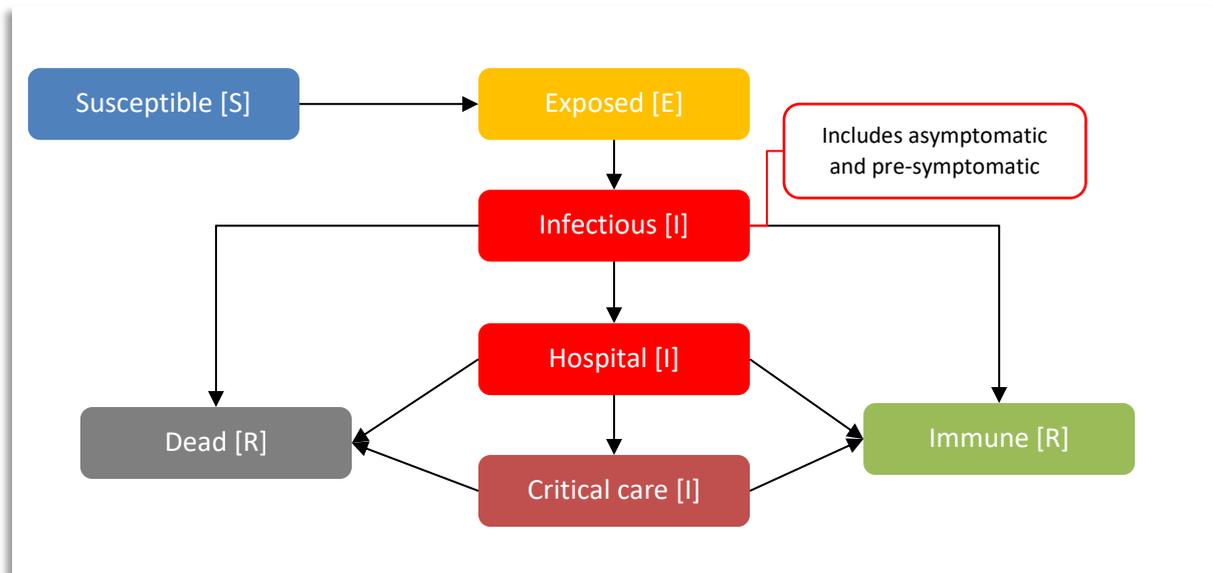


Figure 3: Epidemic state transition paths in the extended SEIR model.

The epidemic is initialised by randomly selecting one person to start in the exposed state. All others are susceptible. The epidemic progresses through a combination of the changes in epidemic states for exposed individuals (see advanced settings at section 5.1 for controls) and transmission from infectious to susceptible people. Interventions modify only the transmission process. Treatments and other interventions that can modify state transitions are not included in the model.

## 2.1 TRANSMISSION CONTROLS

Three sliders are used to control how fast the epidemic spreads (labelled 2 at Figure 1). Most directly, the probability that an infectious person will infect a susceptible person on the same patch during the current time step (day) is given by the **transmission-parameter**, excluding any modifications from the interventions. Increasing this parameter value increases the intensity of the epidemic directly, because a higher proportion of contacts between infectious and susceptible generate a new infection.

The other two sliders control the proportion of people who move a distance of one unit (**prop-move-short**) or three units (**prop-move-long**) each time step. The epidemic will spread spatially more quickly with higher values. Generally, the epidemic will also be more intense because greater movement will allow infectious people to access new susceptible people if they have successfully transmitted to the susceptible people in their current patch.

## 2.2 EFFECT OF HOSPITAL CAPACITY

The epidemic state transition probabilities are controlled through advanced settings (see section 5.1) to reflect the best available real-world values; for example, for the probability of a person who is infectious requiring hospitalisation. By default, these transitions ignore constraints in hospital capacity. However, this approach can be overridden so that a person who needs hospitalisation has a relatively high probability of dying if a bed is not available.

Relevant controls are in the top left section of the interface (labelled 2 at Figure 1). Set the **beds-H** and **beds-C** to the number of general and critical care beds respectively, appropriately scaled for a population of 20,172. Set the **blocked-bed-effect** to 'On', and the **death-no-bed** slider to the desired probability of death if a bed is not available. If there are no available beds (combining general and critical care) when the simulated person attempts to be admitted to hospital, the person remains in the community until they either die (with specified probability) or recover.

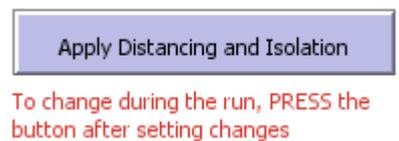
## 3 INTERVENTIONS

The model includes two broad types of interventions to manage the epidemic by disrupting mixing between infectious and susceptible people. Social distancing measures reduce the amount of contact between people. The other measures focus on how people respond to symptoms, limiting the contacts of those known or suspected to be infectious. Both are implemented with controls on the right side of the interface (labelled 3 at Figure 1).

### 3.1 SOCIAL DISTANCING OPTIONS

There are three social distancing policy options in the model: general restriction of activities, voluntary isolation by those with higher risk (sometimes referred to as shielding or shelter-in-place), and restricting movement. These can be applied individually or in combination. To allow government policies to be represented in a reproducible way, the timing of these options can also be controlled by the user so that they are triggered when a set number of cases have arisen, and then maintained for a specific period.

If you adjust the settings during a simulation run, you must also press the **Apply Distancing and Isolation** button to transfer the settings on the interface to the simulated people and ensure all aspects of the different policy are applied.



#### 3.1.1 Timing of social distancing options

The controls displayed in Figure 4 set the timing of the social distancing policies. There are three options for when the policies start, set using a combination of **trigger-type** and **trigger-level**. If the **trigger-type** is set to **days**, the policies will commence after that many simulated days, regardless of the progress of the epidemic. In contrast, the **hospital** and **cases** options apply the social distancing policies from the time at which the specified (cumulative) number of infectious people or hospital admissions has been reached. The **trigger-level** box is used to enter the number of days, cases or hospital admissions as appropriate.

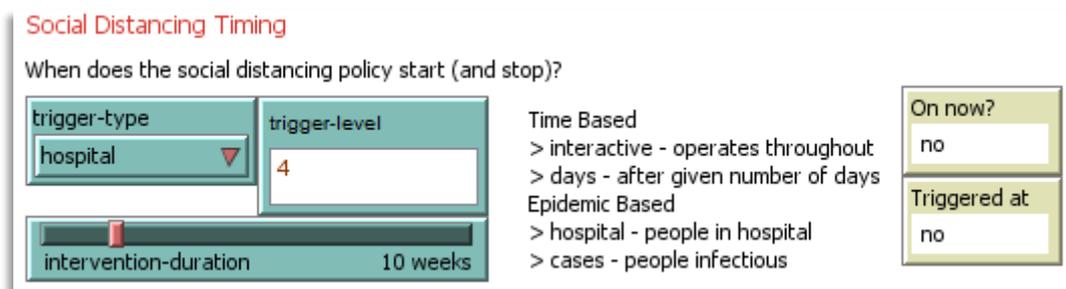


Figure 4: Social distancing policy timing controls.

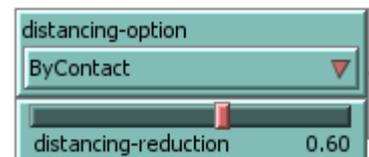
The distancing and movement policies will operate for the length of simulated time controlled by the **intervention-duration** slider and then end. The **Triggered at** monitor reports the day at which the conditions are met and the distancing measures are applied, and the **On now?** monitor reports whether they are currently applied.

The voluntary isolation of high-risk people has a slightly different timing structure. The policy is applied at the same time as other social distancing policies, but has a separate control to allow extended shielding. The duration is set by the **HR-isolation-duration** slider.

There is also an option in the **trigger-type** for **interactive** control. With this option, the policies are constantly applied and can be interactively adjusted, for example starting with a large reduction in social contact and then gradually reducing the strength of the social distancing policy. This is useful for exploring the effect of policies and understanding the model, but it is difficult to compare scenarios where settings are adjusted.

### 3.1.2 Restricting activities and contacts

The **distancing-option** chooser (drop down box) is used to select the way in which social distancing is implemented in the model. The different options represent different assumptions about distancing policies, such as whether they affect all people or just some. The **distancing-reduction** slider controls how much reduction to apply; that is, the strength of the intervention.

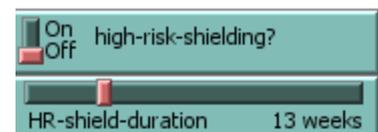


The **ByContact** option reduces the probability of a transmissible contact. For example, if the slider is set to 0.6, then the model does a random draw so that any contact that would have led to transmission without the intervention only has a 40% chance of leading to transmission while the distancing option is operating.

Conceptually, the other options reduce the probability that a person attempts to make a contact. The **AllPeople** scenario reduces every person's activity by the same amount. In contrast, the **AllOrNone** scenario conceptually reduces the activity of some people to nothing (representing isolation) and has no effect on the activity of all other people. If voluntary isolation of high risk people is also operating, they will be included in the people isolating from the **AllOrNone** social distancing policy. Note that, if two people independently reduce their activity by some amount, then the actual number of contacts will reduce by more than that amount (because each contact depends on both people being present). Therefore a specified social distancing reduction will have a larger impact for these two options than for the **ByContact** option.

### 3.1.3 Voluntary isolation by those with higher risk

Those at higher risk can additionally choose to isolate. This option is set by the **high-risk-shielding?** switch and starts with the other social distancing policies. Isolation has both a stronger reduction in activity (set by the **isolation-efficacy** slider) and separate duration (set by the **HR-shield-duration** slider) than the general restriction on activities and contacts.



The actual proportion of the population who are high risk and therefore isolate is controlled by the **prop-high-risk** slider in the advanced settings (lower part of interface, see section 5.3). If the **AllOrNone** option is operating (see section 3.1.2), the high risk people are included in the isolating population, not additional to the isolating population.

### 3.1.4 Restricting movement

The **Movement** scenario changes only the proportion of people moving (short and long). While this does not directly alter the probability of transmission or the number of contacts, it typically reduces the effect of those contacts. This is because people in the infectious state come into contact with the same other people repeatedly, rather than new people who have not yet been exposed to the epidemic. Consequently, infectious people have limited access to susceptible people.



## 3.2 RESPONSES TO SYMPTOMS

This section describes the model settings concerning how people respond to showing symptoms, potentially choosing to self-isolate and inform their recent contacts so those people can isolate themselves. Unlike the social distancing options already described, these options may operate throughout the simulation because they represent individual actions taken by a simulated person rather than restrictions imposed by government. The separate timing is intended to support interactive exploration with multiple changes in implementation as the epidemic situation changes.

Both decisions are controlled by the **isolate-inform?** switch, operating when it is 'On' and not operating when it is 'Off'. In addition, the **mild-asymptomatic** slider is used to control the proportion of the infectious people in the community (not requiring hospital care) that do not develop symptoms and therefore do not take any action under these policy options.



Those people displaying symptoms will do so some number of days after entering the infectious epidemic state. The timing of symptoms is set with a truncated Poisson distribution controlled in the advanced settings (lower part of interface, see section 5.1).

### 3.2.1 Self-Isolate

Under this option, once a simulated person becomes symptomatic, they have some probability of choosing to isolate, with that probability set by the **self-isolators** slider. Note that this should be set to 0 if only the inform contacts option is required as the switch controls both actions. They will isolate for the period set by the **SI-isolation-duration** slider, with the effect of reducing their activity by the amount set by the **isolation-efficacy** slider (which applies to all forms of isolation in the model).

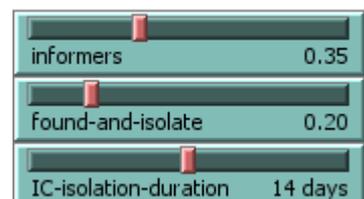
#### Self-Isolate



### 3.2.2 Inform Contacts

Under this option, once a simulated person becomes symptomatic, they have some probability of choosing to inform their recent contacts of their status, with that probability set by the **informers** slider. Note that this should be set to 0 if only the self-isolate option is required as the switch controls both actions.

#### Inform Contacts



If the person decides to inform their contacts, then a further probability is assessed for each contact that combines whether the person is able to identify and notify the contact and whether the contact chooses to isolate. That probability is applied at the time the symptoms emerge for those already exposed, and at the time of exposure for future transmissions.

The contact choosing to isolate will do so for the period set by the **IC-isolation-duration** slider, with the effect of reducing their activity by the amount set by the **isolation-efficacy** slider (which applies to all forms of isolation in the model).

## 4 EXAMPLE SCENARIOS

The example scenarios start with a simple one-off social distancing policy and gradually become more complex with a combination of scenarios changing over time. Together, these demonstrate the most important features of the model that can be accessed through the interface. Interactive modelling facilitates insight, with an experimental learning approach to understand the potential consequences of policy decisions.

As the model incorporates random processes, each run is different, and formal scenario comparison requires multiple simulations of each scenario. Experiments with multiple simulations of pre-defined scenarios are managed with BehaviorSpace (see section 8). Experiments that require changes to model settings during the simulation must be individually programmed and are outside the scope of this manual.

### 4.1 SCENARIO 1: SIMPLE SOCIAL DISTANCING

The United Kingdom introduced several social distancing measures around 23 March 2020, at which time 11,086 cases had been confirmed in England at a rate of 19.8 per 100,000 population.<sup>1</sup> Scaling to the 20,172 simulated people in the model, this is equivalent to 3.99 confirmed cases. At that time, testing only occurred on hospital admission, so a suitable trigger for the social distancing intervention to apply in the model is 4 (cumulative) admissions to hospital. For this scenario, the lockdown will apply for nine weeks with a 70% reduction in activity and then revert to normal (on the equivalent of 25 May 2020).

To construct the scenario, first press the **All Interventions Off** button (top right of the interface). Then set the controllers as shown in Figure 5. The left side of the figure displays the timing controls, with the lockdown starting once the epidemic has led to four hospital admissions and continuing for nine weeks. The right side shows the lockdown effect, with the activity of all people reduced by 0.5. You may wish to confirm that other policies are not operating; with **high-risk-shielding?** set to 'Off', **move-reduction-short** and **move-reduction-long** both set to 0, and **isolate-inform?** switched to 'Off'.

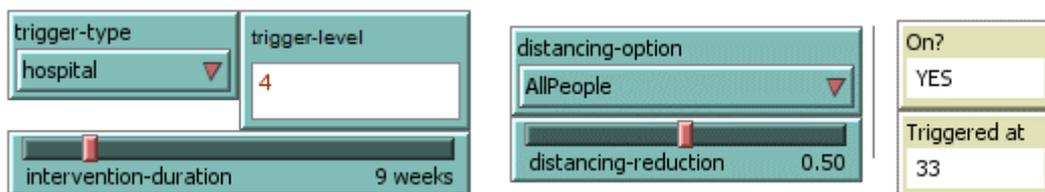


Figure 5: Settings for scenario 1, together with the monitors showing the day on which the social distancing policy starts.

Once the scenario is defined, press the **setup** button (area 1 in Figure 1) to initialise the simulation. This will create simulated people and randomly select one to be exposed to the epidemic. Press the **go** button to start the simulation. The **go** button will remain black (indicating that it is pressed)

<sup>1</sup> UK Government, [Coronavirus \(COVID-19\) in the UK](#), daily cases data available as download.

automatically until the epidemic is over. To pause the simulation, click the **go** button to ‘unpress’ it, then click it again to continue. If the simulation stops before the epidemic takes hold in the population, restart the simulation by pressing the **setup** then **go** buttons again.

The epidemic will spread throughout the world, with each simulated person displayed with a different colour and shape depending on their current epidemic status (see Figure 2). There are also various plots and numbers that summarise the model outputs. The **Epidemic** plot (shown in Figure 6) provides the proportion of the population that are in the exposed or infectious (including hospital) states over time, as well as the newly exposed. For this scenario, the infections are increasing rapidly at the start of the simulation, and then continue to increase but at a much reduced rate once the social distancing starts. The return to normal contact levels after nine weeks leads to a rapid increase in infections.

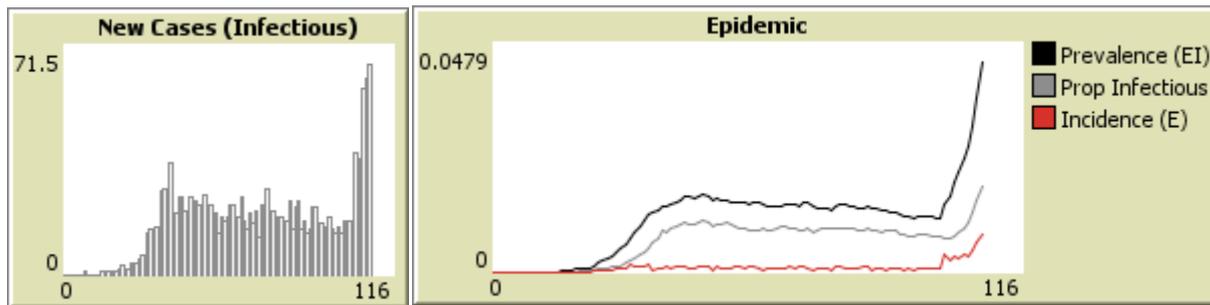


Figure 6: New cases and epidemic plots from the early phase of a typical run with a simple social distancing intervention.

Separate plots also show counts (rather than proportion) of people in different states, such as the newly infectious (Figure 6), daily deaths, new hospital admissions and hospital occupancy (with hospital capacity indicated).

## 4.2 EXAMPLE 2: SEQUENTIAL SOCIAL DISTANCING WITH EXTENDED SHIELDING

More realistically, formal social distancing measures will be gradually eased, and people will at least partly maintain practices such as good hand hygiene, working from home and respecting distance. The second example scenario varies the social distancing policy interactively after the initial lockdown is complete. It also introduces self-isolation for those at higher risk.

Set up the scenario as for the first scenario, with a 0.5 reduction in activity for nine weeks triggered by four hospital admissions. In addition, set the **high-risk-shielding?** switch to ‘On’ and the **HR-shield-duration** slider to 17 weeks (see section 3.1.3 for an image of the appropriate controls). Start the simulation run with the **setup** and **go** buttons as before.

This time, at the end of the lockdown (when the epidemic starts to rise again), pause the model by pressing the **go** button again. Use the dropdown box to change the **trigger-type** to interactive, and then press the **Apply Distancing and Isolation** button. The button is necessary to set the model’s internal variables to new values and is used whenever you change any distancing policy setting.

The purpose of changing to the interactive trigger is that social distancing will continue, regardless of the **intervention-duration** slider (though the high-risk shielding will end on schedule). Press the **go** button again to continue the simulation. Adjust the **distancing-reduction** slider to manage the epidemic (and pressing the **Apply Distancing and Isolation** button), increasing the reduction as infections rise and reducing it again when they are under control. Instead of using the **go** button to

run it indefinitely, you can also use the **week** or **day** button to advance it for seven and one time steps respectively.

### 4.3 EXAMPLE 3: SELF-ISOLATION AND INFORMING CONTACTS

This scenario applies the options concerning how people respond to symptoms. These reflect personal decisions that are conceptually somewhat different than official social distancing policies where governments close certain public spaces or restrict activities. They are started and stopped by a switch rather than the more formal timing of the social distancing options.

To construct the scenario, first press the **All Interventions Off** button to reset all policies to have no effect. In particular, **distancing-option** is set to None so that the slider has no effect. Then set the controllers as shown in Figure 7. With these settings, most of the infectious people will show symptoms and therefore will realise they need to self-isolate and inform their contacts. For those who do show symptoms, the proportions isolating and encouraging others to isolate are set fairly high so that the effect of the policy will be visible in the epidemic outcomes. The scenario starts with the **isolate-inform?** switch set to 'Off' so that the epidemic has a chance to become established in the population.

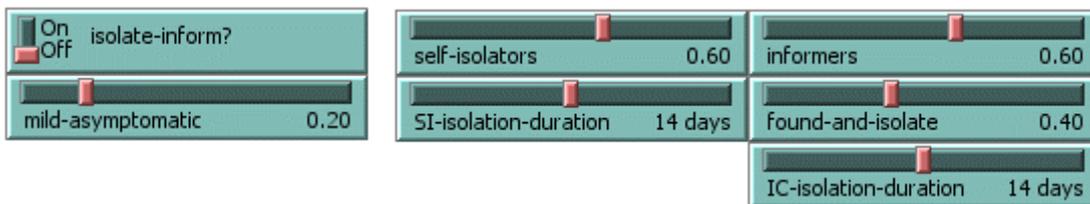


Figure 7: Settings for scenario 3. Note that the **isolate-inform?** switch is set to 'Off'.

Start the simulation running with the **setup** then **go** buttons. Once prevalence is clearly increasing, drag the **inform-isolate?** switch to the 'On' position. The epidemic should slow (or at least increase less dramatically). Switching between 'On' and 'Off' should be clearly visible, as shown in Figure 8.

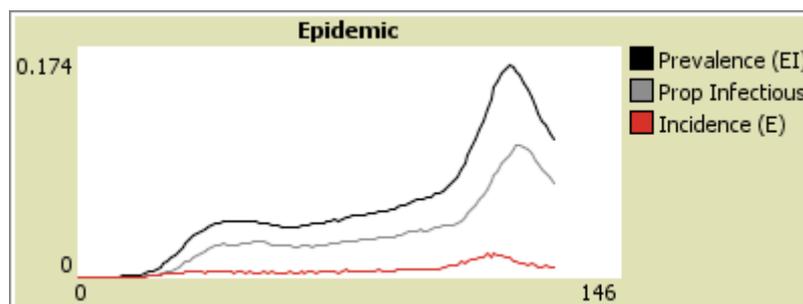


Figure 8: Epidemic plot from part of a run with intermittent isolation and inform responses to becoming symptomatic.

### 4.4 EXAMPLE 4: COMBINING INTERVENTIONS

The final scenario combines elements of the previous scenarios. It is intended to represent a lockdown with strong community spirit leading to self-isolating and informing of contacts as the lockdown is eased. That is, a nine week lockdown is accompanied by 13 weeks of shielding for those at higher risk.

The isolate and inform approach starts during the four weeks of easing while shielding continues, and then continues indefinitely.

This scenario is initialised with the settings shown in Figure 5 and Figure 7, together with setting the *high-risk-shielding?* to On and *HR-shield-duration* to 13 weeks. However, set the *mild-asymptomatic* proportion to 0.5 (rather than 0.2 in the original scenario). Start the simulation (**setup** then **go** buttons). Look at the **Triggered at** monitor to see when the social distancing starts. Add 63 (representing 9 weeks) to that number and pause the simulation before that day is reached. Note that the current day is visible in the menu ribbon at the top of the interface. You can pause early, and then use the **week** and **day** buttons to advance as required.

Once the appropriate day is reached, change the *distancing-option* to interactive (and press the **Apply Distancing and Isolation** button), *distancing-reduction* slider to 0.3, and *isolate-inform?* switch to 'On'. Continue the simulation for another four weeks and pause again. Move the *distancing-reduction* slider to 0.1 but make no other changes and then let the simulation run to the end.

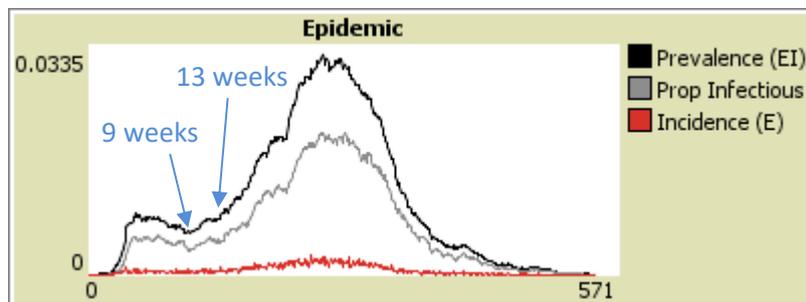


Figure 9: Epidemic plot for scenario with a combination of epidemic control measures. The scenario settings change at 9 weeks and 13 weeks after social distancing is implemented.

Figure 9 shows the epidemic plot for one example of this scenario. The response to symptoms actions are sufficient to control the epidemic as the lockdown eases to 30% activity levels but fail once activity reduction reaches 10%. Nevertheless, the overall epidemic is substantially reduced compared to scenario 1 (or the unmitigated epidemic).

## 5 ADVANCED SETTINGS

Scrolling below the main interface accesses the area for advanced settings. Most of these settings concern the epidemic state progression (see Figure 3), providing the necessary probabilities and durations. In general, these should only be altered as more information emerges about disease progression.

Parameter values can be altered on the interface for a particular run and, if the model is saved, the controls will be set with those revised values when the model is reopened. Pressing the **Use Defaults** button will retrieve the values originally included with the model. Changing the defaults to a new permanent value requires changes to the NetLogo code (see section 6.4).

### 5.1 TRANSITION PROBABILITIES

Following the (potentially very short) infectious state where a person has a mild case and remains in the community, there are three possibilities for the next state: recovery and immunity, death, or

admission to hospital with severe symptoms. The probabilities for the latter two are controlled by the **prob-InfDeath** and **prob-InfHosp** sliders respectively, with the probability of recovery at home calculated as the remainder.

If a person is admitted to hospital, there are three potential paths: recovery and immunity, death, or transfer to critical care. The probabilities for the latter two are controlled by **prob-HospDeath** and **prob-HospCrit** sliders respectively, with the probability of recovery calculated as the remainder. From critical care, a person either dies in hospital (**prob-CritDeath** slider) or is discharged when they recover (calculated as remainder). The model does not allow a person to transfer to the general hospital ward as they improve.

The model also allows for failure of immunity where a person's epidemic status is returned to "susceptible" instead of "immune". This possibility is controlled by the **immune-mild** and **immune-severe** sliders that set the proportion of people who will become immune at the point they recover from their infection in the community or in hospital respectively.

The default values for these transition probabilities are derived from the literature about the epidemiology of COVID-19. The parameter values and sources are summarised at Table 1. Most of these parameters required calculation to obtain the required form for the model.

Table 1: Default values for transition probabilities

Parameter	Value	Source
prob-InfDeath	0.025	ONS death statistics, <sup>2</sup> 1/1.55 of COVID-19 deaths occur in hospital, combined with the value for prob-InfHosp
prob-InfHosp	0.085	ONS Infection survey <sup>3</sup> for 5 Jun 2020 with (lagged) pillar 1 tests <sup>4</sup> and assume 9 days detectable
prob-HospDeath	0.3	ISARIC study <sup>5</sup>
prob-HospCrit	0.17	ISARIC study <sup>5</sup>
prob-CritDeath	0.43	ICNARC report <sup>6</sup> for 29 May 2020
immune-mild	1	Assumed immune when recover
immune-severe	1	Assumed immune when recover

## 5.2 AGE MIXING

In the model, all simulated people interact equally with all other people on the same patch, with the same probability of infectious transmission for any contact. In practice, however, there are strong age structured mixing patterns. As people generally mix with others of their own age (plus cross generation within households), transmission occurs more easily within an age group than between groups.

<sup>2</sup> Office of National Statistics, *Deaths registered weekly in England and Wales*, provisional: [Week ending 29 May 2020](#).

<sup>3</sup> Office of National Statistics, *Coronavirus (COVID-19) Infection Survey pilot*: [5 June 2020](#).

<sup>4</sup> UK.gov, *Number of coronavirus (COVID-19) cases and risk in the UK*. [Testing statistics](#).

<sup>5</sup> Docherty AB et al (2020), 'Features of 20 133 UK patients in hospital with covid-19 using the ISARIC WHO Clinical Characterisation Protocol: prospective observational cohort study', *British Medical Journal*, 369:m1985. DOI: 10.1136/bmj.m1985.

<sup>6</sup> Intensive Care National Audit & Research Centre, *ICNARC report on COVID-19 in critical care*: [29 May 2020](#).

The model emulates age structured mixing by adjusting the transmission probability so that it is higher for the same age group and lower for different age groups. This is turned off by default but can be implemented by setting the *use-age-mixing?* switch to the 'On' position.

This is not recommended, as the model does not include the differences in severity by age group (as this would require a much more complicated interface and would be of limited value as age severity information is limited). The switch is provided for education reasons, to observe the effect of age mixing. The impact on transmission is displayed in the *Prevalence by age* plot.

### 5.3 HIGH-RISK POPULATION

Many COVID-19 policies emphasise the need to shield high risk populations such as those with compromised immune systems or certain existing respiratory diseases. This option is available in the model as discussed in section 3.1.3. This requires two parameters: the proportion of the population in the high risk category who are to be shielded (controlled by the *prop-high-risk* slider) and the risk ratio of that group relative to the remainder of the population (controlled by the *relative-risk* slider).

The parameter values and sources are summarised at Table 2. No information was available concerning the relative risk of the shielded population. Instead, the risk of hospitalisation for those aged at least 70 was used as this is the age at which people are encouraged to take extra precautions.

Table 2: Default values for high risk group probabilities

Parameter	Value	Source
prop-high-risk	0.04	Proportion of the population in the shielded patients list <sup>7</sup>
relative-risk	4	Relative hospitalisation risk of those aged at least 70 from ISARIC study <sup>5</sup> (admission numbers) and ONS population projections <sup>8</sup> (age group numbers)

### 5.4 ADDITIONAL OUTPUTS

Other outputs concern the reproduction.  $R_0$  is the standard notation for the basic reproduction ratio. This is the expected number of new cases directly created by one infectious person placed within a population of susceptible people. This single number incorporates aspects of the disease (eg how contagious it is, how long people are infectious) and the population in which it is spreading (eg number of contacts, general resistance to disease due to health status, whether people go to work when sick). The model reports the  $R_0$  calculated from the input parameters but also tracks the number of exposed people that each infectious person creates. For those people who have recovered or died, the average (over the previous week) and distribution of the number of people they infected are shown in the Reproduction plots.

<sup>7</sup> NHS Digital, *Coronavirus shielded patient list open data set, England*, [dashboard](#)

<sup>8</sup> Office of National Statistics, *Estimates of the population for the UK, England and Wales, Scotland and Northern Ireland: Mid-2019: April 2020 local authority district codes*, sheet MYE1.

---

# TECHNICAL REFERENCE

---

The technical reference is intended for any future developers to understand how to modify the code (in contrast to simply using the model). It describes the operation of the main procedures and derivations for important model parameters. There is also a section that briefly describes the way in which specific scenarios can be added to the model for use in BehaviorSpace experiments.

## 6 MODEL DESIGN AND MECHANISMS

---

Each of the two processes in the model is implemented with a key procedure. Transmission from infectious to susceptible people occurs in the *transmit-infection* procedure, and this is also the where some aspects of social distancing policies are applied. Interventions that do not directly interfere with the transmission process, such as those to do with informing contacts about potential exposure, are implemented with other procedures. The *transition-status* procedure progresses simulated people through the epidemic states. These procedures are commented, but this section provides a description of the general approach of each of these critical procedures to assist with interpretation (and further development as necessary).

### 6.1 SPREADING THE EPIDEMIC

The *transmit-infection* procedure first identifies all agents (simulated people) who are infectious, those in any of three epidemic states: “infectious”, “hospital” and “critical”. If the agent is currently in isolation, they are dropped from the infectious group (for this time step) with a probability determined by the *isolation-efficacy* slider. Those agents in hospital (“hospital” or “critical” epidemic state) are treated as if in isolation, as well as those randomly assigned to high risk status if the shielding policy is active.

Each agent in the transmitting group interacts with every agent in the “susceptible” epidemic state (excluding those in isolation with the same probability) on the same patch. With some probability assessed by a random draw, that contact results in exposure (change of state to “exposed”) of the susceptible agent.

The baseline probability for exposure is set with the *transmission-parameter* slider, but is adjusted for each contact in accordance with the social distancing policy. That is, a reduction in probability of contact arising from social distancing is implemented as an equivalent reduction in the probability of transmission but the actual contacts are not affected.

Where the social distancing policy is set to *ByContact* (and currently operating), the probability of transmission is directly reduced by the *distancing-reduction* (which controls the intensity of the policy). Both the *ByActivity* and *AllOrNone* policies are implemented at the level of individual agents, with a variable (*prop-contact*) that determines the transmission probability for the agent compared to baseline.

## 6.2 SOCIAL POLICY IMPLEMENTATION

The social distancing policies (section 3.1) are implemented in the ***transmit-infection*** procedure through a reduction in the probability a contact leads to exposure, or by removing an isolated person from the infectious or susceptible population. In contrast, the response to symptoms policies (section 3.2) are implemented in a separate ***symptomatic-actions*** procedure.

In all simulations, each agent keeps a list of the agents that it transmits the infection to. Once the agent becomes symptomatic, it checks whether the policies to inform and/or self-isolate are in place and takes the appropriate action accordingly. For the inform contacts policy, the agents already exposed are notified (with some probability) at that time, and future exposures are notified in the same time step as the contact occurs.

## 6.3 EXPOSURE TO RECOVERY OR DEATH

The ***transition-status*** procedure updates the epidemic states for all agents each time step (see Figure 3 for available paths). There are three variables for each agent that control the epidemic state transition. The variable ***epi-status*** describes the agent's current epidemic state. The variables ***next-status*** and ***next-when*** respectively identify the next state in the agent's specific path and when the transition is to occur. These variables are calculated each time the agent adopts a new state.

The ***transition-status*** procedure first identifies all the agents that are to change epidemic state in the current time step (using the ***next-when*** variable). Each randomly draws for its next state from the possible paths and, once that is determined, draws from the appropriate duration distribution to determine the number of time steps to remain in the current state. For example, an agent with a current state of "exposed" and next state of "infectious" reaches the end of its exposed duration and is identified for transition in the ***transition-status*** procedure. It transfers control to the ***Explnf*** procedure, which updates its state to "infectious", randomly selects between "immune", "dead" and "hospital" as the next state and assigns the choice to the ***next-status*** variable. Assume that is "immune", the agent will draw a duration from the weighted probability distribution in the ***get-days-tobe-infectious-toImmune*** procedure.

The default probability values for the transition paths and the sources for those values are provided at Table 1. Duration distributions are summarised at Table 3. In addition, the ***when-symptoms-if-I*** slider value is the mean of a Poisson distribution (truncated with ***max-presymptomatic*** slider) to set the number of days after entering the infectious state that a person shows symptoms.

## 6.4 UPDATING THE DEFAULT PARAMETER VALUES

It is likely that many of the parameter values in the model will need to be revised as further information is discovered about COVID-19. The only option to change the duration distributions is editing the code. Each procedure to draw from a distribution is named in the same way, start with ***get-days-*** so they are easily searchable.

Many other parameters are controlled by sliders or other interface widgets. For these, the value can be adjusted on the interface before saving the model. The new value will be set when the model is opened again. However, this approach can create problems if the model is unintentionally saved with different values, for example at the end of a session investigating different scenarios. There is therefore a ***Use Defaults*** button available to retrieve the current default values for key parameters. A

better solution to permanently updating the model is to revise the **apply-defaults** procedure, where the default values are set.

Table 3: Derivations of probability distributions for the number of days in a epidemic state, given the successor state

State	Successor	Source
exposed	infectious	Lognormal distribution fitted to early cases (Lauer et al 2020) <sup>9</sup>
infectious	immune	ECDC Q&A states 7 to 12 days without attribution <sup>10</sup>
infectious	hospital	Digitised from ECDC symptom onset to hospitalisation figure <sup>11</sup> (section 6), with a day added as infectious before symptoms.
infectious	dead	Replicates the infectious to dead distribution with an extra day. Based on assumption that should have gone to hospital and died the following day.
hospital	critical	ICNARC report <sup>12</sup> for 29 May 2020, Table 2.
hospital	immune	Digitised from Figure 2A (discharged) of Imperial Report 17 <sup>13</sup>
hospital	dead	Digitised from Figure 2A (died) of Imperial Report 17 <sup>13</sup>
critical	immune	Digitised from ICNARC Report <sup>12</sup> Figure 11 (discharge)
critical	dead	Digitised from ICNARC Report <sup>12</sup> Figure 11 (survive)

## 7 CALIBRATION

The model settings for the epidemic process are observable, such as the distribution of lengths of stay in hospital. Others are estimated in published literature using COVID-19 case histories and data collections, such as the proportion of cases that require hospitalisation. As discussed in section 5, the values should be adjusted as more is known about the epidemiology of COVID-19. The number of agents per patch cannot be adjusted but is derived from the average number of effective contacts per day in the POLYMOD<sup>14</sup> and BBC<sup>15</sup> studies, which report 11.74 and 10.47 days respectively.

However, there is limited information to guide the variables that spread the epidemic:

- **transmission-parameter:** probability that a susceptible person will become exposed given a single contact with an infected person;
- **prop-move-short:** proportion of people that move one unit of distance in a random direction in a tick (time step or day); and
- **prop-move-long:** proportion of people that move three units of distance in a random direction in a tick (time step or day).

These variables allow the modelled epidemic to spread spatially throughout the world. The general approach to calibrating these three variables was to best match the epidemic curve generated by a

<sup>9</sup> Lauer SA et al (2020), 'The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application', *Annals of Internal Medicine*, 172(9):577-582. DOI: 10.7326/M20-0504.

<sup>10</sup> European Centre for Disease Prevention and Control, [Q & A on COVID-19: When is a person infectious?](#)

<sup>11</sup> European Centre for Disease Prevention and Control, [COVID-19 Surveillance Report: Week 22, 2020.](#)

<sup>12</sup> Intensive Care National Audit & Research Centre, *ICNARC report on COVID-19 in critical care: 29 May 2020.*

<sup>13</sup> Perez-Guzman PN et al (2020), 'Clinical characteristics and predictors of outcomes of hospitalised patients with COVID-19 in a London NHS Trust: a retrospective cohort study', Imperial College Covid-19 Response Team Report 17.

<sup>14</sup> Mossong J et al (2008), 'Social contacts and mixing patterns relevant to the spread of infectious diseases', *PLoS Medicine*, 5(3):e74. DOI: doi:10.1371/journal.pmed.0050074.

<sup>15</sup> Klepac P et al (2020), 'Contacts in context: large-scale setting-specific social mixing matrices from the BBC Pandemic project', *medRxiv preprint*. DOI: 10.1101/2020.02.16.20023754.

simple deterministic model over three measures of fit: time to epidemic peak, magnitude of that peak, and total population ever infected. We followed a calibration method<sup>16</sup> that identifies the set of objectively best candidates on the Pareto efficient frontier, where the fit on one measure can only be improved at the expense of the fit on another measure.<sup>17</sup> The final choice between these candidates is subjective.

Deterministic compartment models use differential equations to estimate population numbers in each epidemic state.<sup>18</sup> These models typically assume full mixing (or mass action), where each pair of people in the model have an equal probability of contact, excluding such real-world mixing constraints as space and social networks. This assumption is also embedded within other estimates of  $R_0$  and other epidemic features.

An established online deterministic model<sup>19</sup> was used to generate the target epidemic curve. It was set to population 20 172 (the number of agents in the NetLogo model), 6 days and 9 days for incubation and mild infection durations to replicate the population and average durations in the NetLogo model. Severe and critical infections were set to the minimums available (1% and 0%) respectively, with transmission rates of 0 per day for severe infections, and no chance of death.

A modified version of this NetLogo model was used for calibration that matched these settings. It had fixed durations for the epidemic states, set to 6 for exposed and 9 for infectious. In addition, model variables were set so that 1% of infections led to hospitalisation (prop-InfHosp) and no chance of critical care or death (prob-InfImm, prob-HospImm).

The deterministic model was adjusted so that the calculated  $R_0$  was 3.0 (setting of 0.33/day transmission from mild infections). The generated epidemic curve peaked with an infected population of 3518 on day 118, with 5 people in non-susceptible state at day 16, and final size of 18 929 people recovered. These are the target values for the calibration.

The three NetLogo model input variables of interest were varied systematically with fifty runs of each of 1400 parameter value combinations (see Table 4). The model reported the tick (day) at which 5 people were exposed and the tick for the prevalence peak, as well as the standard output of peak prevalence and impact (final size). The tick at which 5 people had become exposed or infectious and the tick at which maximum prevalence (epidemic states of infectious or hospital) occurred were used to derive the time taken to reach the peak. The two other output values were used directly but expressed as the proportion of the population rather than absolute numbers of people.

Table 4: Experimental design for calibration simulations

Parameter	#Values	Values tested
transmission-parameter	20	0.005 to 0.1 by 0.005
prop-move-short	10	0.05 to 0.5 by 0.05
prop-move-long	7	0, 0.025, 0.05, 0.075, 0.1, 0.15, 0.2
repetitions		50
Simulations run:	70 000	20 x 10 x 7 x 50
Simulations checked:	59 000	Removed implausible movement
Simulations analysed	52 546	Removed those without 5 infections

<sup>16</sup> Badham J et al (2017), 'Calibrating with Multiple Criteria: A Demonstration of Dominance', *Journal of Artificial Societies and Social Simulation*, 20(2). DOI: 10.18564/jasss.3212.

<sup>17</sup> Keeney RL and Raiffa H (1993), *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*.

<sup>18</sup> Diekmann O and Heesterbeek JAP (2000), *Mathematical Epidemiology of Infection Diseases*.

<sup>19</sup> Hill A (2020), [Modeling COVID-19 Spread vs Healthcare Capacity](#).

Implausible parameter combinations, with *prop-move-short* less than or equal to *prop-move-long* were removed. Only the remaining 59 000 simulation runs were retained for further analysis. The average was calculated over the fifty runs (excluding 6 454 that never reach the critical value of 5 people) for each of the three calibration measures. Those averages were compared to the target values from the deterministic model.

The rPref package<sup>20</sup> was used to calculate Pareto fronts for the minimum difference between model output and target result over each of the three measures. While the Pareto dominance approach identifies the candidates that generates output that is objectively closest to the target curve, choosing between those candidates is subjective. Figure 10 displays the trade-offs that must be made, improving the fit of one measure at the expense of a second measure. It can also be used to identify candidates that can be rejected as requiring an unusually large loss in one measure.

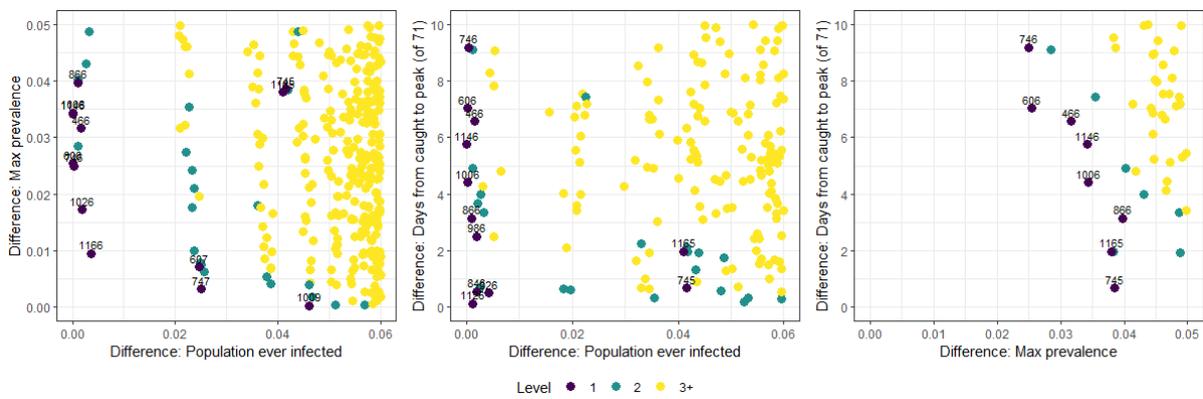


Figure 10: Pareto efficient frontiers over pairs of calibration measures

Table 5: Pareto efficient parameter setting candidates

ID	Parameters			Measures			Comment
	tr	short	long	$\Delta$ when	$\Delta$ max	$\Delta$ final	
326	0.03	0.20	0.15	0.5	0.004	0.052	Reject: $\Delta$ final
466	0.03	0.25	0.2	6.6	0.002	0.032	Reject: $\Delta$ when
606	0.03	0.30	0.2	7.0	0.000	0.025	Reject: $\Delta$ when
745	0.025	0.35	0.2	0.7	0.042	0.039	Reject: $\Delta$ max
746	0.03	0.35	0.2	9.2	0.000	0.025	Reject: $\Delta$ when
846	0.03	0.40	0.1	0.6	0.002	0.051	Reject: $\Delta$ final
866	0.03	0.40	0.15	3.1	0.001	0.040	Selected
986	0.03	0.45	0.1	2.5	0.002	0.050	Reject: $\Delta$ final
1006	0.03	0.45	0.15	4.4	0.000	0.034	
1126	0.03	0.50	0.1	0.1	0.001	0.055	Reject: $\Delta$ final
1146	0.03	0.50	0.15	5.8	0.000	0.034	Reject: $\Delta$ when, $\Delta$ final
1165	0.025	0.50	0.2	2.0	0.041	0.038	Reject: $\Delta$ max

There were 19 parameter combinations with average output measures along the Pareto efficient frontier. Of these, 7 were immediately excluded because the error in time taken to reach the peak

<sup>20</sup> Rocks P (2016), 'Computing Pareto Frontiers and Database Preferences with the rPref Package', *The R Journal*, 8(2):393-404.

was greater than 10 days. The remaining 12 non-dominated parameter combinations were assessed individually (see Table 5). Of these remaining candidates, parameter set 866 was selected as having the best compromise between the three measures of fit.

## 8 PREDEFINED SCENARIOS

---

The user manual describes how to use the model interactively, adjusting policy scenarios to explore the outcome. However, that approach is not suitable for formal scenario comparison where repeated identical adjustments are needed to calculate average outcomes and the uncertainty associated with those outcomes. Instead, any scenarios that are needed for more formal comparison are coded as specific changes at specific times, and then the BehaviorSpace tool is used to run the batch simulation, retrieving the required scenarios.

### 8.1 SCENARIO SELECTION

Each scenario must be named in the *scenario-selector* drop down chooser. The description of the scenario is placed within the *revise-scenario* procedure. That description should be a set of conditions defined by the time step (*tick*) and what settings to change at that time step. Procedures can also be called within a scenario, for example to reset specific parameters, but this has the potential to introduce errors.

### 8.2 INTERVENTION LOCALISATION

As epidemic spreads spatially, the epidemic occurs at different times in different places. However, intervention policies are implemented nationally. The scenario includes an option to choose between different triggers for the social distancing policy. These are named in the *region-selector* chooser and applied in the *scenario-trigger* procedure.

## 9 MODEL HISTORY

---

This documentation is for version 1.0 (and the first public release) of the Social Interventions COVID-19 ABM. The model (with documentation) is available from <https://github.com/jbadham/covid-social>. It was primarily developed by Dr Jennifer Badham, with contributions by Professor Brian Castellani and Dr Peter Barbrook-Johnson (CRESS, University of Surrey). All three are associated with the University of Durham COVID-19 Community Health and Social Care Modelling Team.

The model is released under the terms of the open source GNU General Public License version 3 (see the software for details).

This documentation is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/> or send a letter to:

Creative Commons  
PO Box 1866  
Mountain View, CA 94042  
USA



## 9.1 ACKNOWLEDGEMENTS

We would also like to acknowledge the valuable assistance of Dr Ruth Meyer (Centre for Policy Modelling, Manchester Metropolitan University) and Dr Corinna Elsenbroich and Dr Kavin Narasimhan (both from Centre for Research in Social Simulation, University of Surrey) who tested various draft versions of the model and made valuable suggestions for improvement.

The COVID-19 Community Health and Social Care Modelling Team at Durham University is under the guidance of our respective Executive Deans, Jacqui Ramagge (Science) and Charlotte Clarke (Social Sciences and Health) as part of the contribution of the Wolfson Research Institute for Health and Wellbeing and the Institute of Data Sciences to the University's Health@Durham strategy, as well as supported by the Research and Innovation Services, Marketing and Communications and CECAN (Centre for the study of Complexity Across the Nexus). The team is led by Dr Camila Caiado and Professor Brian Castellani, with the purpose of creating a series of tools and dashboards that Trusts and Councils can use to help support decision and planning accordingly. We would also like to acknowledge the outstanding contribution of Professor Amanda Ellison, Dr Andrew Iskauskas, Dr Bernard Piette, Dr Rachel Oughton, and Dr Corey Schimpf.