

Agent-Based Modeling Grants Insight into the Effect of Policy Responses to COVID-19

Alec Mason

Advised by Brandilyn Stigler

Department of Mathematics
Southern Methodist University
Dallas, Texas 75205

September 2020

Abstract

This project aimed to create a complex agent-based model for pattern analysis in the spread of SARS-CoV-2, the virus that causes COVID-19. The model incorporates a variety of adjustable parameters pertaining to social distancing, testing, contact tracing, mortality, infectiousness, and hospital capacity. Variables that can be studied include cumulative infections, peak infections, peak day, death count, and case-fatality rate. This model was then used to conduct experiments involving thousands of simulated runs to analyze the effects of movement reduction and testing procedures on preventing and slowing the spread of the virus. These experiments found that the cumulative and peak case counts roughly follow a sigmoidal curve, with a range of movement reduction in which the rate of decline is maximized. In addition, the data suggested that while less stringent reductions in movement are not as effective in reducing cumulative case counts, they can still be effective in lowering the peak case count. In regards to testing, the experiments found that while contact tracing is slightly effective in reducing cumulative case counts, it is more effective in reducing peak case counts if positive individuals severely limit their contact with others. Incorporation of isolation mechanisms, whether mandatory or voluntary, can magnify these reductions in cases. We conclude that an optimal strategy involves a combination of significant reductions in movement, measures to reduce infectiousness, maximizing testing, contact tracing, and mandatory or voluntary isolation mechanisms for the confirmed infected. The model was implemented in NetLogo version 4.1.3.

1 Introduction

The spread of the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) and its corresponding disease, COVID-19, has had a profound impact on economic and social operations around the world. The high infectiousness of SARS-CoV-2 has led to exponential growth, with some studies estimating the average reproductive number to be as high as 3.28 [5]. This means that each infected individual will infect more than three others on average. While the case-fatality rate of SARS-CoV-2 is nearly universally estimated to be lower compared to recent epidemic coronaviruses such as SARS-CoV and MERS-CoV, the stronger tendency towards exponential growth due to faster spread has led to case numbers, and subsequent death counts, unlike any recent epidemic [3]. A recent study places the international infection fatality rate of SARS-CoV-2 at 0.66%, with rates increasing with age [7]. However, hospital capacity may play a role in the case-fatality rate, with the rate of fatality increasing as access to care becomes more scarce.

Agent-Based Models (ABMs) can be an effective tool for modeling the collective effects of individual behaviors [6]. Each individual agent in the model is given a set of rules that they will follow in reaction to their environment. Randomness is sometimes incorporated to account for the complexity of real-life behaviors and environmental factors. Each individual model run represents a single resultant path created by the combination of behaviors and factors occurring in that run. While this means that individual model runs should not be interpreted as quantitatively predictive, averaging repeated model runs with varying initial starting conditions can grant insight into averages and larger patterns [1].

This project aims to create a model capable of capturing the effect of various parameters on the spread of COVID-19 through a population using the NetLogo agent-based modeling software [10]. The model, based loosely upon the “Virus” model created by U. Wilensky [9], allows for the adjustment of parameters like movement reduction, social distancing regulation thresholds, and testing availability. Outcomes of the model that can be analyzed include cumulative cases, peak-case count, cumulative deaths, and case-fatality rate.

In Section 2, the parameters, variables, and processes within the model are presented and described. In Section 3, experiments are conducted via aggregation of multiple runs to test the effects of various parameter settings, and recommendations for policy are made based on the data. Finally, the results, implications, and utility of the model are further discussed in Section 4.

2 Model

The model is described using the Overview-Design Concepts-Details protocol [4] and implemented in NetLogo version 4.1.3.

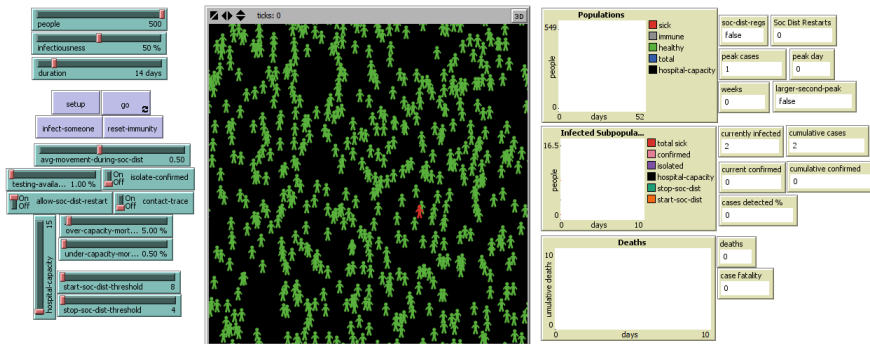


Figure 1: Screenshot of the user interface of the COVID-19 model described in this research

2.1 Purpose

This model aims to study the effect of various parameters on the spread and impact of the SARS-CoV-2 virus in a community. One of the parameters that this model simulates is participation in social distancing efforts. Characterizing the effects of different levels of movement reduction can elucidate the threshold at which the efforts become most effective. In addition to depicting individual participation, community thresholds for implementing, relaxing, and potentially reinstating social distancing regulations are also parameters that can be studied in this model. Another goal of this model is to analyze the effect of testing and quarantine programs on viral spread. The effects of these parameters will show the effectiveness of different levels of daily randomized testing within the community, and can potentially grant insight into which actions are best taken after a positive test. The model is also capable of simulating a contact tracing program. Viral factors such as infectiousness and mortality rate can be adjusted, and community factors like population density and hospital capacity can also be altered and studied in this model, allowing their impacts on spread and fatality to be studied.

2.2 Entities, state variables, and scales

The entities within the model represent individuals within a community experiencing an outbreak of SARS-CoV-2. These entities, hereinafter referred to as “turtles” in the terminology of NetLogo, can exist in one of five states:

- healthy (represented by a green coloration)
- infected but not confirmed (red)
- confirmed infected (pink)
- infected and isolated (purple)

- immune (gray)

The number of turtles present at the start of the simulation can be adjusted via the `people` slider, on a scale of 0 to 500. The only slider directly affecting the behavior of the turtles is `avg-movement-reduction`, which is a scale from 0.1 to 1 that represents the average number of movements per turtle per day. It is assumed that turtles practicing social distancing will move an average of once every 7 days. So the model calculates what percentage of turtles will need to social distance in order to achieve this average daily movement. This calculated parameter becomes the `social-distance-participation`. When each turtle is created, they generate a number from 0 to 1. If the number they generate is lower than the `social-distance-participation` parameter, they will always follow social distancing guidelines. If they roll higher, they will always ignore the guidelines.

The properties of the disease itself can be adjusted via 4 sliders. The infectiousness of the virus, expressed as the likelihood that an infected turtle will infect a healthy turtle on contact, can be adjusted via the `infectiousness` percentage scale from 0 to 100. The duration of the disease, which is interpreted as the time in which a turtle is infectious, can be adjusted via the `duration` scale from 0 to 99 days. Disease fatality is determined by a combination of two sliders, `under-capacity-mortality` and `over-capacity-mortality`, both of which are on a percentage scale from 0 to 100, though it should be noted that the former slider can never exceed the latter. The `under-capacity-mortality` slider determines the likelihood that a turtle infected when the number of infected turtles does not exceed the hospital capacity will die. Likewise, the `over-capacity-mortality` slider determines the likelihood of death for an individual infected when the number of infected turtles exceeds the hospital capacity.

The remainder of the sliders and switches pertain to community regulations in regards to social distancing, healthcare availability, and testing. The `testing-availability` slider can be adjusted on a percentage scale from 0 to 100 at increments of 0.25%, and it determines the percentage of the population that will be tested daily at random. Turtles who are randomly selected to be tested and are currently in the `sick?` but not in the `confirmed?` state will be assigned to the `confirmed?` state. If the global `isolate-confirmed` state is set to true, the `confirmed?` individual will be put into the `isolated?` state. The infected subpopulations act as a hierarchy, so turtles in the `isolated?` state are also in the `confirmed?` state, and all `confirmed?` turtles are also in the `sick?` state. The availability of healthcare resources in the community can be adjusted via the `hospital-capacity` slider, which ranges from a minimum of 0 and a maximum equal to the `people` slider at initialization. It should be noted that this slider represents the number of infections within the community at which the hospital system will be overloaded, including cases who do not necessarily require hospitalization. For example, if one quarter of cases require hospitalization, and the hospital has 20 beds, then the `hospital-capacity` should be set

at 80.

$$\text{hospital capacity} = \frac{\text{number of beds}}{\text{fraction requiring hospitalization}}$$

The final sliders and settings refer to the enactment of social distancing regulations. The threshold of active confirmed cases for starting social distancing is determined by the `start-soc-dist-threshold` slider, which can be adjusted on a scale of 0 to the size of the population. Likewise, the `stop-soc-dist-threshold` slider determines the number of confirmed cases at which social distancing regulations will end, and the slider can be adjusted from 0 to the size of the population. The threshold for stopping social distancing cannot exceed the threshold for starting social distancing. It is also important to note that these thresholds are based on the number of confirmed cases, not total cases. This is done to account for the fact that community governments only have data for the cases that they have tested for. Therefore, if `testing-availability` is set to 0%, social distancing regulations will never start unless the threshold for starting is set to zero confirmed cases.

2.3 Process overview and scheduling

After the `go` procedure is initiated, a series of procedures are sequentially repeated each tick: `progress`, `move`, `infect`, `recover`, `test`, `update-colors`, `update-global-variables`, `tick`, and `update-plot`. The `progress` procedure increases the `sick-count` variable on all of the currently infected turtles by 1. This procedure also resets the color of all patches to black to allow the radius indication for contact tracing to show.

In the `move` procedure, the first step is to determine if social distancing regulations are currently in effect. If the global `soc-dist-regs` variable is currently set to false, all turtles who are not confirmed to be infected will rotate randomly up to 100 degrees left or right and move forward one patch. Turtles who are confirmed infected will generate a random number from 0 to 7. If this number is less than 1, the turtle will rotate randomly up to 100 degrees in either direction and move forward one patch. If the generated number is greater than 1, the turtle will not move on that tick. If the global `soc-dist-regs` variable state is currently set to true, turtles who were assigned the `no-soc-dist` state upon initialization and are not confirmed to be infected will move as normal, rotating up to 100 degrees in either direction and moving forward one patch. If the turtle has `no-soc-dist` set to false or is confirmed infected, they will generate a random float number from 0 to 7. If the generated number is less than 1, they will rotate randomly 100 degrees in either direction and move forward one patch. Otherwise, the turtle will not move on that tick. See Figure 2 for a flowchart.

The `infect` procedure tells all non-isolated infected turtles to check for other turtles on the same patch. For each healthy, non-immune turtle on the same patch, a random float number is generated between 0 and 100. If this number is less than the `infectiousness` parameter, the healthy turtle will carry out the `get-sick` procedure and become infected. In the `get-sick` procedure, the

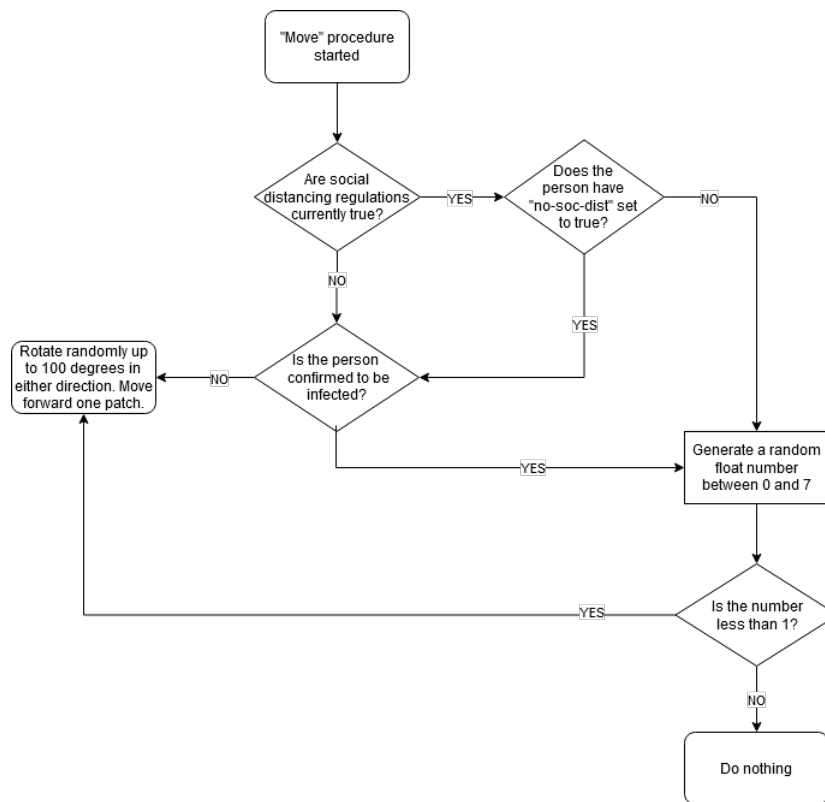


Figure 2: Flowchart depicting the move procedure steps run by each turtle.

`sick?` turtle variable is set to true, the turtle's color is changed to red, and the turtle is assigned a mortality likelihood. If the current count of infected turtles is higher than the `hospital-capacity` parameter, the turtle will be assigned the `over-capacity-mortality` as their mortality variable. If the current count of infected turtles is lower than the `hospital-capacity` parameter, the turtle will be assigned the `under-capacity-mortality`.

Next, the `recover` procedure asks all infected turtles to compare their `sick-count` variable to the `duration` global parameter. If a turtle's `sick-count` is greater than the `duration`, they are asked to generate a random float number from 0 to 100. If the number is lower than the turtle's `mortality` variable, the turtle will die. If the number is higher, the turtle will carry out the `become-immune` procedure. In this procedure, all infection-related variables are set to false, the `sick-count` and `mortality` variables are set to 0, the `immune?` state is set to true, and the turtle color is changed to gray.

The `test` procedure asks a certain fraction of the turtles, equivalent to the `testing-availability` parameter, the status of their `sick?` variable. If the selected turtle has `sick?` set to true, they will start the `confirm-case` procedure. If the global `contact-trace` procedure is set to true, this will also trigger testing of all turtles within a three patch radius. A cyan indicator will flash to indicate the radius. In the `confirm-case` procedure, the infected turtle has their `confirmed?` variable set to true, in addition to the `sick?` variable which is already set to true. This extra variable will slow their movement in the `move` procedure discussed earlier. If the global `isolate-confirmed` procedure is set to true, the turtle will also have their `isolated?` variable set to true. As discussed in the `infect` procedure explanation, the `isolated?` variable prevents the turtle from infecting others.

The `update-colors` procedure is a housekeeping procedure that ensures that turtle colors are correctly assigned by the end of the tick, specifically within the infected subgroups. Turtles who are sick but not confirmed are colored red. Turtles who are confirmed sick but not isolated are colored pink. Turtles who are confirmed sick and isolated are colored purple.

The `update-global-variables` procedure ensures that reporting variables and variables that cause changes in model behavior are adjusted to reflect current conditions. The percent of turtles infected, percent of turtles immune, and the peak information reporting variables are updated. In addition, the important `social-dist-regs` variable state is reassessed at this step. First, the procedure asks if the current number of confirmed infections is under the `stop-soc-dist-threshold` parameter. If this is true, the procedure will set `soc-dist-regs` to false. Next, the procedure will ask a complex if-statement. If the current count of confirmed infected turtles is greater than or equal to the `start-soc-dist-threshold` parameter, and the `soc-dist-res-gauge` variable is not set to 1 when the `allow-soc-dist-restart` global state is set to false, then `soc-dist-regs` will be set to true and `soc-dist-res-gauge` will be set to 1. The `soc-dist-res-gauge` variable allows the model to know if social distancing has already been started once in the run, and will prevent it from starting the regulations again in the case that `allow-soc-dist-restart` is set

to false. After the `update-global-variables` procedure, the model moves forward one `tick` and continues.

Finally, the `update-plots` procedure updates the graphs that can be seen on the model interface. The model then repeats from the beginning of the `go` procedure until stopped.

2.4 Design concepts

2.4.1 Emergence

There are a variety of measures that can be used to evaluate the results of each model run. One of the most important measures is the cumulative case count, which allows for a broad interpretation of the effectiveness of virus containment. A low number on this measure could indicate that global parameters were effective in preventing viral propagation. However, many real responses to COVID-19 have focused on slowing the spread of the virus rather than completely containing it. In order to interpret the effectiveness of “flattening the curve,” the measures of peak cases, peak day, and deaths are useful. The peak cases measure allows the model interpreter to conclude how high the peak of active cases was, and the peak day measure allows for the interpretation of peak delay that usually results from efforts to “flatten the curve.” The death count measure, in conjunction with the case fatality measure, is useful in that it helps interpret how overloaded the hospital system was during the model run. If the model spent There are also some measures that are useful in analysis of specific parameters. The Boolean `larger-second-peak` becomes true when a new case record is set after a period of lower case counts.

2.4.2 Adaptation

Individual turtles in this model exhibit adaptation most prominently in the `move` procedure. If they are assigned to the group who follows social distancing regulations upon generation, they will slow their movement when the global social distancing regulations are on. The same is true for any turtle who is confirmed to be infected. There is no specific goal from the turtle’s perspective, and turtles do not actively avoid infected turtles.

2.4.3 Sensing

Turtles in this model are assumed to be able to sense when social distancing regulations are in effect and adjust their movement accordingly. They are also assumed to know when they have been tested for the virus. Infected turtles can sense when a healthy turtle is on the same patch, and they will potentially pass the infection on.

2.4.4 Interaction

Turtles interact with one another in the event of spreading the infection. When a healthy turtle occupies the same patch as an infected turtle, they will interact and the healthy turtle has a chance to also become infected. In the event that contact tracing is enabled, the newly confirmed turtles will interact with turtles within a three patch radius to refer them for testing as well.

2.4.5 Stochasticity

This model assumes that movement, social distancing participation, selection of turtles for testing, infectiousness, and mortality are at least somewhat random. At setup, turtles are assigned to groups determining participation in social distancing precautions based on generated random numbers. For confirmed infected turtles or turtles who are practicing social distancing, the likelihood of movement on that tick is determined by a random number. The randomness of assignment to social distancing groups and movement determination are used to simulate an average decrease in movement during social distancing. This average is determined by the `avg-movement-during-soc-dist` slider. When movement occurs, it can happen in any random direction within a 100-degree range, centered on the current facing direction. When an infected turtle occupies the same patch as a healthy turtle, a random number is used to determine whether or not the infection will be passed to the healthy turtle. This is used to simulate the contagiousness of the virus. In addition, this randomness accounts for somewhat random factors in real life like exposure time and viral load. While the assignment of mortality likelihoods is not random, the chance a turtle will die at the end of its infection course is based on a random number generation. Since mortality varies from place to place in reality and can vary based on a variety of personal and environmental factors, this randomness allows for simulation of a rate without strictly enforcing it.

2.4.6 Collectives

A variety of collectives are formed in this model, each with varying impacts on the model. Turtles are assigned into collectives that determine whether or not they will comply with social distancing regulations at setup. Individually, the group that one is in will impact the individual's movement behavior. Globally, these collectives have a direct effect on average movement and an indirect effect on infection propagation.

Probably the most impactful collectives are the population groups: healthy, sick, confirmed, isolated, and immune. When in the sick group, turtles will pass the infection on to healthy turtles. Turtles in the confirmed group will slow their movement, no matter if they normally obey social distancing. The number of turtles in the confirmed group will determine whether or not the social distancing regulations are on or off. Turtles in the isolated group will slow their movement and be unable to infect other turtles. As `isolated?` is

a subgroup of `confirmed?`, which itself is a subgroup of `sick`, turtles in this group will be assigned to all three groups.

Another set of collectives is formed in regards to mortality. Turtles who become sick when hospitals are at or over capacity will be assigned to the high mortality group. Likewise, turtles who become sick when the hospitals are under capacity will be assigned to the low mortality group. The only effect each collective has is during the `recover` procedure, in which those in the high mortality group will be more likely to die.

2.5 Initialization

When the model is initialized via the command `setup`, all monitors are set to the default values, and a number of agents, equal to the `people` slider, are generated. The agents, referred to as turtles, are placed at random on a 35×35 grid of patches. Each turtle generates a number from 0 to 1. If the number they generate is greater than the `social-distance-participation` parameter, calculated from the `avg-movement-during-soc-dist` slider, their `no-soc-dist` variable will be set to true, which means they will ignore any regulations on social distancing until they test positive. All turtles execute the `get-healthy` procedure, which sets their color to green and sets the `sick?` and `immune?` variables to false. One turtle is selected at random to execute the `get-sick` procedure. This procedure sets their `sick?` variable to true. The `get-sick` procedure also asks the for the global count of turtles with the `sick?` variable set to true. If this count is higher than the `hospital-capacity` parameter setting, the `mortality` variable for the newly sick turtle is set to `over-capacity-mortality`. Otherwise, the `mortality` is set to the `under-capacity-mortality`. Once the setup of the turtles is complete, the plots are cleared, and the global variables are updated to reflect current counts of infected, immune, and healthy turtles.

3 Analyses and Recommendations

Agent-based models are useful in analyzing patterns and overarching concepts. Since each model run represents one possibility resulting from the combination of various random events, analysis of averages and fit lines after multiple runs is a useful tool. The BehaviorSpace tool within NetLogo allows for automation of the setup, run, and data collection processes of the model which permits simulation experiments to be performed to test the effect of parameter values on the model.

For the first experiment, the BehaviorSpace rules were set up such that only the `avg-movement-during-soc-dist` parameter would vary while the rest remained constant. The setting of the movement parameter was set to change every 500 runs, for a total of 6000 runs. The thresholds for social distancing were both set at zero cases so that social distancing would remain true throughout the run. Each run was set to automatically terminate when the count of infected

Parameter	Constant or Varied?	Value(s)
avg-movement-during-soc-dist	varied	0.1, 0.15, 0.2, 0.25, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
people	constant	500
infectiousness	constant	50
duration	constant	14
testing-availability	constant	0
isolate-confirmed	constant	FALSE
allow-soc-dist-restart	constant	FALSE
contact-trace	constant	FALSE
hospital-capacity	constant	15
over-capacity-mortality	constant	5
under-capacity-mortality	constant	0.5
start-soc-dist-threshold	constant	0
stop-soc-dist-threshold	constant	0

Figure 3: Parameter settings for the social-distancing movement experiment. See Figures 4 and 5 for results.

turtles reaches zero. Data on cumulative infections, deaths, peak infections, and peak day were collected; see Figure 3 for a full list of parameter settings.

The raw comma-separated value data was exported to OriginPro 2020 [8], and a scatter plot was generated depicting the cumulative infections in each run against the social distancing movement parameter setting. The resultant graph is shown in Figure 4, fitted with a sigmoidal function.

The plot shows that reductions in movement have an effect on overall disease spread, but the effect on overall case counts is slight until movement is reduced to below half of normal movement. There is a steep decline in overall case counts between 0.2 and 0.4. From this, it can be interpreted that containment of the virus would require intense reductions in movement.

Anonymized cell phone data collected by Johns Hopkins University researchers shows that few counties reached this reduction amount [2]. At the peak of infections in April, many major counties in the United States were at a higher ratio such as Philadelphia County, PA, at a ratio of 0.52; Los Angeles County, CA, at a ratio of 0.56; and Harris County, TX, at a ratio of 0.62. One of the locations with the most reduced movement in April was New York City with a ratio of 0.35. According to the modeling data, this puts the city within the range of maximal change in effectiveness in preventing viral spread.

As complete containment of COVID-19 has proven to be difficult, many governments have opted instead to focus on slowing the spread of the virus. The hope is that such actions will prevent hospital overload and minimize mortality rate during the period of drug and vaccine development. Figure 5 shows the plot of peak infections versus the average movement parameter, fitted with a sigmoidal function.

While both the cumulative and peak-infection plots appear to follow a sigmoidal shape, the upper asymptote on the peak-infections plot does not appear

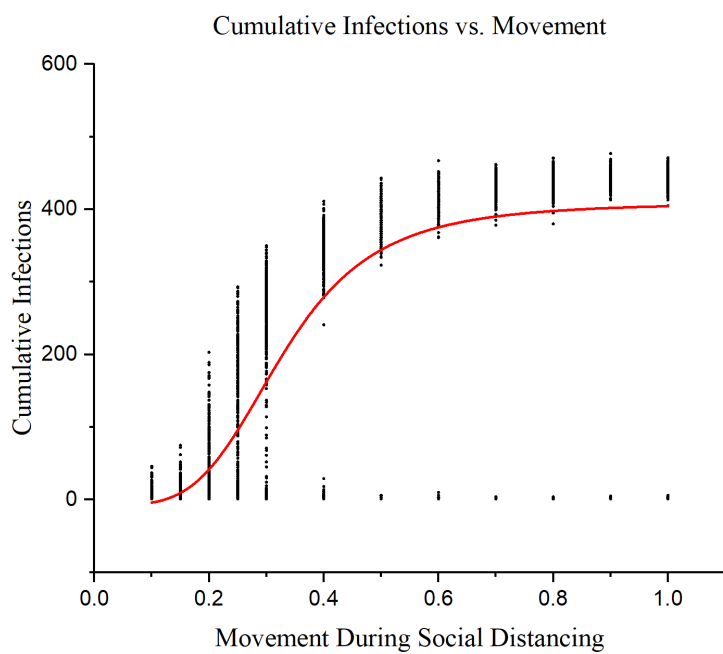


Figure 4: Scatter plot of cumulative infections in the social-distancing movement experiment; see Figure 3 for parameter settings. Each black dot represents a run. The data is fitted with a sigmoidal function, shown as a red solid curve.

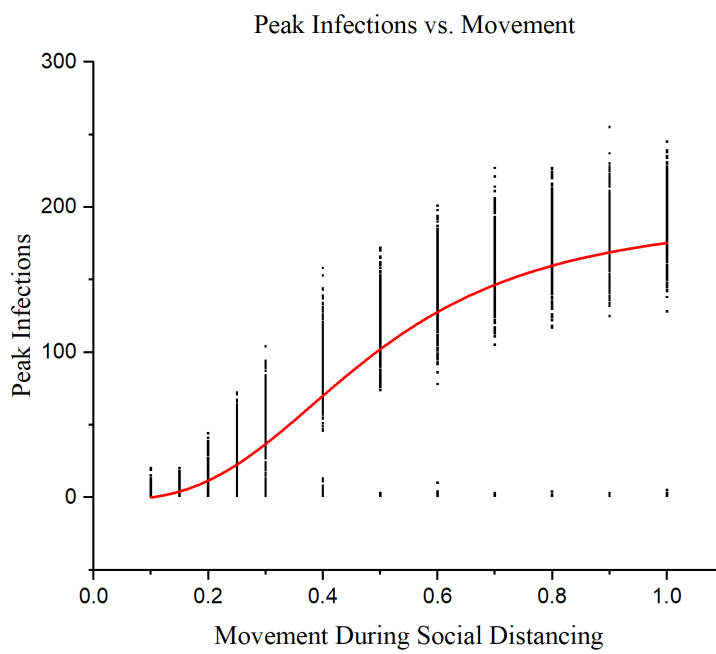


Figure 5: Scatter plot of peak infections in the social-distancing movement experiment; see Figure 3 for parameter settings. Each black dot represents a run. The data is fitted with a sigmoidal function, shown as a red solid curve.

as flat. This means that smaller reductions in movement can have a significant effect on the infection count at the peak, with a higher steep range of 0.3 to 0.5. This supports the idea that slowing the spread of COVID-19 is easier to accomplish, and many counties in the United States might have achieved at least some reduction in peak infections even with a small reduction in movement. However, it is important to note that the range in which these steeper declines occur is dependent on the infectiousness of the virus and the population density. For example, lowering the infectiousness will increase the width of the range and raise the values necessary. Lowering the population density will have a similar effect. This means that measures that reduce infectiousness such as mask wearing may allow communities to achieve the same reduction in cases with less reduction in movement.

One noticeable feature in Figures 4 and 5 is that both have runs with very low infection rates even at high movement parameter settings. This suggests that there may be two major paths that the disease curve can take from the beginning of the outbreak: a high infection path and a low infection path, with very few runs falling in between.

Based on the results of the movement variation experiment, we suggest that communities reduce their movement rate to the minimum possible while maintaining some degree of community economic health. Ideally, this rate should be 0.5 or lower if slowing the spread is the goal, and lower than 0.4 if containment is the goal. These goals should be even lower in communities with low hospital availability or high senior populations.

The second experiment that was conducted analyzed the effect of various testing and follow-up methods. By default, individuals in this model who test positive reduce their movement to one-seventh of their normal rate until they are no longer sick, mimicking once a week movement. Selection of testing recipients is random among turtles which are not currently `confirmed?`. If the `contact-trace` global procedure is set to true, each positive test triggers an automatic test on all turtles within a 3 patch radius of the confirmed turtle. If the `isolate-confirmed` global procedure is set to true, confirmed infected turtles will lose their ability to infect other turtles.

This experiment aims to see what effect each of these settings has on the infection rate at different values for the `testing-availability` parameter; see Figure 6 for the parameter settings in the BehaviorSpace. The BehaviorSpace experiment was run four times, one for each combination of the `contact-trace` and `isolate-confirmed` global procedures. In each experiment, 500 runs were conducted for each `testing-availability` setting, resulting in 7000 runs per BehaviorSpace experiment and 28000 runs all together.

The raw data from these experiments was then exported to Microsoft Excel, and a custom macro was used to calculate averages for cumulative and peak infections at each level of testing availability. The average data for each experiment was plotted in OriginPro 2020, with the full range shown in the left panel of Figure 7.

The data suggest that the addition of contact tracing alone grants slight improvement in cumulative infection counts. Complete isolation of the confirmed,

Parameter	Constant or Varied?	Value(s)
testing-availability	varied	0, 0.25, 0.5, 0.75, 1, 1.5, 2, 2.5, 4, 5, 10, 25, 50, 90
contact-trace	constant (between runs) varied (between exps)	FALSE (exps 1 and 3) TRUE (exps 2 and 4)
isolate-confirmed	constant (between runs) varied (between exps)	FALSE (exps 1 and 2) TRUE (exps 3 and 4)
people	constant	500
infectiousness	constant	50
duration	constant	14
avg-movement-during-soc-dist	constant	1
allow-soc-dist-restart	constant	FALSE
hospital-capacity	constant	15
over-capacity-mortality	constant	5
under-capacity-mortality	constant	0.5
start-soc-dist-threshold	constant	0
stop-soc-dist-threshold	constant	0

Figure 6: Parameter settings for the testing availability and procedure experiments. Each experiment tested the full range of the `testing-availability` parameter. Experiments 2 and 4 had `contact-tracing` set to true and experiments 3 and 4 had `isolate-confirmed` set to true. Experiment 1 had neither procedure enabled as a control. See Figures 7 and 8 for results.

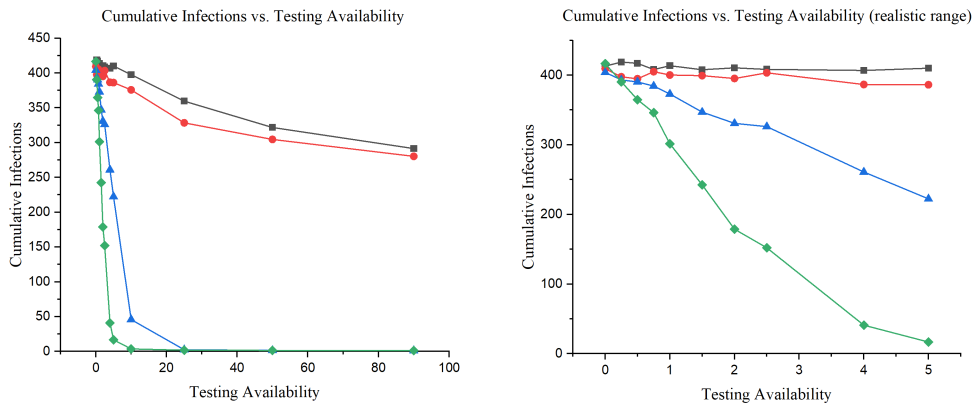


Figure 7: Plots showing the average cumulative infections in the four testing experiments; see Figure 6 for parameter settings. Standard testing protocol is shown in black with square markers (experiment 1); with contact tracing is shown red with circular markers (experiment 2); with isolation of the confirmed is shown in blue with triangular markers (experiment 3); and the experiment with both contact tracing and isolation is shown in green with diamond markers (experiment 4). The left plot shows the average cumulative infections versus testing availability for the four experiments with full range. The right plot shows the average cumulative infections versus testing availability for the four experiments, zoomed into the range from 0-5% testing availability.

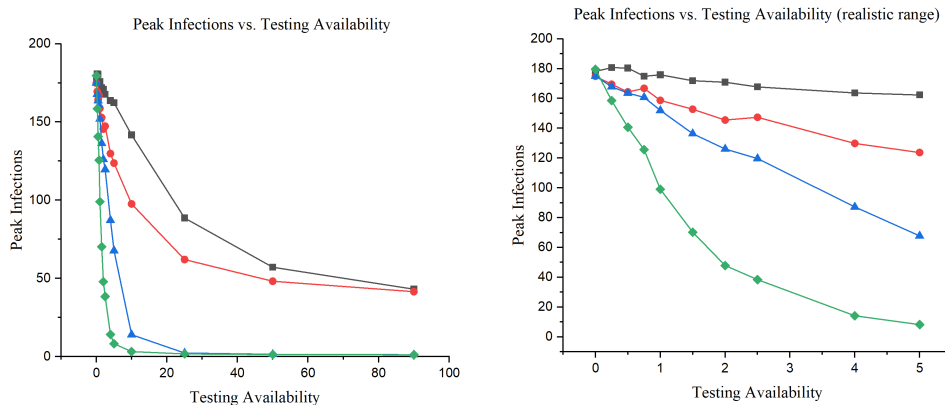


Figure 8: Plots showing the average peak infections versus testing availability for the four experiments. The left plot shows the average peak infections versus testing availability for the four experiments with full range. The right plot shows the average peak infections versus testing availability for the four experiments, zoomed into the range from 0-5% testing availability. Standard testing protocol is shown in black with square markers; contact tracing is shown red with circular markers; isolation of the confirmed is shown in blue with triangular markers; and the experiment with both contact tracing and isolation is shown in green with diamond markers.

however, appears to be very effective in stopping the disease from spreading completely. It is important to note, however, that this experiment tested ranges that are likely too high for human societies to achieve. Testing more than 5% of the population per day is not realistic for most communities and governments. The right panel of Figure 7 shows the same graph zoomed into the range of 0-5% testing availability.

The data suggest that, in realistic ranges, the addition of contact tracing alone will not lower overall infection counts. However, if those who test positive completely isolate themselves, it is possible to significantly reduce the spread of the disease even in lower testing availability ranges.

Plots from the same experiments for average peak infections are shown in Figure 8, with full range shown in the left panel and realistic range shown in the right. Unlike with cumulative infection counts, peak infection counts appear to be much more strongly impacted by the mere implementation of contact tracing. This corresponds well to the social distancing variation data that suggested less stringent policies may still have a significant effect on the speed of spread despite having little influence on cumulative infection counts. As with the cumulative infection data, isolation of the confirmed can lead to even larger improvements in peak infection counts.

Based on the data from the testing availability variation experiment, we suggest that communities work to maximize testing rates while implementing

contact tracing. Communities should also conduct randomized survey testing regularly to find asymptomatic carriers. Forced isolation of the confirmed may not be politically or logistically feasible, especially in countries that value individual freedom. However, providing voluntary quarantine locations such as repurposed hotels to allow infected individuals to self-isolate could provide a similar effect on infection counts. Governments could also provide no-contact delivery services for necessities to allow infected individuals to isolate at home.

4 Discussion

This model and the experiments conducted on it aimed to provide insight into the effectiveness of policies such as social distancing and testing in curbing the spread of SARS-CoV-2. The data indicate that while stringent responses like severe movement reduction and complete isolation of all confirmed infections are most effective in containing the virus, less intense policies still have the potential to slow the spread and prevent hospital overload and unnecessary deaths. The experimental data leads us to recommend that communities reduce their average movement by at least half, and that governments maximize testing availability and implement contact tracing and some degree of infection isolation. Where it is not politically or logistically expedient to enforce a quarantine, governments should provide venues for voluntary isolation and provide services like no-contact delivery of essentials for confirmed infected individuals. In order to reduce unnecessary exposure, governments might also consider providing financial support for non-essential workers to stay home and businesses to remain closed. Combining various policy responses, including reducing movement, increasing testing, contact tracing, encouraging infection isolation, and reducing infectiousness by wearing masks, will allow for more flexibility while still achieving the same reduction in cases.

It should be noted that this modeling process has limitations. First, the experiments that were run cannot provide concrete numbers for movement reduction and testing availability that apply to every community or situation. Factors like population density and infectiousness reducing measures like mask wearing may shift the threshold for maximal reduction in case counts. Communities should adjust the model parameters to match their specific conditions. Further research is needed to quantify the average chance of infection upon contact under various circumstances. However, the main assertions still stand that less movement usually leads to an improvement in case counts, and reducing cumulative infections requires much more stringent measures than reducing peak infections.

Another limitation of this model that testing is random under standard protocol and that results are processed instantly. An improved model may include more subclasses in the `sick?` category that accounts for symptomaticity, with symptomatic individuals having a certain probability of seeking out a test. The model could also be improved by adding a delay between testing and confirmation to simulate real logistical issues.

This model also cannot simulate high-risk groups such as front-line workers and the elderly when determining infection risk and mortality. These values rather operate on averages, meaning that community leaders will need to adjust parameters to account for this for their specific population breakdown.

Models are some of the most effective tools in understanding current and potential future conditions during a disease outbreak. They help leaders and individuals make decisions that can influence future outcomes for the better. While no model is perfect, studying patterns present across multiple modeling strategies can be effective in making informed decisions.

References

- [1] A. Auchincloss and L. Garcia. Brief introductory guide to agent-based modeling and an illustration from urban health research. *Cadernos de Saúde Pública*, 31(1):65–78, April 2017.
- [2] H. Badr, H. Du, M. Marshall, E. Dong, M. Squire, and L. Gardner. Association between mobility patterns and COVID-19 transmission in the USA: a mathematical modelling study. *The Lancet Infectious Diseases*, July 2020.
- [3] M. Fani, A. Teimoori, and S. Ghafari. Comparison of the COVID-2019 (SARS-CoV-2) pathogenesis with SARS-CoV and MERS-CoV infections. *Future Virology*, May 2020.
- [4] V. Grimm, U. Berger, F. Bastiansen, S. Eliassen, V. Ginot, J. Giske, J. Goss-Custard, T. Grand, S. Heinz, G. Huse, A. Huth, J. Jepsen, C. Jørgensen, W. Mooij, B. Müller, G. Pe’er, C. Piou, S. Railsback, A. Robbins, M. Robbins, E. Rossmanith, N. Rüger, E. Strand, S. Souissi, R. Stillman, R. Vabø, U. Visser, and D. DeAngelis. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2):115–126, June 2006.
- [5] Y. Liu, A. A. Gayle, A. Wilder-Smith, and J. Rocklöv. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine*, 27(2):1–4, March 2020.
- [6] R. Robeva and T. Hodge, editors. *Mathematical Concepts and Methods in Modern Biology: Using Modern Discrete Models*. Academic Press, 2013.
- [7] R. Verity, L. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, G. Cuomo-Dannenburg, H. Thompson, P. Walker, H. Fu, A. Dighe, J. Griffin, M. Baguelin, S. Bhatia, A. Boonyasiri, A. Cori, Z. Cucunubá, R. FitzJohn, K. Gaythorpe, W. Green, A. Hamlet, W. Hinsley, D. Laydon, G. Nedjati-Gilani, S. Riley, S. van Elsland, E. Volz, H. Wang, Y. Wang X. Xi, C. Donnelly, A. Ghani, and N. Ferguson. Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*, 20(6):669–677, March 2020.

- [8] Origin(Pro), Version 2020. OriginLab Corporation, Northampton, MA, USA.
- [9] U. Wilensky. NetLogo virus model. <http://ccl.northwestern.edu/netlogo/models/virus>, 1998. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.
- [10] U. Wilensky. NetLogo. <http://ccl.northwestern.edu/netlogo/>, 1999. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.