Effect of Behavioural Strategies on Individual Quality of Life in Agent-Based Social Simulations of COVID-19

Sterre Lutz[1] Dr. Neil Yorke-Smith[1]
[1]TU Delft
s.lutz@student.tudelft.nl, n.yorke-smith@tudelft.nl

June 2020

Abstract

Since the onset of the COVID-19 pandemic, many models have been made to predict the spread and responses to it. Although moral decision-making during the uncertainty pandemics is suggested to be more motivated by individual incentive than collective incentive, decision-making in COVID-19 agent-based models is often modelled implicitly by cultural, social and economic characteristics, making it difficult to investigate the relation between decision-making and individual pay-off. The main research question of this paper is: "What behavioural strategies can maximise an individual's quality of life in an agent-based social simulation of COVID-19?". This is answered for two basic behavioural strategies, Isolate and Ignore, modelled in an existing agent-based model of COVID-19 called *ASSOCC*. Patterns in quality of life per strategy are investigated by what they look like over time and how they are influenced by the degree of isolation and distribution of isolation strategies. The two main conclusions from the experiments in this paper are (1) isolation behaviour can be an acceptable or even a preferable strategy if and only if it is done full-time and collectively and (2) although collective isolation is preferable over collective non-isolation, there is an individual incentive not to join in on this collective isolation. These results can only be seen within the confines of the agent-based model and any proof towards its extension to real life is left for future work.

1 Introduction

Since the onset of the COVID-19 pandemic, research on infectious diseases has surged, and many different models have been created to analyse the spread of this disease and the effectiveness of preventative measures. Research in social and behavioural science suggests that the uncertainty caused by a pandemic promotes moral decision-making based on individual incentives instead collective incentives (Van Bavel et al., 2020; Gino et al., 2016; Garcia et al., 2020). Therefore, understanding what drives individuals can help predict or even steer collective behaviour, which is extremely important when attempting to mitigate the spread and impact of an infectious disease such as COVID-19.

In most agent-based models (ABM) the heterogeneity of moral decision-making is achieved by assigning agents cultural, economic and social characteristics in a similar distribution to the society being modelled and consequently use these to predict behaviour under certain circumstances or policies (Dignum et al., 2020; Gomez et al., 2020; Drogoul et al., 2020). This implicit modelling of moral decision-making can make it difficult to evaluate the (lack of) personal incentives to adopt certain preventative measures such as self-isolating behaviour. Therefore, this paper investigates the individual incentives behind moral decision-making in the COVID-19 pandemic by explicitly modelling behavioural strategies and measuring their relative success by individual quality of life.

The main research question of this work is: "What behavioural strategies can maximise an individual's quality of life in an agent-based social simulation of COVID-19?" This question is answered with help of three sub-questions. First, how successful is full-time isolating behaviour compared to non-isolating behaviour over time? Second, is moderate isolation preferable over full-time isolation behaviour? Third, does the success of isolation behaviour depend on the amount of isolation behaviour shown by the environment? These sub-questions are answered by applying behavioural strategies to the existing agent-based model *ASSOCC* by Dignum et al. (2020). Background to Agent-Based Modelling and the *ASSOCC* model in specific can be found in section 2.

The concrete contribution of this paper is to extend the existing *ASSOCC* ABM with explicitly modelled behavioural strategies and subsequently measure their relative success by resulting quality of life per strategy. The answers to the sub-questions are then combined to reach a general conclusion with regard to behavioural strategies in times of COVID-19. In the final two sections, limitations and reproducibility of the study are discussed, as well as possibilities for future work.

2 Background

In this section related work on modelling of infectious diseases is discussed, as well as background knowledge on agent-based modelling and the *ASSOCC* model in particular.

2.1 Modelling of Infectious Diseases

A whole range of different modelling techniques has been applied to COVID-19 and similar infectious diseases. The difference between the models lies mostly in the degree of heterogeneity in the model population. On the complete homogeneous side of the spectrum for example, Hellewell et al. (2020) created a stochastic transmission model to investigate the effectiveness of contact tracing and isolation of cases. They determined conditions for which these measures are sufficient to control an outbreak within 3 months, but the model is limited by the assumption that these measures are executed consistently by every individual. Compartmental models introduce some more heterogeneity, like how Hou et al. (2020) differentiates between four distinct population groups (susceptible, exposed, infectious, recovered) with their own particular behaviour. The groups are still internally homogeneous, however, and the model does not consider personal action taken such as social distancing or

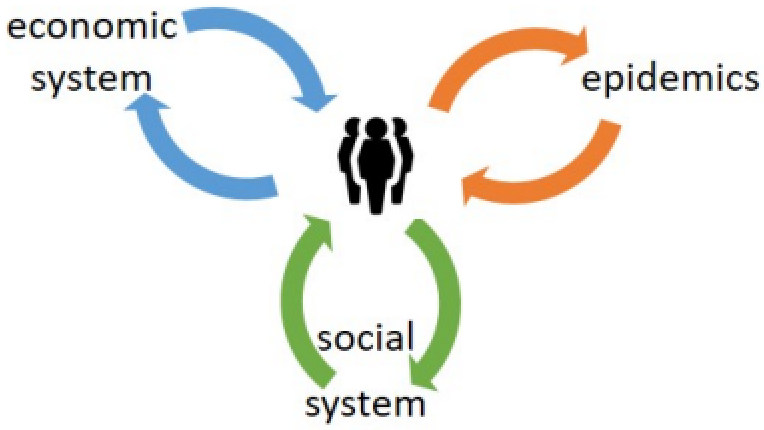

Figure 1: High-level view of *ASSOCC* from Dignum et al. (2020)

mask wearing. Taking it another step further, Gomez et al. (2020) uses Agent-Based Modelling (ABM) to simulate the population of Bogota, Colombia. In ABM every individual can be programmed with a unique set of desires or rules to monitor their interactions with the environment, making it the most heterogeneous technique in this list.

In general, ABMs are more reflective of complex interactions and individual decision making processes, where stochastic models are easier to understand and adapt. Both methods are extremely valuable to predict the effects COVID-19 and preventive measures, provided they are used appropriately.

2.2 ASSOCC

Agent-based Social Simulation of the COVID-19 Crisis (*ASSOCC*) is a complex ABM modelling tool "...that supports decision makers gain insights on the possible effects of policies, by showing their interdependencies, and as such, making clear which are the underlying dilemmas that have to be addressed." (Dignum et al., 2020, p. 3). Instead of attempting to make completely realistic mathematical models for all epidemic, economic and social processes combined, all these processes are modelled by agent-to-agent and agent-environment interactions, as illustrated in Figure 1.

2.2.1 Architecture

Using NetLogo (Tisue and Wilensky, 2004), *ASSOCC* simulates 300-2500 agents on a grid. These agents can autonomously decide to take actions based on characteristics and perceptions of the outside world, but their possible range of actions is constrained by their

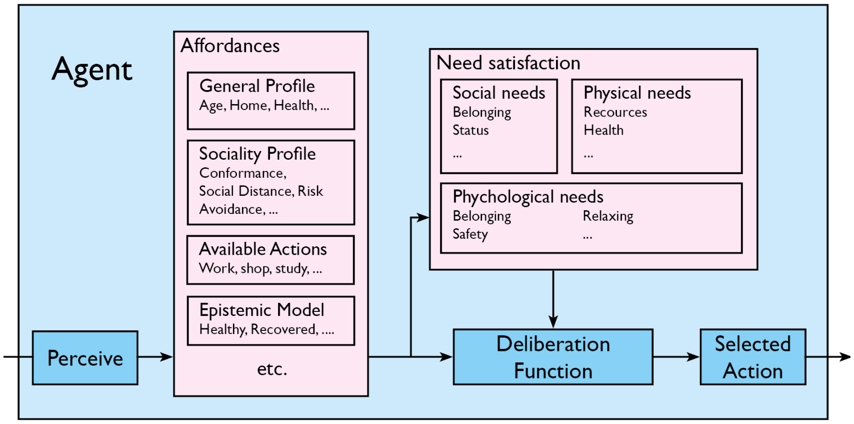

Figure 2: Schematic view of agent architecture by Dignum et al. (2020)

environment. These constraints from the environment can be physical, but can also come from social norms or regulations.

The architecture is made up of roughly four parts: agents, places, global functions and policies.

- **Agents** represent people. They have *needs* such as health, wealth and conformity, *capabilities* that are restricted by for example their age, jobs or family situation, and *characteristics* such as individualism, respect for authority and hedonism. A schematic overview of agents can be seen in Figure 2
- **Places** represent physical locations such as homes, shops, hospitals and workplaces that agents can move towards and perform actions in.
- **Global functions** implement epidemiological, social and economic processes, such as the contagiousness at different type of locations and economic rules like tax and subsidies.
- **Policies** are similar to global functions, but instead of modelling normal circumstances, they model government interventions that are triggered by events (e.g. more than half of the population is infected).

Both the the parameters that can be changed in the setup of a simulation and output from the simulation are visible on a NetLogo Tisue and Wilensky (2004) dashboard as can be seen in Figure 3.

3 Methodology

In this section we will go over definitions for the main research question: "What behavioural strategies can maximise an individual's quality of life in an agent-based social simulation of

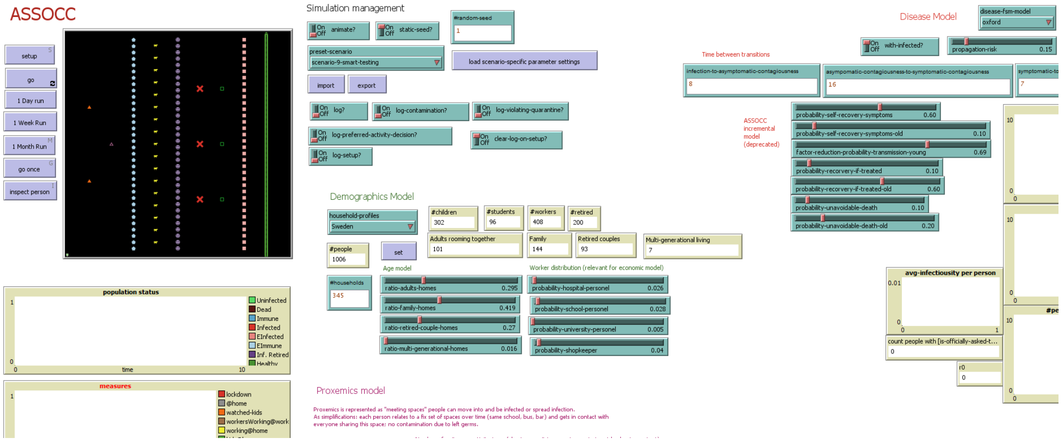

animate? Off
static-seed? On
#random-seed: 1
preset-scenario: scenario-9-smart-testing
log? Off
log-contamination? Off
log-violating-quarantine? Off
log-preferred-activity-decision? Off
clear-log-on-setup? On
log-setup? Off

household-profiles: Sweden
#children: 302
#students: 96
#workers: 408
#retired: 200
#people: 1006
Adults rooming together: 101
Family: 144
Retired couples: 93
Multi-generational living: 7
#households: 345
ratio-adults-homes: 0.295
ratio-family-homes: 0.419
ratio-retired-couple-homes: 0.27
ratio-multi-generational-homes: 0.016
probability-hospital-personel: 0.026
probability-school-personel: 0.028
probability-university-personel: 0.005
probability-shopkeeper: 0.04

Proxemics is represented as "meeting spaces" people can move into and be infected or spread infection.
As simplifications: each person relates to a fix set of spaces over time (same school, bus, bar) and gets in contact with everyone sharing this space; no contamination due to left germs.

disease-fsm-model: oxford
with-infected? On
propagation-risk: 0.15
infection-to-asymptomatic-contagiousness: 8
asympomatic-contagiousness-to-symptomatic-contagiousness: 16
symptomatic-to: 7
probability-self-recovery-symptoms: 0.60
probability-self-recovery-symptoms-old: 0.10
factor-reduction-probability-transmission-young: 0.69
probability-recorvery-if-treated: 0.10
probability-recorvery-if-treated-old: 0.60
probability-unavoidable-death: 0.10
probability-unavoidable-death-old: 0.20
count people with [is-officially-asked-t...: 0
r0: 0

Figure 3: Part of the *ASSOCC* NetLogo dashboard

COVID-19?". Then we will describe the experiments for the sub-questions step by step for reproducibility.

3.1 Definitions

Before it is possible to answer this question, the terms *behavioural strategies* and *quality of life* must be defined in a quantifiable way. In the interest of time only definitions within the existing model *ASSOCC* (Dignum et al., 2020) are considered, which is an agent-based model of COVID-19. A general overview of the *ASSOCC* model is given in Figure 1, with the particulars of the agent architecture in Figure 2: each individual agent represents a person that has a range of different needs and values, for example: health, sleep, risk avoidance and compliance. Each section of the day this agent chooses an action to take from the available ones, based on how much this action will satisfy their range of needs. The model has an extensive range of global processes to simulate real life - such as economic, cultural and governmental - that influence the affordances of the agents and are in turn influenced by the actions of the agents. Within the context of this model the the terms *behavioural strategies* and *quality of life* are defined in depth in the next two subsections.

3.1.1 Quality of Life

There are many ways to define quality of life, which makes it a tricky metric to work with. In the *ASSOCC* model, a composite value `quality-of-life-indicator` quantifies this abstract metric in a specific point in time by giving a weighted mean of all the need satisfaction levels of an agent at that time in the simulation. The respective weights assigned to the needs are decided by both a global need prioritisation and an individual one. For example, there is a global prioritisation of the need of survival, safety and belonging over the need of esteem, according to the Maslow Pyramid (McLeod, 2007). There can also be an individual prioritisation of needs such as luxury and leisure, due to the variation from agent to agent in the importance given to hedonism.

Since `quality-of-life-indicator` is a value that is specific to a certain point in time, in order to use it as a metric for quality of life, one must decide on which point(s) in time to consider. Since the need levels are added to and subtracted from over time, one could argue that it is an aggregate value, and that looking at the final value is sufficient. However, for the purposes of this research it is more informative to look at the quality of life over time, and summarise if necessary (for example when showing relation to different strategy ratios) by taking the mean of all `quality-of-life-indicator` values over time. This will take into account "rough patches" in an individual's life as well, and not just the quality of life at the end of the simulation.

3.1.2 Behavioural Strategy

The (composite) agent attribute that represents a behavioural strategy should:

1. (indirectly) cause or prevent behaviour that (indirectly) influences the quality of life of agents and
2. be represented by a simple value (boolean, numerical, or string) from a finite set of possibilities.

The second requirement is mostly in the interest of time, since, realistically not all values between 0 and 100 can properly be investigated individually, but investigating the effect of a simple yes/no is much more achievable.

With the help of those criteria, the following promising (composite) agent attributes from *ASSOCC* are considered: ratio of time spent in social distance, likeliness to self-quarantine when symptomatic and preventive isolation. In consideration of future research we will go over definitions for all three options (a summary of findings can be found in Table 1).

- **Ratio of Time Spent in Social Distance** is a metric that results from the agent attribute `is-I-apply-social-distancing?` by the formula given below.

  $$\text{Ratio of Time Spent in Social Distance} = \frac{\#\text{ticks where } \texttt{is-I-apply-social-distancing?}}{\#\text{ticks}}$$

  Only after a full simulation is run for each agent this metric can be determined, and then the agents with a similar ratio can be grouped together to see the mean quality of life of this group. Some example results of this can be found in Figure 4a. The biggest issue for this metric is that the variable `is-I-apply-social-distancing?` is merely descriptive of past behaviour, and does not cause or prevent anything, therefore not satisfying requirement 1. Moreover, the metric can be any rational value between 0 and 100, thus also not satisfying requirement 2.

- **Likeliness to Self-Quarantine when Symptomatic** is a global variable represented by a value between 0 and 100. It being a global variable means that agents do

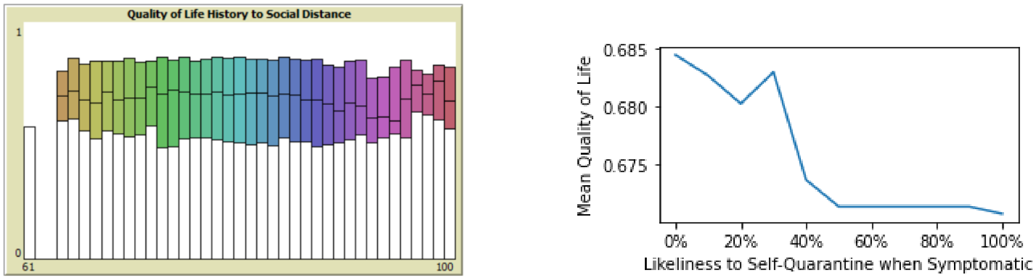

(a) Quality of life against ratio of time spent in social distance.

(b) Quality of life against likeliness to self-quarantine when symptomatic.

Figure 4: Example results for different agent attributes.

Table 1: Summary of possible strategy definitions

Attribute	**Causation**	**Representability**
Ratio of time spent in social distance	✗	✗
Likeliness to self-isolate when symptomatic	✓	✗
Preventive Isolation	✓	✓

not have a predetermined strategy, but every day and infected agent self-quarantines with a certain (global) likeliness. An example output of the relation between this metric and the quality of life can be seen in Figure 4b. This homogeneity of the metric is not ideal for the purpose of this investigation, but this could be remodelled to ratio of people that are predetermined to self-quarantine when symptomatic. Similar to the previous metric, this would not be ideal by the second requirement, but it does satisfy the first: the metric influences the actions of the agents and resulting quality of life.

- **Preventive Isolation** Preventive Isolation is not directly modelled in the *ASSOCC*, except for in the need of risk-avoidance being best satisfied when staying home. My approach to this is to assign each agent a preventive isolation strategy, that dictates whether or not to self-isolate preventively. Two simple strategies are implemented: Ignore and Isolate. The Ignore strategy changes nothing in the agent architecture shown in Figure 2, while the Isolate strategy adds a new step to the deliberation function. For any possible action that takes place outside of the agent's home, we subtract a global `isolation-tendency` value from their expected need satisfaction. Thus, the higher the `isolation-tendency`, the more an agent will be more inclined to do actions at home instead of outdoors. Once the infection has died out and no agent is infected anymore, the isolation behaviour is dropped.

As can be seen in Table 1 where the different possible strategy definitions are summarised, the most suitable attribute defined is preventive isolation strategies, which is why it will be the only behavioural strategies discussed in the rest of the paper.

3.2 Experiment Setup

In order to answer the sub-questions, we run several experiments in the adapted *ASSOCC* model that can be exactly replicated using the corresponding BehaviorSpace (Tisue and Wilensky, 2004) experiments provided in the codebase. The ASSOCC model is mostly left untouched, except for where the original deliberation function is influenced by the newly modelled behavioural strategies. Every unique combination of parameters is run 5 times in the default *ASSOCC* environment (294 people, 100 households, Swedish demographic and cultural model). The attributes that vary per experiment are summarised in Table 2.

3.2.1 Quality of Life of Full-Time Isolating Behaviour over Time

To answer the question "How successful is full-time isolating behaviour compared to non-isolating behaviour over time?", we compare Isolate and Ignore by comparing the mean `quality of life indicator` of all agents with that strategy. The distribution of the strategies over the agents is 50/50. To make sure agents can reap the benefits of an disease-free environment if they manage to make a swift end to the epidemic, we let every simulation run for at least 4 months and until no agent is infected anymore.

3.2.2 Relative Success of Moderate Isolation over Full-time Isolation Behaviour

The experiments we will do to answer the question "Is moderate isolation preferable over full-time isolation behaviour?" are very similar to the first experiment, except that we will be comparing the success of different degrees of isolation. To achieve this the `isolation-tendency` parameter will be set to different values in the range of 0.0 (no isolation behaviour) to 0.5 (full isolation behaviour).

3.2.3 Influence of Other Agent's Strategies on Quality of Life of Isolators

For the third question, "Does the quality of life of agents with isolation behaviour depend on the amount of isolation behaviour shown by others?", we test the success of full-time isolation (i.e. isolation tendency of 0.5) in different strategy distribution settings. The strategy distribution in the model is controlled by a parameter called `ratio-isolate-of-basic-strategies`, which ranges from 0% to 100%. Since we want to measure the quality of life of both Isolate and Ignore in all simulations, we avoid the extreme values, and work only with a range of 1% to 99%. All parameter values that are not specifically mentioned in this experiment description are the same as in Experiment 1 (see Table 2).

4 Results

In this section we discuss the results of the experiments described in Section 3 and summarised in Table 2. After, we answer their respective sub-questions.

Table 2: Overview of experiments.

	% with Strategy Isolate	**Isolation Tendency**
Experiment 1	50%	0.5
Experiment 2	50%	{0.025, 0.05, 0.075, 0.1}
Experiment 3	{1%, 25%, 50%, 75%, 99%}	0.5

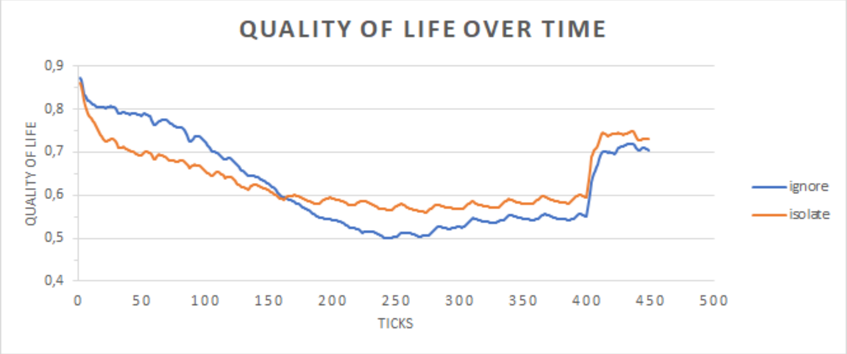

Figure 5: Success of full-time isolating behaviour over time: 50/50 division of strategies and isolation-tendency value of 0.5

4.1 Quality of Life of Full-Time Isolating Behaviour (Experiment 1)

To answer the question “How successful is full-time isolating behaviour compared to non-isolating behaviour over time?”, we look at the quality of life over time per strategy in the default Swedish environment with a 50/50 division of Isolate/Ignore. The simulation is run 5 times to divide the strategies differently over the agents each time (and thus a different social/financial situation). As we can see in Figure 5, there is no obvious superior strategy in this environment, since both strategies have the upper hand for only part of the simulation. Taking the average over time for both strategies gives little extra insight: 0.624 and 0.632 for Ignore and Isolate respectively. This relatively small difference cannot be used to conclude that full-time isolation behaviour is more successful than no isolation behaviour. However, what we can do with these results is identify two key points in the graph that can be used in the next two sub-questions:

1. A turning point at which those with strategy Isolate start outperforming Ignore (in this experiment at ±150 ticks).
2. A spike in quality of life for both strategies when the virus dies out (in this experiment at ±400 ticks).

In the next two experiments we see how these two points are influenced by moderate behaviour and different strategy distributions.

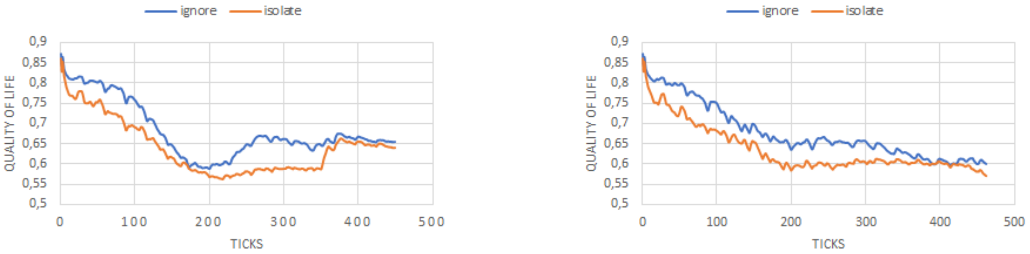

(a) Isolation tendency value of 0.025

(b) Isolation tendency value of 0.05

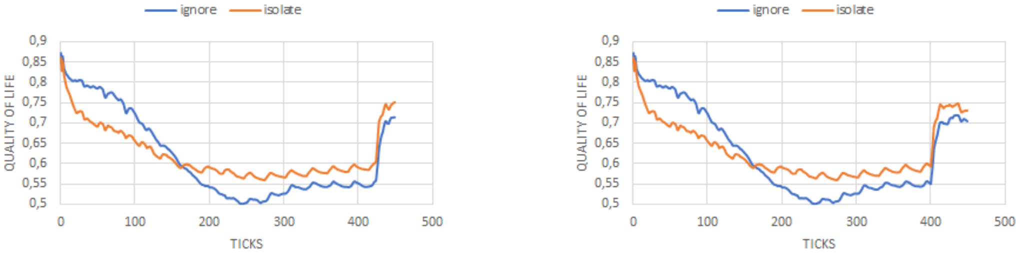

(c) Isolation tendency value of 0.075

(d) Isolation tendency value of 0.1

Figure 6: Quality of life over time per strategy for different isolation tendency values

4.2 Result of Moderate Isolation Behaviour (Experiment 2)

Now that we have seen the pattern of how a full-time isolation strategy pays off compared to regular behaviour, we experiment with different degrees of isolation. Since an isolation-tendency of 0.0 causes Isolate agents to act exactly the same as Ignore agents, we ignore this setting and start the experiments at 0.025. As can been seen in Figure 6a, when isolation tendency is as small as 0.025, the quality of life of agents with strategy Isolate is unquestionably worse, never showing the turning point identified in Experiment 1. The same is true for 0.05, shown in Figure 6b. Somewhere between values 0.05 and 0.075 however, the long-term benefits of Isolate start to appear: the turning point around 150 ticks starts appearing again, as does the spike at 400 ticks caused by being disease-free.

The straightforward answer to “Is moderate isolation preferable over full-time isolation behaviour?” is no, it is not. Moderate isolation seems to be the least successful when compared to full-time and no isolation. An interesting additional observation is that isolation tendency does not seem to influence the timing of the two points identified in Experiment 1, just their existence: the of quality of life over time with any isolation tendency ≥ 0.075 shows the turning point at ±150 ticks and disease-free spike at ±400 ticks consistently, and any simulation with isolation tendency ≤ 0.075 does not show the points at all.

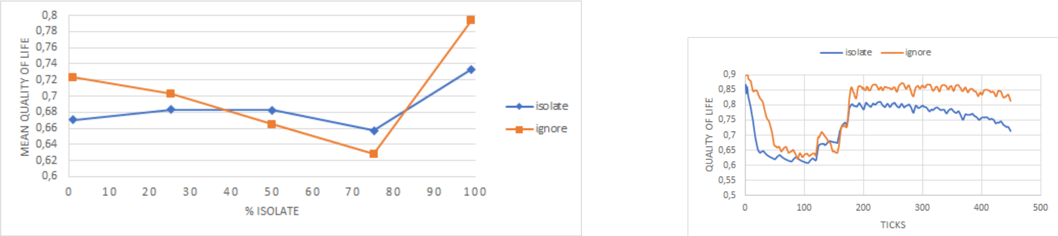

(a) Quality of life per strategy for different strategy distributions

(b) Quality of life over time when 99% of agents has strategy Isolate

Figure 7: Quality of life over time per strategy for different distributions with constant isolation tendency of 0.5

4.3 Influence of Other Agent's Strategies (Experiment 3)

In order to investigate the influence of different strategy distributions on the success of isolation behaviour, we run experiments similar to Experiment 1, only with varying distributions of isolation strategies. The model parameter `ratio-isolate-of-basic-strategies` represents this distribution, and is set to 1, 25, 50, 75 and 99 sequentially. To limit the search space for more concrete results, only full-time isolation (i.e. isolation tendency of 0.5) is considered. As shown in Figure 7a, the maximum achievable quality of life for both strategies lies at the same distribution: 99% isolation behaviour. When we look at the quality of life over time for this distribution (as shown in Figure 7b), it is easily discovered why: the lifespan of the virus is very short in this scenario. Around 150 ticks already, the environment is virus-free and regular life can be picked up for the remaining time of the simulation.

What these results show is that isolation behaviour is very costly, and only if done collectively the benefits outweigh the costs. There is an interesting catch to this ideal situation: everyone's personal best interest is that the collective situation is in the 99% isolation situation, and yet their own optimal strategy is Ignore in that situation. Thus, if everyone chooses their strategy based on the highest possible personal reward, the collective behaviour is more likely to end up in a sub-optimal situation of unanimous non-isolation, even though they would have been better off being in the 99% of isolators.

5 Discussion and Responsible Research

This section contains a reflection on the limitations of the modelling technique used, the ethical aspects of the research done and the reproducibility of the experiments.

5.1 Limitations of Modelling Technique

The consequences and correctness of the conclusions drawn in the Results section are limited to the correctness of the underlying *ASSOCC* model and how reflective the underlying mode is of real life. This limitation is inherent to using this ABM (Agent-Based Model) as a backdrop for experiments, since this ABM was never meant to answer quantitative questions such as the precise amount of people who will be infected in Sweden by July 2020, or how much money the crisis will cost the government exactly. Its purpose is rather to test whether certain patterns emerge and relations between parameters can be discovered, which is exactly the purpose of this research as well. It is meant to investigate the relation between an agents own well-being and the isolation behaviour of itself and others.

5.2 Ethical Aspects

Since COVID-19 and isolation behaviour can have far-stretching consequences, some extra ethical considerations were necessary when writing this report. The limitations of modelling techniques illustrated in the previous sub-section are not always repeated clearly in every section individually to avoid redundancy, which could lead to some results being misunderstood or taken out of context and to be used as motivation against advice given by health professionals or government. However, these experiments are merely conducted to show how collective behaviour can be steered by individual motivation in a synthetic environment similar to the COVID-19 pandemic, and not to make any suggestions about best practice in times of a pandemic. The word "agent" instead of "person" is used in every method description or result to repeatedly show that it is limited to modelled behaviour and does not necessarily extend towards real life.

5.3 Reproducibility of Experiments

Every experiment described in the Method section can be reproduced by using the corresponding BehaviorSpace (Tisue and Wilensky, 2004) experiments that are included in the NetLogo (Tisue and Wilensky, 2004) model (accessible by request). The data used to answer sub-questions 1, 2 and 3 can be reproduced by running "experiment-1", "experiment-2" and "experiment-3" respectively. The experiments are run with a different random-seed each time, which does hurt the exact reproducibility of the results, but since all the results are averages of 5 runs with different random-seeds, the reproducibility should be sufficient.

6 Conclusions and Future Work

This section contains a summary of the research questions, experiments and results and an overview of open questions that are fit for future work.

6.1 Research Questions and Results

The main research question "What behavioural strategies can maximise an individual's quality of life in an agent-based social simulation of COVID-19?" is answered through three sub-questions, all investigating different patterns in quality of life: (1) what it looks like over time, (2) how it is influenced by the degree isolation and (3) how it is influenced by the distribution of isolation strategies.

To answer the question “How successful is full-time isolating behaviour compared to non-isolating behaviour over time?”, we ran several *ASSOCC* simulations in the default environment with even distribution of isolating and non-isolating agents and compared their respective quality of life to measure success. The results show a pattern where those who do not isolate seem happier at first, while after a turning point, the isolators gain the upper hand instead. Also, we see that both strategies benefit from the end of the virus strongly and equally.

The experiment corresponding to the second sub-question (“How successful is full-time isolating behaviour compared to non-isolating behaviour over time?”), repeated the first experiment, but with a range of different degrees of isolation. It shows that when applying a very moderate type of isolation (still performing activities outside if they are expected to be very satisfactory) there is no turning point at which the cost of isolation starts paying off, nor is there a quick end to the virus. This implies that isolation is only a preferable or even acceptable strategy when doing full-time isolation.

The last experiment answers the question “Does the success of isolation behaviour depend on the amount of isolation behaviour shown by others?” and repeats the first experiment with different strategy distributions each time. It shows that the highest possible pay-off for both the Isolate and Ignore strategy is achieved when 99% of the agents isolates due to having a long virus-free period afterwards. However, in this situation, the pay-off for isolators, although the highest possible value for this strategy, is still lower than that of the ignorers. This means that optimal situation for every individual is almost impossible to achieve unless people settle for a slightly lower pay-off and choose the Isolate strategy collectively.

The two main conclusions that can be drawn towards the main research question are (1) isolation behaviour can be an acceptable or even preferable strategy if and only if it is done full-time and collectively and (2) collective isolation (although the preferable situation) can be difficult to achieve if every individual chooses their strategy only according to the highest possible pay-off.

6.2 Future Work

There are two possible avenues to take this research: applying the modelled behavioural strategies or extending them. We will go over my suggestions for both options separately since they are both valuable and can be done independently.

6.2.1 Application of Model

An interesting possible research question that can be answered using this model is how government restrictions influence the success of behavioural strategies. Many different government restrictions and policies are already modelled in *ASSOCC* and it would be interesting to see if a government could steer or motivate people to isolate by understanding the personal motivations involved. Another interesting research subject would be looking at the satisfaction values that make up the quality of life. Which satisfaction levels are the main cause for fluctuations seen in the results of this paper? And which of these are reflective of real life and which of these are caused merely by technical constraints/misjudgements of the model?

Another interesting research topic would be to see if the conclusions of this paper can also be drawn if the experiments are replicated with a different underlying model than *ASSOCC*, for example with the rather similar COMOKIT model Drogoul et al. (2020). The strategies added in this paper are straightforward and kept in a separate file from the original model, so they are well suited for being recreated in a different model.

6.2.2 Extension of Model

When extending the model, there is more to do than improving the underlying *ASSOCC* model (which is still a work in progress at the time of publishing this paper). It would be valuable to look into making the behavioural strategies more nuanced and realistic. Currently, the isolation strategy is very straight-forward: anything outside is bad and the agent does not do it. This means that households who isolate will lock the doors until the epidemic is over, despite not having eaten in weeks. This unrealistic behaviour could be remodelled by creating a subdivision within the actions and treating actions that are absolutely necessary for survival differently, even if they are outside of the house.

Another option for making behavioural strategies more realistic is making them dynamic. In real life, if you are bankrupt and/or starving, it is likely that you will make some changes to your life. Therefore it would make sense to let people change their strategy if their survival is threatened or if they are very unhappy. Once this is implemented it would be interesting to see if everyone ends up with the same strategy, which strategy this is, and how the initial distribution of strategies at setup influences this outcome.

References

Dignum, F., Dignum, V., Davidsson, P., Ghorbani, A., van der Hurk, M., Jensen, M., Kammler, C., Lorig, F., Ludescher, L. G., Melchior, A., et al. (2020). Analysing the combined health, social and economic impacts of the corovanvirus pandemic using agent-based social simulation. *arXiv preprint arXiv:2004.12809*.

Drogoul, A., Taillandier, P., Gaudou, B., Choisy, M., Chapuis, K., Quang Nghi, H., Doanh, N., Philippon, D., BrugiÃšre, A., and Larmande, P. (2020). Designing social simulation to (seriously) support decision-making: Comokit, an agent-based modeling toolkit to analyze and compare the impacts of public health interventions against covid-19. *RofASSS*.

Garcia, T., Massoni, S., and Villeval, M. C. (2020). Ambiguity and excuse-driven behavior in charitable giving. *European Economic Review*, page 103412.

Gino, F., Norton, M. I., and Weber, R. A. (2016). Motivated bayesians: Feeling moral while acting egoistically. *Journal of Economic Perspectives*, 30(3):189–212.

Gomez, J., Prieto, J., Leon, E., and Rodriguez, A. (2020). Infekta: A general agent-based model for transmission of infectious diseases: Studying the covid-19 propagation in bogotá-colombia. *medRxiv*.

Hellewell, J., Abbott, S., Gimma, A., Bosse, N. I., Jarvis, C. I., Russell, T. W., Munday, J. D., Kucharski, A. J., Edmunds, W. J., Sun, F., et al. (2020). Feasibility of controlling covid-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health*.

Hou, C., Chen, J., Zhou, Y., Hua, L., Yuan, J., He, S., Guo, Y., Zhang, S., Jia, Q., Zhao, C., et al. (2020). The effectiveness of the quarantine of wuhan city against the corona virus disease 2019 (covid-19): well-mixed seir model analysis. *Journal of medical virology*.

McLeod, S. (2007). Maslow's hierarchy of needs. *Simply psychology*, 1:1–8.

Tisue, S. and Wilensky, U. (2004). Netlogo: A simple environment for modeling complexity. In *International conference on complex systems*, volume 21, pages 16–21. Boston, MA.

Van Bavel, J. J., Baicker, K., Boggio, P. S., Capraro, V., Cichocka, A., Cikara, M., Crockett, M. J., Crum, A. J., Douglas, K. M., Druckman, J. N., et al. (2020). Using social and behavioural science to support covid-19 pandemic response. *Nature Human Behaviour*, pages 1–12.