A data-driven agent-based model of the SARS-CoV-2 epidemic in the Netherlands

A Master Thesis for the Degree of
Economics of Markets and Organisations

by

P.A.J. Versfelt
375614

Supervisor: Dr. SV Kapoor
Second assessor: Dr. V Karamychev

Erasmus School of Economics
ERASMUS UNIVERSITY ROTTERDAM

June 2021

1 Introduction

A novel coronavirus (so-called SARS-COV-2 or COVID-19) emerged in Wuhan, China at the end of 2019 (Chen et al., 2020). Cases quickly spread to other countries as the outbreak reached global proportions. On 11 March 2020, the World Health Organization declared the outbreak as a pandemic (WHO, 2020b). Many countries' healthcare systems were not prepared for outbreaks of such a scale, causing healthcare systems to operate at maximum capacity and not being able to provide adequate care (Kandel et al., 2020). At the end of 2020, after nearly a year after the beginning of the pandemic, roughly 1.8 million people have reportedly died by COVID-19 worldwide (WHO, 2020a). During this time, countries reacted by implementing non-pharmaceutical interventions (NPIs) to combat the rapid spread of the virus. However, the effectiveness of NPIs vary and thus a need arises to evaluate NPI strategies to avert future health crises.

On 12 March 2020, the Dutch government introduced the first NPI by banning gatherings of more than 100 people. On subsequent days, the government introduced more stringent NPIs and on 23 March, the government initiated a self-proclaimed 'intelligent' lockdown, a set of additional NPIs (RIVM, 2020). In this lockdown, all gatherings were canceled, all citizens were asked to stay home as much as possible and adhere to health regulations. These NPIs reduced the number of contacts between people and lessened the pressure on the Dutch healthcare system, and on 11 May some NPIs were relaxed. Given the complexity of society, the prediction of the impact that interventions strategies have on a potential epidemic is not a trivial task. For this purpose, policy makers can be informed by model-based predictions of the impact of different NPIs on the course of the epidemic.

Computational models can provide insights and evaluate effectiveness of interventions (Furguson et al., 2020). In epidemiology, the two main approaches are equation-based model and agent-based models. Equation-based models simulate the spread of a disease on a population level, while agent-based models simulate interactions and behavior of a population at an individual level.

In this paper, I develop and present results of a stochastic agent-based model (ABM) of the COVID-19 epidemic in Dutch municipalities. The model can be applied to the Netherlands as a whole, however due to computational limitations, it is implemented for two municipalities: Delft and Apeldoorn during a two-month time span. The model is calibrated using observed data on hospitalizations and validated on number of infections, hospitalizations and deaths. Hereafter, the impact of various NPIs on the course of the epidemic is evaluated in terms of hospitalizations. Finally, a sensitivity analysis is performed to better understand the ABM dynamics and address potential limitations.

The proposed model calibrated and validated the number of hospitalized patients for both municipalities well. While the model was not trained on deaths, it could explain the simulated deaths well for both municipalities. However, tests showed a significant difference between the model-predicted and observed age distribution of diseased people in hospitals. In line with expectations, the model overestimated the observed number of infected individuals. This difference is likely caused by an under-representation of asymptomatic individuals in the observed data. The model shows that NPIs targeting children have a large impact on the spread of cases in the two municipalities studied. Furthermore, the model showed that a delay on enacting NPIs resulted in a large spread of the disease in the Dutch municipalities, showing that implementing NPIs early is of great importance. The sensitivity analysis revealed, by varying values of parameters with uncertainty, that population features (e.g., classroom size and social network size) did not have a large impact on the outcome of the model. While uncertainties concerning the transmissibility of the disease (e.g., symptomatic or asymptomatic duration) affected the model outcome substantially.

The main contributions of the paper are the following. First, I contribute to a literature that examines COVID-19 NPIs. Numerous models have evaluated various NPIs and showed that school closure and early NPIs are effective measures (Soltesz et al., 2020), (Flaxman et al., 2020). My results support these findings for Dutch municipalities. Second, this paper contributes to a literature that examines agent-based models concerning diseases in

Europe. ABMs exist for epidemic spreading in Switzerland (Hackl and Dubernet, 2019), France (Hoertel et al., 2020), Italy (Ciofi degli Atti et al., 2008), (Merler et al., 2009), (Sjödin et al., 2020) and Ireland (Hunter and Kelleher, 2021). However, no ABM is yet developed for any infectious diseases in the Netherlands. Thus, this paper provides a first overview on available data for ABMs in Dutch municipalities. Furthermore, as NPIs impact countries differently due to social and cultural differences (WHO, 2019), this paper analyzes the effectiveness of NPIs specifically in Dutch municipalities. Third, the proposed model differs from the above-mentioned studies, in that data of 2.260 unique individuals' activity patterns of Dutch citizens are included in the model. An activity patterns dictates what activity individuals execute at different time periods for each day (e.g., going to work at 9AM or going to bed at 10PM). Most data-driven models on COVID-19 do not include empirical information about human activity patterns, but rather use aggregated activity patterns. Which limits the external validity, as infection rates are highly dependent on patterns of individuals (Wang et al., 2021).

The rest of the paper is organized as follows: Section 2 provides a brief overview of related work. Section 3 describes the model and calibration process. Section 4 details the synthetic geography and population generation process. Section 5 discusses the calibration and validation results. In section 6, the sensitivity analysis results are presented and finally in Section 7, the paper is concluded and limitations and future directions are discussed.

2 Related work

In epidemiology, the most famous approach to model the spread of epidemics is the equation-based model (EBM) proposed by Kermack and McKendrick (1927). Generally, EBMs for diseases assume that a population (of size N) is compartmentalized into a small number of groups. These separated groups represent various phases of a disease, and a set of equations specify the transition rates between the compartmentalized health states. Commonly used

phases are those that can contract the disease (called susceptible, denoted by S), those that are infected and thus contagious (infectious, I) and those who are recovered from the disease (recovered, R). The Susceptible-Infected-Susceptible (SIS) model by Iannelli et al. (1992) only contains two transitions, either from susceptible to infected, or the other way around. Another model by Shulgin (1998) is the Susceptible-Infected-Recovered (SIR) model, where individuals cannot be infected more than once by the disease. The Susceptible-Exposed-Infectious-Recovered (SEIR) model by d'Onofrio (2002) extends the standard model by including a phase where an individual was exposed to the disease, while not being infectious (the so-called incubation period). The SIS, SIR, SEIR models form the basis of EBMs for diseases, other models are proposed that include or exclude variations of the before mentioned disease phases.

EBMs fail to capture the heterogeneous mechanisms that are present in the spread of a disease. For instance, EBMs, generally, cannot model spatial heterogeneity, contact patterns and individual behaviors (Di Stefano et al., 2000). ABMs however, are capable of handling these mechanisms (Hunter et al., 2018). For this reason, and the increase of computational power available, the ABM framework has become popular for disease modeling. ABMs can include spatial (e.g., locations such as schools, offices and homes) and population characteristics (e.g., activity patterns, employment rate) and offer a high-resolution representation of a population (Bonabeau, 2002).

In ABMs, spatial aspects of epidemics can be studied by permitting interactions between individuals or "agents". Generally, ABMs consist of a population, an environment and a set of rules that dictate the actions of individuals (Epstein et al., 1996). In agent-based approaches for epidemiology, health states of individuals obey compartmentalized frameworks, such as SEIR. However, unlike EBMs, changes in health states of individuals in agent-based approaches come from interactions between individuals. As ABMs track contacts of every individual in physical locations, the transmission of the disease follows through these contacts. Actions of agents are determined by their surroundings and a set of rules. The rules

that dictate the movement of individuals (e.g., going to work/school, visiting friends) and the transmission of the disease are explicitly modeled (Gilbert, 2005).

The purpose of existing ABMs, for epidemiology, is mostly to analyze the behavior of a system under certain conditions or to evaluate the impact of policies to slow the spread of a disease in an epidemic. ABMs have been implemented for many diseases (e.g., influenza, HIV or Ebola) and in many countries (Willem et al., 2017). Furthermore, ABMs have been applied in various fields, such as flows (e.g., evacuation, traffic), ecology, economics, and finance (Wallace et al., 2015).

Since the start of the COVID-19 pandemic, the scientific community had a fast response by dedicating studies using models to focus on treatment and forecasting. Most studies using EBMs implemented SEIR-based approaches, as COVID-19 has a known incubation period (McAloon et al., 2020). Some studies added new phases to the SEIR approach, for instance: isolated, hospitalized, or asymptomatic (He et al., 2020), (Tang et al., 2020), (Wu et al., 2020). Several ABMs are implemented to explore the impact of interventions and on the epidemic spread of COVID-19. Some studied the impact of NPIs, such as Furguson et al. (2020) in GB and USA, Hoertel et al. (2020) in France, Silva et al. (2020) in Brazil. Others studied contact tracing with ABMs, such as Shamil et al. (2021), Truszkowska et al. (2021) and Kerr et al. (2020) in USA and Bicher et al. (2020) in Austria. These studies show that ABMs calibrate and validate well for the COVID-19 pandemic.

3 Model description

In this section, the agent-based model is further elaborated. First, the transmission dynamics for COVID-19 in the model are described. Hereafter, the disease model is discussed. Finally, the process of simulating the COVID-19 spread is explained.

3.1 Disease transmission

In the model, a set of $A = \{1, ..., n\}$ agents are able to move between locations of set $L = \{1, ..., m\}$. Each location $l_q \in L$ has a corresponding location type q. Appendix 8.1 denotes all location types. Interactions between individuals arise when multiple individuals are present at locations (e.g., schools, households or workplaces). The probability that any susceptible individual i becomes infected when residing at a location of type q at any time t, is equal to

$$p_i(t) = 1 - e^{-\Delta(t)\lambda_q(t)}, \tag{1}$$

based on the paper by Ajelli et al. (2010). Where $\Delta(t)$ denotes the time spent in location l_q. Furthermore, $\lambda_q(t)$ specifies the infection risk associated with location l_q at time t. Note that $p_i(t) \to 1$ as $\Delta(t) \cdot \lambda_q(t) \to \infty$. Thus, individual i is more likely to get infected if the individual resides longer at a certain location or if the risk associated with that location is higher.

The parameter $\lambda_q(t)$ specifies the infection risk at the associated location of type q (e.g., household, school, or workplace) and is equal to

$$\lambda_q(t) = \frac{1}{n_q} \sum_{k=1}^{n_q} (E_k + Sy_k \cdot c) \cdot P_k \cdot \beta_q \cdot \gamma \cdot \theta_{t,q}. \tag{2}$$

The infection risk is calculated by performing a sum over all (n_q) individuals present at time t at the corresponding location l_q. E_k is an indicator variable that is equal to 1 if individual k is at time t in the exposed (E) health state, and 0 otherwise. If individual k is symptomatic, Sy_k is equal to 1 or 0 otherwise. The parameter P_k depicts the infectiousness of individual k. The infectiousness varies among individuals and is drawn from Gamma(0.25, 4) distribution with mean 1 (Furguson et al., 2020). The infectiousness difference between symptomatic and exposed individuals is captured in the term c. Current research estimates that symptomatic individuals are more infectious compared to

	Home	Work	Healthcare	Shop	Religious	School	Leisure	Public transport	Horeca
β_q	0.8	1	1	1.5	1.5	1.5	1	1.5	3

Table 1: Location dependent transmissibility factor.

asymptomatic individuals. However, there exists some heterogeneity in the estimates of the relative infectiousness (McEvoy et al., 2021). Therefore, the term c is assumed to be equal to 2, in line with the recommended studies by McEvoy et al. (2021). The relative infectiousness estimate of 2 indicates that symptomatic individuals are twice as infectious relative to asymptomatic individuals. The sensitivity of c will be analyzed in the sensitivity analysis. To capture differences between activities at locations, each location is assigned a location type dependent transmission factor (β_q) based on Rădulescu et al. (2020). These location transmission factors are depicted in Table 1.

The infection risk, $\lambda_q(t)$, represents the risk of infection that each susceptible individual is exposed to, while residing at a location l_q at time t. Each susceptible individual present at location l_q at time t experiences the same risk, as $\lambda_q(t)$ approximates the chance of interactions between susceptible and infectious individuals at locations. Furthermore, each susceptible individual i is subject to the same probability of infection $p_i(t)$, as recent studies suggest that no large differences of susceptibility by age exist for COVID-19 (Omori et al., 2020).

The parameters β_q, P_k and c of the infection risk are not calibrated, as estimates for these parameters exist in the literature on COVID-19 and the inclusion of these parameters can lead to overfitting the model. However, two parameters of the infection risk are calibrated by the model: γ and $\theta_{t,q}$. By calibrating the parameter γ, the model is able to adjust the infection risk caused by infectious individuals and thus alter the probability of infection that susceptible individuals experience when in contact with infectious individuals. The second calibrated parameter, $\theta_{t,q}$ is included to captures the change in transmissibility caused by the implementation of a NPI on 14 October. From this date onward, the Dutch government forced individuals to wear face masks in schools and public transport, reducing the conta-

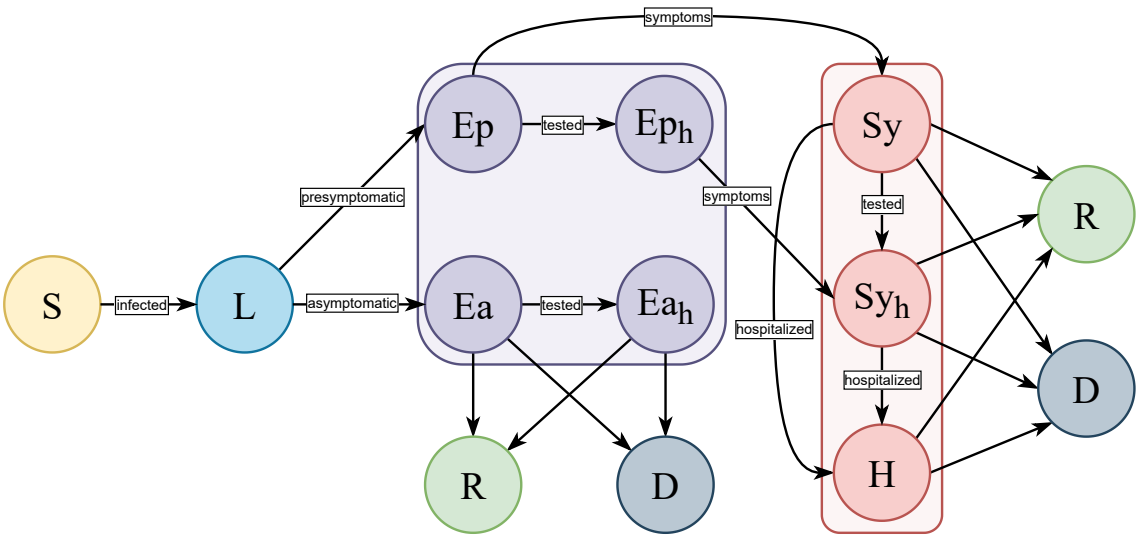

Figure 1: Schematic representation of health states and its transitions for individuals.

giousness of individuals in these locations. In the model, $\theta_{t,q}$ is equal to 1 before 14 October, equal to the calibrated value for subsequent time periods for schools and public transport locations or 1 otherwise.

3.2 Disease model overview

The disease model adopts a compartmentalized standard SEIR transmission dynamics. Additional states are included as COVID-19 is known to have a latent period and cause asymptomatic illness, studied in meta-analysis studies by McAloon et al. (2020) and Byrne et al. (2020). The following states are included in the model: susceptible (S), latent - where an individual has been infected, but is not yet infectious (L), exposed asymptomatic - an individual is infectious but does not develop symptoms (Ea), exposed presymptomatic - where an individual is infectious, but has not yet developed symptoms (Ep), symptomatic infectious (Sy), and permanently recovered (R) or dead (D) individuals.

In the model, exposed and symptomatic individuals can get tested on their own accord. The probability of getting a test is based on the average willingness (45%) of Dutch citizen in the period 30 September till 4 October as based on the research of RIVM (2021e). All asymptomatic individuals that receive a positive test result, will be placed into home

	Value	Reference
Transition duration		
Latent period, days	Lognormal(1.63, 0.5)	(McAloon et al., 2020)
Presymptomatic period, days	2.5	(Byrne et al., 2020)
Sympotmatic period, days	7.5	(Byrne et al., 2020)
Asymptomatic period, days	Gamma(3, 2)	(Byrne et al., 2020)
Testing parameters		
Time for PCR result, days	3	(RIVM, 2021d)
Test willingness symptoms	45.0%	(RIVM, 2021e)
Test willingness housemate positive	83.5%	(RIVM, 2021e)

	Age-dependent value					
	0-14	15-24	25-44	45-64	65+	Reference
Probability						
Developing symptoms	10%	30%	50%	50%	50%	(Furguson et al., 2020)
Hospital admission	0.30%	0.51%	1.45%	5.21%	20.00%	(Papst et al., 2021)
Death	0.00%	0.01%	0.03%	0.48%	5.42%	(Papst et al., 2021)
Duration, days						
Time until hospitalization	2	6	6	6	4	(CDC, 2021)
Hospitalization	2	3	3	4	5	(CDC, 2021)

Table 2: Parameter values and references of disease characteristics.

isolation (Ea_h) (RIVM, 2021e). Furthermore, if an individual is hospitalized or receives a positive test result, its housemates can get themselves tested (83.5% do so). Thus, presymptomatic individuals can get tested if their housemate received a positive test result. These presymptomatic individuals will be placed into home isolation (Ep_h) after receiving its result. If an individual develops symptoms, while already in home isolation, it is transferred to the symptomatic home isolation state (Sy_h). Finally, some symptomatic individuals develop such severe symptoms that they need to be hospitalized (H). An overview of health states and transitions are shown in Figure 1.

All infected individuals start their latent period after infection occurred. When the latent period has expired, the individual progresses into the exposed state. All individuals can develop symptoms, with the probability based on its age. Any individual that developed symptons, is transferred from the latent health state into the presymptomatic state (E_p). After their presymptomatic period has elapsed, the individual is transitioned into the symptomatic infectious state (Sy). If an individual does not develop any symptoms,

the individual remains in the asymptomatic exposed state until the recovery duration has elapsed (E_a). Some of the individuals need to be hospitalized, based on probabilities from an age-distribution. If an individual is hospitalized, it is placed into an isolated location at hospitals. In these isolated locations, the individual cannot infect others as the RIVM (2021a) projected that less than 1 percent of individuals are infected in Dutch hospitals. Finally, an hospitalized individual can recover from the disease or die, based on probabilities from an age-distribution. After an individual died or recovered, it cannot be infectious or infected. The disease parameters on the age-dependent transitions, the duration of health states and test willingness are depicted in Table 2.

3.3 Simulation process

The agent-based model proposed in this research is driven by a discrete-event simulation (DES). A DES models a system as a sequence of events, chronological in time. Within this research, events are transitions between locations or changes in health states of individuals, as these events produce interactions between susceptible and infectious individuals. The DES assumes that no changes occur in the time between events and as such, the simulation time can jump to the time of the next event in the chronological sequence (hence the term discrete). In the DES, individuals are modeled as independent entities with attributes. These attributes dictate the movement of individuals and affect the disease progression of individuals. Furthermore, the DES keeps track of variables that form the output of the model. For this model, the output variables are the cumulative infections, hospitalizations and deaths of individuals.

Before the simulation is started, all individuals with their unique attributes and the environment are initialized. This procedure is further elaborated in Section 4. Hereafter, at the start of the simulation, the entire population is initialized at home and their health states are set to be susceptible (S). Hereafter, for each individual an event is created that signifies the time when that individual leaves their home. All these events of the entire population

Algorithm 1: Pseudocode for the Discrete Event Simulation

```
Initialize agents and locations;
Create initial events of agents;
Infect random agents;
while simulation clock does not exceed simulation clock limit do
    while event sequence contains events for the current simulation clock do
        Remove next event from sequence and process the event;
    end
    Update output variables;
    Increment simulation clock;
end
```

Algorithm 2: Pseudocode for an activity event subroutine

```
Parameter: Agent
Depart Agent from its current location;
Add Agent to new location;
Schedule new event at end time of activity;
```

are added to the event sequence. Then, a number of individuals are randomly selected and set to the exposed (E) health state, to form the set of initially infectious individuals. These infectious individuals will subsequently infect other individuals and so forth. Hereafter, events are chronologically processed. Algorithm 1 depicts the pseudocode for the general DES.

When an event is processed, a subroutine is executed. Algorithm 2 displays the pseudocode for the subroutine of an activity event. In such an event, an individual is departed from its current location, and will be added to a new location. Finally, a new event will be scheduled to signify the departure of its new location. For example, let us consider an empty location, an event sequence containing two events and a simulation clock at time t_0. In the first event (E_1), an infectious individual, i_1, arrives at time t_1 at the location, the individual will subsequently depart at t_3. The second event (E_2) contains the arrival of a susceptible individual, i_2, at time t_2. Individual i_2 will departure at time t_4. The processing of three events will be further elaborated below. Figure 2 depicts the three events of the DES for the arbitrary location, where E_i denotes the ith processed event.

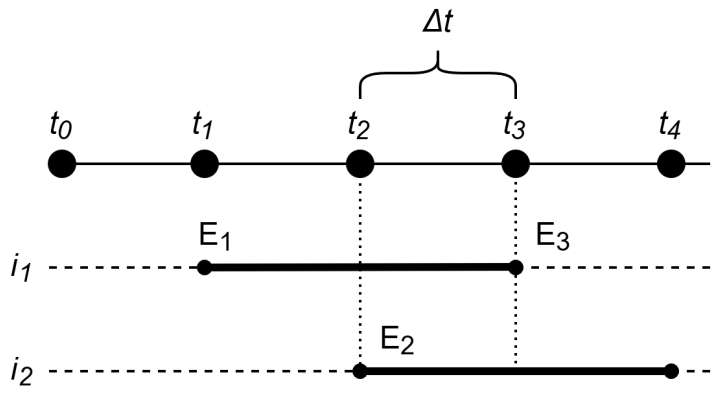

Figure 2: Example of processing events in the DES for a location and two individuals.

At the start, the simulation clock is incremented to t_1, as no events are available at t_0. Hereafter, the event E_1 is removed from the sequence for processing, as it is the first event in chronological order. Thus the individual i_1 is removed from its current location, and added to the empty location. The infection risk at the empty location is not calculated, as there were no individuals present at the location before t_1. Finally, a new event E_3 is created for the departure of i_1 from the location at time t_3. Hereafter, the simulation clock is incremented to t_2 and the next event E_2 is removed from the sequence. In the processing of E_2, the individual i_2 is added to the location. No infection risk is calculated, as only one person was present between t_1 and t_2. A new event is subsequently created for the departure of i_2 at t_4. The simulation clock is incremented to t_3 and finally, the next event E_3 is removed from the sequence. Individual i_1 is removed from the location and the infection risk is calculated, as two individuals were present between t_2 and t_3, one being susceptible and one infectious. As such, individual i_2 is exposed to the disease with probability $p(t_3) = 1–e^{-\Delta(t_3)\lambda(t_3)}$ (as described in Equation 1). Where $\Delta(t_3) = t_3–t_2$.

Travel is explicitly modeled in the simulation. However, if an individual travels by car, bike or foot, it is assumed that this type of travel is executed in isolation. This is achieved by placing the individual into a unique location only containing the individual for the entire duration of that travel activity. If the specified mean of transport is public transportation, interactions between individuals are allowed. Any individual that performs a travel activity

using public transport, enters a fictitious location containing other individuals using public transport at that time.

Social activities between individuals are modeled through social events. At the start of a social event, an organizer invites other individuals from its social network to its home. The invited individuals can only accept an invitation if their current activity allows for social interaction with other agents. Furthermore, the organizer can only invite other agents if the number of individuals visiting its home does not exceed the maximum allowed number by the government. Research by RIVM showed that the compliance regarding visitors was 99.1% during the considered period (RIVM, 2021e).

To simulate the COVID-19 epidemic in Dutch cities, restrictions initiated by the Dutch government are implemented in the simulation. At the start of the simulated period, eating establishments are closed from 8 pm, and households could invite at most 3 individuals over per day. On 14 October, the Dutch government closed all eating establishments all day and closed all shops from 8pm onward. Furthermore, face masks were made mandatory in secondary and higher education institutions and public transport. On 4 November, the government closed all public buildings (e.g., museums, theaters) and limited the number of visitors a household could receive per day to 1.

3.4 Calibration

The parameter configuration of the model is identified by using simulation optimization. Simulation optimization integrates an optimization technique into the simulation modeling, creating a loop of: simulation, evaluation and parameter updating. This loop is terminated when a certain criterion is met. The objective of the optimization problem is to closely match the simulated and observed hospitalization data.

In the ABM, a parameter configuration evaluation is computationally expensive, as multiple independent replications of the simulation are performed to obtain the output of a single configuration. Furthermore, the function of the proposed model does not have an ana-

lytical form, the parameter configuration can only be evaluated with nonlinear optimization methods. For this purpose, a direct search approach called Nelder-Mead (NM) is implemented (Nelder and Mead, 1965). The NM framework is well suited for the proposed model as only a couple parameters need to be evaluated for each iteration. The objective function of the model consists of two components: the Kolmogorov-Smirnov (K-S) statistic and the percentage difference between the predicted and observed hospitalized on the last day. The K-S statistic serves as a goodness of fit of the distribution. The K-S statistic is defined as

$$D = \sup_x |F_1(x) - F_2(x)|, \tag{3}$$

where $F_t(x)$ denotes the distribution function of either the observed or simulated sample. The K-S statistic (D) quantifies the maximum distance between the simulated and observed cumulative distribution function of hospitalizations at any point. This maximum distance can be observed at the final day of the calibration period, while preferably the simulated hospitalizations are close to the observed values on the final day. For example, multiple parameter configuration could posse the same K-S statistic value, while some of these configurations could closely match the simulated to the observed cumulative hospitalizations on the final day. For this purpose, a percentage difference between the predicted and observed hospitalized on the final day is included in the objective value.

The key parameters of the calibration are: initial infections, transmissibility probability (γ) and lockdown transmission transmissibility ($\theta_{t,q}$). The number of initial infections is calibrated to account for reporting errors in number of infected individuals, and is calibrated over the range [150, 250] for both Delft and Apeldoorn. As Apeldoorn has a larger population, while Delft suffered a more severe COVID-19 outbreak. Transmissibility is calibrated over the range [2x10-4, 3x10-4]. The transmissibility reduction parameter ($\theta_{t,q}$) is calibrated over the range [0.90 – 1.00] for schools and public transport after the 14 October, and is equal to 1 for all other locations during this period. The time-span of the model is 61 days,

ranging from 1 October until 30 November. The model is calibrated on the first 31 days, starting from the 1 October. While the last 30 days serve as a validation period.

In the validation process, three measures are used to evaluate the goodness of fit: the coefficient of determination (R^2), Pearson's r and the Kolmogorov-Smirnov two sample test. The coefficient of determination, is defined as

$$R^2 = 1 - \frac{\sum_t (y_t - \hat{y}_t)^2}{\sum_t (y_t - \bar{y})^2}, \tag{4}$$

where y_t depicts the observed value at time t, $\hat{y}_t$ the estimated value and $\bar{y}$ the average of observed values. The R^2 represents the proportion of the variance explained by the model. Pearson correlation coefficient (Pearson's r), defined as

$$r_{x,y} = \frac{\sum_t (x_t - \bar{x})(y_t - \bar{y})}{\sqrt{\sum_t (x_t - \bar{x})^2 \sum_t (y_t - \bar{y})^2}}, \tag{5}$$

and represents the correlation between two sets of data. A value of 0 would imply that no linear correlation is present in between the sets, while a value of 1 would suggest that the observed and simulated data are perfectly correlated. Finally, the K-S two sample test is used to compare the cumulative distribution functions between two samples. The null hypothesis of the K-S two sample test states that both distributions are drawn from the same distribution. If the null hypothesis is rejected, it concludes that both distributions are not drawn from the same underlying distribution. All following tests are evaluated a significance level of 0.05.

The used optimization algorithm is prone to be stuck in local minima due to the initialization values of the parameters. To combat this, the algorithm is restarted to escape these local minima, by randomizing parameter configurations within their ranges. The optimization of the parameter configuration is terminated when the number of iterations exceeds a

user-specified number, or if the similarity between observed and simulated data converged sufficiently. The returned parameter configuration is used to validate the model on the last 30 days of the 61 day period for each municipality.

4 Data

In this section, the data and the generation process of the synthetic geography and population are described. Hereafter, the activity patterns and the creation of social networks for individuals is outlined. The disease data is discussed and finally, the results and the limitations of the generation process are discussed.

4.1 Synthetic geography

All 355 municipalities are divided into multiple districts by the Central Bureau of Statistics (CBS) of the Netherlands. This subdivision provides a high resolution view of each municipality. A geographical information system software (QGIS version 3.16.2) was used to obtain administrative bounds for all (12.814) districts, provided by CBS (2021f). The Basic Registration of Addresses and Buildings (BAG) from the Dutch cadaster provides a dataset on all buildings in the Netherlands (Kadaster, 2021). Each building in the BAG dataset has its function categorized (residential, office, retail, industrial, educational, healthcare).

Hereafter, geographical coordinates of amenities (e.g., restaurants, cinemas, or supermarkets) were obtained from OpenStreetMap (OSM) by using the software Osmosis (Osmosis, 2021). A full list of amenities included in the model is provided in Appendix 8.1. OSM is a volunteer geographic information dataset, where individuals can submit information about the environment which they are living in (OSM, 2021). Numerous scientific studies, studied the data quality of OSM. These studies show that around urban areas, a vast amount of positional accurate data can be obtained (Ali and Schmid, 2014). The geographical coordinates of amenities obtained from OSM were mapped to buildings. To this extent, a

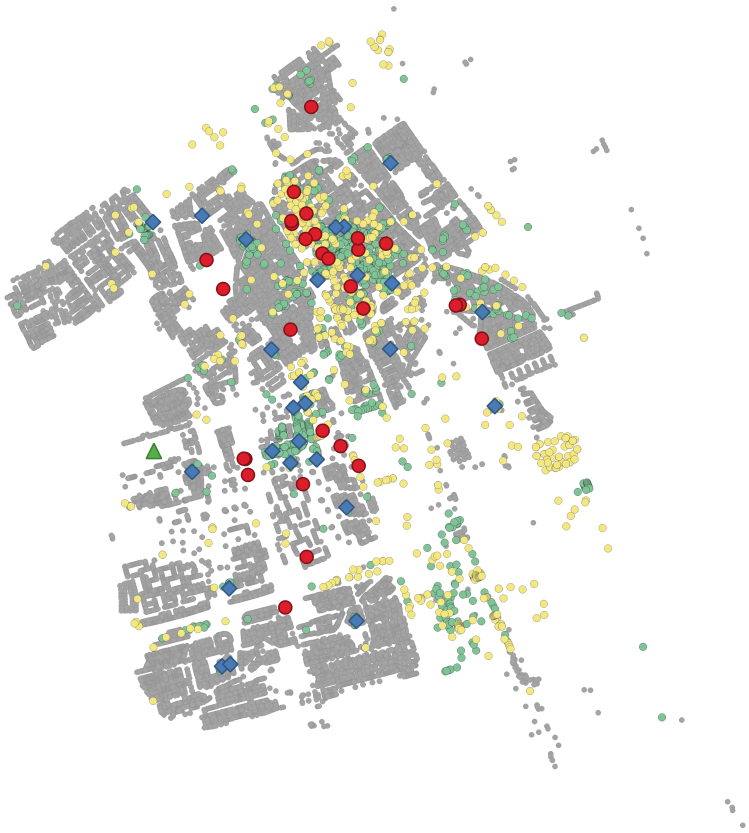

Figure 3: Map of Delft. Grey circles represent households, yellow circles office buildings, green circles shops, red circles religious buildings, blue diamond supermarkets and green triangle hospitals.

QuadTree data structure was employed to determine the closest building to a coordinate (Aizawa et al., 2008). Some amenities in the OSM dataset specified an address, while others did not. Those that did not contain an address, were matched based on nearest relevant building, while matching by address was used in cases whenever possible.

School data is obtained from the Dutch office of education (DUO, 2021). The dataset includes all locations and number of students of all types of schools in the Netherlands. These schools are matched to buildings according to their address. Finally, all buildings were matched to their corresponding district in QGIS. Figure 3 depicts the distribution of some building types throughout Delft.

4.2 Synthetic population

For each district in the Netherlands, data is available on the number of people living in the district categorized into age groups (0-14, 15-24, 25-44, 45-64 and 65+ years) and types of household (single parent, living alone, living together without children, living together) (CBS, 2021i). The age-distributions per household type was not available on a district-level, but was obtainable on municipality-level (CBS, 2021h). Using these datasets, synthetic households were created and synthetic individuals were distributed over these households. The resulting households were hereafter assigned to residential buildings within the corresponding districts.

Data on the number of students for all education types was obtained for each municipality (CBS, 2021j). Individuals were assigned to education types (primary, secondary, vocational, college or university) based on their age. Hereafter, students were assigned to nearby schools based on the distance between the school and the student's home. Finally, all students were assigned to classes in each school. For primary and secondary schools class sizes are on average 25 students per class (Slob, 2021), while for higher types of education, school class size was assumed to be 60.

The number of workers per industry type (industry, healthcare, culture, retail, education, or office) was obtained on a municipality-level (CBS, 2021d). Each individual, not a student or retired, is assigned to be a worker based on the employment rate corresponding to its household position (being a parent, single, single parent, child, adult without children). Data on household position employment was not obtainable on a district or municipal-level, but only available on a national-level from CBS (2021c). Commute data was obtained from CBS (2021l), it specified the number of commuters between all municipalities in the Netherlands. Finally, workers are assigned to companies with varying sizes (5, 15, 35 and 100+) and companies to buildings matching the industry type. The distribution of company size for each industry type is based on national-level data (CBS, 2021e).

In the COVID-19 epidemic, the government has urged individuals to work from home.

However, not all type of jobs are suitable to work from home (e.g., industrial jobs or eating establishments). To account for these differences, data was obtained from TNO (2021) about the percentages of employees working from home for each industry type.

4.3 Activity patterns and social networks

An activity pattern dataset is obtained from the Netherlands Institute for Social Research (SCP, 2021). This dataset provides sequences of activity patterns of 2,260 individuals that kept a diary of their activities for one week in 2016. Activities in a sequence denote the type of activity (e.g., working, traveling, or eating), location (e.g., home, at work or restaurant) and duration in 10 minute intervals. To assign activity patterns to individuals, the dataset is categorized into three categories: age, social role, and household role. The social role of an individual signifies its occupancy (unemployed, school-going, worker or retired), while the household role denotes the position of the individual in a household (e.g., child or parent). Based on the categorized dataset, each individual is assigned a random activity pattern from their respective categories.

Furthermore, each individual is assigned a network that consist of people with who the individual discussed personal matters or would visit in their leisure time. This network is the so-called 'core discussion network' studied by Mollenhorst (2009) by analyzed the *Survey of the Social Networks of the Dutch* (Röper et al., 2009). These networks for individuals in the Netherlands, have an average size of 9.75 (SD = 4.65). The network consists of 4 types of relationships: friends (24.6%), relatives (21.7%), colleagues (17.4%), and neighbors (36.2%).

In the model, all individuals are assigned a social network size sampled from a Gamma(4.33, 2.70) distribution, with corresponding mean 9.75 and standard deviation 4.65. However, differences exist between individuals (age, gender, household position and social role) that influence the probability that two individuals will be included in each others social network. For instance, an 80-year-old retiree is not likely to have a 20-year-old student included in their network. For this purpose, a social similarity score is calculated between individuals

based on the approach used in the ABM by Zhang (2016).

The similarity between two individuals A and B, is evaluated by a weighted Euclidean distance, expressed by

$$S(A, B) = \sqrt{\sum_{f} \mu_f (A_f - B_f)^2}, \tag{6}$$

where f denotes a feature of set $F = \{\text{age, gender, social role and household position}\}$ and μ_f is the weight factor for feature $f \in F$. This score signifies the percentage of similarity between two agents.

The similarity score is used to determine whether individuals are included in each others social network. The process of assigning individuals to a social network differs between relationship types. Specifically, colleagues are randomly selected from an individuals' place of work, and accepted in its network if they are sufficiently similar. Neighbors are determined by searching for the 100 nearest households utilizing a nearest-neighbor search in the QuadTree originating from an individual's home. Friends and relatives are randomly selected individuals within a range of 20 km from an individuals' home, as on average 87% of social trips occur within 20km (CBS, 2021g). Moreover, the relatives in a social network include all individuals that live together with a family relationship (e.g., child-parent, or parents in the some household).

4.4 Disease data

The daily cumulative number of positively tested individuals, hospitalized patients and deaths for each municipality was obtained from RIVM (2021c). The considered period is from 1 October until 30 November 2020. This period is commonly regarded as the 'second' wave in the Netherlands. Like the first wave, the second wave crippled the Dutch healthcare systems and had a potential to cumulate into a catastrophe. While testing was not widely available in the first wave, it was available for all individuals by appointment during the second wave (Algemene Rekenkamer, 2021). At the start of October, only a very limited

number of NPIs were enacted by the government. A combination of tests being widely available, the low number of NPIs and the low number of hospitalizations at the beginning of the second wave, provides the model a good position to estimate the effects of alternative NPIs.

The RIVM notes that the actual number of positive tests, hospitalizations and deaths are under-reported as not all individuals get tested. Especially the daily number of positive tests lacks accuracy, as individuals can be asymptomatic and thus do not get tested as they are unaware of their COVID-19 infection. Furthermore, the RIVM notes that the data regarding hospitalized individuals and deaths could be delayed due to delays in registration by medical professionals.

4.5 Synthetic data generation results and limitations

The proposed synthetic population generation process and its data, as described in paragraph 4.2, can be applied to all municipalities in the Netherlands. However, due to limited computation power, this research focuses on two Dutch municipalities: Delft and Apeldoorn. Individuals (so-called agents in the model) are assumed to be contained in the municipality (i.e., no individual can enter or leave a municipality). This assumption restricts the individuals commuting between municipalities for work or schools, and limits the model in its ability to simulate the spread of the disease throughout neighboring municipalities by omitting interactions between individuals from various municipalities. However, neighboring municipalities show similar infection rates to the studied municipalities (RIVM, 2021f), and thus individuals are exposed to relatively the same infection risks even if they cannot commute. The scalability of the model to all municipalities in the Netherlands is discussed in Appendix 8.2. Delft and Apeldoorn consist of 101,474 and 157,198 residents respectively. The two municipalities were selected based on their relevance in the Dutch epidemic. Apeldoorn is regarded the most average Dutch municipality, as its population composition closely resembles the whole of Netherlands (Whooz, 2021). Moreover, of all municipalities, the results of Apeldoorn for the 2021 national elections closest resemble those of the Netherlands

	Delft		Apeldoorn	
Category	Synthetic	Census	Synthetic	Census
Population				
Total individuals	101,474	101,910	157,198	160,995
Male	53.2%	53.2%	49.7%	49.7%
Female	46.8%	46.8%	50.3%	50.3%
Age distribution				
0-14 years	12.5%	12.5%	16.0%	16.1%
15-24 years	20.2%	20.2%	11.5%	11.2%
25-44 years	28.4%	28.4%	23.1%	22.7%
45-64 years	23.3%	23.3%	29.1%	28.4%
65+ years	15.6%	15.6%	20.4%	20.0%
Household characteristics				
Number of households	58,488	58,565	72,760	72,035
Average household size	1.85	1.76	2.16	2.23
One person households	56.8%	57.0%	35.5%	35.9%
Households with kids	20.9%	21.3%	32.2%	33.4%
Household without kids	22.2%	21.7%	32.3%	30.7%
Student count per school type				
Primary	6,600	6,600	12,989	12,989
Secondary	4,203	4,203	9,436	9,436
Vocational	1,839	1,839	5,096	5,096
College	2,967	2,967	3,445	3,445
University	13,480	13,480	932	932

Table 3: Population characteristics of the created synthetic population.

(ANP, 2021), further indicating that Apeldoorn is a good representation of the Dutch average. While Delft, at the start of the selected time span, had the most infections per 100.000 individuals of all municipalities in the Netherlands (RIVM, 2021b). Thus, the municipality of Delft provides additional information on the effect of NPIs on municipalities struggling with a COVID-19 outbreak.

Every building obtained through the cadaster dataset was mapped to a district of a municipality in the Netherlands. Hereafter, the geography generation process successfully matched all 215,179 amenities and 8,668 schools in the Netherlands to these buildings. In the geography data generation process, these amenities and schools were matched to buildings with an offset of 6 meters if their addresses were specified. While, the offset was 81 meters if the address of the amenity of school could not be extracted from its data. The coordinates

Category	Delft Synthetic	Apeldoorn Synthetic	Census
Social network			
Average size (std. dev.)	9.75 (4.65)	9.82 (4.59)	9.75 (4.65)
Friends	24.3%	23.7%	24.6%
Colleagues	22.0%	19.7%	21.7%
Neighbors	17.7%	19.6%	17.4%
Relatives	36.0%	37.1	36.2%
% employment by household type			
One-person household	63.1%	63.1%	61.3%
Single parent household	74.1%	74.1%	72.0%
Two parent household	89.7%	89.8%	87.2%
Without children household	75.3%	75.4%	73.2%
Living with parents	78.9%	79.0%	76.7%
% working at home by industry type			
Industry	27.2%	25.3%	26.2%
Shop	24.7%	24.0%	24.9%
Horeca	25.2%	24.8%	24.9%
Office	70.8%	70.2%	70.6%
Education	26.4%	27.1%	27.1%
Healthcare	13.2%	12.9%	13.0%
Culture	46%	44.1%	44%

Table 4: Social network and employment characteristics of the created synthetic population.

lacking addresses were matched to the nearest building by employing the QuadTree data structure to find nearest neighbors. The small offset of assignment of buildings indicate that amenities and schools, especially with addresses, were assigned to nearby buildings in the model, thus providing high positional accuracy of facilities in the model.

The results of the synthetic population generation are presented in Table 3. For both municipalities, a column of their synthetic generated population and a column on the relevant census data is shown. The synthetic population matched the census data well for both Delft and Apeldoorn. Some differences in census data between the municipalities can be noted, as Delft has a vast number of university students compared to Apeldoorn. This difference in students results in an increase of individuals aged 15-24 years, an increase of males due to the technical university located and an increase of one person households in Delft.

The results of the social network generation and employment assignment is displayed in

Education type	Average travel distance (km)	Average occupancy compared to census data
Primary	0.83	96.7%
Secondary	4.04	101.1%
Vocational	19.28	101.3%
College	16.35	104.0%
University	12.93	95.5%

Table 5: Education characteristics of the synthetic population.

Table 4. The generation process closely adheres the data regarding social networks, indicated by the closely matched relative group sizes. Furthermore, the average size and standard deviation of the social network size is similar to the observed data for both municipalities. Moreover, the employment per household type and the percentage of individuals working from home per industry type matches the census data.

The results of the student assignment is depicted in Table 5. The table shows that for each education type, the synthetic schools closely matched the census data on number of students for each school. For example, primary schools in the model are assigned on average 96.7% of the students they are expected to have per their census data. Furthermore, the table depicts the average distances that students have to travel to their schools for each education type. The average distance that a student has to travel to primary and secondary schools is 0.8 and 4.0 km respectively. The Dutch average distance from a students' home to the *nearest* primary school is 0.7 km (CBS, 2021a), and 2.4 km to the *nearest* secondary school (CBS, 2021b). The nearest distance serves as a lower bound of the actual average distance that a student has to travel, as not all students attend the nearest school to their home. Only 0.2% of primary school students need to travel more than 3 km to the nearest primary school (Eurydice, 2021), while 15% of secondary school students need to travel more than 5 km (CBS, 2021b). Therefore, if a student does not attend the nearest school, the average distance to travel to the second nearest school is lower for primary schools than secondary schools. Thus, the student assignments for primary and secondary school shows that it adheres to the lower bounds and the simulated travel differences between primary and

secondary schools can be explained by the differences in nearest distance to these schools.

Higher education lacked data on the average distance to students' homes or nearest school distance. However, the generated differences in distances for higher educations can be explained by their relative difference in percentage of independent living. Only 17.8% of the students enrolled in vocational schools live on their own. College students have a higher average of 46.8%. While 64.5% of the students enrolled in universities live, on average, independent (CBS, 2021k). These difference between independent living, can explain the reduction of average distances from students' homes to their assigned schools per education type. For example, university students are more likely to live independently compared to students enrolled in vocational schools, and thus university students need to travel a shorter distance to their schools on average. The results on average distance traveled and occupancy of schools, indicate that the model's student assignment performs well.

While the generation process is successful, there are limitations regarding the data. Firstly, by using the OSM dataset, it is possible that not all relevant locations of amenities are included in their dataset as it is an volunteered geographic information project. However, Delft and Apeldoorn are urban areas, and the Netherlands is one of the countries with the highest number of contributors per area (Osmstats, 2021). Secondly, some data regarding the population was not obtainable on a district or municipal-level. Thus national-level data was used, while this data represents the average of Dutch municipalities, it does not capture difference between municipalities. Nevertheless, national-level data provides differences within the population that would otherwise not be captured. Moreover, the national-level data is still relevant to the Dutch setting, as it represents Dutch averages. Thirdly, the disease data can be inaccurate as healthcare professionals can experience delays with registering the data. To combat this, all disease data is averaged over 3 days. Fourthly, the activity patterns are not differentiated other than on age, social role and household position. Thus, the heterogeneity between individuals is not fully captured, as for example, individuals working in different industry types exhibit different activity patterns. Furthermore, the ac-

tivity patterns of individuals do not vary between the municipalities, omitting social-cultural differences between different municipalities. Lastly, the assumptions made regarding similarity for social networks, class size and company size are further analyzed in the sensitivity analysis to analyze the impact of their uncertainty on the output of the model.

5 Results

The results of the simulated ABM are presented in this section. First, the results of the calibration and validation periods are presented for Delft and Apeldoorn. Second, the results of various NPI strategies are discussed. The simulations performance results and the coefficient of variation are shown in Appendix 8.2.

5.1 Calibration and validation results

The model is calibrated on the cumulative hospitalization between 1 October and 31 October. The municipality of Delft calibrated well, with a R^2 and Pearson's r of 0.976 and 0.997 respectively for the number of hospitalizations before 31 October. Similarly, the model validated well on the number of hospitalizations between 1 November and 30 November, supported by the R^2 and Pearson's r of 0.970 and 0.989 respectively. Apeldoorn calibrated well with a R^2 and Pearson's r of 0.966 and 0.985, and validated with 0.821 and 0.966 respectively. The hospitalization for both municipalities is shown in Figure 4. The validation R^2 of Apeldoorn decreased substantially more than the validation R^2 of Delft. The lower R^2 can be explained by the relatively large number of observed hospitalized individuals right before the end of the calibration period and at the start of the validation period. Finally, the K-S two-sample statistic did not reject the null hypothesis for both the calibration and validation period for both municipalities, depicted in Table 6. Indicating that the simulated and observed hospitalizations distribution functions did not significantly differ from each other.

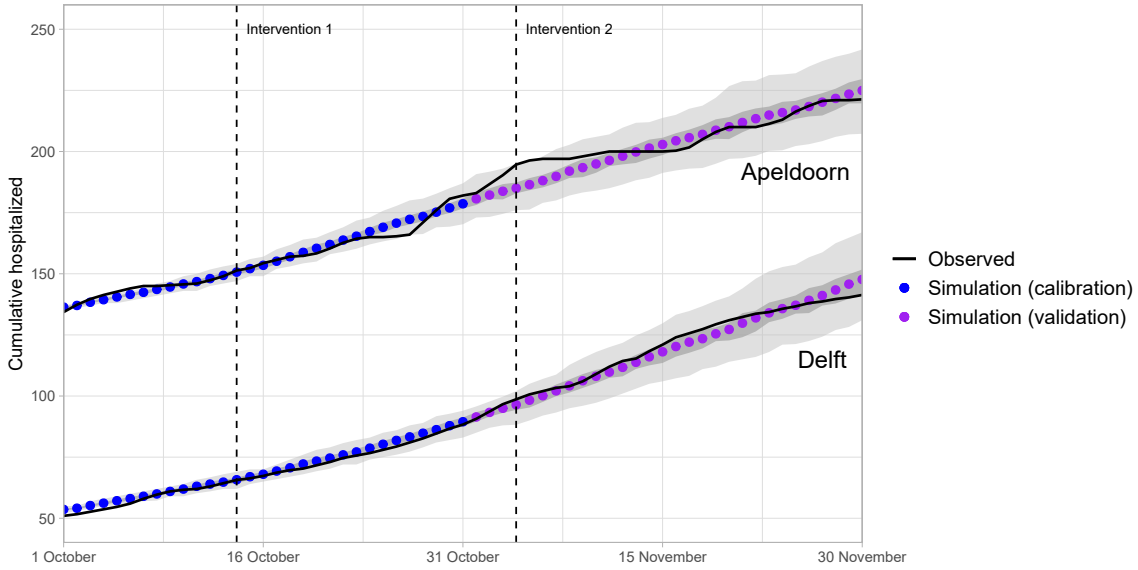

Figure 4: The cumulative hospitalization of Delft and Apeldoorn for the calibration (1 October until 31 October) and validation (1 November until 30 November) period. Vertical lines depict government enacted NPIs days (14 October and 4 November). Dark grey shaded represent the 95% bootstrap C.I., while the light grey shaded area depicts the IQR.

The model was not trained using data on deaths, therefore the data on the complete period between 1 October and 30 November can be used to validate the model in terms of goodness of fit using data on deaths. Figure 5 depicts the fit of predicted deaths for both Delft and Apeldoorn. The model provided a good fit for the number of deaths between 1 October and 30 November, with a R^2 and Pearson's r of 0.926 and 0.981 for Delft. The R^2 and Pearson's r for Apeldoorn are 0.382 and 0.982 respectively. Indicating that the deaths for Delft are well explained by the model, while the model does not explain the deaths

		Delft		**Apeldoorn**	
	Days	Kolmogorov-Smirnov statistic	p-value	Kolmogorov-Smirnov statistic	p-value
Hospitalization	1 - 31	0.097	0.999	0.129	0.959
	32 - 61	0.100	0.999	0.200	0.586
Deaths	1 - 61	0.197	0.189	0.279	0.018**
	15 - 61	0.213	0.238	0.255	0.093*

Table 6: KS-statistic for cumulative hospitalization and deaths for Delft and Apeldoorn. Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

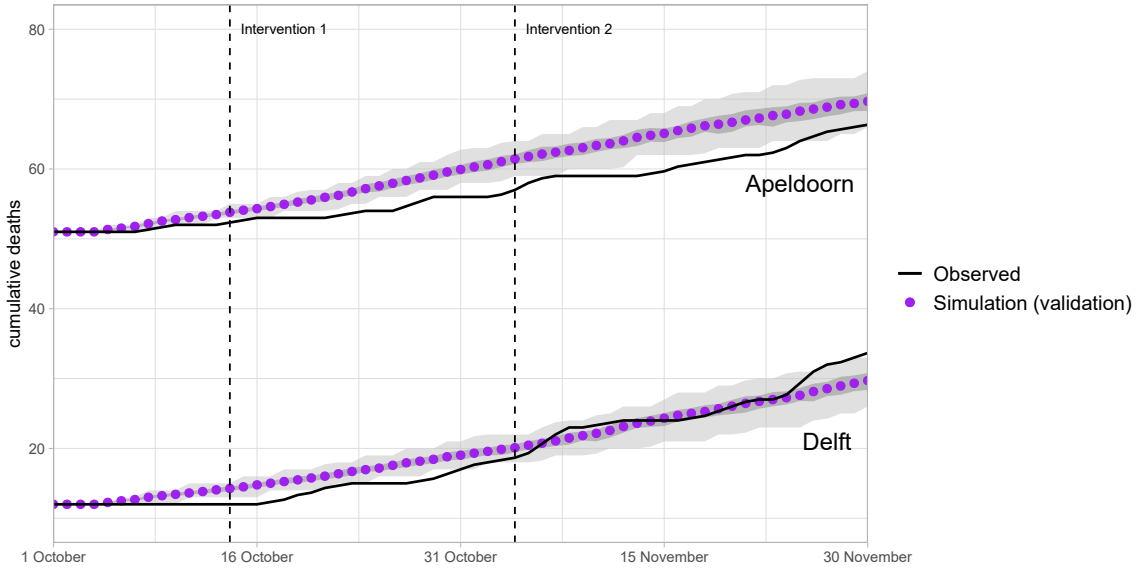

Figure 5: The cumulative deaths of Delft and Apeldoorn for the validation (1 October until 30 November) period. Vertical lines depict government enacted NPIs days (14 October and 4 November respectively). Dark grey shaded depicts the 95% bootstrap C.I., while the light grey shaded area depicts the IQR.

well for Apeldoorn. Although the model could not predict the deaths in Apeldoorn well, the observed and simulated deaths were highly correlated as indicated by the Pearson's r. This can be explained by the overestimation of the deaths in Apeldoorn at the start of the simulation, while later it matches the observed deaths growth. This relation was further confirmed by the KS two-sample test of Apeldoorn shown in Table 6. The KS two-sample test showed that the probability functions significantly differed when looking at all 61 days for Apeldoorn. However, when the first two weeks are omitted, the distributions of the observed and simulated growth would not significantly differ at a 5% significance level.

Furthermore, the age-distribution of deaths in Delft and Apeldoorn are depicted in Figure 6. The Chi-squared test null hypothesis of no differences between the simulated and observed age-distribution is rejected for both Delft and Apeldoorn, as the age-distribution of deaths showed significant differences, $\chi^2(df = 4) = 48.7$, $p < .01$ and $\chi^2(df = 4) = 57.2$, $p < .01$ respectively. While the null hypothesis of the Chi-squared test is rejected, the model still explained the observed deaths for the individuals above the age of 45 years relatively well

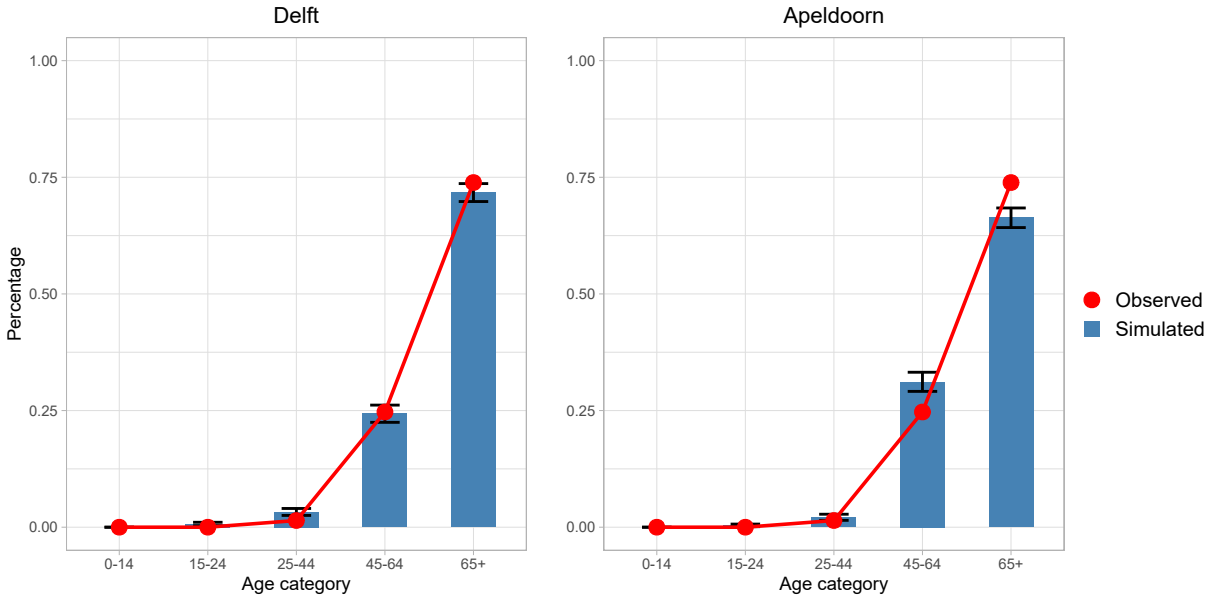

Figure 6: The age-distribution of deaths for Delft and Apeldoorn. Error bars denote the 95% confidence interval.

for Delft and Apeldoorn. The model does poorly for younger individuals, and the high χ^2 value can be attributed to the age categories below 45 years. However, the individuals aged 45 or younger only account for 1.4% of the total deaths in the Netherlands (RIVM, 2021c). Nonetheless, this difference might be caused by differences in underlying characteristics of municipalities (e.g., Delft has many students) that are not being captured in the observed age-distribution of deaths as this distribution comes from the Netherlands as a whole.

Finally, the simulated daily number of infections and the daily number of positive test per day for both Delft and Apeldoorn are shown in Figure 7. While Delft has a smaller population, its daily number of positive tests is comparable to Apeldoorn. This can be explained by the fact that Delft had the most infections of all municipalities per person at the start of the considered period. Furthermore, as expected, the simulated number of infections per day considerably outnumbers the observed positive test results on most days. This difference can be attributed infected individuals that are not willing to get tested and individuals unaware of their COVID-19 infection.

Concluding, the model could explains the cumulative hospitalization well for two Dutch

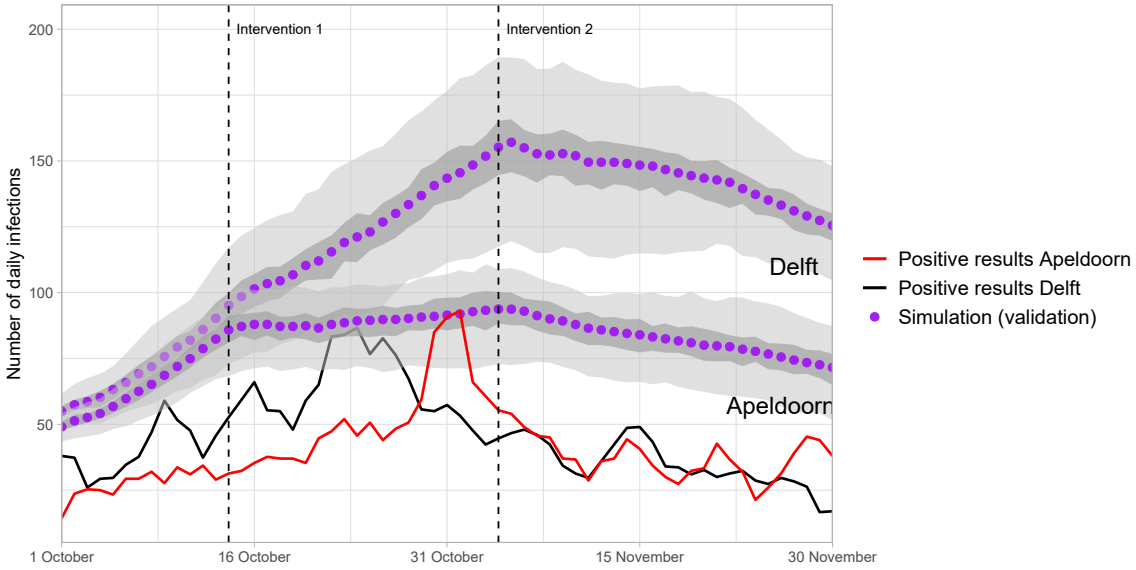

Figure 7: The daily infections for Delft and Apeldoorn as validated (purple) and 3-day averaged observed for Apeldoorn (red and black respectively). Vertical lines depict days on which the dutch government enacted NPIs (14 October and 4 November).

municipalities. Furthermore, deaths could be explained well for Delft. The model overestimated the number of deaths in the first two weeks for Apeldoorn, however it correctly modeled the deaths after these two weeks. As expected, the model overestimated the observed positive test results per day. While the model could not explain the observed age-distribution of deaths in the studied municipalities over all age groups, it estimated the age-distribution of deaths for individuals aged 45 or older to some extent.

5.2 NPI strategies evaluation

Next, varying intervention strategies are simulated. First, a strategy is adopted where no interventions are enacted. Hereafter, the enacted NPIs by the Dutch government, are implemented on other days than the original NPIs. Specifically, the enacted NPIs are scheduled 7 days earlier and 7 days later. Where the original dates were 14 October and 4 November for the original two enacted NPIs. Thirdly all enacted NPIs are scheduled on a single day, 14 October, instead on the original two separate dates. Finally, the following additional NPIs

are investigated: closure of religious locations, closure of primary and secondary schools, closure of higher education, and the closure of all schools. All enacted on 14 October, the first stringent intervention date of the second wave. The results of the varying intervention strategies are displayed in Figure 8 for Delft and Apeldoorn. The output of each intervention strategy is based on 100 simulations for both municipalities, as discussed in the performance analysis (Appendix 8.2).

The 'what-if' intervention strategies display the importance of closures. For each strategy, percentage differences are calculated between the strategy and the original simulated government strategy in terms of the cumulative hospitalizations on the final day (30 November) for Delft and Apeldoorn.

The top panels of Figure 8 depict the spread of COVID-19 when no intervention would have been enacted and the effect of closing religious places in Delft and Apeldoorn. The number of hospitalized individuals would have been drastically higher for both municipalities if no NPIs were to be enacted by the Dutch government (Delft: +51.7%, Apeldoorn: +59.1%). The additional intervention of closing religious places on 14 October resulted in a decrease in the hospitalized individuals for Delft (-5.9%), while a minor decrease was observed for Apeldoorn (-0.5%).

The middle panels show the effect of different timings of the NPIs implemented by the Dutch government. If the government had acted sooner and thus implemented the NPIs on 14 October and 4 November 7 days earlier, the model predicted a significant decrease in the number of hospitalizations on the 30th of November (Delft: -10.9%, Apeldoorn: -7.4%). Would the government have waited 7 days more for both intervention responses, the model predicted an significant increase in hospitalization for both municipalities (Delft: +10.0%, Apeldoorn: +11.1%). If the Dutch government enacted the original NPI responses on a single day, 14 October, the cumulative hospitalization were projected to decrease for both municipalities (Delft: -14.7%, Apeldoorn: -6.1%). The difference in effect from the early or stringent NPI strategy between Delft and Apeldoorn can be explained by the difference in

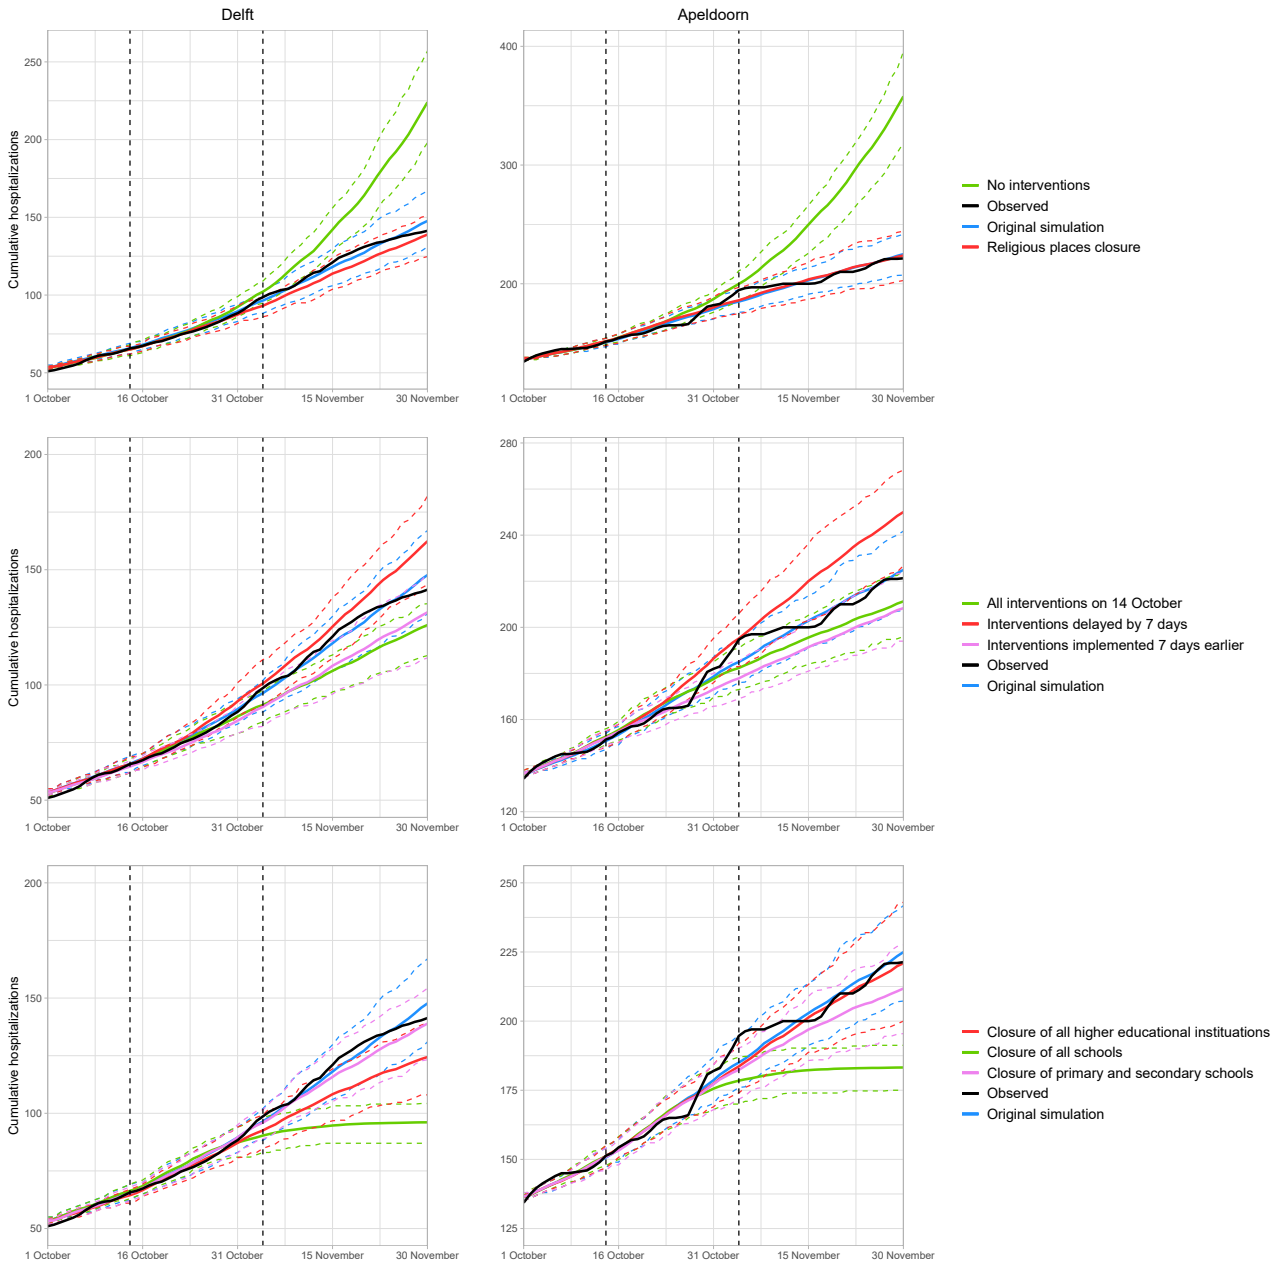

Figure 8: NPI strategy evaluation for Delft (left) and Apeldoorn (right). Vertical lines depict government original NPI dates (14 October and 4 November). The dashed lines represent the IQR of the estimates.

number of students. As Delft has many students, and thus benefits more from the mandatory face mask wearing in education institutions compared to Apeldoorn.

The bottom panels in Figure 8 depict the strategies related to education. The model predicts that interventions related to schools had a significant impact on the reduction of

hospitalizations for both municipalities. In Delft, the closure of higher education type schools (-15.8%) had a larger impact on the hospitalization than primary and secondary schools (-5.9%). This difference can be explained by the relative large number of students living in Delft that attend higher education, as discussed in section 3. In Apeldoorn, the closure of primary and secondary schools was related to a larger reduction in hospitalization (-5.9%) than higher education institutions (-1.8%). This result is in line with the expectations, as a relatively low number of higher education students compared to primary and secondary school students are living in the municipality of Apeldoorn. When all schools are closed additional to the enacted NPIs on 14 October, the spread is drastically reduced for both municipalities (Delft: -34.9%, Apeldoorn: -18.5%). This result is in line with related literature on NPIs that show that interventions targeting education institutions produce large impacts.

6 Sensitivity analysis

The sensitivity of parameters are analyzed by performing an one-factor-at-a-time (OAT) sensitivity analysis. In OAT, parameters are varied one at a time, while keeping all others fixed to generate output diversity (Pianosi et al., 2016). The variation in output caused by varying a single parameter can be used to investigate the type of relationship between a parameter and the output. For instance, the relationship could be linear, nonlinear or a tipping point where the output drastically varied due to a minor change in a parameter value (ten Broeke, 2017). Investigating these dependencies, the mechanism of the model can be better understood and guide future research.

In the OAT sensitivity analysis, 9 equidistant points within $\pm 20\%$ of the normal parameter value are used to analyze the sensitivity of each parameter. For each of the 9 parameter values, 25 replications of the simulation are used to estimate the output due to stochasticity. Earlier results on the coefficient of variation in Appendix 8.2 indicated that 25 replications

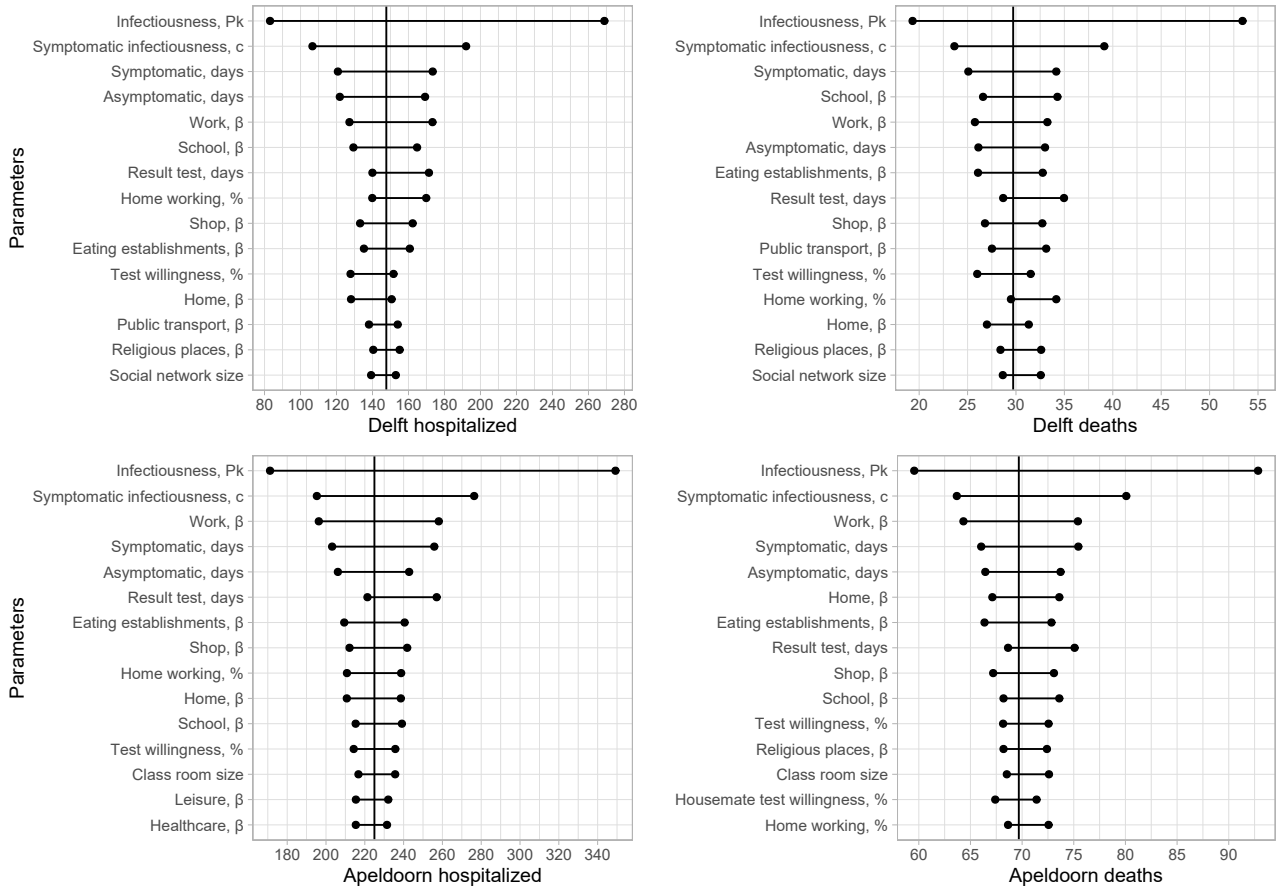

Figure 9: Sensitivity analysis results by varying parameters containing uncertainty between -20% and +20% of their nominal value denoted by the vertical line.

would suffice to roughly estimate the output as the variance stabilizes around 25. During the sensitivity analysis, the model is not recalibrated when varying a single parameter. As the purpose of the OAT is to investigate the change in output by small variations of single parameters.

In the model, 20 parameters were identified to contain some uncertainty. All 20 parameters were analyzed, the 15 parameters that produced the largest variation of output are depicted in Figure 9.

Parameters related to the disease progression produced the largest variation in output of the model. Specifically, the infectiousness of an individuals (P_k) and the difference in contagiousness between symptomatic and asymptomatic individuals (c) impacted the out-

put the most. The duration of the symptomatic and asymptomatic period had a lesser, but still large impact on the output. These results were to be expected, as an increased infectiousness of individuals (P_k) accelerates the rate of which the disease spreads exponentially. Furthermore, changes in the symptomatic period produced larger differences in the output than changes in the asymptomatic period. This can be explained by the difference in infectiousness between asymptomatic and symptomatic individuals. Thus a change in the symptomatic infectious period produces larger differences, as symptomatic individuals contribute more to the contagion risk at locations.

The OAT showed that parameters related to the contagion risks of different locations (β) had a substantial impact on the variation of the output for both municipalities. Specifically, contagion risk of workplaces had a large impact on the output. Contagion risks of schools was identified to have a large impact in Delft, while it had a smaller impact in Apeldoorn. These results are in line with expectations, as Delft has a large amount of students. Furthermore, research by (RIVM, 2021a) projected that most infections occur at workplaces and schools, which is reflected in the sensitivity of the parameters related to workplaces and schools. Similar, but to a lesser extent, the contagion risk in households, eating establishments, shops and public transport had a impact on the output variation. While healthcare, religious and leisure related locations had the smallest impact on the output of the parameters related to location dependent contagion risk.

Uncertainty surrounding parameters on testing was found to have some impact on the output. The time until an individuals receives its test result had the largest impact of these parameters. The willingness of individuals to get tested had a smaller impact. This difference can be attributed to the number of individuals affected by changes in the parameters. As a change in the time until a test result is received, impacts all individuals that get a test in the model. Which in turn affects the time that individuals spent infecting others. While a change in the percentage of individuals that get tested, only affect a relative small proportion of this population. Thus, varying the time until a test result is received, affects the propagation of

the disease more than varying the percentage of people willing to get tested. The willingness of housemates to get tested when their housemate got a positive test result, was found to have a small impact on the output.

The analysis revealed that parameters related to the initialization of individuals, such as social network size, classroom size and social similarity threshold values of social networks, did not produce large differences in the output for both municipalities. Thus, the model is quite robust regarding changes in the initialization of individuals. While changes in the parameters regarding the disease progression, testing and contagiousness at locations affected the output substantially.

7 Discussion

In this paper an ABM was presented to simulate the disease spread of COVID-19 through Dutch municipalities. The model initialized the synthetic locations and households of two Dutch municipalities successfully. The synthetic buildings and facilities were mapped to districts of Dutch municipalities with high positional accuracy. Furthermore, the generated households closely matched the census data of municipalities, and the assignment individuals to workplaces and students to schools is in accordance with employment and student data. With the use of real world activity patterns, the model simulated the spread of COVID-19 through the synthetic population.

The model was trained using hospitalization data of one month, and validated on the data of the subsequent month. The proposed ABM was able to closely match the recorded hospitalizations for both municipalities during the validation period. Furthermore, while the model was not trained using data on deaths, it could explain the deaths well for Delft during the two month period. The model overestimated the deaths in Apeldoorn in the first two weeks, while it correctly estimated the number of deaths in the following six weeks. The age-distribution of deaths significantly differed from the observed distribution. However, the model explained the age-distribution of deaths of older individuals well, whom make up the

largest portion. Various NPI strategies were analyzed, and the closure of religious places were found to have no large impact on reducing the spread of COVID-19. Furthermore, the closure of schools was related to large decreases of COVID-19 hospitalizations in both municipalities. The sensitivity analysis revealed which parameters had strong relationships with the output of the model, especially parameters related to the disease progression produced large variation in the output of the model. All relationships and their sizes could be explained, indicating that the model performs as expected.

However, the proposed ABM has its shortcomings. First, in the model only isolated municipalities are analyzed, omitting travel between municipalities. Many employees commute between municipalities for their work in the Netherlands, thus by omitting these dependencies between neighboring municipalities the model fails to capture the propagation of the disease between municipalities. However, neighboring municipalities displayed similar infection rates per 100.000 residents compared to the studied municipalities. Thus, while the studied individuals were not able to contract the disease outside their municipality, they were roughly exposed to the same infection risk. Second, individuals did not differ in behavior other than their activity patterns. All individuals had the same probabilities of getting tested, or working from home for an industry type. While some individuals would be more likely to not adhere to multiple social distancing guidelines (e.g., getting tested, working from home and abiding the maximum visitors allowed). Third, while travel is explicitly modeled, public transportation is aggregated and not specified into different transportation possibilities (e.g., bus, train, or subway). Thus the model is currently unable to capture the spread of the disease through parts of the population that uses a specific type of public transport. Finally, locations are aggregated in a number of categories, omitting difference between location types and thus their inherent differences of contagious risk. For instance, theatres and cinemas are both categorized as leisure. While people in theatres could be reacting to a show by laughing or cheering, increasing the contagious risk in the venue.

Finally, some directions of future research are presented. First, the ABM could be im-

proved by including more municipalities in the simulation. As shown in the computational results (Appendix 8.2), this can easily be achieved by employing more computational power as the model scales linearly in time and space complexity. Second, various vaccination strategies could be included in the model to analyze their impact on the hospitalizations in Dutch municipalities. For example, vaccinating younger individuals first, as the NPI evaluation showed that these individuals had a significant impact on the disease propagation. Third, as changes in parameters of the disease progression are shown to have strong relationships with the output, more research into these parameters can make the model more robust. For instance, the infectiousness variability between individuals and the difference of infectiousness between symptomatic and asymptomatic cases were found to have the largest impact on the output.

References

Aizawa, K., Motomura, K., Kimura, S., Kadowaki, R., and Fan, J. (2008). Constant time neighbor finding in quadtrees: An experimental result. pages 505–510.

Ajelli, M., Gonçalves, B., Balcan, D., Colizza, V., Hu, H., Ramasco, J., Merler, S., and Vespignani, A. (2010). Comparing large-scale computational approaches to epidemic modeling: Agent-based versus structured metapopulation models. *BMC Infectious Diseases*, 10(1).

Algemene Rekenkamer (2021). Testen op corona: Wat er in het voorjaar gebeurde. `https://www.rekenkamer.nl/binaries/rekenkamer/documenten/rapporten/2020/09/23/testen-op-corona/Focusonderzoek+Testen+op+corona.pdf`. Accessed June, 2021.

Ali, A. L. and Schmid, F. (2014). Data quality assurance for volunteered geographic information. pages 126–141.

ANP (2021). De meest gemiddelde stemmers. `https://www.rtlnieuws.nl/nieuws/nederland/artikel/5220528/uitslagen-kaart-meest-gemiddeld-partijen-door-het-land-gemeente`. Accessed June, 2021.

Bicher, M., Rippinger, C., Urach, C., Brunmeir, D., Siebert, U., and Popper, N. (2020). Evaluation of Contact-Tracing Policies Against the Spread of SARS-CoV-2 in Austria – An Agent-Based Simulation. *medRxiv*.

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7280–7287.

Byrne, A., McEvoy, D., Collins, A., Hunt, K., Casey, M., Barber, A., Butler, F., Griffin, J., Lane, E., McAloon, C., O'Brien, K., Wall, P., Walsh, K., and More, S. (2020). Inferred duration of infectious period of SARS-CoV-2: rapid scoping review and analysis of available evidence for asymptomatic and symptomatic COVID-19 cases. *BMJ Open*, 10(8):e039856.

CBS (2021a). Afstand tot dichtstbijzijnde basisschool. `https://www.cbs.nl/nl-nl/achtergrond/2017/34/afstand-tot-dichtstbijzijnde-basisschool`. Accessed June, 2021.

CBS (2021b). Afstand tot voortgezet onderwijs, 2014. `https://www.clo.nl/indicatoren/nl213006-woonafstand-tot-voortgezet-onderwijs`. Accessed June, 2021.

CBS (2021c). Arbeidsdeelname; positie huishouden. `https://opendata.cbs.nl/portal.html?_la=nl&_catalog=CBS&tableId=82956NED&_theme=8`. Accessed June, 2021.

CBS (2021d). Banen van werknemers in december; economische activiteit (sbi2008),regio. `https://opendata.cbs.nl/statline/portal.html?_la=nl&_catalog=CBS&tableId=83582NED&_theme=244`. Accessed June, 2021.

CBS (2021e). Bedrijven; bedrijfstak. `https://opendata.cbs.nl/statline/portal.html?_la=nl&_catalog=CBS&tableId=81589NED&_theme=41`. Accessed June, 2021.

CBS (2021f). Dataset: CBS Gebiedsindelingen. https://www.pdok.nl/introductie/-/article/cbs-gebiedsindelingen. Accessed June, 2020.

CBS (2021g). Het onderzoek verplaatsingen in nederland - 2013. Accessed June, 2021.

CBS (2021h). Huishoudens; personen naar geslacht, leeftijd en regio, 1 januari. https://opendata.cbs.nl/portal.html?_la=nl&_catalog=CBS&tableId=71488ned&_theme=267. Accessed June, 2021.

CBS (2021i). Kern cijfers wijken en buurten. https://opendata.cbs.nl/statline/portal.html?_la=nl&_catalog=CBS&tableId=84286NED&_theme=235. Accessed June, 2021.

CBS (2021j). Leerlingen, deelnemers en studenten; onderwijssoort, woonregio. https://opendata.cbs.nl/statline/portal.html\?_la=nl\&_catalog=CBS\&tableId=71450ned\&_theme=333. Accessed June, 2021.

CBS (2021k). Studenten gaan minder op kamers. https://www.cbs.nl/nl-nl/nieuws/2018/04/studenten-gaan-minder-op-kamers. Accessed June, 2021.

CBS (2021l). Woon-werkafstand werknemers; regio. https://opendata.cbs.nl/statline/portal.html?_la=nl&_catalog=CBS&tableId=83628NED&_theme=244. Accessed June, 2021.

CDC (2021). Pandemic planning scenarios. https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html. Accessed June, 2021.

Chen, Y., Liu, Q., and Guo, D. (2020). Emerging coronaviruses: Genome structure, replication, and pathogenesis. *Journal of Medical Virology*, 92(4):418–423.

Ciofi degli Atti, M. L., Merler, S., Rizzo, C., Ajelli, M., Massari, M., Manfredi, P., Furlanello, C., Scalia Tomba, G., and Iannelli, M. (2008). Mitigation measures for pandemic influenza in italy: An individual based model considering different scenarios. *PLOS ONE*, 3(3):1–11.

Di Stefano, B., Fuks, H., and Lawniczak, A. (2000). Object-oriented implementation of CA/LGCA modelling applied to the spread of epidemics. *2000 Canadian Conference on Electrical and Computer Engineering. Conference Proceedings. Navigating to a New Era (Cat. No.00TH8492)*.

d'Onofrio, A. (2002). Stability properties of pulse vaccination strategy in SEIR epidemic model. *Mathematical Biosciences*, 179(1):57–72.

DUO (2021). Open education data. https://duo.nl/open_onderwijsdata/databestanden/. Accessed June, 2020.

Epstein, J., Axtell, R., and Project, . (1996). *Growing Artificial Societies*. Amsterdam University Press, Amsterdam, Netherlands.

Eurydice (2021). Organisatie van primair onderwijs. https://eacea.ec.europa.eu/national-policies/eurydice/content/organisation-primary-education-32_nl. Accessed June, 2021.

Flaxman, S., Mishra, S., Gandy, A., Unwin, H., Mellan, T., Coupland, H., Whittaker, C., Zhu, H., Berah, T., Eaton, J., Monod, M., Ghani, A., Donnelly, C., Riley, S., Vollmer, M., Ferguson, N., Okell, L., and Bhatt, S. (2020). Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, 584(7820):257–261.

Furguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunuba Perez, Z., and Cuomo-Dannenburg, G. (2020). Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Technical report.

Gilbert, N. (2005). *Simulation for the Social Scientist*. Open University Press, 2 edition.

Hackl, J. and Dubernet, T. (2019). Epidemic spreading in urban areas using agent-based transportation models. 11(Future Internet).

He, S., Peng, Y., and Sun, K. (2020). SEIR modeling of the COVID-19 and its dynamics. *Nonlinear Dynamics*, 101(3):1667–1680.

Hoertel, N., Blachier, M., Blanco, C., Olfson, M., Massetti, M., Rico, M., Limosin, F., and Leleu, H. (2020). A stochastic agent-based model of the SARS-CoV-2 epidemic in France. *Nature Medicine*, 26(9):1417–1421.

Hunter, E. and Kelleher, J. (2021). Using a hybrid agent-based and equation based model to test school closure policies during a measles outbreak. *BMC Public Health*, 21(1).

Hunter, E., Namee, B., and Kelleher, J. (2018). A Comparison of Agent-Based Models and Equation Based Models for Infectious Disease Epidemiology. Technological University Dublin.

Iannelli, M., Milner, F., and Pugliese, A. (1992). Analytical and Numerical Results for the Age-Structured S-I-S Epidemic Model with Mixed Inter-Intracohort Transmission. *SIAM Journal on Mathematical Analysis*, 23(3):662–688.

Kadaster (2021). Basisregistratie Adressen en Gebouwen. https://www.kadaster.nl/zakelijk/registraties/basisregistraties/bag. Accessed June, 2020.

Kandel, N., Chungong, S., Omaar, A., and Xing, J. (2020). Health security capacities in the context of COVID-19 outbreak: an analysis of International Health Regulations annual report data from 182 countries. *The Lancet*, 395(10229):1047–1053.

Kermack, W. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721.

Kerr, C., Stuart, R., Mistry, D., Abeysuriya, R., Rosenfeld, K., Hart, G., Núñez, R., Cohen, J., Selvaraj, P., Hagedorn, B., George, L., Jastrzebski, M., Izzo, A., Fowler, G., Palmer, A., Delport, D., Scott, N., Kelly, S., Bennette, C., Wagner, B., Chang, S., Oron, A., Wenger, E., Panovska-Griffiths, J., Famulare, M., and Klein, D. (2020). Covasim: an agent-based model of COVID-19 dynamics and interventions. *medRxiv*.

McAloon, C., Collins, I., Hunt, K., Barber, A., Byrne, A., Butler, F., Casey, M., Griffin, J., Lane, E., McEvoy, D., Wall, P., Green, M., O'Grady, L., and More, S. (2020). Incubation period of COVID-19: a rapid systematic review and meta-analysis of observational research. *BMJ Open*, 10(8):e039652.

McEvoy, D., McAloon, C., Collins, A., Hunt, K., Butler, F., Byrne, A., Casey-Bryars, M., Barber, A., Griffin, J., Lane, E. A., Wall, P., and More, S. J. (2021). Relative infectiousness of asymptomatic sars-cov-2 infected persons compared with symptomatic individuals: a rapid scoping review. *BMJ Open*, 11(5).

Merler, S., Ajelli, M., and Rizzo, C. (2009). Age-prioritized use of antivirals during an influenza pandemic. *BMC Infectious Diseases*, 9(1).

Mollenhorst, G. (2009). Networks in contexts:How meeting opportunities affect personal relationships. Technical report.

Nelder, J. and Mead, R. (1965). Errata. *The Computer Journal*, 8(1):27–27.

Omori, R., Matsuyama, R., and Nakata, Y. (2020). The age distribution of mortality from novel coronavirus disease (COVID-19) suggests no large difference of susceptibility by age. *Scientific Reports*, 10(1).

OSM (2021). Openstreetmap. `https://www.openstreetmap.org`. Accessed June, 2020.

Osmosis (2021). Osmosis. `https://github.com/openstreetmap/osmosis/releases/tag/0.48.3`. Accessed June, 2020.

Osmstats (2021). Edits per country. `https://osmstats.neis-one.org/?item=countries`. Accessed June, 2021.

Papst, I., Li, M., Champredon, D., Bolker, B., Dushoff, J., and D. Earn, D. (2021). Age-dependence of healthcare interventions for COVID-19 in Ontario, Canada. *BMC Public Health*, 21(1).

Pianosi, F., Beven, K., Freer, J., Hall, J., Rougier, J., Stephenson, D., and Wagener, T. (2016). Sensitivity analysis of environmental models: A systematic review with practical workflow. *Environmental Modelling Software*, 79:214–232.

RIVM (2020). Tijdlijn van maatregelen voor bestrijding covid-19. `https://www.rivm.nl/gedragsonderzoek/tijdlijn-maatregelen-covid`. Accessed June, 2020.

RIVM (2021a). Aantal en aandeel van positieve tests stijgt. `https://www.rivm.nl/nieuws/aantal-en-aandeel-van-positieve-tests-stijgt`. Accessed June, 2021.

RIVM (2021b). Actuele covid-19 cijfers per gemeente, provincie, land en continent. https://www.databronnencovid19.nl/Bron?naam=Actuele-Covid-19-cijfers-per-gemeente\%2C-provincie\%2C-land-en-continent. Accessed June, 2021.

RIVM (2021c). Covid-19 ziekenhuisopnames (volgens nice registratie) per gemeente per ziekenhuisopnamedatum en meldingsdatum. https://data.rivm.nl/geonetwork/srv/dut/catalog.search#/metadata/4f4ad069-8f24-4fe8-b2a7-533ef27a899f. Accessed June, 2021.

RIVM (2021d). De pcr-test. https://www.rivm.nl/coronavirus-covid-19/testen/pcr-test. Accessed June, 2021.

RIVM (2021e). Naleven gedragsregels. https://www.rivm.nl/gedragsonderzoek/maatregelen-welbevinden/naleven-gedragsregels. Accessed June, 2021.

RIVM (2021f). Verspreiding covid-19. https://www.rivm.nl/nieuws/verspreiding-nieuwe-coronavirus-gaat-onverminderd-door. Accessed June, 2021.

Röper, A., Völker, B., and Flap, H. (2009). Social networks and getting a home: Do contacts matter? *Social Networks*, 31(1):40–51.

Rădulescu, A., Williams, C., and Cavanagh, K. (2020). Management strategies in a SEIR-type model of COVID 19 community spread. *Scientific Reports*, 10(1).

SCP (2021). Tijdsbestedingsonderzoek 2016 -tbo 2016. https://www.monitorarbeid.tno.nl/nl-nl/publicaties/de-impact-van-de-covid-19-pandemie-op-werknemers/. Accessed June, 2021.

Shamil, M., Farheen, F., Ibtehaz, N., Khan, I., and Rahman, M. (2021). An Agent-Based Modeling of COVID-19: Validation, Analysis, and Recommendations. *Cognitive Computation*.

Shulgin, B. (1998). Pulse vaccination strategy in the SIR epidemic model. *Bulletin of Mathematical Biology*, 60(6):1123–1148.

Silva, P., Batista, P., Lima, H., Alves, M., Guimarães, F., and Silva, R. (2020). COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions. *Chaos, Solitons Fractals*, 139:110088.

Sjödin, H., Wilder-Smith, A., Osman, S., Farooq, Z., and Rocklöv, J. (2020). Only strict quarantine measures can curb the coronavirus disease (COVID-19) outbreak in Italy, 2020. *Eurosurveillance*, 25(13).

Slob, A. (2021). Kamerbrief over groepsgrootte en leerling-leraarratio in basisonderwijs 2018. https://www.rijksoverheid.nl/onderwerpen/basisonderwijs/documenten/kamerstukken/2019/06/17/kamerbrief-over-groepsgrootte-en-leerling-leraarratio-in-basisonderwijs-2018. Accessed June, 2021.

Soltesz, K., Gustafsson, F., Timpka, T., Jaldén, J., Jidling, C., Heimerson, A., Schön, T., Spreco, A., Ekberg, J., Dahlström, R., Bagge Carlson, F., Jöud, A., and Bernhardsson, B. (2020). The effect of interventions on COVID-19. *Nature*, 588(7839):E26–E28.

Tang, B., Wang, X., Li, Q., Bragazzi, N., Tang, S., Xiao, Y., and Wu, J. (2020). Estimation of the Transmission Risk of the 2019-nCoV and Its Implication for Public Health Interventions. *Journal of Clinical Medicine*, 9(2):462.

ten Broeke, G. (2017). Sensitivity analysis methodologies for analysing emergence using agent-based models.

TNO (2021). De impact van de covid-19 pandemie op werknemers. `https://www.monitorarbeid.tno.nl/nl-nl/publicaties/de-impact-van-de-covid-19-pandemie-op-werknemers/`. Accessed June, 2021.

Truszkowska, A., Behring, B., Hasanyan, J., Zino, L., Butail, S., Caroppo, E., Jiang, Z., Rizzo, A., and Porfiri, M. (2021). High-Resolution Agent-Based Modeling of COVID-19 Spreading in a Small Town. *Advanced Theory and Simulations*, 4(3):2000277.

Wallace, R., Geller, A., and Ayano Ogawa, V. (2015). Assessing the Use of Agent-Based Models for Tobacco Regulation. *National Academies Press (US)*.

Wang, Y., Li, B., Gouripeddi, R., and Facelli, J. (2021). Human activity pattern implications for modeling SARS-CoV-2 transmission. *Computer Methods and Programs in Biomedicine*, 199:105896.

WHO (2019). *Non-pharmaceutical public health measures for mitigating the risk and impact of epidemic and pandemic influenza*. World Health Organization.

WHO (2020a). Who coronavirus (covid-19) dashboard. `https://covid19.who.int/table`. Accessed June, 2020.

WHO (2020b). Who director-general's opening remarks at the media briefing on covid-19 - 11 march 2020. `https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19`. Accessed June, 2021.

Whooz (2021). De meest gemiddelde gemeente. `https://www.whooz.nl/dashboard-meest-gemiddelde-gemeente`. Accessed June, 2021.

Willem, L., Verelst, F., Bilcke, J., Hens, N., and Beutels, P. (2017). Lessons from a decade of individual-based models for infectious disease transmission: a systematic review (2006-2015). *BMC Infectious Diseases*, 17(1).

Wu, J., Leung, K., and Leung, G. (2020). Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *The Lancet*, 395(10225):689–697.

Zhang, M. (2016). Large-scale agent-based social simulation. Technical report.

8 Appendix

8.1 Location types

Category	OSM amenity key	Comment
Home	N/A	Obtained from Dutch Cadaster data
Work	N/A	Obtained from Dutch Cadaster data
Healthcare	Clinic Dentist Doctors Hospital Pharmacy	
Shop	Supermarket Beauty Hairdresser Bank Post Office Shop	
Religion	Place Of Worship	
School	N/A	Obtained from Dutch school data
Leisure	Arts Centre Museum Theatre Library Zoo Theme Park Sauna Casino Cinema	
Public transport	N/A	Modeled as synthetic locations
Horeca	Nightclub Restaurant Fast Food Pub Café Bar Ice Cream	

8.2 Performance and coefficient of variation

The performance of the ABM is determined by simulating municipalities with varying population sizes for 7 days. The performance is evaluated in terms of time and memory usage depicted in Figure 10. The simulations are executed on a i5-4670K processor with 32 GB RAM using Java version 8. The results show that the simulation of the ABM scales linearly in computational time and memory usage for various population sizes. These results suggest that the model scales well, as the time and space complexity scale linearly in the population size and thus populations at the size of the Netherlands can be simulated by linearly increasing the computational power.

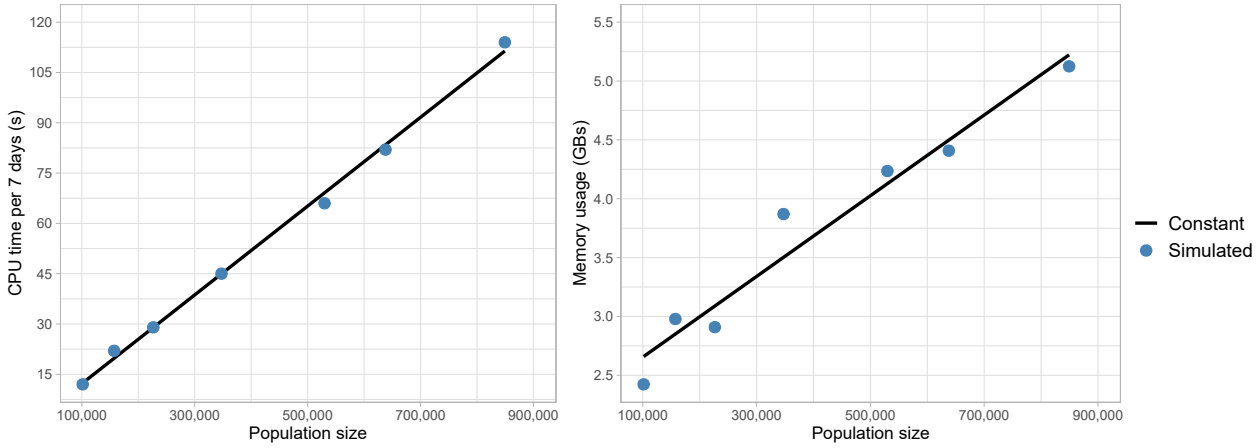

Figure 10: The ABM performance in terms of CPU computational time (left) and memory usage (right) for 7 simulated days for various population sizes.

The stability of the simulation results are analyzed by performing a number of replications with a default parameter set for the municipalities of Delft and Apeldoorn. The coefficient of variation is equal to

$$c_v = \frac{\sigma(n)}{\mu(n)}, \tag{7}$$

where $\sigma(n)$ is the standard deviation and $\mu(n)$ the mean of the output, where n indicates the number of simulations with the parameter set. This stability is investigated for varying values of n for Delft and Apeldoorn on the calibration period of 31 days for the cumulative hospitalization. Figure 11 shows that the coefficient stabilizes after 25 simulation runs for both municipalities. The difference between coefficient values for Delft and Apeldoorn can be attributed to the difference in municipality population. Apeldoorn experiences more hospitalizations, resulting in a higher $\mu(n)$, which results in a decrease in the c_v. Based on this outcome, 100 replications is deemed to be sufficient to estimate the output of the model.

The output of the model, the number of infected, hospitalized and deaths are averaged over all 100 replications of a simulation run for each simulated day. Furthermore, the percentile bootstrap method with replacement using 1000 resamples is applied to estimate the 95% confidence interval of the averages. Increasing the number of replications is computationally expensive, while applying resampling with bootstrapping is computationally inexpensive.

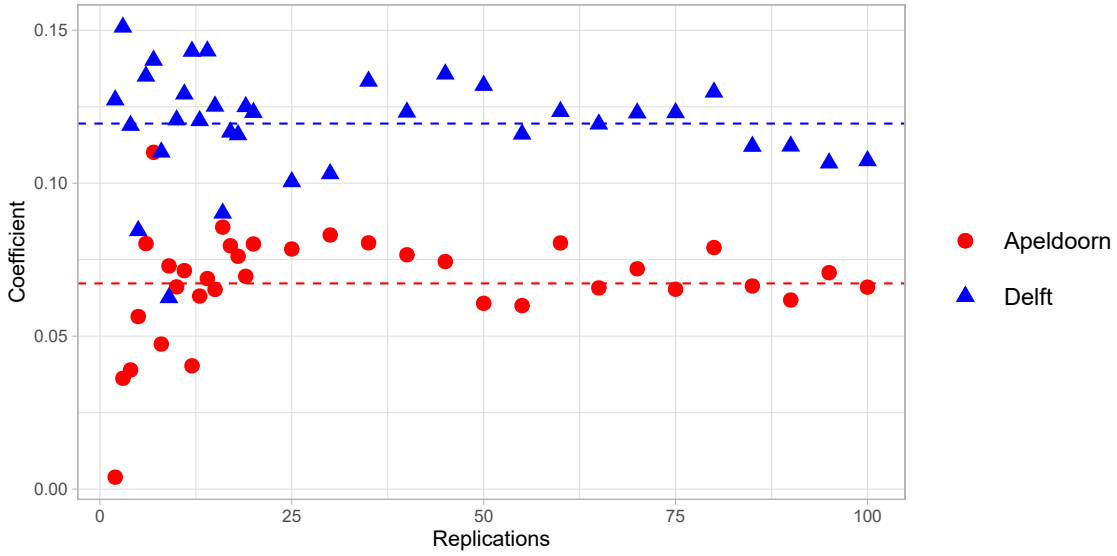

Figure 11: The coefficient of variation for both Delft and Apeldoorn.