



OPEN

## On urban maladaptation in times of epidemics

Mikhail Sirenko , Alexander Verbraeck &amp; Tina Comes

Epidemics are long-lasting and transboundary crises that challenge traditional approaches. Given the complexity and interconnectedness of modern cities, interventions can lead to unintended consequences or maladaptation. Although adaptation is central to resilience, crisis management often focuses on short-term response, leaving a gap in understanding urban adaptation and maladaptation. This study examines the impacts of uniform interventions across diverse urban districts to assess this (mal)adaptive process. We use the COVID-19 pandemic in The Hague, Netherlands, as a case study, employing a large-scale agent-based model. We find that without an intervention, the high-contact city centre becomes an infection hotspot due to the transient population it attracts. Conversely, the outer residential district, with fewer amenities, experiences infections primarily among its residents. A uniform lockdown policy significantly reduces infections in the city centre by limiting mobility and social interactions, but inadvertently increases risk in the outer residential district. Using the Urban Adaptation Index (UAI), we demonstrate that these uncontextualised policies can constitute maladaptation, confirming the unintended consequences of 'one-size-fits-all' approaches. Our results underscore this need, leading us to propose an updated, equity-oriented crisis management framework that accounts for the heterogeneous nature of modern cities.

**Keywords** Urban, Adaptation, Maladaptation, Simulation, Behaviour, Epidemic

Driven by factors like climate change and biodiversity loss, epidemics are expected to take an increasing toll in the coming years<sup>1</sup>. These events often become long-term, transboundary crises<sup>2</sup> with modern cities as their centres of gravity. However, cities are not uniform; they are increasingly heterogeneous systems<sup>3</sup>. Due to segregation and socioeconomic inequalities, two adjacent districts can be vastly different<sup>4</sup>. This heterogeneity, which concentrates vulnerable populations with pre-existing health conditions or lower incomes<sup>5,6</sup>, fundamentally challenges traditional crisis management, which has historically focused on short-term, sector-specific events.

In response to crises like the COVID-19 pandemic, governments often implement uniform, country-scale interventions, such as lockdowns. While appearing fair by applying the same rules to all, these 'one-size-fits-all' strategies often ignore the diverse urban and social fabrics that shape behaviour and vulnerability<sup>7,8</sup>. This can lead to unintended consequences, or *m maladaptation* - a concept borrowed from climate change literature referring to interventions that inadvertently worsen the situation or increase vulnerability for certain groups<sup>9,10</sup>. Instead of fostering adaptation, a uniform policy may shift risk, amplifying inequalities rather than mitigating them<sup>11</sup>.

Understanding *how* and *why* such maladaptation occurs at the urban scale requires appropriate methods, but current approaches face a methodological gap. Global or national-level studies offer broad comparisons but lack the high spatial granularity needed for effective local action<sup>12,13</sup>. Conversely, local case studies provide essential context but struggle with generalisability<sup>14</sup>. To bridge this gap, computational models are needed - specifically, large-scale spatially explicit simulation models that move beyond national averages<sup>15-18</sup>. Because virus spread is driven by human mobility and contact behaviours<sup>19,20</sup>, models must be centred on this behaviour at a high spatial and temporal resolution to be able to analyse the complex, emergent impacts of interventions<sup>21,22</sup>.

This study follows an exploratory approach<sup>23,24</sup> to develop a framework for identifying and measuring urban maladaptation during an epidemic. Instead of seeking predictions, our work aims to build the theoretical and conceptual foundations for understanding the mechanisms that drive these unintended outcomes. To do this, our research is threefold. First, we develop a large-scale agent-based model (ABM), a full-scale 'digital replica' of a city, to serve as a 'computational laboratory' for testing intervention scenarios. Second, we conceptualise how maladaptation, as defined by<sup>9</sup>, unfolds in response to an urban crisis, focusing on citizen adaptive behaviour and supplementing it with a simple and interpretable Urban Adaptation Index (UAI) to quantify how uniform policies spatially redistribute risk across diverse districts. We apply this framework to the COVID-19 pandemic

Faculty of Technology, Policy and Management, Delft University of Technology, Delft, Netherlands. email: m.sirenko@tudelft.nl

in The Hague, Netherlands. Finally, we propose an updated adaptation and equity-oriented crisis management cycle.

## Results

### Central vs outer residential: equal access but unequal distribution of essential services

The Hague, the third most populous city in the Netherlands, presents a diverse urban landscape with varying patterns of service distribution and citizen behaviour across its districts. Our analysis focuses on two exemplar but distinct districts (Fig. 1): Centrum, located in the city centre (further Central), and Ypenburg, an affluent Outer residential district (further Outer residential). Despite having comparable population sizes and areas, these districts exhibit different urban and social characteristics.

We find that the Central district, with a population of 26,833 residents over two km<sup>2</sup>, and the Outer residential district, with 21,643 residents over five km<sup>2</sup>, both have relatively high access to essential services such as grocery shops. The median distance to the closest supermarket is 0.3 km for the Central and 1.1 km for the Outer residential district. Central has a higher percentage of elderly residents, a significant proportion of individuals with chronic health conditions, and a larger segment with low income and low educational attainment. Conversely, the Outer residential district has a population with fewer chronic health conditions and psychological distress, higher income, and higher educational attainment.

The Hague's Central district serves as a hub of activity, hosting a significant proportion of the city's amenities. With 2,440 amenities (2019 data; 17% of the city's total), Central has 113 amenities per 1,000 residents, significantly higher than Outer residential with 292 amenities (2% of the total) or 11 per 1,000 residents. This uneven distribution extends to essential services such as supermarkets, where Central has ten compared to the Outer residential district with just one. This centralisation of services in the Central district is arguably designed to cater not only to its residents but also to the transient population, visitors, and tourists passing through the city centre.

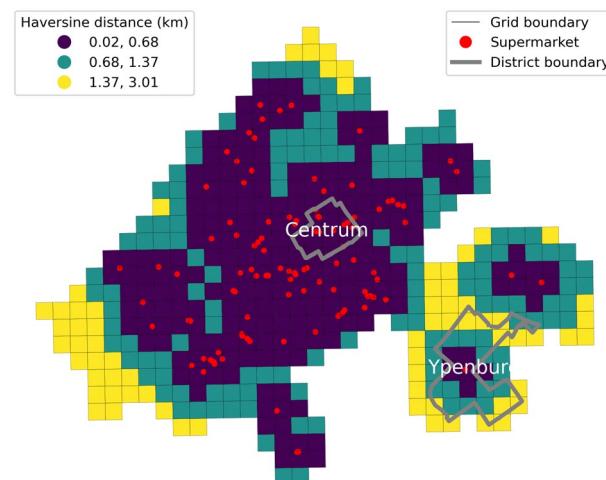
Previous research has found that the behaviour of citizens in terms of transportation and shopping habits is an important driver of the spread of a virus<sup>25,26</sup>. From Fig. 2, we see that before the pandemic, a significant proportion of grocery shopping trips by residents of highly urbanised areas in the Netherlands were made by bicycle or on foot (49%), with the remainder split between car (37%) and public transport (13%). The median distances travelled for groceries are 0.6 km on foot, 1.5 km by bicycle, and 5 km by car. These distances indicate that most shopping activities occur within the area where the resident is: either at home, work, or recreation, forming trip-chaining.

Another important aspect of the virus spread is how much time is spent with potentially infected individuals in indoor spaces<sup>27</sup>. We can see from Fig. 3 that the median time spent on grocery shopping is 30 minutes, typically involving one visit per week (44%).

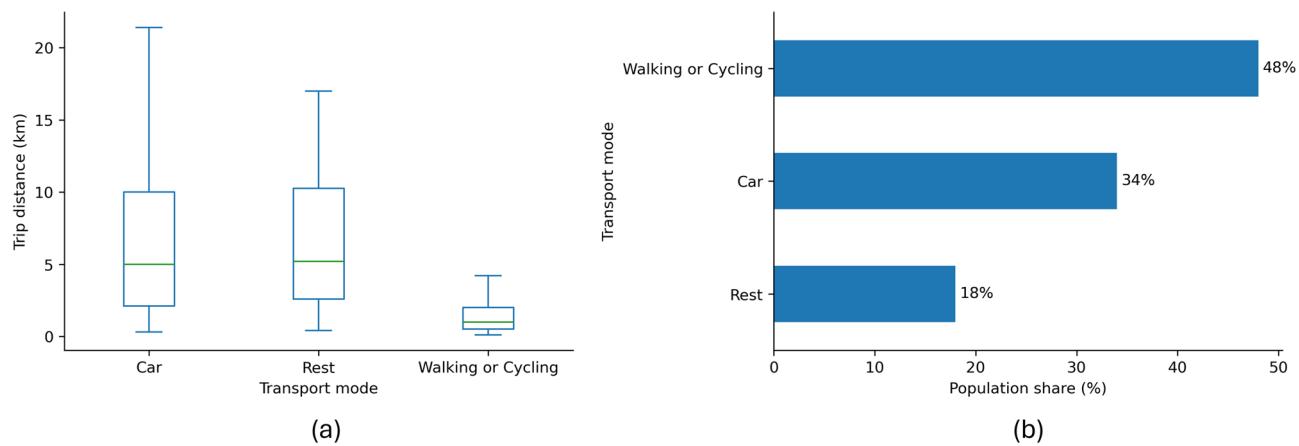
Given the two districts' very different urban and social fabric, we hypothesise that the adaptation to government interventions during COVID-19 will also differ. Understanding adaptive citizen behaviour in these districts is key to analysing potential maladaptation.

### The city centre is a hotspot under the no-response scenario

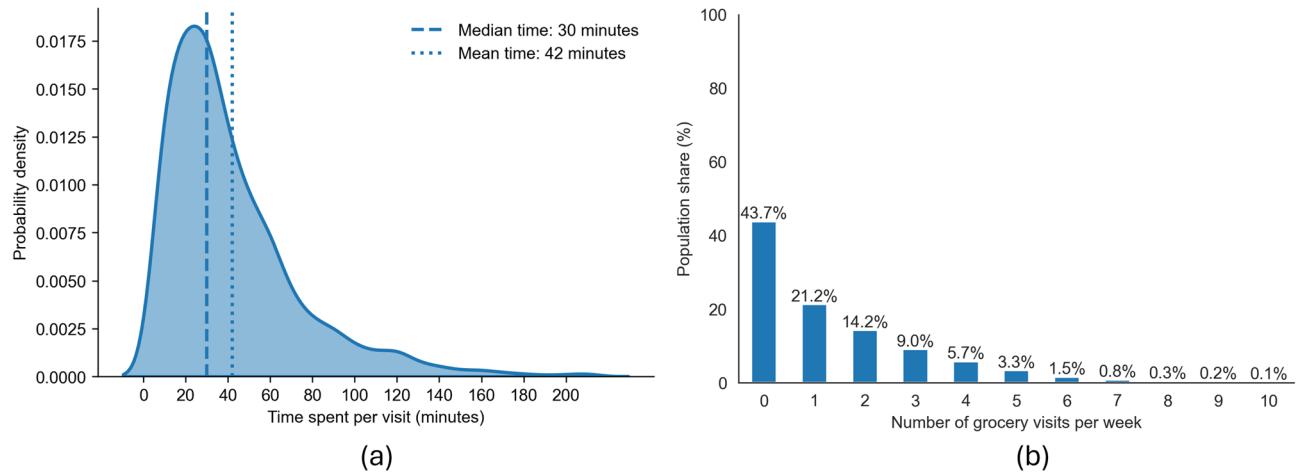
Few countries worldwide have approached the COVID-19 pandemic without imposing any restrictions on their populations<sup>28,29</sup>. Figure 4 visualises the potential consequences of such a no-restriction approach across the urban districts of The Hague. In this scenario, the city does not implement any measures to control the spread of the virus. Thus, we assume that citizens would continue their daily routines, businesses and services remain open, and other activities remain unchanged. This 'business-as-usual' approach leads to a rapid spread of the



**Fig. 1.** Supermarket proximity in Central and Outer residential districts. Red dots mark supermarkets; district boundaries are grey; grid cells are colour-shaded by Haversine distance to the nearest supermarket (shorter = closer). Black grid lines indicate analysis cells.



**Fig. 2.** Grocery trips: distance by mode and population share (pre-pandemic baseline). **(a)** Trip-distance distributions by transport mode; boxes show median, interquartile range, and whiskers at  $1.5 \times \text{IQR}$ . **(b)** Share of residents by preferred transport mode for grocery shopping.

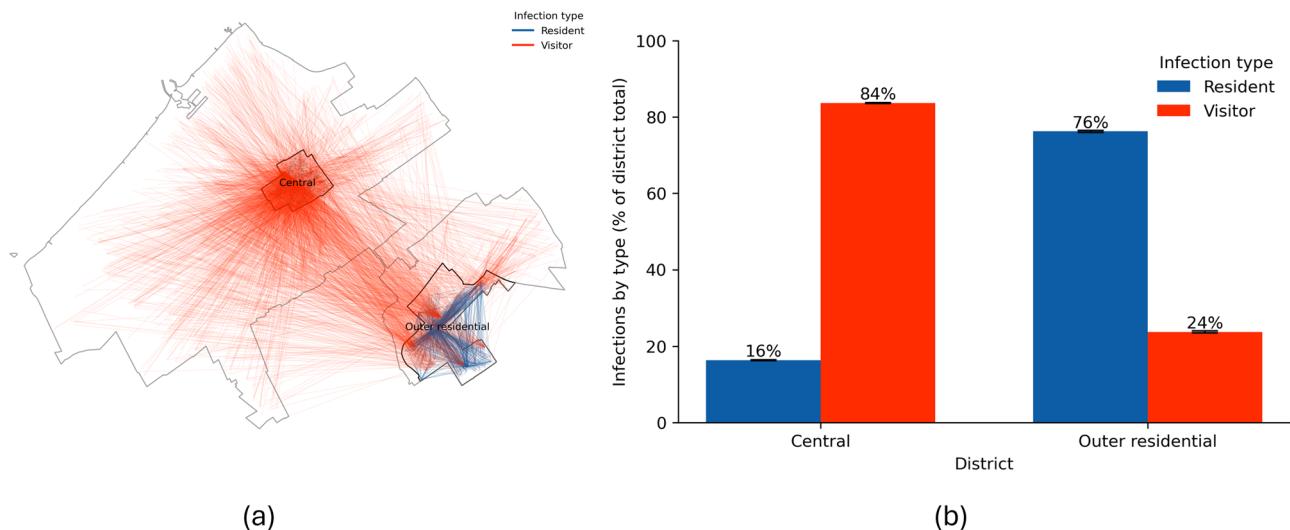


**Fig. 3.** Grocery shopping behaviour (pre-pandemic baseline). **(a)** Probability density of per-visit duration: median 30 minutes and mean 42 minutes. **(b)** Weekly visit frequency as population shares.

virus (see Supplementary Information for the temporal progression of the virus). In both districts, the spread peaks at around 25 days. Following this peak, the number of infections declines as most of the population is infected and subsequently recovers. Panel a in Fig. 4 captures a snapshot on day 25 of a single simulation run. This visualisation uses lines to connect the residences of agents with the locations where they contracted the virus. Red lines indicate that the infected agent was a visitor to the district, while blue lines denote infections occurring within the agent's district of residence. We focus on displaying infections happening only in these districts of interest. However, as can be seen from the map, the visitors come from all over the city.

Despite the similar number of residents in both districts (26,833 residents in Central, and 21,643 in Outer residential on 1 January 2020,) the patterns of virus spread exhibit striking differences. The Central district, characterised by its numerous amenities and attractions, becomes an infection hub and ranks as the first across all the districts by the number of infections (mean of 38,772 total infections or 8.19% of all infections across ten simulation model runs, see Supplementary Information for the complete overview of ranking). People from across the city travel to the Central district for work, leisure, and shopping. During these visits, some individuals contract the virus and then return to their home districts, potentially spreading the infection to household members.

Panel b demonstrates an almost opposite cumulative trend in the proportion of infections in the two study districts at the end of the model run. In the Central district, most infections (84%) occur among visitors, further driving the spread of the virus in their home districts, while only 16% of infections involve the district's residents. This is because the Central district's work locations, amenities, and other attractions draw numerous visitors, who then become vectors for the virus. In contrast, the Outer residential district shows a reverse trend. The Outer residential district, with fewer work locations, amenities, and attractions, is less frequently visited by



**Fig. 4.** Impact of the No-response on infection spread. **(a)** Snapshot on day 25 of one simulation: lines connect an infected agent's residence to the location of infection; red = *visitor infections* (infected outside home district), blue = *resident infections*. **(b)** Mean proportion of infections by agent type within each district at the end of the simulation, averaged across ten runs; error bars indicate  $\pm 1$  SD. Percentages are relative to total infections within each district.

outsiders. Consequently, social interactions primarily occur among residents. As a result, most infections (76%) in Outer residential are among its residents, while only 24% of infections involve visitors.

In summary, the No-response strategy makes the city centre a hotspot for the virus. From this central area, the virus spreads quickly to other districts, making it crucial to address the outbreak at its source to prevent a citywide epidemic. Additionally, the lack of a response exacerbates the vulnerability of the Central district's residents, who are already at risk due to a combination of socioeconomic and health factors. Their increased exposure to interactions with visitors from other districts further heightens their exposure, vulnerability, and the corresponding risk of infection.

#### Maladaptation and flipped risk from the city centre to the periphery

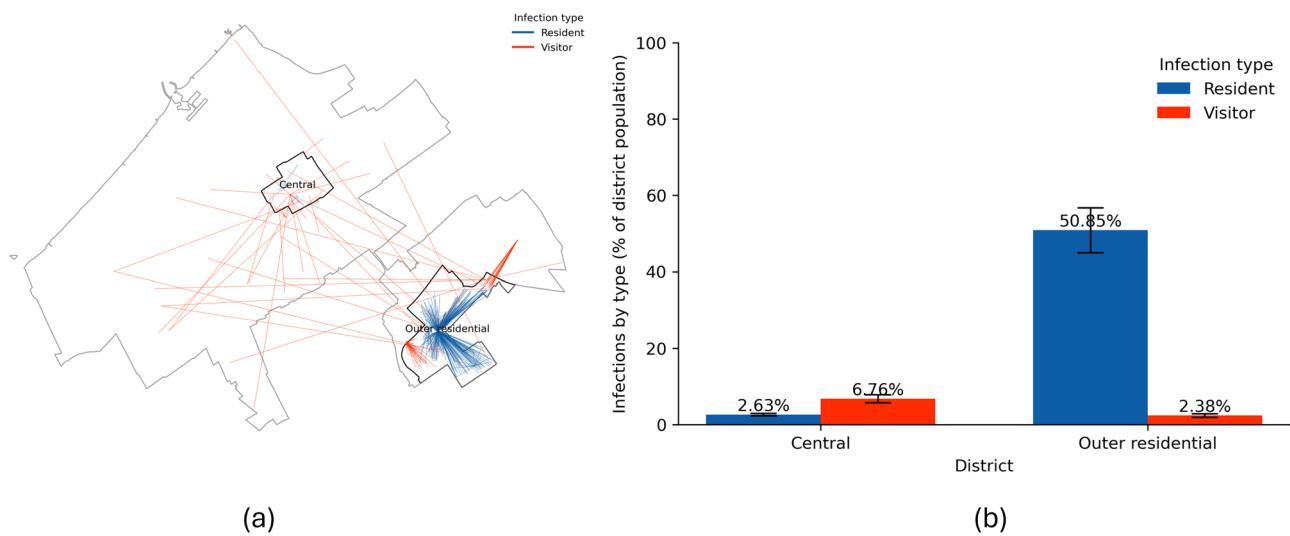
During the pandemic, the Netherlands has implemented several interventions to combat the pandemic<sup>30</sup>. All interventions were the same for the entire country. Only a few countries, e.g. Vietnam, tried to consider the difference in the spread of the virus at the urban and community scale, as opposed to the country scale, which resulted in district lockdowns.

Figure 5 illustrates the potential impact of an extended set of measures, termed a 'Hard lockdown'. This intervention involves closing schools, childcare centres, bars, and restaurants and extending to cultural and recreational amenities and non-essential shops. Essential services such as supermarkets, hospitals, police, and fire stations remain operational. The motivation behind such a combination of measures is as follows. The Netherlands' initial response to the pandemic did not occur immediately at the onset of the pandemic. The first COVID-19 case was reported on 27 February 2020, but stringent measures were not introduced until 15 March 2020. Following this, we activate the Hard lockdown on day 20 of the simulation model run. On this date, the government mandated the closing of schools, childcare facilities, bars, and restaurants. Despite these efforts, the infection rates continued to rise. Importantly, these policies were applied uniformly across the country without accounting for variations in vulnerability, behavioural patterns or social contact rates. Thus, by extending the types of amenities that must be closed, we wanted to explore whether these measures would slow down the spread of the virus.

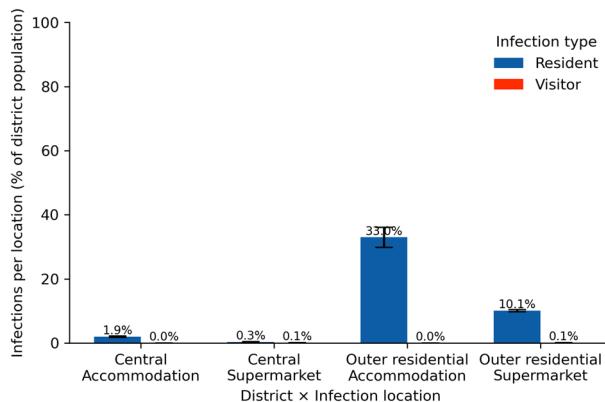
The primary benefit of the Hard lockdown is a sharp reduction in the number of infected agents (see Supplementary Information for the temporal progression under the Hard lockdown). City-wide, mean total infections fell from 473,678 to 89,852 - a reduction of 383,826 (-81.0%). The upper tail decreased as well, with P90 dropping from 474,157 to 96,924 (-79.6%), indicating substantially fewer extreme outcomes and a lower burden on healthcare services.

In the Central district, our simulation shows that infections dropped significantly under the Hard lockdown compared to the No-response scenario: from 38,773 to 2,520 (-93.5%). Visitor infections declined from 32,430 to 1,815 (-94.4%), and their share of infections decreased from 83.6% to 72.0% (-11.6 pp), indicating that reducing mobility and social interactions in densely populated areas can effectively mitigate the spread of the virus.

Conversely, in the Outer residential district, where social interactions were less frequent under the No-response scenario, and the population density is lower, the simulation shows that while the number of infections decreased in the absolute sense, the reduction was less pronounced compared to the Central district: from 26,008



**Fig. 5.** Impact of the Hard lockdown on infection spread. **(a)** Snapshot on day 25 of one simulation: lines connect an infected agent's residence to the infection location; red = *visitor infections*, blue = *resident infections*. **(b)** Mean proportion of infections by agent type within each district at the end of the simulation, averaged across ten runs; error bars indicate  $\pm 1$  SD. Percentages are relative to the total population of each district.



**Fig. 6.** Infections by location type under Hard lockdown. Bars show the mean proportion of infections at home (*Accommodation*) or in supermarkets, split by residents (blue) versus visitors (red) and district at the end of the simulation, averaged across ten runs. Error bars indicate  $\pm 1$  SD. Percentages are relative to the total population of each district.

to 11,520 infections (a drop of 55.7%). The number of infections by the district's visitors dropped significantly from 6,184 to 515 (91%); however, the residents' infections decreased only by 44.5%, from 19,822 to 11,005.

In summary, the Hard lockdown, while reducing overall city-wide infections, reveals a significant maladaptation. This uniform, 'one-size-fits-all' intervention inadvertently flips the risk profile of the city. The Central district, once the epidemic's epicentre, benefits immensely: closing its numerous amenities and workplaces starves the virus of its primary transmission routes, causing infections to decrease by over 93%. However, this same policy is far less effective in the Outer residential district. While total infections drop here as well, the reduction is much smaller, and this district surpasses the city centre in absolute infection numbers. This outcome suggests that the lockdown, by eliminating central hubs, fails to address the different social and mobility patterns in the periphery, effectively shifting the primary infection burden from the urban core to residential areas.

To understand this pattern, we need to examine the role of the essential services that remained open: supermarkets. Figure 6 shows the infections by agent and location type: supermarkets versus agent homes. The figure reveals that supermarkets are pivotal in the transmission chain. Individuals adapt their behaviour and do not visit any other location anymore. Additionally, the trip chaining stops: no more shopping on the go or close to the workplaces. As a result, individuals are more likely to get infected during their supermarket visits and subsequently spread the virus within their homes. This pattern is especially pronounced in the Outer residential district, which has a higher overall number of infections compared to the Central district, indicating that

supermarkets are crucial in both districts. However, due to the limited number of supermarkets serving Outer residential residents, these locations become infection hotspots. According to our findings, the Hard lockdown alters daily routines, making supermarkets the primary place where people, except for essential workers, interact with others. On the urban scale, supermarkets have served as transmission hubs for the virus, similar to airports' role on a national and international level<sup>31</sup>.

### Measuring urban (mal)adaptation with an index

To assess whether a policy measure is effective or maladaptive at the urban scale, we introduce the *Urban Adaptation Index (UAI)*. The UAI is designed to capture how the **spatial distribution** of policy impacts evolves over time, specifically by tracking changes in the relative burden and ranking of different districts.

Because our simulation model generates multiple realisations (runs) for each scenario, we first aggregate infection counts across runs by computing their daily mean for each district and population group. This ensures that districts with differing numbers of valid runs are weighted equally and that the resulting trajectories represent the average expected behaviour under each policy condition.

The UAI's key feature is its focus on the *absolute number* of infections in each district, which is then used to determine ranking. This approach allows us to track how the hierarchy of districts, based on their total infection burden, changes under a policy. This is distinct from measuring the *relative contribution* (or share) of each district to the citywide total. For each day  $t$  and district  $d$ , we take the mean daily infections  $y_b^d(t)$  and  $y_p^d(t)$  (calculated as described above) for both the baseline ( $b$ ) and policy ( $p$ ) scenarios. Districts are then ranked by this descending absolute number  $r_b^d(t) \in \{1, \dots, D\}$  and  $r_p^d(t) \in \{1, \dots, D\}$ , where  $r_{b,p}^d(t) = 1$  indicates the district with the highest number of mean infections, and  $D$  is the number of districts. The day-by-day change in rank quantifies how the district's position in the urban risk hierarchy evolves under the policy:

$$\Delta r^d(t) = r_b^d(t) - r_p^d(t). \quad (1)$$

A positive  $\Delta r^d(t)$  indicates that the district moved *up* the ranking (toward rank 1) under the adaptive measure, signalling potential maladaptation. This means that even if the policy reduced total infections everywhere, this district's position in the hierarchy of hotspots worsened. Conversely, a negative value indicates an improvement (the district's rank moved down, away from rank 1).

To smooth day-to-day fluctuations, we evaluate the cumulative mean:

$$\overline{\Delta r}^d(t) = \frac{1}{t} \sum_{\tau=1}^t \Delta r^d(\tau). \quad (2)$$

which represents the average change in rank up to time  $t$ . The final value  $\overline{\Delta r}^d(T)$  at the end of the evaluation window provides a summary measure of each district's adaptive trajectory:

- $\overline{\Delta r}^d(T) < 0$ : the district's average rank improved - *adaptive response*;
- $\overline{\Delta r}^d(T) \approx 0$ : the district's average rank remained stable - *neutral response*;
- $\overline{\Delta r}^d(T) > 0$ : the district climbed in the ranking - *maladaptive response*.

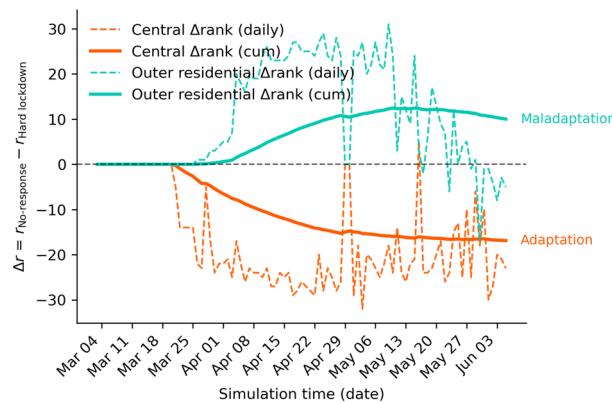
To illustrate, imagine a city with three districts: A, B, and C. In the 'No-response' baseline, their infection ranking is: 1. District A (100 cases), 2. District B (50 cases), and 3. District C (10 cases). A policy is then implemented. In the 'Policy' scenario, total infections drop, but unevenly: 1. District B (30 cases), 2. District A (20 cases), and 3. District C (5 cases). The UAI highlights this shift. District A's rank improved from 1 to 2 (a positive adaptive outcome). However, District B's rank worsened from 2 to 1. Even though District B saw an absolute reduction in cases (from 50 to 30), it became the city's new number one hotspot, indicating a potential maladaptive outcome in terms of spatial equity.

It is important to note that the UAI is not intended as a standalone metric. It is designed to reveal shifts in the spatial hierarchy of risk, not to capture the overall (absolute) effectiveness of a policy. It should therefore be used alongside other metrics that measure the potential impact of a measure.

Applied to our case city, The Hague, the UAI reveals that while the Hard lockdown reduced overall infections citywide, it altered the spatial hierarchy of risk (Fig. 7). The Central district, initially having the highest number of infections, saw a sharp decline and moved down the ranking (negative  $\overline{\Delta r}^d = -16.89$ ), reflecting strong adaptation. In contrast, the Outer residential district moved up the ranking (positive  $\overline{\Delta r}^d = +10.01$ ), meaning its position in the absolute infection hierarchy worsened relative to other districts. Although the policy was effective in aggregate, the UAI highlights a form of spatially uneven adaptation: the policy's benefits were strongest in the city core, while peripheral areas experienced a relative worsening of their rank.

### Discussion

This study examines how uncontextualised, uniform interventions during an epidemic can unintentionally result in maladaptation in heterogeneous urban environments. Our findings reveal a critical emergent pattern: a uniform 'Hard lockdown' 'flips' the risk profile of the city. The Central district (Centrum), the original epidemic hotspot due to its high density of amenities and transient visitor population, sees a significant reduction in infections. Conversely, the Outer residential district (Ypenburg), which had fewer infections in the 'No-response' scenario, becomes the new city hotspot, surpassing the centre in absolute infections.



**Fig. 7.** Temporal evolution of the Urban Adaptation Index (UAI) based on district rank changes under the Hard lockdown relative to the No-response scenario. dashed lines show the daily change in rank and solid lines the cumulative change over time:  $\Delta r(t) = r_{\text{No-response}}(t) - r_{\text{Hard lockdown}}(t)$ . Positive values indicate movement *up* the ranking under the policy (greater relative burden; maladaptation), while negative values indicate movement *down* (reduced relative burden; adaptation); the horizontal dashed line marks  $\Delta r = 0$ .

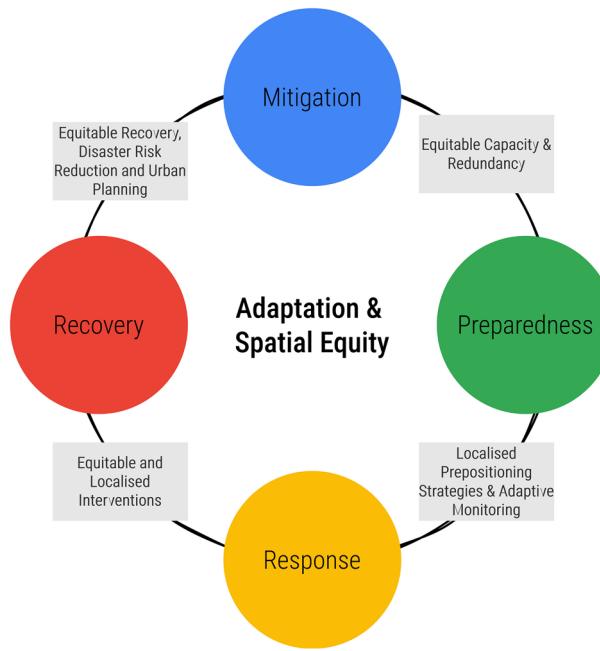
The primary mechanism for this 'flipped risk' is the clash of behavioural adaptation with inequitable local service capacity. In the Central district (Centrum), the lockdown is highly effective because it shut down the amenities and workplaces that attracted a large transient population, cutting the virus from its main transmission routes. In both districts, residents adapted their essential travel, stopping 'trip-chaining' (e.g., shopping near work) and defaulting to their closest local supermarket. This is where the maladaptation occurred. In the Outer residential district (Ypenburg), which has very few supermarkets, this rational, adaptive behaviour inadvertently concentrated the entire district's population into a handful of essential 'hubs'. These locations, with a low service capacity relative to the population they serve, became new infection hotspots, driving an increase in cases that surpassed the city centre.

Our findings align with research highlighting the heterogeneous spread of viruses in cities<sup>32,33</sup> and the role of places people visit<sup>19,21</sup>, but add a crucial distinction. Much recent literature has advocated for the '15-minute city' as a pandemic-resilient model, focusing on equal access to services<sup>34</sup> or diversity<sup>35,36</sup> to reduce infections. Our results challenge this, suggesting that local service capacity may be more critical than mere access during a crisis. While Ypenburg's residents had access to other supermarkets across the city, their adaptive behaviour, a form of 'threat rigidity' where individuals revert to the simplest, most habitual response<sup>37</sup>, drove them to the closest option. When local capacity is low, this 'rigid' individual adaptation becomes maladaptive for the district population.

This directly exposes the pitfalls of uniform policies and raises critical questions of spatial equity<sup>38</sup>. A policy that appears 'equal' by applying the same rules to all is, in fact, spatially 'inequitable'. It ignores the underlying heterogeneity that dictates how the policy is experienced, amplifying vulnerability in already underserved areas<sup>39,40</sup>. While previous work has noted that crises disproportionately affect vulnerable individuals<sup>41</sup>, our model shows how interventions can create vulnerable districts, underscoring the need for interventions centred on local capacity, not just access<sup>42</sup>.

To address this, we propose integrating adaptation and principles of spatial equity directly into the crisis management cycle (Fig. 8). This begins with conducting more equitable risk assessments during the Mitigation and Prevention phases, ensuring that all urban populations are considered and their vulnerabilities are not exacerbated. These assessments should identify and address systemic inequalities that often leave marginalised communities disproportionately exposed to risks. During the Preparedness phase, it is crucial to consider the broader social, economic, and environmental contexts that arise from existing urban inequalities. This means developing preparedness plans that are not only technically sound but also socially inclusive, taking into account the unique needs and capacities of different urban neighbourhoods. Incorporating spatial equity principles into the Response phase involves developing flexible strategies that adapt to changing conditions and behavioural responses while ensuring equitable resource allocation. This approach helps prevent the reinforcement of existing inequalities, flipping the risks from one area to another and promoting a more unified and effective crisis response. Finally, the Recovery phase should emphasise lessons learned from various social groups' behavioural responses and system vulnerabilities, focusing on addressing these disparities to build a more resilient urban environment. Recovery efforts should prioritise rebuilding in ways that enhance long-term resilience for all communities, ensuring that no group is left behind.

It is important to position this research within its methodological context. This study is exploratory: its purpose is not prediction but to build the theoretical and conceptual foundations for understanding urban maladaptation<sup>23,24</sup>. The agent-based model serves as a 'computational laboratory' fit for this purpose: to identify the underlying mechanisms and emergent patterns that result from complex interactions between policy, behaviour, and urban form. To this end, we deliberately selected a 'most different cases' design<sup>43</sup>. By comparing two districts with fundamentally different urban characteristics (central vs. peripheral, high vs. low amenity



**Fig. 8.** Adaptation and spatial equity-oriented crisis management cycle. The diagram illustrates the four critical stages - Mitigation, Preparedness, Response, and Recovery - each guided by the proposed principles that could foster successful adaptation and promote spatial equity throughout the crisis management process.

density, transient vs. residential population), we could isolate how and why the same uniform policy produces spatially divergent outcomes.

This exploratory approach necessarily has limitations. First, any model is an abstraction<sup>44</sup>. Therefore, multiple real-life objects and processes are either simplified or aggregated. We limit ourselves to modelling only citizens, their behaviour, and the businesses and services they use. Previous research has demonstrated that these components are critical for understanding the disease's spread<sup>6,21,30</sup>. However, others also stressed the importance of public transport as the key driver of COVID-19<sup>26</sup>. Second, we focus on the first wave of the pandemic when the residents' response is rather uniform. Further into the pandemic, we have observed much more complex response patterns, e.g., 'pandemic fatigue'<sup>45,46</sup>. Given the district's characteristics and the crisis phase, further work can be centred around incorporating other types of adaptive behaviour into the model.

In conclusion, our study demonstrates that uniform, uncontextualised interventions during a pandemic can lead to maladaptation in spatially heterogeneous urban environments, flipping the risk from the urban core to peripheral residential areas. This outcome is not random but a 'predictable' result of behavioural rigidity interacting with an inequitable distribution of local essential service capacity. These findings, generated from our exploratory approach, can be framed as testable propositions for future research:

- Districts with higher local capacity of essential services demonstrate greater positive adaptation (i.e., risk reduction) to lockdown-style interventions.
- The effectiveness of uniform interventions is fundamentally tied to the primary population of a district (e.g., transient visitors vs. permanent residents).
- Local service capacity, not just spatial access, is a key determinant of intervention effectiveness during a mobility-restricting crisis.

These findings underscore the critical need for adaptive, context-specific, and equity-aware crisis strategies that account for the heterogeneous nature of modern cities. By adopting such strategies, policymakers and crisis managers can build stronger urban resilience against future epidemics and other crises, ensuring more equitable outcomes for everyone.

## Methods

This section starts with a high-level overview of the model. To facilitate its understanding, we follow a common approach in the agent-based modelling studies: Overview, Design concepts, Details (ODD) protocol's summary. The rest of the section dives deeper into each of the model's 'building blocks', and Supplementary Information elaborates on each of them.

### Introduction to the model

The purpose of the model is to evaluate the impact of different interventions on the spread of the COVID-19 virus throughout urban districts. The basic idea is that to understand how an airborne virus spreads, we must simulate human behaviour at the individual level and interactions among people in cities at a high spatio-

temporal resolution<sup>20</sup>. Specifically, the model addresses the following question: What is the impact of 'inaction' (business-as-usual or No-response) and a 'hard lockdown' (only essential services remain operational) on the two exemplary districts, Central and Outer residential? That said, while focusing on two exemplary districts, we simulate an entire city. Thereby, we avoid all sorts of scaling artefacts. The model comprises four submodels: *person*, *location*, *transport*, *activity*, and *disease*. The latter consists of *transmission* and *progression* submodels. To consider our model realistic enough for its purpose, we derive patterns in daily citizen activities that form weekly routines, spatially allocated demographics and socio-economics, built environment and places of interest from existing data. We base the epidemiological submodel on peer-reviewed literature and available empirical data.

The model includes the following entities: person and location. A person entity models a resident of a city of a certain age. Each person is associated with a household of one of four types: *single* (one agent), *couple without children* (two agents), *couple with children* (more than two agents) and *single-parent* (two or more agents). Importantly, a household is not an entity but a person's attribute. People are part of one of 11 social groups (student, pensioner, etc.) based on age and other factors. Each person has a set of activities (schedules) they perform on weekdays and weekends at different locations. Depending on a person's social group, they have different activities (study, shopping, etc.) to perform and, therefore, various schedules. Assuming all family members stay together, all persons belonging to the same household are placed in a single location of *home* type. Additionally, most persons are also associated with either a *work*- or *studyplace* depending on their social group. Finally, a person has a *disease phase*, e.g. susceptible, exposed, etc., defined in line with the corresponding progression model.

The location entity models residential objects or homes (apartments or houses) and non-residential objects or places of interest (POI). Each object has a pair of coordinates and an area in m<sup>2</sup>. If it is a POI, it has a category and a subcategory: education category and school subcategory, shopping category and supermarket subcategory, etc. A POI can be indoors or outdoors, such as a park.

As for spatial and temporal scales, the model is spatially and temporally explicit and has high resolution. Persons and locations have a set of latitude and longitude pairs; therefore, the model can report back the outcomes on the building scale (opposite to the district or neighbourhood scale). However, we aggregate the outcomes on the district scale due to the nature of the input data: it is open and anonymised, but district-scale, and, thus, to avoid a false sense of precision.

The model does not have ABM's timestep concept. It follows the discrete-event formalism, and we can examine a phenomenon or count attributes at a specific moment. For example, instead of checking every 10 minutes the number of citizen-agents in a supermarket, we count the number only when an agent comes or leaves. However, for simplicity's sake, we collect and aggregate model outcomes by hour. As for the extent, the model represents a single city: The Hague, the Netherlands. By default, there are no incoming and outgoing persons from and to other cities. Thus, the model is a *closed system*.

The most important processes of the model, which are repeated every hour, are the movement of person agents between various locations and their potential infection. Each person has two sets of daily schedules: one for weekdays and one for weekends. A schedule is a combination of routine activities repeatedly executed by an agent over a day. The total number of available activities is 11. For example, sleep at home, travel by public transport, work at a university, shop at the supermarket, travel on foot, and sleep at home. Each activity has a start time, end time and location where it must be executed. The schedules differ by social group. A student agent has a different set of schedules from a pensioner. This process results in the following: multiple agents happen to simultaneously carry out an activity at the same location of a particular floor area. If there is an infectious agent among them, the susceptible agents present may also get infected.

The most important design concept of the model is the spatially and temporally explicit heterogeneous behaviour of agents. Person agents are heterogeneous to an extent and form 11 social groups, each with an associated set of weekday and weekend schedules. The model is spatially explicit, and each agent has a home, might have a work or study place and a schedule, along with the mode of transport and destination of each activity. For example, an agent of a student social group lives in a district called Central and has an activity to do in the evening: meeting with others. The model is designed to first sample the activity's start time, which in turn depends on how long the previous activities lasted. Next, the model samples the type of location: a bar, restaurant or cafe. Finally, depending on where the agent lives and other attributes (transport), it will sample the exact location and how the agent will get there. Such a design allows us to generate realistic enough dynamics of modern urban areas.

The key process in the model is how the virus spreads. The intuition behind this process is as follows. We integrate the viral load progression within an infected person over time with the likelihood of transmitting the virus to others in closed spaces. The amount of virus (viral load) a person carries changes over time, peaking and then declining, which directly impacts their contagiousness. We translate this understanding into a probability of transmission, factoring in the effectiveness of masks, the duration of exposure between individuals, and adherence to social distancing measures.

### Modelling citizen behaviour via synthetic city

The unique behaviour of each individual citizen collectively forms discernible patterns, often referred to as urban routines<sup>47</sup>. These behaviours are not only temporal (e.g., sleeping at night and working during the day) but also spatial, as people typically perform daily activities, like grocery shopping, in nearby locations, thus forming a *social contact*. Previous research highlighted the importance of social contacts in various places of interest (POI)<sup>21</sup>. Thus, we aim to simulate the virus's spread dynamics by modelling urban routines and assigning them to agents of different social groups, then situating these activities in specific spaces. The process of modelling urban routines involves five main steps: creating a synthetic population, defining locations, modelling activities, identifying transport choices, and integrating these components in time and space.

### Creating synthetic population

The foundation of any data-driven agent-based model is the synthetic population<sup>48</sup>, forming the *People* submodel. Methods for generating this population vary greatly<sup>49</sup>, with some relying on microdata for accuracy and precision and others using more aggregate or open data due to privacy concerns or data availability, as is the case in the Netherlands. Following the methodology proposed by<sup>50</sup>, we introduce a modified approach utilising open data to generate households and individuals through a stepwise approach informed by aggregate district-level statistics provided by The Hague's municipal data portal [www.denhaagincijfers.nl](https://www.denhaagincijfers.nl):

1. **Generate households.** For each district and seed, tabulated counts from district-level census tables are expanded into household records via `generate_households`, then enriched with socio-economic attributes (e.g., household size, income) by `assign_attribute_to_households`.
2. **Generate individuals.** Households are decomposed into persons based on their size with `generate_individuals`, and ages are sampled within age brackets using `assign_age_to_individuals`.
3. **Assign social group.** Each individual is classified into groups (e.g. pupil, worker, pensioner) by `assign_social_groups_to_individuals`, reflecting age-dependent participation rates.
4. **Spatial anchoring.** Homes and workplaces are assigned: `assign_home_ids` attaches home locations and `assign_workplaces` allocates workplaces based on the number of businesses and institutional establishments data.
5. **Reproducibility.** For each seed, compressed CSVs of households and individuals are saved.

The resulting synthetic population for The Hague, numbering 553,667 individual agents, is summarised in Fig. 9. This overview validates our generation process against key demographic indicators. The age distribution, shown in panel 9a, reflects a relatively young city, with the largest cohort being adults in their 20–44 years. House composition-wise, the city is dominated by 'smaller units'; while two-person households are the most common type (Fig. 9b, 25.1%), single-person households are nearly as prevalent (22.9%). When categorised by type (Fig. 9c), families with children form the largest group (42.0%), followed by single-person households (22.9%) and couples without children (21.6%). The social group (Fig. 9d) is anchored by a large working population (52.9%) and a significant pensioner group (15.1%), with the remainder composed of students, pupils, and other groups. Finally, panel 9e illustrates the spatial anchoring step, displaying the successful distribution of these synthetic residents across The Hague's districts, highlighting the population density variations from the city centre to its periphery (for comparison of real and synthetic data see Supplementary Information).

### Defining locations

Businesses and corresponding places of interest (POI) are essential for city life and disease spread<sup>21</sup>, forming the *Locations* submodel. We integrate POI data from OpenStreetMap (OSM) with local urban data from the Municipality of The Hague <https://denhaagincijfers.nl/>, detailing the number of businesses by type and their respective employee counts. Previous studies have demonstrated OSM's utility in crises, particularly in large cities like The Hague<sup>51</sup>.

More specifically, we extract all POIs within The Hague from an OSM data dump dated early January 2020. While OSM's default schema includes approximately 319 distinct POI types across categories like 'shop', 'tourism', and 'leisure', we aggregate these into 20 broader categories (e.g., 'bars and restaurants', 'colleges', see Supplementary Information for a complete overview) for the purposes of this study. Each POI is geolocated with latitude and longitude coordinates and shows significant spatial heterogeneity (Fig. 10a). For example, the city centre contains a higher concentration of supermarkets compared to other districts (Fig. 1). A similar variance is observed in business sizes; after assigning each business to a size category, we sample a specific number of employees from within that category's range, revealing notable geographic variations (Fig. 10b). Furthermore, we account for sublocations within a single POI, assuming that, for example, a location such as a school comprises multiple classrooms visited by different groups of pupils (for a complete overview of Locations' statistics, see Supplementary Information).

### Modelling activities

Citizen behaviour is captured through urban routines, which form the *Activities* submodel. We utilise the Time Use Survey (TUS) in the Netherlands<sup>52</sup>, where respondents from different social groups record their activities in 10-minute intervals over a week. This data creates a weekly calendar detailing sleep, work, study and other activities. These schedules inform the activities of agents from each social group, specifying the day and duration of each activity.

A schedule is a nested data structure defined by seven hierarchical levels: policy, epidemiological state, day of the week, social group, activity, time, and location. We construct these schedules for each of the 10 active social groups (infants are excluded and assumed to remain at home). The schedules capture distinct temporal patterns, differentiating between weekdays and weekends to reflect significant variations in daily routines. Drawing from the work of<sup>52</sup>, we model urban life using nine primary activities: personal care, travel, work or school, lunch, dinner, shopping, and social activities. Each activity is linked to a specific location type; for instance, studying must occur at a school or university (Fig. 11).

The duration of each activity is determined by stochastically sampling start and end times from distributions derived from the TUS data. The choice of location is also partially stochastic. While locations for routine commitments like work or school are sampled once at the simulation's start and remain fixed, locations for other activities like shopping or socialising are re-sampled for each occurrence. This sampling process is typically guided by proximity; an agent will likely select the closest supermarket, but may choose a random bar or restaurant within a 2.5 km radius.

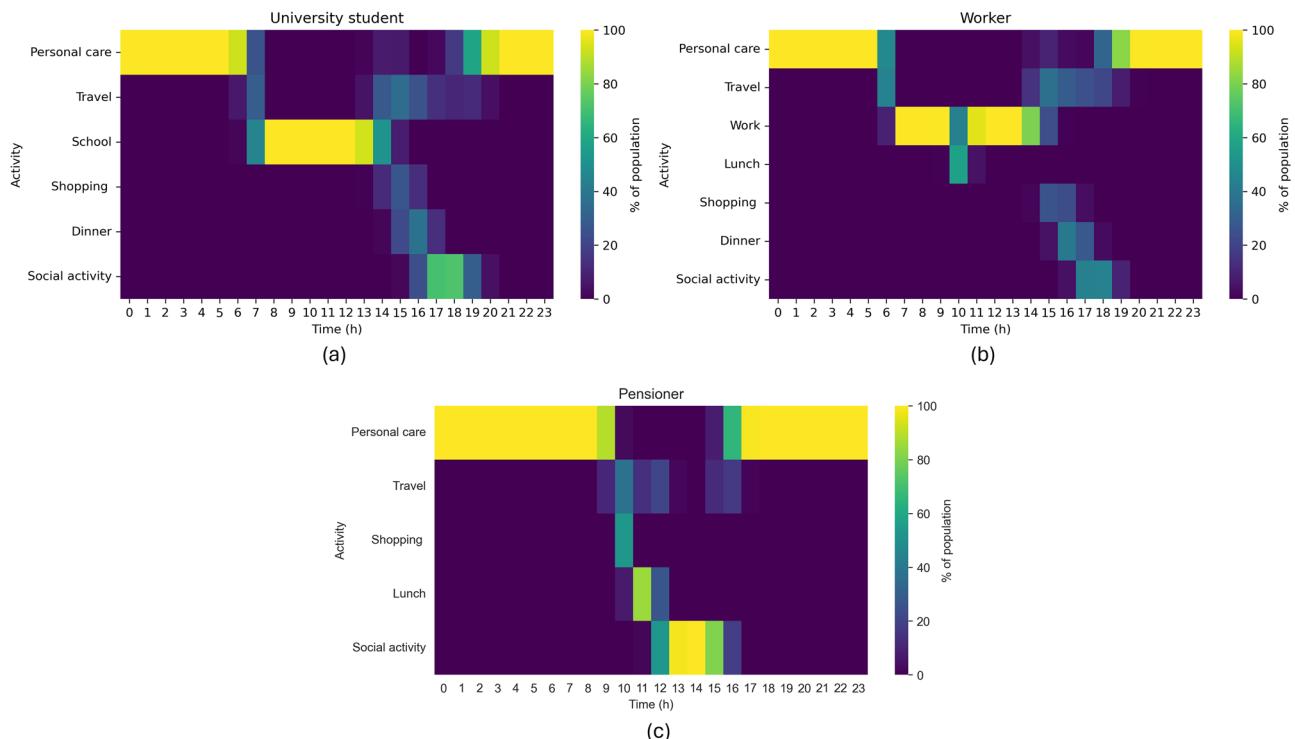


**Fig. 9.** Overview of generated synthetic population for The Hague. **(a)** Age distribution showing the number and share of individuals by age. **(b)** Distribution of household sizes. **(c)** Household type composition. **(d)** Social group structure by age, educational status and/or occupation. **(e)** Number of generated agents-residents in The Hague's districts.

Finally, this baseline schedule is dynamically overridden by two key factors. First, the active policy can alter behaviour. Under a baseline scenario with no policy intervention, agents follow their normal schedules. However, the Hard lockdown policy restricts activities to essential ones, such as grocery shopping, though essential workers also continue to work. Second, an agent's epidemiological state, aligned with an extended SEIRD model, can supersede the schedule. Agents who become 'Hospitalised' or are in the 'ICU' are removed from the general population and confined to a hospital location. Upon entering an infected state, symptomatic agents-individuals stay home, whereas asymptomatic individuals continue their normal routines (schedules for the rest of the social groups can be found in Supplementary Information).



**Fig. 10.** Non-residential locations and employment across The Hague districts. **(a)** Count of places of interest (POIs) per district. **(b)** Number of employees per district.



**Fig. 11.** Daily activity schedules (Monday) by social group in the pre-pandemic baseline. Heatmaps show the share of each group engaged in an activity by hour. **(a)** University students. **(b)** Workers. **(c)** Pensioners.

### Identifying travel patterns

Modelling realistic citizen behaviour relies critically on understanding their travel patterns. In our simulation model, these patterns are essential in connecting the *People*, *Activities*, and *Locations* submodels, and they collectively form the *Transport* submodel. We define a travel pattern as a combination of three core components: the activity an agent intends to perform, the travel mode they choose, and the distance they travel.

To parameterise these patterns with empirical data, we utilise the 2019 Onderweg in Nederland (ODiN) survey, a comprehensive national travel behaviour dataset. For this study, we zoom in on the urban travel patterns within The Hague to derive representative travel distances and population shares.

To align the data with the objectives of our epidemiological model, we performed two key simplifications. First, we aggregated the 14 original activity types found in the ODiN survey into five broader categories relevant

to disease spread (e.g., 'Work or school', 'Shopping'). Second, we collapsed the numerous transport modes into three primary groups: 'Car', 'Walking/Cycling', and a 'Rest' category. This 'Rest' group includes public transport as well as other miscellaneous modes.

Table 1 summarises the results of this data processing. It details the population share for each of our three simplified transport modes and provides key descriptive statistics for trip distance, broken down by our five aggregated activity types.

### Integrating components

The final and most critical phase of the model's construction is to integrate these distinct, data-driven submodels - *People*, *Locations*, *Activities*, and *Transport* into a single, coherent simulation of daily urban life. This integration works in two stages: static 'anchoring' at initialisation and dynamic execution at runtime.

First, we perform the static anchoring of the population. The synthetic agents generated in the *People* submodel are spatially anchored within the built environment defined by the *Locations* submodel. Each household is assigned to a specific residential building within its designated district, establishing a fixed 'home' location. Likewise, 'workers' or 'students' agents are assigned to a fixed workplace or educational institution, respectively, sampled from the available non-residential POIs.

Second, the dynamic integration takes place at runtime as agents execute their daily schedules derived from the *Activities* submodel. While 'home', 'work', and 'school' locations are fixed, other flexible activities (e.g., 'shopping', 'social activity') require a dynamic location choice. This process synthesises within the *Activities* and *Transport* submodels. When an agent initiates a flexible activity, the model does not select from a predefined list. Instead, it consults the empirical travel pattern data to sample a travel distance from that activity's specific distribution. The agent then selects a POI of the correct category (e.g., 'supermarket') from the *Locations* submodel that is reachable within that sampled distance.

This sequential process - linking *People* to *Locations* at initialisation, and then using *Activities* and *Transport* data to dynamically guide movement between them ensures that all components are fully integrated. This synthesis generates the realistic spatio-temporal mobility patterns and subsequent social contacts that are essential for simulating disease transmission.

### Disease transmission and progression models

The *Disease* submodel in our agent-based model consists of two primary components: the *transmission* model and the *progression* model. The transmission model captures the process by which an infected agent can transmit the virus to a non-infected agent within a closed space. The progression model describes how the disease progresses in an infected agent based on age and COVID-19-specific factors. Together, these models simulate the spread of COVID-19 within a city.

One of the key challenges in modelling disease transmission is accurately representing how the virus spreads in a closed environment. This task is complicated by numerous uncertainties, including factors related to droplet transmission, types of activities, and environmental conditions. While previous work, such as<sup>53</sup>, has addressed some of these issues, their models are often difficult to scale to the level of an entire city.

To address this, we extend the work of<sup>54</sup> by incorporating the newly available epidemiological insights (e.g., development of viral load over time, etc.). At the core of our transmission submodel is the infection probability  $p_i$  of person  $i$  (where persons  $j \in \{1, \dots, M_k\}$  are infectious, and persons  $i \in \{1, \dots, N_k\}$  are susceptible), which we define as follows:

$$p_i = 1 - e^{-\frac{\beta \cdot p_B \cdot t_{i,j}}{\sigma_T \cdot A_k} \sum_{j=1}^{M_k} p_j(t_e)}$$

where:

- $M_k$  represents the number of infectious individuals in the  $k$ -th location,
- $\beta$  is a correction factor for mask wearing and other personal protection measures (range: [0, 1]),
- $p_B$  is the base contagiousness parameter of the virus variant,
- $t_{i,j}$  is the time that person  $i$  and person  $j$  spent together in hours,
- $\sigma_T$  is a correction factor for ventilation and social distancing for location type  $T$  (range: (0, 1]),
- $A_k$  is the area of the  $k$ -th location in square meters,

Activity	Population share (%)			Trip distance (km)						
	Car	Rest	Walking/Cycling	Mean	SD	Min	P25	Median	P75	Max
Work or school	25.2	42.3	32.5	19.39	29.31	0.1	3.1	7.8	21.5	272.9
Shopping	33.6	19.2	47.2	6.63	14.07	0.1	1.0	2.5	5.9	151.0
Social activities	30.5	16.2	53.3	15.15	28.77	0.1	2.0	5.0	12.0	237.8
Personal care	33.7	28.0	38.3	15.48	29.94	0.1	2.0	5.0	13.0	269.4

**Table 1.** Travel behaviour summary for grocery shopping and other key activities. Left: population share by transport mode; right: trip-distance distribution (km) by activity (mean, SD, min, P25, median, P75, max). 'Rest' groups public transport and other less-common modes.

- $p_j(t_e)$  is the infectiousness of person  $j$  at  $t_e$  hours since exposure, following a triangular distribution defined by:
  - $t_{e,\min}$ : the first time after exposure when a person becomes contagious,
  - $t_{e,\text{mode}}$ : the time after exposure when a person is most contagious (peak infectiousness),
  - $t_{e,\max}$ : the last time after exposure when a person is contagious,

$$\text{where } p_j(t_e) = \begin{cases} \frac{t_e - t_{e,\min}}{t_{e,\text{mode}} - t_{e,\min}} & \text{if } t_{e,\min} \leq t_e < t_{e,\text{mode}} \\ \frac{t_{e,\max} - t_e}{t_{e,\max} - t_{e,\text{mode}}} & \text{if } t_{e,\text{mode}} \leq t_e \leq t_{e,\max} \\ 0 & \text{otherwise} \end{cases}$$

This *area-based* model simplifies the transmission dynamics by directly using the physical area of locations rather than computing interpersonal distances. It assumes uniform mixing within a (sub)location and models the infection risk as inversely proportional to the available area. The infectiousness over time follows a triangular distribution.

Once an infection event occurs between a susceptible and an infected agent, the susceptible agent may become infected, depending on the probability defined by the transmission model. Following infection, the agent transitions through various states (compartments) as the disease progresses, such as from infected to hospitalised or recovered.

We employ a modified SEIRD (Susceptible-Exposed-Infected-Recovered-Dead) model<sup>55</sup>. This model does not consider reinfection (Fig. 12), which is particularly suited for our study's focus on the first wave of the pandemic over a short-term period of 90–120 days.

A key feature of our modified SEIRD model is the distinction between symptomatic and asymptomatic infections<sup>56</sup>. This distinction is crucial because symptomatic agents are assumed to alter their behaviour by staying at home, while asymptomatic agents continue their routines, potentially spreading the virus further.

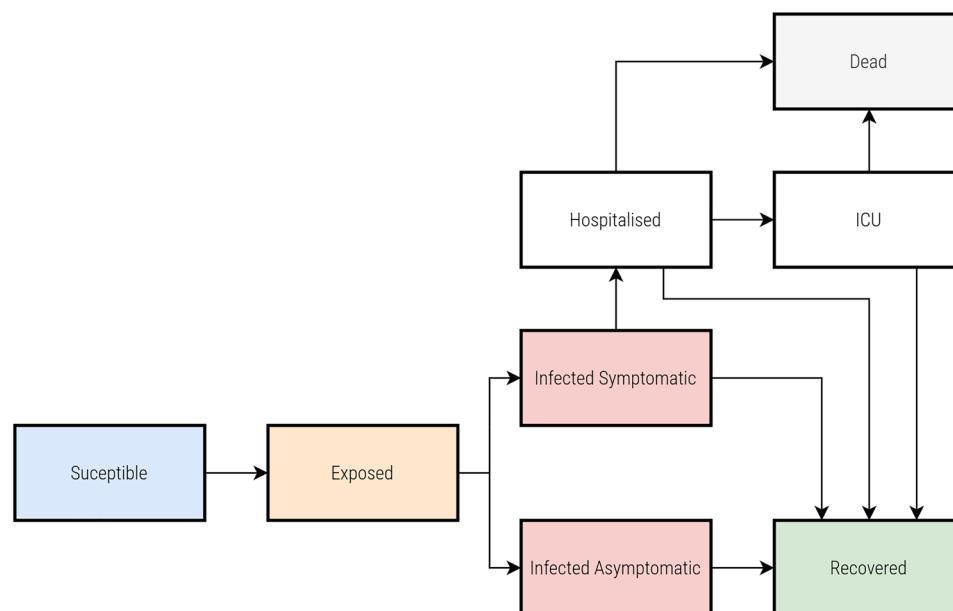
We note that the parameters for both the transmission and progression models are based on the best available data as of June 2020, including works of<sup>56–64</sup>. A key aspect of our methodology, common for large-scale, data-driven agent-based models of this nature, is that we do not perform fine-grained calibration of individual submodels; we prioritise keeping these components as close to the empirical data as possible. Such an approach is especially relevant given the significant uncertainty inherent in modelling human behaviour under an unprecedented crisis like COVID-19.

Consequently, our approach to the epidemiological model is one of validation rather than calibration. We aim to validate the emergent outcomes (i.e., infection dynamics) that arise from the interaction of these data-driven components, rather than calibrating model parameters to fit a specific observed infection curve. This validation-centric approach aligns with the model's primary purpose: to explore the dynamics of virus spread at an urban scale under different interventions, not to predict the exact number of infections.

Default values and corresponding references of each parameter can be found in the Supplementary Information.

### Experimental setup and output analysis

To evaluate the model, we examine two primary policy scenarios: (1) a 'No-response' baseline and (2) a 'Hard lockdown'. The baseline scenario simulates 'business-as-usual', where all urban locations remain open and



**Fig. 12.** A modified SEIRD model that accounts for asymptomatic infection.

Location type	No-response	Hard lockdown	Infection risk
Accommodation	Open	Open	Yes
Bar & restaurant	Open	Closed	Yes
College	Open	Closed	Yes
Fire station	Open	Open	Yes
Food & beverage	Open	Closed	Yes
Healthcare	Open	Closed	Yes
Hospital	Open	Open	No
Kindergarten	Open	Closed	Yes
Mall	Open	Closed	Yes
Park	Open	Closed	No
Pharmacy	Open	Closed	No
Police	Open	Open	Yes
Primary school	Open	Closed	Yes
Recreation	Open	Closed	Yes
Religion	Open	Closed	Yes
Retail	Open	Closed	Yes
Secondary school	Open	Closed	Yes
Supermarket	Open	Open	Yes
University	Open	Closed	Yes
Workplace	Open	Closed	Yes

**Table 2.** Operational status and infection risk by location type under No-response and Hard lockdown policy scenarios. 'Infection risk' indicates whether the model allows within-location transmission in that category.

agents follow their complete, normal activity patterns. To mimic the policy measure which was in place in the Netherlands during the first wave of the pandemic (March 2020), we activate it on day 20 of the simulation model run.

In contrast, the Hard lockdown scenario restricts all activity to essential services, which directly alters agent behaviour in two fundamental ways. First, agents stop activities related to closed amenities. For instance, students no longer attend school, and the time they would have spent elsewhere is now spent at home, increasing potential exposure among household members. Second, the lockdown induces a key behavioural adaptation by changing the spatial pattern of remaining activities. Previously, agents engaged in 'trip-chaining', such as shopping at a supermarket near their workplace. With work and school travel eliminated, agents now default to shopping at the supermarket closest to their home. This spatial redistribution of essential activities is a critical dynamic in understanding the disease's spread.

Table 2 details the operational status of each location type for both scenarios. Under the Hard lockdown, non-essential locations (e.g., schools, restaurants, non-food retail, and most workplaces) are closed, while essential services (e.g., supermarkets, pharmacies) and residential accommodations remain open and potential sites of transmission.

Please note two key considerations define the scope of our model: we exclude infections at parks, assuming open spaces and physical distancing mitigate transmission, and we do not model disease spread in hospitals and pharmacies due to their unique environment and specialised safety measures.

All simulations are initialised using epidemiological parameters for the Alpha variant of COVID-19, which was prevalent during our March-May 2020 case study period (see Supplementary Information for a complete list of parameters). To initiate the epidemic, we set 100 randomly selected agents to an 'infected' state. These agents are distributed across the city's residential areas, representing a realistic early outbreak and ensuring the simulation does not terminate prematurely due to stochastic effects.

Each simulation runs for a maximum of 120 days. We determined this duration through preliminary runs as sufficient to observe whether the epidemic wave infects a significant portion of the city or dies out. While this timeframe is not designed to capture the long-term epidemiological consequences of the virus, it is ideal for focusing on the short-term impacts of the non-pharmaceutical interventions (NPIs) on urban districts. The simulation may conclude earlier than 120 days if the virus is eradicated (i.e., no more infected agents) or if no susceptible agents remain.

The model generates a high-resolution dataset of all transmission events. For this study, we focus on four primary outputs for each infection:

- **Infection location:** The specific POI category where transmission occurred (e.g., home, supermarket, work).
- **Infection district:** The district where the infection location is situated.
- **Residence district:** The home district of the newly infected agent.
- **Infection time:** The simulation day on which the infection occurred.

This analytical setup allows us to differentiate between infections of 'residents' (where the infection district and residence district are the same) and 'visitors' (where an agent is infected in a district other than their own). We focus our analysis on two representative districts of The Hague: one central business district, Centrum ('Central'), and one 'Outer residential' Ypenburg district. To ensure robustness and account for stochastic variability, all reported outcomes represent the average of ten simulation runs, and we include minimum/maximum ranges and error bars where applicable.

### On the validity of the synthetic city

Assessing the validity of any simulation model is a well-known challenge<sup>65,66</sup>, especially when modelling complex human systems and unprecedented scenarios, such as a novel pandemic, for which historical validation data does not exist<sup>67,68</sup>. This challenge is amplified when building a model at the scale of an entire city. Many computational models developed during the COVID-19 pandemic operated at a much smaller scale, for example, simulating a single district or a stylised network, which can introduce scaling artefacts and, more importantly, cannot capture the crucial inter-district dynamics that define urban life<sup>69,70</sup>.

Given these challenges, our take on validity is pragmatic and threefold, rooted in the principles of exploratory modelling for policy analysis<sup>24</sup>. Unlike conventional simulation modelling that aims for precise prediction or optimal solutions, exploratory modelling seeks to understand the range of possible outcomes under deep uncertainty. During the early stages of COVID-19, uncertainty was profound: from the virus's epidemiological characteristics to human behavioural responses and the long-term impacts of interventions<sup>71,72</sup>. While models that prescribe optimal solutions are useful in stable conditions<sup>73</sup>, our setup focuses on exploring 'what if' scenarios to identify robust policies and potential vulnerabilities.

Our take on validity supports this exploratory goal and is threefold. First, the model is fundamentally data-driven. As detailed in the preceding sections, it is composed of multiple submodels (e.g., *People*, *Locations*, *Activities*) that are independently parameterised using the best available empirical data (e.g., spatial statistics, time-use and travel surveys). We have demonstrated that each of these components replicates real-world statistics and patterns to a high degree.

Second, we focus on validating the model's emergent properties. Rather than aiming for point-prediction (e.g., matching the exact number of cases on a specific day), we ensure that the macro-level patterns generated by the simulation, such as mobility flows, activity hotspots, are plausible and resonate with empirical findings reported across various real-world sites.

Finally, a crucial component of our model's validity is its scale. By simulating the entire city of The Hague, we explicitly avoid the artefacts that arise from modelling a single district in isolation. This comprehensive scale is what allows us to study the heterogeneous impacts of uniform policies and the complex interactions between different urban districts (e.g., a central business district and an outer residential one).

By explicitly modelling the heterogeneity of modern cities, with diverse social groups spread across distinct urban districts, this model serves as a useful tool for estimating the potential district-scale impacts of various interventions. Its primary purpose is not to predict the spread of a virus following an intervention. Instead, it functions as a 'computational laboratory' to study the interrelation between interventions and their intended, as well as unintended, effects<sup>74</sup>.

### Key assumptions and limitations of the model

Like any simulation model, the one presented in this study is an abstraction of reality and is based on a set of simplifying assumptions<sup>44</sup>. It is important to acknowledge these limitations, as they define the scope and interpretation of our findings.

First, the model's structure involves several key boundaries. We simulate The Hague as a 'closed system', meaning we do not account for agent mobility to or from other cities. While this approach avoids scaling artefacts by modelling the entire city, it overlooks the significant commuter and visitor flows that characterise a major urban centre, which could influence infection dynamics. Furthermore, our in-depth analysis focuses on two representative districts. While this 'most different cases' design is intentional for exploring mechanisms, the findings may not be generalizable to all 44 districts within The Hague, each with its own unique urban and social fabric.

Second, our model of citizen behavior relies on necessary simplifications. The synthetic population is generated from aggregate, district-level open data rather than individual microdata, which inherently limits its granularity. Agent schedules are based on pre-pandemic (2019) time-use and travel surveys. This provides a robust baseline but does not capture any voluntary behavioural changes that may have occurred in the pandemic's earliest days before a formal lockdown. We also assume a binary response to illness: symptomatic agents stay home, while asymptomatic agents continue their normal routines. This omits more nuanced behaviours, such as partial self-isolation.

Third, the epidemiological and spatial models have specific exclusions. The disease progression model is a modified SEIRD that does not account for reinfection, a simplification considered acceptable for the short 120-day 'first wave' scenario. Spatially, we excluded transmission in open spaces like parks (assuming low risk) and in hospitals (due to their unique protocols and non-public nature). Moreover, while public transport is implicitly included in the 'Rest' travel category, it is not explicitly modelled as a distinct, high-contact transmission environment, which other studies have identified as an important driver of spread<sup>26</sup>.

Finally, our experimental design is exploratory. We focus on two starkly different policy scenarios (No-response and Hard lockdown) over a 120-day period. This timeframe is suitable for assessing the short-term impacts of interventions but precludes the analysis of long-term dynamics such as 'pandemic fatigue'<sup>45</sup>, waning immunity, or the impact of subsequent waves and variants.

## Data availability

The datasets generated and analysed during the current study are available at 4TU.ResearchData portal (<https://doi.org/10.4121/33c01ff0-d3af-4293-8690-339bbca2bb37>).

## Code availability

Current study is available in the Medlabs-HERoS GitHub repository (<https://github.com/averbraeck/medlabs-heros>). The software used to analyse generated datasets in the current study is available in the urban-maladaptation GitHub repository (<https://github.com/miskh/urban-maladaptation>).

Received: 24 September 2024; Accepted: 16 December 2025

Published online: 06 January 2026

## References

1. Marani, M., Katul, G. G., Pan, W. K. & Parolari, A. J. Intensity and frequency of extreme novel epidemics. *Proceedings of the National Academy of Sciences* **118**, e2105482118. <https://doi.org/10.1073/pnas.2105482118> (2021).
2. Comes, M. *et al.* Strategic crisis management in the european union: Evidence review report. Tech. Rep., SAPEA ( Publisher: SAPEA, 2022).
3. Brelsford, C., Lobo, J., Hand, J. & Bettencourt, L. M. A. Heterogeneity and scale of sustainable development in cities. *Proceedings of the National Academy of Sciences* **114**, 8963–8968. <https://doi.org/10.1073/pnas.1606033114> (2017).
4. Nicoletti, L., Sirenko, M. & Verma, T. Disadvantaged communities have lower access to urban infrastructure. *Environment and Planning B: Urban Analytics and City Science* **50**, 831–849. <https://doi.org/10.1177/23998083221131044>. (2023) (Publisher: SAGE Publications Ltd STM).
5. Chen, M. *et al.* Rising vulnerability of compound risk inequality to ageing and extreme heatwave exposure in global cities. *npj Urban Sustainability* **3**, 38 (2023). <https://www.nature.com/articles/s42949-023-00118-9>.
6. Kalocsányiová, E., Essex, R. & Fortune, V. Inequalities in Covid-19 Messaging: A Systematic Scoping Review. *Health Communication* **38**, 2549–2558. <https://doi.org/10.1080/10410236.2022.2088022> (2023).
7. Tintori, A., Cerbara, L., Ciancimino, G. & Crescimbene, M. Adaptive behavioural coping strategies as reaction to COVID-19 social distancing in Italy. *European Review for Medical and Pharmacological Sciences* (2020).
8. Duque-Calvache, R., Torrado, J. M. & Mesa-Pedrazas, A. Lockdown and adaptation: residential mobility in Spain during the COVID-19 crisis. *European Societies* **23**, S759–S776. <https://doi.org/10.1080/14616696.2020.1836386> (2021).
9. Schipper, E. L. F. Maladaptation: When Adaptation to Climate Change Goes Very Wrong. *One Earth* **3**, 409–414 (2020). <https://linkinghub.elsevier.com/retrieve/pii/S259033220304838>.
10. Ahmed, F., Ahmed, N., Pissarides, C. & Stiglitz, J. Why inequality could spread COVID-19. *The Lancet Public Health* **5**, e240 (2020). <https://linkinghub.elsevier.com/retrieve/pii/S2468266720300852>.
11. Nundy, S., Ghosh, A., Mesloub, A., Albaqwy, G. A. & Alnaim, M. M. Impact of COVID-19 pandemic on socio-economic, energy-environment and transport sector globally and sustainable development goal (SDG). *Journal of Cleaner Production* **312**, 127705 (2021). <https://linkinghub.elsevier.com/retrieve/pii/S0959652621019235>.
12. IFRC. World Disasters Report: Trust, Equity and Local Action. Tech. Rep., International Federation of Red Cross and Red Crescent Societies, Geneva (2023). OCLC: 1381186351.
13. UNDRR. GAR Special Report: Measuring Resilience for the Sustainable Development Goals. Tech. Rep., United Nations Office for Disaster Risk Reduction, Geneva (2023).
14. Vega, D. Case studies in humanitarian logistics research. *Journal of Humanitarian Logistics and Supply Chain Management* **8**, 134–152 (2018). <https://www.emerald.com/insight/content/doi/10.1108/JHSLCM-01-2018-0005/full/html>.
15. Truszkowska, A. *et al.* High-Resolution Agent-Based Modeling of COVID-19 Spreading in a Small Town. *Advanced Theory and Simulations* **4**, 2000277 (2021). <https://pmc.ncbi.nlm.nih.gov/articles/PMC7995144>.
16. Lopes, P. H. *et al.* Measuring the impact of nonpharmaceutical interventions on the SARS-CoV-2 pandemic at a city level: An agent-based computational modelling study of the City of Natal. *PLOS Global Public Health* **2**, e0000540 (Publisher: Public Library of Science, 2022). <https://journals.plos.org/globalpublichealth/article?id=10.1371/journal.pgph.0000540>.
17. Tatapudi, H., Das, R. & Das, T. K. Impact assessment of full and partial stay-at-home orders, face mask usage, and contact tracing: An agent-based simulation study of COVID-19 for an urban region. *Global Epidemiology* **2**, 100036 (2020). <https://www.sciedirect.com/science/article/pii/S2590113320300201>.
18. Dong, T., Dong, W. & Xu, Q. Agent Simulation Model of COVID-19 Epidemic Agent-Based on GIS: A Case Study of Huangpu District, Shanghai. *International Journal of Environmental Research and Public Health* **19**, 10242 (2022).
19. Kato, H. & Takizawa, A. Human mobility and infection from Covid-19 in the Osaka metropolitan area. *npj Urban Sustainability* **2**, 20 (2022). <https://www.nature.com/articles/s42949-022-00066-w>.
20. Zhang, M., Verbraeck, A., Meng, R., Chen, B. & Qiu, X. Modeling Spatial Contacts for Epidemic Prediction in a Large-Scale Artificial City. *Journal of Artificial Societies and Social Simulation* **19**, 3 (2016). <http://jasss.soc.surrey.ac.uk/19/4/3.html>.
21. Liu, L. *et al.* Infectiousness of places – Impact of multiscale human activity places in the transmission of COVID-19. *npj Urban Sustainability* **2**, 28 (2022). <https://www.nature.com/articles/s42949-022-00074-w>.
22. Lorig, F., Johansson, E. & Davidsson, P. Agent-Based Social Simulation of the Covid-19 Pandemic: A Systematic Review. *Journal of Artificial Societies and Social Simulation* **24**, 5 (2021). <http://jasss.soc.surrey.ac.uk/24/3/5.html>.
23. Eisenhardt, K. M. & Graebner, M. E. Theory Building From Cases: Opportunities And Challenges. *Academy of Management Journal* <https://doi.org/10.5465/amj.2007.24160888>. (2007) (Publisher: Academy of Management Briarcliff Manor, NY 10510).
24. Bankes, S., Walker, W. E. & Kwakkel, J. H. Exploratory Modeling and Analysis. In *Encyclopedia of Operations Research and Management Science*, 532–537 (Springer US, Boston, MA, 2013). [http://link.springer.com/10.1007/978-1-4419-1153-7\\_314](http://link.springer.com/10.1007/978-1-4419-1153-7_314).
25. Guo, Y., An, S. & Comes, T. From warning messages to preparedness behavior: The role of risk perception and information interaction in the Covid-19 pandemic. *International Journal of Disaster Risk Reduction* **73**, 102871 (2022). <https://linkinghub.elsevier.com/retrieve/pii/S2212420922000905>.
26. Tapiador, L., Gomez, J. & Vassallo, J. M. Exploring the relationship between public transport use and COVID-19 infection: A survey data analysis in Madrid Region. *Sustainable Cities and Society* **104**, 105279 (2024). <https://linkinghub.elsevier.com/retrieve/pii/S2210670724001070>.
27. Atamer Balkan, B. *et al.* The multi-dimensional challenges of controlling respiratory virus transmission in indoor spaces: Insights from the linkage of a microscopic pedestrian simulation and SARS-CoV-2 transmission model. *PLOS Computational Biology* **20**, e1011956 (2024). <https://dx.plos.org/10.1371/journal.pcbi.1011956>.
28. Åslund, A. Responses to the COVID-19 crisis in Russia, Ukraine, and Belarus. *Eurasian Geography and Economics* **61**, 532–545. <https://doi.org/10.1080/15387216.2020.1778499> (2020).
29. Ludvigsson, J. F. How Sweden approached the COVID-19 pandemic: Summary and commentary on the National Commission Inquiry. *Acta Paediatrica* **112**, 19–33. <https://doi.org/10.1111/apa.16535> (2023).

30. Van Den Broek-Altenburg, E. & Atherly, A. Adherence to COVID-19 policy measures: Behavioral insights from The Netherlands and Belgium. *PLOS ONE* **16**, e0250302 (2021). <https://dx.plos.org/10.1371/journal.pone.0250302>.
31. Silva, G. C. & Ribeiro, E. M. S. The impact of Brazil's transport network on the spread of COVID-19. *Scientific Reports* **13**, 2240 (2023). <https://www.nature.com/articles/s41598-022-27139-1>.
32. Plümper, T. & Neumayer, E. The pandemic predominantly hits poor neighbourhoods? SARS-CoV-2 infections and COVID-19 fatalities in German districts. *European Journal of Public Health* **30**, 1176–1180 (2020). <https://academic.oup.com/eurpub/article/30/6/1176/5894972>.
33. Mena, G. & Aburto, J. M. Unequal impact of the COVID-19 pandemic in 2020 on life expectancy across urban areas in Chile: a cross-sectional demographic study. *BMJ Open* **12**, e059201 (Publisher: British Medical Journal Publishing Group Section: Epidemiology, 2022). <https://bmjopen.bmjjournals.org/content/12/8/e059201>.
34. Guzman, L. A., Oviedo, D. & Cantillo-Garcia, V. A. Is proximity enough? A critical analysis of a 15-minute city considering individual perceptions. *Cities* **148**, 104882 (2024). <https://linkinghub.elsevier.com/retrieve/pii/S0264275124000969>.
35. Falagara Sigala, I., Sirenko, M., Comes, T. & Kovács, G. Mitigating personal protective equipment (PPE) supply chain disruptions in pandemics – a system dynamics approach. *International Journal of Operations & Production Management* **42**, 128–154 (2022). <https://www.emerald.com/insight/content/doi/10.1108/IJOPM-09-2021-0608/full.html>.
36. Walker, B. et al. Response diversity as a sustainability strategy. *Nature Sustainability* **6**, 621–629 (2023). <https://www.nature.com/articles/s41893-022-01048-7>.
37. Staw, B. M., Sandelands, L. E. & Dutton, J. E. Threat Rigidity Effects in Organizational Behavior: A Multilevel Analysis. *Administrative Science Quarterly* **26**, 501 (1981). <https://www.jstor.org/stable/2392337?origin=crossref>.
38. Talen, E. & Anselin, L. Assessing Spatial Equity: An Evaluation of Measures of Accessibility to Public Playgrounds. *Environment and Planning A: Economy and Space* **30**, 595–613. <https://doi.org/10.1068/a300595>. (1998) (Publisher: SAGE Publications Ltd).
39. Perry, B. L., Aronson, B. & Pescosolido, B. A. Pandemic precarity: COVID-19 is exposing and exacerbating inequalities in the American heartland. *Proceedings of the National Academy of Sciences* **118**, e2020685118. <https://doi.org/10.1073/pnas.2020685118> (2021).
40. Elers, C., Jayan, P., Elers, P. & Dutta, M. J. Negotiating Health Amidst COVID-19 Lockdown in Low-income Communities in Aotearoa New Zealand. *Health Communication* **36**, 109–115. <https://doi.org/10.1080/10410236.2020.1848082>. (2021) (Publisher: Routledge).
41. Chen, J. T. & Krieger, N. Revealing the Unequal Burden of COVID-19 by Income, Race/Ethnicity, and Household Crowding: US County Versus Zip Code Analyses. *Journal of Public Health Management and Practice* **27**, S43 (2021). [https://journals.lww.com/jphmp/fulltext/2021/01001/revealing\\_the Unequal\\_burden\\_of\\_covid\\_19\\_by8.aspx](https://journals.lww.com/jphmp/fulltext/2021/01001/revealing_the Unequal_burden_of_covid_19_by8.aspx).
42. Felicotti, A., Romice, O. & Porta, S. Design for Change: Five Proxies for Resilience in the Urban Form. *Open House International* **41**, 23–30 (2016). <https://www.emerald.com/insight/content/doi/10.1108/OHI-04-2016-B0004/full.html>.
43. Seawright, J. & Gerring, J. Case Selection Techniques in Case Study Research: A Menu of Qualitative and Quantitative Options. *Political Research Quarterly* **61**, 294–308. <https://doi.org/10.1177/1065912907313077>. (2008) (Publisher: SAGE Publications Inc).
44. Thompson, E. L. & Smith, L. A. Escape from model-land. *Economics* **13**, 20190040 (2019). <https://www.degruyter.com/document/doi/10.5018/economics-ejournal.ja.2019-40/html>.
45. Petherick, A. et al. A worldwide assessment of changes in adherence to COVID-19 protective behaviours and hypothesized pandemic fatigue. *Nature Human Behaviour* **5**, 1145–1160 (2021). <https://www.nature.com/articles/s41562-021-01181-x>.
46. MacIntyre, C. R. et al. Mask use, risk-mitigation behaviours and pandemic fatigue during the COVID-19 pandemic in five cities in Australia, the UK and USA: A cross-sectional survey. *International Journal of Infectious Diseases* **106**, 199–207 (2021). <https://www.sciencedirect.com/science/article/pii/S1201971221002745>.
47. Champlin, C., Sirenko, M. & Comes, T. Measuring social resilience in cities: An exploratory spatio-temporal analysis of activity routines in urban spaces during Covid-19. *Cities* **135**, 104220 (2023). <https://www.sciencedirect.com/science/article/pii/S026427512300032X>.
48. Zhou, M., Li, J., Basu, R. & Ferreira, J. Creating spatially-detailed heterogeneous synthetic populations for agent-based microsimulation. *Computers, Environment and Urban Systems* **91**, 101717 (2022). <https://www.sciencedirect.com/science/article/pii/S0198971521001241>.
49. Chapuis, K., Tailandier, P. & Drogoul, A. Generation of Synthetic Populations in Social Simulations: A Review of Methods and Practices. *Journal of Artificial Societies and Social Simulation* **25**, 6 (2022).
50. Ge, Y., Meng, R., Cao, Z., Qiu, X. & Huang, K. Virtual city: An individual-based digital environment for human mobility and interactive behavior. *SIMULATION* **90**, 917–935. <https://doi.org/10.1177/0037549714531061>. (2014) (Publisher: SAGE Publications Ltd STM).
51. Gebremedhin, E. T., Basco-Carrera, L., Jonoski, A., Iliffe, M. & Winsemius, H. Crowdsourcing and interactive modelling for urban flood management. *Journal of Flood Risk Management* **13**, e12602. <https://doi.org/10.1111/jfr3.12602> (2020).
52. Roeters, A. Time use in the Netherlands. Tech. Rep. 2, The Netherlands Institute for Social Research (2019). <https://digital.scp.nl/timeuse2/about-the-time-use-survey>
53. Li, X. et al. A spatiotemporally resolved infection risk model for airborne transmission of COVID-19 variants in indoor spaces. *Science of The Total Environment* **812**, 152592 (2022). <https://linkinghub.elsevier.com/retrieve/pii/S0048969721076701>.
54. Stroud, P., Del Valle, S., Sydoriak, S., Riese, J. & Mniszewski, S. Spatial Dynamics of Pandemic Influenza in a Massive Artificial Society. *Journal of Artificial Societies and Social Simulation* **10** (2007).
55. Menda, K., Laird, L., Kochenderfer, M. J. & Caceres, R. S. Explaining COVID-19 outbreaks with reactive SEIRD models. *Scientific Reports* **11**, 17905 (2021). <https://www.nature.com/articles/s41598-021-97260-0>.
56. Nishiura, H. et al. Estimation of the asymptomatic ratio of novel coronavirus infections (COVID-19). *International Journal of Infectious Diseases* **94**, 154–155 (2020). <https://linkinghub.elsevier.com/retrieve/pii/S1201971220301399>.
57. Chu, D. K. et al. Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: a systematic review and meta-analysis. *The Lancet* **395**, 1973–1987 (2020). <https://www.thelancet.com/article/S0140-67362031142-9/fulltext>. Publisher: Elsevier.
58. De Bruin, J., Voorvaart, R., Menger, V., Kocken, I. & Phil, T. Novel Coronavirus (COVID-19) Cases in The Netherlands (2020). <https://zenodo.org/records/4278891>.
59. Ferguson, N. et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Tech. Rep. (2020). <http://spiral.imperial.ac.uk/handle/10044/1/77482>. Publisher: Imperial College London.
60. He, X. et al. Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine* **26**, 672–675 (2020). <https://www.nature.com/articles/s41591-020-0869-5>. Publisher: Nature Publishing Group.
61. Linton, N. M. et al. Incubation Period and Other Epidemiological Characteristics of 2019 Novel Coronavirus Infections with Right Truncation: A Statistical Analysis of Publicly Available Case Data. *Journal of Clinical Medicine* **9**, 538 (2020).
62. Setti, L. et al. Airborne Transmission Route of COVID-19: Why 2 Meters/6 Feet of Inter-Personal Distance Could Not Be Enough. *International Journal of Environmental Research and Public Health* **17**, 2932 (2020).
63. Sun, C. & Zhai, Z. The efficacy of social distance and ventilation effectiveness in preventing COVID-19 transmission. *Sustainable Cities and Society* **62**, 102390 (2020). <https://www.sciencedirect.com/science/article/pii/S2210670720306119>.
64. Verity, R. et al. Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet. Infectious Diseases* **20**, 669–677 (2020).

65. Sargent, R. G. Verification and validation of simulation models. *Journal of Simulation* **7**, 12–24. <https://doi.org/10.1057/jos.2012.20> (2013).
66. Balci, O. Validation, verification, and testing techniques throughout the life cycle of a simulation study. *Annals of Operations Research* **53**, 121–173. <https://doi.org/10.1007/BF02136828> (1994).
67. Grimm, V. et al. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science (New York, N.Y.)* **310**, 987–991 (2005).
68. Windrum, P., Fagiolo, G. & Moneta, A. Empirical Validation of Agent-Based Models: Alternatives and Prospects. *Journal of Artificial Societies and Social Simulation* **10**, 1–8 (2007). <https://ideas.repec.org/a/jas/jasss/2006-40-2.html>. Publisher: Journal of Artificial Societies and Social Simulation.
69. Tracy, M., Cerdá, M. & Keyes, K. M. Agent-Based Modeling in Public Health: Current Applications and Future Directions. *Annual Review of Public Health* **39**, 77–94 (2018).
70. Polhill, J. G. et al. Using Agent-Based Models for Prediction in Complex and Wicked Systems. *Journal of Artificial Societies and Social Simulation* **24**, 2 (2021).
71. Liu, Q. & Cao, L. Modeling time evolving COVID-19 uncertainties with density dependent asymptomatic infections and social reinforcement. *Scientific Reports* **12**, 5891 (2022). <https://www.nature.com/articles/s41598-022-09879-2>.
72. Swallow, B. et al. Challenges in estimation, uncertainty quantification and elicitation for pandemic modelling. *Epidemics* **38**, 100547 (2022). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7612598/>.
73. Teerawattananon, Y. et al. Recalibrating the notion of modelling for policymaking during pandemics. *Epidemics* **38**, 100552 (2022). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8889889/>.
74. Epstein, J. M. Agent-Based Computational Models and Generative Social Science. *Complexity* **4**, 41–60 (1999).

## Acknowledgements

This research was supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 101003606. Mikhail Sirenko is extremely grateful to Daan van Bilsen, Sahiti Sarva, Jin Rui Yap, Fabio Tejedor, Anmol Soni, Hidde Bijlard, and Srijith Balakrishnan for their enormous contribution to the foundation of this work.

## Author contributions

M.S. initiated and designed the study, prepared the input data, developed the agent-based model, performed the simulations and analysed experiments. T.C. and M.S. developed the conceptual framework, and T.C. provided guidance on resilience and adaptation theories. A.V. and T.C. supervised the research. A.V. developed the agent-based model, provided expertise in simulation modelling, and contributed to the interpretation of results. M.S. wrote the original draft of the manuscript, and all authors contributed to the reviewing and editing of the final manuscript. All authors have approved the final version of the manuscript.

## Declarations

### Competing interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-33158-5>.

**Correspondence** and requests for materials should be addressed to M.S.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026