Simulating Mitigation Strategies for COVID-19 Hot Spots

Location Specific Activity Restrictions in EpiSim

Jakob Rehmann

Submitted as partial fulfilment of the degree requirements for the
M.Sc. Planung und Betrieb im Verkehrswesen

on
December 13$^{\text{th}}$, 2021

under the supervision of
Dr. Kai Nagel
M.Sc. Sebastian Müller

at the
Technische Universität Berlin
Institut für Land- und Seeverkehr
Verkehrssystemplanung und Verkehrstelematik

Declaration of Independent Work

I hereby declare that I have prepared this thesis independently and by my own hand, without any unauthorized outside help, and exclusively using the sources and auxiliary materials listed.

Selbstständigkeitserklärung

Hiermit erkläre ich, dass ich die vorliegende Arbeit selbständig und eigenhändig sowie ohne unerlaubte fremde Hilfe und ausschließlich unter Verwendung der aufgefürhten Quellen und Hilfsmittle angefertigt habe.

Jakob Rehmann, December 13, 2021
E-Mail: j.rehmann@tu-berlin.de
Matrikelnummer: 400998

Abstract

The backdrop for this thesis is the ongoing coronavirus pandemic. EpiSim is an agent-based epidemiological simulation software, which was developed to simulate the virus' spread and the impact of mitigation measures (e.g. mask-mandates, lockdowns, and vaccinations).

The functional contribution of this thesis is the localization of EpiSim; the software was extended to allow different regions within the study-area to have diverging parameters that influence virus spread. This thesis introduces a location-based remaining fraction, which allows pandemic-induced activity reductions (with respect to pre-coronavirus mobility patterns) to vary between sub-regions of the study area. Residents of one neighborhood could have a 20% reduction in leisure activities on a given day, while an adjacent neighborhood could have a 30% reduction.

The functionality was tested on the 12 boroughs that make up Berlin, Germany. The study's time-frame was the second wave of infections that began in Fall 2020. The first research focus was to see whether the localized EpiSim could better capture the actual infection dynamics of Berlin's boroughs; instead of using mobility reductions aggregated on a city level, the reductions were differentiated by borough. Additionally, regionally diverging average home-sizes were added to EpiSim and smaller homes were assigned higher chances of infection. Neither of those modifications showed significant improvements in capturing local infection dynamics. This indicates that mobility reduction and home-size aren't the best factors for explaining diverging dynamics between Berlin's boroughs.

The second research focus was to evaluate the impact of pinpointed lockdowns during the second wave of infections in Berlin. EpiSim's `AdaptivePolicy` was localized such that individual sub-regions can be locked down (and opened up) independently of other sub-regions. Restrictions are automatically applied if the incidence of a sub-region surpasses a threshold. The adaptive policy with the most lenient parameters resulted in 62,000 fewer infections without the lockdowns being more restrictive than what occurred in reality. Compared to a global adaptive policy, the local alternative was shown to be particularly effective when the incidence threshold was low.

Zusammenfassung

Der Hintergrund für diese Arbeit ist die aktuelle Coronavirus-Pandemie. EpiSim ist eine agentenbasierte epidemiologische Simulationssoftware, die entwickelt wurde, um die Ausbreitung des Virus und die Auswirkungen von Schutzmaßnahmen (z. B. Maskenpflicht, Lockdowns und Impfungen) zu simulieren.

Der funktionale Beitrag dieser Arbeit ist die Lokalisierung von EpiSim. Die Software wurde so erweitert, dass verschiedene Regionen innerhalb des Untersuchungsgebiets abweichende Parameter aufweisen können, die die Virusausbreitung beeinflussen. In dieser Arbeit wird eine ortsbezogene "remainingFraction" eingeführt, die es ermöglicht, die pandemiebedingte Aktivitätsverringerung (im Vergleich zu den Mobilitätsmustern vor dem Coronavirus) zwischen den Unterregionen des Studiengebiets zu variieren. Die Bewohner*innen eines Stadtteils könnten an einem bestimmten Tag eine Verringerung der Freizeitaktivitäten um 20% erfahren, während ein benachbarter Stadtteil eine Verringerung um 30% aufweisen könnte.

Die Funktionalität wurde anhand der 12 Bezirke von Berlin getestet. Der Zeitrahmen der Studie war die zweite Infektionswelle, die im Herbst 2020 begann. Der erste Forschungsschwerpunkt bestand darin herauszufinden, ob das lokalisierte EpiSim die tatsächliche Infektionsdynamik in den Berliner Bezirken besser abbilden kann. Anstelle der auf Stadtebene aggregierten Mobilitätsreduktionen wurden die Reduktionen nach Bezirken differenziert. Außerdem wurde EpiSim mit regional abweichenden durchschnittlichen Wohnungsgrößen ergänzt, wobei kleineren Wohnungen höhere Infektionswahrscheinlichkeiten zugeordnet wurden. Keine dieser Änderungen zeigte signifikante Verbesserungen bei der Erfassung der lokalen Infektionsdynamik. Dies deutet darauf hin, dass die Verringerung der Mobilität und die Wohnungsgröße nicht die besten Faktoren sind, um die unterschiedliche Dynamik in den Berliner Bezirken zu erklären.

Der zweite Forschungsschwerpunkt war die Bewertung der Auswirkungen von gezielten Lockdowns während der zweiten Infektionswelle in Berlin. Die `AdaptivePolicy` von EpiSim wurde so lokalisiert, dass einzelne Bezirke unabhängig von anderen Bezirken gesperrt (und geöffnet) werden können. Lockdowns werden automatisch angewendet, wenn die Inzidenz eines Bezirkes einen Schwellenwert überschreitet. Die `AdaptivePolicy` mit den mildesten Parametern führte zu 62.000 weniger Infektionen, ohne dass die Lockdowns restriktiver waren als in der Realität. Im Vergleich zu einer globalen `AdaptivePolicy` erwies sich die lokale Alternative als besonders wirksam, wenn die Inzidenzschwelle niedrig war.

Acknowledgements

To the VSP team: thank you for nurturing me over the last two and a half years into a budding transportation planner and software developer. When I first started, Gregor Leich had to teach me to write unit tests or run a debugger; now, I can run simulations on a super computer (how sci-fi is that?!). In particular, I want to acknowledge the EpiSim team for guiding me through this thesis. I am so grateful to Sebastian Müller and Kai Nagel for going above-and-beyond as supervisors.

Thank you also to my family—Ditte • Janni • Tucci • Succi—for accompanying me through all ups and downs of my life so far. In high school you dealt with the full brunt of my procrastination; pulling all-nighters to help me finish my Dostoevsky papers. You have taught me to find meaning and joy in my work; I've also learned the importance of boundaries between work and life from you.

Last but not least, I want to thank my roommates and friends for giving me a stable and love-filled home in Berlin; you've always lifted my mood (and literally fed me) when I was "brain-fried" after hours of debugging. A very special shout-out to Dominik: you've accompanied my intellectual, political, and moral progression since day 1 in Berlin, and I am grateful to have you in my corner for years to come. Also, a deep thanks to Brendan: not only for letting me re-purpose your bedroom as a home-office, but also for being my most reliable friend and partner since we met in Harlem eight years ago.

Contents

List of Tables

List of Figures

List of Acronyms

ABM agent-based model.

COVID-19 coronavirus disease 2019.

LOR lebensweltlich orientierter Raum.

NPI non-pharmaceutical interventions.

RKI Robert Koch Institut.

SARS-CoV-2 severe acute respiratory syndrome coronavirus 2.

TTI test-trace-isolate.

TUB Technische Universität Berlin.

VOC variant of concern.

VSP Fachgebiet Verkehrssystemplanung und Verkehrstelematik.

WHO World Health Organization.

Glossary

EpiSim an agent-based software used to simulate epidemiological dynamics; it was designed to evaluate mitigation policies for COVID-19.

incidence weekly infections per 100,000 residents.

infectivity the degree to which an infectious person transmits the virus to susceptible agents.

MATSim an agent-based simulation framework used to simulate transportation-related policies.

susceptibility the likelihood that a susceptible person becomes becomes infected when exposed to the virus.

Chapter 1

Introduction

This Introduction gives a brief background of the coronavirus disease 2019 (COVID-19) and the non-pharmaceutical interventions (NPI) that governments around the world implemented. Following, the epidemiological simulation software used in and expanded for this thesis will be presented: EpiSim. Then, the Introduction will lay out the two research focuses and motivations for this thesis. The following two sections will present the research questions and case studies corresponding to those motivations. The chapter will be rounded off with an outline for the further chapters.

1.1 COVID-19 and NPIs

> “Mr. Li is one of 59 people in the central city of Wuhan who have been sickened by a pneumonia-like illness, the cause of which is unclear [1].”

This quote from the *New York Times* was published on January 6, 2020; the unknown “illness” and “cause” were identified and named in a matter of weeks—coronavirus disease 2019 (COVID-19) and severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), respectively [2]. Containing the virus, however, proved to be more difficult; by March 11, 2020, there were 118,000 reported cases of COVID-19 in 114 countries and more than 4,000 people had died of the disease [3]. The director-general of the World Health Organization (WHO) declared COVID-19 a pandemic and made the following plea:

> “We cannot say this loudly enough, or clearly enough, or often enough: all countries can still change the course of this pandemic. If countries detect, test, treat, isolate, trace, and mobilize their people in the response, those with a handful of cases can prevent those cases becoming clusters, and those clusters becoming community transmission [3].”

At the time, pharmaceutical interventions—vaccines to prevent COVID-19 transmission and treatments to mitigate the disease's effects—were not yet in development. In fact, they are still unavailable for most of the world, as demonstrated in [4]. Thus, the primary defence against COVID-19 was and potentially still is the implementation of non-pharmaceutical interventions (NPI), which Perra [5] describes as: "a wide range of both top-down (i.e., governmental) and bottom-up (i.e., self-initiated) measures aimed at interrupting infection chains by altering key aspects of our behavior." In January 2021, Chinese authorities locked down the cities of Wuhan and Huanggang and restricted travel to and from a total of 10 cities [6]. As seen in [7], most governments globally had implemented some combination of NPIs by April 1, 2020. Implemented measures include mask mandates, school closing, travel bans, stay-at-home restrictions, and bans of public events. Google mobility data presented in [7] also shows that people spent a lot more time at home and forwent recreation, education, work and shopping activities. This reduction was attributable to to government interventions, as well as the decisions of individuals seeking to protect themselves and their community.

1.2 EpiSim

A few weeks after COVID-19 spread in Germany, a research group at the Technische Universität Berlin (TUB) in Germany began developing an agent-based simulation software (EpiSim) to track the development of COVID-19 cases and prognosticate the impact of various government-imposed mitigation strategies [8]. For every simulated day, agents complete a set of activities. If a susceptible and infectious agent meet at an activity location or in public transit, there is a chance that the susceptible agent is infected.

To determine what activities an agent completes on a given day, we start with the typical daily (pre-coronavirus) activity trajectories of all individuals living in the study area. Due to government restrictions or personal decision-making, agents will only complete a fraction of their pre-coronavirus activities: "Remaining Fraction" (Rf). The Rf varies per activity type, and changes throughout the pandemic. The core-functionality of EpiSim, as well as it's applications, will be covered in Chapter 2.

1.3 Contributions

The fact that EpiSim is an agent-based model (ABM) allows infection dynamics to vary between sub-populations; e.g. older agents are given more severe disease progressions (e.g. higher likelihoods of needing hospitalization) and higher prioritization in vaccine distribution. The goal of this master's thesis is to further enrich EpiSim by allowing infection dynamics to depend on the home or activity locations of agents.

The cornerstone of added functionality was the localization of the remaining fraction.

EpiSim was extended to allow one sub-region of the study area to have a different remaining fraction than another: "locationBasedRf." On this basis, a local adaptive policy was designed to dynamically restrict or open up an area based on the incidence[1] of that area. The added functionality will further be described in Chapter 3.

The second contribution of this thesis is a set of four case studies, which make use of the added functionality. There were two broad research focuses explored by the case studies:

A) Improve EpiSim's ability to capture the infection dynamics of individual geographic regions within the study area.

B) Evaluate the merits of imposing localized lockdowns, wherein some regions are restricted while others remain open.

The first research focus is geared towards improving EpiSim's ability in capturing the virus progression, while the second is geared towards expanding EpiSim's toolbox of policy recommendations to evaluate. The following two sections present the two research focuses separately. In each section, the motivation for the research focus is detailed and the corresponding research questions are developed. Each research question is accompanied by a case study; the corresponding results will also be summarized.

1.4 Research Focus A: Local Infection Dynamics

Research focus A explores whether the addition of regional characteristics into EpiSim will improve EpiSim's ability to produce more accurate incidence curves for sub-regions within the study-area.

1.4.1 Motivation

The first motivation of localizing EpiSim is to improve the model's ability to capture infection dynamics of sub-regions. While EpiSim simulations are calibrated to mirror the actual infections and hospitalizations in a general study region, they do not necessarily produce realistic infection curves for individual neighborhoods.

If the addition of neighborhood characteristics improves EpiSim's ability to reproduce the infection curves of individual neighborhoods, this can help researchers understand how the virus spreads. These realizations could aid policy makers in improving and focusing disease mitigation efforts. For instance, Endt et al. [9] show that Berlin boroughs with higher unemployment rates also face higher infection curves. This information could

[1] Incidence indicates the number of weekly infections per 100,000 residents. This term will be frequently used throughout this thesis.

be used by the government to tailor its messaging or vaccination campaign to jobless individuals.

Using the functionality developed as part of this thesis, two location-based factors were added to the EpiSim simulation: activity reduction and home-size. The following two subsections introduce two research questions (RQ), which explore the two added factors.

1.4.2 A1: Localized Activity Reduction

In EpiSim, the daily mobility of agents must be scaled down with respect to the typical pre-coronavirus daily activities. This is done by incorporating daily activity reductions, gleaned from cell-phone mobility data, into EpiSim. In the standard EpiSim model, a single value for activity reduction is applied to the entire study area (per day and activity type). RQ A1 explores the application of different values of activity reduction for different parts of the study area.

RQ A1: *Does the addition of localized activity reductions improve EpiSim's ability to capture local infection dynamics?*

Case Study A1 explores RQ A1 by applying different activity reductions to the 12 neighborhoods of Berlin. On a certain day, the borough of Spandau may reduce leisure activities by 30%, while Mitte only reduces it by 20%. The borough-based activity reductions were also gleaned from historical cell phone mobility data.

The results of Case Study A1 show that the inclusion of borough-based remaining fractions did change the incidence curves of individual boroughs. However, these changes did not systematically improve EpiSim's ability to capture the boroughs' infection dynamics.

1.4.3 A2: Localized Contact Intensity

As will be described in the Chapter 2, if an infectious person meets with a susceptible person, the `InfectionEventHandler` calculates the chance of infection based on various characteristics of the individuals involved. Another important variable in this calculation is "contact intensity"; this describes the infection characteristics of the room in which the agents meet, including the room-size and ventilation. In the standard EpiSim model, the contact intensity for home events is equal for all agents. The next research question attempts to localize EpiSim by applying varying home-sizes to different sub-regions, and varying the contact intensity accordingly.

RQ A2: *Does the inclusion of localized contact intensity for home activities (based on varying home-size) improve EpiSim's ability to capture local infection dynamics?*

This research question was explored in context of Berlin with Case Study A2. The average home size per person was available per lebensweltlich orientierter Raum (LOR), a Berlin-specific planning unit of which there are 448 in the city. Case Study A2 assigns the average home-size per person to all agents living in an LOR. New contact intensities were then assigned to home events; agents with smaller home-sizes had higher contact intensities (leading to a higher chance of infection).

The incidence curves of individual boroughs deviated insignificantly between the base case and the policy case, where localized contact intensities were implemented. Thus, the addition contact intensities based on home-size did not improve the simulation's ability to capture the infection dynamics of the boroughs.

1.5 Research Focus B: Local Lockdowns

The second research focus was the evaluation of localized lockdowns. This is meant to expand the types of NPIs that EpiSim can model, in order to facilitate more diverse policy recommendations.

1.5.1 Motivation

When I joined the EpiSim team in early 2021, there was a growing unhappiness with ongoing lockdowns of uncertain length. In this context, alternative restriction regimes were discussed, such as the no-COVID proposal. Bauman et al. [10] explain that the goal of no-COVID is to eradicate COVID-19 through targeted lockdowns along with a robust test-trace-isolate (TTI) program. The general idea of the no-COVID strategy was to 1) use a hard lockdown to reduce the incidence; 2) remove restrictions step-wise and use a vigorous TTI program to find all contacts of infected people and place them into quarantine; 3) if the infections of an area surpass the capacity the TTI program, put that region into lockdown, while allowing free movement in other areas [10]. People can move freely within and between the unrestricted areas ("green zones"). Movement in to or out of lockdown areas ("red zones"), however, is strongly curtailed. These green zones are meant to expand progressively until they can encompass multiple countries [10].

Most countries that touted policies similar to no-COVID have since abandoned them; in October 2021, New Zealand "gave up" on the goal of eradicating the virus within its borders, partially due to the increased infectiousness of the Delta virus [11]. As this is being written, China is the last country retaining this type of policy [12]. Although the eradication of the virus may be untenable, pinpointed lockdowns may remain a promising tool in keeping the virus spread in check.

Inspired by the red and green zones of no-COVID [10], research focus B attempts to evaluate the merit of implementing local restrictions in an EpiSim simulation. The central goal was to explore whether pinpointed restrictions can mitigate the extent of the pandemic and/or reduce the amount of time people spend in lockdown, in comparison to global restrictions.

1.5.2 B1: Pinpointed Lockdown

Before examining the effects of a local adaptive policy, we will evaluate the effects of a single local lockdown. This is meant to demonstrate that reducing the location-based remaining fraction of a sub-region has an impact on local infection dynamics. The first research question is rather general; it's goal is to check whether the pinpointed restrictions work as intended.

RQ B1: *How does a local lockdown affect the infection dynamics of the restricted region and the un-restricted regions?*

Case Study B1 restricts the Berlin borough of Mitte for one month: October 2020. The first purpose is to demonstrate that a pinpointed restriction has an effect on the incidence curves of Mitte, as well as the other boroughs. Additionally, the goal of Case Study B1 is to compare the effects of two restriction regimes: a) locking down residents of Mitte from conducting leisure activities and b) preventing all Berliners from conducting leisure activities within Mitte.

The results from Case Study B1 show that a local lockdown in Mitte (for both restriction regimes) significantly reduces infections in Mitte itself; however, other boroughs are also have reduced infections. As to be expected, the non-restricted boroughs benefit more (in terms of reduced infections) from regime b, which bars Berliners from entering Mitte to complete leisure activities.

1.5.3 B2: Local Adaptive Restrictions

The final research question attempts to explore the effects of a local adaptive policy, which was a functionality introduced by this thesis. If the incidence in a borough surpasses a threshold, a lockdown is imposed in that borough; residents of other boroughs are free to continue completing activities (corresponding to regime a from the previous subsection). This policy only partially matches the no-COVID proposal; the local adaptive policy still allows residents of unrestricted regions to enter restricted regions.

RQ B2: *How does a local adaptive policy affect incidences and time uses, compared to a global adaptive policy? How do the parameters of the adaptive policy affect its benefit?*

RQ B2 evaluates one benefit and one cost of the local adaptive policy. The benefit is a reduction in incidence. The cost is a reduction in time-use: the amount of time people spend outside of their homes each day.

Case Study B2 explores the application of a local adaptive policy in Berlin during the second wave of the coronavirus. The case study varies two parameters regarding the automatic lockdowns: the threshold at which a lockdown is imposed or lifted ("Trigger") and the percentage of activities can still take place during a lockdown ("Remaining fraction for restricted policy"). The local adaptive policy was first compared to the base case: the standard EpiSim model which is calibrated for Berlin to match the actual historical incidences. The local adaptive policy is also compared to the global adaptive policy—wherein all of Berlin is locked down if the Berlin-wide incidence surpasses a threshold—to examine at what parameters the local policy is beneficial.

The results for Case Study B2 show that both the local and global adaptive policies reduce more infections with the more stringent parameters: lower Trigger and lower Rf. This benefit is offset by the social cost of the more stringent policies: people spend less time outside of their homes every day, on average. The most lenient adaptive restrictions are most efficient in terms of this cost-benefit relationship; they significantly reduce infections, while barely reducing the amount of time people can spend outside of their homes. Finally, the results of Case Study B2 show that the local adaptive policy performs best in Berlin, as compared to the global one, when the Trigger is low.

1.6 Outline

This chapter has laid out the motivations for this thesis. It has also introduced the research questions and corresponding case studies, which will be explored in this thesis. Chapter 2 will introduce epidemiological simulations in general. It will then give an in-depth description of how EpiSim functions, and show how it can be utilized to explore the research questions presented above.

The Methodology (see Chapter 3) aims to highlight the software contributions of this thesis. It will also describe how the four case studies were prepared. The results of the case studies will be described in Chapters 4 and 5. Chapter 4 presents the case studies for research focus A, while Chapter 5 presents the results for research focus B. Both chapters end with a discussion. The paper is rounded off with a Conclusion (see Chapter 6), which summarizes the work in this thesis and gives a general outlook.

Chapter 2

Modelling Background

The first section of this chapter gives a brief history of epidemiology and different model types utilized by it. The following sections describe EpiSim, the epidemiological simulation software used in and expanded upon as part of this thesis.

2.1 Epidemiological Models

Whereas most medical sciences study disease in individuals, epidemiology studies how disease propagates within groups of humans; it examines what factors lead to the spread of infections and what mitigating measures can be implemented [13]. Mulner [13] describes that many regard the birth of modern epidemiology to stem from London, during the 1854 cholera outbreak. John Snow, a physician, found that cholera victims' homes were clustered around a public drinking fountain; thus, he hypothesized that the infections were caused by contaminated drinking water [13].

John Snow used rudimentary spatial analysis in his epidemiological research, meaning he drew dots on a map to identify clusters [13]. Since then, the capabilities of models have significantly increased. Luyao et al. [14] discuss three general types of disease spreading models: A) statistical models that use regression or machine learning to predict spread—these models can generally only make short-term predictions and aren't very sensitive to interventions; B) equation based models, which aggregate the population and use complex differential equations to model spread between them; C) an agent-based model (ABM), which can depict a more heterogeneous population with complex social interactions. As EpiSim is agent-based, ABMs will be described in greater detail.

An agent-based model (ABM), as described by [15], runs simulations on a heterogeneous population. They can simulate how synthetic individuals ("agents"), who have different demographics, interact with each other and the environment. Gilbert and Troitzch [16] define agents as having autonomy, social ability, reactivity, and proactivity.

Thus, ABMs can portray more realistic movement and contact patterns than equation based models, which generally assume uniform mixing [15]. This allows epidemiological ABMs to zoom in on individual infection chains or examine the infection dynamics for particular sub-groups.

ABMs are frequently used in the context of epidemiology. For instance, Merler et al. [17] developed a spatial ABM to predict the spread of Ebola virus in Liberia. It allowed them to differentiate between transmission rates at/in funerals, hospitals and households. The researchers [17] could then estimate the effect of NPIs in combating the spread of Ebola in those locations. In a study by Zhou et al. [18], a combination of an ABM and a susceptible-exposed-infected-recovered model is used to show that COVID-19 transmission patterns vary throughout the city of Guangzhou, China. The authors [18] find that immunity rate required to end the pandemic is spatially heterogeneous (i.e. higher rate required in urban center). Furthermore, if the supply of vaccine is limited, the most efficient distribution involves spatial (as well as age-based) prioritization [18].

2.2 EpiSim Overview

EpiSim [8] was developed at Fachgebiet Verkehrssystemplanung und Verkehrstelematik (VSP)—english: the Department of Transport Systems Planning and Transport Telematics—at the Technische Universität Berlin (TUB), which is led by Dr. Kai Nagel. Before COVID-19 arrived in Germany, the team primarily worked on software development of an agent-based transport model, MATSim [19]. This software is used to simulate the mobility of synthetic individuals in a study region. MATSim is primarily used to simulate the effects of different mobility developments; for instance, changes to infrastructure (i.e. bike highways [20]), changes to policy (i.e. congestion pricing [21]) or the introduction of new forms of mobility (i.e. shared taxis or "demand responsive transit" [22]).

The output of a MATSim simulation is an events file for a typical pre-coronavirus day; this contains the activities each individual completes (including where, what, when, and with whom), as well as how they traveled between those locations [19]. This was an invaluable starting point for developing an epidemiological simulation because the places where people meet are also the places where people can become infected. Once COVID-19 became relevant in Germany, VSP was able to put together an epidemiological simulation software (written in Java) based on these activity chains and real-time data within two weeks [8].

2.3 Simulation Setup

At the beginning of an EpiSim simulation, various input files have to be imported and the settings must be configured. Doing this manually for each simulation would be work-intensive and prone to human-error. Thus, each study area has a separate class, a so-called "production scenario," which sets all the regionally applicable input files and default configuration values. For instance, the school vacation periods would be differ between Berlin and Cologne (the class `SnzBerlinProductionScenario` is used for Berlin). The following subsection describes the input files that the production scenario reads before the simulation begins. Afterwards, the data structure containing the configuration options—episim-config—will be introduced. Finally, the way in which activity reductions are configured will be detailed.

2.3.1 Inputs

Below, I introduce the most important input files for an EpiSim simulation.

Population and Facilities Inputs: EpiSim has agents that move about the city, and meet in facilities to complete activities (where there exists the chance of infection). As EpiSim is an ABM, each individual and each facility can have different characteristics which influence infection dynamics. Within the production scenario, a population file is imported: this contains the demographic information for every individual simulated in EpiSim. In the standard Berlin model, this includes age, sex, and details to their home: longitude and latitude, facility ID, and county. The ability to integrate this demographic information into the simulation makes this ABM powerful; e.g., for age-prioritized vaccine distribution strategies, the age of an individual is read from the population file to determine eligibility. Similarly, the facilities file enriches the model by providing information about each activity location: the longitude and latitude, as well as information about what activities can occur in that location.[1]

Events Files: Vital to EpiSim is the mobility behavior of individuals; three event files must be imported that show a typical weekday, Saturday, and Sunday. These are generated by filtering the MATSim output events to only contain events relevant to the epidemiological simulation. If an agent travels per public transit, we know when they got on and off the vehicle and the ID of that vehicle. When the agent conducts an activity, we know the activity type (e.g. home, leisure, work), the activity location, and the time-frame. During the simulation, the contact model can use this information to detail all the contacts between individuals, whether in a bus or university class. Within

[1] The standard facilities and population files do not specify what Berlin borough the locations are in. Home-size is also not included by default. This information is required for the case-studies; the addition of these parameters will be described in Chapter 3.

the events files, IDs of the involved agent, activity facility, or public transit vehicle are specified. These IDs can be matched to the population and facilities files, to extract pertinent information.

Activity Reduction: While the events files provide the typical daily plans of individuals, an activity reduction input is required to scale back mobility during the coronavirus pandemic. Every day, VSP is provided with mobility data from cell-phone providers. The information includes the number and average duration of activities (for various activity types) that residents of a given zip-code complete on a given day. The class `AnalyzeSnzData` aggregates this data for a given study area, and calculates the daily activity reductions, as compared to a pre-coronavirus baseline. As will be shown in Section 2.3.3, these activity reductions will be incorporated into the `Restrictions` for certain activity types.

2.3.2 EpiSim Configuration

EpiSim is designed to be extensively configurable and extendable. The class `EpiSim-ConfigGroup` specifies most of the aspects of EpiSim that can be configured by the user. Based on this class, a container called the episim-config is created, which holds the parameters and options selected for a given run. It includes parameters relevant to simulation mechanics, such as the simulation start-date and what outputs should be produced. It also includes parameters relevant to infection dynamics, including SARS-CoV-2 susceptibility and infectivity. Infectivity describes the degree to which the infectious host can cause new infections while susceptibility describes the likelihood for the susceptible agent to contract the disease [23]. These parameters are configured in the episim-config container for different age groups. The episim-config also specifies how many individuals are infected with SARS-CoV-2 at the beginning of the simulation.

Nested within the episim-config are two very important configuration groups. The progression-config specifies the transition probabilities between different disease states: i.e. what percent of young adults will become “seriously sick” after “showing symptoms” for 4 days. It also includes the policy-config, which specifies how restrictions will be applied to the population throughout the simulation (see Section 2.4.3). These restrictions can be set at the beginning of the simulation—“Fixed Policy”—or imposed during the simulation based on simulated incidences—“Adaptive Policy” (see Section 2.3.3).

2.3.3 Policy

As described earlier, the episim-config contains a policy-config, which specifies how activity restrictions will be applied to the population throughout the simulation. To begin, the `Restriction` container will be described. This is a basic data structure that contains activity reduction information. Afterwards, two types of restriction policies will be described: `FixedPolicy` and `AdaptivePolicy`.

Restriction: The activity reduction information is stored in a data structure called `Restriction`. One `Restriction` container is applicable to one activity type; the `Restriction` generally applies for one week.[2] Within the `Restriction`, the `remainingFraction` field dictates what percent of activities will take place that day. On a given day, the `remainingFraction` for work activities could be 66%, meaning that one third of work events will be skipped. The `Restriction` container also contains information on mask usage—what percent of people wear what type of mask—maximum group sizes, whether activities locations close early or are closed altogether. The policy-config specifies how the `Restriction` container is filled.

Fixed Policy: The fixed policy is the standard policy used in EpiSim simulations. Therein, all `Restriction` containers are created and filled at the beginning of the simulation. For a work or leisure `Restriction`, the `remainingFraction` is filled using the activity reduction data from the cell-phone providers. The production scenario delegates the task of reading the activity reduction input files to `CreateRestrictionsFromCSV`, which then populates the `remainingFraction` field of each `Restriction`. This class can also extrapolate the activity reductions into the future, where no cell phone data is available. For educational activities, the `remainingFraction` is generally set manually within the production scenario. If schools are closed due to holidays or government lockdowns, the `remainingFraction` within the `Restriction` containers for educational activities is set to 0.0.

Adaptive Policy: As shown above, the fixed policy defines activity restrictions at the beginning of the simulation. In contrast, the adaptive policy can change restrictions during the simulation based on the current incidence. There are three predefined restriction states: initial, open, and restricted. Each state sets one `Restriction` per activity type. A `Restriction` for the restricted phase should have a lower `remainingFraction` than the `Restriction` for the open phase.

The initial policy is applied at the beginning of the simulation. If the incidence surpasses a certain threshold ("restricted-trigger"), the restricted policy is implemented. Then, if the incidence stays below a certain level ("open-trigger") for a set number of days (default is two weeks), the open policy is applied. The triggers can differ between activity types, which means that the `Restriction` for the initial phase could be in effect for educational activities, while the `Restriction` for the restricted phase could be simultaneously active for leisure activities. Additionally, the open-trigger and restricted-

[2]This time-frame could also be shortened, e.g. to apply different reductions for holidays.

trigger for a given activity type can differ from one another.

2.4 EpiSim Simulation

So far, the scenario preparation process has been described; the episim-config data structure, the various inputs, and the restriction policies have been detailed. Now, it is time for the simulation to begin. At the beginning of the simulation, a random set of individuals is infected. The class `EpisimRunner` iterates through all the days (iterations) in the simulation. One simulated week involves replaying the weekday events five times, followed by the Saturday events and, finally, the Sunday events.

At the beginning of each day, there is a certain initialization phase: e.g. vaccinations and tests are administered to portions of the population. This is also where the disease progression model changes the disease states of individuals (see Section 2.4.3). Based on transition probabilities, some infectious individuals will e.g. start showing symptoms or some seriously sick individuals will recover.

At the beginning of a day, the `DefaultParticipationModel` ascertains which activities should be skipped. Since the events file corresponds to pre-coronavirus activity trajectories, a certain portion of planned daily activities should be removed due to pandemic-induced mobility reductions. For each event in the day's events file, the applicable `remainingFraction` is extracted from the `Restriction` for the activity type and date in question. A random number between 0.0 and 1.0 is then generated; if it lies above the `remainingFraction`, then the activity will not occur on this date. Thus, for a `remainingFraction` for work of 66%, e.g., one third of all work activities on that day will not occur, on average. If an activity is skipped, then the associated public transit travel to that activity will also not be realized.[3]

Each day is simulated as follows: the corresponding events file is replayed minute-by-minute, by processing the events chronologically. Episim [8] predicts daily infection dynamics using three interlocking sub-models: the contact model identifies situations where agents meet; given an agent is infectious, the infection model predicts whether SARS-CoV-2 is transmitted to a susceptible agent; if an agent is infected, the disease progression model determines the course and severity of the disease. These models will be described in the following subsections.

[3]This is the process when the episim-config option of `activityParticipation` is set to "startOfDay." If, instead, it is set to "duringContact," the activity participation for a given day is ascertained in the contact model. This is computationally less efficient than the process described here.

2.4.1 Contact Model

The `DefaultContactModel` starts the day with a set of empty containers, corresponding to all the enclosed spaces where agents can meet (activity facilities and public transit vehicles). The `ReplayHandler` class then replays all the events that occur in a day in chronological order.

If an activity occurs (as specified by the `DefaultParticipationModel`), the `DefaultContactModel` will then add the agent in question to the corresponding container. When the events file specifies that an agent ends an activity, the `DefaultContactModel` checks whether at least one person in the container is susceptible and one is infectious. If so, this agent-pair is passed along to the `InfectionEventHandler`, which will check whether an infection actually occurs. This same process also occurs when an agents enters or leaves a public transit vehicle [8].

2.4.2 Infection Model

If the contact model renders that a contagious person met with a susceptible person, the infection model calculates the probability that an infection occurs. A mechanical infection model developed by Smieszck [24] [25] is used in EpiSim. The equation used in EpiSim to calculate the probability of infection along with a detailed explanation can be found in [8]. The following is a simplified version, which is useful for explanatory purposes:

$$p(\text{infect}|\text{contact}) \approx \Theta \times sh \times ci \times in \times \tau$$

As detailed in [8], the "shedding rate", sh, is the viral load that the infected person produces through exhalation; this is dependent on what stage of the disease the person is in and whether they are wearing a mask. Conversely, the "intake rate", in, is how much viral load a susceptible agent breathes in, which is also dependent on mask-usage. Age-dependent susceptibility and infectivity, which were defined in the episim-config, also play a role in the shedding and intake rates. The "contact intensity", ci, describes the concentration of the virus in the air, and is dependent on the size of the room and how much air exchange takes place (i.e. through the opening of windows). Finally, the duration of contact (τ) is an important factor, which can be gleaned from the activity trajectories. The calibration factor, Θ, is determined through model calibration; this also absorbs all units of the individual factors. All of this together predicts whether susceptible individuals become infected with SARS-CoV-2 during an interaction with an infectious person. This explanation was paraphrased from [8].

2.4.3 Disease Progression Model

After an agent is infected by SARS-CoV-2, the disease progression model is used to determine if and when certain disease stages occur. When the infection model reports that a `susceptible` person is infected, the agent will initially be assigned the state of `infectedButNotContagious`.

At the beginning of each iteration (simulated day), there will be a certain chance for the infection state of the agent to change. The worst possible progression would be as follows:[4] `infectedButNotContagious` → `infectious` → `showingSymptoms` → `seriouslySick` → `critical` → `seriouslySick` → `recovered`. However, starting with `infectious` state, there is always a chance of becoming `recovered`. The progression model uses age-dependent transition probabilities to determine whether the agent moves on to a more dire phase or to the `recovered` phase. The progression model also sets the duration of each phase based on medians and standard deviations found in the literature. The agent can only infect other agents in the infectious state or within the first 4 days of the showing symptoms state; in later states, it is assumed that the agent is either isolating at home or in the hospital. Readers wanting more information on the sources of the transition probabilities are encouraged to read Section 2.3 of [8].

2.5 EpiSim Outputs and Analysis

All infections and disease state changes are documented by `EpisimReporting`. This class is used to produce the output files useful for analysis. One of the most important outputs created is infections.txt, which documents the infection state counts for each date and each county: e.g. number of individuals who are susceptible, infected but not contagious, contagious, seriously sick, recovered, quarantined, tested, vaccinated, etc. Other important outputs of `EpisimReporting` include the prevalence of different virus strains and timeUse.txt: how many minutes individuals spend outside of their homes per day.

Both infection dynamics and political discourse are in constant flux; for EpiSim to make an impact on decision-makers, the turnover must be rapid. Covid-Sim is a visualization tool that was developed to synthesize and plot the outputs of hundreds or thousands of simulations. The user simply uploads the output files to a server, and then can immediately visit a custom URL to see their results. The website is populated with an array of plots that show, e.g., incidence, hospitalizations, virus strains, and testing rates for a given scenario. As shown in a screenshot of Covid-Sim (See Figure 2.1), the user can change the parameters on the left side of the screen, to see how the results

[4] There is no death in EpiSim.

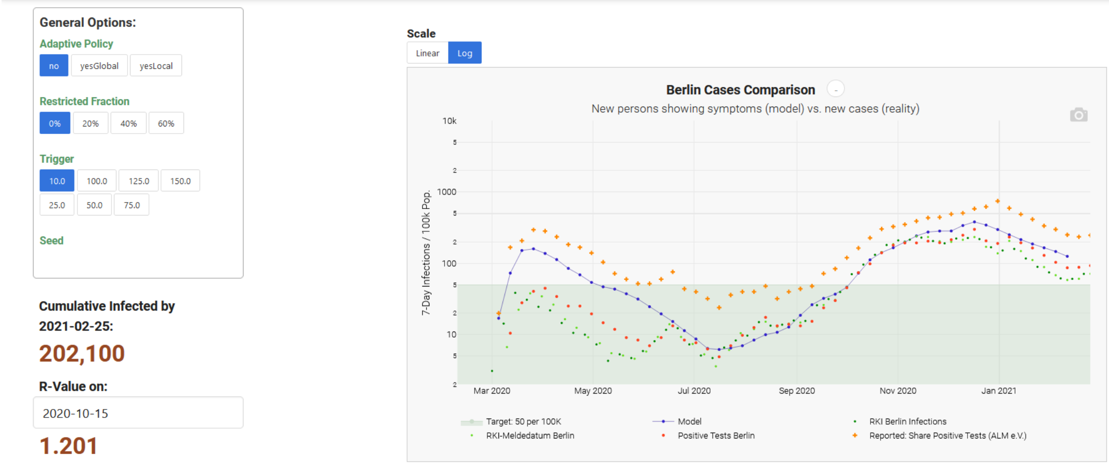

Figure 2.1: **Web Interface of Covid-Sim:** The upper left panel ("General Options") shows three variables of an illustrative case study, each with several options. The right side shows incidence plot corresponding to the selected options (scrolling down would reveal more plots). The screenshot, taken on December 9, 2021, corresponds to case study B2 of this thesis: `https://covid-sim.info/jakob/master/b2`

change.

Readers are encouraged to visit `https://covid-sim.info` to view recent simulation results and read the corresponding reports.[5]

2.6 Uses of EpiSim

The last subsection presented Covid-Sim, which allowed the rapid visualization of EpiSim's outputs. There are several aspects of EpiSim that make it suitable to the ever-changing conditions of the pandemic.

First, EpiSim efficiently stays up-to-date by using dynamic data, as described in [8]. Primarily, the activity reductions of individuals—self-initiated and due to government interventions—is calculated on a daily basis using cell-phone provider data. Secondly, the number of infectious people entering the study area every day ("disease import") varies throughout the year, due to e.g. vacation times or travel restrictions. Thirdly, the adherence to mask mandates also varies year-round. Finally, the weather and temperature changes seasonally, which has an impact what proportion of activities occur outdoors (lower chance of infection) [8]. The ever-changing base of data can be easily inputted into EpiSim, allowing current epidemiological situations to be shown.

[5]The Covid-Sim links will be posted as footnotes for each case study in the results: Chapters 4 and 5.

Also, the mechanical interactions between agents in EpiSim allow new policy proposal to be quickly integrated into the software: e.g a rapid testing program. At the beginning of a simulated day, a portion of the population is tested; the infection status of each individual is checked, and the "test result" is attached as attribute to the agent (some test results will be false negatives). A positive test result could lead to the application of a follow-up PCR test or a quarantine for the individual. Since each agent has certain attributes, the testing program can also be applied differently between sub-populations; e.g. school children are tested three times a week, while adults only get tested once.

Due to its data-driven and mechanical nature, EpiSim was able to be at the fore-front of political and epidemiological debates in real-time, particularly in the context of Berlin. The VSP team submitted its first report [26] to the German Ministry of Research on April 8, 2020, where the effects of school closures were examined. Since then, VSP has consistently delivered reports on COVID-19 development and policy recommendations every couple of weeks [8]. The next two reports incorporated mask-usage [27] and contact tracing programs [28].

Over the winter of 2020/2021, two major factors with opposing impacts were added: vaccination strategies [29] and a new variant of concern (VOC): B.1.1.7 [30]. In the new year, curfew options were explored [30] as well as rapid testing strategies [31]. As the thesis is being written, booster vaccines are being simulated [32] and a new study region is being explored: Cologne [33]. VSP's research findings, which use EpiSim, have been cited in Germany's supreme court ruling on the constitutionality of government-imposed lockdowns [34].

2.7 Discussion

This overview of research applications is meant to show that EpiSim is actively used to report on the ever-developing coronavirus progression and provide policy recommendations. The software design is also well-suited to explore the research questions proposed in Chapter 1. The trajectory of this thesis is to localize coronavirus projections and recommendations by varying infection dynamics by the home or activity locations of individuals. In the following chapter, EpiSim will be shown to be readily expandable to meet this goal. On the one hand, it is agent-based; the home and activity locations are known for each individual. On the other hand, there is plenty of data available on neighborhood characteristics. For instance, the cell phone mobility data is available on a zip-code level; while this is aggregated on a city scope in the standard EpiSim model, it can also be aggregated on a neighborhood basis.

The following chapter will describe the functionality added as part of this thesis, as well as show how the case studies are set up.

Chapter 3

Methodology

The contributions of this thesis are of two types: 1) added functionality to EpiSim and 2) applications of that functionality in the form of case studies. This chapter is split up in this same manner; first, the software extensions are described and then the configuration and setup of the four case studies will be detailed.

3.1 Functionality

With the aim of localizing EpiSim, the research for this thesis produced two main functional expansions to the software. The centerpiece of thesis is the location-based remaining fraction, which allows for differentiated activity reductions depending on an individual's home or activity location. Building on that, the second significant expansion was to localize the `AdaptivePolicy`, allowing dynamic restrictions to be applied to sub-regions of the study area.

3.1.1 Location-Based Remaining Fraction

As described in Section 2.3.3, the "Remaining Fraction" is the percent of activities people perform on any given day, compared to the pre-pandemic times. The reduction of activities is influenced by government mandates, as well as personal risk-mitigating decisions. Previously, there was a global daily remaining fraction value that would be applied to the entire study area. With the addition of location based remaining fractions, differing activity reduction could be applied depending on the home or activity location. While the following case studies split up Berlin into 12 boroughs, this functionality can be applied to any type of geographic scope: e.g. zip-code, neighborhood, state, country, etc.

As described in Section 2.3.3, the container that holds the remaining fraction for a given day and activity type is `Restriction`. To allow for different sub-regions to have varying remaining fractions, a new component was added to the Restriction container: locationBasedRf. This data structure (a "map") can be filled with location-based restrictions; each sub-region's name linked with the corresponding local remaining fraction. The manner in which this novel data-structure is filled will be described in detail later in this chapter, when the case studies are presented.

Section 2.4 described how the DefaultParticipationModel determines which activities an individual completes or skips at the start of each day (iteration). This is done based on the remainingFraction. To force the participation model to use the local remaining fractions instead of the general ones, a new class LocationBasedParticipationModel was introduced. This checks the episim-config to see which option for "DistrictLevelRestrictions" is active.[1] If "no" is selected, the global remaining fraction is used and the locationBasedRfs are disregarded. If yesForActivityLocation is active, the participation model checks whether the activity's location is within one of the boroughs; if so, the remaining fraction for that borough is used instead of the global remaining fraction. If, instead, yesForHomeLocation is turned on, the same process occurs using the home location of the individual in question.

For the participation model to know what borough an individual's home or activity location is in, we need to attach the corresponding borough to each individual and each facility, in the population and facilities files, respectively. For this purpose, two classes were written: DistrictLookupBerlinPopulation and DistrictLookupBerlinFacilities. They both function as follows: A) read a shape file of German zip-codes; B) for each location (home or facility), match the coordinates to the corresponding zip-code; C) if a zip-code is part of one of Berlin's boroughs, assign the corresponding borough to the location. The borough name is added as an attribute to each individual or facility in the population and facilities files, respectively. The name of the attribute is stored in the episim-config, so that the LocationBasedParticipationModel knows where to look to find borough information ("subdistrict" is the attribute name used in this thesis).

3.1.2 Local Adaptive Policy

The global adaptive policy, as introduced in Section 2.3.3, institutes a lockdown in the entirety of Berlin when the incidence surpasses a threshold. The local adaptive policy, which was developed as part of this thesis, institutes a targeted lockdown in an individual city borough if the incidence of that borough surpasses a threshold.

To calculate the incidence of a borough, the number of infections that occur in that borough must first be recorded. While EpisimReporting is already set up to keep

[1] EpisimConfigGroup had to be modified to hold configurations options relating to location-based restrictions; DistrictLevelRestriction is one such addition.

track of the county-level infections, it had to be modified to also record infection states for different scopes. The class was modified to record infection states for each date and borough.[2] `EpisimReporting` provides this information to the `AdaptivePolicy`, to allow it to calculate incidences of the individual boroughs.

A new configuration option was added to the configuration of the class `AdaptivePolicy`: `restrictionScope`. If set to "local", the `AdaptivePolicy` will apply restrictions on individual boroughs of Berlin instead of the city as a whole. At the beginning of each iteration, the `AdaptivePolicy` calculates the incidence for each borough using the infection data provided by `EpisimReporting`. If the borough is either in its initial state or open state, and the incidence surpasses the pre-defined threshold, then the borough will be placed in the restricted state. This causes the borough's remaining fraction within the map `locationBasedRf` to be replaced by the restricted remaining fraction, as specified in the adaptive policy configuration. Conversely, if the borough is restricted and the incidence is below the incidence threshold (for a certain period of time), the open remaining fraction will be applied.

`EpisimReporting` was also modified to create a new output file, adaptiveRestrictions.tsv, which reports the restriction status (initial, open, or closed) for each day, borough, and activity type throughout the simulation.

3.2 Case Studies

In order to address the four research questions posed in Chapter 1, four case studies were designed. This section describes the study area and time-frame for the case studies. It will give a general description of how the run classes for the case-studies are set up. These are needed to start the simulation. Following, Sections 3.3 and 3.4 will present the unique configurations of the four case studies.

3.2.1 Study Area and Time-Frame

Berlin, the capital of Germany, was chosen as the study-area for the case studies presented in Chapters 4 and 5. The main reason Berlin was chosen is that the VSP team had already created and calibrated an EpiSim scenario for the city. Thus, the time allotted for this master's thesis could be spent expanding EpiSim, rather than calibrating a scenario for a different city or region.

Figure 3.1 shows the city of Berlin, including the 12 administrative boroughs that it is composed of. The city-wide population is over 3.7 million; the population of the boroughs range between 245,000 (Spandau) and 410,000 (Pankow) [37]. If these boroughs

[2] This information is also saved as an output—subdistrict_infections.txt—to be used for analysis after the simulation is complete.

Figure 3.1: **Study Area—Map of Berlin:** Map of 12 administrative boroughs that make up Berlin, CC BY-SA 3.0 [35]

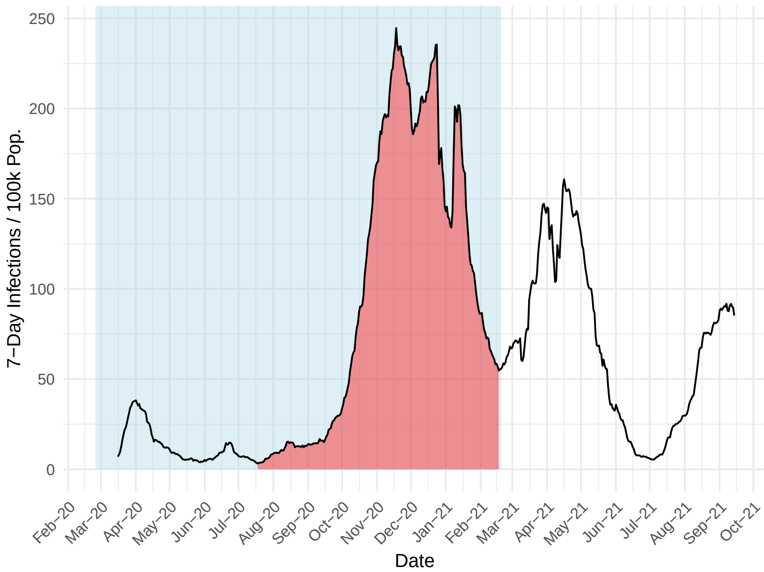

Figure 3.2: **Incidences in Berlin—Time-Frames of Case Studies:** Case Studies A1 and A2 examine the time-frame indicated by the light blue rectangle: February 25, 2020 to February 19, 2021. Case Studies B1 and B2 only examine the period encompassing the second wave of infections, which is indicated by the red shading: July 18, 2020 to February 17, 2021. Data retrieved from [36].

were cities themselves, they would each be considered a "Großstadt"—large city—in Germany [38].

The simulations in the four case studies each ran 360 iterations. The first simulated day was February 25, 2020 and the final day was exactly 360 days later: February 19, 2021. For the results presented in Chapter 4 (case studies A1 and A2), the entire study frame is examined (see blue rectangle in Figure 3.2). For Chapter 5 (case studies B1 and B2), a shorter time frame was selected: the period encompassing the second wave of infections. The study period is flanked by the days with lowest recorded incidences on either side of the second wave: July 18, 2020 had an incidence of 3.2, while February 17, 2021 had incidence of 54.7 (see red area in Figure 3.2).

3.2.2 Run Class Setup

Each case study has a unique run class, but we will begin by describing the parts that are the same for all case studies. Each run class starts by initializing the `SnzBerlin-ProductionScenario`, which sets up a Berlin scenario and populates the episim-config with default values for Berlin. This is where the activity reduction input file is read and processed to fill the `Restrictions`; it also includes, e.g., vacation times, weather patterns, lockdowns of schools, and vaccination rates.[3] For computational purposes, the sample size is set to 25%; this means that only a quarter of the population of Berlin is simulated. The results presented in Chapters 4 and 5 scale all outputs back up to 100%.

The run classes are generally set up to run a batch of simulations; this allows users to compare the results of simulations with varying configurations. Thus, each run class contains a "Parameters" section: here, the parameters are defined, and the arguments for each parameter are specified.

For example, if batch of simulations was meant to evaluate a nightly curfew, one important parameter could be the time at which the curfew begins. There could be three arguments for start time: 18:00, 20:00, and 22:00. Another parameter could be the percentage of the population that adheres to the curfew. The arguments for curfew compliance could be 60% and 80%. In order to compare every combination of those two parameters, the run class would produce 6 scenarios.

There are stochastic processes in EpiSim, which influence the simulation results. When comparing the results of two scenarios, we want to be sure that differing incidence curves can be attributed to difference in the arguments rather than stochastic effects. Thus, we run each scenario multiple times with different random seeds. Going back to the nightly curfew example, 6 scenarios with 10 seeds each would result in 60 simulations.

[3] As part of this thesis, a new section was added to `SnzBerlinProductionScenario` to allow new localized input files to be read: events, population, and facilities.

After all simulations are have run, the 60 sets of results will be aggregated back down to the 6 scenarios.

Each of the following two sections corresponds to one research focus, presented in Sections 1.4 and 1.5. Each research focus is associated with two research questions. A total of four case studies were designed to answer the four research questions. Each of the case studies will explore different parameters and arguments, which will be presented within the respective subsections. All case studies are run with 10 random seeds.

The localized functionality of EpiSim is located on the "locationBasedRestrictions" branch of the software on GitHub.[4] Simulations were run using the commit version `a7988322f0994471b0f00ca7b31021d35487a179` on the same branch.

3.3 Research Focus A: Local Infection Dynamics

As described in Section 1.4, research focus A explores EpiSim's ability to capture the infection dynamics of sub-regions within the study area. The two research questions ask whether the addition of localized activity reductions (RQ A1) or localized contact intensities (RQ A2) yield more accurate incidence curves for the sub-regions. The following subsections present the case studies designed to answer these two research questions.

3.3.1 Case Study A1: Localized Activity Reduction

As outlined in Section 1.4.2, Case Study A1 is designed to answer the following research question:

RQ A1: *Does the addition of localized activity reductions improve EpiSim's ability to capture local infection dynamics?*

Case Study A1 attempts to better capture the infection dynamics of individual boroughs of Berlin by adding location-based remaining fractions. The base case is the standard EpiSim scenario, which uses global remaining fractions. The policy case is a scenario where where the global remaining fraction is superseded by the local remaining fraction for individual boroughs.

To prepare this case study, the `locationBasedRf` fields must be filled with the remaining fractions for Berlin's boroughs for the entire time-frame of the case study. To start, the class `AnalyzeSnzData` was used to parse the cell-phone mobility data and produce an input file per borough containing the daily activity reductions.

At the start of the simulation, the methods within class `CreateRestrictionsFromCSV`

[4] See `https://github.com/matsim-org/matsim-episim-libs/tree/locationBasedRestrictions`.

read Berlin's activity reduction input, and populate the remaining fraction field in each `Restriction`. The methods within this class had to be expanded to also read all 12 borough-specific activity reduction files and populate the `locationBasedRf` field as well.

Case Study A1: 2 scenarios $\times$ 10 seeds = 20 simulations

Parameter	Arguments
DistrictLevelRestriction	no, yesForHomeLocation

The only parameter that varied in this case study is the episim-config option "DistrictLevelRestriction." The base case, which uses global remaining fractions, corresponds to the "no" setting. "yesForHomeLocation" is used for the policy case, where borough-based activity reductions are applied to agents living in the respective boroughs. Since the activity reduction data relates to the home location of individuals, the option "yesForActivityLocation" was superfluous in this case study. The results of Case Study A1 are discussed in Section 4.2.

3.3.2 Case Study A2: Localized Contact Intensity

As described in Section 1.4.3, Case Study A2 explores whether the addition of regionally differentiated contact intensities into EpiSim improve the model's ability to capture local infection dynamics. The intuition is that smaller home-sizes should have higher chances of infection because the contact intensity is higher. Case Study A2 attempts to answer the following research question:

RQ A2: *Does the inclusion of localized contact intensity for home activities (based on varying home-size) improve EpiSim's ability to capture local infection dynamics?*

In the standard EpiSim model, a different value for contact intensity is applied for each activity type. All home events have a contact intensity of 1.0. The variation of contact intensities based on home-size was completed in two steps: 1) Replace standard home events with home_XX events, which indicate the living space per person (e.g. home_25 indicates 25m^2 per person) and 2) specify a varied contact intensity for each of the size-differentiated home_XX events.

Varying home sizes were applied to agents based on their home-location. Data [39] on the average home-size per person is available on the level of a Berlin-specific planning unit named lebensweltlich orientierter Raum (LOR). Since there are 448 LORs in Berlin, this planning unit has a higher resolution than the 12 boroughs in Berlin. An R script

home_15:	1.0 + 3 ∗ ciModifier
home_25:	1.0 + 2 ∗ ciModifier
home_35:	1.0 + 1 ∗ ciModifier
home_45:	1.0
home_55:	1.0 - 1 ∗ ciModifier
home_65:	1.0 - 2 ∗ ciModifier
home_75:	1.0 - 3 ∗ ciModifier

Table 3.1: Modified Contact Intensity Equations

was written to clean this data, and prepare it to be used for EpiSim.

The java class `DifferentiateHomeSizeInEvents` was written to translate the default home events into home_XX events; if an agent lives in an LOR where the average home size per person is between 40 and 50m^2, their home events will be transformed into home_45 events. By splitting up the general home event type into 7 activity types (between 15m2 and 75m^2), differing contact intensities can be applied to each home_XX event.

Case Study A2: 88 scenarios × 10 seeds = 880 simulations

Parameter	Arguments
ciModifier	0.0, 0.1, 0.2, 0.3
thetaFactor	0.95, 0.955, 0.96, 0.965, 0.97, 0.975, 0.98, 0.985, 0.99, 0.995, 1.0
DistrictLevelRestriction	no, yesForHomeLocation

The first parameter, ciModifier, controls to what degree the contact intensities for home events are skewed. As described earlier, home events occurring in areas with below average m^2 per person receive higher contact intensities (thus, a higher chance of infection) and vice versa. The way in which ciModifier was used to skew contact intensities is delineated in Table 3.1. When the ciModifier is 0.0, contact intensities for all home sizes is 1.0, just like the standard EpiSim model. When the ciModifier is 0.3, the contact intensities vary more: between 0.1 for home_75 events and 1.9 for home_25 events.

In addition to ciModifier, two more parameters were added to this case study. Changing the contact intensities for home events disrupts calibration of the model. A calibration adjustment factor (thetaFactor) was therefore also included to slightly decrease the overall chance of infection for the entire model. In the simulation analysis, we will have to compare the base case (with uniform contact intensities) and a policy case (with skewed contact intensities). In the analysis, we will choose the thetaFactor which causes the policy case to have a similar number of total infections in Berlin as the base case. We want to examine the effects of localized contact intensity on the borough level, but don't want the Berlin-wide incidence curves to differ significantly.

We also included the DistrictLevelRestrictions, to see whether the combination of local activity reduction (see previous subsection) and varied contact intensities would improve the results. The results of Case Study A2 are discussed in Section 4.3.

3.4 Research Focus B: Local Lockdowns

The case studies presented in the last section attempt to improve EpiSim's ability in capturing local infection dynamics. Research focus B attempts to evaluate the merit of implementing local restrictions in an EpiSim simulation.

For the following cases, the global `remainingFraction` is applied to all boroughs in Berlin, unless a lockdown is instituted in a borough. Once a localized restriction is applied to a borough, the borough is added to the `locationBasedRf`, and the corresponding Rf will be used instead of the global one. This means that the `locationBasedRf` will no longer be filled with historical borough-based activity reductions from cell-phone data.

The following two case studies respond to RQ B1 and RQ B2.

3.4.1 Case Study B1: Pinpointed Lockdown

Case Study B1 examines the application of a localized lockdown to answer the following research question:

RQ B1: *How does a local lockdown affect the infection dynamics of the restricted region and the un-restricted regions?*

Case Study B1 restricts the Berlin borough of Mitte for the month of October 2021 (during the rise of the second wave). The choice to restrict Mitte for a month is relatively arbitrary; the purpose of this case study is to demonstrate the effects of a local lockdown. As described earlier, the global remaining fraction applied for all boroughs and all activities. Only for leisure activities, did the `locationBasedRf` contain a single entry between October 1 and October 31, 2020: Mitte - 0.0.

Case Study B1: 3 scenarios × 10 seeds = 30 simulations

Parameter	Arguments
DistrictLevelRestriction	no, yesForHomeLocation, yesForActivityLocation

The only parameter for Case Study B1 is DistrictLevelRestriction. Two modes of location based restrictions were tested: a) all agents who live in Mitte do not conduct leisure activities and b) all leisure activities that are supposed to be performed in Mitte, regardless of home location, are banned. These two options correspond to "yesForHomeLocation" and "yesForActivityLocations," respectively. This is the only case study where the argument of "yesForActivityLocation" is explored. When DistrictLevelRestriction is set to "no," Mitte will not be restricted; this corresponds to the base case. The results of Case Study B1 are discussed in Section 5.1.

3.4.2 Case Study B2: Local Adaptive Restrictions

Finally, Case Study B2 explores the effects of a local (and global) adaptive policy in Berlin.

RQ B2: *How does a local adaptive policy affect incidences and time uses, compared to a global adaptive policy? How do the parameters of the adaptive policy affect its benefit?*

Previously in this chapter, the development of the local adaptive policy was described. Case Study B2 compares the effects of a local adaptive policy for the 12 boroughs of Berlin to a global adaptive policy for the entirety of Berlin.

As described in Chapter 2, parameters for adaptive policy include the remaining fractions for the initial, restricted, and open policies. The initial policy was set to reflect the historical daily remaining fractions gleaned from the mobility behaviour from cell phone data. For this case study, the global `remainingFraction` is used for the initial policy. The `locationBasedRf` was only filled if a boroughs enters a restricted or open phase. In this case study, DistrictLevelRestriction fixed to "yesForHomeLocation" for all simulations; thus, local restrictions apply to the residents of the borough in question.

The remaining fraction for the open policy was 0.9 for all activity types. This assumes that even if a borough is unrestricted, pre-coronavirus mobility levels will not be reached (e.g. because employees still make use of home-office options). For a borough to open up, the incidence must be below the Trigger for two weeks.

Case Study B2: 84 scenarios $\times$ 10 seeds = 840 simulations

Parameter	Arguments
adaptivePolicy	no, yes-local, yes-global
Rf	0.0, 0.2, 0.4, 0.6
Trigger	10, 25, 50, 75, 100, 125, 150

The parameter adaptivePolicy controls whether an adaptive policy should be applied; if so, whether it should be a global or local adaptive policy. The parameter Rf specifies the remaining fraction for the restricted policy. For an argument of 0.0, no activities will occur during a restricted phase. The final parameter, Trigger, specifies the incidence threshold; if the incidence of a borough surpasses the Trigger, it will enter the restricted phase. The results of Case Study B2 are discussed in Section 5.2.

Chapter 4

Results A: Local Infection Dynamics

This chapter presents the results of case studies A1 and A2, which are geared to respond to RQ A1 and RQ A2, respectively. The research focus behind these two case studies is to improve EpiSim's ability to capture the local infection dynamics of individual boroughs. The need for this improvement will be shown in Section 4.1; herein the incidence curves of the standard EpiSim model will be compared to the incidences reported by the Robert Koch Institut (RKI). Sections 4.2 and 4.3 will present two attempts to improve the localized analysis: first, by varying the activity reductions by borough; second, by assigning higher contact intensities to areas with small home sizes (resulting in higher chances of infection).

4.1 Status-Quo

Figures 4.1 and 4.2 show the incidence curves for the second wave on a city and borough level, respectively; the black plot shows the cases reported to the health authority and the blue plot shows simulation output of the standard EpiSim setup. Starting with Figure 4.1, it can be seen that the simulation projects higher incidences than reported to the health authority; particularly during the peak of the second wave in December 2020. Looking at the borough level in Figure 4.2, the same observation can be made in most boroughs. A possible explanation for this is that under-reporting becomes more common when the health care system is overloaded; e.g. when contact tracing and testing programs are over capacity, less cases will be captured.

Looking again at the simulated incidences per borough (blue plot in Figure 4.2), we notice that the boroughs generally follow similar progressions; all 12 peak in mid-

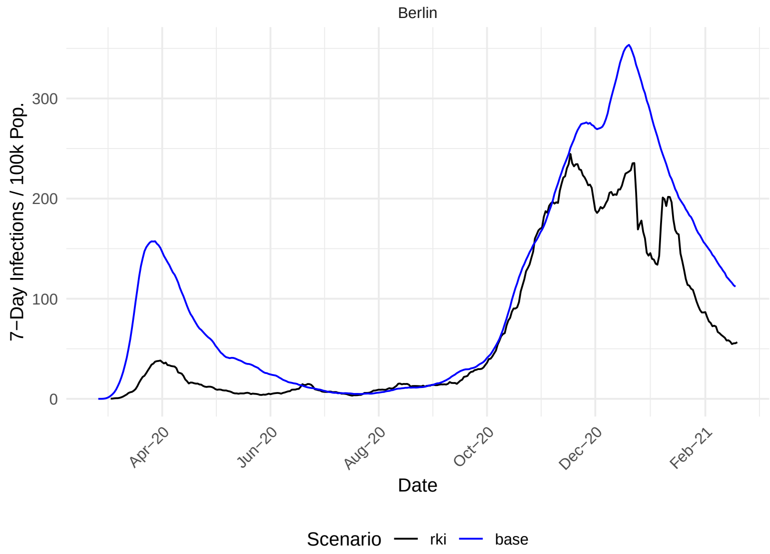

Figure 4.1: **Incidence in Berlin—EpiSim vs. RKI:** Comparison between standard EpiSim model (blue plot) and the reported cases to RKI (black plot) for entire Berlin. The gap between plots during wave peaks could indicate under-reporting.

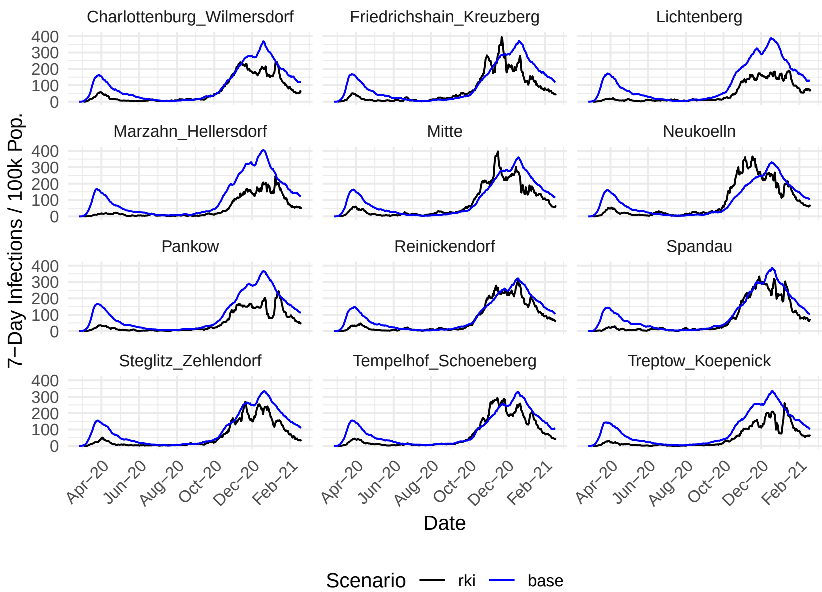

Figure 4.2: **Incidence in Berlin's Boroughs—EpiSim vs. RKI:** Comparison of standard EpiSim model with RKI cases, faceted by Berlin Borough.

December with incidences between 300 and 400. The cases reported to the RKI (black plot), on the other hand, vary more drastically: for instance, Neukölln consistently has a higher incidence than Lichtenberg.

To summarize, whereas the reported incidences vary significantly between boroughs, the simulated incidences vary less between boroughs. This indicates that while the standard EpiSim simulation is calibrated to model the infection dynamics of the entire study region, the dynamics of sub-regions are harder to capture. The following two sections attempt to improve the analysis of local virus spreading by individualizing the simulation: first, by applying activity reduction rates per borough rather than a Berlin-wide remaining fraction; second, by varying the chance of infection by an individual's home size (contact intensity).

If we do see changes between the base case and the policy case in the following subsections, it will be difficult to ascertain whether the change is an improvement. Intuitively, we would like to see the gap between the policy case and RKI cases be constant between boroughs. However, we do not know whether all boroughs are equally impacted by under-reporting.

4.2 Case Study A1: Localized Activity Reduction

Case Study A1 incorporates `locationBasedRfs` for Berlin's boroughs in order to improve the simulated infection curves for the individual boroughs. The broader purpose is to answer the following research question:

RQ A1: *Does the addition of localized activity reductions improve EpiSim's ability to capture local infection dynamics?*

The research question will be addressed in two parts: do the residents of borough A reduce their activities on average more than the residents of borough B (for whatever reason)? If so, does this decrease borough A's simulated incidence curve such that it is closer to reality?

Figure 4.3 shows the simulation results for Case Study A1.[1] The base case (blue plot) uses the global `remainingFraction`, while the policy case (red plot) uses the `locationBasedRf` (as discussed in Section 3.3.1). In many boroughs, the localized activity reduction doesn't produce a large change in the incidence curve with respect to the the standard EpiSim model. This indicates that activity reduction in most boroughs is

[1] More simulation results for Case Study A1 can be found here: `https://covid-sim.info/jakob/master/a1`.

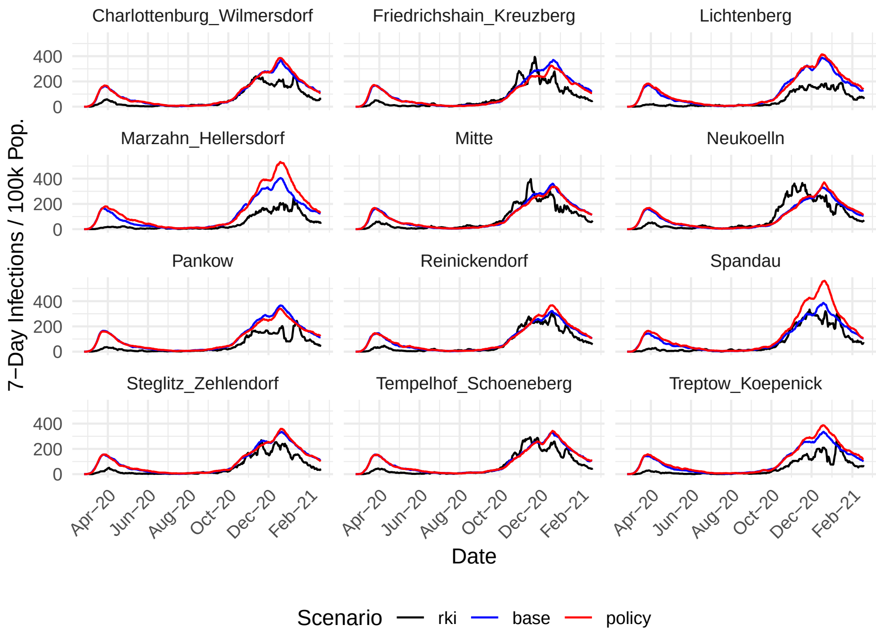

Figure 4.3: **Incidence in Berlin's Boroughs—Borough-Based Activity Reduction:** Set up analogously to Figure 4.2, this plot shows incidence curves for each borough. The policy case (red plot), uses local remaining fractions for each borough. The base case (blue plot) uses a single global remaining fraction for all boroughs.

similar to the Berlin average. The major exceptions are two outer-boroughs—Marzahn-Hellersdorf and Spandau—where the policy case produces a higher second wave than the base case, indicating that residents completed more activities out-of-home than the average Berliner. Perhaps their populations' behaviors differ more strongly from that of the other boroughs because they far from the city-center.

In those boroughs where the policy case produces a different incidence curve than the base case, the policy case generally doesn't actually show an improvement. In Kreuzberg-Friedrichshain, the author expected the policy case to have higher incidences than the base case (so as to better match the RKI numbers); however the opposite is true. In Marzahn-Hellersdorf, a reduction in infections was expected; again, this was not realized.

The inclusion of location-based remaining fractions in EpiSim does change the incidence curves in some boroughs. This change, however, is not generally very large, which indicates that activity reductions don't vary significantly between most boroughs. Responding to RQ A1, the inclusion of borough-based remaining fractions doesn't seem to improve EpiSim's ability to capture the boroughs' infection dynamics.

4.3 Case Study A2: Localized Contact Intensity

Case Study A2 attempts to answer the following research question:

RQ A2: *Does the inclusion of localized contact intensity for home activities (based on varying home-size) improve EpiSim's ability to capture local infection dynamics?*

In this case study we skewed the contact intensity based on average home-size per LOR; agents living in an area with small home sizes will have higher contact intensities (and thus, a higher chance of infection). RQ A2 will be explored in two parts: does the inclusion of localized contact intensity have an impact on the incidence curves of the boroughs? If there is an impact, does it show an improvement in the simulation's ability to capture the infection dynamics of individual boroughs?

Figure 4.4 shows the policy case, where the contact intensity is differentiated by home-size.[2] For the policy case shown, a ciModifier of 0.3 is used (to produce maximum skewed contact intensities). The calibration adjustment factor (thFactor) is set to 0.96.

Looking at Figure 4.4, the difference between the base case and policy case is minuscule for all boroughs. Thus, adding localized contact intensities didn't show any

[2]More simulation results for Case Study A2 can be found here: `https://covid-sim.info/jakob/master/a2`.

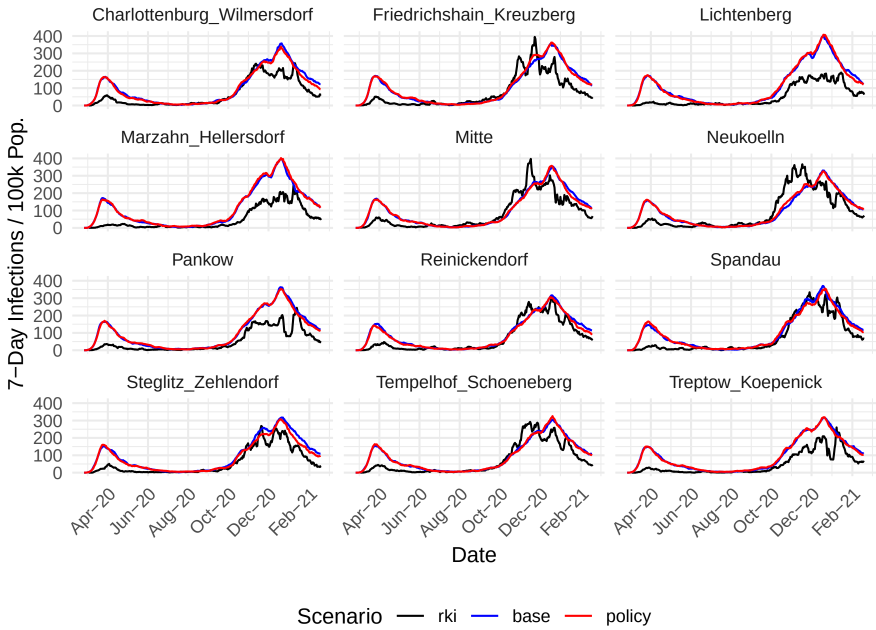

Figure 4.4: **Incidences in Berlin's Boroughs—Localized Contact Intensity:** Shows incidence curves for each Berlin borough. Black plot shows cases reported to the RKI. Blue plot shows the output of the standard EpiSim model, wherein contact intesnity is uniform across the study area. The red plot shows the policy case, where agents living in LORs with small home-sizes have higher contact intensities. Parameters used for policy case: ciModifier = 0.3, thetaFactor = 0.96, and DistrictLevelRestriction = no (see Section 3.3.2 for more information).

significant effects. Figures 4.5 and 4.6 show that although there are significant differences in average home-sizes between LOR, the difference is not large between boroughs. Regarding RQ A2, the inclusion of localized contact intensity does not improve EpiSim's ability to capture the infection dynamics on the borough level. If we were to look at the infection curves per LOR, we might see more interesting results; unfortunately, data on the reported cases is not available from RKI on the resolution of LOR.

4.4 Discussion

The first section showed that there is a overall divergence between the base case of EpiSim for individual boroughs and the reported numbers. The purpose of this chapter was to add location-based factors to the simulation in order to improve simulation's ability to capture the incidence curves of individual boroughs. Two factors were added, varied

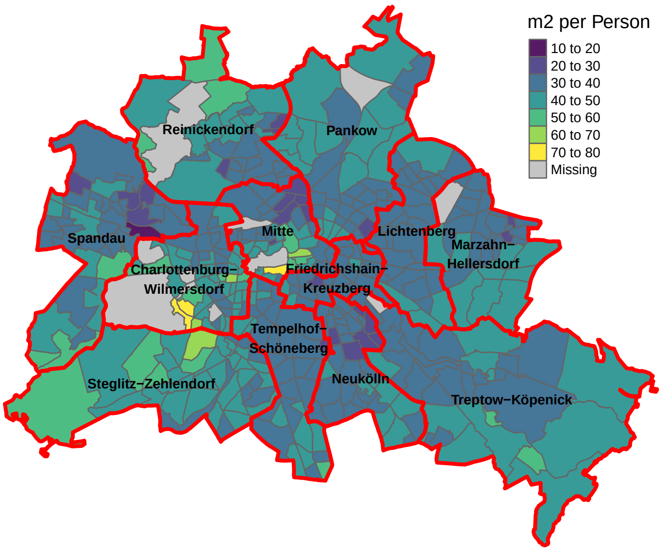

Figure 4.5: **Home Size in Berlin's LORs:** Shows average home size (meters squared) per person. Red lines mark the borough boundaries, which are also named within the figure.

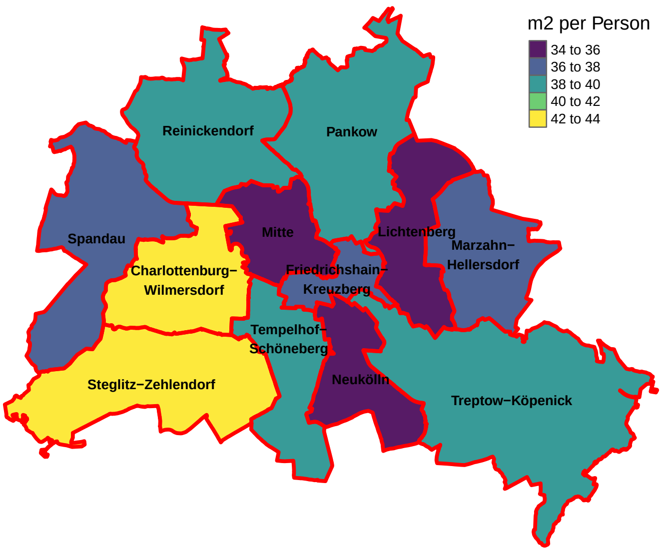

Figure 4.6: **Home Size in Berlin's Boroughs:** Aggregation of 4.5 on the borough level; weighted average was taken using the population of each LOR.

across geographical areas: activity reduction per borough, and home size per LOR. Neither approach showed a significant improvement of the model's ability to capture the local infection dynamics.

This is, in itself, is a significant result. The divergence between borough's incidence curves cannot be explained by the fact that one borough's residents act differently (e.g. less "responsibly") than another borough's residents. The divergence can also not be explained by different home sizes. Further work could examine whether there is a geographic dichotomy in other demographic factors, which could improve EpiSim's local projections. Income, for example, could serve as a proxy for factors that have an impact on infection chances: e.g. access to healthcare or ability to work from home.

This chapter also shows that localization on the scale of borough doesn't yield promising results in Berlin. Each Berlin borough has a population of between 200,000 and 400,000. At that size, the internal composition of individual boroughs is probably rather heterogeneous. Thus, average activity reduction and average home-size doesn't vary significantly between boroughs (as shown in Figure 4.6). Further work should explore the infection dynamics of more homogeneous areas. Can the infection dynamics of individual neighborhoods be improved through the addition of localized activity reductions or home size? The incidence data for smaller scopes was unfortunately not available to the author.

The next chapter will continue to examine localized activity reductions; however, these will not stem from the real mobility patterns gleaned from cell phone data. Instead, artificial activity reductions will be imposed in the form of localized lockdowns.

Chapter 5

Results B: Local Lockdowns

This chapter presents the results of case studies B1 and B2, which are geared to respond to RQ B1 and RQ B2, respectively. Research focus B is concerned with evaluating the merit of localized lockdowns.

Section 5.1, which shows Case Study B1, examines the effects of a month-long local lockdown in a single borough (Mitte). Then, Section 5.2 evaluates the local adaptive policy; this means lockdowns will be dynamically imposed on boroughs if their incidence surpasses a certain threshold. This chapter is rounded off with a discussion of the results.

5.1 Case Study B1: Pinpointed Lockdown

Case Study B1 applies a local lockdown to the borough of Mitte for October 2020 in order to answer the following research question:

RQ B1: *How does a local lockdown affect the infection dynamics of the restricted region and the un-restricted regions?*

Figure 5.1 shows the incidence curves for each Berlin borough.[1] As described in 3.4.1, two types of local restrictions are simulated: A) all residents of Mitte are barred from completing leisure activities and B) all Berliners are barred from conducting leisure activities in Mitte. The following two paragraphs evaluate the effects of each restriction scheme.

Restricting Mitte's Residents: Restricting people living in Mitte from conducting leisure activities (policy-home, red plot) makes a significant dent in Mitte's incidence

[1] More simulation results for Case Study B1 can be found here: `https://covid-sim.info/jakob/master/b1`.

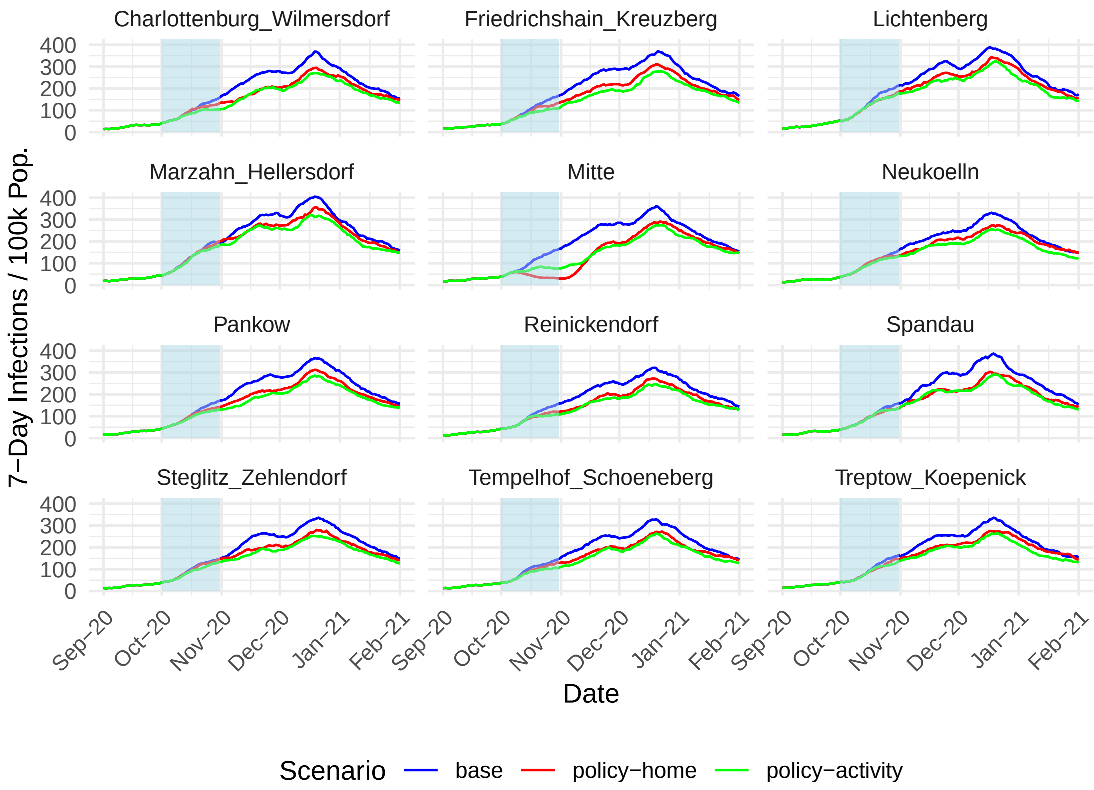

Figure 5.1: **Incidence in Berlin's Boroughs—Local Lockdown of Mitte:** Incidence plot for each borough during the second wave. Base case (blue) shows the standard EpiSim model without a local lockdown. Policy-home (red) shows the simulation output when residents of Mitte are barred from conducting any leisure activities. Policy-activity (green) is the case, where no Berliners are allowed to conduct leisure activities within Mitte. The lockdowns of both policy cases occur during the month of October 2020, indicated by the blue rectangle.

curve. These results make sense, since Mitte's residents have significantly fewer contacts.[2] In other boroughs, the infections continue to increase; however, their incidence curves have more shallow slopes and do no reach the same peak as in the base case. Even though residents of other boroughs are not restricted in any way, they benefit from meeting fewer Mitte residents, who could potentially be infectious.

Restricting Leisure Activities in Mitte: Restricting Berliners from doing leisure activities in Mitte (policy-activity, blue plot) lead to a drop in incidences in all boroughs. The most significant drop occurs in Mitte. A greater proportion of Mitte's residents, compared with residents of other boroughs, probably complete leisure activities in Mitte itself. This leads to an out-sized impact of the policy on Mitte's incidence.

In most non-Mitte boroughs, the policy-activity makes a bigger dent in their incidence curves than policy-home. While policy-home in no way impacted the movement of most Berliners, policy-activity applies to a portion of all boroughs' residents: those who want to eat dinner or visit a friend in Mitte. Thus, while a portion of the policy-activity's benefit is applied in Mitte, the rest is distributed over the other boroughs. Overall, policy-activity has a greater Berlin-wide impact in reducing infections: 13.6% for policy-home and 19.2% for policy-activity reduction with respect to the base-case.

5.2 Case Study B2: Local Adaptive Restrictions

Case Study B2 is designed to answer the following research question:

RQ B2: *How does a local adaptive policy affect incidences and time uses, compared to a global adaptive policy? How do the parameters of the adaptive policy affect its benefit?*

The final investigation of this thesis was to examine how effective an adaptive restriction policy could have been in mitigating the effects of the second wave that hit Berlin starting in Fall 2020. The global adaptive policy institutes a lockdown in the entirety of Berlin when the incidence surpasses a threshold. The local adaptive policy, which was developed as part of this thesis, institutes a targeted lockdown in an individual city borough if the incidence of that borough surpasses a threshold; for example, the boroughs of Spandau and Mitte could be in lockdown while all other boroughs are open.[3]

The use of lockdowns to combat the pandemic comes with a trade-off: public health vs. freedom of movement. Thus, the benefit and cost of policy case will be evaluated

[2] Note: the incidence curves are based on the home location of the infected individuals.

[3] Simulation results for case study B2 can be found here: `https://covid-sim.info/jakob/master/b2`.

using two proxies: a reduction of infections (benefit) and a decrease average time each person spends outside of their home (cost). The results of Case Study B2 will be explored in 5 steps, each with a corresponding subsection. The key findings of each subsection are summarized here:

1. The dynamics of an adaptive policy are such that an incidence increase leads to the imposition of restrictions; those, in turn, will reduce the incidence.
2. The temporal pattern of the restriction phases depend on the Trigger and remaining fraction for the restricted phase.
3. Adaptive policies carry an inherent trade-off; more stringent adaptive policies will significantly reduce infections, at the cost of time spent outside of the home.[4]
4. The benefit/cost ratio of the adaptive policies is highest in the more lenient scenarios.[5]
5. The local adaptive policy performs better than the global adaptive policy in the more stringent scenarios.

5.2.1 Relationship between Trigger and Lockdown:

Before aggregating over the 10 seeds, we will begin by looking at the incidence curves for a single local adaptive policy run. The purpose is to show how the incidence curves activate lockdowns, and, in turn, how those lockdowns influence the incidence curves. For this illustrative run, a Trigger[6] of 25 and a Rf of 0.4 for the restricted policy was chosen (and the seed 4711).

Figure 5.2 shows results of a local adaptive policy; one incidence curve per Berlin borough. The overlaid timeline shows the restriction policy active in that borough. As we can see, whenever the incidence curve surpasses the trigger of 25, the borough enters its restricted phase ("red"). Soon after, the incidences plummet. After the incidence is below the Trigger for approximately two weeks, the borough opens up once again ("blue").

Figure 5.2 shows that during some restricted periods, the incidence continues to increase after the lockdown has been established. This effect is magnified starting in December 2020. At the end of the study time-frame, the so-called Delta variant of concern (VOC) became more prevalent; a restriction regime that was able to handle the

[4]Stringent, for this section, indicates that the Rf and Trigger are low.
[5]Lenient, for this section, indicates that the Rf and Trigger are high.
[6]As described in Section 3.4.2, the Trigger serves as the incidence threshold between initial/open phase and the restricted phase, as well between the restricted phase and open phase.

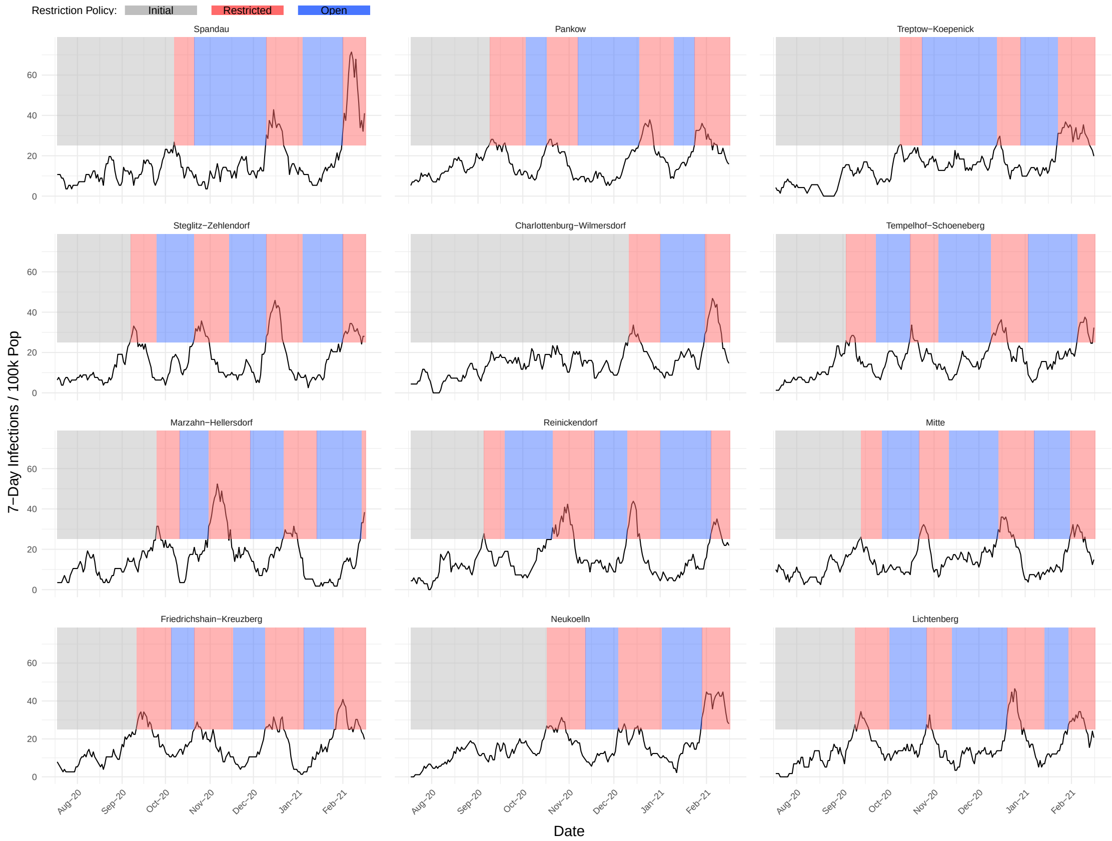

Figure 5.2: **Incidences and Restriction Timelines for Berlin's Boroughs—Local Adaptive Policy:** Each facet shows the incidence of a Berlin borough for a single seed (4711) where the remaining fraction for restricted policy is 0.4 and the trigger is 25. Timelines indicate when a borough has an initial, restricted, and open policy. The bottom edge of timeline indicates trigger value.

wild type of SARS-CoV-2 may not be sufficiently tough to stop the Delta VOC in its tracks.

Looking at the discrepancy in lockdowns, one might wonder why Charlottenburg-Wilmersdorf stays below the trigger until early December. As we saw in the previous subsection, the incidence of a borough is influenced by the lockdown measures of other regions; thus Charlottenburg-Wilmersdorf's infections could be dampened because surrounding boroughs are restricted. Additionally, while in the initial phase, boroughs follow the data-based activity reductions gleaned from cell-phone data. In late 2020, the government enacted restrictions which reduced the activity participation in Berlin. Thus, Charlottenburg-Wilmersdorf's incidence was dampened due to real-world restrictions rather than the simulated adaptive ones.[7]

This subsection has shown how the local adaptive policy functions. While all boroughs are impacted by at least two restricted phases, the lockdowns occur at different times. This demonstrates the promise of local adaptive restrictions: some boroughs with lower incidences do not need to be restricted while other boroughs are in lockdown.

5.2.2 Temporal Patterns

The last subsection showed that lockdowns occur at different times in different boroughs. As a next step, the temporal patterns of lockdowns will be explored in greater detail, as they relate to Rf and Trigger inputs. Figure 5.3 shows the policy timeline for each borough, faceted by the extreme values for Rf and Trigger.

Trigger: To begin, we will examine the effect of the Trigger; the left and right facets of Figure 5.3 show Triggers of 10 and 150, respectively. For a Trigger of 10 (left side), most boroughs go into restricted states very soon; this makes sense, since the Berlin-wide incidence was 3.2 at beginning of the study time-frame (July 18), so it doesn't take much time to reach 10. The incidences in Berlin's boroughs don't reach 150 until September, which causes the boroughs in the right facets to remain in their initial states for longer.

Figure 5.3 also shows that the lockdown phases are more synchronized between boroughs when the Trigger is high. A single borough doesn't usually reach an incidence of 150 alone; there is usually exponential growth throughout the entire city. This result is reasonable because high incidences in one borough will fuel the exponential growth in other boroughs, especially in a city as interconnected as Berlin. At lower Triggers, a local outbreak (e.g. super-spreader event) could spike a borough's incidence above the Trigger, while other boroughs are not close to the limit yet.

Rf: Next we will compare the top and bottom facets of Figure 5.3, which show a Rf

[7] Further work should consider decoupling the initial policy from the activity reductions provided by the cell phone data.

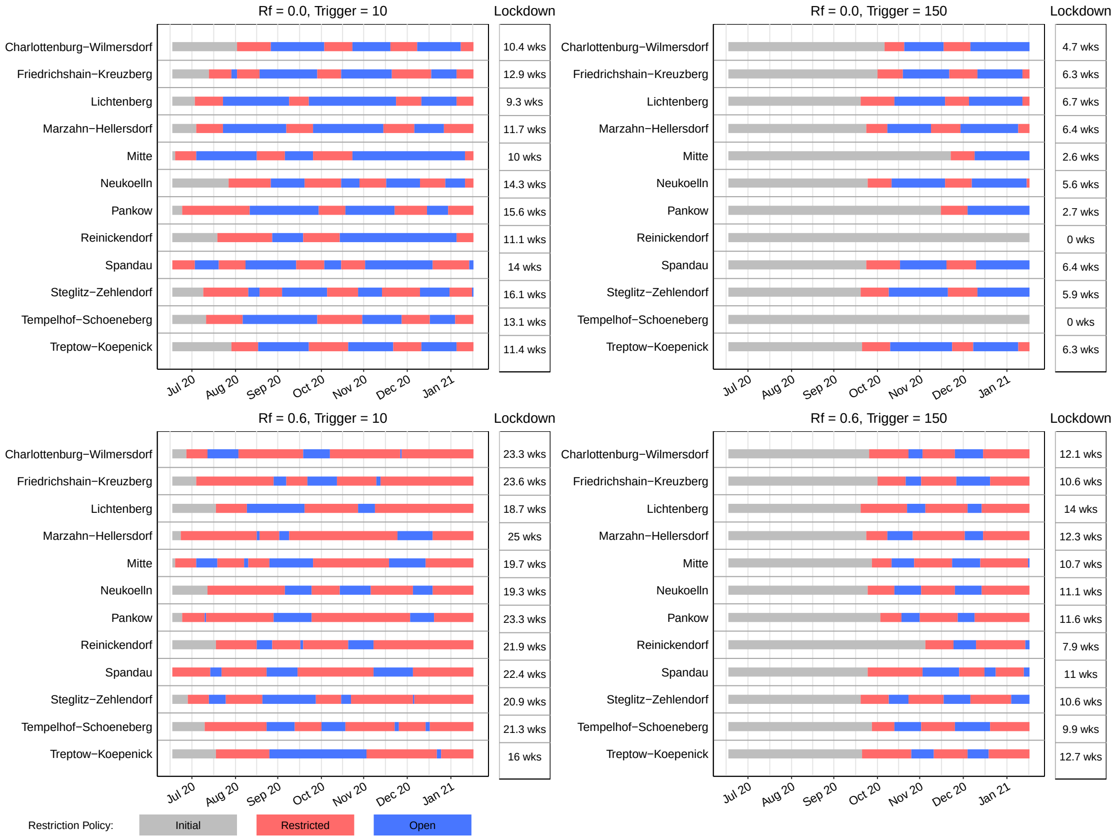

Figure 5.3: **Restriction Timelines for Berlin's Boroughs—Local Adaptive Policy:** The four facets show the combinations of the extreme values for two variables: Trigger (min is left, max is right) and Rf (min in top, max is bottom). Each facet contains a restriction timeline for each Berlin borough; the gray phase shows the initial policy, the red phase shows a lockdown policy, and the blue phase shows an open policy. The number to the right of each timeline indicates the number of weeks the borough is in lockdown. This plot uses a single seed: 4711.

for restricted policy of 0.0 and 0.6, respectively. Here we can see that the lower Rf leads to shorter lockdowns. This can be explained by the fact that more intense lockdowns will push the incidence below the trigger threshold more rapidly, allowing the borough to open up sooner. Additionally, the open phases between lockdowns last longer with a lower Rf. For a borough to open up, it needs to be below the trigger for 2 weeks; in that time, the lower Rf will force the incidence further down than with a higher Rf. From that lower level, it will take the borough longer to surpass the trigger again, leading to longer open phases.

To summarize, runs with higher triggers lead to later and more synchronized lockdowns than runs with lower triggers. Runs with higher Rfs have longer lockdowns than ones with lower Rf's.

5.2.3 Effect of Local Adaptive Policy on Incidence and Time-Use

So far, some characteristics of the restriction policies and the incidence curves have been described; now, the benefits and costs of the adaptive policies will be evaluated. Restrictions are implemented to reduce the number of people who get seriously sick, must be hospitalized, and/or die; in the context of this thesis, however, the reduction in overall infections will be considered as the main goal. A major cost of restrictions is that people cannot go to places they want to go to, or do things they want to do. The following paragraphs will consider the reduction in average time spent outside of the home as a proxy for this cost.

Figure 5.4 shows Berlin's incidences, while Figure 5.5 shows the average minutes per day spent outside of the home. Both plots are faceted by extreme values for Trigger and Rf, and each facet shows the base case, local adaptive policy, and global adaptive policy. The plots and all further analysis use aggregated data; for each metric, the mean over 10 runs is taken to remove some of the stochastic effects.

Incidence: Figure 5.4 shows the incidence progressions in Berlin for the extreme values of those two parameters. Several initial observations can be made: the local maximum of the infection waves depends heavily on the Trigger, while the local minimum depends more on the Rf of the restricted policy. All adaptive policies lead to a significant reduction in infections as compared to the base case (note the logarithmic scale on the y-axis). In the most stringent scenario—Trigger=10 and Rf=0.0—the local adaptive reduces the total infections by 96% (155,000). In the most lenient scenario—Trigger=150 and Rf=0.6—the reduction is still significant: 38% or 62,000 avoided infections.

Table 5.1 shows the percent reduction in infections of the local adaptive policy vs. the base case; not just for the extreme values of Trigger and Rf, but for all the simulated values. The trend shown in Figure 5.4 is validated: lower Triggers and Rfs result in the greatest reductions of infections. The same trend can also be seen when comparing the global adaptive policy to the base case; this is shown in Table A.1 in the Appendix.

Time-Use Reduction: The goal of reducing infections is not the only criteria for implementing lockdowns; the time lost must also be considered as a second criteria. Figure 5.5 shows average time spent outside of the home, faceted analogously to the previous figure. The most stringent adaptive policy led to the least time spent outside of the home—a 29% reduction of about 86 minutes per day, as compared to the base-case. The most lenient scenario only results in a 2 minute reduction of out of home activities.

Table 5.2 shows the percent reduction in time spent outside of the home for all simulated arguments for the parameters of Trigger and Rf. Here again, the trend is confirmed: the lower the Trigger and Rf, the less time people spend outside of their homes. Table A.2 in the Appendix shows the same trend for the global adaptive policy.

Summary: The more stringent adaptive policies lead to the greater reductions in infections; the associated cost is that people spend the less time outside of their homes. The following subsection attempts to weigh the costs against the benefits.

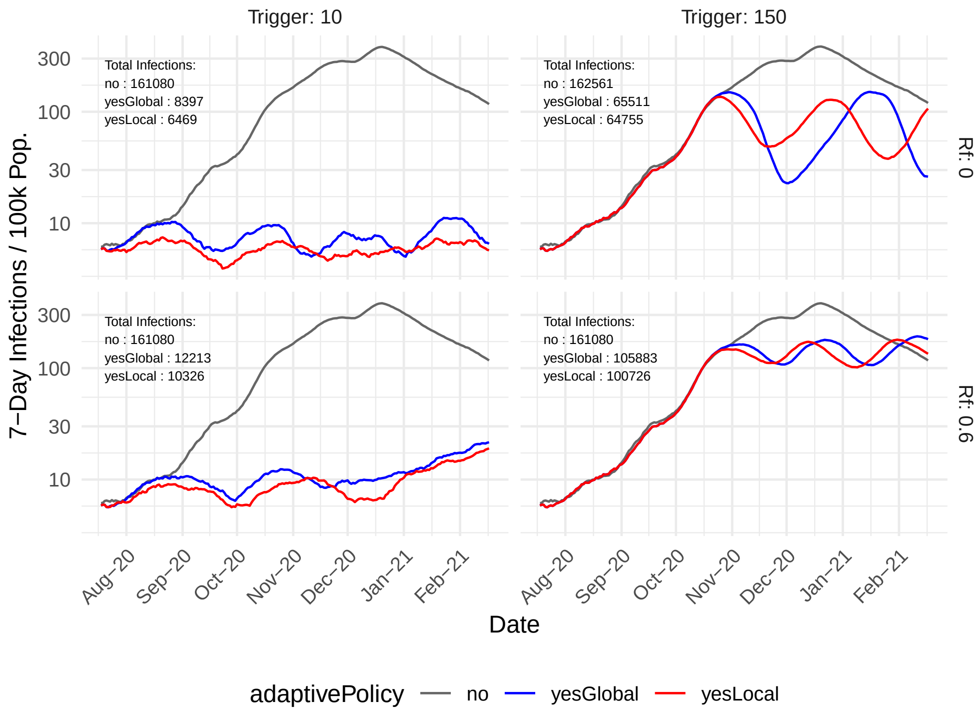

Figure 5.4: **Incidences For Berlin—Adaptive Policies:** Incidence curves (logarithmic scale) for Berlin, faceted for the extreme values of Rf and Trigger. The black plot shows the base case (standard EpiSim model). The blue and red plots show the two policy cases: the global and local adaptive policies, respectively. The text box shows the total infections in the study-time frame for each case.

	Trigger						
Rf	10	25	50	75	100	125	150
0.0	96.0	90.8	83.7	79.1	74.2	67.1	60.2
0.2	95.7	90.7	82.4	77.0	71.6	64.9	58.7
0.4	95.1	89.2	80.1	72.2	67.1	58.3	50.1
0.6	93.6	86.8	75.3	64.7	55.2	47.4	37.5

Table 5.1: **Percent Decrease in Infections, Local Adaptive Policy vs. Base:** Percent reduction calculated as follows: $(\text{Inf}_{base} - \text{Inf}_{local}) \div \text{Inf}_{base} \times 100\%$, where Inf indicates the total number of infections in the study time-frame; *local* and *base* indicate the local adaptive policy and base case, respectively. For example, for Trigger $= 10$ and Rf $= 0.0$, the local adaptive policy reduced 96% percent of infections with respect to the base case.

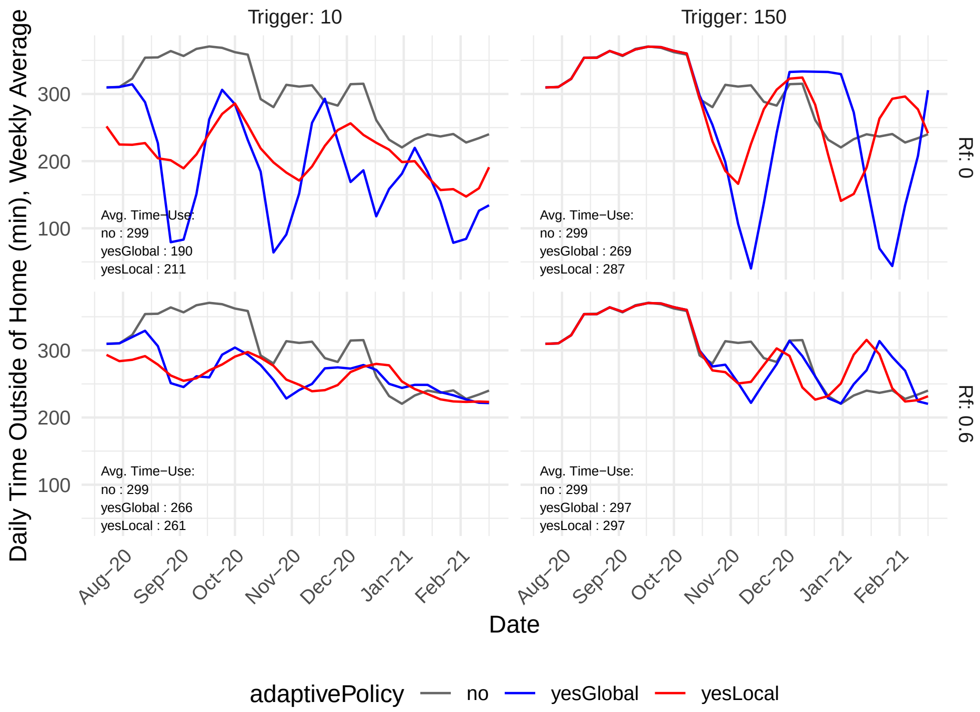

Figure 5.5: **Time-Use For Berlin—Adaptive Policies:** The plot shows the average minutes agents spend outside of their home every day (smoothed using weekly average). The text box shows the average over the entire study time frame (second wave). The plot colors and facets are designated to match Figure 5.4.

	Trigger						
Rf	10	25	50	75	100	125	150
0.0	29.3	14.1	10.9	8.1	7.3	4.6	4.3
0.2	23.3	9.9	8.4	6.0	4.6	3.5	2.5
0.4	18.5	7.0	5.6	3.7	1.9	2.2	1.6
0.6	12.7	4.0	2.8	2.2	2.1	1.3	0.8

Table 5.2: **Percent Decrease of Time-Use, Local Adaptive Policy vs. Base Case:** Shows percent reduction in average time agents spend outside of their homes. Calculated as follows: $(\mathrm{T}_{base} - \mathrm{T}_{local}) \div \mathrm{T}_{base} \times 100\%$, where T indicates the mean daily minutes that Berlin's residents spend outside of their homes, averaged over the study time-frame; *local* and *base* indicate the local adaptive policy and base case, respectively. For example, for Trigger = 10 and Rf = 0.0, Berlin's residents spent 29.3% less time outside of their homes than in the base case.

	Trigger						
Rf	10	25	50	75	100	125	150
0.0	8.2	16.0	19.3	24.2	25.6	36.8	35.4
0.2	10.3	22.9	24.6	31.9	39.0	46.1	59.4
0.4	12.8	31.7	35.6	48.9	91.3	67.2	78.1
0.6	18.4	53.7	66.8	74.3	66.8	92.0	114.8

Table 5.3: **Impact of Local Adaptive Policy with Respect to Base Case:** Average daily infections avoided by the local adaptive policy (with respect to base case) divided by the average number of out-of-home minutes lost through local adaptive policy (again, with respect to base case). For example, for Rf = 0.0 and Trigger = 10, 8.2 daily infections are avoided per sacrificed out-of-home minute. For the more lenient cases, the impact is much higher.

5.2.4 Incidence vs. Time-Use

It is difficult to weigh the benefit in incidence reduction against the cost of lost activities. It is most striking, however, to look at the most lenient scenario: there is only a 2 minute average daily reduction of out-of-home activities. Looking again at Figure 5.5, we see that in the most lenient parameters (top-right facet), the adaptive policies (red and blue plots) implemented earlier lockdowns than occurred in reality (black plot); this lead to less out-of-home time in October and November, as compared to the base case. However, the local adaptive policy balances this loss out by opening up around Christmas, leading to a spike in out-of-home time in the beginning of January. This indicates that if the Berlin lockdowns in the second wave were more strategically placed, 62,000 people would have avoided infection at almost no cost to out-of-home time.

In attempt to compare the costs to the benefits of each scenario in a more systematic way, a ratio was calculated for each scenario: average daily infections avoided by the adaptive policy divided by the average daily out-of-home time lost.

$$Impact = \frac{\bar{I}_{base} - \bar{I}_{policy}}{\bar{T}_{base} - \bar{T}_{policy}}$$

where $\bar{I}$ is the average daily infections and $\bar{T}$ is the average daily time spent outside of the home; *base* and *policy* denote the standard EpiSim simulation and the policy case utilizing adaptive restrictions. Table 5.3 shows this metric for every argument combination for the local adaptive policy. The greatest impact is made when Rf and Trigger are high: every minute that an individual stays at home every day avoids 114.8

new daily infections. For lower Triggers and Rfs, more time has to be sacrificed to reduce the infection curves.

As described above, the lenient scenario doesn't significantly add new lockdowns; rather, it institutes lockdowns earlier leading to a significantly dampened second wave without impacting the net time spent out of home. At lower triggers, you attack smaller waves, which doesn't have the same efficiency as attacking big waves. Lets imagine two towns, each with 10,000 residents: town A has 10 infectious people and town B has 150 infectious people. If no restrictions were applied, a lot more people would be infected in town B than in town A. If, instead, equal lockdowns had been imposed in both cases, a lot more people in town B would be avoided infection than in town A. As the time lost would be equal in both towns, the lockdown in town B would have more impact. However, town A would still have fewer total infections than town B.

The benefit/cost ratio of the adaptive policies is highest in the more lenient scenarios. The same trend can be observed for the global adaptive policy, which is shown in Table A.3 in the Appendix.

5.2.5 Local Adaptive Policy vs. Global Adaptive Policy

We have now summarized the general trends of the local adaptive policy, and explained that the global policy follows the same trend. The next step is to compare the two adaptive policies. Does the local adaptive policy have an advantage over the global one? If so, at what parameters (Trigger and Rf)?

We do this by comparing the impact of the two policies for each scenario. However, we didn't use the standard time-frame for this analysis. The results of this analysis depended heavily on what end-date was chosen for the for the time frame. At a fixed end date, the scenario with the local adaptive policy may be in an open phase while the global adaptive policy scenario might be in a restricted phase (as shown in upper right facet of Figure 5.4). This distorts the results when comparing the local and global policy.

Thus, we chose a variable time frame, depending on when scenarios are in a similar point of restriction policy. For each scenario: the start date of the variable time frame is when the first lockdown is lifted; the end date is when the second lockdown is lifted. Thus, the data analyzed for each scenario includes one open phase and one restricted phase. The impact for each scenario was then calculated for each unique time-frame. Finally, Table 5.4 compares the impacts of the local vs global adaptive policies in terms of percent difference; a positive values indicates that the local adaptive policy performs better than the global one.

	Trigger						
Rf	10	25	50	75	100	125	150
0.0	135.2	−58.9	74.6	15.2	57.3	19.5	57.1
0.2	48.5	−49.0	83.1	29.6	18.3	−50.2	−49.2
0.4	15.0	−39.3	13.4	−5.4	27.6	−52.3	−79.8
0.6	54.1	−37.2	24.3	12.9	0.0	−45.7	−54.9

Table 5.4: **Impact of Local Adaptive Policy vs. Impact of Global Adaptive Policy, Percent Change:** Calculated as follows: $(Impact_{local} - Impact_{global}) \div Impact_{global} \times 100\%$. Positive values indicate that local adaptive policy performs better than global adaptive policy. The Impact is calculated using a different time-frame for each scenario; the start and end dates of the variable time frames are set such that the time-frame encompass one restricted phase and one open phase.

The negative values in Table 5.4 indicate that the local adaptive policy does not always perform better than the global one. However, a relatively robust trend can be identified: when the Trigger is lower—between 10 and 100— the local adaptive policy outperforms the global policy.[8] At a Trigger of 125 or 150, the global policy performs better than the local one. It is reasonable that at a lower trigger, the local policies are advantageous; a pinpointed restriction will keep the infections in check without imposing a significant time loss in all boroughs.

When Rf=0.0, the local policy is advantageous (except for Trigger = 25). For the higher Rf's, the trend is less clear. The local policy is most advantageous in the most stringent scenario, where Trigger=10 and Rf=0.0; the local adaptive policy's impact has a percent change of 135% as compared to the global one.

5.3 Discussion

This chapter has been concerned with examining the effects of instituting local lockdowns in Berlin's boroughs. Case Study B1 aimed to demonstrate the effects of a local lockdown for the borough in lockdown and the non-restricted boroughs. Two types of lockdowns were tested: A) restricting resident's of Mitte from conducting leisure activities and B) barring all Berliners from conducting leisure activities in Mitte. In both cases, all boroughs had a reduction their incidences; however, the incidence of the restricted borough was most reduced. This demonstrated how interconnected Berlin is; restriction policies in one region will impact the infection dynamics of surrounding regions.

The impact of Mitte's lockdown on the other boroughs is not constant. We can notice

[8]Notable exception: when Trigger is 25.

that two outer boroughs—Lichtenberg and Marzahn-Hellersdorf—are least impacted by either restriction of Mitte. This could be explained by their geographical distances from Mitte, which might reduce the mobility between them. The counter example is Spandau, on the western edge of Berlin, which does, in fact, have a significant reductions in infections. Further work could explore the reasons why some boroughs are impacted more than others. It would be interesting to study the mobility patterns between boroughs and whether they can help us understand divergent infection dynamics. What locations serve as significant meeting points for residents of different boroughs? Which boroughs originate more leisure activities than they attract (residents go to other boroughs to have fun), or vice versa?

The red zones in the no-COVID proposal [10] both curtail people from entering a restricted area and from leaving it. The functionality developed as part of this thesis can only restrict either direction of movement separately. To better evaluate the no-COVID proposal, a restriction of bi-directional movement should be implemented.

Case Study B2 examines the local adaptive policy. We saw that the lockdowns patterns for boroughs depended strongly on the two parameters that were varied between scenarios: Trigger and Rf. We also saw that the more stringent policies reduce the infections more drastically, while also reducing the time people spend outside of their homes. When looking at the cost-benefit of these two factors, we see that the more lenient policy can remove many infections without a significant time loss. The more stringent the scenario, the fewer infections are avoided per minute lost due to the restriction. It is unclear how to compare the value of 1 minute per day lost to the value of avoiding one infection; this is a task for policy makers. Making policy recommendations are out of the scope of this thesis. The benefits of the lenient adaptive policies are abundantly clear: 62,000 fewer people would have been infected in the second wave, while the reduction in time spent outside of the home would have been minimal.

The local adaptive policy, which was developed as part of this thesis, outperforms the global one when the Trigger was low (especially when the Rf is also 0.0). Thus, if policy makers decide that the benefits of a stringent adaptive policy outweigh the costs, then a local adaptive policy would be an promising option.

The parameters at which the local adaptive policy outperforms the global policy—a low trigger and low Rf—approach the recommendations of the no-COVID proposal [10]. Baumann et al. [10] argue for an even lower Trigger: a region should be restricted whenever there are any local cases that aren't contained by quarantines or isolation policies.

The results of the local adaptive policies showed many cycles of imposing and lifting restrictions. This is a dynamic that the no-COVID proposal was designed to avoid. An important element of the no-COVID proposal [10] is not explored in Case Study B2: a robust test-trace-isolate (TTI) program to keep the incidence low. While EpiSim implements TTI functionality, the capacity of the TTI is not increased as part of this thesis. Future work should examine if increasing the capacity of a TTI program could avoid the cycles of lockdowns. This could improve the evaluation of the no-COVID proposal.

It is insightful to see that as Delta VOC became more prevalent, the local restrictions weren't able to suppress incidences as well as with the wild-type. This observation fits with a point made in Chapter 1: New Zealand abandoned the goal of complete eradication of the virus after suffering a wave of Delta infections [11]. Future work should examine the effect of localized restrictions in later time-frames, where VOCs gain prevalence.

It is unclear how tenable borough-based lockdowns would be in Berlin. Berlin's boroughs are the not self-contained; there is a lot of travel between boroughs. Thus, the enforcement of travel restrictions between boroughs would be difficult. In general, I think that larger study areas would be more beneficial to test this theory. In the Ruhr-area of Germany, for instance, there are many closely situated cities. I think it would be useful to explore how local restrictions of separate cities would work; since they are more self-contained, the incidence of one region would be less dependent on another. I believe that a local adaptive policy would perform better than the global one in such a scenario.

This chapter has shown that pinpointed lockdowns—local adaptive policies, in particular—are a promising tool to mitigate the effects of COVID-19.

Chapter 6

Conclusion

The motivations of this thesis was to A) simulate how regional characteristics influence regional infection dynamics and B) to evaluate the merits of pinpointed lockdowns. The contributions of this thesis are two-fold: A) the added functionality to EpiSim, which allows localized simulations and B) a set of case studies to explore the motivations mentioned above. These two contributions will be summarized in the following two sections.

6.1 EpiSim

EpiSim [8] is an epidemiological simulation model developed to model infection dynamics of the COVID-19 pandemic. It was implemented to prognosticate the effects of various mitigation strategies, including lockdowns, vaccination campaigns and contact tracing. Chapter 2 gave a detailed explanation of the data that EpiSim requires and how the software functions.

A significant contribution of this thesis is the added functionality to EpiSim, which allows infection parameters to vary geographically and localized restrictions to be imposed. The standard EpiSim model applies a daily activity reduction to the pre-coronavirus activity trajectories of the residents of the study area. This thesis extends EpiSim to have separate activity reductions for different regions instead of a single global one.

On this basis, a local adaptive policy was designed to dynamically restrict or open up an area based on the incidence of that area. The functional contributions were described in detail in Chapter 3.

6.2 Case Studies

The functional contributions allowed a set of research questions to be posed; case studies were developed to respond to those questions. The manner in which the case studies were set up was described in Chapter 3. Chapters 4 and 5 analyze the results of the case studies. The findings are also summarized below.

6.2.1 Research Focus A: Local Infection Dynamics

The first motivation was to improve EpiSim's ability to capture the infection dynamics of individual geographic regions within the study area. This led to the following two research questions:

RQ A1: *Does the addition of localized activity reductions improve EpiSim's ability to capture local infection dynamics?*

RQ A2: *Does the inclusion of localized contact intensity for home activities (based on varying home-size) improve EpiSim's ability to capture local infection dynamics?*

The results, as shown in Chapter 4, showed that varying activity reductions and contact intensity didn't cause significant changes in the incidences of most boroughs. When there was a difference, it didn't necessarily close the gap to the reported cases from the RKI. This indicates that the regional differences in activity reduction and home size are not the best factors for explaining diverging incidence curves between Berlin boroughs. This, itself, is an notable result: it means, for instance, that diverging incidence curves cannot be explained by the residents of one Berlin borough of behaving more "irresponsibly" than residents of another. Further work should explore whether these factors are more relevant for larger study regions. Additionally, further work could explore the addition of other regionally diverging factors into the model.

6.2.2 Research Focus B: Local Lockdowns

The second motivation was to explore the merits of imposing localized lockdowns, wherein some regions are restricted while others remain open. This research focus was inspired by the no-COVID proposal [10], which attempts to eradicate infections through local lockdowns. The following research questions were posed:

RQ B1: *How does a local lockdown affect the infection dynamics of the restricted region and the un-restricted regions?*

RQ B2: *How does a local adaptive policy affect incidences and time uses, compared to a global adaptive policy? How do the parameters of the adaptive policy affect its benefit?*

Case Study B1 showed that a lockdown in one Berlin borough not only dampens the incidence of the restricted borough, but also decreases the incidences of other boroughs. Case Study B2 explored the effects of a local adaptive policy for the 12 boroughs that make up Berlin. When the policy was configured stringently, the incidences plummeted; the downside of this policy was that people spent less time outside of their homes. The most lenient of the adaptive policies was able to also prevent 62,000 people from contracting SARS-CoV-2 during the second wave, while the reduction in time outside of the home was minimal.

The local policy, developed by this thesis, proved to be a promising restriction policy at certain configurations. It performed better than the global adaptive policy when the incidence threshold for a lockdown (Trigger) was low.

6.3 Outlook

As this is being written, one week before thesis submission, a new variant of concern (VOC) is spreading around the world: B.1.1.529 or "Omicron" [40]. While many properties of Omicron are still unknown, many countries have imposed travel restrictions to South Africa, where the VOC was first identified. Japan and Israel have closed their borders to stave off the variant [41]. The EpiSim team is currently working furiously to model the spread of Omicron in Cologne, Germany, and testing the effectiveness of new lockdowns.

Modeling the spread of COVID-19 with epidemiological simulations will continue to be vital undertaking in the coming months and years. While SARS-CoV-2 continues to shape our daily reality, reducing the amount of people stricken by this dangerous disease while also regaining normalcy in our mobility and interactions remains a monumental task. This thesis introduced functionality to EpiSim, which allows the simulation of pinpointed restrictions. The case studies showed that a local lockdown regime in Berlin could be a promising tool in fighting the pandemic.

Bibliography

[1] S.-L. Wee and V. Wang, "China Grapples With Mystery Pneumonia-Like Illness," *The New York Times*, Jan. 2020. [Online]. Available: https://www.nytimes.com/2020/01/06/world/asia/china-SARS-pneumonialike.html [Accessed: 2021-10-19].

[2] Y.-C. Liu, R.-L. Kuo, and S.-R. Shih, "COVID-19: The first documented coronavirus pandemic in history," *Biomedical Journal*, vol. 43, no. 4, pp. 328–333, Aug. 2020. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7199674/ [Accessed: 2021-10-19].

[3] T. Adhanom, "WHO Director-General's opening remarks at the media briefing on COVID-19 - 11 March 2020," Mar. 2020. [Online]. Available: https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---11-march-2020 [Accessed: 2021-10-19].

[4] B. Sen-Crowe, M. McKenney, and A. Elkbuli, "Disparities in global COVID-19 vaccination rates & allocation of resources to countries in need," *Annals of Medicine and Surgery*, vol. 68, p. 102620, Aug. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2049080121005707 [Accessed: 2021-10-20].

[5] N. Perra, "Non-pharmaceutical interventions during the COVID-19 pandemic: A review," *Physics Reports*, vol. 913, pp. 1–52, May 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0370157321000624 [Accessed: 2021-10-19].

[6] J. Duhon, N. Bragazzi, and J. D. Kong, "The impact of non-pharmaceutical interventions, demographic, social, and climatic factors on the initial growth rate of COVID-19: A cross-country study," *Science of The Total Environment*, vol.

760, p. 144325, Mar. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0048969720378566 [Accessed: 2021-10-19].

[7] H. Ritchie, E. Mathieu, L. Rodés-Guirao, C. Appel, C. Giattino, E. Ortiz-Ospina, J. Hasell, B. Macdonald, D. Beltekian, and M. Roser, "Coronavirus Pandemic (COVID-19)," Mar. 2020. [Online]. Available: https://ourworldindata.org/policy-responses-covid [Accessed: 2021-10-19].

[8] S. A. Müller, M. Balmer, W. Charlton, R. Ewert, A. Neumann, C. Rakow, T. Schlenther, and K. Nagel, "A realistic agent-based simulation model for COVID-19 based on a traffic simulation and mobile phone data," *arXiv:2011.11453 [physics]*, Nov. 2020. [Online]. Available: http://arxiv.org/abs/2011.11453 [Accessed: 2021-08-09].

[9] C. Endt, L. Fischer, C. Grefe-Huge, M. Klack, and J. Tröger, "Soziale Ungleichheit und Corona: Das sind die Corona-Hotspots in deutschen Großstädten," *Die Zeit*, Jun. 2021. [Online]. Available: https://www.zeit.de/wissen/2021-05/soziale-ungleichheit-corona-infektionen-aermere-stadtteile-datenanalyse-soziale-brennpunkte [Accessed: 2021-11-29].

[10] M. Baumann, M. Beier, M. Brinkmann, H. Bude, C. Fuest, M. Hallek, I. Kickbusch, M. Mayer, M. Meyer-Hermann, A. Peichl, E. Rosert, and M. Schneider, "A proactive approach to fight SARS-CoV-2 in Germany and Europe," Jan. 2021. [Online]. Available: https://nocovid-europe.eu/assets/doc/nocovid_framework.pdf [Accessed: 2021-07-12].

[11] N. Frost, "Battling Delta, New Zealand Abandons Its Zero-Covid Ambitions," *The New York Times*, Oct. 2021. [Online]. Available: https://www.nytimes.com/2021/10/04/world/australia/new-zealand-covid-zero.html [Accessed: 2021-11-29].

[12] V. Wang, "Why China Is the World's Last 'Zero Covid' Holdout," *The New York Times*, Oct. 2021. [Online]. Available: https://www.nytimes.com/2021/10/27/world/asia/china-zero-covid-virus.html [Accessed: 2021-11-29].

[13] R. Mullner, "Epidemiology | medicine," Aug. 2014. [Online]. Available: https://www.britannica.com/science/epidemiology [Accessed: 2021-10-20].

[14] L. Kou, X. Wang, Y. Li, X. Guo, and H. Zhang, "A multi-scale agent-based model of infectious disease transmission to assess the impact of vaccination

and non-pharmaceutical interventions: The COVID-19 case," *Journal of Safety Science and Resilience*, vol. 2, no. 4, pp. 199–207, Dec. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666449621000335 [Accessed: 2021-10-20].

[15] E. Hunter, B. M. Namee, and J. Kelleher, "A Comparison of Agent-Based Models and Equation Based Models for Infectious Disease Epidemiology," *CEUR workshop proceedings*, vol. 2259, pp. 33–44, Jan. 2018.

[16] G. N. Gilbert and K. G. Troitzsch, "Multi-agent models," in *Simulation for the social scientist*, 2nd ed. Maidenhead, England ; New York, NY: Open University Press, 2005, pp. 172–198, chapter 8.

[17] S. Merler, M. Ajelli, L. Fumanelli, M. F. C. Gomes, A. P. y. Piontti, L. Rossi, D. L. Chao, I. M. Longini, M. E. Halloran, and A. Vespignani, "Spatiotemporal spread of the 2014 outbreak of Ebola virus disease in Liberia and the effectiveness of non-pharmaceutical interventions: a computational modelling analysis," *The Lancet Infectious Diseases*, vol. 15, no. 2, pp. 204–211, Feb. 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1473309914710746 [Accessed: 2021-10-19].

[18] S. Zhou, S. Zhou, Z. Zheng, and J. Lu, "Optimizing Spatial Allocation of COVID-19 Vaccine by Agent-Based Spatiotemporal Simulations," *GeoHealth*, vol. 5, no. 6, p. e2021GH000427, 2021. [Online]. Available: http://onlinelibrary.wiley.com/doi/abs/10.1029/2021GH000427 [Accessed: 2021-10-19].

[19] A. Horni, K. Nagel, and K. W. Axhausen, *The Multi-Agent Transport Simulation MATSim.* Ubiquity Press, Aug. 2016. [Online]. Available: https://www.ubiquitypress.com/site/books/e/10.5334/baw/ [Accessed: 2021-10-21].

[20] A. Agarwal, D. Ziemke, and K. Nagel, "Bicycle superhighway: An environmentally sustainable policy for urban transport," *Transportation Research Part A: Policy and Practice*, vol. 137, pp. 519–540, Jul. 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S096585641731162X [Accessed: 2021-10-21].

[21] I. Kaddoura, "Marginal Congestion Cost Pricing in a Multi-agent Simulation: Investigation of the Greater Berlin Area," *Journal of Transport Economics and Policy*, vol. 49, no. 4, pp. 560–578, 2015, publisher: University of

Bath. [Online]. Available: http://www.jstor.org/stable/jtranseconpoli.49.4.0560 [Accessed: 2021-12-12].

[22] J. Bischoff, M. Maciejewski, and K. Nagel, "City-wide shared taxis: A simulation study in Berlin," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Oct. 2017, pp. 275–280, iSSN: 2153-0017. [Online]. Available: https://doi.org/10.1109/ITSC.2017.8317926 [Accessed: 2021-10-21].

[23] M. L. Barreto, M. G. Teixeira, and E. H. Carmo, "Infectious diseases epidemiology," *Journal of Epidemiology and Community Health*, vol. 60, no. 3, pp. 192–195, Mar. 2006. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2465549/ [Accessed: 2021-12-12].

[24] T. Smieszek, "A mechanistic model of infection: why duration and intensity of contacts should be included in models of disease spread," *Theoretical Biology and Medical Modelling*, vol. 6, no. 1, p. 25, Dec. 2009. [Online]. Available: https://tbiomed.biomedcentral.com/articles/10.1186/1742-4682-6-25 [Accessed: 2021-12-09].

[25] T. Smieszek, M. Balmer, J. Hattendorf, K. W. Axhausen, J. Zinsstag, and R. W. Scholz, "Reconstructing the 2003/2004 H3N2 influenza epidemic in Switzerland with a spatially explicit, individual-based model," *BMC infectious diseases*, vol. 11, p. 115, May 2011.

[26] S. A. Müller, W. Charlton, R. Ewert, C. Rakow, T. Schlenther, and K. Nagel, "MODUS-COVID Vorhersage vom 8.4.2020," Technische Universität Berlin, Report, Apr. 2020. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/11125 [Accessed: 2021-12-09].

[27] S. A. Müller, W. Charlton, R. Ewert, C. Rakow, T. Schlenther, and K. Nagel, "MODUS-COVID Vorhersage vom 24.4.2020," Technische Universität Berlin, Report, Apr. 2020. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/11126 [Accessed: 2021-12-09].

[28] S. A. Müller, W. Charlton, R. Ewert, C. Rakow, T. Schlenther, and K. Nagel, "MODUS-COVID Vorhersage vom 7.5.2020," Technische Universität Berlin, Report, May 2020. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/11127 [Accessed: 2021-12-09].

[29] S. A. Müller, W. Charlton, N. D. Conrad, R. Ewert, C. Rakow, H. Wulkow, T. Conrad, K. Nagel, and C. Schütte, "MODUS-COVID Bericht vom 18.12.2020," Technische Universität Berlin, Report, Dec. 2020. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/12235 [Accessed: 2021-12-09].

[30] S. A. Müller, W. Charlton, N. D. Conrad, R. Ewert, C. Rakow, H. Wulkow, T. Conrad, K. Nagel, and C. Schütte, "MODUS-COVID Bericht vom 15.01.2021," Technische Universität Berlin, Report, Jan. 2021. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/12447 [Accessed: 2021-12-09].

[31] S. A. Müller, W. Charlton, N. D. Conrad, R. Ewert, D. Jefferies, C. Rakow, H. Wulkow, T. Conrad, C. Schütte, and K. Nagel, "MODUS-COVID Bericht vom 19.03.2021," Technische Universität Berlin, Report, Mar. 2021. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/12878 [Accessed: 2021-12-09].

[32] S. A. Müller, W. Charlton, N. D. Conrad, R. Ewert, S. Paltra, C. Rakow, H. Wulkow, T. Conrad, C. Schütte, and K. Nagel, "MODUS-COVID Bericht vom 22.10.2021," Technische Universität Berlin, Report, Oct. 2021. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/13734 [Accessed: 2021-12-09].

[33] S. A. Müller, W. Charlton, N. D. Conrad, R. Ewert, S. Paltra, C. Rakow, H. Wulkow, T. Conrad, C. Schütte, and K. Nagel, "MODUS-COVID Bericht vom 24.09.2021," Technische Universität Berlin, Report, Sep. 2021. [Online]. Available: https://depositonce.tu-berlin.de/handle/11303/13639 [Accessed: 2021-12-09].

[34] BVerfG, "Beschluss des Ersten Senats vom 19. November 2021," Nov. 2021, 1 BvR 781/21, Rn. 1-306. [Online]. Available: https://www.bundesverfassungsgericht.de/e/rs20211119_1bvr078121.html [Accessed: 2021-12-11].

[35] TUBS - Wikimedia Commons, "Map of administrative divisions of Berlin," Aug. 2012. [Online]. Available: https://commons.wikimedia.org/wiki/File:Berlin,_administrative_divisions_(%2Bdistricts_%2Bboroughs_-pop)_-_de_-_colored_(less_colors).svg [Accessed: 2021-12-01].

[36] LAGeSo, "COVID-19 in Berlin, Verteilung in den Bezirken - Gesamtübersicht | Offene Daten Berlin," Dec. 2021. [Online]. Available: https://daten.berlin.de/datensaetze/covid-19-berlin-verteilung-den-bezirken-gesamt%C3%BCbersicht [Accessed: 2021-10-21].

[37] Amt für Statistik Berlin-Brandenburg, “Einwohnerbestand Berlin.” [Online]. Available: https://www.statistik-berlin-brandenburg.de/kommunalstatistik/einwohnerbestand-berlin [Accessed: 2021-12-01].

[38] BBSR, “Laufende Stadtbeobachtung - Raumabgrenzungen.” [Online]. Available: https://www.bbsr.bund.de/BBSR/DE/forschung/raumbeobachtung/Raumabgrenzungen/deutschland/regionen/Grossstadtregionen/Grossstadtregionen.html [Accessed: 2021-12-01].

[39] Senatsverwaltung für Stadtentwicklung und Wohnen Berlin, “Monitoring Soziale Stadtentwicklung 2019,” 2019. [Online]. Available: https://www.stadtentwicklung.berlin.de/planen/basisdaten_stadtentwicklung/monitoring/download/2019/tab/4.1.KontextInd_Anteile_PLR_MSS2019.pdf [Accessed: 2021-09-14].

[40] WHO, “Classification of Omicron (B.1.1.529): SARS-CoV-2 Variant of Concern.” [Online]. Available: https://www.who.int/news/item/26-11-2021-classification-of-omicron-(b.1.1.529)-sars-cov-2-variant-of-concern [Accessed: 2021-12-09].

[41] V. Harrison, “Omicron variant: G7 to hold emergency Covid meeting as Japan closes its borders,” *The Guardian*, Nov. 2021. [Online]. Available: https://www.theguardian.com/world/2021/nov/29/omicron-variant-g7-to-hold-emergency-covid-meeting-as-japan-closes-its-borders [Accessed: 2021-12-08].

Appendix A

Global Adaptive Policy Tables

The following tables are auxiliary to Chapter 5. While Case Study B2 presented tables for the local adaptive policy, the following tables show the corresponding results for the global adaptive policy.

	Trigger						
Rf	10	25	50	75	100	125	150
0.0	94.8	89.8	83.8	78.6	73.2	67.1	59.7
0.2	94.6	89.5	81.9	76.8	72.3	65.5	58.0
0.4	94.0	88.4	78.5	72.2	66.2	58.6	52.9
0.6	92.4	85.1	73.0	61.9	52.4	43.5	34.3

Table A.1: **Percent Decrease in Infections, Global Adaptive Policy vs. Base Case:** Percent reduction calculated as follows: $(\mathrm{Inf}_{base} - \mathrm{Inf}_{global}) \div \mathrm{Inf}_{base} \times 100\%$, where Inf indicates the total number of infections in the study time-frame; *global* and *base* indicate the global adaptive policy and base case respectively.

Rf	Trigger 10	25	50	75	100	125	150
0.0	36.0	23.1	10.9	10.5	10.4	10.2	10.6
0.2	26.3	16.0	7.7	5.9	5.3	5.4	6.3
0.4	18.7	9.5	5.4	2.7	1.8	1.6	1.4
0.6	11.1	4.5	2.8	2.1	2.0	1.4	0.5

Table A.2: **Percent Decrease of Time-Use, Global Adaptive Policy vs. Base Case:** Shows percent reduction in average time agents spend outside of their homes. Calculated as follows: $(\mathrm{T}_{base} - \mathrm{T}_{global}) \div \mathrm{T}_{base} \times 100\%$, where T indicates the mean daily minutes that Berlin's residents spend outside of home, averaged over the study time-frame; *global* and *base* indicate the global adaptive policy and base case respectively.

Rf	Trigger 10	25	50	75	100	125	150
0.0	6.6	9.7	19.3	18.6	17.7	16.6	14.2
0.2	9.1	14.0	27.0	32.4	34.7	30.1	23.1
0.4	12.6	23.1	36.6	66.9	91.4	93.6	97.1
0.6	20.7	47.5	65.0	73.1	67.2	79.9	157.8

Table A.3: **Impact of Global Adaptive Policy with Respect to Base Case:** Average daily infections avoided by the global adaptive policy (with respect to base case) divided by the average number of out-of-home minutes lost through global adaptive policy (again, with respect to base case).