A Novel Queue-based Stochastic Epidemic Model with Adaptive Stabilising Control

Edilson F. Arruda[a,*], Rodrigo e A. Alexandre[b], Marcelo D. Fragoso[c], João B. R. do Val[d], Sinnu S. Thomas[e]

[a]*Department of Decision Analytics and Risk, Southampton Business School, University of Southampton, 12 University Rd, Southampton SO17 1BJ, UK*
[b]*Alberto Luiz Coimbra Institute-Graduate School and Research in Engineering, Federal University of Rio de Janeiro. CP 68507, Rio de Janeiro 21941-972, Brasil*
[c]*National Laboratory for Scientific Computation, Av. Gettúlio Vargas 333, Quitandinha, Petrópolis RJ 25651-075, Brasil*
[d]*School of Electrical Engineering, University of Campinas, Av. Albert Einstein 400, Cidade Universitária, Campinas SP 13083-852, Brasil*
[e]*School of Computer Science and Engineering, Digital University Kerala, Technocity, Mangalapuram Thonnakkal PO Thiruvananthapuram, Kerala-695317, India*

Abstract

The main objective of this paper is to propose a novel SEIR stochastic epidemic model. A distinguishing feature of this new model is that it allows us to consider a setup under general *latency and infectious period distributions*. To some extent, queuing systems with infinitely many servers and a Markov chain *with time-varying transition rate* comprise the very technical underpinning of the paper. Although more general, the Markov chain is as tractable as previous models for exponentially distributed latency and infection periods. It is also significantly more straightforward and tractable than semi-Markov models with a similar level of generality. Based on stochastic stability, we derive a sufficient condition for a shrinking epidemic regarding the queuing system's occupation rate that drives the dynamics. Relying on this condition, we propose a class of ad-hoc stabilising mitigation strategies that seek to keep a balanced occupation rate after a prescribed mitigation-free period. We validate the approach in the light of the COVID-19 epidemic in England and in the state of Amazonas, Brazil, and assess the effect of different stabilising strategies in the latter setting. Results suggest that the proposed approach can curb the epidemic with various occupation rate levels if the mitigation is timely.

Keywords: Stochastic epidemic models, Markov processes, queuing theory, stabilising control.

1. Introduction

Classical epidemic models [e.g., 1, 2, 3] are powerful tools to understand the spread of diseases and support public health policies. These models often divide the population into compartments. For example, a classical SIR model includes *susceptible* (S), *infected* (I) and *recovered* (R) individuals. The SEIR model is a generalisation whereby E represents the *exposed* individuals who have been contaminated but are not yet infectious due to the disease's *latency* period. Classical compartmental models were extensively applied to describe the evolution of the COVID-19 pandemic, differentiated only by the number and types of compartments. While Liu et al. [4] utilised a classical SIR model, McQuade et al. [5] applied a SEIR model that differentiated symptomatic and asymptomatic infected individuals. Prasad [6] utilised a QSIR model, where Q represents the quarantined population. Mandal et al. [7] introduced an analogous QSEIR model, whereas Upadhyay et al. [8] applied an age-structured QSIR model that differentiated the disease evolution across different age groups. Although nearly ubiquitous in epidemic modelling, classical compartmental models are deterministic and do not capture the underlying uncertainties or the random duration of the disease cycle [9]. Stochastic epidemic models [10, 11] were designed to better capture some of these uncertainties and provide more practical support for decision-making.

Network models are a variety of stochastic epidemic models that consider the population density and mobility [e.g., 12, 13, 14]. However, they are computationally intensive, limiting their applicability to low-dimensional and sometimes artificial problems. Perhaps due to its simplicity, SIR is arguably the most utilised class of stochastic epidemic models. Trapman and Bootsma [15] used this framework to demonstrate the advantage of an M/G/1 queuing

*Corresponding author. Tel.: +44 023 8059 7677
Email addresses: E.F.Arruda@southampton.ac.uk (Edilson F. Arruda), alvim.rodrigo@yahoo.com.br (Rodrigo e A. Alexandre), frag@lncc.br (Marcelo D. Fragoso), jbosco@fee.unicamp.br (João B. R. do Val), sinnu.thomas@duk.ac.in (Sinnu S. Thomas)

model to estimate the size of an epidemic at the time of detection. Sometime later, classical M/M/S queues were utilised to estimate the whole outbreak of the Ebola virus [16]. More recently, Barraza et al. [17] employed pure Birth processes to fit data from the initial stages of the COVID-19 epidemic. Underpinning the analysis is the theory of Markov processes [18], used to model the transition of individuals among the SIR populations.

Markov models provide a powerful analytical framework for SIR models, allowing, for example, the treatment of non-homogeneous populations [19]. However, one of the limitations is that the infectious period's duration is assumed exponential, thus narrowing the technique's applicability [20, 21]. A possible alternative is to develop more complex block-structured Markov chains that mimic certain non-exponential infection times [22, 23]. Albeit limited, the added flexibility comes at the price of less interpretable and tractable models, with limited applicability to small populations.

A thorough treatment of general infection times demands a class of semi-Markov processes known as piecewise-deterministic processes (PDP) [24]. Clancy [20] uses such a class and martingale theory to derive the distribution of the number of infected individuals throughout the pandemic under generally distributed infection times. Later, Gómez-Corral and López-García [21] utilised a similar model to obtain the distribution of the number of secondary cases due to an infected individual. PDP models require memory of the disease progression of all currently infected individuals, which impacts the model's analytical and computational tractability and limits its use. Hence, these models are typically applied to find the structural properties of a spread rather than to model and control an outbreak.

More appropriate for epidemics with non-negligible latency period, such as COVID-19 [25], the stochastic SEIR models in the literature [e.g., 26, 27, 28] are limited to exponentially distributed latency and infectious periods and also focus on deriving structural properties of the epidemic spread rather than describing the epidemic spread and underpinning control strategies. Therefore, inspired by the recent COVID-19 outbreak and considering the lack of a realistic and computationally tractable epidemic model that incorporates control and appropriately describes the inherent uncertainty, this works develops a novel stochastic SEIR model that is parsimonious enough to underpin the design of control strategies, while keeping

a level of generality that allows for arbitrary latency and infection period distributions.

This paper proposes an innovative $M_t/G/\infty$ model to describe the epidemic's stochastic behaviour during the outbreak. It is the first queuing model of the kind, as the $M/G/1$ queue in [15] is a specialised model to evaluate the size of the epidemic solely at the time of detection and therefore does not cover the epidemic evolution. The approach hinges on two unique novelties. Firstly, we describe the input process as a time-varying Poisson process, which enables us to describe the variation of infection rates as a function of the system's dynamics. Secondly, we select a queuing model with infinitely many servers by realising that there are no limits to the number of new infections and considering that conditions progress in parallel.

To our knowledge, this is the first work to introduce a time-varying Markov chain capable of emulating a stochastic epidemic model with general latency and infection times. This innovation relies on the results of Eick et al. [29], which demonstrate that the output process of $M_t/G/\infty$ queues comprises a series of time-indexed Poisson variables. Albeit as tractable as the models that assume exponential (memoryless) infection and latency periods [e.g., 26, 27, 28], the proposed framework is more general than the complex block-structured models [e.g., 22, 23] as it does not impose any assumption on the infection and latency period distributions.

Even if our approach's generality level is partially matched by previous PDP models within the more restricted SIR framework [20, 21], these require memory of the disease progression of all individuals in the *infected* population. Of course, this is infeasible for all but tiny population sizes and limits such models' applicability to support decision-making. In contrast, the proposed framework considers general latency and infection periods by keeping track solely of the new expositions within a complete disease cycle: from catching the virus to entering the removed population. Table 1 below briefly summarises the properties of existing approaches and the characteristics of our method. Please notice that the computational aspects, mentioned before, are not included.

In addition to the methodological innovations, we propose a novel strategy to curb the epidemic based on the classical occupation rate of the $M_t/G/\infty$ exposition-to-removal queue.

Table 1: Summary of the characteristics of existing models. The mnemonic **O** represents a characteristic that occasionally appears in the literature.

			SIR	SEIR	Stochastic SIR	Stochastic SEIR	Stochastic network	Queuing models	Present paper
Latency period	Deterministic			✓					
	Stochastic	Exponential				O			
		General				O			✓
Recovery period	Deterministic		✓	✓			O		
	Stochastic	Exponential			O	O	O	O	
		General			O	O		O	✓
Analysis of spread	Structural properties only				✓	✓	✓		
	Only epidemic outset							✓	
	Full epidemic simulation		✓						✓
Control capability			O	O	No	No	O		✓

This measure provides a dynamic estimate of the epidemic's short-term growth. The strategy relies on mitigating actions specially tailored to maintain the system's occupation rate ρ as close as possible to a prescribed level $0 < \bar{\rho} < 1$ at all times, thus ensuring a shrinking epidemic. We validate the model's predictive capability in light of the COVID-19 epidemic in England. We also conducted experiments using data from the COVID-19 dynamics in Amazonas, Brazil. The experiments illustrate the approach's accuracy and utility while providing invaluable insights into the system's response to the proposed mitigation policies. The results demonstrate that the epidemic can be conquered by maintaining appropriate occupation-level targets as long as the mitigation is not excessively delayed.

Although parameter fitting and data-driven analysis are beyond the scope of this paper, it is worth mentioning that insufficient testing and asymptomatic infections may compromise parameter fitting from epidemic data, regardless of the model's level of detail. Consider, for instance, the state of Amazonas, Brazil: a largely unmitigated epidemic event whereby a serological study estimated an accumulated incidence of 76% [30, 31] in October 2020. In contrast, official epidemic reports featured a cumulative incidence of about 3.9% up to the end of that month [32]. These figures suggest reconciling conflicting reports beyond ensuring adequate parameter fitting [33]. One advantage of the proposed model is that it can directly use the latency and infectious period distributions from medical studies, such as [25, 34].

Thus, only the infection rate β can be inferred from epidemic data. For a comprehensive discussion of the relationship between the reproductive number, the growth rate, and the shape of the disease cycle distribution, refer to [9].

To summarise, we present a novel SEIR-based queuing model to describe epidemic evolution and underpin epidemic control policies. The model is simpler and more computationally parsimonious than competing stochastic epidemic models, such as network models. In contrast to classical epidemic models, our approach is stochastic, dynamic, and adaptive; furthermore, it allows for general latency and infection periods. In contrast to existing stochastic models, our approach allows the decision maker to adjust the level of control according to a prescribed expected decrease in infection levels. It is also less susceptible to poor data from insufficient testing and asymptomatic infections than data-driven approaches.

This work is organised as follows. Section 2 introduces the mathematical formulation of the proposed SEIR model, starting with the proposed $M_t/G/\infty$ queues. We then use these queues' input and output processes to propose a continuous time Markov chain to describe the epidemic's evolution. In Section 3, we present a stochastic stability analysis that gives rise to a class of ad-hoc mitigation strategies to curb the epidemic. Section 4 features a set of experiments designed to illustrate the model's performance in real-world settings. The experiments also highlight the effectiveness of prescribed mitigation policies belonging to the class introduced in Section 3. Finally, Section 5 concludes the manuscript.

1.1. A Few Notations

We will use the following notation throughout this paper. Let $\mathbb{Z}_+$ be the usual set of non-negative integers and $\mathcal{N} = \{0, 1, 2, 3, ...N\} \subset \mathbb{Z}_+$, where N is a finite number. Consider $\Omega = \mathcal{N}^4$ and define the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. In addition, $\mathcal{E}(\cdot)$ stands for the mathematical expectation.

2. Mathematical formulation

This section proposes a stochastic dynamic formulation in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. To some extent, it is inspired by the classical deterministic SEIR epidemiological model

[1, 2, 3], depicted in Figure 1 and further explained below.

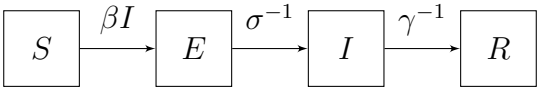

Figure 1: The classical SEIR model

The model comprises four compartments, namely: (*i*) susceptible, (*ii*) exposed, (*iii*) infected and (*iv*) removed. *Susceptible* individuals can be infected if they come in contact with infected individuals. *Exposed* individuals have been infected, but the disease is still latent, so they cannot spread the disease and do not manifest any symptoms. Once the latency period is over, exposed persons become *infected* and may present symptoms. Finally, infected individuals either die due to the disease or become immune to it. In both cases, they migrate to the *removed* population, which comprises individuals who can no longer contract or spread the disease.

In the deterministic model in Figure 1, susceptible individuals become ill at a rate $\beta > 0$ upon making contact with infected (infectious) individuals, thus resulting in an infection rate $\delta = \beta SI$. The individuals who acquire the condition immediately become exposed and initiate their latency period, which averages σ units of time. Upon completing the latency period, exposed individuals become infectious for an average of γ time units. Afterward, they enter the removed population due to acquired immunity or death. Table 2 below presents the parameters of the deterministic model:

Table 2: Parameters of SEIR dynamics.

Parameter	**Description**	**Unit**
β	Transmission rate	transmissions/encounter
σ	Latency period	days
γ	Recovery period	days

2.1. Stochastic formulation

Although the model in Figure 1 is invaluable to understand the underlying process, it is evident that the real-life dynamics are fundamentally stochastic. Albeit both the latency and the recovery periods are, indeed, stochastic variables, the deterministic SEIR model can only deal with the corresponding average rates. As expressed by Wallinga and Lipsitch [9], the reproductive number R_0 does not suffice to determine the growth rates as the epidemic evolves. They demonstrated that, for a given R_0, the generation interval distribution (latency period plus infectious period) determines the epidemic's growth rates. Therefore, considering the actual latency period and infectious period distributions is also vital for depicting the epidemic. See, for example, Backer et al. [25] and Verity et al. [34] for estimations of the latency and recovery periods at the early stages of the COVID-19 epidemic. This Section models the SEIR dynamics as a queue with infinite service capacity.

Queue 1 in Figure 2 represents an individual's trajectory from acquiring the disease and therefore becoming exposed to their removal of the system. Because the trajectories of different individuals are assumed independent, the system can be modelled as a $M_t/G/\infty$ queue, as there is no upper limit on the number of simultaneous infections [29]. The service time is the sum of two generally distributed random variables σ and γ, representing the latency and recovery periods, see Table 2; the latter corresponds to the length of the infectious period.

Essentially, the proposed model depicted in Figure 2 considers each infection as a service provision that includes a *latency* service that takes a random time (described by random variable σ) to be completed, followed by a random *infection* service time that takes γ units of time to be completed. Like in the deterministic model, the individual becomes infectious after the latency period and immune or removed after the infection period. However, the proposed stochastic model is more realistic because it allows for random latency and infectious periods and makes no assumption regarding their distributions, which can be promptly inferred from data and included inputted in the model.

We model the arrival rate as a Poisson random variable with a rate proportional to the maximum number of encounters between healthy and infected individuals, recalling that

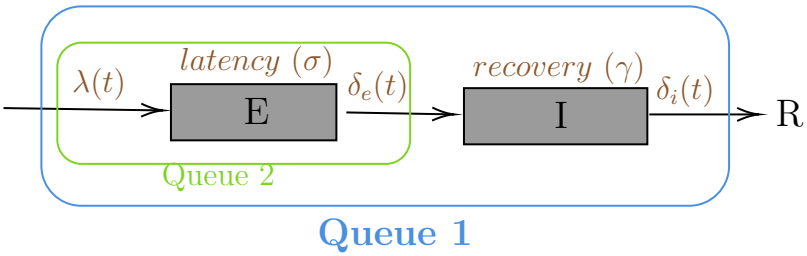

Figure 2: **Queue 1:** $M_t/G/\infty$ Queue from Exposed to Removed. **Queue 2:** $M_t/G/\infty$ Queue from Exposed to Infected

each new arrival represents a new contagion. The proposed model innovates by considering a time-varying arrival rate $\lambda(t) = \beta S(t)I(t)$ to cover the epidemic's evolution over time. Figure 2 shows that each departure from Queue 1 represents a decrease of one infected individual. Each arrival corresponds to a new exposition, i.e., a susceptible individual contracting the disease in its latent stage. Each compartment (E, I, and R) denotes the corresponding SEIR population.

Also, in Figure 2, **Queue 2** is a second $M_t/G/\infty$ queue that covers the *latency* period σ only. Each departure of this queue increases the infected population by one.

The queues in Figure 2 are inspired by the work of Eick et al. [29]. However, the proposed epidemic model is more general since the input rate $\lambda(t)$ is not solely a function of time. Instead, it depends on the number of infected and susceptible individuals at time $t \geq 0$. As previously stated, stochastic jumps in the number of infected individuals depend upon the departure processes of the two queues introduced in Figure 2.

Remark 1. *Deterministic models typically use differential equations to describe the dynamics of the system in Figure 1 [e.g., 35], considering both σ and γ as deterministic values. In contrast, the approach in Figure 2 regards a disease cycle as the sequential provision of two activities of random duration (latency and infection), which can be seamlessly represented by a queuing model [e.g., 36]. In the next section, we tailor the queuing parameters to model an epidemic spread appropriately.*

2.1.1. The arrival and departure processes

We start by modelling the arrival process $\lambda(t)$ shared by both queues, that represents the number of new infections over time. At any given time $t \geq 0$, the rate of contagion

(exposition) is given by

$$\lambda(t) = \beta S(t) I(t),\ t \geq 0, \tag{1}$$

where β is a scalar parameter, $S(t)$ and $I(t)$ are the sizes of the susceptible and infected populations at time t, respectively. This means that $\lambda(t)$ is the rate at which previously health individuals will randomly catch the disease at time t, where $\lambda(t)$ is given by Eq. (1).

Recall that σ is the random length of the latency time, after which a newly contaminated individual becomes infectious. Therefore, After experiencing a random *latency time* σ, exposed individuals become *infected*. As **Queue 2** represents the patient's journey through the latency phase (Fig. 2), the output of this queue $\delta_e(t)$ represents the number of new infections at time t. According to Eick et al. [29], $\delta_e(t)$ is a Poisson random variable with mean:

$$\overline{\delta_e}(t) = \mathcal{E}(\lambda\,(t - \sigma)\,). \tag{2}$$

This is because the latency period σ is the service time of Queue 2.

Now let us recall that γ is the random recovery time after the infection is manifested, i.e., after an individual becomes infectious. Therefore, after experiencing the *recovery time* (γ), the infected patient becomes *removed*. As **Queue 1** represents an individual's whole journey, including latency and recovery (Fig. 2), the output of this queue $\delta_i(t)$ represents the number of new removals at time t. Recalling that the total time from exposition to removal is the sum of random variables σ and γ, the results of Eick et al. [29] implies that $\delta_i(t)$ is a Poisson random variable with mean:

$$\overline{\delta_i}(t) = \mathcal{E}(\lambda(\,t - (\sigma + \gamma)\,)). \tag{3}$$

Remark 2. *Note that, at time t, both $S(t)$ and $I(t)$ are known; therefore, so is $\lambda(t)$, which is a deterministic function of these quantities. Moreover, since the system has a single realisation, all past values of $\lambda(t)$ are known at time t. Therefore, the results of Eick et al. [29] can be applied as they require that $\lambda(t)$ be a known function of time up to time t; they do not require the knowledge of the future realisations of $\lambda(t)$.*

We develop an equivalent time-varying Markov model to our queuing formulation in the following subsection. This will allow us to describe the system's dynamics via simple

transition probabilities, enabling a parsimonious simulation of an epidemic across a large population.

2.1.2. Markov formulation with time-varying transition rate

We make use of the arrival and departure rates of Queues 1 and 2 in Eq. (1)-(3) to define a time-varying Markov process X_t, $t \geq 0$ that describes the evolution of the populations in the SEIR compartments. At any time $t \geq 0$, $X(t) = (\, S(t),\, E(t),\, I(t), R(t)\,) \in \Omega$ is the state of process X_t, $t \geq 0$. By definition, X_t, $t \geq 0$ will be subject to random jumps that happen whenever either of these events occurs: *i*) a new exposition, represented as a new arrival in the queuing system, *ii*) a new infection - represented as a departure from Queue 2, or *iii*) a new removal, i.e., a new departure from Queue 1 - see Fig. 2.

To describe the dynamics of the process X_t, $t \geq 0$, we use Poisson variables to represent the input and output of both queues [29] at any given time $t \geq 0$. Consequently, the time until the next event will be exponentially distributed, and the total jump rate at time $t \geq 0$ is:

$$\Lambda(t) = \lambda(t) + \overline{\delta_e}(t) + \overline{\delta_i}(t). \tag{4}$$

Let $\{\tau_0,\, \tau_1, \ldots\}$ be the sequence of jumps in the system, with $\tau_0 \equiv 0$ and $\tau_{k+1} > \tau_k$, $\forall k \geq 0$. Now, assume a jump occurs at time $t = \tau_k$. The following holds:

$$P(\, X(t^+) = Y | X(t) = (S(t),\, E(t),\, I(t),\, R(t))),\, t = \tau_k\,)$$
$$= \begin{cases} \dfrac{\lambda(t)}{\Lambda(t)} & \text{if } Y = (S(t) - 1,\ E(t) + 1,\ I(t),\ R(t))\,) \\ \dfrac{\overline{\delta_e}(t)}{\Lambda(t)} & \text{if } Y = (S(t),\ E(t) - 1,\ I(t) + 1,\ R(t))\,) \\ \dfrac{\overline{\delta_i}(t)}{\Lambda(t)} & \text{if } Y = (S(t),\ E(t),\ I(t) - 1,\ R(t) + 1)\,) \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The first expression on the right-hand side of Eq. (5) corresponds to a new contagion/exposition, which happens at rate $\lambda(t)$, see Eq. (1), and implies the transference of

an individual from the susceptible to the exposed population. The second expression corresponds to a new infection that happens upon the departure of an individual from Queue 2, which occurs at rate $\overline{\delta_e}(t)$, see Eq. (2). In that case, this individual moves from the exposed to the infected population. Finally, the third possibility is the departure of an individual from Queue 1, which happens at rate $\overline{\delta_i}(t)$, see Eq. (3). In that case, this individual migrates from the infected to the removed population.

Finally, after the jump at $t = \tau_k$, $k \geq 0$, the value of the exposition rate $\lambda(t)$ also changes, and becomes:

$$\lambda(t^+) = \beta S(t^+) I(t^+), \tag{6}$$

where the new values on the right-hand side of the expression above vary as a function of the transition probabilities in (5), clearly, $\lambda(t)$ remains unaltered between successive jumps. Given an initial distribution and the transition probability in (5), the process X_t, $t \geq 0$ is a continuous-time Markov chain [18].

3. Stochastic stability and reproduction number

Whereas deterministic models rely on the evaluation of trivial equilibrium to derive stability conditions and the so-called reproduction number R_0 [e.g. 37], the notion of stochastic stability [e.g., 38] provides an ideal framework to evaluate the conditions for a receding epidemic as time elapses. Queuing theory connects this notion with the so-called occupation rate, a measure of the input-to-output ratio as time elapses [36], deriving conditions to ensure that the infected population consistently decreases towards zero.

Consider Queue 1 in Figure 2. Assuming a considerable population, the stability of such a queue hinges on the occupation rate [e.g., 36]:

$$\rho(t) = \frac{\lambda(t)}{\overline{\delta_i}(t)}, \tag{7}$$

and can be ascertained if a finite time $\bar{t} \geq 0$ exists such that $\rho(t) < 1, \forall t \geq \bar{t}$. This guarantees that the system stabilises and the number of customers in the queue remains finite. However, stability is guaranteed with a finite population because the number of infected individuals

will remain within finite, albeit possibly large, bounds. In that case, we are interested in the trend of Queue 1, which indicates whether the epidemic is increasing or receding. Theorem 1 below establishes the condition for a receding epidemic.

Theorem 1. *Consider the Markov process described in Section 2.1.2 and assume that* $\rho(t) < 1$ *for all* $t > \bar{t} \geq 0$*, with* $\bar{t} < \infty$*. Then, it follows that:*

$$\mathcal{E}\left(E(t^+) + I(t^+) \mid X(t)\ ;\ t = \tau_k\right) < E(t) + I(t),\ \forall t = \tau_k > \bar{t},\ k \geq 0, \tag{8}$$

where $\{\tau_0, \tau_1, \ldots\}$ *is the sequence of jumps in the system, as defined in Section 2.1.2.*

Proof From (5) we have that:

$$\begin{aligned}\mathcal{E}\left(E(t^+) + I(t^+) \mid X(t)\ ;\ t = \tau_k\right) =& \frac{\lambda(t)}{\Lambda(t)}\left(E(t) + I(t) + 1\right) + \\ & \frac{\overline{\delta_e}(t)}{\Lambda(t)}\left(E(t) + I(t)\right) + \frac{\overline{\delta_i}(t)}{\Lambda(t)}\left(E(t) + I(t) - 1\right) \\ =& E(t) + I(t) + \frac{\lambda(t) - \overline{\delta_i}(t)}{\Lambda(t)},\end{aligned}$$

for all $t = \tau_k$. The last equality holds because $\Lambda(t) = \lambda(t) + \overline{\delta_e}(t) + \overline{\delta_i}(t)$ - Eq. (4). Now, Eq. (8) follows by assuming $\rho(t) < 1$ for all $t > \bar{t} \geq 0$.

Notice that Eq. (8) means that the expected total number of infected and exposed individuals will decrease at each new jump time τ_k, $k \geq 0$. As a consequence of Theorem 1, we can interpret $\rho(t)$ as the reproduction number of the system at time t and $\rho(t) < 1$ for all $t > \bar{t} \geq 0$ as a sufficient condition for the epidemic to stabilise and shrink after time $\bar{t}$. It intuitively states that if more people are recovering than getting infected, we have a decreasing trend of infections.

Theorem 1 implies that we must consider the disease's whole cycle, from exposition to removal, to evaluate the epidemic trend. It underscores the importance of accurately tracking the epidemic's evolution through an efficient testing strategy. The accurate evaluation of such a cycle's length will also be essential to evaluate the lags between mitigating actions and decreases in infections and expositions. Although mitigation can prevent future expositions and infections, it does not affect recent transmissions still in the latent stage, which will continue to manifest and may drive infection up in the first stages of the mitigation.

3.1. Mitigation strategies

Theorem 1 provides a basis for developing mitigation strategies to stabilise the epidemic, whereas the Markov model in Section 2.1.2 enables us to evaluate the long-term effects of such strategies, for example via simulation. Following the literature, we introduce the control (mitigating action) in the form of non-pharmaceutical interventions to limit the spread of the disease [e.g., 39, 40, 35]. A control level $0 \leq u(t) \leq 1$ attains a reduction of $100u(t)$ percentage points in the transmission rate at time $t \geq 0$, thus resulting in a controlled exposition rate:

$$\lambda(t, u(t)) = \beta(1 - u(t))S(t)I(t),\ t \geq 0. \tag{9}$$

Let $\pi = \{u(t),\, t \geq 0\}$ be a mitigation strategy and let Π denote the set of all feasible strategies $\Pi = \{\pi : 0 \leq u(t) \leq 1,\, \forall t \geq 0\}$. The dynamics of the system under any strategy $\pi \in \Pi$ can be evaluated by running the Markov chain X_t, $t \geq 0$ with the same underlying dynamics, but making $\lambda(t) = \lambda(t, u(t))$, $t \geq 0$. This means at any time t, the mitigation policy will prevent $100u(t)\%$ of the new contagions that would happen without such an intervention.

We are particularly interested in the class of stabilising policies $\Pi^S \in \Pi$, such that:

$$u(t) = \begin{cases} 0, & \text{if } t \leq \bar{t}, \\ \min\left[\left(1 - \dfrac{\bar{\rho}}{\rho(t)}\right),\ 0\right], & \text{if } t > \bar{t}, \end{cases} \tag{10}$$

where $\rho(t)$ is the uncontrolled occupation rate in Eq. (7), designed to ensure that the expected rate of new contagions after the control is at most $100\bar{\rho}\,\%$ of the expected number of recoveries, thus ensuring a consistent decrease of the epidemic levels.

It is easy to see that such policies lead to a controlled occupation rate $\rho_c(t) = \frac{\lambda(t,u(t))}{\bar{\delta}_i(t)} \leq \bar{\rho}$, $\forall t > \bar{t}$, see Eq. (9). Henceforth, $\bar{\rho}$ will be called *target occupation level*, and $\pi = (\bar{t}, \bar{\rho})$ will describe any stabilising policy $\pi \in \Pi^S$ that satisfies Eq. (10). In the next Section, we will evaluate these policies in light of the COVID-19 epidemic and explore the stabilising effect of parameters $\bar{\rho}$ and $\bar{t}$ in Eq. (10). Observe that because the population is finite, the effect of the control can be limited by a delayed start of mitigation.

Remark 3. *Note from Eq.* (9)–(10) *that it is always possible to find a control* $u \in [0, 1]$ *that leads to a particular target occupation rate* $\bar{\rho} \in [0, 1]$. *Hence, finding mitigation levels to drive infections down and stabilise the epidemic is also possible. Note also that the mitigating levels can act as a guide for policy-making. They determine how effective the set of mitigating actions should be to drive infection levels down. For example,* $u(t) = 0.5$ *means that the set of actions mandated—such as mask-wearing and social isolation should prevent at least 50% of the new contagions. In possession of prevention targets, one can evaluate different combinations of actions to attain such targets, dynamically incorporating newly discovered courses of action.*

Observe, however, that the Markov model in Section 2.1.2 is not limited to the proposed class of mitigation strategies. One can substitute any arbitrary mitigation policy $\pi \in \Pi$ in Eq. (9) and run the controlled Markov chain to evaluate the effect of such a policy. To demonstrate that, our experiments will also include the *on-off* lockdown policies proposed by Tarrataca et al. [35], but evaluated there using a deterministic model with average rates. Designed to control hospital bed occupation, these policies trigger a full-scale lockdown when infections surpass a prescribed upper bound; conversely, mitigating actions are lifted when infections fall below a prescribed lower bound.

4. Numerical Experiments

This Section validates the proposed stochastic SEIR dynamic model in Section 2.1.2 in light of data from the COVID-19 epidemic evolution in England and the state of Amazonas, Brazil.

4.1. The case of England

Given the identified inconsistencies in public datasets reporting positive tests [e.g. 31, 41] and the actual number of infections, we will first validate our model with data from the weekly COVID-19 survey that estimates by random sampling the total levels of infection in England since mid-2020 [42]. The survey provides a reliable benchmark for model validation before we investigate ad-hoc control strategies in the following subsection. Table 3 lists the population and transmission rate parameters estimated by Arruda et al. [43].

Table 3: Parameters of the Simulation.

Parameter	**Description**	**Value**
P	Total population	$56,000,000$
$I(0)$	Initial number of infected people	$50,400$
$E(0)$	Initial number of exposed people	$16,800$
$S(0)$	Initial number of susceptible people	$55,932,799$
β	Transmission rate	$2.55 \cdot 10^{-9}$

To complete the necessary data to run the model, we also need the distributions of the latency period σ and the latency plus infection period $(\sigma+\gamma)$, representing the entire disease cycle. The distributions are compatible with previous studies on the random duration of the disease cycle [25, 34] and assumed discrete to cope with the typical daily data collection.

Tables 4 and 5 unveil the distributions of σ and $(\sigma + \gamma)$, respectively.

Table 4: Latency Period Distribution (σ).

Day	0	1	2	3	4	5	6	7
Probability	-	0.0009	0.0056	0.0222	0.0611	0.1222	0.1833	0.2095
Day	8	9	10	11	12	13	14	
Probability	0.1833	0.1222	0.0611	0.0222	0.0056	0.0009	$6 \cdot 10^{-5}$	

Remark 4. *It is worth reinforcing that latency and recovery periods are disease-specific parameters that are not deterministic nor observable from epidemic reports. Hence, it is unsurprising that fitting these parameters from epidemic data compromises the results, as infinitely many parameter combinations can be found [33]. The latency period is non-observable, and the daily number of new infections is often imprecise in epidemic reports as these depend on a myriad of factors (e.g., availability of tests, people volunteering to test, logistics, among others). Fortunately, latency and recovery periods can be recovered from medical studies of the disease, such as [25, 34], which are independent of epidemic data and can therefore give rise to identifiable models [33].*

Table 5: Latency and Infection Period Distribution ($\sigma + \gamma$).

Day	0	1	2	3	4	5	6	7
Prob.	-	-	-	-	-	-	-	-
Day	8	9	10	11	12	13	14	15
Prob.	-	-	$3.87 \cdot 10^{-7}$	$4.84 \cdot 10^{-6}$	$3.87 \cdot 10^{-5}$	0.0002	0.0010	0.0034
Day	16	17	18	19	20	21	22	23
Prob.	0.0098	0.0233	0.0466	0.0792	0.1151	0.1439	0.1550	0.1439
Day	24	25	26	27	28	29	30	31
Prob.	0.1151	0.0792	0.0466	0.0233	0.0098	0.0034	0.0010	0.0002
Day	32	33	34	35				
Prob.	$3.87 \cdot 10^{-5}$	$4.84 \cdot 10^{-6}$	$3.87 \cdot 10^{-7}$	$1.49 \cdot 10^{-8}$				

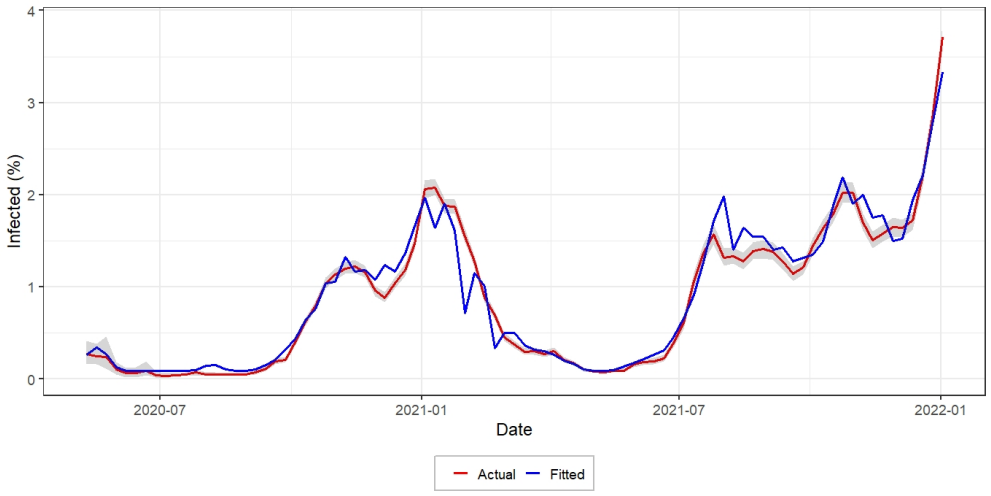

Figure 3: **Epidemic evolution in England. The grey area surrounding the red line corresponds to the reported 95% confidence interval centered in the median (red line).**

Figure 3 depicts the simulated infection levels as time elapses against the 95% confidence interval of England's weekly estimated infection levels [42]. The mean square error (MSE) between the predicted and observed values was 0.0313. To adjust the data, we needed to infer the control levels $u(t)$ during the epidemic evolution, resulting from the many mitigation policies applied in the country in response to the epidemic evolution. Figure 4 depicts the

control effect as time elapses.

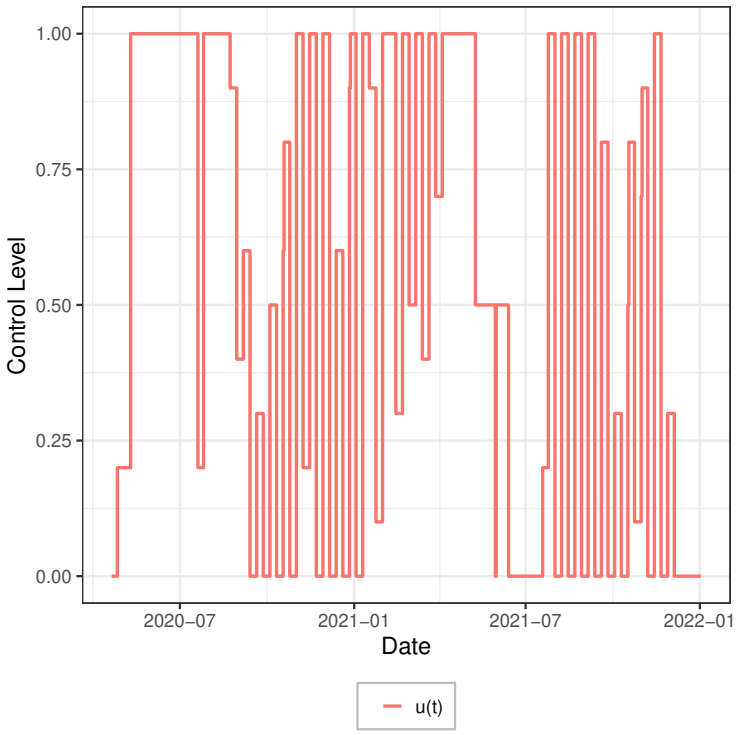

Figure 4: **Estimated preventive effect of epidemic mitigating policies in England**

It is evident in Figure 3 that our approach can effectively follow the epidemic trend as time elapses. Meanwhile, Figure 4 illustrates the approach's ability to infer the effect of mitigation in terms of infection prevention under uncertainty. One can notice an increase in control levels when the epidemic is decreasing - the effect of the mitigating policies applied in response to surges in infection, as well as a gradual decrease in control levels at times when relaxed mitigation led to increased infection levels.

Figure 5 depicts the error evolution between the predicted infection levels and the median of England's weekly estimated infection levels. One can see that the proposed approach closely follows the observed trends, keeping the errors limited as time elapses.

Finally, Figure 6 depicts the 95% confidence interval of the predicted infection levels in England across 100 simulations. One can see that the predictions are robust and that the confidence interval remains limited as time elapses. That behaviour is maintained in the

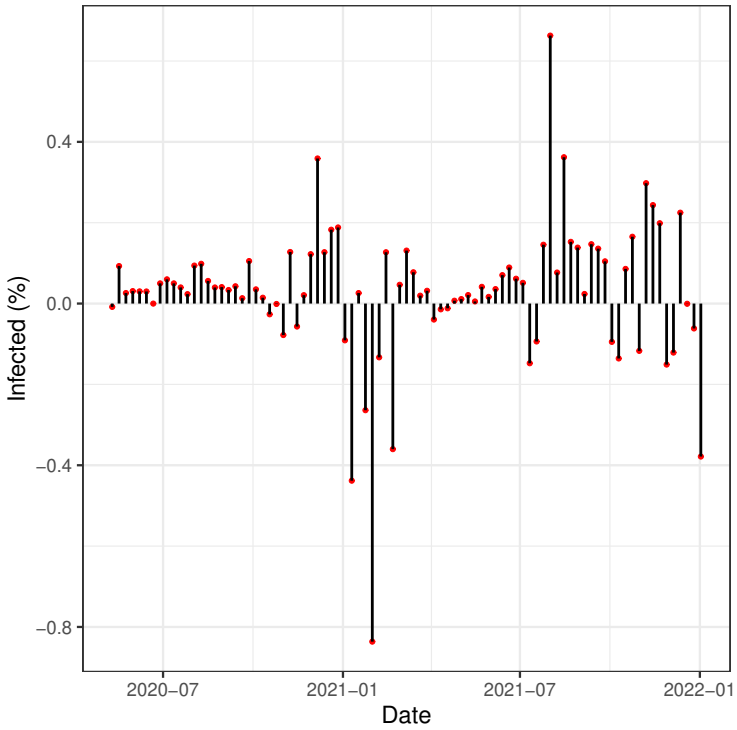

Figure 5: **Prediction Error**

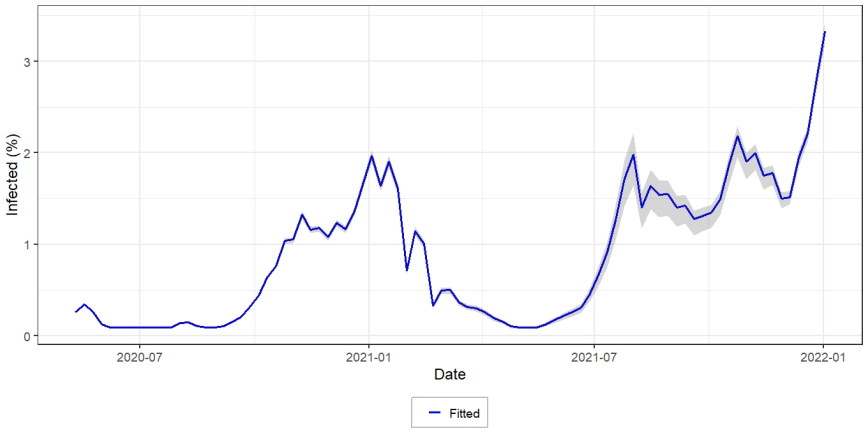

Figure 6: **95% confidence interval for the model**

simulations of the following subsections; therefore, we only plot the median of the simulated values, as the variation remains limited, and our analysis focuses on the mitigation policies' effects.

4.2. The epidemic in the state of Amazonas, Brazil

To simulate the epidemic evolution in the State of Amazonas, we employ parameters listed in Table 6, as well as the latency and infectious period distributions of Tables 4 and 5. Without mitigation, that results in a total infection rate of about 80% as the system stabilises, in line with the 76% rate estimated from a serological study in that state [30].

Table 6: Parameters of the Simulation.

Parameter	**Description**	**Value**
P	Total population	$4,144,596$
$I(0)$	Initial number of infected people	2
$E(0)$	Initial number of exposed people	252
$S(0)$	Initial number of susceptible people	$4,144,342$
β	Transmission rate	$3.447 \cdot 10^{-8}$

4.3. The effect of delayed mitigation

The first experiment series examines policies $\pi = (\bar{t}, \bar{\rho}) \in \Pi^S$ that follow Eq. (10) with a fixed target occupation level $\bar{\rho} = 0.95$. The objective is to evaluate the effect of delaying the start of mitigating actions by varying the first parameter. Cases A-F in Table 7 feature different triggering times for the mitigating actions, starting with the case where no mitigation is enforced ($\bar{t} = \infty$). For each experiment, we ran 100 simulations of the system for two years, starting from the outset of the epidemic. The values plotted in the following figures correspond to the median of the simulated trajectories.

Remark 5. *The target occupation level is a parameter that determines how fast the epidemic will reduce and can be adjusted depending on the decision maker's preferences. We vary the target rates in Section 4.4 to expand the analysis.*

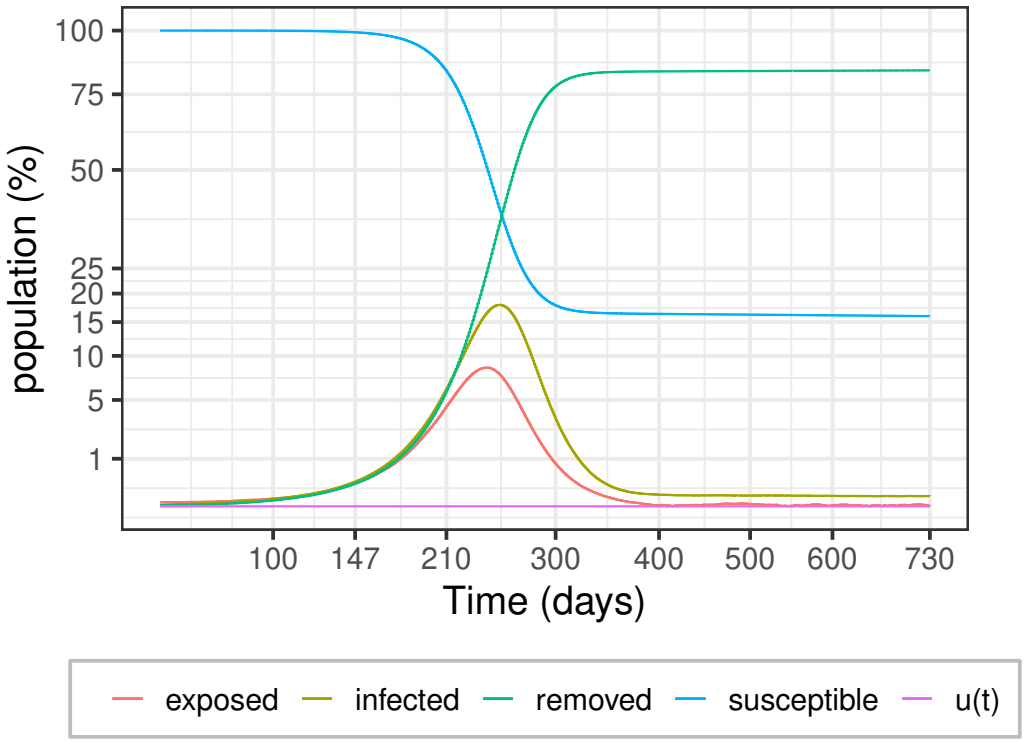

Figure 7: **Epidemic evolution for Case A: the unmitigated spread**

Figure 7 shows the daily evolution of the epidemic for Case A, i.e., in the absence of mitigating actions. The number of individuals in each SEIR compartment is depicted as a percentage of the total population to facilitate the interpretation. One can see that the epidemic spreads rapidly, with the number of infected individuals peaking at around 19% of the population after 250 days. The epidemic shrinks due to the decrease in the susceptible population until it subsides around the 400^{th} day. Overall, about 85% of the population catch the disease at some point during the two years.

Figure 8 shows that Case B starts mitigating early in the epidemic, with low infection

Table 7: Simulated cases.

Case	A	B	C	D	E	F
$\bar{t}$	∞	147	168	189	210	231
$\bar{\rho}$	0.95	0.95	0.95	0.95	0.95	0.95

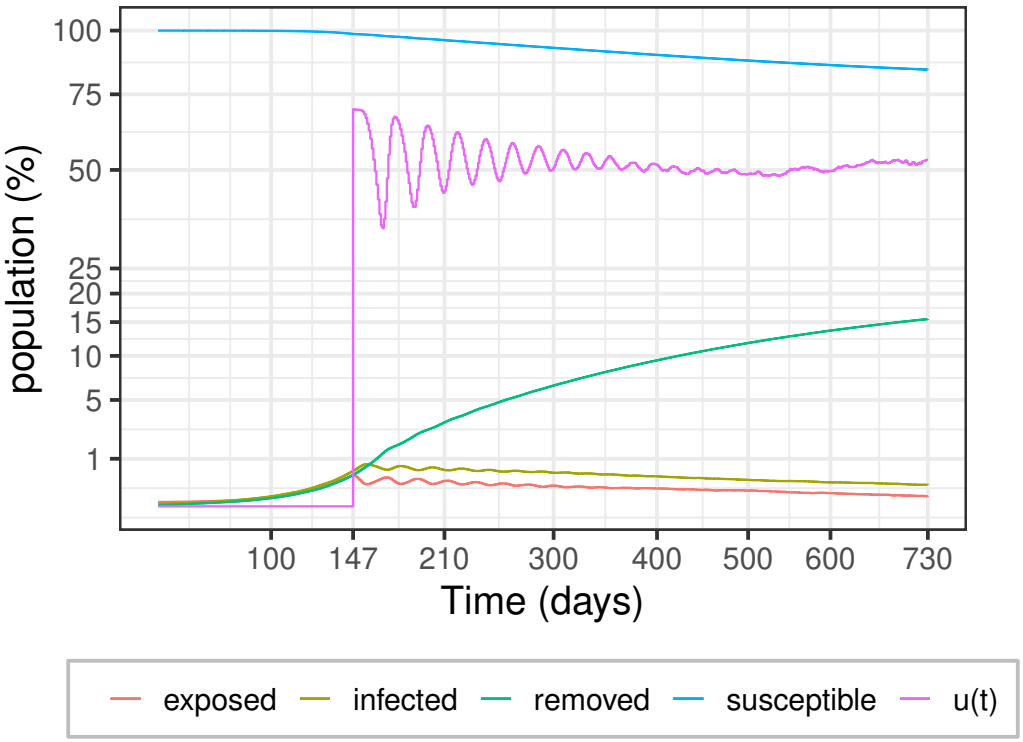

Figure 8: **Epidemic evolution for Case B**

levels. The results illustrate the effectiveness of the mitigation policy. Observe that the mitigation policy successfully prevents the spread of the disease, maintaining 85% of the population healthy over the whole two-year interval. In contrast, the removed population reaches 15% at the end of that interval. Controlling the epidemic, however, demands high levels of mitigation that oscillate and stabilise around 0.50.

In Case C, the mitigation strategy is delayed for an extra three weeks concerning Case B. The mitigating actions are similar, tend to stabilise around the same value, and are equally sufficient to curb the epidemic. The difference concerning Case B is an increase in the final number of individuals that catch the disease at some point, represented by the removed population. This number increases from 15% in Case B to around 26% in Case C, thus highlighting the sensitivity of the spread concerning mitigation delay. The results also show that a three-week delay produces a higher level of infection throughout the epidemic, with the peak of infections just exceeding 1% of the population and decreasing slowly over time.

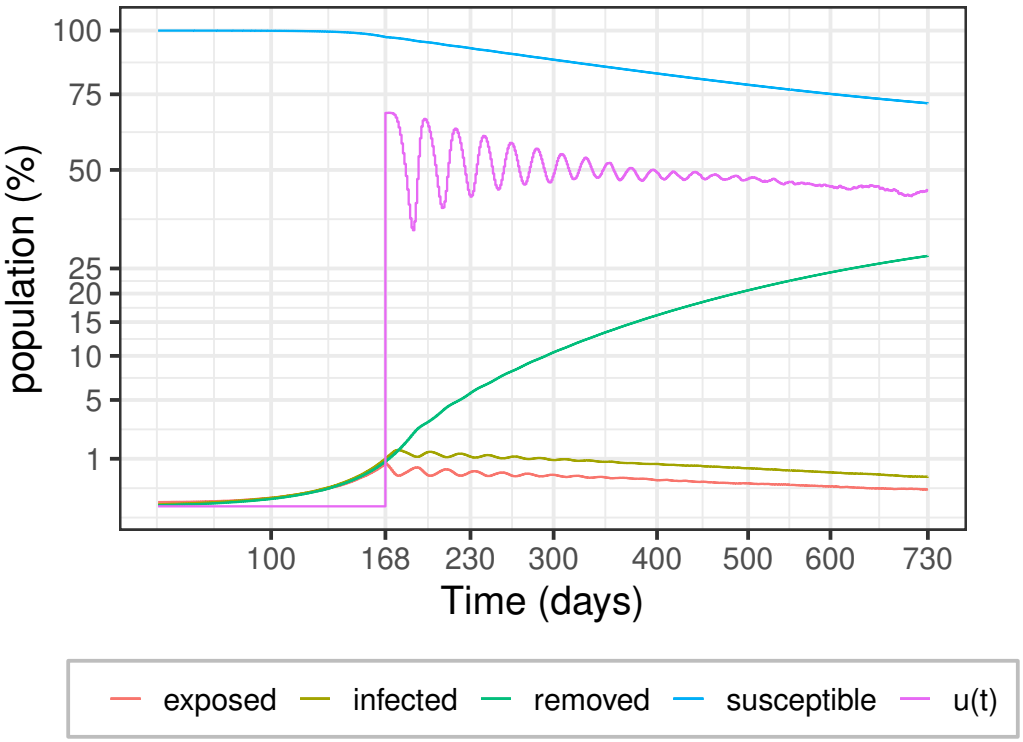

Figure 9: **Epidemic evolution for Case C**

Figures 10 feature the results for Case D. Despite an extra delay of three weeks concerning Case C, it is still possible to flatten the curve and prevent an uncontrollable epidemic increase. Nonetheless, the delay produces a peak of infections five times as large as in Case C, with 5% of people being infected 200 days after the outset of the epidemic. The total removed population skyrockets from about 26% in Case C to about 70% in Case of D, thus reinforcing the importance of early mitigation to prevent the spread of the disease and the exponential effect of extra mitigation delays. Due to the spread of the disease and the corresponding decrease in the number of susceptible individuals, we can slowly relax the mitigation until the disease has contaminated about 50% of the population. As the disease spreads further, we can rapidly relax the mitigation while maintaining a steady decrease in infected individuals. The mitigation is completely lifted just before the anniversary of the outset of the epidemic. After about 600 days, the control levels are increased again to prevent a recrudescence.

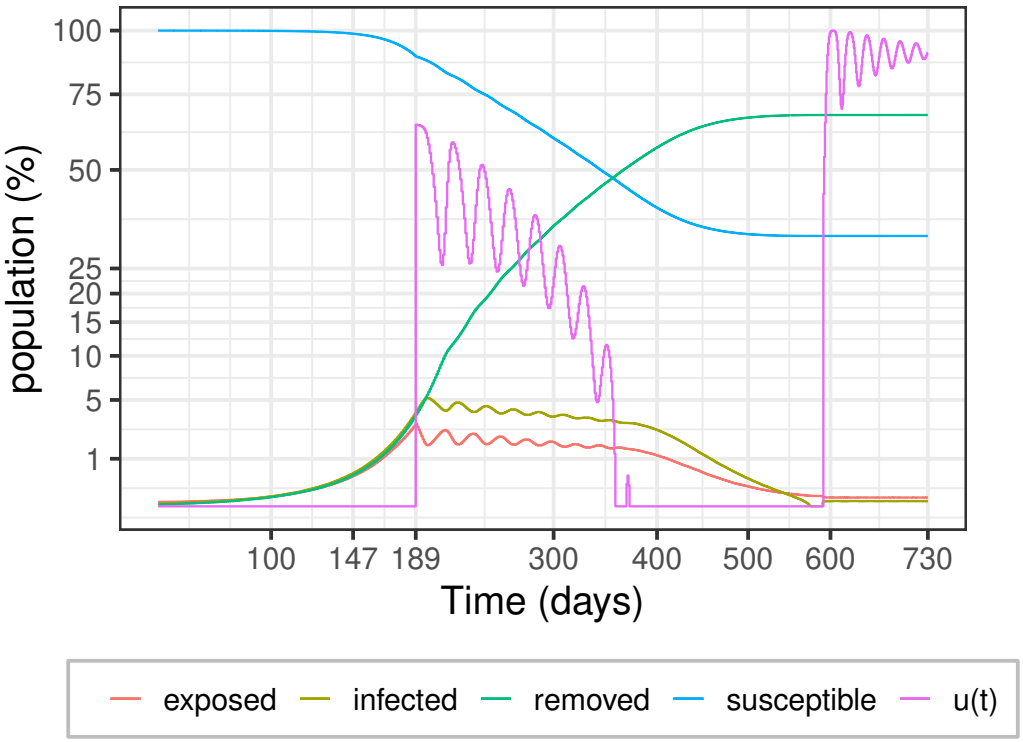

Figure 10: **Epidemic evolution for Case D**

Figure 11 shows the results for Case E that provide further evidence of the deleterious effect of mitigation delay. A further delay of three weeks concerning Case D results in an infection peak of around 9% closely after the start of mitigation. Due to the more prominent peak, we can relax the mitigation earlier and more rapidly, reaching complete relaxation just before 300 days. The final number of removed individuals indicates that the epidemic affects about 75% of the population within two years.

Finally, Figure 12 presents the results for Case F, demonstrating that delaying the mitigation by 12 weeks effectively renders it meaningless. Indeed, on day 231, the uncontrolled epidemic is so widespread that the infection levels are dropping due to fewer susceptible individuals remaining in the population. The infection peaks on day 240 when it simultaneously afflicts approximately 17% of the population.

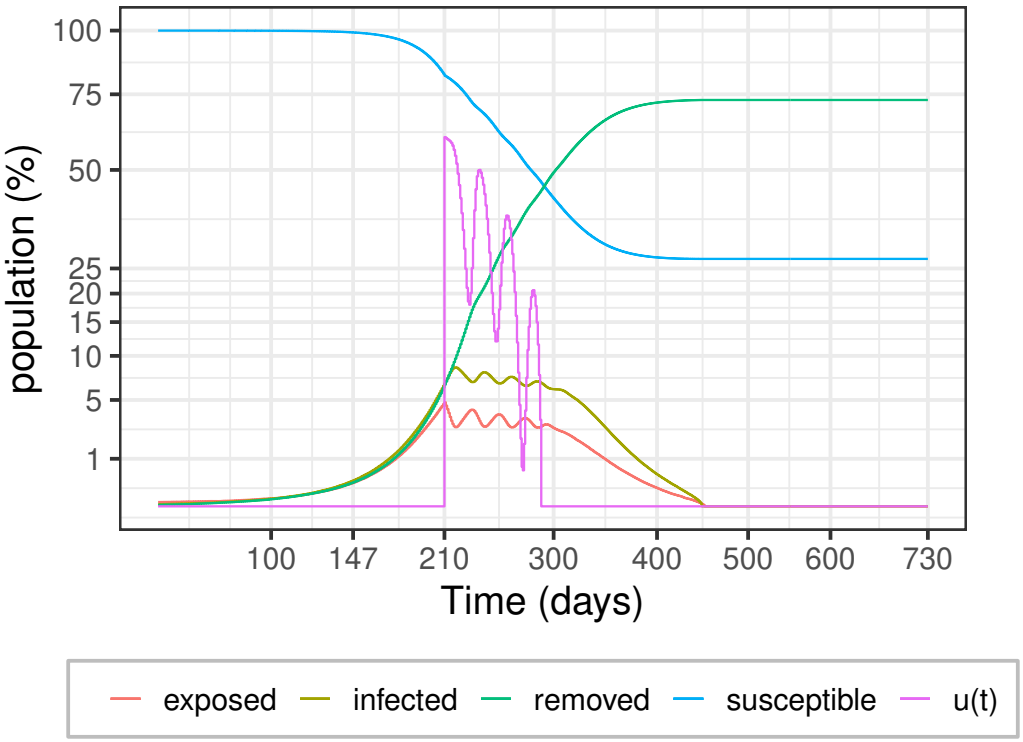

Figure 11: **Epidemic evolution for Case E**

4.3.1. Brief notes on policy implementation

As observed in the experiments, since the stabilising policies are adaptive, they will prescribe mitigation levels that vary over time. However, we can set up fixed average mitigation levels for specified intervals for practical purposes. For example, suppose the decision-maker wishes to revise the mitigation policies bi-weekly. In that case, we can use the model to simulate the next two weeks and find the control levels for the whole period. Then, we can recommend the average value of the control in the next two weeks as a fixed mitigation level for the period. This process would repeat every other week.

4.4. The effect of distinct occupation rates

This Section investigates the effect of different target occupation rates $\bar{\rho}$ in the stabilising policies. To provide an exciting baseline, we start mitigating actions on $\bar{t} = 189$ when the epidemic reaches a significant proportion of the population. However, as in Case D of Figure

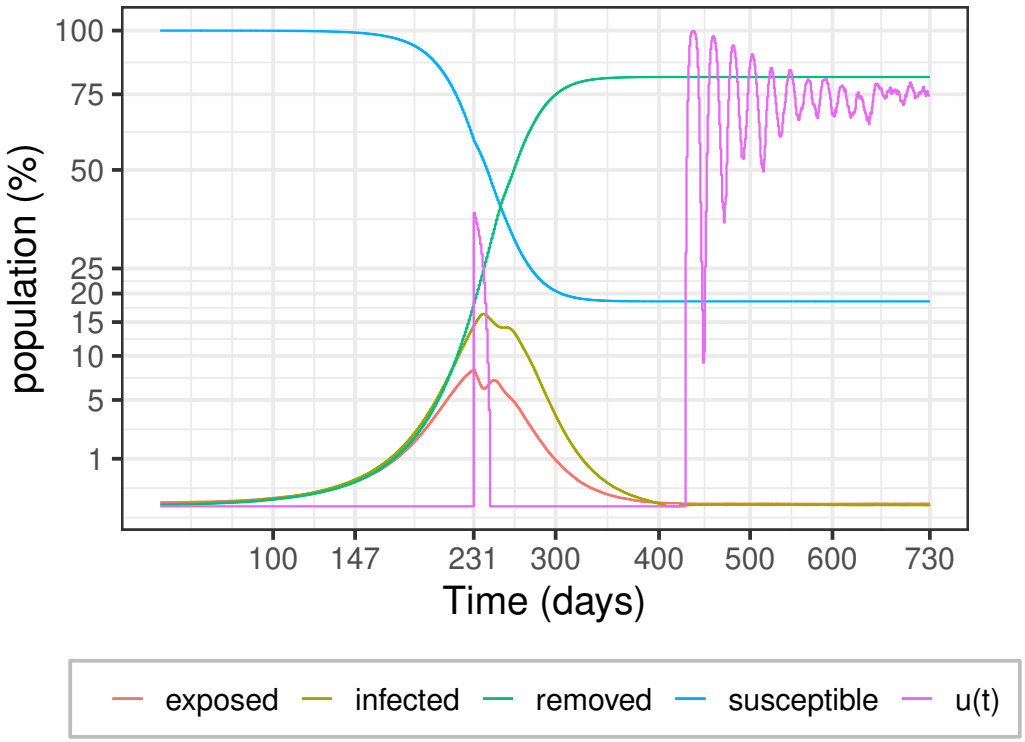

Figure 12: **Epidemic evolution for Case F**

10, the epidemic can still be controlled. The set of experiments in Table 8 assesses the effect of varying the target occupation level from 60% to 90% in regular intervals.

Table 8: Second set of experiments.

Case	H	I	J	K
$\bar{t}$	189	189	189	189
$\bar{\delta}$	0.60	0.70	0.80	0.90

Figure 13 presents the results for Case H; as one can observe, the percentage of infected people peaks around 7%. As expected, the more substantial mitigation has no effect in the early stages and cannot prevent a similar peak as in the baseline (Case D). However, as time elapses, we can notice that the increased mitigation produces a much more pronounced decrease in infections. In effect, infections are virtually extinguished after about 500 days.

The increased control can keep the removed population at about 26% at the end of two years, whereas that level was 70% in the baseline. Notably, the mitigation is kept high after the system stabilises to avoid recrudescence.

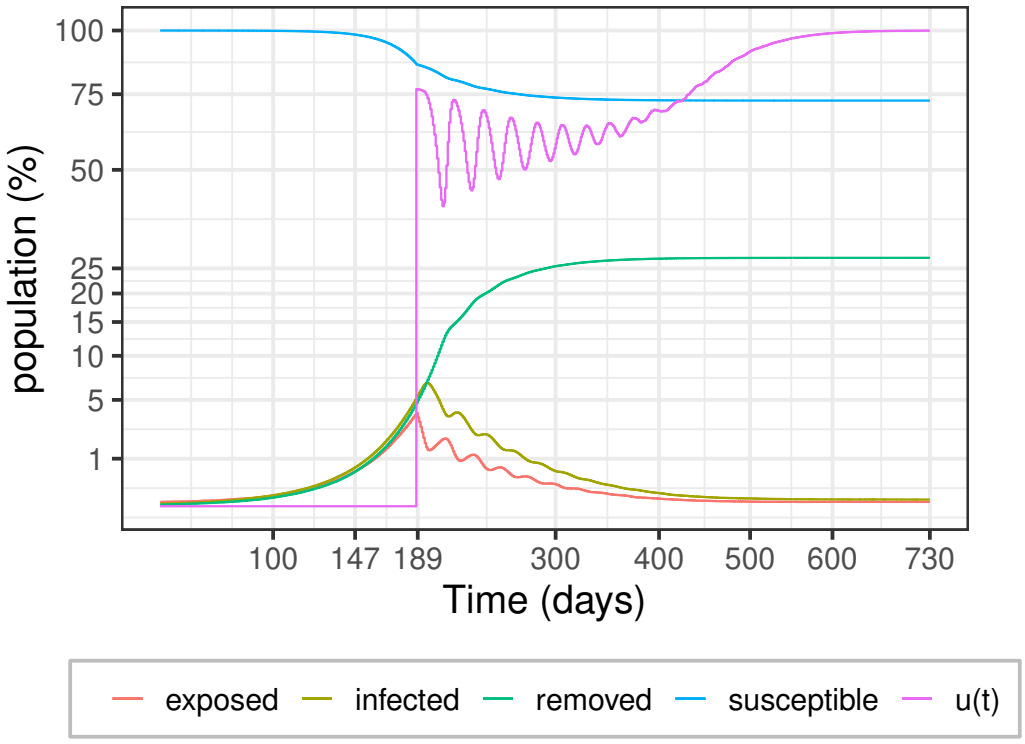

Figure 13: **Epidemic evolution for Case H**

The results for Case I are shown in Figure 14. The increase in the target occupation rate $\bar{\rho}$ concerning Case H has several noticeable effects: a slight decrease in the mitigating actions, a slower rate of decrease of the infection levels, and an increase in the removed population. The infection levels decrease to 0.5% after two years. The reduced control levels also lead to an increase in the removed population, which climbs from 26% in Case H to close to 37% in Case I.

The trend of increasing infection levels gathers speed in Case J, as the decrease in the required mitigating actions leads to about 50% of removals within the two-year horizon. Coupled with the significant decrease in the susceptible population, which reduces the potential spread, the decrease in $\bar{\rho}$ produces a steep decrease in the control levels. However,

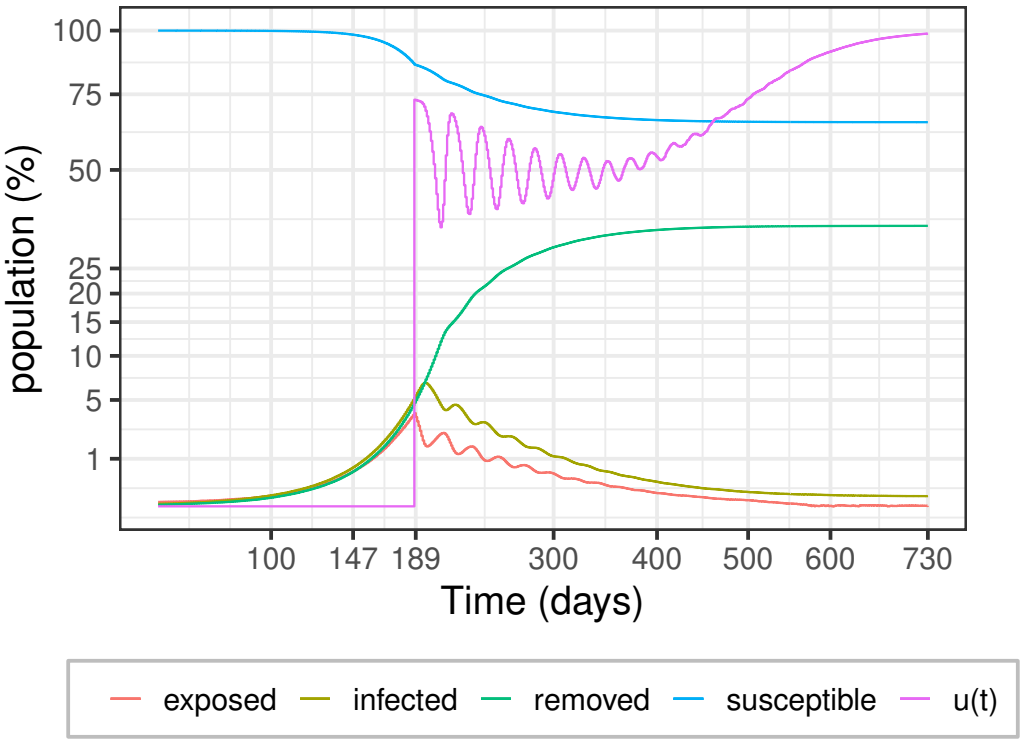

Figure 14: **Epidemic evolution for Case I**

increased levels of mitigation are required later to avoid a second wave.

Finally, Figure 16 presents the results for Case K. As the target occupation level approaches the baseline (Case D), the overall behaviour becomes quite similar. Despite dropping mitigating actions around day 350, the reduced size of the susceptible population causes the infection to fall gradually. After the system stabilises, late mitigation efforts are needed to keep the infection at bay just before the 500^{th} day.

4.5. The effect of on-off lockdown policies

This Section briefly evaluates the effect of *on-off* lockdown policies. Proposed by Tarrataca et al. [35], these policies demand a full-scale lockdown ($u(t) = 1$, $\bar{\rho} = 0$) when the number of infections surpasses a prescribed upper limit; conversely, all measures are lifted ($u(t) = 0$, $\bar{\rho} = \infty$) as soon as these numbers drop below a lower bound. We simulate the two policies whose upper and lower bounds appear in Table 9.

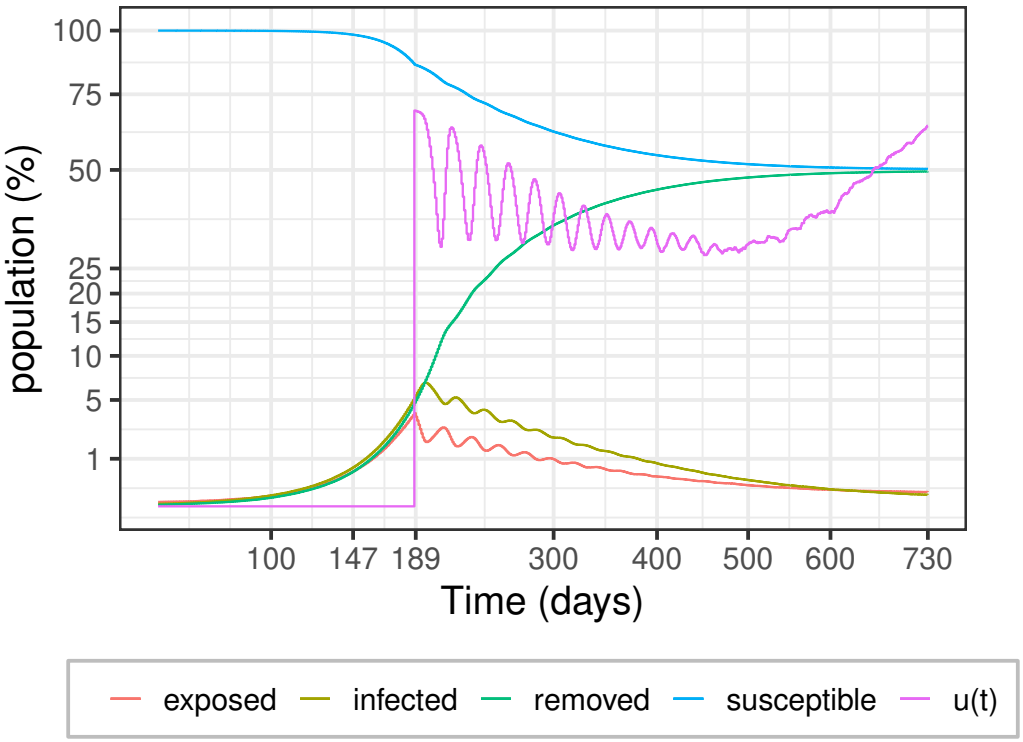

Figure 15: **Epidemic evolution for Case J**

Table 9: The third set of experiments.

Case	L	M
Start	1,5%	2%
Stop	0.3%	1%

Figure 17 shows that Case L features five lockdowns in two years, the first starting on day 160. Infections recurrently peak just above 1.5% after the start of the first lockdown; the removed population amounts to 22% at the end of the second year. Consistently with the results reported by Tarrataca et al. [35], the system alternates between rapid increases and decreases in infection, and the lockdowns become shorter and more spaced over time.

A similar behaviour is observed in Figure 18 for Case M. But because the upper and lower bounds are more significant, the total number of removals increases from 22% to over 40% in two years. In addition, the peak of infections increases to 3%. Although Case M

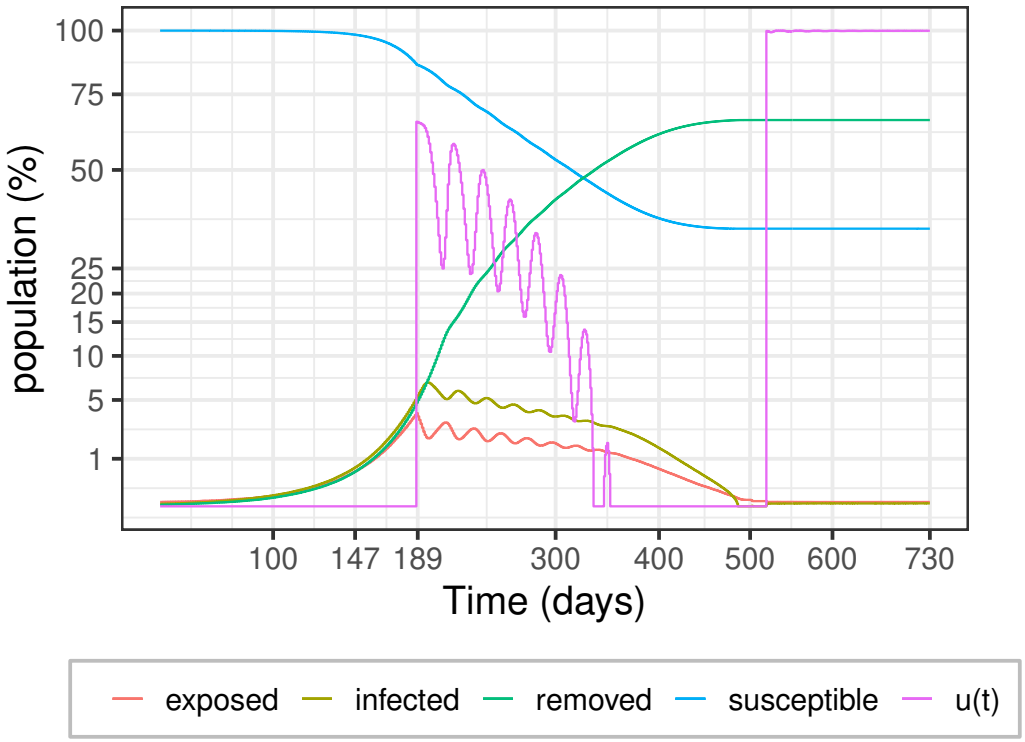

Figure 16: **Epidemic evolution for Case K**

has two more lockdowns in two years, the total time in lockdown is similar, as illustrated in Figure 19, which shows the cumulative time in lockdown on the left-hand side and the duration of individual lockdowns on the opposite half. The mean duration of a lockdown is about 21 days in case L and 17 days in case M.

Comparing with on-off strategies, one can realise that, among other things, our stochastic SEIR model allows for smoother control and can avoid high infection levels as time elapses. That results in consistently lower hospital occupation levels and, consequently, lower societal costs. It also allows for a smooth and planned decrease in infection levels while considering the uncertainty in the spread.

5. Conclusions

Inspired by the unprecedented COVID-19 epidemic, this paper explored innovative stochastic modelling alternatives to describe the evolution of an epidemic using the classical SEIR

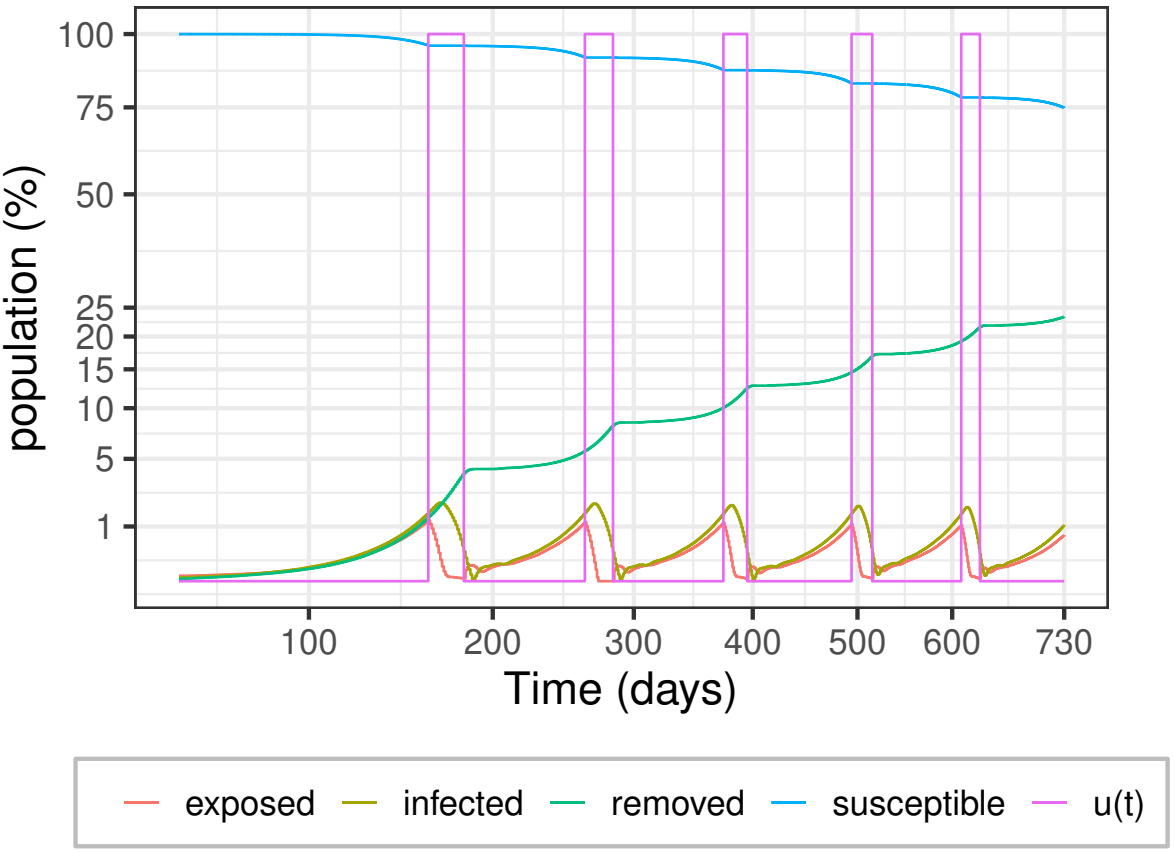

Figure 17: **Epidemic evolution for Case L**

framework. The proposed stochastic model innovates by addressing the spread of the disease within a realistic model that considers uncertainty and is compatible with random latency and infection times, regardless of their distribution. The model is dynamic and adaptive, but our approach ensures that despite its generality, it remains simple, tractable, and valuable to support decision-making as the decision-maker can set up the level of control to enforce a prescribed and sustainable decrease in the infected population.

By drawing a parallel between stochastic stability and the traditional reproduction number, the paper introduced stabilising strategies to curb an epidemic under general conditions, provided that mitigation is initiated promptly. Beyond the academic contributions, the generality of the proposed approach renders it invaluable to supporting real-world decision-making in the face of future epidemics. The methodology was validated using official COVID-

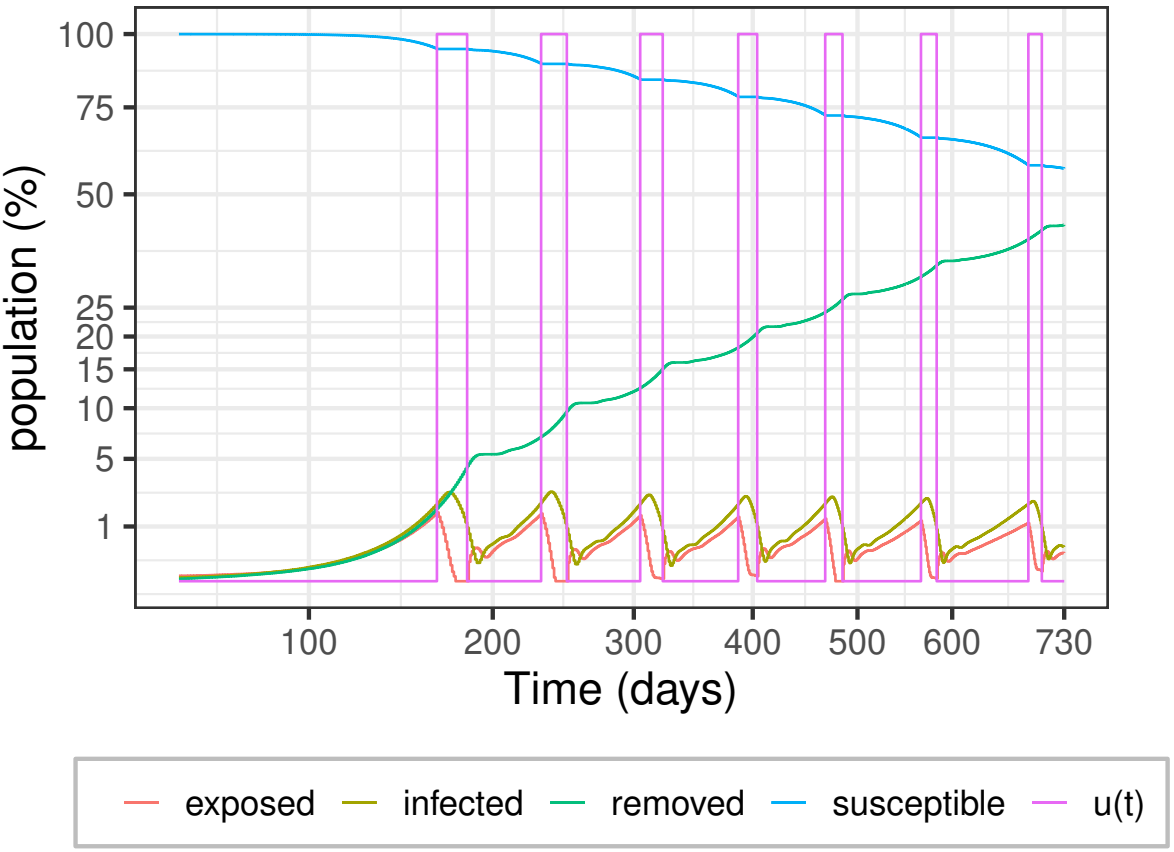

Figure 18: **Epidemic evolution for Case M**

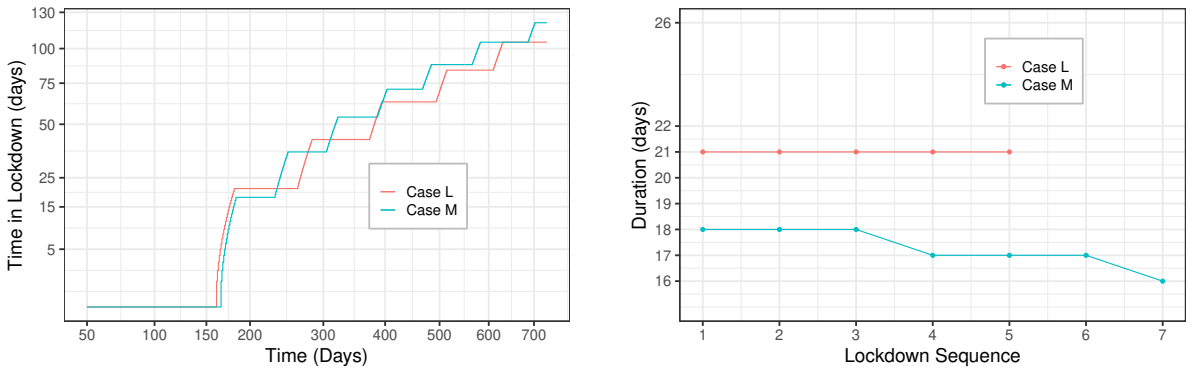

Figure 19: **Time in Lockdown in Cases L and M**

19 data from England and literature reports on the epidemic spread in Amazonas, Brazil. The results provide a panorama of insights into the pros and cons of distinct stabilising mitigation strategies over a two-year horizon. Additionally, the fact that the model could

closely replicate England's epidemic and estimate the equivalent control levels validates the approach in a reliable real-world setting, highlighting its value in supporting the design of epidemic mitigation strategies.One limitation is that the approach does not include multiple viral variants, reinfection, or vaccination. Future work should contemplate the inclusion of these aspects, either individually or combined.

As expected, the experiments illustrate that early intervention is vital to prevent the disease from affecting a large portion of the population. However, to effectively prevent the spread, we need very high levels of mitigation over the whole two-year horizon. The results also suggest that delayed mitigation leads to a dramatic increase in the overall number of infections. In effect, delays tend to produce persistent infection levels over time, even with mitigation, increasing the spread even though the infection curves are flattened. Beyond flattening the curves, this behaviour suggests that we also need to ensure, by acting swiftly, that their summit is tolerable from both societal and healthcare perspectives.

The proposed approach leaves many research avenues to be explored in future works. One possibility is to evaluate the model with multiple virus variants, reinfection, and vaccination. Another obvious route is to pursue stochastic optimal control policies that somehow address the compromise between healthcare aspects, societal issues, and the economic burden of mitigation strategies. The challenges involve finding a meaningful trade-off among the different elements that decision-makers need to consider and proposing effective formulations that avoid the curse of dimensionality [44, 45] to ensure that the problem remains tractable. Another branch goes into developing filtering and analytical approaches to estimate a system's parameters considering that the available information is delayed and biased since dependent on local testing and reporting policies.

Acknowledgements

The Brazilian Research Council partly supported this study—CNPq, under grants #311075/2018-5, #303352/2018-3 and 304801/2015-1, and by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES) [Finance Code 001].

References

[1] R. Ross. An application of the theory of probabilities to the study of a priori pathometry-part I. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 92(638):204–230, 1916. doi: 10.1098/rspa.1916.0007.

[2] W. O. Kermack, A. G. McKendrick, and G. T. Walker. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, 1927. doi: 10.1098/rspa.1927.0118.

[3] H. W. Hethcote. The Mathematics of Infectious Diseases. *SIAM Review*, 42(4):599–653, 2000. doi: 10.1137/S0036144500371907.

[4] Fangfang Liu, Zheng Ma, Ziqing Wang, and Shaobo Xie. Trade-off between covid-19 pandemic prevention and control and economic stimulus. *International Journal of Environmental Research and Public Health*, 19(21):13956, 2022.

[5] Sean T McQuade, Ryan Weightman, Nathaniel J Merrill, Aayush Yadav, Emmanuel Trélat, Sarah R Allred, and Benedetto Piccoli. Control of covid-19 outbreak using an extended seir model. *Mathematical Models and Methods in Applied Sciences*, 31(12):2399–2424, 2021.

[6] Peri Subrahmanya Hari Prasad. Covid-19 disease spread modeling by qsir method: The parameter optimal control approach. *Clinical epidemiology and global health*, 13:100934, 2022.

[7] Manotosh Mandal, Soovoojeet Jana, Swapan Kumar Nandi, Anupam Khatua, Sayani Adak, and TK Kar. A model based study on the dynamics of covid-19: Prediction and control. *Chaos, Solitons & Fractals*, 136:109889, 2020.

[8] Ranjit Kumar Upadhyay, Sourin Chatterjee, Parimita Roy, and Dyuti Bhardwaj. Combating covid-19 crisis and predicting the second wave in europe: an age-structured modeling. *Journal of Applied Mathematics and Computing*, pages 1–21, 2022.

[9] J. Wallinga and M. Lipsitch. How generation intervals shape the relationship between growth rates and reproductive numbers. *Proceedings of the Royal Society B: Biological Sciences*, 274(1609):599–604, 2007. doi: 10.1098/rspb.2006.3754.

[10] L. J. S. Allen. An Introduction to Stochastic Epidemic Models. In F. Brauer, P. van den Driessche, and J. Wu, editors, *Mathematical Epidemiology*, pages 81–130. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISBN 978-3-540-78911-6. doi: 10.1007/978-3-540-78911-6_3.

[11] T. Britton. Stochastic epidemic models: A survey. *Mathematical Biosciences*, 225(1):24 – 35, 2010. ISSN 0025-5564. doi: 10.1016/j.mbs.2010.01.006.

[12] Abu Quwsar Ohi, MF Mridha, Muhammad Mostafa Monowar, and Md Abdul Hamid. Exploring optimal control of epidemic spread using reinforcement learning. *Scientific reports*, 10(1):22106, 2020.

[13] Askat Kuzdeuov, Daulet Baimukashev, Aknur Karabay, Bauyrzhan Ibragimov, Almas Mirzakhmetov,

Mukhamet Nurpeiissov, Michael Lewis, and Huseyin Atakan Varol. A network-based stochastic epidemic simulator: controlling covid-19 with region-specific policies. *IEEE journal of biomedical and health informatics*, 24(10):2743–2754, 2020.

[14] Dror Meidan, Nava Schulmann, Reuven Cohen, Simcha Haber, Eyal Yaniv, Ronit Sarid, and Baruch Barzel. Alternating quarantine for sustainable epidemic mitigation. *Nature communications*, 12(1):220, 2021.

[15] P. Trapman and M. C. J. Bootsma. A useful relationship between epidemiology and queueing theory: The distribution of the number of infectives at the moment of the first detection. *Mathematical Biosciences*, 219(1):15 – 22, 2009. ISSN 0025-5564. doi: 10.1016/j.mbs.2009.02.001.

[16] C. O. Dike, Z. M. Zainuddin, and I. J. Dike. Queueing Technique for Ebola Virus Disease Transmission and Control Analysis. *Indian Journal of Science and Technology*, 9(46), 2016. doi: 10.17485/ijst/2016/v9i46/107077.

[17] N. R. Barraza, G. Pena, and V. Moreno. A non-homogeneous Markov early epidemic growth dynamics model. application to the SARS-CoV-2 pandemic. *Chaos, Solitons & Fractals*, 139:110297, 2020. ISSN 0960-0779. doi: 10.1016/j.chaos.2020.110297.

[18] P. Brémaud. *Gibbs fields, monte carlo simulation, and queues.* Springer-Verlag, New York, 1999.

[19] M. López-García. Stochastic descriptors in an SIR epidemic model for heterogeneous individuals in small networks. *Mathematical Biosciences*, 271:42 – 61, 2016. ISSN 0025-5564. doi: 10.1016/j.mbs.2015.10.010.

[20] D. Clancy. SIR epidemic models with general infectious period distribution. *Statistics & Probability Letters*, 85:1 – 5, 2014. ISSN 0167-7152. doi: 10.1016/j.spl.2013.10.017.

[21] A. Gómez-Corral and M. López-García. On SIR epidemic models with generally distributed infectious periods: Number of secondary cases and probability of infection. *International Journal of Biomathematics*, 10(02):1750024, 2017. doi: 10.1142/S1793524517500243.

[22] C. Lefèvre and M. Simon. SIR-Type Epidemic Models as Block-Structured Markov Processes. *Methodology and Computing in Applied Probability*, 22(2):433–453, 2020. doi: 10.1007/s11009-019-09710-y.

[23] Z. G. İşlier, R. Güllü, and W. Hörmann. An exact and implementable computation of the final outbreak size distribution under Erlang distributed infectious period. *Mathematical Biosciences*, 325:108363, 2020. ISSN 0025-5564. doi: 10.1016/j.mbs.2020.108363.

[24] M. H. A. Davis. *Markov models and optimization.* Chapman and Hall, London, 1993.

[25] J. Backer, D. Klinkenberg, and J. Wallinga. Incubation period of 2019 novel coronavirus (2019-nCoV) infections among travellers from Wuhan, China, 20–28 january 2020. *Eurosurveillance*, 25, 02 2020. doi: 10.2807/1560-7917.ES.2020.25.5.2000062.

[26] J. R. Artalejo, A. Economou, and M. J. Lopez-Herrero. The stochastic SEIR model before extinction:

Computational approaches. *Applied Mathematics and Computation*, 265:1026 – 1043, 2015. ISSN 0096-3003. doi: 10.1016/j.amc.2015.05.141.

[27] Mariajesus Lopez-Herrero. Epidemic Transmission on SEIR Stochastic Models with Nonlinear Incidence Rate. *Mathematical Methods in the Applied Sciences*, 40(7):2532–2541, 2017. doi: 10.1002/mma.4179.

[28] J. Amador and M. Lopez-Herrero. Cumulative and maximum epidemic sizes for a nonlinear SEIR stochastic model with limited resources. *Discrete & Continuous Dynamical Systems - B*, 23:3137, 2018. ISSN 1531-3492. doi: 10.3934/dcdsb.2017211.

[29] S. G. Eick, W. A. Massey, and W. Whitt. The Physics of the $M_t/G/\infty$ Queue. *Operations Research*, 41(4):731–742, 1993. doi: 10.1287/opre.41.4.731.

[30] Lewis F. Buss, Carlos A. Prete, Claudia M. M. Abrahim, Alfredo Mendrone, Tassila Salomon, Cesar de Almeida-Neto, Rafael F. O. França, Maria C. Belotti, Maria P. S. S. Carvalho, Allyson G. Costa, Myuki A. E. Crispim, Suzete C. Ferreira, Nelson A. Fraiji, Susie Gurzenda, Charles Whittaker, Leonardo T. Kamaura, Pedro L. Takecian, Pedro da Silva Peixoto, Marcio K. Oikawa, Anna S. Nishiya, Vanderson Rocha, Nanci A. Salles, Andreza Aruska de Souza Santos, Martirene A. da Silva, Brian Custer, Kris V. Parag, Manoel Barral-Netto, Moritz U. G. Kraemer, Rafael H. M. Pereira, Oliver G. Pybus, Michael P. Busch, Márcia C. Castro, Christopher Dye, Vítor H. Nascimento, Nuno R. Faria, and Ester C. Sabino. Three-quarters attack rate of SARS-CoV-2 in the brazilian amazon during a largely unmitigated epidemic. *Science*, 371(6526):288–292, 2021. ISSN 0036-8075. doi: 10.1126/science.abe9728.

[31] E. C. Sabino, L. F. Buss, M. P. S. Carvalho, C. A. Prete Jrt, M. A. E. Crispim, N. A. Fraiji, R. H. M. Pereira, K. V. Parag, P. S. Peixoto, M. U. G. Kraemer, M. K. Oikawa, T. Salomon, Z. M. Cucunuba, M. C. Castro, A. A. S. Santos, V. H. Nascimento, H. S. Pereira, N. M. Ferguson, O. G. Pybus, A. Kucharski, M. P. Busch, C. Dye, and N. R. Faria. Resurgence of COVID-19 in Manaus, Brazil, despite high seroprevalence. *The Lancet*, 397(10273):452–455, 2021. doi: 10.1016/S0140-6736(21)00183-5.

[32] Fundação de Vigilância em Saúde do Amazonas. COVID-19 no amazonas. dados epidemiológicos e financeiros das ações de combate à covid-19. publicações, 2021. URL `http://www.fvs.am.gov.br/publicacoes`. accessed 20 May 2021.

[33] N. Tuncer and T. T. Le. Structural and practical identifiability analysis of outbreak models. *Mathematical Biosciences*, 299:1–18, 2018. ISSN 0025-5564. doi: https://doi.org/10.1016/j.mbs.2018.02.004.

[34] R. Verity, L. C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, G/ Cuomo-Dannenburg, H. Thompson, P. G. T. Walker, H. Fu, A. Dighe, J. T. Griffin, M. Baguelin, S. Bhatia, A. Boonyasiri, A. Cori, Z. Cucunubá, R. FitzJohn, K. Gaythorpe, W. Green, A. Hamlet, W. Hinsley, D. Laydon, G. Nedjati-Gilani, S. Riley, S. van Elsland, E. Volz, H. Wang, Y. Wang, X. Xi, C. A. Donnelly, A. C.

Ghani, and N. M. Ferguson. Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*, 2020. ISSN 1473-3099. doi: 10.1016/S1473-3099(20)30243-7.

[35] L. Tarrataca, C. M. Dias, D. Haddad, and E. F. Arruda. Flattening the curves: on-off lock-down strategies for COVID-19 with an application to Brazil. *Journal of Mathematics in Industry*, 11(1):2, 2021. doi: 10.1186/s13362-020-00098-w.

[36] J. Shortle, J. Thompson, D. Gross, and C. Harris. *Fundamentals of Queueing Theory*. Wiley Series in Probability and Statistics. Wiley, New York, 5 edition, 2018. ISBN 9781118943526. doi: 10.1002/9781119453765.

[37] P. van den Driessche and J. Watmough. Reproduction numbers and sub-threshold endemic equilibria for compartmental models of disease transmission. *Mathematical Biosciences*, 180(1):29 – 48, 2002. ISSN 0025-5564. doi: https://doi.org/10.1016/S0025-5564(02)00108-6.

[38] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, New York, 1993.

[39] N Ferguson, D Laydon, G Nedjati Gilani, N Imai, K Ainslie, M Baguelin, S Bhatia, A Boonyasiri, ZULMA Cucunuba Perez, G Cuomo-Dannenburg, A Dighe, I Dorigatti, H Fu, K Gaythorpe, W Green, A Hamlet, W Hinsley, L Okell, S Van Elsland, H Thompson, R Verity, E Volz, H Wang, Y Wang, P Walker, P Winskill, C Whittaker, C Donnelly, S Riley, and A Ghani. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Technical report, Imperial College London, 03 2020.

[40] M. Kantner and T. Koprucki. Beyond just "flattening the curve": Optimal control of epidemics with purely non-pharmaceutical interventions. *Journal of Mathematics in Industry*, 10:23, 2020. doi: 10.1186/s13362-020-00091-3.

[41] Owen Dyer. Covid-19: Peru's official death toll triples to become world's highest. *BMJ*, 373, 2021. doi: 10.1136/bmj.n1442.

[42] ONS. Coronavirus (COVID-19) infection survey, UK statistical bulletins. `https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/bulletins/coronaviruscovid19infectionsurveypilot/previousReleases`, 2023. Accessed: 2023-01-10.

[43] E. F. Arruda, S. S. Das, C. M. Dias, and D. H. Pastore. Modelling and optimal control of multi strain epidemics, with application to COVID-19. *PLOS ONE*, 16(9):e0257512, 2021. doi: 10.1371/journal.pone.0257512.

[44] W.B. Powell. *Approximate Dynamic Programming Solving the Curses of Dimensionality*. John Wiley & Sons, Inc., New Jersey, USA, 2011.

[45] W.B. Powell. A Unified Framework for Stochastic Optimization. *European Journal of Operational*

Research, 275(3):795–821, 6 2019. ISSN 0377-2217. doi: 10.1016/J.EJOR.2018.07.014.