




Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

Full Length Article

Robust and optimal predictive control of the COVID-19 outbreak[☆]Johannes Köhler^{*}, Lukas Schwenkel, Anne Koch, Julian Berberich, Patricia Pauli, Frank Allgöwer

Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany



ARTICLE INFO

Keywords:

Epidemic control
COVID-19
Optimal control
Model predictive control
Robustness

ABSTRACT

We investigate adaptive strategies to robustly and optimally control the COVID-19 pandemic via social distancing measures based on the example of Germany. Our goal is to minimize the number of fatalities over the course of two years without inducing excessive social costs. We consider a tailored model of the German COVID-19 outbreak with different parameter sets to design and validate our approach. Our analysis reveals that an open-loop optimal control policy can significantly decrease the number of fatalities when compared to simpler policies under the assumption of exact model knowledge. In a more realistic scenario with uncertain data and model mismatch, a feedback strategy that updates the policy weekly using model predictive control (MPC) leads to a reliable performance, even when applied to a validation model with deviant parameters. On top of that, we propose a robust MPC-based feedback policy using interval arithmetic that adapts the social distancing measures cautiously and safely, thus leading to a minimum number of fatalities even if measurements are inaccurate and the infection rates cannot be precisely specified by social distancing. Our theoretical findings support various recent studies by showing that (1) adaptive feedback strategies are required to reliably contain the COVID-19 outbreak, (2) well-designed policies can significantly reduce the number of fatalities compared to simpler ones while keeping the amount of social distancing measures on the same level, and (3) imposing stronger social distancing measures early on is more effective and cheaper in the long run than opening up too soon and restoring stricter measures at a later time.

1. Introduction

Social distancing is an effective way to contain the spread of a contagious disease, particularly when little is known about the virus and no vaccines or other pharmaceutical interventions are available. Social distancing and isolation (together with other non-pharmaceutical measures such as hygiene and face masks) have a direct influence on the infection rates and hence on the spread of the virus (Kissler, Tedi-janto, Lipsitch, & Grad, 2020; Maharaj & Kleczkowski, 2012; Maier & Brockmann, 2020). While this combination has proven effective during the last weeks, e.g. in the German outbreak of COVID-19, strict social distancing is also very costly in terms of economical and psychological damage, which naturally leads to a multi-objective decision problem.

There have been numerous approaches to model the COVID-19 outbreak and to predict future behavior for different distancing policies in simulation studies. The most commonly used modeling approaches are different extensions of the SIR (susceptible–infected–removed) model

formulated either as system dynamics or as agent-based simulations (e.g. for Germany (Barbarossa et al., 2020; Dehning et al., 2020; German, Djanatljev, Maile, Bazan, & Hackstein, 2020)). In many such studies, different policies are simulated and compared with respect to the goals that both the health care system is not overwhelmed such that every patient in need receives treatment and the mortality rate is kept low, and also such that the majority of people can resume social interaction as soon as possible. However, in line with Alleman, Torfs, and Nopens (2020) and others, we advocate to go from mere model predictions to (model predictive) control, since control generally offers the theory to develop and apply optimal or robust decision making under uncertainties.

While mathematical modeling and control of epidemics is a topic with rich history (see, e.g., the survey in Nowzari, Preciado, and Pappas (2016) and the references therein), there have also been numerous approaches to apply control theory to the COVID-19 spread. In Casella

[☆] This work was supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Grants GRK 2198/1 - 277536708, AL 316/12-2 - 279734922. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Lukas Schwenkel, Anne Koch, Julian Berberich, and Patricia Pauli.

^{*} Corresponding author.

E-mail addresses: johannes.koehler@ist.uni-stuttgart.de (J. Köhler), lukas.schwenkel@ist.uni-stuttgart.de (L. Schwenkel), anne.koch@ist.uni-stuttgart.de (A. Koch), julian.berberich@ist.uni-stuttgart.de (J. Berberich), patricia.pauli@ist.uni-stuttgart.de (P. Pauli), frank.allgower@ist.uni-stuttgart.de (F. Allgöwer).

(2021), the author applies control theoretic principles and insights to a simple model of the outbreak to point out the difficulties of the system at hand: fast unstable dynamics with significant delays. In more recent literature, multiple works have addressed the problem of open-loop optimal control for the COVID-19 pandemic. In Bin et al. (2020), Tsay, Lejarza, Stadtherr, and Baldea (2020), for example, the authors argue in favor of 'on and off'-policies of the social distancing measures, yielding a bang–bang like optimal control strategy. Such 'on and off'-policies, where the control input switches between two states, however, could pose great challenges, amongst others, for the society, but also for production lines, supply chains and the economy in general.

In this paper, we propose optimal open-loop and feedback control strategies to handle the German COVID-19 outbreak. We employ the recently developed SIDARTHE model (Giordano et al., 2020) in order to design control policies which minimize the number of fatalities within a time horizon of two years, without using excessive social distancing measures. We also address robustness of our policies w.r.t. model and measurement uncertainties via a (robust) model predictive control (MPC) feedback strategy. Note that the following discussion and results are all based on the information and data available prior to the initial submission of this paper (May 2020).

Similar to the setup in this paper, the authors in Djidjou-Demasse, Michalakos, Choisy, Sofonea, and Alizon (2020) explore the best policy to implement while waiting for the availability of a vaccine. In their paper, they also distinguish between varying severity of symptoms ('mild' or 'severe') and seek a solution to the multi-objective optimization problem of minimizing fatalities and costs due to the implementation of the control strategy itself. Their main outcome of the open-loop input strategy is qualitatively similar to our results in Section 3.2: Start with a loose strategy, soon increase all distancing measures such that the health care system capacities are never exceeded and then relax the social distancing measures gradually and slowly. Another example for an open-loop optimal policy applied to the COVID-19 pandemic is presented in Kantner and Koprucki (2020) where the authors consider optimal control of the German outbreak using a slightly simpler model as the one chosen in the present paper (without distinguishing between detected and undetected individuals), which also includes an increased mortality rate if the ICU capacity is exceeded. Therein, the objective is not only to minimize the number of fatalities but also the number of susceptible individuals at the end of the time horizon, thus aiming for herd immunity. Our investigations in Section 3, however, indicate that herd immunity cannot be reached in a reasonable amount of time without overwhelming the hospital capacities. Therefore, our approach minimizes the number of fatalities after two years, with the underlying assumption that a vaccine will be available thereafter.

However, an open-loop optimal policy cannot suffice to control the COVID-19 pandemic given all the uncertainties in the spreading of the virus and the disease progression, as we will see in the numerical results. We argue, similar to Alleman et al. (2020), that an MPC-based feedback strategy is the right tool to develop optimal and robust social distancing policies, especially in the presence of model inaccuracies. By using online measurements of the current outbreak, feedback inherently introduces robustness with respect to uncertainties and disturbances to the policy. We also robustify the feedback mechanism by introducing a robust MPC-based feedback strategy for uncertain state measurements which is crucial in a situation where only a limited amount of data is available and, for example, the number of the currently infected persons can only be estimated roughly by applying different studies.

Our results are also in accordance with a very recently published joint strategy paper for Germany by authors from different German research institutions (Fraunhofer-Gesellschaft, Helmholtz Association, Leibniz Association and Max Planck Society) (Meyer-Hermann, Pigeot, Priesemann, & Schöbel, 2020). Firstly, they also state that reaching herd immunity without the availability of a vaccine would either exceed the health care capacities (with a resulting high mortality rate)

or take several years (cf. our results in Section 3). Secondly, they state that the goal of wiping out the virus can only be a robust solution if this eradication would be a worldwide effort with very high social and economic costs (cf. our results in Section 3.1.1), which seems impossible to realize. Finally, they suggest an adaptive strategy for all policies influencing the infection rates with the goal to keep the spread of COVID-19 at bay while requiring the least possible restraints on the society and economy. With exactly this reasoning, we develop suitable control approaches in Section 3.2, Section 4, and Section 5 for such an 'adaptive' strategy.

To summarize, our key contributions are the following:

- We extend the model in Giordano et al. (2020) by a mortality rate dependent on the state of the health care capacity and fit the parameters with data from Germany (Section 2).
- We develop an optimization problem for finding the optimal input (in terms of setting infection rates) that minimizes the number of fatalities while keeping the costs occurring due to distancing measures low (Section 3). Moreover, we show that such an optimal input has significant advantages compared to simpler baseline policies.
- We show that simply applying a precomputed (optimized) input is dangerous if the model is uncertain and explain why feedback is of utmost importance when dealing with such an unstable and uncertain system. Further, we demonstrate how such feedback can be incorporated via MPC, and we showcase the advantages of this control policy (Section 4).
- We develop a robust MPC-based feedback strategy, which takes model inaccuracies, uncertain state measurements, and inexact inputs into account and can thus handle the COVID-19 outbreak cautiously and safely (Section 5).

Although based on a simple model fitted with limited data, we hope that these high-level insights inspire further investigations, possibly on more complex epidemiological models, and can ultimately help decision makers to improve and optimize their policies to mitigate the spread of epidemics while keeping the toll on the society and economy low.

2. Modeling of the COVID-19 epidemic

In this section, we describe the model of the COVID-19 epidemic that we use for our subsequent control approach. Our model is adapted from the SIDARTHE model proposed in Giordano et al. (2020) with the key differences that (i) we use more recent data to estimate new parameters to model the German COVID-19 outbreak (in Giordano et al. (2020), the Italian outbreak was considered) and (ii) we model the fact that the mortality rate increases if the number of critically ill patients exceeds the capacity of the German health care system. In Section 2.1, we describe the model of Giordano et al. (2020) and explain its ingredients. Thereafter, in Section 2.2, we provide details on our parameter estimation algorithm which fits the model to the German COVID-19 outbreak. Finally, we propose an extension of the model by increasing the mortality rate when the health care system is overwhelmed in Section 2.3.

2.1. The SIDARTHE model

The considered model based on Giordano et al. (2020) is shown in Fig. 1 and includes eight states: **S** - Susceptible, **I** - Infected (asymptomatic, undetected), **D** - Diagnosed (asymptomatic, detected), **A** - Ailing (symptomatic, undetected), **R** - Recognized (symptomatic, detected), **T** - Threatened (symptomatic with life-threatening symptoms, detected), **H** - Healed (immune after prior infection, detected or undetected), **E** - Extinct (dead, detected). In accordance with Fig. 1, the following differential equations describe the SIDARTHE model:

$$\dot{S} = -S(\alpha I + \beta D + \gamma A + \beta R), \quad (1a)$$

$$\dot{I} = S(\alpha I + \beta D + \gamma A + \beta R) - (\epsilon + \zeta + \lambda)I, \tag{1b}$$

$$\dot{D} = \epsilon I - (\zeta + \lambda)D, \tag{1c}$$

$$\dot{A} = \zeta I - (\theta + \mu + \kappa)A, \tag{1d}$$

$$\dot{R} = \zeta D + \theta A - (\mu + \kappa)R, \tag{1e}$$

$$\dot{T} = \mu A + \mu R - (\sigma(T) + \tau(T))T, \tag{1f}$$

$$\dot{H} = \lambda I + \lambda D + \kappa A + \kappa R + \sigma(T)T, \tag{1g}$$

$$\dot{E} = \tau(T)T. \tag{1h}$$

In the Eqs. (1), capital letters describe fractions of the whole population that are currently in the respective state. Since the model represents the whole population, the states sum up to 1, i.e., they must satisfy $S + I + D + A + R + T + H + E = 1$ at all times. Therefore, one equation in (1) is redundant and hence, e.g., the state H can be expressed via the algebraic relation $H = 1 - (S + I + D + A + R + T + E)$ instead of Eq. (1g), as it is common in the field of differential algebraic equations. In most parts of this section, we omit time arguments for simplicity. Further, Greek letters are the model parameters which are briefly summarized in the following:

- α, β, γ describe the infection rates for susceptible individuals, i.e., the rate at which susceptible individuals are infected by the states I, D or R , and A , respectively, and hence join the state I .
- ϵ, θ describe the testing rate, i.e., at which rate (asymptomatic or symptomatic) infected individuals go from undetected to detected.
- ζ describes the rate of asymptomatic (detected or undetected) infected individuals exhibiting symptoms, i.e., going from states I or D to A or R , respectively.
- μ is the rate at which infected individuals in A or R develop life-threatening symptoms, i.e., join the state T .
- $\lambda, \kappa, \sigma(T)$ are recovery rates for individuals affected by COVID-19. The recovery rate for threatened individuals $\sigma(T)$ depends on T , compare Section 2.3.
- $\tau(T)$ is the mortality rate, i.e., the rate at which individuals with life-threatening symptoms decrease, and it depends on T , compare Section 2.3.

Key features of the considered model for the COVID-19 pandemic compared to simpler ones (e.g., SIR models, compare (Kermack & McKendrick, 1927)) are that it distinguishes between detected and undetected cases, symptomatic and asymptomatic individuals, and it includes a separate state T for patients with life-threatening symptoms (compare Giordano et al. (2020) for a more detailed explanation of the key ingredients). The present model, i.e., Eqs. (1) as well as Fig. 1, is a mild modification of the model suggested in Giordano et al. (2020). First, we reduce the number of parameters by including the following assumptions. We assume that the rate for developing (severe) symptoms is the same for detected and undetected cases, since (to this day) no effective medication of COVID-19 is known. More precisely, the transitions from I to A and D to R have the same dynamics with rate ζ , and similarly for the respective recovery rates as well as for the transitions from A to T and R to T . Moreover, we assume that the rate β at which susceptible individuals are infected is the same from states D and R , since the state D is neglected for the parameter identification step (compare Section 2.2). Finally, as a key difference to Giordano et al. (2020), we consider T -dependent rates $\tau(T)$ and $\sigma(T)$ for threatened patients, i.e., the mortality and recovery rates depend on the current number of threatened patients. Essentially, $\tau(T)$ increases and $\sigma(T)$ decreases if T exceeds the capacity of the German health care system (see Section 2.3 for a detailed description of this model refinement).

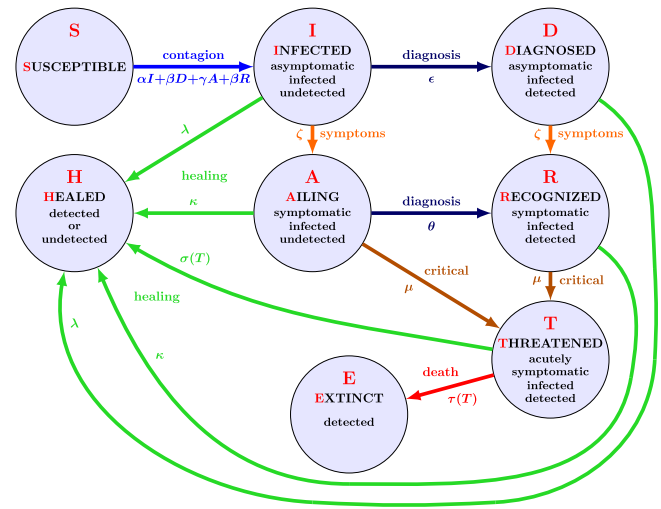


Fig. 1. Scheme of the considered model for the COVID-19 epidemic. Source: Adapted from Giordano et al. (2020).

2.2. Parameters for the German outbreak

In this section, we adjust the model parameters and the initial condition given in Giordano et al. (2020) to the COVID-19 outbreak in Germany. This is necessary, because the outbreaks in Germany and Italy evolved differently due to differences in the testing policy, the testing capacity, the health care system, the reaction of the governmental authorities, and the underlying counting method of confirmed cases.

In order to compute realistic parameters for Germany, we will use a pragmatic approach that enables us to easily include prior knowledge about relations between parameters. The approach is a least squares optimization of the available data, where prior knowledge is incorporated via hard constraints in the optimization problem. The available data is marked by a tilde and is given by:

- the confirmed COVID-19 cases \tilde{C} , deaths \tilde{E} , and recoveries \tilde{H}_c from Dong, Du, and Gardner (2020a, 2020b) for the days $t \in [0, 49]$ from February 28, 2020 ($t = 0$) to April 21, 2020 ($t = 53$). We filtered this data set using the Matlab function `kaiser(7, 3)` with window length 7 and shape factor 3 to reduce the effect of noise corruption and having less confirmed cases during weekends. Further, we have to divide the data set by the total German population $N_{\text{total}} = 8.3 \cdot 10^7$ to ensure all values are normalized and are in the range $[0, 1]$.
- the COVID-19 patients in ICU \tilde{T}_2 and how many of them died \tilde{E}_2 or recovered \tilde{H}_2 from DIVI-IntensivRegister (2020) for $t \in [24, 53]$ from March 23 ($t = 24$) to April 21 ($t = 53$).¹

This data set, however, is rather small compared to the complexity of the model (1) consisting of eight states and 13 parameters. Therefore, we need to leverage additional prior knowledge in order to avoid over-fitting and ensure a realistic resulting parameter set and initial conditions. Based on other studies and the interpretation of our model states and parameters, we enforce the following assumptions.

- The detection rate of asymptomatic cases is negligible, as the current German policy is to test only people showing symptoms (Jakob Simmank, 2020), i.e., $\epsilon = 0$.

¹ There was no reliable data available from before March 23; the data from March 23 to March 26 contains only information about \tilde{T}_2 , not \tilde{E}_2 or \tilde{H}_2 ; from April 5 to April 7, there is a gap in the data due to a server migration of the DIVI Intensivregister (DIVI-IntensivRegister, 2020).

- At February 28, the initial date for our fit, there were 48 confirmed cases, hence, we assume $R(0) = 48/N_{\text{total}}$, $T(0) = D(0) = H(0) = E(0) = 0$, and $I(0)$, $A(0)$, $S(0)$ appear as decision variables with $S(0) = 1 - R(0) - I(0) - A(0)$.
- The test rate θ is approximately constant. Please note that this does not mean that the absolute number of tests is constant per day, as this value is rather proportional to θA than to θ .
- The infection rates α and γ were influenced by the countermeasures that the German authorities installed to fight the spread of the pandemic. According to [Dehning et al. \(2020\)](#), three main events changed the spreading rates: (1) March 9 — canceling large events, (2) March 16 — closing schools and non-essential stores, and (3) March 23 — contact ban (Kontaktsperre) that prohibits groups of more than two people and requires people to maintain a distance of at least 1.5 m in public. Hence, there are four different policies u_i , $i = 1, 2, 3, 4$ monotonically increasing from no countermeasures $u_1 = 0$ to full lockdown $u_4 = 1$ resulting in $\alpha_i = \alpha_{\max} + u_i(\alpha_{\min} - \alpha_{\max})$ and $\gamma_i = \gamma_{\max} + u_i(\gamma_{\min} - \gamma_{\max})$. This yields the following six decision variables α_{\min} , α_{\max} , γ_{\min} , γ_{\max} , u_2 , and u_3 .
- One of the main reasons why the COVID-19 pandemic is spreading so fast is that infectiousness peaks even before the onset of symptoms ([He et al., 2020](#)). As asymptomatic individuals have no indication of their infection, they are on top of that also not as cautious as people with symptoms. Therefore, we require $\alpha \geq \gamma$ when searching for realistic parameters. Further, we want to ensure that people tested positive are significantly less contagious while in quarantine, such that we require $\gamma \geq 5\beta$.
- The percentage of *confirmed* COVID-19 cases is estimated in the study ([Bommer & Vollmer, 2020](#)) as 27.32% in Germany. In our model, this value approaches the constant $\phi = \frac{\zeta}{\lambda + \zeta} \frac{\theta + \mu}{\kappa + \theta + \mu}$, which is the proportion of people that develop symptoms (I to A , I to D can be ignored as $\epsilon = 0$) and get detected (A to R or T); that is the percentage of confirmed accumulated cases in a steady state ($I = D = A = R = T = 0$). To make sure our model coincides with the findings of [Bommer and Vollmer \(2020\)](#), we expect ϕ to be slightly above of the estimated 27.32% as a steady state is not reached yet and the proportion of detected cases increases over time, i.e., we constrain $\phi \in [0.3, 0.45]$.
- The percentage of asymptomatic disease progressions was estimated at 43% in a population screening study in Iceland ([Gudbjartsson et al., 2020](#)), at 43.2% in a comprehensive testing of the whole municipality of Vo', Italy ([Lavezzo et al., 2020](#)) and at 17.9% ([Mizumoto, Kagaya, Zarebski, & Chowell, 2020](#)) in a study regarding the cruise ship Diamond Princess. To ensure that our model has a comparable ratio, we add the constraint $\frac{\lambda}{\lambda + \zeta} \in [0.18, 0.43]$ to the optimization problem.
- The (base) reproduction rate in the beginning of March was estimated as approximately 3 ([an der Heiden & Hamouda, 2020](#)). Thus, for the parameters α_{\max} , γ_{\max} with no active countermeasures we require $R_0(\alpha_{\max}, \gamma_{\max}) \in [2.5, 3.5]$ where $R_0(\alpha, \gamma)$ is given by (see [Giordano et al. \(2020\)](#) for details)

$$R_0(\alpha, \gamma) = \frac{1}{\zeta + \lambda} \left(\alpha + \frac{1}{\theta + \mu + \kappa} \left(\gamma \zeta + \frac{\beta \theta \zeta}{\mu + \kappa} \right) \right). \quad (2)$$

- The median of the incubation time is 5–6 days ([Li et al., 2020](#); [Linton et al., 2020](#); [World Health Organization \(WHO\), 2020](#)), which we identify as the half life period a person is in the state I , i.e. $\log(2)/(\lambda + \zeta) \in [5, 6]$. Further, the median time from onset of symptoms until intensive care is 10–11 days ([Wang et al., 2020](#); [Yang et al., 2020](#)). Hence, we constrain the half life period of a ailing or recognized individuals to $\log(2)/(\kappa + \mu) \in [10, 11]$.

In the state H of (1) the confirmed recovered cases are not distinguished from the undetected ones, thus we define the number of confirmed recovered cases as H_c , with $H_c(0) = 0$ and $\dot{H}_c = \lambda D +$

Table 1

The parameters of the model for Germany.

$\alpha_{\min} = 0.0422$	$u_1 = 0.5816$	$\zeta = 0.0790$	$\mu_1 = 0.0080$	$\mu_2 = 0.0050$
$\alpha_{\max} = 0.3614$	$u_2 = 0.7062$	$\lambda = 0.0596$	$\sigma_1 = 0.0370$	$\sigma_2 = 0.0552$
$\gamma_{\min} = 0.0422$	$\beta = 0.0084$	$\kappa = 0.0563$	$\tau_1 = 0.0159$	$\tau_2 = 0.0242$
$\gamma_{\max} = 0.3614$	$\theta = 0.1981$	$I_0 = 500/N_{\text{total}}$	$A_0 = 304/N_{\text{total}}$	

$\kappa R + \sigma_1 T_1 + \sigma_2 T_2$ and further the number of confirmed accumulated cases as $C = D + R + T + H_c + E$ in order to match the data \tilde{H}_c and \tilde{C} . Considering the COVID-19 patients in intensive care \tilde{T}_2 , a natural choice would be to identify them with threatened state T , however, all deaths in the model have been in T before, but in reality only half to a third of the deaths happens in ICU ([DIVI-IntensivRegister, 2020](#); [Dong et al., 2020a](#)). Hence, as the patients in ICU are only a part of T , we split T into T_1 and T_2 , where T_2 represents the number of people in intensive care and T_1 are all otherwise threatened COVID-19 cases. We assume that there are no transitions from T_1 to T_2 and vice versa, such that $T = T_1 + T_2$ can be modeled as

$$\dot{T}_1 = \mu_1(A + R) - (\sigma_1 + \tau_1)T_1 \quad (3a)$$

$$\dot{T}_2 = \mu_2(A + R) - (\sigma_1(T_2) + \tau_2(T_2)) \quad (3b)$$

with $\mu_1 + \mu_2 = \mu$. This more complex model with T_1 and T_2 will be approximated with a model of the form described in Section 2.1 as sketched in Fig. 2. Further, we define H_2 (E_2) to be the numbers of people that recovered (died) from T_2 .

Finally, we perform the parameter optimization by solving a least squares problem via CasADi ([Andersson, Gillis, Horn, Rawlings, & Diehl, 2019](#)) to fit C , E , H_c , T_2 , E_2 , H_2 to the data \tilde{C} , \tilde{E} , \tilde{H}_c , \tilde{T}_2 , \tilde{E}_2 , \tilde{H}_2 . The best fitting parameters are given in Table 1 and the resulting fit is shown in Fig. 3. Many of the constraints listed above are active at the optimal set of parameters, e.g., $\alpha = \gamma$, which is not surprising since we use the constraints to keep the parameters in a realistic range without further regularization.

This fit further enables us to specify the full model state of today $x(53) =: x_0$, which will be used in the following sections as the initial condition where $t = 53$ corresponds to April 21

$$x_0 = \frac{1}{N_{\text{total}}} \begin{bmatrix} 82\,636\,256 & 20\,581 & 0 & 8\,041 & 41\,931 & 11\,469 & 276\,911 & 4\,810 \end{bmatrix}^T. \quad (4)$$

Please note that the model is quite sensitive to changes in the parameters and one obtains quite different parameter values if, e.g., the estimated range of unknown cases or the percentage of asymptomatic cases deviate from the assumptions.

2.3. Modeling of the mortality rate

It has been recognized as a key difficulty in handling the COVID-19 pandemic that the virus is highly contagious, thus infecting large numbers of individuals. In addition, since many elderly and ill people require hospitalization and/or intensive care ([Verity et al., 2020](#)), large waves of infections can quickly exceed the capacities of local health care systems ([Grasselli, Pesenti, & Cecconi, 2020](#)). Hence, ensuring that health care resources are sufficient is a key issue in handling the outbreak ([IHME COVID-19 health service utilization forecasting team, C. J. L. Murray, 2020](#)), given that an overwhelmed health care system even correlates positively with the mortality rate ([Ji, Ma, Peppelenbosch, & Pan, 2020](#)).

In this section, we describe how the mortality and recovery rates $\tau(T)$ and $\sigma(T)$ in (1) depend on the number of threatened patients T . The basic idea is that they are constant as long as the health care system's capacity is not at its limit, and the mortality rate τ (the recovery rate σ) increases (decreases) significantly if it is overwhelmed. According to [DIVI-IntensivRegister \(2020\)](#), there are (on April 21) 2 908




Fig. 2. The threatened state T is split up in ICU cases T_2 and non-ICU cases T_1 which is later approximated using a lumped model.

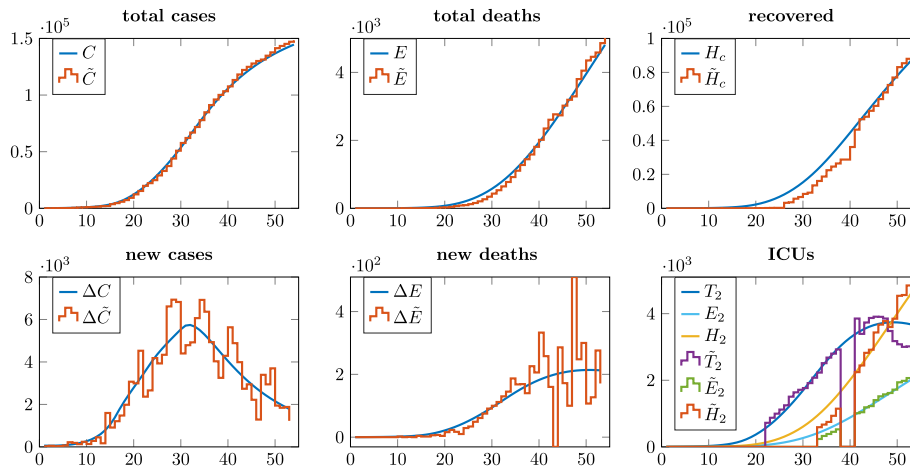


Fig. 3. Simulation of the model (1) with the estimated parameters in Table 1 scaled by N_{total} compared to the actual data (DIVI-IntensivRegister, 2020; Dong et al., 2020a). The horizontal axis represents the time in days, where $t = 1$ is February 28 and $t = 53$ is April 21.

COVID-19 patients in an ICU and 12 623 ICU spots are available. Hence, the overall ICU capacity currently available for COVID-19 patients is $2908 + 12\,623 = 15\,531$, and we define the relative ICU capacity as $T_{\text{ICU}} = \frac{15\,531}{N_{\text{total}}}$, where $N_{\text{total}} = 8.3 \cdot 10^7$. We consider a constant value of T_{ICU} for simplicity, although it is likely that it will further increase in the future.

We assume that the mortality rate increases if the number of individuals requiring treatment in an ICU exceeds T_{ICU} , i.e., if $T_2 > T_{\text{ICU}}$, with T_2 as in 2. More precisely, we assume that if a patient requiring intensive care does not receive it, then the patient deceases (i.e., for such patients, the mortality rate increases and the recovery rate is zero). According to data of deceased individuals from Italy, those who were not admitted to an ICU deceased in median within 4 days (COVID-19 Surveillance Group, 2020). Hence, we model those individuals in T_2 which are not admitted to an ICU via decaying first order dynamics with a half-life period of 4 days, i.e, the corresponding time constant τ_{crit} satisfies $e^{-4\tau_{\text{crit}}} = 0.5$, thus leading to $\tau_{\text{crit}} = 0.173$.

In a first approximation $T_2 \approx \frac{\mu_2}{\mu} T$ and hence we only modify the mortality rate τ in case that $\frac{\mu_2}{\mu} T > T_{\text{ICU}}$. In the model (1), $\tau(T)$ and $\sigma(T)$ only occur jointly with T , which leads us to the following formula for $\tau(T)T$ and $\sigma(T)T$:

$$\tau(T)T = \frac{\mu_1}{\mu} \tau_1 T + \max \left\{ \frac{\mu_2}{\mu} \tau_2 T, \tau_2 T_{\text{ICU}} + \tau_{\text{crit}} \left(\frac{\mu_2}{\mu} T - T_{\text{ICU}} \right) \right\}, \quad (5a)$$

$$\sigma(T)T = \frac{\mu_1}{\mu} \sigma_1 T + \sigma_2 \min \left\{ \frac{\mu_2}{\mu} T, T_{\text{ICU}} \right\}. \quad (5b)$$

If $\frac{\mu_2}{\mu} T \leq T_{\text{ICU}}$, then (5) implies $\tau(T)T = (\frac{\mu_1}{\mu} \tau_1 + \frac{\mu_2}{\mu} \tau_2)T$ and $\sigma(T)T = (\frac{\mu_1}{\mu} \sigma_1 + \frac{\mu_2}{\mu} \sigma_2)T$, i.e., a simple lumped model is recovered as long as the ICU capacity is not exceeded. If however $\frac{\mu_2}{\mu} T > T_{\text{ICU}}$, then the mortality rate increases to τ_{crit} for those $\frac{\mu_2}{\mu} T - T_{\text{ICU}}$ patients which require intensive care but do not receive it. Similarly, for this fraction, the recovery rate is set to zero implicitly in (5b). The individuals $T_1 = \frac{\mu_1}{\mu} T$ not receiving intensive care are not affected by this mechanism.

Clearly, the modified rates in (5) are just a simple approximation of the effect that the mortality rate increases if hospitals are overwhelmed. This modification in the model is crucial when studying the effect of loosening quarantine measures and corresponding optimal policies, as done in the remainder of this paper. Since (fortunately) the German health care system has not been overwhelmed to this date, there are no quantitative data to validate the above modification and in particular, the exact value of τ_{crit} . Nevertheless, the refinement is confirmed qualitatively by experiences in other countries (Grasselli et al., 2020; IHME COVID-19 health service utilization forecasting team, C. J. L. Murray, 2020; Ji et al., 2020). Moreover, even a substantial change of τ_{crit} has little effect on the overall dynamics since it only affects the exact number of fatalities. In particular, changing τ_{crit} does not lead to a qualitative change in an optimal policy to control the outbreak as long as τ_{crit} is sufficiently larger than τ_2 and it is possible not to exceed the ICU capacity.

3. Open-loop optimal control of the COVID-19 outbreak

In this section, we discuss different policies that can be considered to keep the number of fatalities due to COVID-19 low while at the same time also impose as little constraints as possible on the public. The most significant degree of freedom currently is certainly influencing the infection rates α and γ . Measures for influencing the infection rates include hygienic measures, face masks, and different nuances of distancing policies, up to a mandated lockdown. Therefore, we define u as introduced in Section 2.2 as our input, representing distancing policies or other measures that have a direct influence on the infection rates α and γ . We model this influence via

$$\alpha(t) = \alpha_{\text{max}} + (\alpha_{\text{min}} - \alpha_{\text{max}})u(t) \quad (6a)$$

$$\gamma(t) = \gamma_{\text{max}} + (\gamma_{\text{min}} - \gamma_{\text{max}})u(t), \quad (6b)$$

where a value of $u = 1$ hence represents the policies in Germany as of mid April (lockdown) and $u = 0$ represents no social distancing or other measures (i.e. corresponding to infection rates as in the beginning of March). Furthermore, we assume that the policies affecting the infection rates α, γ (i.e. u) stay constant for at least one week and can only be changed every seven days. In the first subsection, we will introduce different baseline policies which can give insights into the effects of different inputs u and which will serve as a comparison to the optimal controller in the following subsection. More specifically, these baseline strategies will be used to define an upper bound on the social distancing measures that the optimal control in Section 3.2 and later on the feedback strategies in Sections 4 and 5 can employ to minimize the fatalities.

Test capacity

In addition to varying the infection rates α and γ , another degree of freedom to influence the model (1) lies in adapting the testing policy. Testing individuals on COVID-19 is represented in the current model by parameters θ and ϵ for symptomatic and asymptomatic individuals, respectively. In the following, we assume that only a fixed number of tests can be carried out every day. If we wish to only test symptomatic individuals, this includes both symptomatic individuals infected with COVID-19 (i.e., members of the state A) and individuals suffering from other illnesses with similar symptoms. In Robert Koch Institut (2017), the Robert Koch Institute estimates numbers on influenza-like illnesses (ILI) in Germany. While the numbers show clear seasonal differences, approximately 1.3% of the population become newly infected with ILI on average per week, and approximately 37% of them see a doctor (an indication for more severe symptoms). Moreover, influenza symptoms usually last 4–5 days leaving us with an approximate average of $p_{\text{sick}} = 0.3\%$ of the population showing significant influenza-like symptoms at an arbitrary time of the year. When testing asymptomatic individuals, this includes infected persons without symptoms but also any other individual not known to be infected or healed and not showing any symptoms (i.e. $S + I + H - H_c - p_{\text{sick}}$, where H_c are the confirmed healed cases, compare Section 2.2). The total amount of resources used for testing is then captured by the following cost:

$$c_{\text{test}}(\epsilon, \theta, A, S, I, H - H_c) = \epsilon(S + I + H - H_c - p_{\text{sick}}) + \theta \cdot (A + p_{\text{sick}}) + A\mu. \quad (7)$$

Denoting the parameter θ from Table 1 as θ_n , we assume a fixed bound $\bar{c} > 0$ on the amount of resources for testing c_{test} and that θ_n corresponds to the nominal value for $A = 0$, i.e., $\bar{c} := c_{\text{test}}(0, \theta_n, 0)$. In the following, we assume that the current policy with respect to testing stays in place: all available tests are used on a daily basis for as many symptomatic people as tests are available. Then the testing policy used throughout this section reads $\epsilon(t) = 0$ (as is current practice) and

$$\theta(t) = (\theta_n \cdot p_{\text{sick}} - \mu A(t)) / (p_{\text{sick}} + A(t)). \quad (8)$$

Note that this also implies that throughout this paper the state $D \equiv 0$. The allocation of tests (with the possibility of also saving test resources for later) can also be modeled as control inputs. However, in the present model the effects of temporarily saving tests (under the current resource constraints) are negligible compared to the effects of changes in the infection rates. Increasing the overall test capacity or improving the choice of test subjects (e.g. with tracing of cases), which corresponds to increasing values of θ_n/\bar{c} , on the other hand, can potentially improve the evolution of the pandemic significantly, since detected individuals are less contagious than undetected ones. However, increasing test capacities or better allocated testing (especially with regard to ϵ , i.e. tracing of also asymptomatic infections) is at the current stage not included in our consideration but could be addressed in future work with the presented model by choosing $\epsilon \neq 0$ and making θ_n/\bar{c} an increasing and time-varying variable.

Control goal

Given the introduced control input u , different control goals can be formulated. One such goal could be to obtain herd immunity. Herd immunity corresponds to the only stable equilibrium given no social distancing measures (i.e. with $\alpha_{\text{max}}, \gamma_{\text{max}}$) and requires a large part of the society to be immune. More precisely, herd immunity is reached if $S < S^*$, where (Giordano et al., 2020) provide a formula for calculating S^* (see Section 3.1.1 for more details). Given our model, we can now calculate the minimum time that is needed to reach herd immunity. For this, we assume that we can choose a policy that utilizes the full health care capacity at all times. With $T \equiv \frac{\mu}{\mu_2} T_{\text{ICU}}, \dot{T} \equiv 0, \dot{A} + \dot{R} \equiv 0, \dot{I} + \dot{D} \equiv 0$, S decreases each day by $\frac{(\zeta + \lambda)(\mu + \kappa)(\sigma + \tau)}{\zeta \mu} \frac{\mu}{\mu_2} T_{\text{ICU}}$. Hence,

$$t_{\text{herd}} = \frac{\zeta \mu_2 (1 - S^*)}{(\zeta + \lambda)(\mu + \kappa)(\sigma + \tau) T_{\text{ICU}}}$$

gives a lower bound on the time required to reach herd immunity without exceeding the health care capacity given the introduced model. The herein identified model parameters yield a time span t_{herd} of more than six years. A stable steady state in the absence of a vaccine (i.e. herd immunity) can hence only be obtained either after many years or by overstraining the health care system and a corresponding significant rise in the number of fatalities. Therefore, our ongoing assumption throughout this section is that prior to herd immunity, a vaccine will be available and we assume the availability of the vaccine in approximately two years. Our goal is thus to find an optimal policy minimizing the number of fatalities for the next two years while imposing as little constraints as possible on the public and the economy.

In the next subsection, we simulate and discuss the following policies:

1. Keeping the social distancing measures in place (or even increasing the measures) until the virus is eradicated in Germany
2. Slowly (or more aggressively) loosening the distancing measures without overwhelming the health care capacities (while possibly risking a second wave).

In fact, the presented baseline policies are similar to the policies suggested by the German ‘‘Helmholtz-Initiative’’ in Helmholtz-Initiative ‘Systemische Epidemiologische Analyse der Covid-19-Epidemie’ (2020). We will discuss our conclusions in comparison to theirs at the end of the section.

In Section 3.2, we will then improve these baseline policies by applying optimal control techniques and we will discuss the importance and significance of the improvements.

3.1. Introducing different baseline policies

3.1.1. Consistent lockdown

In the following, we argue that a consistent lockdown strategy necessitates strong lockdown measures over a long time horizon to fully eradicate the virus as otherwise, dropping the social distancing measures too early leads to a second outbreak wave.

Based on the SIDARTHE model fitted to the German outbreak, described in Section 2.1, we simulate how long we would need to remain in lockdown and simply wait for the virus to disappear. We define the disappearance of the virus as follows: If – most probably – there is less than one active contagious case, i.e., $I + D + A + R + T < 0.5/N_{\text{total}}$, the virus is eradicated. It takes 305 days, which is almost one year, until this condition is fulfilled and clearly the economical and psychological damage caused by a lockdown period this long is excessive such that staying and waiting in lockdown is not an option. With even stricter measures, such as $\alpha_3 = 0.8\alpha_3, \gamma_3 = 0.8\gamma_3$, we could only marginally accelerate this process to 288 days while increasing social distancing is costly, cf. the cost function in Section 3.2. Note that the equilibrium attained under the above lockdown policy is an unstable one that is not robust to uncertainties. In particular, if only

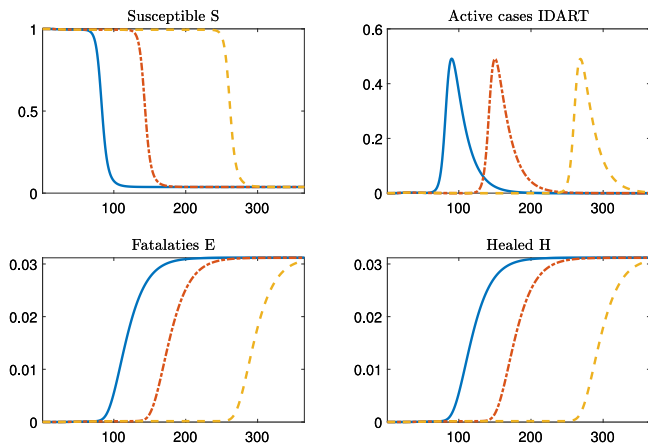


Fig. 4. Simulation of the model (1) for a consistent lockdown of different lengths $T_{l,i}$ with days on the horizontal axis ($T_{l,1} = 0$ - blue, solid; $T_{l,2} = 50$ - red, dash-dotted; $T_{l,3} = 150$ - yellow, dashed).

one person remains infected when the measures are suspended they could cause a new outbreak. Also, the virus may be reimported from other countries or humans might be reinfected by an interim host.

Next, we simulate the following three scenarios in all of which the German population is kept in lockdown for a predefined period of time, followed by no measures at all. The only difference is the length of the lockdown period. In the first scenario, the measures are abolished immediately (April 21). The second one keeps the current strict measures for an additional 50 days. The third variant simulates an even longer lockdown period, ending after 150 days counting from April 21.

In Fig. 4, we compare the three scenarios. We clearly see that in all three cases the number of currently infected people $I + D + A + R + T$ rises drastically a few days after the measures are removed independent of how long the lockdown persisted before. In any case, we experience a second outbreak wave. Staying longer in lockdown slightly delays the following peak of the share of active cases $I + D + A + R + T$, yet the peak amplitude is almost the same in all three scenarios.

This behavior can be explained as follows. If there is no one who currently has the virus, i.e., $I_{eq} = D_{eq} = A_{eq} = R_{eq} = T_{eq} = 0$, such that $S_{eq} + H_{eq} + E_{eq} = 1$, an equilibrium point is attained. The stability of the equilibrium point depends on the value of S_{eq} and the model parameters. In Giordano et al. (2020), the authors show that the IDART subsystem is asymptotically stable if and only if $S_{eq} < S^*$, where S_{eq} is the susceptible state at equilibrium for a given initial condition x_0 and the corresponding parameters, especially for α and γ that are actively adjusted according to an underlying policy. The value of S^* follows from the stability analysis of the linearized IDART subsystem and has the following structure with respect to α and γ

$$S^* = \frac{a_1}{\alpha a_2 + \gamma a_3 + a_4}, \tag{9}$$

where $a_i, i = 1 \dots 4$ are constants, see Giordano et al. (2020) for details and the definition of S^* . Note further that the commonly stated base reproduction rate (2) is directly linked to the value S^* via $R_0 = 1/S^*$. The stability of an equilibrium that depends on the parameters changes once we adjust α and γ . For strict measures ($\alpha_{min}, \gamma_{min}$), the value S^* is high ($S^* = 2.242$), such that a stable equilibrium is attained for any S . This means that only a small number of people is infected by the virus before the equilibrium is attained. With no measures ($\alpha_{max}, \gamma_{max}$), a stable equilibrium requires the share of susceptible people to be smaller than $S^* = 0.292$, i.e., herd immunity.

We can hence conclude that if $S_{eq}(x_0, \alpha_3, \gamma_3) > S^*(\alpha_0, \gamma_0)$, the equilibrium attained during lockdown is unstable with no measures. This means, there inevitably is a second outbreak wave once the lockdown

ends. For the fitted model of the German outbreak, $S_{eq}(x_0, \alpha_3, \gamma_3) = 0.9956$ is attained after the first wave. Hence, at least another 70.4% of the German population get infected in the second wave before a stable equilibrium is attained. Altogether, this leaves us with the following conclusion of two possible outcomes when choosing a consistent lockdown strategy:

- Strong lockdown measures over long time horizons have to be taken to eradicate the virus in Germany. However, this takes a big toll on the public and any infected person, e.g. from abroad, could spark a second wave at any point.
- Any lockdown strategy that does not fully wipe out the virus inevitably yields a second outbreak wave once all measures are suspended.

3.1.2. Iterative loosening of the distancing policies

As argued in the previous subsection, keeping the lockdown policy strict can never lead to a stable equilibrium when ending the lockdown, no matter how long it did take place before all measures were suspended. Hence, many countries are now discussing or have even already started to loosen the lockdown in very small steps. Indeed, experts consulting the German government ("Nationale Akademie der Wissenschaften Leopoldina") have recently published their recommendations concerning a possible strategy for loosening the lockdown gradually in small steps (Leopoldina, 2020).

They name the following conditions for loosening the lockdown in small steps:

- The number of new infections remains at a low value.
- The capacity of the health care system must not be exceeded.
- Precautions (such as hygienic measures, face masks, distancing) remain in force.

In the following, we try to translate these recommendations into a policy for our simplified model to first analyze the results and second, to use this as a baseline policy for the optimizer in the following subsection. We implement the conditions presented above via the following policy strategy:

- u can only be decreased if, over the last n_{stab} days, the number of newly infected persons (i.e. $S(t-1) - S(t)$) is decreasing and u has not been increased
- u can only be decreased if less than X_{lower} of the ICUs are occupied
- The decrease in u can only be a small decrease at a time and therefore, the interval between $u_{max} = 1$ (lockdown) and $u_{min} = 0$ (no measures) is divided into n_{steps} equidistant steps.

Additionally, we add that u will be increased again (with the same step size as the decrease) if more than X_{upper} of the ICUs are occupied and no decrease in the amount of necessary ICU is witnessed. This policy results in four 'tuning parameters' of the policy: $n_{stab}, X_{lower}, X_{upper}$ and n_{steps} . In fact, it turns out that the outcome of the simulation is not very sensitive to the tuning parameters of the policy, but can be tuned to be slightly more careful or more aggressive. In the following subsection, we choose two different sets of parameters as baseline policies for the optimal control approach.

3.2. Optimal control strategy

In this section, we contrast the baseline policy from Section 3.1.2 with an optimal control policy, under idealized assumptions (exact model and state measurement). The purpose of this section is twofold:

- Understand how an optimal policy differs qualitatively from the baseline policies.
- Quantify the loss of performance (in terms of increased fatalities and/or unnecessary social policy u) resulting from using a suboptimal baseline policy. The degree of freedom is the input

$u \in [0, 1]$ affecting the social policy and we consider the fact that the policy can only be changed every $T_s = 7$ days.

Multi-objective optimal control problem

In the following, we consider the problem only for the next $N = 100$ weeks, assuming that thereafter a vaccine might be developed that would ideally prevent (almost all) further fatalities in the future. The overall control problem can be seen as a multi-objective optimal control problem, where we wish to simultaneously minimize the number of fatalities E and the societal and economical cost of the social policy measures, which will be measured by the function $c_{\text{policy}}(u) = 1/\alpha(u)$. We point out that due to the parametrization (6) this cost also inherently considers the infection rate γ . This cost function is such that the social cost of achieving an arbitrarily small infection rate α grows unbounded, while for large values of u incremental differences are less relevant. In order to suitably characterize an *optimal* solution to this multi-objective problem we use the baseline policies in Section 3.1. The resulting optimal control problem is given by (11) below, which will be explained in the following. In particular, our goal is to find an input policy that minimizes the number of accumulated fatalities, while using less resources than the baseline policy in terms of accumulated social impact of c_{policy} (c.f. (11d)). We point out that similar “stabilization” problems subject to resource constraints for the control of epidemic outbreaks can be found in the literature, also using a fractional cost c_{policy} , compare e.g. Köhler, Enyioha, and Allgöwer (2018), Preciado, Zargham, Enyioha, Jadbabaie, and Pappas (2014).

When minimizing the number of accumulated fatalities, it is important to consider not only the extinct individuals $E(N \cdot T_s)$ at the end of the two year horizon, but to account as well for the part of the already infected individuals that will decrease *after* the two year horizon. The reason for this is, while the availability of a vaccine at the end of the horizon might prevent future infections, it cannot cure already infected people. Hence, if we do not account for the inevitable fatalities among the individuals infected at the end of the prediction horizon, the optimal controller does not take any efforts to keep them low and as a result a lot of people would die shortly after the two year horizon. Therefore, we propose an optimization objective that includes all past and inevitable future fatalities. Based on the model (1), we know that a total of $\frac{\zeta}{\zeta+\lambda}(I+D)$ of the infected people $I+D$ will develop symptoms in the future and further that a total of $\frac{\mu}{\mu+\kappa}\left(\frac{\zeta}{\zeta+\lambda}(I+D)+A+R\right)$ will become threatened. Thus, assuming the capacity T_{ICU} is not exceeded afterwards, i.e., setting constant values $\tau = \tau(0) = \frac{\mu_1}{\mu}\tau_1 + \frac{\mu_2}{\mu}\tau_2$ and $\sigma = \sigma(0) = \frac{\mu_1}{\mu}\sigma_1 + \frac{\mu_2}{\mu}\sigma_2$, the amount of inevitable fatalities is exactly given by

$$F = E + \frac{\tau}{\tau + \sigma} \left(\frac{\mu}{\mu + \kappa} \left(\frac{\zeta}{\zeta + \lambda} (I + D) + A + R \right) + T \right). \tag{10}$$

Hence, given a baseline solution u^b, X^b from Section 3.1.2, the corresponding optimal control problem reads as follows:

$$\min_{u(\cdot)} F(N \cdot T_s) \tag{11a}$$

$$0 \leq u(k \cdot T_s) \in [0, 1] \tag{11b}$$

$$k = 0, \dots, N - 1 \tag{11c}$$

$$\sum_{k=0}^{N-1} c_{\text{policy}}(u(k \cdot T_s)) \leq \sum_{k=0}^{N-1} c_{\text{policy}}(u^b(k \cdot T_s)). \tag{11d}$$

Since we only change the policy every week, the index k in (11) corresponds to weeks and $F(N \cdot T_s)$ corresponds to the objective function in (10), where the states result from simulating the system (1) with the parameters and the initial condition from Section 2 and the input $u(\cdot)$ over k -weeks. Condition (11d) ensures that the encountered social cost is smaller than the cost of the baseline policy. This optimal control problem is such that the baseline policy $u = u^b$ is a feasible solution and thus the resulting fatalities $F(N \cdot T_s)$ will always be




Fig. 5. Optimal control strategy (blue, solid), baseline policy (red, dashed) and ICU capacity (dotted, black).

lower than that of the baseline policy. We point out that it is possible to consider a more restrictive transient constraint on the policy cost instead of (11d), which is discussed in detail in Appendix B. The optimal control problem (11) can be formulated as a nonlinear program (NLP) and is subsequently solved using CasADi (Andersson et al., 2019).

Numerical results

For comparison and to implement the constraint (11d), we use the baseline policy from Section 3.1.2 with $X_{\text{lower}} = 0.4, X_{\text{upper}} = 0.7, n_{\text{steps}} = 14, n_{\text{stab}} = 14$, which overall is rather cautious and does not exceed the ICU capacity. The corresponding results for the baseline policy and the optimal control strategy can be seen in Fig. 5. Although the optimal control input yields initially (first 100 days) a slightly larger number of infected individuals and thus slightly more fatalities in the first 200 days, the number of infected individuals is subsequently significantly lower and the overall number of fatalities is reduced to only 26%. The optimal controller allows for a smooth increase of the infection rate α , while keeping the number of threatened individuals (T) consistently below the corresponding value of the baseline policy (after the first 200 days), thus yielding a small number of fatalities. The rising number of infected individuals (I) at the end results from the finite-horizon and will be considered later in more detail.

We also consider a second baseline policy using $X_{\text{lower}} = 0.6, X_{\text{upper}} = 0.85, n_{\text{steps}} = 12, n_{\text{stab}} = 14$, which slightly relaxes the social policy, but also exceeds the ICU capacity. The result is shown in Fig. 6. We can see that in comparison to this second baseline policy the optimal policy significantly reduces the number of fatalities to only 39%. The optimal strategy is a lot more cautious in reducing the social policy, while the baseline is more aggressive and goes back and forth between increasing and decreasing α , resulting in significant violations of the ICU capacity. Furthermore, the simple baseline policy results in a second wave as the restrictions are loosened too quickly, while the optimal strategy slowly but steadily increases α after the first 200 days, and thereby avoids a second wave.

In both examples, a further observation should be highlighted. After an initial phase of containing the outbreak, the measures are slowly but steadily relaxed until a larger release at the end of the horizon. Similar behavior can be observed for many optimal control problems with finite horizons and is commonly referred to as “turnpike” behavior, which goes back to Dorfman, Samuelson, and Solow (1987). An explanation for this is that the consequences of decisions taken at later points in




Fig. 6. More aggressive baseline: Optimal control strategy (blue, solid), baseline policy (red, dashed) and ICU capacity (dotted, black).

time mainly occur after the end of the horizon, such that a more aggressive policy towards the end is optimal, when only considering the finite two year horizon. Of course, one would not implement the “leaving arc” if the development of a vaccine would not be finalized after two years, since implementing such a policy may lead to an uncontrollable increase of infections towards the end of the time horizon in case that the model is inaccurate and should thus be avoided in practice. In Appendix A, we therefore discuss how adding “terminal constraints” to the optimization problem can prevent this turnpike behavior of the optimal solution, at the price of an increasing number of fatalities.

3.3. Discussion

If we compare the results in Section 3.2 with the consistent full lockdown from Section 3.1.1, we can see that it is possible to appropriately increase α without exceeding the ICU capacities, while the consistent full lockdown strategy would require a lockdown that takes approximately a year to be effective. Hence, while a consistent lockdown can effectively minimize the number of deaths, this strategy is only viable in case this lockdown can be prolonged over the corresponding time horizon, unless a vaccine is developed earlier. On the other hand, both the optimal controller and the baseline controller allow for a significant relaxation of the lockdown (on average a doubling of α), without significantly increasing the number of fatalities.

In comparison to the baseline policy suggested in Section 3.1.2, the optimal control policy results in a slower but smooth loosening of the distancing policies. Without increasing the social cost over the full time horizon, this optimal policy avoids any violation of the maximum ICU capacity and hence results in a significantly smaller fatality rate. We point out that the resulting optimal policy of slowly increasing α is qualitatively similar to the resulting optimal policy in Kantner and Koprućki (2020), albeit for a different control goal. It can be seen that an initially “stronger” lockdown (i.e., a smaller value of α) over a longer time period with subsequent loosening leads to a better handling of the pandemic, compared to repeated tightening and loosening of distancing measures. Moreover, a smooth and monotone loosening of distancing policies is also desirable from an economic aspect since repeated lockdowns after interim-periods of relaxed distancing guidelines may be even more damaging to the economy, compared to an initially longer lockdown.

Comparing our results with the proposed scenarios by the Helmholtz Association (Helmholtz-Initiative ‘Systemische Epidemiologische Analyse der Covid-19-Epidemie’, 2020), we find that we agree that the goal of herd immunity without overwhelming the health care capacity would require years. Our results further agree with Helmholtz-Initiative ‘Systemische Epidemiologische Analyse der Covid-19-Epidemie’ (2020) that the contact restrictions can only be loosened slowly if the health care capacity must not be overwhelmed. However, since the authors in Helmholtz-Initiative ‘Systemische Epidemiologische Analyse der Covid-19-Epidemie’ (2020) do not consider the availability of a vaccine, their conclusion is keeping or even increasing the lockdown until the number of infected persons is small and all infections can be traced efficiently and effectively via strategically allocated (and increased) testing. With the assumption of a vaccine within the next two years, we argue that a slow and smooth loosening of the lockdown does not lead to many more fatalities while decreasing the social and economic cost significantly according to our model. Concurrently, increased and strategically better allocated testing, e.g., via a contact tracing mobile app (Oliver et al., 2020), is of course highly beneficial and greatly advisable to improve the performance (even if this was not accounted for in our model).

To summarize, it seems possible to reduce the current restrictions and thus allow α to increase without exceeding the ICU capacity. Furthermore, optimized policies can significantly improve the outcome (in terms of fewer deaths and/or less social restrictions). However, the result is highly sensitive w.r.t. the change in the infection rate, while an accurate control of the infection rate α (e.g. $\pm 5\%$) through governmental policies seems difficult/unrealistic. In the next section, we will therefore deal with these issues by formulating a robust control strategy that takes uncertainty in our COVID-19 model into account and uses feedback based on uncertain state information.

4. Optimal feedback control of the COVID-19 outbreak

Section 3.2 shows that an optimal control policy can significantly reduce the number of fatalities compared to a baseline policy that allows for iterative loosening of social distancing measures. This optimal control policy is computed by optimizing over all possible policies to find the one minimizing the number of fatalities predicted by the model Eqs. (1) without using stronger shutdown measures than the baseline. Hence, the policy proposed in Section 3.2 strongly relies on the accuracy of the model identified in Section 2 and thus may fail to effectively control the outbreak in case of a model mismatch. However, such a model mismatch is inevitable in practice, especially since the model itself is a simplification of a much more complex reality and the identification outlined in Section 2.2 strongly depends on the (sparse) available data and the additional prior knowledge based e.g. on further studies concerning COVID-19, which also provide only estimates. In addition, the optimal control policy relies on exact knowledge of all states and on the assumption that values for α and γ can be exactly imposed up to arbitrary precision via social distancing measures, both of which are unrealistic assumptions when applying the policy in practice.

In this section, we show how *online measurements* can be utilized via *feedback* to effectively and robustly control the German COVID-19 outbreak in the presence of uncertain parameters. More precisely, we illustrate that the optimal open-loop policy of Section 3.2 may lead to poor performance when applied to validation models with a different set of parameters, although these validation models result from adjusting only one prior assumption in the identification and still fit the past data well. On the other hand, we show that a model predictive control (MPC) feedback strategy, based on repeatedly computing an open-loop policy for the nominal model from Section 2, is inherently robust w.r.t. model inaccuracies and successfully handles the outbreak.

Basic idea

At each time step $k = 1, \dots, N$, where k corresponds to weeks and $N = 100$ as in Section 3.2, we solve the optimization problem (11) over the time horizon k, \dots, N using the current measurements as initial condition at week k . Then, we apply the computed optimal policy over one week before solving the problem for the new measurements again. In this way, since the initial conditions in the optimal control problem are updated, a feedback mechanism is included as is standard in MPC (Rawlings, Mayne, & Diehl, 2017). As a result, the prediction horizon $N - k$ of the optimization problem is shrinking with each time step k , such that it never exceeds the considered total time horizon of $N = 100$ weeks, after which we assume the availability of a vaccine. Hence, since the constraint (11d) needs to hold over the whole time horizon N , we replace it by

$$\sum_{j=k}^{N-1} c_{\text{policy}}(u(j \cdot T_s)) \leq c_{\text{policy}}^1 - \sum_{j=0}^{k-1} c_{\text{policy}}(u(j \cdot T_s)), \tag{12}$$

with $c_{\text{policy}}^1 = \sum_{k=0}^{N-1} c_{\text{policy}}(u^b(k \cdot T_s))$ being the cost of the first baseline policy in Section 3.1.2. As a second modification, we adapt the bound on the social distancing cost online, depending on the predicted states, as is detailed in the following.

Online adaptation of social policy constraint

In Section 3.2, we proposed an open-loop optimal control strategy, where the inputs were the infection rates α and γ . Loosely speaking, the control goal was to achieve a minimum number of fatalities without imposing stronger social distancing measures than a simple baseline policy (compare (11d)). Since this constraint heavily depends on the model to which the baseline is applied, a realistic setting with imperfect model knowledge should allow to adapt the constraints on the policy online in case that the nominal model is overly optimistic or pessimistic. Instead of simply requiring that the cost of the MPC-based feedback cannot exceed c_{policy}^1 , we increase the maximum cost in case that the predicted number of patients requiring intensive care lies above 90% of the maximum capacity T_{ICU} at least once during the horizon, and we decrease it in case that the number consistently lies below 10% of T_{ICU} . This adaption is natural, as in reality one would increase the efforts to contain the outbreak if the current measures are insufficient, and on the other hand, the population cannot be expected to accept strict measures when there are only few (severe) cases across the country. Therefore, the maximum cost in week k , denoted by $c_b(k)$, varies online with k and is initialized as $c_b(0) = c_{\text{policy}}^1$. The amount by which we change c_b online is $\pm \Delta_u \frac{N-k}{N}$, where $\Delta_u = \frac{1}{\alpha_{\text{min}}} - \frac{1}{\alpha_{\text{max}}}$ with α_{min} and α_{max} as in Section 2.2. If, for instance, the model predicts large numbers of future ICU patients, then the cost bound c_b is increased by the difference between the minimum and the maximum social distancing cost, scaled by the remaining time horizon. This increase corresponds to the social cost of an additional week in full lockdown, scaled by the remaining time horizon via the factor $\frac{N-k}{N}$.

MPC-based feedback strategy

The proposed MPC-based feedback strategy is summarized in Algorithm 1. In the algorithm, $T(j \cdot T_s \mid k \cdot T_s)$ denotes the number of threatened individuals at time $j \cdot T_s$, predicted by the optimal solution of (11) at time $k \cdot T_s$. Essentially, the algorithm repeatedly applies the open-loop optimal control policy of Section 3.2 with the key difference that, at time k , all past measurements $j = 1, \dots, k$ are used in the optimization problem, thus including an online feedback. In addition, in Step 3 of the algorithm, the social policy constraints is adapted as described above.

Validation

To assess the improved robustness of Algorithm 1 compared to open-loop optimal control, we produce two validation models. More

Algorithm 1. MPC-based feedback strategy

1. Given state measurements up to time k , solve the following problem

$$\min_{u(\cdot)} F(N \cdot T_s) \tag{13a}$$

$$0 \leq u(j \cdot T_s) \in [0, 1] \tag{13b}$$

$$j = k, \dots, N - 1 \tag{13c}$$

$$\sum_{j=k}^{N-1} c_{\text{policy}}(u(j \cdot T_s)) \leq c_b(k) - \sum_{j=0}^{k-1} c_{\text{policy}}(u(j \cdot T_s)), \tag{13d}$$

with the state-dependent cost F as in (10), based on simulating the model (1) over the remaining horizon $N - k$ subject to the input u , starting at the current measured state at time k .

2. Apply the optimal policy $u^*(k \cdot T_s)$ for the next $T_s = 7$ days.

3. Update the social policy cost as

$$c_b(k+1) = \begin{cases} c_b(k) + \Delta_u \frac{N-k}{N} & \text{if } \frac{\mu_2}{\mu} \max_{j \in [k, N]} T(j \cdot T_s \mid k \cdot T_s) \geq 0.9T_{\text{ICU}} \\ c_b(k) - \Delta_u \frac{N-k}{N} & \text{if } \frac{\mu_2}{\mu} \max_{j \in [k, N]} T(j \cdot T_s \mid k \cdot T_s) \leq 0.1T_{\text{ICU}} \\ c_b(k) & \text{otherwise} \end{cases}$$

where $\Delta_u = \frac{1}{\alpha_{\text{min}}} - \frac{1}{\alpha_{\text{max}}}$.

4. Set $k = k + 1$.

precisely, we identify two new sets of parameters A and B by proceeding exactly as in Section 2.2 with the only difference that we change the prior assumption that the stationary ratio of confirmed COVID-19 cases is in the interval [0.3, 0.45]. Instead, we assume that this value is in [0.3, 0.6] for set A and in [0.3, 0.4] for set B. When performing parameter identification based on these modified prior assumptions, we also obtain sets of parameters that can accurately explain the existing past data on COVID-19 cases in Germany. However, the resulting models have different dynamics and different reproduction rates for the same lockdown policy. Increasing the above ratio as in parameter set A decreases the number of infected and undetected individuals resulting in a higher reproduction rate to explain the same amount of confirmed cases. Hence, if an open-loop policy based on the nominal model (i.e., with parameters described in Section 2.2) is applied to the validation model with parameters A, then the number of infections and thus the number of fatalities increases significantly.

To illustrate this effect, we apply the open-loop optimal control policy based on the model with parameters as in Section 2.2 to the new models with parameter sets A and B. The control effort of this policy, i.e., the amount of social distancing, is constrained as in (11d) by the cost of the baseline policy when applied to the nominal model identified in Section 2.2. The results for the model with parameters A can be seen in Fig. 7. Since this validation model has a higher reproduction rate for similar inputs as explained above, the number of fatalities after $N = 100$ weeks increases significantly compared to the simulations in Section 3. This is due to the fact that the open-loop policy is only computed once, at the beginning of the time horizon, and is then applied over the whole time span of two years without any online adaption based on new measurements. Therefore, it cannot handle the model mismatch and thus has a significantly worse performance. In addition, Fig. 7 shows the evolution under the proposed MPC-based feedback, which leads to a significantly lower number of fatalities compared to the open-loop policy. We point out that the feedback (partially) compensates the fact that the control action is computed based on the nominal model parameters from Section 2.2, which differ significantly from parameter set A. Due to the larger number of infected individuals, the maximum social cost c_b is increased at multiple time

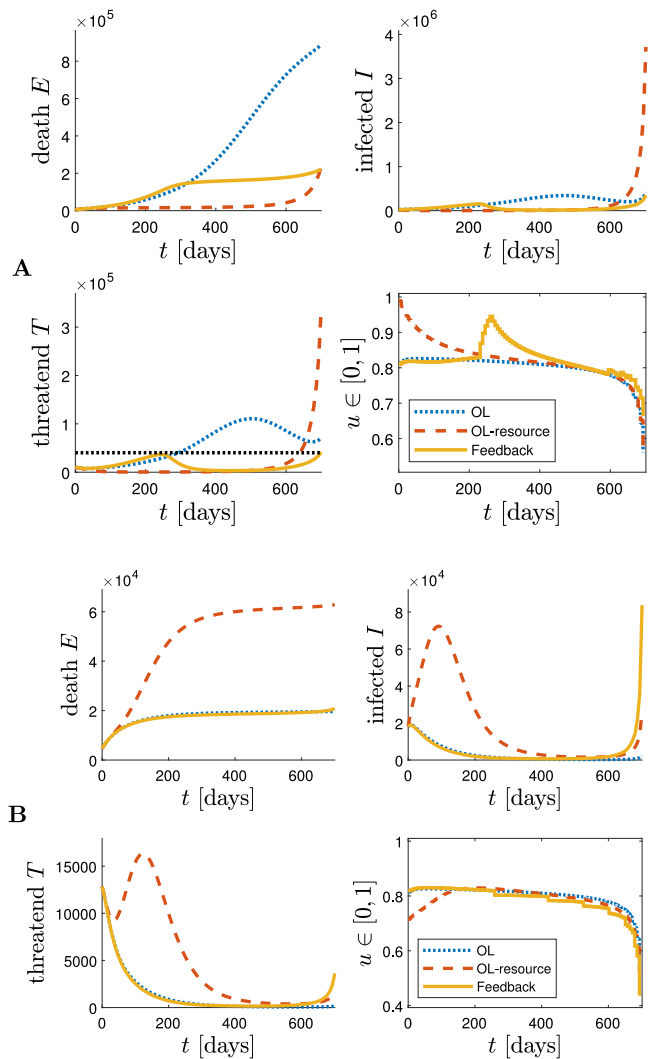


Fig. 7. Different policies designed based on the nominal model identified in Section 2.2, applied to the validation model with parameter set A (top) and parameter set B (bottom): Open-loop optimal control strategy ('OL', blue dotted), Open-loop optimal control strategy with the same amount of resources as the MPC-based feedback ('OL-resource', red dashed), MPC-based feedback strategy as in Algorithm 1 (yellow solid), and ICU capacity (black dotted).

steps, which is indicated by the step-like increases of the input. Finally, an open-loop optimal control policy is computed which is allowed to use the same amount of resources as the MPC-based feedback (in hindsight), i.e., the adapted social and economical cost $c_b(N)$. While this policy performs better than the initial open-loop policy with fewer resources, it leads to a similar number of fatalities at the end of the horizon compared to the feedback controller. However, the number of threatened patients is very large at time $k = N$, which would lead to a significant increase in fatalities after the considered time period, even if a vaccine is available.

Fig. 7 also shows the same comparison for the model with parameter set B. In this case, since the reproduction rate is lower, the open-loop optimal policy leads to fewer fatalities than in Section 3.2. The MPC-based feedback leads to almost identical performance, but it reduces the cost budget at several time instants, i.e., it can handle the outbreak similarly well but with significantly lower social and economic costs. When restricting the budget of the open-loop optimal policy to the one of the feedback strategy, i.e., $c_b(N)$, then it leads to a dramatic increase in fatalities.

Table 2

Uncertainties in the states.

$I \pm 50\%$	$D \pm 1\%$	$A \pm 20\%$	$R \pm 1\%$	$H \pm 50\%$	$E \pm 1\%$
$T \pm (\frac{\mu_1}{\mu} \cdot 1\% + \frac{\mu_2}{\mu} \cdot 5\%)$					

To conclude, the above discussion reveals that a combination of open-loop optimal control with feedback is inherently robust in the sense that it effectively controls the German COVID-19 outbreak even if the employed model is inaccurate. When comparing the result to an open-loop strategy, then the MPC-based feedback strategy can dramatically decrease the number of fatalities or the necessary amount of social distancing, respectively. Such robustness is an important property for applying any control strategy in a real-world scenario, where accurate model knowledge is rarely available. In the next section, we propose a more systematic robust MPC approach which explicitly takes model inaccuracies as well as uncertain state measurements and control inputs into account in order to safely and cautiously control the COVID-19 pandemic.

5. Robust and optimal feedback control of the COVID-19 outbreak

While the MPC-based feedback policy proposed in Section 4 is significantly more successful in handling the outbreak compared to a simple open-loop policy, it relies on the assumption that exact measurements of the state in (1) are available at each time step. In this section, we consider a more realistic scenario of uncertain measurements in terms of biased state estimates, and we analyze the impact on the closed-loop operation. In particular, we develop a robust MPC-based feedback strategy using interval arithmetic that takes the uncertainty into account during the predictions and thus leads to a safer policy minimizing the number of fatalities.

Biased state measurements

In the following, we consider the case where at each day k instead of the true state $x(k)$ we only obtain an estimated state $\hat{x}(k)$, which is subject to an additional bias. In Table 2, we summarize the uncertainties in the states. For individuals in states D and R , the disease COVID-19 was detected by tests. Hence, their values are well known, nevertheless, we assume that they can slightly differ from the true states by $\pm 1\%$, as there might be cases on the borderline between D and R that are hard to assign to either of the states. The number of people in ICUs is well documented. However, the state T contains not only patients in ICUs ($T_{ICU} \pm 1\%$) but also other infected members of the risk group ($T_2 \pm 5\%$), cf. Section 2.1, such that the uncertainty we use is $\pm(\frac{\mu_1}{\mu} \cdot 1\% + \frac{\mu_2}{\mu} \cdot 5\%)$. We assume that the number of deaths is certain by $\pm 1\%$ as it includes some people that died of different causes. As the undetected cases can by definition not be measured, they must be estimated using random sampling or strategies like (Bommer & Vollmer, 2020). Therefore, the states I and A are much less certain, especially without symptoms ($I \pm 50\%$, $A \pm 20\%$). Recovering from the disease is a resulting state from both rather certain states, D and R , and highly uncertain states, I and A such that overall it is uncertain itself ($H \pm 50\%$). The uncertainty of the state of susceptible persons S results from the other states: $\hat{S} = \hat{x}_1 = 1 - \sum_{i=2}^8 \hat{x}_i$.

It is possible to directly use this biased state estimate $\hat{x}(k)$ in Algorithm 1 and compensate the bias through the inherent robustness in the feedback implementation. In the following, we derive an alternative *robust* formulation that explicitly considers the uncertainty in the prediction.

Interval predictions

First, given a biased state estimate $\hat{x}(k)$ and known bounds on the bias (Table 2), it is possible to compute interval bounds $\underline{x}(k)$, $\bar{x}(k)$ such that the true state is guaranteed to lie in that interval, i.e., $x_i(k) \in$

$[\underline{x}_i(k), \bar{x}_i(k)]$. The following formulation will predict the interval bounds \underline{x}_i and \bar{x}_i instead of using some nominal prediction. This methods is similar to interval arithmetic employed in robust MPC (Limon, Bravo, Alamo, & Camacho, 2005) and the robust moment enclosure for an SEIV epidemic model in Watkins, Nowzari, and Pappas (2019). Using the fact that the system is positive (x_i and all the parameters are positive), it is possible to derive an interval prediction of the form

$$\dot{\underline{x}} = \underline{f}(\underline{x}, \bar{x}, u), \tag{14a}$$

$$\dot{\bar{x}} = \bar{f}(\underline{x}, \bar{x}, u), \tag{14b}$$

which ensures that $x(t) \in [\underline{x}(0), \bar{x}(0)]$ implies $x(t) \in [\underline{x}(t), \bar{x}(t)]$ for all $t \geq 0$, given suitable bounds on the uncertain parameters in the system model (1). The detailed derivation of the interval prediction model (14) can be found in Appendix C (more precisely, Eqs. (C.4)). Since deriving reliable bounds on all parameters in the model (1) is rather difficult or unnecessarily complex, we only focus on the uncertainty associated with the infection rate α . In particular, we consider an uncertainty of $\pm 5\%$ on the infection rate α . Thereby, we explicitly consider the problem that the infection rate cannot be precisely specified via social distancing measures. We will see later in the simulations that, although we do not account for all possible mismatches in the prediction model, we nevertheless obtain the desired properties in closed loop.

Given this interval prediction model, the proposed robust formulation now predicts an interval for the different state variables and minimizes the worst-case number of fatalities \bar{F} based on $\bar{x}(N \cdot T_s)$. The overall procedure is summarized in Algorithm 2.

Algorithm 2. Robust MPC strategy using interval arithmetic

1. Given biased state estimate $\hat{x}(k \cdot T_s)$, compute set $[\underline{x}(k \cdot T_s), \bar{x}(k \cdot T_s)]$.
2. Solve the following problem

$$\min_{u(\cdot)} \bar{F}(N \cdot T_s) \tag{15a}$$

$$0 \leq u(j \cdot T_s) \in [0, 1] \tag{15b}$$

$$j = k, \dots, N - 1 \tag{15c}$$

$$\sum_{j=k}^{N-1} c_{\text{policy}}(u(j \cdot T_s)) \leq c_b(k) - \sum_{j=0}^{k-1} c_{\text{policy}}(u(j \cdot T_s)), \tag{15d}$$

with \bar{F} based on (10) using \bar{x} , which results from the interval predictions of the model (14) over the remaining horizon $N - k$ subject to the input u , starting at the current set estimate $[\underline{x}(k \cdot T_s), \bar{x}(k \cdot T_s)]$.

3. Apply the optimal policy $u^*(k \cdot T_s)$ for the next $T_s = 7$ days.
4. Update the social policy cost as in Algorithm 1 using \bar{T} instead of T .
5. Set $k = k + 1$.

Numerical results

In the following simulations, we consider the extreme case where the number of estimated infected or previously infected individuals (I, D, A, R, T, H, E) is underestimated. The results for the robust MPC and the nominal MPC in comparison to open-loop optimal control strategies for the two validation parameter sets A and B (compare Section 4) can be seen in Fig. 8. Due to the worst-case prediction in the robust MPC, at $t = 0$ the robust MPC already increases the resources two times, for both parameter sets, such that at $t = T_s$ the predictions satisfy $\frac{\mu_1}{\mu_2} \bar{T}(k \cdot T_s) \leq 0.9 T_{ICU}$.

In the simulation with the model based on parameter set A, we can directly see that both the nominal MPC and the robust MPC reduce the number of fatalities compared to an open-loop optimal control strategy. Furthermore, if we compare the robust MPC and the nominal MPC, we can see that after $t = 140$ days the nominal MPC

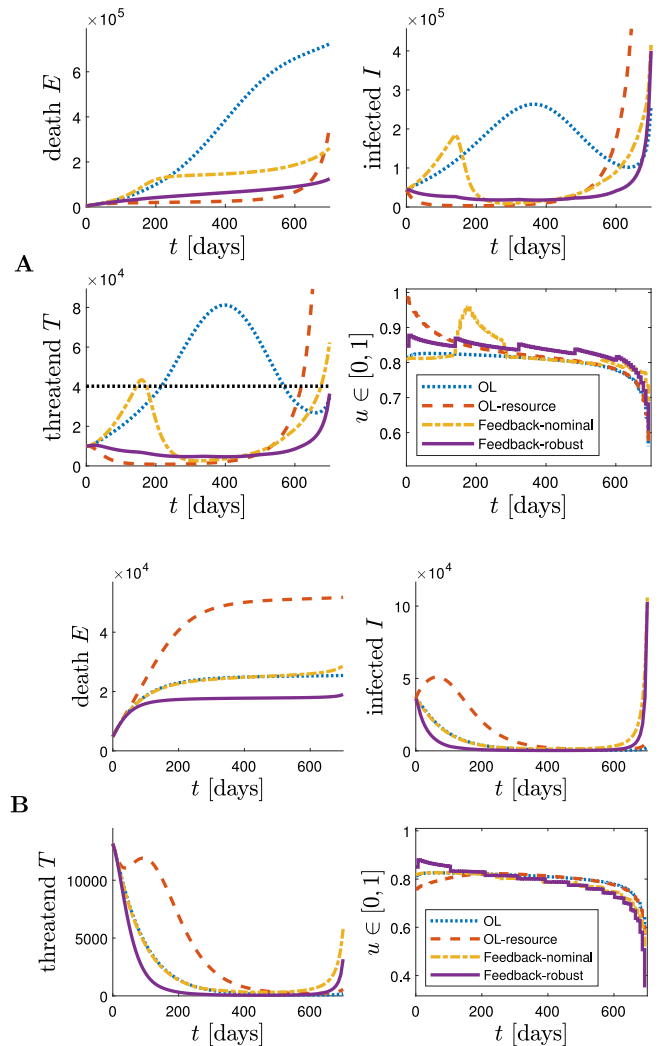


Fig. 8. Different policies designed based on the nominal model identified in Section 2.2, applied to the validation model with parameter set A (top) and parameter set B (bottom) with biased state measurements: Open-loop optimal control strategy (‘OL’, blue dotted), Open-loop optimal control strategy with the same amount of resources as the robust MPC-based feedback (‘OL-resource’, red dashed), nominal MPC-based feedback strategy as in Algorithm 1 with biased estimated state $\hat{x}(k)$ (yellow, dash-dot), robust MPC-based feedback strategy using Algorithm 2 (purple, solid), and ICU capacity (black dotted).

implementation realizes that the spread is worse than initially assumed. This leads to a strong increase in social measures u at $t = 140$. With the robust formulation, u decreases almost monotonically. Furthermore, the nominal implementation results in twice the number of fatalities, while the applied resources c_{policy} over the two year horizon differ by less than Δ_u , which corresponds to one week of lockdown. This indicates that the robust MPC, planning cautiously from the beginning, exploits its resources more efficiently by imposing stricter social distancing measures early on, which results to be beneficial in the long run.

For the second parameter set B, we can see that the worst-case robust formulation uses initially an unnecessarily high control effort. Nevertheless, overall the applied resources c_{policy} differ by less than Δ_u compared to the nominal MPC, while the number of fatalities is still reduced by 33%. In comparison to the open-loop policies, both MPC policies require less or equally restrictive policy measures u , while the number of fatalities are significantly reduced.

6. Conclusions and discussion

In the following, we summarize our findings on a high-level and highlight the main take-away messages:

- Our results in Section 3 confirm the conclusions in Meyer-Hermann et al. (2020) that neither eradication of the virus nor herd immunity without the availability of a vaccine are viable solutions to handle the current COVID-19 outbreak.
- Applying an optimizer to the mathematical model describing the outbreak, one can significantly reduce the number of fatalities without increasing the costs associated to decreasing the infection rate (social distancing policies, closing schools, etc.), compare Section 3.2.
- Since the proposed model can never exactly predict the COVID-19 pandemic, applying a nominal optimal policy introduces unnecessary conservatism, at best, up to posing a great danger (i.e. overwhelming the health care capacities risking high mortality rates). Therefore, our findings in Section 4 support (Meyer-Hermann et al., 2020) by showing that any policy to control the COVID-19 outbreak successfully has to be an *adaptive* strategy. This means we need to constantly measure, monitor and estimate the current numbers and adapt our policy accordingly, i.e., feedback is necessary for reliably handling the outbreak.
- If we already a priori take into account that our model includes mismatches and that all measured and estimated numbers are not exact and can have a bias, we can further improve the outcome, as shown in Section 5. More specifically, we developed a robust MPC-based feedback strategy using interval arithmetic. The application of feedback without the robust description of the considered model can lead to intermediate increases in the number of new infections necessitating another period of lockdown. On the contrary, a robust feedback strategy can take these model mismatches and other uncertainties into account and is hence able to avoid such behavior, thus significantly reducing the number of fatalities.
- When looking at the qualitative results the robust MPC-based feedback offers, one can see that, accounting for the instability and uncertainty of the spread of the virus, the controller suggests a rather strict policy at the beginning and only then allows for a gradual increase in the infection rate. Keeping this loosening slow at the beginning shows a beneficial effect in the long run. This qualitative result of the robust MPC underpins also the German policy and reaction to the outbreak of COVID-19 in Germany where initially strong measures (what we here refer to as lockdown) were applied. Only very recently the German government started to loosen these measures slowly and gradually.

There are also influences on the course of the outbreak that were not taken into account in the present paper but which are important in an overall strategy towards the spread of COVID-19 (e.g. increasing testing capacities, tracking of infections, as well as investigating which measures lead to the desired infection rate). However, controlling the infection rate is certainly one of the key factors and hence, this paper contributes towards mitigating the spread of COVID-19 under manageable societal and economic costs. We hope that the proposed feedback strategies inspire further investigations in this direction and offer qualitative and high-level insights that underpin the current policies or strategy papers.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.




Fig. A.9. Cautious baseline: Optimal control strategy with terminal constraints (blue, solid), baseline policy (red, dashed) and ICU capacity (dotted, black).

Appendix A. Optimal control formulation using terminal constraints

In order to avoid artifacts of considering a finite-horizon problem (e.g. a lot of infected people at the end of the horizon), an alternative to considering the modified cost function F from (10) is the inclusion of additional terminal constraints for the contagious population $IDART = (I, D, A, R, T) \in \mathbb{R}^5$. In particular, we require that at the end of the control horizon N , the number of contagious individuals in each category (I, D, A, R, T) should be smaller than the corresponding number from the baseline policy (c.f. (A.1a)). In addition, at the end of the horizon the number of contagious individuals should be non-increasing, which is implemented as (A.1b).

Hence, in the following, we replace the cost F by the number of fatalities E , and we add the following constraints to the optimal control problem (11) from Section 3.2:

$$IDART(N \cdot T_s) \leq IDART^b(N \cdot T_s), \tag{A.1a}$$

$$IDART(N \cdot T_s) \leq IDART((N - 1) \cdot T_s). \tag{A.1b}$$

Again, the index k in (11) corresponds to weeks and the states $IDART(k \cdot T_s)$ correspond to the result of simulating the system (1) with the parameters and initial condition from Section 2. These terminal conditions (A.1a)–(A.1b) (which should be interpreted element-wise) ensure that the final state after the finite horizon N is “better” than the baseline solution (c.f. (A.1a)) and the outbreak can be contained (c.f. (A.1b)). The simulation results with the two baseline policies shown in Figs. A.9 and A.10 demonstrate that the terminal constraints indeed effectively prevent the turnpike behavior. However, the additional constraints also lead to a slight increase in the number of fatalities.

Appendix B. Alternative average constraint formulation

We wish to briefly mention a stronger restriction on the societal cost of the optimal control strategy. In particular, instead of only restricting the cost over the considered horizon of $N = 100$ weeks, a stronger property is to ensure that at any time t , the previously accumulated policy cost is smaller than the corresponding cost of the baseline policy. This can be done by replacing condition (11d) with the following

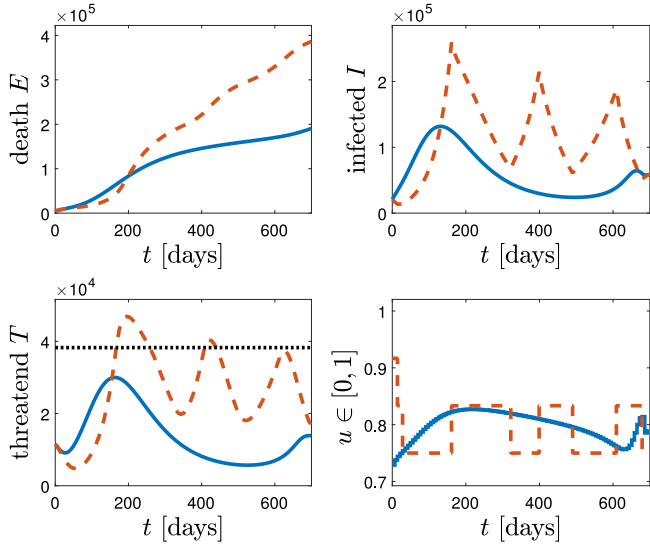


Fig. A.10. Aggressive baseline: Optimal control strategy with terminal constraints (blue, solid), baseline policy (red, dashed) and ICU capacity (dotted, black).

transient constraint:

$$\sum_{k=0}^{i-1} c_{\text{policy}}(u_1(k)) \leq \sum_{k=0}^{i-1} c_{\text{policy}}(u_1^b(k)), \quad (\text{B.1})$$

$$i = 1, \dots, N. \quad (\text{B.2})$$

The corresponding results for both baselines considered in Section 3.2 can be seen in Fig. B.11. In this case the number of fatalities are reduced by 33% and 37%, respectively. Thus, also for this more restrictive setting, the optimal controller can significantly reduce the number of fatalities. In addition, in the comparison to the more aggressive baseline we also see that *early* measures are absolutely crucial, since the two policies are essentially equivalent in the period from 110 days–150 days, i.e., including the critical time period where the ICU capacity is exceeded, but differ significantly over 30 weeks prior to the violation of the ICU capacity.

Appendix C. Interval predictions

In the following, we derive the dynamics of the interval predictions \bar{f}, \underline{f} used in (14). Note that the following property holds for any scalars $a \in [a, \bar{a}]$, $b \in [b, \bar{b}]$:

$$a \cdot b \in [ab, \bar{a}\bar{b}], \quad (\text{C.1})$$

if $\underline{a}, \underline{b} \geq 0$. Furthermore, to ensure that $x(t) \in [\underline{x}(t), \bar{x}(t)]$ given $x(0) \in [\underline{x}(0), \bar{x}(0)]$, we only require

$$\dot{\bar{x}}_i \geq \dot{x}_i \quad \text{if } \bar{x}_i = x_i, \quad (\text{C.2})$$

and similarly

$$\dot{\underline{x}}_i \leq \dot{x}_i \quad \text{if } \underline{x}_i = x_i, \quad (\text{C.3})$$

for all $i = 1, \dots, 8$. Essentially, we use the property (C.1) together with the fact that x_i and the parameters are positive to ensure that (C.2) and (C.3) hold in order to derive the differential equations for the interval. More precisely, assuming that every parameter is uncertain in some bounds (e.g. $\beta \in [\underline{\beta}, \bar{\beta}]$) yields the following $2 \cdot 8$ ODEs:

$$\dot{\bar{S}} = -\bar{S}(\alpha I + \beta D + \gamma A + \beta R), \quad (\text{C.4a})$$

$$\dot{\underline{S}} = -\underline{S}(\bar{\alpha} I + \bar{\beta} D + \bar{\gamma} A + \bar{\beta} R), \quad (\text{C.4b})$$

$$\dot{\bar{I}} = \bar{S}(\bar{\alpha} I + \bar{\beta} D + \bar{\gamma} A + \bar{\beta} R) - (\underline{\epsilon} + \underline{\zeta} + \underline{\lambda}) \bar{I}, \quad (\text{C.4c})$$

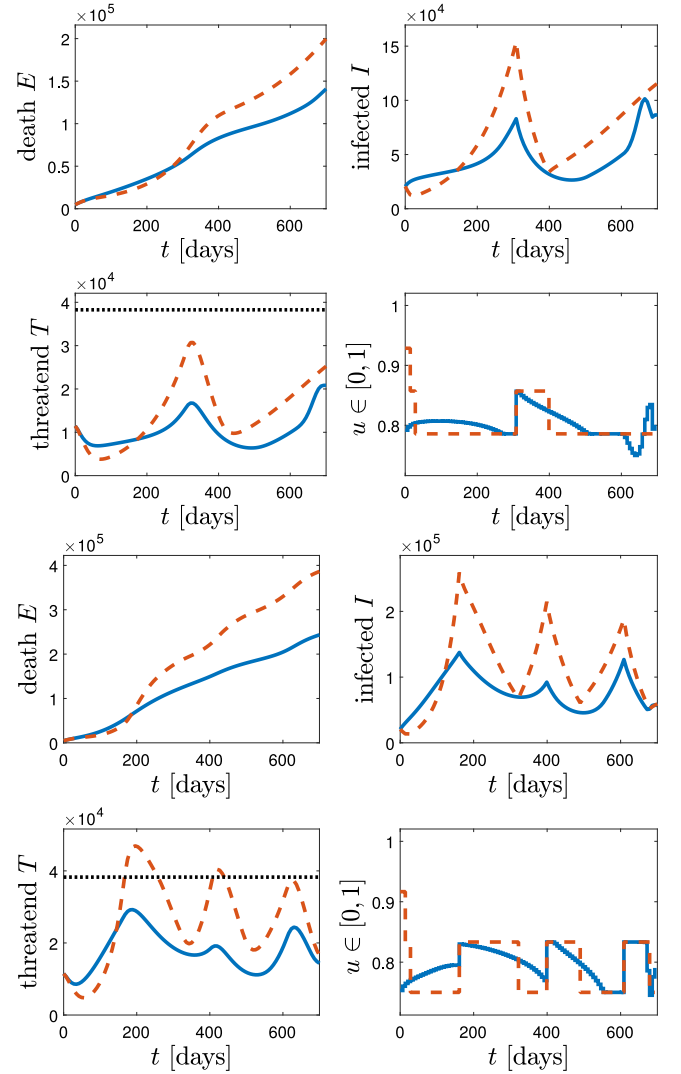


Fig. B.11. Transient constraint on social policy: Top: cautious baseline; Bottom: aggressive baseline. Optimal control strategy (blue, solid), baseline policy (red, dashed) and ICU capacity (dotted, black).

$$\dot{\bar{I}} = \bar{S}(\alpha I + \beta D + \gamma A + \beta R) - (\bar{\epsilon} + \bar{\zeta} + \bar{\lambda}) \bar{I}, \quad (\text{C.4d})$$

$$\dot{\underline{I}} = \underline{S}(\bar{\alpha} I + \bar{\beta} D + \bar{\gamma} A + \bar{\beta} R) - (\underline{\epsilon} + \underline{\zeta} + \underline{\lambda}) \underline{I}, \quad (\text{C.4e})$$

$$\dot{\bar{D}} = \bar{\epsilon} \bar{I} - (\bar{\zeta} + \bar{\lambda}) \bar{D}, \quad (\text{C.4f})$$

$$\dot{\underline{D}} = \underline{\epsilon} \underline{I} - (\underline{\zeta} + \underline{\lambda}) \underline{D}, \quad (\text{C.4g})$$

$$\dot{\bar{A}} = \bar{\zeta} \bar{I} - (\bar{\theta} + \bar{\mu} + \bar{\kappa}) \bar{A}, \quad (\text{C.4h})$$

$$\dot{\underline{A}} = \underline{\zeta} \underline{I} - (\underline{\theta} + \underline{\mu} + \underline{\kappa}) \underline{A}, \quad (\text{C.4i})$$

$$\dot{\bar{R}} = \bar{\zeta} \bar{D} + \bar{\theta} \bar{A} - (\bar{\mu} + \bar{\kappa}) \bar{R}, \quad (\text{C.4j})$$

$$\dot{\underline{R}} = \underline{\zeta} \underline{D} + \underline{\theta} \underline{A} - (\underline{\mu} + \underline{\kappa}) \underline{R}, \quad (\text{C.4k})$$

$$\dot{\bar{T}} = \bar{\mu} \bar{A} + \bar{\mu} \bar{R} - (\bar{\sigma}(\bar{T}) + \bar{\tau}(\bar{T})) \bar{T}, \quad (\text{C.4l})$$

$$\dot{\underline{T}} = \underline{\mu} \underline{A} + \underline{\mu} \underline{R} - (\underline{\sigma}(\underline{T}) + \underline{\tau}(\underline{T})) \underline{T}, \quad (\text{C.4m})$$

$$\dot{\bar{H}} = \bar{\lambda} \bar{I} + \bar{\lambda} \bar{D} + \bar{\kappa} \bar{A} + \bar{\kappa} \bar{R} + \bar{\sigma}(\bar{T}) \bar{T}, \quad (\text{C.4n})$$

$$\dot{\underline{H}} = \underline{\lambda} \underline{I} + \underline{\lambda} \underline{D} + \underline{\kappa} \underline{A} + \underline{\kappa} \underline{R} + \underline{\sigma}(\underline{T}) \underline{T}, \quad (\text{C.4o})$$

$$\dot{\bar{E}} = \bar{\tau}(\bar{T}) \bar{T}, \quad (\text{C.4p})$$

$$\dot{\underline{E}} = \underline{\tau}(\underline{T}) \underline{T}. \quad (\text{C.4q})$$

Since these dynamics only correspond to possibly conservative overapproximations, we can use $\sum_{i=1}^8 x_i = 1$ to possibly improve the resulting

bounds for \underline{S} using the following projections: $\bar{S} \leq 1 - \sum_{i=2}^8 x_i$ and $\underline{S} \geq 1 - \sum_{i=2}^8 \bar{x}_i$. In principle it would also be possible to directly set \bar{S}, \underline{S} using the other states \bar{x}_i, x_i instead of simulating (C.4a)–(C.4b), but this may not necessarily ensure $\bar{S} \leq 1$ and $\underline{S} \geq 0$.

References

- Alleman, T., Torfs, E., & Nopens, I. (2020). COVID-19: from model prediction to model predictive control. Available online https://biomath.ugent.be/sites/default/files/2020-04/Alleman_et_al_v2.pdf. (Accessed 30 April 2020).
- an der Heiden, M., & Hamouda, O. (2020). Schätzung der aktuellen Entwicklung der SARS-CoV-2-Epidemie in Deutschland - Nowcasting. *Epidemiologisches Bulletin*, 2020(17), 10–15.
- Andersson, J. A., Gillis, J., Horn, G., Rawlings, J. B., & Diehl, M. (2019). CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1), 1–36.
- Barbarossa, M. V., Fuhrmann, J., Heidecke, J., Varma, H. V., Castelletti, N., Meinke, J. H., et al. (2020). A first study on the impact of current and future control measures on the spread of COVID-19 in Germany. medRxiv preprint <http://dx.doi.org/10.1101/2020.04.08.20056630>.
- Bin, M., Cheung, P., Crisostomi, E., Ferraro, P., Myant, C., Parisini, T., et al. (2020). On fast multi-shot epidemic interventions for post lock-down mitigation: Implications for simple COVID-19 models. arXiv preprint [arXiv:2003.09930](https://arxiv.org/abs/2003.09930).
- Bommer, C., & Vollmer, S. (2020). Average detection rate of SARS-CoV-2 infections is estimated around nine percent. Available online <https://www.uni-goettingen.de/en/606540.html>. (Accessed 15 April 2020).
- Casella, F. (2021). Can the COVID-19 epidemic be controlled on the basis of daily test reports? *IEEE Control Systems Letters*, 5(3), 1079–1084.
- COVID-19 Surveillance Group (2020). Characteristics of COVID-19 patients dying in Italy, Report based on available data on March 30th, 2020. Available online: https://www.epicentro.iss.it/coronavirus/bollettino/Report-COVID-2019_30_marzo_eng.pdf. (Accessed 16 April 2020).
- Dehning, J., Zierenberg, J., Spitzner, F. P., Wibral, M., Neto, J. P., Wilczek, M., et al. (2020). Inferring COVID-19 spreading rates and potential change points for case number forecasts. arXiv preprint [arXiv:2004.01105](https://arxiv.org/abs/2004.01105).
- DIVI-IntensivRegister (2020). Tagesreport 27.03.2020 – 21.04.2020. Available online: <https://www.divi.de/divi-intensivregister-tagesreport-archiv>.
- Dijdjou-Demasie, R., Michalakos, Y., Choisy, M., Sofonea, M. T., & Alizon, S. (2020). Optimal COVID-19 epidemic control until vaccine deployment. medRxiv preprint <http://dx.doi.org/10.1101/2020.04.02.20049189>.
- Dong, E., Du, H., & Gardner, L. (2020a). 2019 Novel Coronavirus COVID-19 (2019-nCoV) data repository by Johns Hopkins CSSE. <https://github.com/CSSEGISandData/COVID-19>. (Accessed 22 April 2020).
- Dong, E., Du, H., & Gardner, L. (2020b). An interactive web-based dashboard to track COVID-19 in real time. *The Lancet Infectious Diseases*, 20(5), 533–534.
- Dorfman, R., Samuelson, P. A., & Solow, R. M. (1987). *Linear programming and economic analysis*. Courier Corporation.
- German, R., Djanatliev, A., Maile, L., Bazan, P., & Hackstein, H. (2020). Modeling exit strategies from COVID-19 lockdown with a focus on antibody tests. medRxiv preprint <http://dx.doi.org/10.1101/2020.04.14.20063750>.
- Giordano, G., Blanchini, F., Bruno, R., Colaneri, P., Filippini, A. D., Matteo, A. D., et al. (2020). Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. *Nature Medicine*, 26, 855–860.
- Grasselli, G., Pesenti, A., & Cecconi, M. (2020). Critical care utilization for the COVID-19 outbreak in Lombardy, Italy: early experience and forecast during an emergency response. *Journal of the American Medical Association*, 323(16), 1545–1546.
- Gudbjartsson, D. F., Helgason, A., Jonsson, H., Magnusson, O. T., Melsted, P., Norddahl, G. L., et al. (2020). Spread of SARS-CoV-2 in the Icelandic population. *New England Journal of Medicine*, 382, 2302–2315.
- He, X., Lau, E. H. Y., Wu, P., Deng, X., Wang, J., Hao, X., Lau, Y. C., et al. (2020). Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine*, 26(5), 672–675.
- Helmholtz-Initiative 'Systemische Epidemiologische Analyse der Covid-19-Epidemie' (2020). Stellungnahme der Helmholtz-Initiative 'Systemische Epidemiologische Analyse der COVID-19-Epidemie'. https://www.helmholtz.de/fileadmin/user_upload/01_forschung/Helmholtz-COVID-19-Papier_02.pdf. (Accessed 15 April 2020).
- IHME COVID-19 health service utilization forecasting team, C. J. L. Murray (2020). Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months. medRxiv preprint <http://dx.doi.org/10.1101/2020.03.27.20043752>.
- Jakob Simmank, J. W. (2020). So testet Deutschland. ZEIT online, <https://www.zeit.de/wissen/gesundheit/2020-03/coronatests-deutschland-coronavirus-covid-19-who-pandemie#wer-soll-in-deutschland-getestet-werden>. (Accessed 15 April 2020).
- Ji, Y., Ma, Z., Peppelenbosch, M. P., & Pan, Q. (2020). Potential association between COVID-19 mortality and health-care resource availability. *The Lancet Global Health*, 8(4), Article e480.
- Kantner, M., & Koprucki, T. (2020). Beyond just "flattening the curve": Optimal control of epidemics with purely non-pharmaceutical interventions. *Journal of Mathematics in Industry*, 10(1), 1–23.
- Kermack, W., & McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London*, (115), 700–721.
- Kissler, S., Tedijanto, C., Lipsitch, M., & Grad, Y. H. (2020). Social distancing strategies for curbing the COVID-19 epidemic. medRxiv preprint <http://dx.doi.org/10.1101/2020.03.22.20041079>.
- Köhler, J., Enyioha, C., & Allgöwer, F. (2018). Dynamic resource allocation to control epidemic outbreaks - A model predictive control approach. In *Proc. Annual American control conference (ACC)* (pp. 1546–1551).
- Lavezzo, E., Franchin, E., Ciavarella, C., Cuomo-Dannenburg, G., Barzon, L., Del Vecchio, C., et al. (2020). Suppression of a SARS-CoV-2 outbreak in the Italian municipality of Vo'. *Nature*, 584(7821), 425–429.
- Leopoldina (2020). Dritte Ad-hoc-Stellungnahme: Coronavirus-Pandemie - Die Krise nachhaltig überwinden. Available online https://www.leopoldina.org/uploads/tx_leopublication/2020_04_13_Coronavirus-Pandemie-Die_Krise_nachhaltig.%C3%BCberwinden.final.pdf. (Accessed 15 April 2020).
- Li, Q., Guan, X., Wu, P., Wang, X., Zhou, L., Tong, Y., et al. (2020). Early transmission dynamics in wuhan, China, of novel coronavirus-infected pneumonia. *New England Journal of Medicine*, 382(13), 1199–1207.
- Limon, D., Bravo, J., Alamo, T., & Camacho, E. (2005). Robust MPC of constrained nonlinear systems based on interval arithmetic. *IEEE Proceedings D (Control Theory and Applications)*, 152(3), 325–332.
- Linton, N. M., Kobayashi, T., Yang, Y., Hayashi, K., Akhmetzhanov, A. R., Mok Jung, S., et al. (2020). Incubation period and other epidemiological characteristics of 2019 novel coronavirus infections with right truncation: A statistical analysis of publicly available case data. *Journal of Clinical Medicine*, 9(2), 538.
- Maharaj, S., & Kleczkowski, A. (2012). Controlling epidemic spread by social distancing: Do it well or not at all. *BMC Public Health*, 12(679), 1–16.
- Maier, B. F., & Brockmann, D. (2020). Effective containment explains subexponential growth in recent confirmed COVID-19 cases in China. *Science*, 368(6492), 742–746.
- Meyer-Hermann, M., Pigeot, I., Priesemann, V., & Schöbel, A. (2020). Adaptive Strategien zur Eindämmung der COVID-19-Epidemie. https://www.helmholtz.de/fileadmin/user_upload/01_forschung/28-04-2020_Strategien_zur_Eindaemmung.pdf. (Accessed 29 April 2020).
- Mizumoto, K., Kagaya, K., Zarebski, A., & Chowell, G. (2020). Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan. *Eurosurveillance*, 25(10), Article 2000180.
- Nowzari, C., Preciado, V. M., & Pappas, G. J. (2016). Analysis and control of epidemics. *IEEE Control Systems Magazine*, 36, 26–46.
- Oliver, N., Letouzé, E., Sterly, H., Delataille, S., Nadai, M. D., Lepri, B., et al. (2020). Mobile phone data and COVID-19: Missing an opportunity?. arXiv preprint [arXiv:2003.12347](https://arxiv.org/abs/2003.12347).
- Preciado, V. M., Zargham, M., Enyioha, C., Jadbabaie, A., & Pappas, G. J. (2014). Optimal resource allocation for network protection against spreading processes. *IEEE Transactions on Control of Network Systems*, 1(1), 99–108.
- Rawlings, J. B., Mayne, D. Q., & Diehl, M. (2017). *Model predictive control: Theory, computation, and design*. Nob Hill Publishing.
- Robert Koch Institut (2017). *Epidemiologisches Bulletin Nr.27* (p. 243). Available online https://www.rki.de/DE/Content/Infekt/EpidBull/Archiv/2017/Ausgaben/27_17.pdf?__blob=publicationFile.
- Tsay, C., Lejarza, F., Stadtherr, M. A., & Baldea, M. (2020). Modeling, state estimation, and optimal control for the US COVID-19 outbreak. arXiv preprint [arXiv:2004.06291](https://arxiv.org/abs/2004.06291).
- Verity, R., et al. (2020). Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*, 20(6), 669–677.
- Wang, D., Hu, B., Hu, C., Zhu, F., Liu, X., Zhang, J., et al. (2020). Clinical characteristics of 138 hospitalized patients with 2019 Novel Coronavirus-infected Pneumonia in Wuhan, China. *Journal of the American Medical Association*, 323(11), 1061–1069.
- Watkins, N. J., Nowzari, C., & Pappas, G. J. (2019). Robust economic model predictive control of continuous-time epidemic processes. *IEEE Transactions on Automatic Control*, 65, 1116–1131.
- World Health Organization (WHO) (2020). Report of the WHO-China Joint Mission on Coronavirus Disease 2019 (COVID-19). Available online <https://www.who.int/docs/default-source/coronavirus/who-china-joint-mission-on-covid-19-final-report.pdf>. (Accessed 15 April 2020).
- Yang, X., Yu, Y., Xu, J., Shu, H., Xia, J., Liu, H., et al. (2020). Clinical course and outcomes of critically ill patients with SARS-CoV-2 pneumonia in wuhan, China: a single-centered, retrospective, observational study. *The Lancet Respiratory Medicine*, 8, 475–481.