



OPEN

Modeling, state estimation, and optimal control for the US COVID-19 outbreak

Calvin Tsay^{1,3}, Fernando Lejarza^{1,3}, Mark A. Stadtherr¹ & Michael Baldea^{1,2}✉

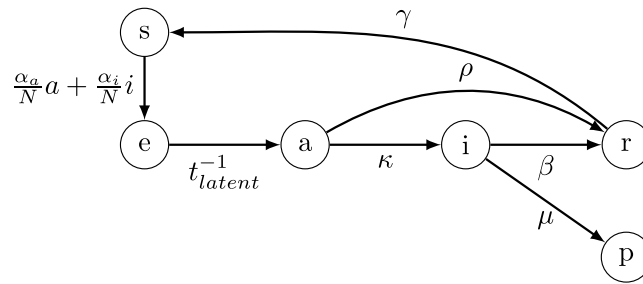
The novel coronavirus SARS-CoV-2 and resulting COVID-19 disease have had an unprecedented spread and continue to cause an increasing number of fatalities worldwide. While vaccines are still under development, social distancing, extensive testing, and quarantining of confirmed infected subjects remain the most effective measures to contain the pandemic. These measures carry a significant socioeconomic cost. In this work, we introduce a novel optimization-based decision-making framework for managing the COVID-19 outbreak in the US. This includes modeling the dynamics of affected populations, estimating the model parameters and hidden states from data, and an optimal control strategy for sequencing social distancing and testing events such that the number of infections is minimized. The analysis of our extensive computational efforts reveals that social distancing and quarantining are most effective when implemented early, with quarantining of confirmed infected subjects having a much higher impact. Further, we find that “on-off” policies alternating between strict social distancing and relaxing such restrictions can be effective at “flattening” the curve while likely minimizing social and economic cost.

Since its first reported case in early December 2019 in Wuhan, Hubei Province, China, the novel Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) and the resulting COVID-19 disease have reached 184 countries/regions causing a total of 149,024 deaths, as of this writing. The rate of spread of the virus is substantially higher than that of similar previously reported epidemics such as SARS and MERS¹. Because of this worldwide and rapid spread, on March 11, 2020, the World Health Organization (WHO) declared the COVID-19 outbreak a global pandemic². The exponential increase in the number of confirmed infectious cases and deaths has caused countries all around the world to respond with severe lock-down, quarantining, and social distancing measures to contain the spread of the disease. For example, northern Italy went into a state of emergency on March 8, 2020, imposing a complete lock-down that was expanded to the rest of the country three days later³. Following Italy’s attempts to decrease the contagion rate, Spain enforced a nationwide lock-down on March 16⁴. In contrast, the response in the US was significantly slower than those of the aforementioned (and most other European) countries, delaying widespread social distancing measures and disease screening initiatives.

Given the exponential growth of virus infections, policy makers face the urgent challenge of determining the appropriate response(s). For example, restrictions on population mobility, increased resource investment (e.g., personal protective equipment, ventilators, hospital beds in intensive care units), as well as improved COVID-19 screening can greatly impact the spread of infection^{5,6}. Here, fundamental epidemiological models⁷, typically comprising sets of coupled, nonlinear ordinary differential equations (ODEs), are valuable tools for simulating the dynamics of the epidemic and investigating suppression strategies. Several works (e.g., Anastassopoulou et al.⁸, Peng et al.⁹, Magal and Webb¹⁰) have addressed the development and fitting of such models for the COVID-19 outbreak using regression techniques. Models were subsequently used to characterize the infection^{11,12}, predict its potential spread^{13–15}, and/or evaluate mitigation strategies^{16,17}.

On the other hand, fewer studies have investigated the dynamic optimization, or *optimal control*, based on epidemiological models in relation to the current COVID-19 outbreak. In this context, the aforementioned epidemiological models are used to determine the policies (i.e., model inputs), such as the social distancing measures and testing rates, that lead to the best outcomes at the population level. Such outcomes include, for example, minimizing the peak number of infected people or total number of deaths, while also accounting for constrained

¹McKetta Department of Chemical Engineering, The University of Texas at Austin, Austin, TX, USA. ²Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX, USA. ³These authors contributed equally: Calvin Tsay and Fernando Lejarza. ✉email: mbaldea@che.utexas.edu



Input/parameter	Range	Description
$\alpha_a(t)$ [days ⁻¹]	[0.05, 0.5]	social distancing
$\alpha_i(t)$ [days ⁻¹]	[0.01, 0.3]	quarantining
$\kappa(t)$ [days ⁻¹]	[0.10, 0.3]	rate of testing
β [days ⁻¹]	[0.001, 0.05]	recovery rate
μ [days ⁻¹]	[0.001, 0.05]	death rate
e_0 [people]	$[0, N \times 10^{-6}]$	initial exposed

Figure 1. Digraph representation of SEAIR model; model inputs and parameters.

resource availability and the extent to which social distancing is feasible. Among works of this type, Djidjou-Demasse et al.¹⁸ investigated the optimal control of a single “intervention” input variable for an ODE model. Moore and Okyere¹⁹ formulated a similar problem with additional inputs, including hospitalization rates and environmental spraying. Both studies solved the optimal control problem using an iterative forward–backward sweep method²⁰ and assumed the model parameters to be fully known. While optimal control of epidemiological models has been well-studied (e.g., Biswas et al.²¹, Neilan and Lenhart²²), a limiting factor for implementing modeling and optimization concepts during early stages of an outbreak is that obtaining accurate estimates of key model parameters can be challenging.

In this work, we report a novel and complete dynamic optimization-based approach to the entire epidemiological modeling and outbreak control workflow for the US COVID-19 outbreak. We first formulate a dynamic optimization strategy for identifying both the values of time-invariant parameters and the historical trajectories of time-varying parameters (i.e., inputs) of an epidemiological model. We then investigate how optimal control of the inputs of the same model, which reflect social distancing and testing, relates to infection mitigation strategies. The optimal control problem is cast as a *simultaneous* dynamic optimization problem (i.e., reformulated as an algebraic problem)²³, enabling the natural use of deterministic global optimization technology²⁴. This in turn provides solutions that are proven to be globally optimal, whereas iterative schemes such as the above²⁰ often only guarantee global optimality under certain conditions (the term “global” here refers to the notion that the best possible disease control policy is identified from a set of “local” solutions that otherwise satisfy optimality conditions. “Global” in this context should not be interpreted as “world-wide”). Furthermore, we show how state estimation should be used to update important hidden model states (e.g., the number unconfirmed infections) to reduce the impact of model inaccuracy. Implementations of the parameter estimation and optimal control problems in open-source software are provided freely at <https://github.com/Baldea-Group/covid-19>.

Results

Mathematical modeling. We consider the compartmental model shown in Fig. 1, which brings a few key modifications to the conventional SEIR (susceptible–exposed–infectious–recovered) structure. Importantly, the $a(t)$ state is added to account for infected subjects that are not included in the count of confirmed cases, either because they are asymptomatic or because of insufficient testing. The virus is thought to be asymptomatic in 20–40% of cases^{25,26} and may be transmitted by asymptomatic carriers²⁷. We also include a $p(t)$ state, which tracks the population that perishes due to the virus.

The modified SEAIR model has six compartments/states: $s(t)$ represents the number of subjects that are susceptible to infection, $e(t)$ the number that have been exposed to the virus, $a(t)$ the number that are infected but asymptomatic/unconfirmed, $i(t)$ the number with confirmed infections, and $r(t)$ the number that have recovered from infection. The states are assumed to sum to a known total population N , and $p(t)$ can be calculated algebraically at all times. Note that N is assumed to be constant. This assumption is suitable since the number of deaths is much smaller than the initial population, and thus considering a time-varying total population would have little impact on our model at the expense of increased complexity. The model structure is shown in Fig. 1.

The SEAIR model is described by the following equations:

$$\frac{ds(t)}{dt} = \frac{-\alpha_a(t)}{N} s(t)a(t) - \frac{\alpha_i(t)}{N} s(t)i(t) + \gamma r(t) \tag{1}$$

$$\frac{de(t)}{dt} = \frac{\alpha_a(t)}{N} s(t)a(t) + \frac{\alpha_i(t)}{N} s(t)i(t) - t_{\text{latent}}^{-1} e(t) \quad (2)$$

$$\frac{da(t)}{dt} = t_{\text{latent}}^{-1} e(t) - \kappa(t)a(t) - \rho a(t) \quad (3)$$

$$\frac{di(t)}{dt} = \kappa(t)a(t) - \beta i(t) - \mu i(t) \quad (4)$$

$$\frac{dr(t)}{dt} = \rho a(t) + \beta i(t) - \gamma r(t) \quad (5)$$

$$\frac{dp(t)}{dt} = \mu i(t) \quad (6)$$

where $\alpha_a(t)$ and $\alpha_i(t)$ are the rates of exposure to the virus from the population of asymptomatic/unconfirmed and confirmed infected subjects, respectively. These two rates of exposure are defined as time-varying and independent model inputs to reflect different measures taken during the course of the pandemic. Specifically, $\alpha_a(t)$ corresponds to exposure from asymptomatic carriers (a) and reflects social distancing and/or shelter-in-place strategies. On the other hand, $\alpha_i(t)$ corresponds to exposure from infected subjects (i) and reflects quarantining of infected subjects. We also model $\kappa(t)$, or the rate at which unconfirmed cases become confirmed, as a time-dependent input to reflect varying levels of screening and testing.

The other parameters are assumed to be constant over the time horizons considered herein. t_{latent}^{-1} is the inverse of the latent period of the virus, or the time before an exposed subject becomes infectious. We assume a value of $t_{\text{latent}}^{-1} = 0.5 \text{ days}^{-1}$ based on Peng et al.⁹ The parameter ρ describes the infectious period for subjects with unconfirmed infections, for which we assume a value of $\rho = 0.1 \text{ days}^{-1}$, based on Rocklöv et al.²⁸ The rates at which subjects with confirmed infections recover and perish are described by β and μ , respectively. Finally, γ describes the rate at which recovered subjects become susceptible to the disease again. Since the immunity period for the virus is unknown, we assume a value of $\gamma = 0$. While these parameters largely describe the virus itself, they may change over longer time horizons, e.g., as new treatments are developed. A potential future avenue of research is to consider parameters β and γ to be time-varying, to reflect decision-making regarding budget allocation for medical supplies and equipment, as well as health care capacity expansion over time.

In summary, the SEAIR model (1)–(6) has three time-varying inputs ($\alpha_a(t)$, $\alpha_i(t)$, and $\kappa(t)$), three parameters (t_{latent}^{-1} , ρ , and γ) whose values are based on available literature, and two parameters whose values must be estimated (β and μ). Note that since $e(t)$ and $a(t)$ are unmeasured, their magnitudes relative to the infected population i are effectively set by (the bounds on) κ . As the prevalence of asymptomatic infection is established (e.g., Bendavid et al.²⁹ estimate up to 4% prevalence in Santa Clara County using serological testing) the magnitudes of e and a may be adjusted by scaling κ , and other parameters as relevant.

Parameter estimation results. This work primarily addresses the current COVID-19 situation in the US, but the still relatively early stage of the outbreak in the US renders the available data insufficient to fit parameters reflective of prevention measures already in place. For this reason, we perform a comparative analysis by solving the parameter estimation problem for Italy, Spain, and Germany, where, while the epidemic has not been fully contained yet, the daily number of new confirmed cases has been on a steady decline.

Model parameters for the US, Italy, Spain, and Germany. Figure 2 shows the predicted values obtained by solving the parameter estimation problem and the historical data by country, retrieved by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (<https://github.com/CSSEGISandData/COVID-19>; accessed April 16, 2020). Day 1 of the dataset corresponds to January 22, while day 85 corresponds to April 16. The solid line in each state trajectory plot shows the mean predicted value. Standard deviations for the predictions were estimated using bootstrapping, with the time-invariant parameters sampled from a multivariate normal distribution. The means and covariances of the estimated parameters for each country are provided in the Supplementary Information. Plots corresponding to infected, recovered and perished (left plots in Fig. 2) were obtained by simulating (1)–(6) using the estimated trajectories of the time-varying inputs, shown in the right-most column of Fig. 2.

Parameter estimation results provide several interesting insights. First we note that the mortality rates, while higher for the European countries, are comparable in magnitude for the four countries considered and are similar to prior estimates³⁰. The higher death rates in Italy and Spain are likely explained by demographic factors (e.g., age), as well as the medical resources available to treat the infected population. However, the simulated trajectories of p in Italy and Spain deviate slightly from the historical data, especially around day 80 (Fig. 2), suggesting that the death rate μ may be time-varying. While μ can vary in reality due to improved treatment, early detection bias, etc., it is difficult to anticipate future changes, and we therefore approximate it as a constant.

From the trajectories of α_a , α_i , and κ , we observe that social distancing measures (i.e., lower values of α_a) were first implemented in Europe, and more recently the US, which agrees with their true chronology. This insight is confirmed by the fitted values for e_0 , with Italy having the highest value and the US the lowest, which is representative of the beginning of the outbreak in each country. In general terms, the identified evolution of containment and testing measures follow similar trends for all four countries, which supports that the model

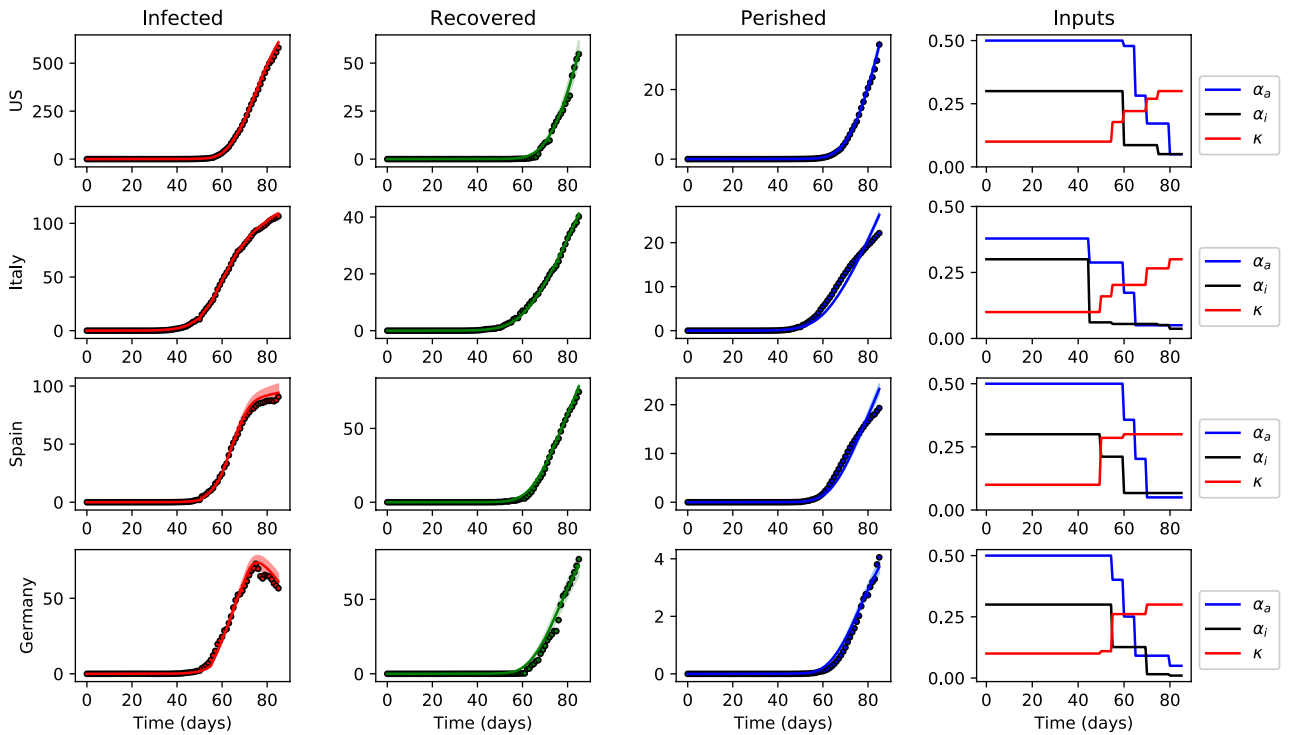


Figure 2. Historical data and model fit for infected ($i(t)$), recovered ($r(t)$), and perished ($p(t)$) subjects (in thousands). Solid lines represent the mean of 500 Monte Carlo simulations, shaded areas represent two standard deviations from the mean, and circle markers are historical data. The right-most column shows the fitted trajectories of the time-varying inputs.

structure is correct and can appropriately reflect (see “Methods”) different dynamics of the same underlying disease. Further, we expect that the parameters obtained for the US, while being fitted using relatively premature information, are likely an adequate representation of the current COVID-19 situation.

State simulations for fitted parameters. For the estimated parameter values as described previously, we simulate the results of implementing two different simplistic control policies: (i) continuing with strict social distancing, quarantining, and testing, policies that result from continuing to lower the asymptomatic (α_a) and infected (α_i) exposures shown in Fig. 2; and (ii) a relaxed policy with more lenient measures and reduced testing, in this case the values of α_a and α_i are increased to 0.2 and 0.02, respectively, while κ is decreased to 0.2. The population levels resulting from implementing these two policies are shown in Fig. 3, where the numbers of recovered subjects are omitted for the sake of brevity. We additionally show the number of new confirmed cases per day, as it is a commonly used metric to illustrate the current spread/containment of the virus. Since the three European countries showed very similar trends under the two policies considered, we only compare the results for the US to those obtained for Italy. On the one hand, from Fig. 3 it is evident that relaxing current control policies can result in an alarming number of infected cases and deaths, particularly in the US. On the other hand, continuing with the strict shelter-in-place measures and maximizing testing can result in earlier and substantially flatter pandemic peaks, with significantly lower numbers of casualties. While the latter approach is the most effective at preventing further exponential spread of the disease, it is also the most socially and economically disruptive policy. In light of this trade-off, we argue that by means of employing optimal control concepts it is possible to find effective policies that maintain the number of infected cases below a given threshold, while minimizing the extent of social and economic disruption.

Dynamic optimization results. Here, we consider the optimal control of the COVID-19 infection in the US using the deterministic SEAIR model (1)–(6) with the mean values of the time-invariant parameters as described above. As such, the optimization problem only considers the mean predicted values (e.g., solid lines in Figs. 2 and 3). Future work may consider a stochastic optimization approach that, at the expense of increased computational effort, leverages information about model uncertainty.

Optimization of future actions. We solve a dynamic optimization problem that minimizes a measure of socio-economic cost, subject to keeping the peak (max) value of the infected population below a given number i_{peak} . The lower bound of $\kappa(t)$ is raised to 0.15, from 0.10 in Fig. 1. This change was made because (i) the historical value of κ for all countries analyzed quickly jumps to its maximum value 0.3, perhaps reflecting public initiatives, and (ii) the $[0.15, 0.3]$ range translates to $\frac{\rho}{\rho+\kappa} \in [25\%, 40\%]$, which is reflective of current estimates of the asymptomatic ratio^{25,26}.

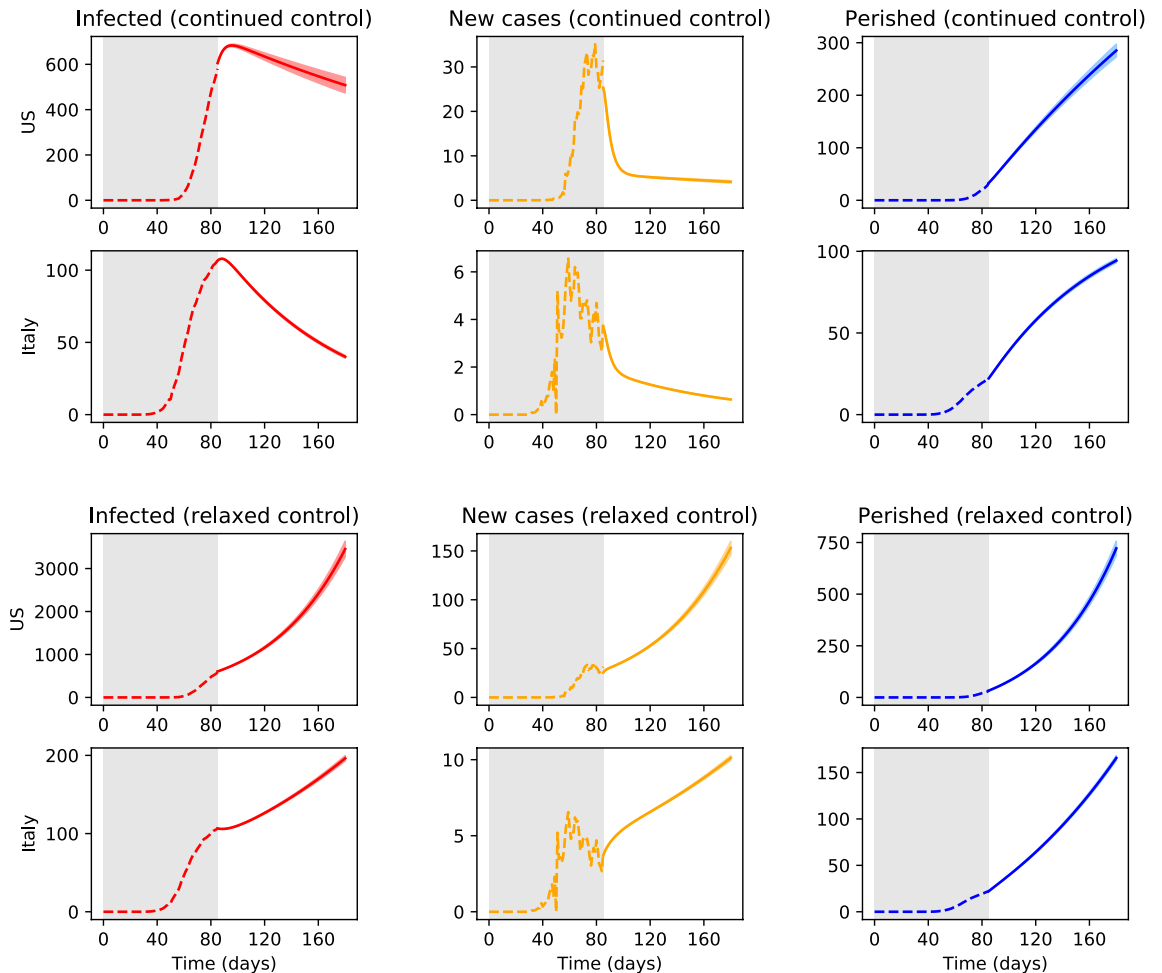


Figure 3. Simulation of future infected (i), new confirmed cases (κa), and perished (p) subjects (in thousands) for current control policies. Shaded grey area indicates historical data, color solid lines represent the mean of 500 Monte Carlo simulations, and color shaded area represents two standard deviations from the mean.

Figure 4 depicts the solution of this problem for the cases where $i_{\text{peak}} = 700,000$ subjects and $i_{\text{peak}} = 1,400,000$ subjects. The optimization problem was solved for the 100 days following the end of the available dataset ($t_0 = 85$, $t_f = 185$), with the initial condition at t_0 computed by simulating days 1–85. Although the optimization problem only considers the mean values (solid lines), two standard deviations are shown in the shaded areas to reflect parametric model uncertainty for the given inputs, computed using the same bootstrapping approach as above.

Considering the current status of the pandemic in the US, keeping the peak below 700,000 infected subjects is very challenging. Note that the minimum feasible value of i_{peak} corresponds to a peak of 612,493 infected subjects. Therefore, maintaining $i(t) \leq 700,000$ requires immediately decreasing α_a and α_i to their lower bounds for approximately the next 25 days, during which the exposed population can decrease significantly. After this period, α_a is relaxed to its upper bound for seven brief periods (Fig. 4, left). However, α_i remains at its lower bound at all times, reflecting a strict quarantining policy.

This policy reveals that the impact of α_i on the infected population size is larger than that of α_a . In turn, this suggests that quarantining of infected people is more important than social distancing, which mitigates exposure to unconfirmed cases. The optimization problem penalizes κ to account for the cost of testing, and κ is decreased during periods of decreased α_a , suggesting that testing is less important during times of social distancing/lockdown. Intuitively, when the exposure to asymptomatic subjects is already low due to social distancing, there is less benefit to testing asymptomatic subjects (and transferring them to the similarly quarantined infected population). On the other hand, κ is increased preemptively for each period of increased α_a , suggesting that testing is most important in the days leading up to a period of relaxed social distancing. This testing moves asymptomatic subjects to the confirmed-infected population, which can remain quarantined.

Figure 4 (right) depicts the solution obtained for a value of $i_{\text{peak}} = 1,400,000$ people. While α_i again remains at its lower bound at all times, re-emphasizing the importance of quarantining infected subjects, there are more frequent and longer periods of increased α_a . In this case, κ remains at its upper bound for most of the control horizon. In both cases, the optimization problem does not account for the effect(s) of the control policy after 100 days, and growth of the infected population near the end of the time horizon may be concerning. This can

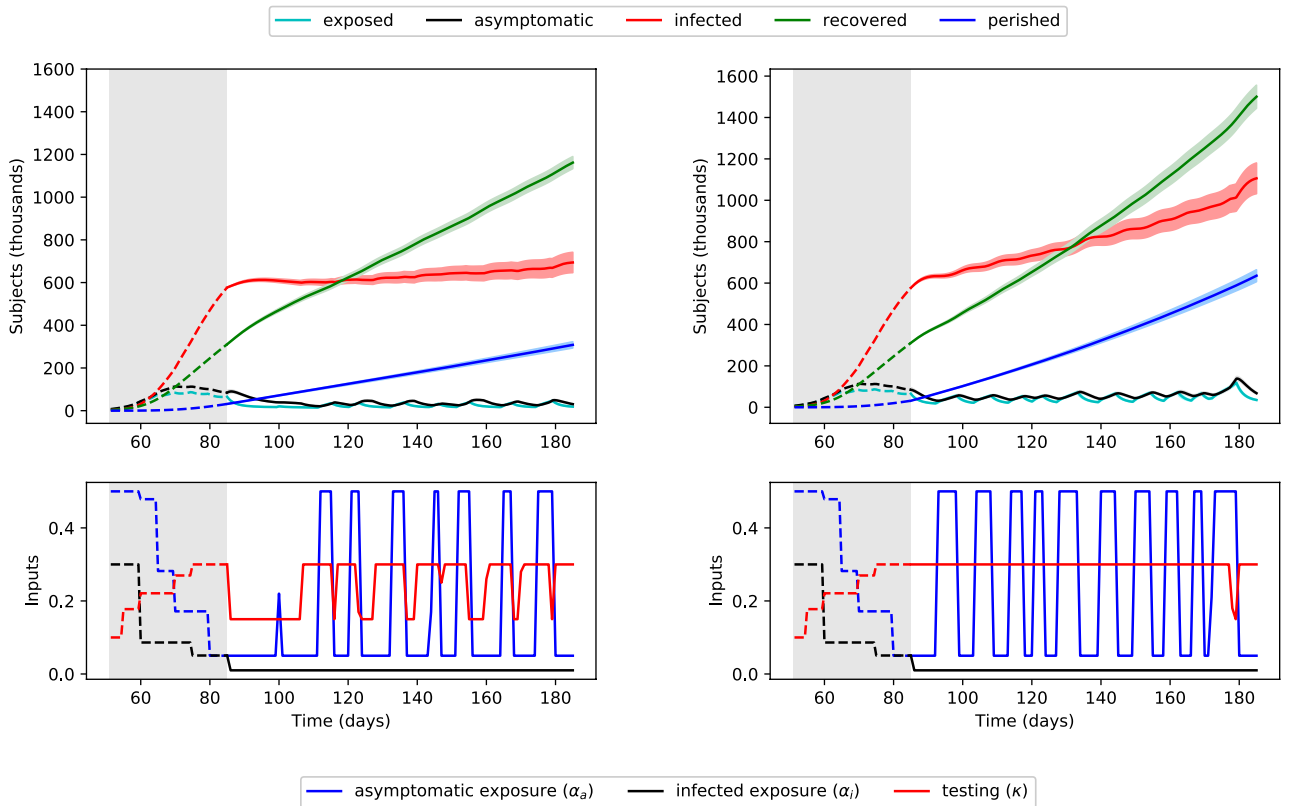


Figure 4. Optimal control policy to limit peak infections to 700,000 (left) or 1,400,000 (right) in the next 100 days. Top: population numbers, with two standard deviations shaded. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized).

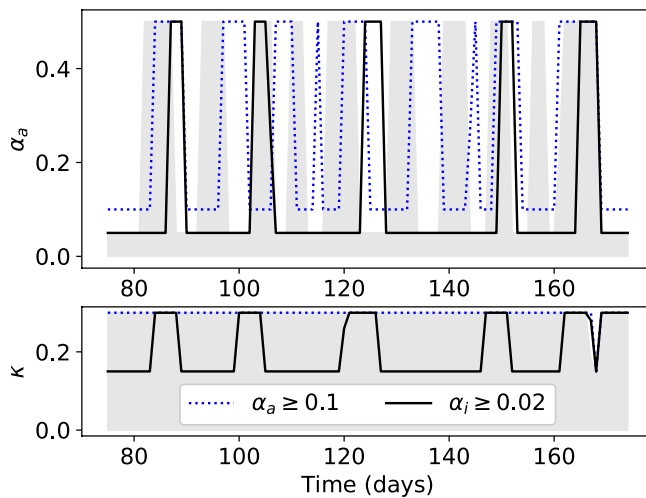


Figure 5. Optimal containment and testing strategies to limit peak infections to 1,000,000 in the next 100 days for different constraints on α_a and α_i . Top: $\alpha_a(t)$. Bottom: $\kappa(t)$. The shaded grey area indicates the solution found using the normal bounds, replicated from Fig. 4 (right).

be addressed, e.g., using a moving horizon control strategy, where policy measures are revised periodically, as discussed below.

The lower bound used for α_i is lower than estimated values for the US for the past (Fig. 4), corresponding to a new level of quarantining. Furthermore, the current value of α_a , or social distancing, may not be economically sustainable over longer periods. To investigate the effect of not achieving these proposed levels, we solve the same problem with $i_{\text{peak}} = 1,400,000$ people, using different lower bounds for the inputs. Figure 5 shows the optimal

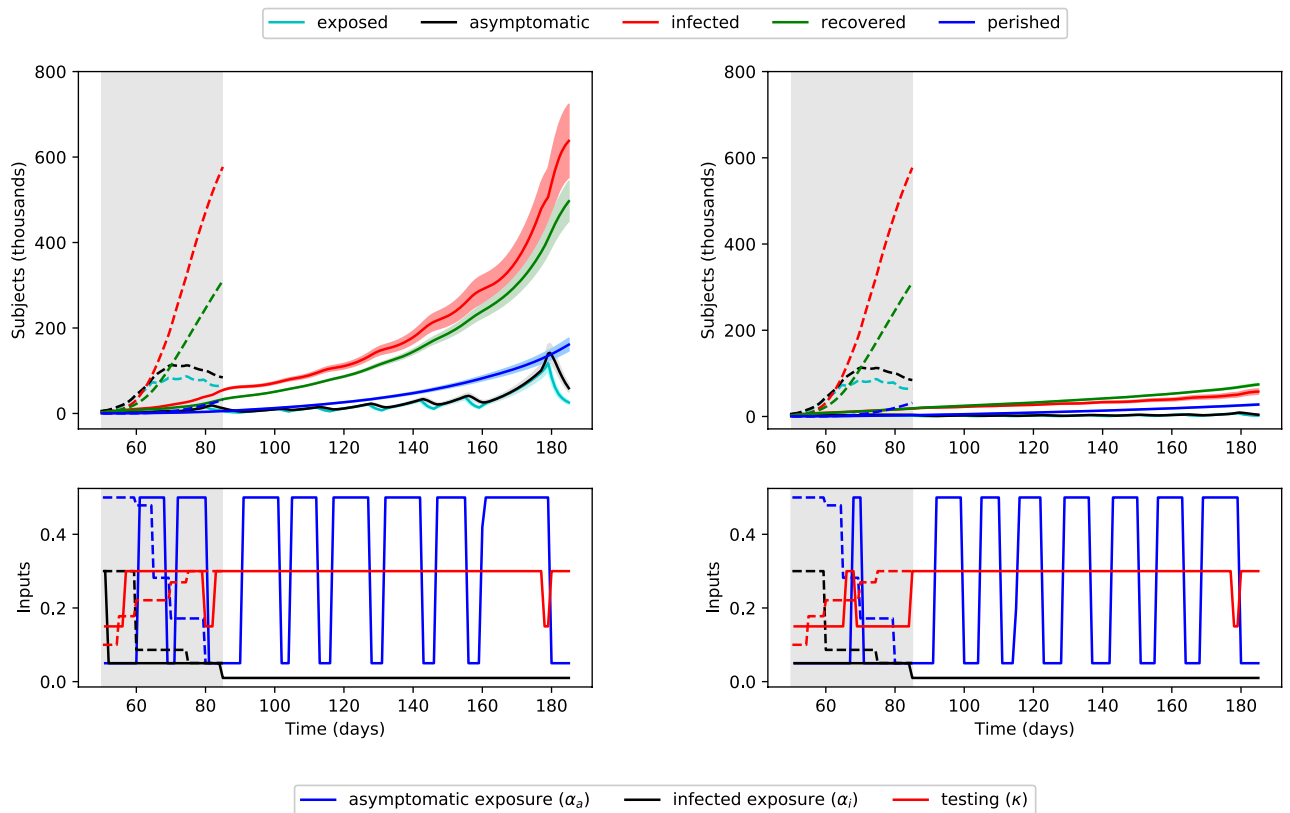


Figure 6. Optimal control policy to limit peak infections to 700,000 (left) or 70,000 (right) in the past 35 days and next 100 days. Top: population numbers, with two standard deviations shaded. Bottom: containment and testing profiles. The shaded grey area indicates past days, for which the true historical inputs and outputs are shown as dashed lines and the optimized are shown as solid lines.

input profiles found for the cases where the lower bound of α_a is doubled from 0.05 to 0.1 (dotted blue lines), and where the lower bound of α_i is doubled from 0.01 to 0.02 (solid black lines). When the lower bound of α_a is doubled (less social distancing is achieved during lock-down periods), the same infected population peak can be maintained by increasing the frequency of social distancing periods. However, when the lower bound of α_i is doubled (less quarantining is achieved), maintaining the same number of peak infections requires social distancing at almost all times, closely resembling the control policy found for a lower peak (Fig. 4, left). In either case, the optimal value of α_i remains at its lower bound (either 0.01 or 0.02) at all times and is therefore not shown. Periods of social distancing are further increased in both frequency and duration if both bounds are doubled.

Compared to the decreases seen in Fig. 3 for the cases of continued control, the solutions in Fig. 4 seek to merely flatten the growth of the infected population. The associated optimal control policies resemble “bang-bang” control (the inputs are always at the lower or upper bound), which can be expected from a systems-theoretic point of view for (1)–(6), as the equations are linear functions of the inputs α_a , α_i , and κ . This result supports the strategy proposed by the Imperial College COVID-19 Response Team³¹, which involves periodic suppression measures. Their strategy triggers the start of a social distancing period by when the number of weekly ICU cases increases past an “on” threshold, and the end of the period when the same number decreases below an “off” threshold. Figure 4 shows that this type of strategy maintains a low exposed population, and therefore flattens the growth of the infected population while still allowing periods of social mobility.

Revisiting past actions. We examine the same optimal control problem with a time horizon starting from day 50 (i.e., covering 35 days in the past). Rolling the time horizon of the optimization problem backwards to $t_0 = 50$, allows us to investigate the optimal inputs for days 50–85, corresponding roughly to the latter half of March and the first half of April. Figure 6 (left) depicts the solution obtained for a value of $i_{\text{peak}} = 700,000$ people. The initial condition at $t_0 = 50$ was computed by simulating days 1–50. The minimum historical values for α_a and α_i were used as their respective lower bounds for days 50–85, such that the optimal control policy for days 50–85 only involves an extent of quarantining and social distancing already experienced by the general population.

In this case, maintaining the number of infected subjects below 700,000 appears easy. The minimum feasible peak of infected subjects is 16,543, compared to 612,493 when the optimization starts at day 85. Therefore, maintaining $i(t) \leq 700,000$ can involve extended periods of no social distancing $\alpha_a = 0.5$. Similar to the result in Fig. 4 (right), testing and quarantine are important in this scenario: κ remains at its upper bound and α_i its lower bound. The solution for an order-of-magnitude lower minimum peak, corresponding to 70,000 infected subjects is shown in Fig. 6 (right). The optimal control policy for this scenario involves more frequent periods of social

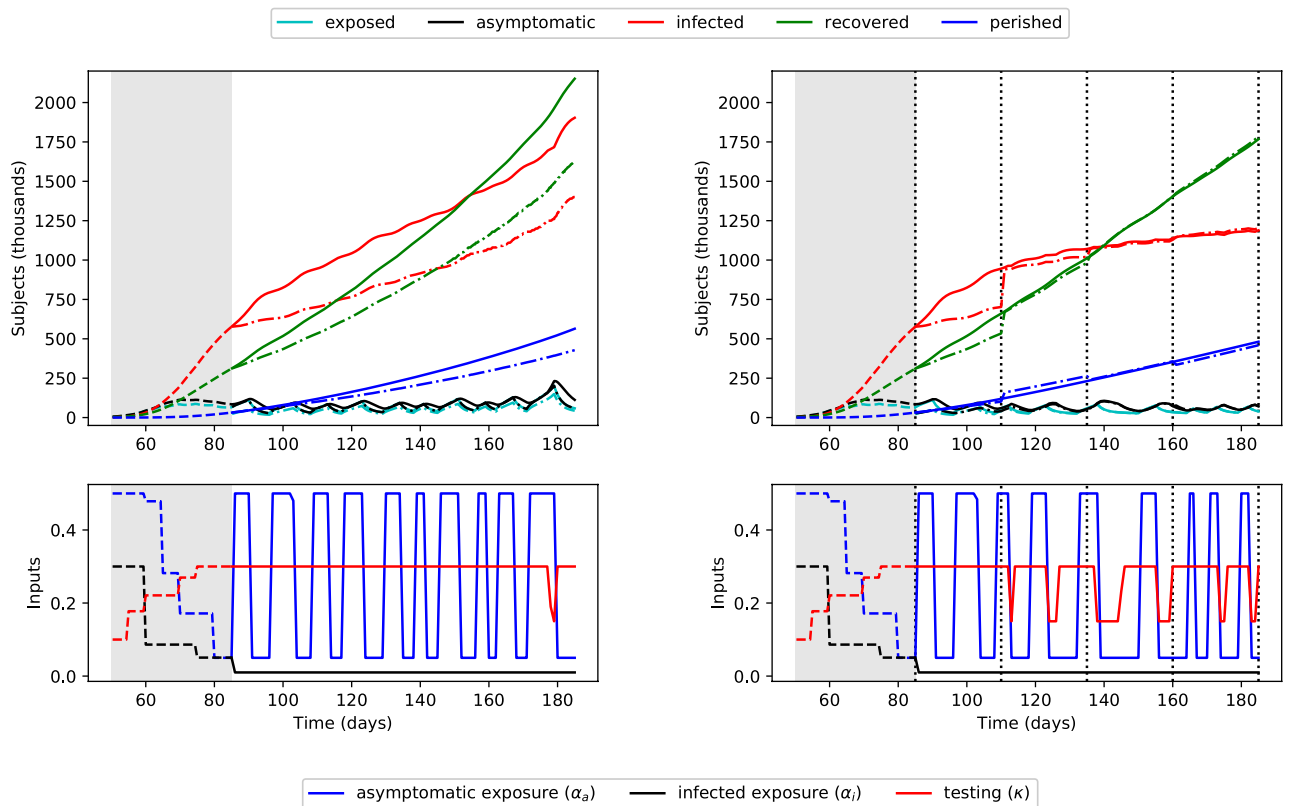


Figure 7. Optimal moving horizon control policy (right) to limit peak infections to 1,400,000, with e and a underestimated by a factor of three at $t = 74$, and comparison to the same situation without a moving horizon strategy (left). Top: predicted (dash-dotted) and true (solid) population numbers. Bottom: containment and testing profiles. The shaded grey area indicates past days, which were simulated using historical inputs (not optimized). The policies are updated every 25 days with daily state estimation.

distancing. Early implementation of testing and quarantining strategies clearly has an enormous effect, manifest in the large difference between the historical inputs (dashed) and optimal inputs (solid) in days 50–85 of Fig. 6.

Estimation of hidden states and moving horizon control. The SEAIR model (1)–(6) includes two *hidden* states, $e(t)$ and $a(t)$, which are not measured in practice (note that widespread serological testing²⁹ may eventually reveal the true levels of the asymptomatic population a). Nevertheless, solutions for the optimal control problem based on (1)–(6) are highly dependent on the values of $e(t)$ and $a(t)$, and therefore the initial conditions $e(t_0)$ and $a(t_0)$. Fig. 7 (left) depicts the optimal control policies for $t_{\text{peak}} = 1,400,000$ people over the next 100 days (similar to Fig. 4, right), where $e(t_0)$ and $a(t_0)$ are now underestimated by a factor of three. This is a practically motivated scenario, as the number of asymptomatic cases is considered largely uncertain^{6,32}. The dashed-dotted lines show the state profiles predicted by the optimization problem (with incorrect initial conditions), while the solid lines show their true evolution for the given inputs.

The error in the hidden states $e(t)$ and $a(t)$ causes the predicted and actual profiles to diverge over time. By the end of the 100 days, there are 1.90 million actual infected subjects, almost 40% greater than the predicted 1.40 million. To account for discrepancies between the modeled and true systems, the optimal control inputs should be periodically updated. We propose to achieve this with a *moving horizon* decision-making strategy: after a given length of time, the dynamic optimization problem is re-solved, setting the measured values of the states as their initial conditions. The values of the un-measured, hidden states can be estimated using state estimation techniques, such as the Kalman filter (the optimal linear estimator) or its nonlinear extensions.

We consider a relatively long time span of 25 days before each revision of policy decisions (and re-optimization), as planned suppression strategies likely cannot be altered quickly. Figure 7 (right) shows the optimal trajectories found for the same problem (initial conditions for e and a underestimated by a factor of three) using the moving horizon strategy. The vertical dotted lines indicate times when the optimization problem was resolved. Values of the hidden states were estimated daily with the (discrete) unscented Kalman filter³³, implemented using the pykalman library (<https://pykalman.github.io/>; Accessed April 7, 2020). Measurements for state estimation and control updates were simulated by adding independent, normally distributed noise to the true state values. A standard deviation of 5,000 people was used for the measured states (i , r , and p) based on the size of residuals during parameter estimation.

The predicted values (dash-dotted lines in Fig. 7) are first updated when the optimization problem is re-solved at $t = 110$ days. The solutions to the re-optimization problems are shown in the Supplementary Information. The

effect of this moving horizon strategy is quite significant, in comparison to the optimal policy shown in Fig. 7 (left). Periods of social distancing are quickly increased after day 110 to account for the larger-than-expected increase in infected population. Testing levels, reflected by κ , are also updated accordingly. The moving horizon approach exhibits three clear benefits:

- Errors between the model predictions and reality are minimized by incorporating new measurements as they become available
- The infected population does not grow exponentially at the end of the time horizon, as is the case when considering only a single time horizon.
- The control policies can be computed indefinitely into the future by shifting the time window forward in time, while the problem size remains tractable (only 100 days are considered during optimization).

The effect of state estimation is most notable in the first moving horizon window, between days 85–110. Here, while the values predicted by the optimization problem are consistently under-predicted, the estimated values approach the (unknown) true values as daily measurements are added. The estimated values over time are shown in the Supplementary Information. The new solution of the optimization problem at day 110 uses the population levels measured/estimated based on new data as initial conditions, and therefore predicts more accurate trajectories.

Discussion

This work investigated dynamic optimization strategies to characterize and control the US COVID-19 outbreak, by minimizing the socioeconomic cost associated with containment strategies and testing. The results provide several overarching conclusions.

The quarantining of infected subjects is the most important of the considered mitigation strategies and should be maximized at all times. Additionally, periods of social distancing/lock-down help to flatten the peak by preventing exposure from asymptomatic and unconfirmed cases. Screening and testing for the disease are key immediately preceding periods of relaxed social distancing, in order to minimize the number of unconfirmed infections during periods of social mobility. Early action has much larger effects than later interventions, even as the later interventions are more drastic. Optimal policies are highly dependent on estimates of “hidden states,” i.e., the asymptomatic and unconfirmed cases. Moving horizon (periodically updated) policies and state estimation should be used to mitigate inaccuracies in the model and counts of asymptomatic/unconfirmed cases, by accounting for new data as they becomes available.

The “on-off” policies identified as optimal are characteristic of the class of problems considered (linear objective function and input-affine nonlinear model). Further, these policies are likely the easiest to implement in practical scenarios and to convey to the general population (as opposed to a policy where the social distancing parameters would take values between their upper and lower bounds). Their implementation would simply alternate between the strictest possible limitations, followed by periods of relative freedom of movement. The frequency of these periods may be restricted, e.g., by including rate-of-change constraints in the optimization problem, to account for the ability of a country or local authority to intermittently enforce social distancing policies.

The model used in this work does not include population influx to and/or outflux from the given system. Therefore, the optimal containment and testing strategies found do not account for new cases that may enter from outside the US. We also assumed that the recovery and death rates for the virus are constant and equal to their historical values for the next 100 days. Thus the results do not account for the possibility of, e.g., improved medical treatments, vaccine development, and/or viral mutation. Similarly, the model does not account for possible surveillance/detection bias in the historical data (e.g., increased likelihood of testing for subjects with more severe symptoms). The impacts of these assumptions are topics for future work. Also, while Sameni¹⁶ determined fixed points and did linear stability and observability analyses for a similarly-structured SEIR model with time-invariant parameters, another topic for future work is the extension of this mathematical analysis to the SEAIR model with time-dependent parameters used here.

Methods

Parameter estimation. We follow nonlinear (least-squares) regression for parameter estimation, which can be cast as an optimization problem, where the objective is to find the parameter values that minimize the mean squared error (MSE) between the predicted states and their measured values. The measured values correspond to data retrieved by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (<https://github.com/CSSEGISandData/COVID-19>; accessed April 16, 2020) and are provided in the Supplementary Information for several countries. The data consist of the total number of infected (\hat{i}_j), recovered (\hat{r}_j) and dead (\hat{p}_j) reported subjects, where j represents each day during the time period from January 22 to April 16, where the index j represents the time period each value was recorded. The MSE is given by:

$$L_0 = \sum_{j=0}^N (i(t=j) - \hat{i}_j)^2 + (r(t=j) - \hat{r}_j)^2 + (p(t=j) - \hat{p}_j)^2 \quad (7)$$

We note that when solving the parameter estimation problem, the contribution of the term $\rho a(t)$ to the recovered state $r(t)$, is removed from the system (1)–(6). This is because the data obtained only accounts for those

recovered that were confirmed to have the disease, meaning the only contribution to the observed recovered is the term $\beta i(t)$.

The parameters with values/trajectories to estimate are $\alpha_a(t)$, $\alpha_i(t)$, $\kappa(t)$, β , and μ . Furthermore, we must estimate the initial condition of the unreported states e and a at $t = 0$. We assume that there are initially no asymptomatic infections. The “seed” of the outbreak is therefore the initial number of exposed subjects $e_0 = e(t = 0)$, which we include as a parameter in the least-squares problem. We include bounds on the possible parameter values, based on values reported for similar models fitted to data from other regions⁹. Note that, while κ reflects the level of testing, it also affects the predicted asymptomatic ratio, which cannot be controlled. Therefore, tighter bounds should be used for κ than for α_a and α_i . The estimated parameters and their bounds are summarized in Fig. 1.

To prevent over-fitting the available data, we restrict the time-varying inputs (α_a , α_i , and κ) to be piece-wise constant over five-day intervals. These constraints reflect the fact that policies implemented have significant time delays in steering the states towards the desired values, and therefore should not and practically cannot be manipulated too frequently before their effect is observed on the population. Additionally, we constrain $\alpha_a(t)$ and $\alpha_i(t)$ to be monotonically decreasing over time, reflecting the enforcement of increasingly stricter disease control measures. Similarly, $\kappa(t)$ was constrained to be monotonically increasing over time to account for increased COVID-19 screening availability.

To easily interface with raw data sources, the least-squares regression problem was implemented in Python using the Pyomo modeling and optimization package³⁴. We discretized the dynamical system (1)–(6) with respect to the time domain using orthogonal collocation on finite elements²³, with one finite element per day. We used IPOPT³⁵ to solve the resulting nonlinear dynamic optimization problem to local optimality. Standard deviations for the time-invariant parameters were approximated by fixing the time-varying inputs to their estimated trajectories, and solving a maximum likelihood problem for the time-invariant parameters using the Model Validation tool in gPROMS (general PROcess Modeling System) v5.1.4.³⁶ Residuals were assumed independent and normally distributed, with variances estimated in the same maximum likelihood problem. The parameter covariances (see Supplementary Information) are then computed from the diagonal entries of the Hessian of the objective function.

Optimal control. The aim of the optimal control problem is to find the trajectories of “handles” $\alpha_a(t)$, $\alpha_i(t)$ and $\kappa(t)$, that minimize (or maximize) the value of a certain objective function. Most response measures for the COVID-19 outbreak seek to “flatten” the epidemic curve, that is, to contain the growth rate of the number of infected subjects via a combination of social distancing and testing. Clearly, social distancing and isolation/quarantining carry significant social and economic costs. We thus formulate the optimal control problem as a dynamic optimization problem, aiming to minimize a measure of social and economic cost subject to ensuring that the maximum number of infected subjects remains under a given peak value, i_{peak} . The optimization problem is expressed mathematically as:

$$\begin{aligned} \min_{\alpha_a(t), \alpha_i(t), \kappa(t)} & \int_{t_0}^{t_f} C(t) dt \\ \text{s.t.} & \text{SEAIR model (1)–(6)} \\ & \max_t(i(t)) \leq i_{\text{peak}} \\ & C(t) = -\alpha_a(t) - \alpha_i(t) + \lambda_{\kappa} \kappa(t) \\ & \alpha_a(t) \in [0.05, 0.5] \\ & \alpha_i(t) \in [0.01, 0.3] \\ & \kappa(t) \in [0.15, 0.3] \end{aligned} \quad (8)$$

where $C(t)$ is the cost function, and λ_{κ} is the relative cost of testing (increasing κ). Testing costs are assumed to be relatively small in comparison to isolation measures, and we selected $\lambda_{\kappa} = 0.1$. While the current cost function equally penalizes decreases in α_a and α_i , an additional weighting factor could be introduced to distinguish between the socioeconomic costs of social distancing and quarantining.

The solution to (8) can provide control policies over the full time horizon $[t_0, t_f]$ as in the case of the initial studies presented in this work. Alternatively, the policies can be updated after a shorter horizon (before t_f is reached), as in the case of the moving horizon approach. In the latter case, the time window considered, $[t_0, t_f]$ is shifted forward at a pre-determined time interval. Note that t_f is typically longer than the frequency at which the optimization problem is recomputed, and only the solution for the first time step(s) is implemented. For example, the moving horizon scenario presented here is solved with $t_f = 100$ days, but with policies updated after each 25-day period.

While the problem (8) identifies the minimum societal cost required to achieve a certain peak value of infections, a related, but different, problem could be formulated by minimizing i_{peak} subject to an upper bound on $C(t)$, i.e., finding the lowest achievable peak for a given total societal cost. For this study we investigate the relationship between C and i_{peak} by solving (8) for varying values of i_{peak} , provided in the Supplementary Information.

The optimization problem (8) was again solved by discretization of the time domain using orthogonal collocation, with one finite element per day. To report the best possible solutions, the resulting algebraic optimization problem was implemented in GAMS and solved using the commercial global optimization solver BARON v18.5.8³⁷. Problems were solved to a 0.1% optimality gap (i.e., the reported solution is proven to have an objective function value within 0.1% of that of the best possible solution). We additionally provide an implementation

of the same problem using open-source tools (Pyomo³⁴ and IPOPT³⁵), which can be used to solve (8) to local optimality.

Data availability

Models are available at <https://github.com/Baldea-Group/covid-19>.

Received: 18 April 2020; Accepted: 3 June 2020

Published online: 01 July 2020

References

1. Peeri, N. C. *et al.* The SARS, MERS and novel coronavirus (COVID-19) epidemics, the newest and biggest global health threats: what lessons have we learned?. *Int. J. Epidemiol.* <https://doi.org/10.1093/ije/dyaa033> (2020).
2. Sohrabi, C. *et al.* World Health Organization declares global emergency: a review of the 2019 novel coronavirus (COVID-19). *Int. J. Surg.* **76**, 71–76 (2020).
3. Paterlini, M. Lockdown in Italy: personal stories of doing science during the COVID-19 quarantine. *Nature: Career Feature* <https://doi.org/10.1038/d41586-020-01001-8> (2020).
4. Tobías, A. Evaluation of the lockdowns for the SARS-CoV-2 epidemic in Italy and Spain after one month follow up. *Sci. Total Environ.* <https://doi.org/10.1016/j.scitotenv.2020.138539> (2020).
5. Kraemer, M. U. *et al.* The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science* <https://doi.org/10.1126/SCIENCE.ABB4218> (2020).
6. Gostic, K., Gomez, A. C., Mummah, R. O., Kucharski, A. J. & Lloyd-Smith, J. O. Estimated effectiveness of symptom and risk screening to prevent the spread of COVID-19. *eLife* **9**, e55570 (2020).
7. Keeling, M. J. & Rohani, P. *Modeling Infectious Diseases in Humans and Animals* (Princeton University Press, Princeton, 2011).
8. Anastassopoulou, C., Russo, L., Tsakris, A. & Siettos, C. Data-based analysis, modelling and forecasting of the COVID-19 outbreak. *PLoS One* **15**, e0230405 (2020).
9. Peng, L., Yang, W., Zhang, D., Zhuge, C. & Hong, L. Epidemic analysis of COVID-19 in China by dynamical modeling. arXiv preprint [arXiv:2002.06563](https://arxiv.org/abs/2002.06563) (2020).
10. Magal, P., & Webb, G. Predicting the number of reported and unreported cases for the COVID-19 epidemic in South Korea. *medRxiv* [2020.03.21.20040154](https://doi.org/10.1101/2020.03.21.20040154) (2020).
11. Kucharski, A. J. *et al.* Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *Lancet Infect. Dis.* **20**, 553–558 (2020).
12. Liu, Y., Gayle, A. A., Wilder-Smith, A. & Rocklöv, J. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *J. Travel Med.* **27**, taaa021 (2020).
13. Liu, Z., Magal, P., Seydi, O. & Webb, G. Predicting the cumulative number of cases for the COVID-19 epidemic in China from early data. arXiv preprint [arXiv:2002.12298](https://arxiv.org/abs/2002.12298) (2020).
14. Park, S. W., Cornforth, D. M., Dushoff, J. & Weitz, J. S. The time scale of asymptomatic transmission affects estimates of epidemic potential in the COVID-19 outbreak. *medRxiv* [2020.03.09.20033514](https://doi.org/10.1101/2020.03.09.20033514) (2020).
15. Wu, J. T., Leung, K. & Leung, G. M. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *Lancet* **395**, 689–697 (2020).
16. Sameni, R. Mathematical modeling of epidemic diseases; a case study of the COVID-19 coronavirus. arXiv preprint [arXiv:2003.11371](https://arxiv.org/abs/2003.11371) (2020).
17. Pan, J. *et al.* Effectiveness of control strategies for coronavirus disease 2019: a SEIR dynamic modeling study. *medRxiv* [2020.02.19.20025387](https://doi.org/10.1101/2020.02.19.20025387) (2020).
18. Djidjou-Demasse, R., Michalakis, Y., Choisy, M., Sofonea, M. T. & Alizon, S. Optimal COVID-19 epidemic control until vaccine deployment. *medRxiv* [2020.04.02.2004918](https://doi.org/10.1101/2020.04.02.2004918) (2020).
19. Moore, S. E. & Okyere, E. Controlling the transmission dynamics of COVID-19. arXiv preprint [arXiv:2004.00443](https://arxiv.org/abs/2004.00443) (2020).
20. Lenhart, S. & Workman, J. T. *Optimal Control Applied to Biological Models* (CRC Press, Boca Raton, 2007).
21. Biswas, M. H. A., Paiva, L. T. & De Pinho, M. A SEIR model for control of infectious diseases with constraints. *Math. Biosci. Eng.* **11**, 761–784 (2014).
22. Neilan, R. M. & Lenhart, S. An introduction to optimal control with an application in disease modeling. *Model. Parad. Anal. Dis. Transm. Models* **67–82** (2010).
23. Biegler, L. T. An overview of simultaneous strategies for dynamic optimization. *Chem. Eng. Process. Process Intensif.* **46**, 1043–1053 (2007).
24. Floudas, C. A. *Deterministic Global Optimization: Theory, Methods and Applications* Vol. 37 (Springer Science and Business Media, New York, 2013).
25. Mizumoto, K., Kagaya, K., Zarebski, A. & Chowell, G. Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan, 2020. *Eurosurveillance* **25**, 2000180 (2020).
26. Nishiura, H. *et al.* Estimation of the asymptomatic ratio of novel coronavirus infections (COVID-19). *Int. J. Infect. Dis.* **94**, 154–155 (2020).
27. Bai, Y. *et al.* Presumed asymptomatic carrier transmission of COVID-19. *JAMA* **323**, 1406–1407 (2020).
28. Rocklöv, J., Sjödin, H. & Wilder-Smith, A. COVID-19 outbreak on the Diamond Princess cruise ship: estimating the epidemic potential and effectiveness of public health countermeasures. *J. Travel Med.* **27**, taaa030 (2020).
29. Bendavid, E. *et al.* COVID-19 Antibody Seroprevalence in Santa Clara County, California. *medRxiv* [2020.04.14.20062463](https://doi.org/10.1101/2020.04.14.20062463) (2020).
30. Lai, C.-C., Shih, T.-P., Ko, W.-C., Tang, H.-J. & Hsueh, P.-R. Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) and corona virus disease-2019 (COVID-19): the epidemic and the challenges. *Int. J. Antimicrob. Agents* **55**, 105924 (2020).
31. Ferguson, N. M. *et al.* Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand (Tech. Rep, Imperial College London, 2020).
32. Anderson, R. M., Heesterbeek, H., Klinkenberg, D. & Hollingsworth, T. D. How will country-based mitigation measures influence the course of the COVID-19 epidemic?. *Lancet* **395**, 931–934 (2020).
33. Wan, E. A. & Van Der Merwe, R. The unscented Kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*, 153–158 (IEEE, 2000).
34. Hart, W. E. *et al.* *Pyomo-optimization Modeling in Python* Vol. 67 (Springer, New York, 2017).
35. Wächter, A. & Biegler, L. T. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.* **106**, 25–57 (2006).
36. Process Systems Enterprise. general PROcess Modeling System (gPROMS) <http://www.psenterprise.com/products/gproms> (1997–2020).
37. Tawarmalani, M. & Sahinidis, N. V. A polyhedral branch-and-cut approach to global optimization. *Math. Program.* **103**, 225–249 (2005).

Acknowledgements

Partial financial support from the National Science Foundation (NSF) through CAREER Award 1454433 is acknowledged with gratitude. The authors also thank Dr. Richard Pattison (Apeel Sciences) for his insightful comments throughout this study.

Author contributions

All authors conceived the experiments; C.T. and F.L. implemented the optimisation models, conducted the computational experiments, and analysed the results; All authors prepared and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-67459-8>.

Correspondence and requests for materials should be addressed to M.B.

Reprints and permissions information is available at www.nature.com/reprints.

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



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

© The Author(s) 2020