A Tale of Two Doses:
Model Identification and Optimal Vaccination for COVID-19

Raphael Chinchilla [a], Guosong Yang [a], Murat K. Erdal [a], Ramon Costa [b], João P. Hespanha [a]

Abstract— We address the model identification and the computation of optimal vaccination policies for the coronavirus disease 2019 (COVID-19). We consider a stochastic Susceptible–Infected–Removed (SIR) model that captures the effect of multiple vaccine treatments, each requiring a different number of doses and providing different levels of protection against the disease. We show that the inclusion of vaccination data enables the estimation of the state of the model and key model parameters that are otherwise not identifiable. This estimates can, in turn, be used to design strategic approaches to vaccination that aim at minimizing the number of deaths and the economic cost of the disease. We illustrate these results with numerical examples.

I. INTRODUCTION

The COVID-19 pandemic has resulted in one of the fastest vaccine development efforts and largest vaccination programs in history [1]. As of the writing of this paper, three vaccines have received emergency authorization in the US [2], [3], [4]. Two of these require two doses for maximum efficacy, each dose providing some marginal effectiveness [5]; whereas the third one requires a single dose. More candidate vaccines are still in development or in the approval process worldwide [6], [7]. A huge vaccination effort by national governments and international organizations is currently underway [8], [9].

In this paper, we study two distinct but related topics emerging from this massive effort. Vaccination data provides information that we can use to better understand the epidemic. As it affects the behavior of the disease, we can take advantage of it to (uniquely) identify unknown parameters of a Susceptible–Infected–Removed (SIR) model and determine the current state of the pandemic for a given region/county/city. In turn, we can use this information to design optimal vaccination strategies that could lead to an earlier end to the pandemic.

The main contribution of this paper is to show the role that vaccination has both for model identification and for control. We consider an SIR model with time varying parameters. We adopt the approach in our previous work and simultaneously estimate the states and parameters of the SIR model [10] to account for the fact that the daily infection rate, the removal rate, the fraction of reported case and the mortality rate cannot be transferred from one region to the other. This is

because these parameters depend not only on intrinsic biological properties of the virus, but also on societal behaviors that change both in time and space; including whether lockdown measures are being enforced, how willing the population is to wear masks and engage in social distancing, the quality of the treatment dispensed, the efficacy of contact tracing measures, and how often the population is being tested.

In contrast with the epidemic parameters above, the efficacy of a particular vaccine is intrinsically a biological parameter, and can be determined through controlled studies. Knowledge of vaccines efficacy parameters, along with the daily number of vaccinations, can be used to uniquely identify an SIR model, overcoming the lack of identifiablility noted in [10], [11]. As a consequence, we can compute a relatively precise estimate of the current state of the pandemic, which enables the design optimal vaccination.

Having available distinct vaccines that differ in efficacy and number of doses, we want to decide when to use each vaccine. We model this as an optimization problem in which we minimize the sum of the daily number of infected people over an interval of time. Using this optimization criteria turns out to minimize the total number of people who will die of COVID-19. This criteria also minimizes the economic impact of the pandemic as people who are infected are unable to work. This optimization problem is nontrivial because of the multiple constraints involved in a large-scale vaccination effort: a finite daily supply of the different vaccines dues to manufacturing limitations, a finite daily number of administered vaccines due to limitations of the healthcare system, and the waiting time between vaccine doses for multi-dose vaccines.

The rest of the paper is organized as follows. In Section II, we extend the SIR model in [10] to incorporate vaccination with multiple doses. In Section III, we show that when a certain number of parameters are unknown (which include the infection rate, fraction of infected people reported as new cases, etc.), the model is only identifiable if we have a known (and strictly positive) number of daily vaccinations. In Section IV, we present the optimal vaccination problem and the multiple constraints related to it. In Section V we illustrate the identifiability results and the optimal vaccination problem in two scenarios that are particularly relevant as of the writing of this article: when to give the first and the second dose of a two-dose vaccine and how to optimally balance between vaccines with different types of efficacy.

Related Work: Identifiability and state estimation for epidemiological models have been well studied in the literature. It is known that an SIR model with constant rates is

[a] Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, USA; {raphaelchinchilla, guosongyang,m_erdal,hespanha}@ucsb.edu
[b] Federal University of Rio de Janeiro, Brazil; ramon@coep.ufrj.br

not structurally identifiable [12] and this problem continues to hold for time-varying models [10], [11]. However, many models capturing the ever-changing course of the COVID-19 pandemic have been developed despite this identifiability problem [13]. Some of these models deal with the lack of identifiability by using external data such as mobility and mask use [14]. While others use weighted regression models [15], [16], or introduce nonlinearities to capture subtle dependencies in data to help in projection [17].

Inclusion of vaccination in epidemiological models has become a necessity in the face of the afore-mentioned vaccination efforts and can take many different forms. For example, it has been modeled as a percentage of the susceptible population [18], [19] and as a percentage of the newborn population [20].

Availability of the vaccination statistics can also help in identification of the pandemic dynamics. However, this impact largely depends on how the vaccination is modeled. For example, it was shown that the addition of birth-targeted vaccination does not make the model structurally identifiable for unknown vaccine efficacy [20].

The number of vaccines administered can be viewed as a control input to optimize an epidemiological model. The total number of infected and susceptible people is a common choice for the criteria used to develop optimal vaccination strategies. In addition, quadratic penalties have also been included in such criteria to account for side-effects [18], [19].

II. SIR MODEL WITH VACCINATION

Our stochastic SIR model with vaccinations takes the form

$$S(t+1) = S(t) - \nu(t) - \Psi(t), \tag{1a}$$
$$I(t+1) = I(t) + \nu(t) - \rho(t), \tag{1b}$$
$$R(t+1) = R(t) + \rho(t) + \Psi(t), \tag{1c}$$

where $S(t)$ denotes the number of susceptible individuals on day t, $I(t)$ the number of infected individuals, and $R(t)$ the number of removed individuals either through cure or death. These update equations involve three key variables that account for the transitions between states: $\nu(t)$ denotes the number of new infections on day t, $\rho(t)$ the number of patients that transitioned on day t from the infected to the removed state by either cure or death, and $\Psi(t)$ the number of people that transitioned on day t from the susceptible to the removed state through vaccination. The daily numbers of infections and removed patients take the form

$$\nu(t) := \frac{\beta(t)I(t)S(t)}{N_0} + d_\nu(t),$$
$$\rho(t) := \gamma I(t) + d_\rho(t) \tag{2}$$

where N_0 is the total population, $\beta(t)$ the daily infection rate, and γ the removal rate. Both numbers include independent identically distributed zero-mean Gaussian random perturbations $d_\nu(t)$ and $d_\rho(t)$, with standard deviations σ_ν and σ_ρ, respectively, to account for stochastic variability.

We assume available a set of M distinct vaccines and that each of them can be applied multiple times to improve immunity. Specifically, we denote by θ_{ij} the fraction of people that develop enough antibodies with the jth dose of vaccine i to become immune, but had not yet developed enough antibodies with the $(j-1)$th dose of this. For a vaccine i that can be applied up to m_i times, we should thus have $\sum_{j=1}^{m_i} \theta_{ij} \leqslant 1$, with equality if and only if the vaccine is 100% effective after the m_i doses. The definition of θ_{ij} is more clear with a numerical example: if for a given vaccine $\theta_{i1} = 50\%$ and $\theta_{i2} = 40\%$ than the total efficacy of the vaccine after two shots will be 90%.

Denoting by $\psi_{ij}(t)$ the number of people that got the jth dose of the vaccine i on day t, we thus have that

$$\Psi(t) := \sum_{i=1}^{M} \sum_{j=1}^{m_i} \theta_{ij} \psi_{ij}(t),$$

where $\theta_{ij}\psi_{ij}(t)$ denotes the number of people that moved from the susceptible to the removed compartment after receiving the jth dose of vaccine of type i on day t, meaning that they will not spread the disease. For simplicity, we neglect any stochastic component to $\Psi(t)$ since it would simply add to the stochastic components that already appear in (1) through $\nu(t)$ and $\rho(t)$.

The measurement model is the same as in [10]: in most regions of the globe, we have access to the daily number of new cases $y_C(t)$ and the daily number of new deaths $y_D(t)$, modeled as

$$y_C(t) = \phi(t)\nu(t) + w_C(t) \tag{3a}$$
$$y_D(t) = \omega I(t) + w_D(t) \tag{3b}$$

where $w_C(t)$ and $w_D(t)$ are independent identically distributed zero-mean Gaussian random variables with standard deviations σ_C and σ_D; $\phi(t)$ is the fraction of newly infected patients that are reported as new cases on day t; and ω is the fraction of infected people who died on day t. The fraction $\phi(t)$ is typically strictly smaller than one because evidences suggest that COVID-19 has a substantial number of asymptomatic patients and because some people may not report cases to the agencies responsible for collecting the data.

We model the time-varying parameters $\beta(t)$ and $\phi(t)$ as random walks of the form

$$\beta(t+1) = \beta(t) + d_\beta(t) \tag{4a}$$
$$\phi(t+1) = \phi(t) + d_\phi(t) \tag{4b}$$

where $d_\beta(t)$ and $d_\phi(t)$ are independent identically distributed zero-mean Gaussian random variables with standard deviations σ_β and σ_ϕ. It is important to emphasize that (4) is not meant to represent the actual dynamics of β and ϕ but rather a priori distributions consistent with the observations that, in the absence of measurements, the most likely values for $\beta(t+1)$ and $\phi(t+1)$ are $\beta(t)$ and $\phi(t)$, respectively. The specific realizations of β and ϕ depend on a multitude of parameters and will be estimated based on the measurements in (3). Time varying parameters ω and γ were also considered in [10], but do not appear to significantly improve estimation.

III. IDENTIFICATION

The SIR model enforces the conservation of the total number of individuals $S(t) + I(t) + R(t) =: N_0$, $\forall t \geqslant 1$ and therefore its state can be unambiguously represented by just two state variables:

$$U(t+1) = U(t) + \nu(t) + \Psi(t)$$
$$R(t+1) = R(t) + \rho(t) + \Psi(t) \tag{5}$$

where $U(t) := N_0 - S(t) = I(t) + R(t)$ corresponds to the number of "unsusceptible" individuals.

A. Deterministic Identifiability

In this subsection, we study conditions under which the initial state of the SIR model (5) and the unknown parameters in (2)–(5) can be uniquely identified from the measurements in (3). For this analysis, we ignore all stochastic fluctuations due to disturbances and measurement noise, and ask the question of whether the initial state and the SIR parameters could be uniquely determined from noise-free measurements collected over a given small interval of time, so all the parameters can be considered constant over the same interval.

We denote by $\Theta := (\gamma, \beta, \phi, \omega) \in \mathbb{R}^4$ a constant vector of unknown parameters in (2)–(5), and by $\xi(t; R_1, U_1, \Theta, \Psi) \in \mathbb{R}^2$ the tuple of outputs (y_C, y_D) in (3) at time $t \in \mathbb{N}$ associated with the solution to (5) for the initial state $R(1) = R_1 \in \mathbb{R}$ and $U(1) = U_1 \in \mathbb{R}$, parameters Θ, and input sequence $\Psi : \mathbb{N} \to \mathbb{R}$.

Definition 1: A tuple of initial state and parameters $(R_1^*, U_1^*, \Theta^*) \in \mathbb{R}^6$ is said to be *locally identifiable on an interval* $\{1, \ldots, T\}$ *for an input sequence* $\Psi : \mathbb{N} \to \mathbb{R}$ when there exists a neighborhood $W \subset \mathbb{R}^6$ of (R_1^*, U_1^*, Θ^*) such that, for all $(R_1, U_1, \Theta) \in W$, if the outputs satisfy

$$\xi(t; R_1, U_1, \Theta, \Psi) = \xi(t; R_1^*, U_1^*, \Theta^*, \Psi),$$
$$\forall t \in \{1, \ldots, T\}, \quad (6)$$

then $(R_1, U_1, \Theta) = (R_1^*, U_1^*, \Theta^*)$. $\qquad\square$

We first provide a negative result showing that the SIR model is not locally identifiable in the absence of vaccination.

Lemma 1: For all $T \in \mathbb{N}$, if $\Psi(t) = 0$, $\forall t \in \{1, \ldots, T-1\}$, then no tuple of initial states and parameters $(R_1, U_1, \Theta) \in \mathbb{R}^6$ is locally identifiable on the interval $\{1, \ldots, T\}$. $\qquad\square$

Proof of Lemma 1. This result follows from the identifiability result in [10, Sec. 2.3], where it is essentially shown that when $\Psi(t) = 0$, $\forall t \in \{1, \ldots, T-1\}$, then (6) holds for all tuples $(R_1^*, U_1^*, \Theta^*), (R_1, U_1, \Theta) \in \mathbb{R}^6$ with

$$R_1 = cR_1^* + (1-c)N_0, \qquad U_1 = cU_1^* + (1-c)N_0,$$
$$\gamma = \gamma^*, \qquad \beta = \beta^*/c, \qquad \phi = \phi^*/c, \qquad \omega = \omega^*/c$$

for every $c \in (0, 1)$, and by noticing that (R_1, U_1, Θ) can be made arbitrarily close to (R_1^*, U_1^*, Θ^*) by selecting a c close enough to 1. $\qquad\blacksquare$

Next, we provide mild conditions for local identifiability while vaccination is taking place.

Theorem 1: A tuple of initial states and parameters $(R_1, U_1, \Theta) \in \mathbb{R}^6_{\geqslant 0}$ is locally identifiable on the interval $\{1, 2, 3\}$ for an input sequence $\Psi : \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ if

$$\beta, \phi, \omega > 0, \tag{7a}$$
$$S(1) > 0, \tag{7b}$$
$$I(t) > 0 \qquad \forall t \in \{1, 2, 3\}, \tag{7c}$$
$$\Psi(2) \neq \frac{I(2)S(2)}{I(1)S(1)}\Psi(1), \tag{7d}$$

where $S(t) = N_0 - U(t)$ and $I(t) = U(t) - R(t)$ are the susceptible and infected state variables at $t \in \mathbb{N}$. $\qquad\square$

To prove Theorem 1, we study the observability of an augmented system with state $x := (\bar{S}, \bar{I}, \gamma, \beta, \bar{\phi}, \bar{\omega}) \in \mathbb{R}^6$ that includes the state of (5) and the unknown parameters in (2)–(5), transformed by

$$\bar{S} := 1 - \frac{U}{N_0}, \quad \bar{I} := \frac{U - R}{N_0}, \quad \bar{\phi} := \phi N_0, \quad \bar{\omega} := \omega N_0 \tag{8}$$

with input $u := \Psi/N_0 \in \mathbb{R}$ and with output $y := (y_C, y_D) \in \mathbb{R}^2$. The correspondence between the augmented system's state x and (R, U, Θ) is bijective and the transformation (8) is only used to simplify derivation. Ignoring all the stochastic disturbances and noise, the dynamics of this augmented system are given by

$$x(t+1) = f(x(t), u(t)), \tag{9a}$$
$$y(t) = h(x(t)), \tag{9b}$$

where $h(x) := \left[\bar{\phi}\beta\bar{I}\bar{S}, \bar{\omega}\bar{I}\right]'$ and

$$f(x, u) := \left[(1 - \beta\bar{I})\bar{S} - u, \bar{I}(1 - \gamma + \beta\bar{S}), \gamma, \beta, \bar{\phi}, \bar{\omega}\right]'.$$

Consistent with the previous terminology, we denote by $\xi(t; x_1, u) = (\xi_C(t; x_1, u), \xi_D(t; x_1, u))$ the output $y \in \mathbb{R}^2$ in (9b) at time $t \in \mathbb{N}$ associated with the solution to (9a) for the initial state $x(1) = x_1 \in \mathbb{R}^6$ and the input sequence $u : \mathbb{N} \to \mathbb{R}$.

The notion of identifiability introduced in Definition 1 can be equivalently restated as local observability for the augmented system as follows.

Definition 2: An initial state $x_1^* \in \mathbb{R}^6$ of (9) is said to be *locally observable on an interval* $\{1, \ldots, T\}$ *for an input sequence* $u : \mathbb{N} \to \mathbb{R}$ when there exists a neighborhood $W \subset \mathbb{R}^6$ of x_1^* such that, for all $x_1 \in W$, if the outputs satisfy

$$\xi(t; x_1, u) = \xi(t; x_1^*, u) \qquad \forall t \in \{1, \ldots, T\}, \tag{10}$$

then $x_1 = x_1^*$. $\qquad\square$

Definition 2 is strongly inspired by the definition of local weak observability from [21, Def. 2.3]. However, while [21, Def. 2.3] considers a notion of local weak observability for which there is no neighborhood such that (10) holds for all input sequences that take values in an open set, here we consider a stronger property for which there is no neighborhood such that (10) holds for a fixed input sequence. Our choice is motivated by the fact that, in our numerical results, the input is obtained from known vaccination data.

In the following, we provide a sufficient condition for local observability, and then use it to prove Theorem 1. For a fixed time $T \geqslant 1$ and input sequence $u : \mathbb{N} \to \mathbb{R}$, we define the following set of C^∞ functions $\Xi_{u,T} := \{\xi_C(t;\cdot,u), \xi_D(t;\cdot,u) : 1 \leqslant t \leqslant T\}$, and denote by $d\Xi_{u,T}$ the codistribution spanned by the differentials of all functions in $\Xi_{u,T}$ [22, p. 19]. The next result is established along the same lines as [21, Th. 3.1].

Lemma 2: If $\dim d\Xi_{u,T}(x_1^*) = 6$ at an initial state $x_1^* \in \mathbb{R}^6$ of (9), then x_1^* is locally weakly observable on the interval $\{1, \ldots, T\}$ for the input sequence $u : \mathbb{N} \to \mathbb{R}$. $\square$

Proof of Lemma 2. If $\dim d\Xi_{u,T}(x_1^*) = 6$, then the functions in $\Xi_{u,T}$ define a map from $\mathbb{R}^6$ to $\mathbb{R}^{2T}$ that is an immersion [23, p. 96] at x_1^*. Hence there exists a neighborhood $W \subset \mathbb{R}^6$ of x_1^* such that this map, restricted to W, is injective. Consequently, for all $x_1 \in W$, if (10) holds then $x_1 = x_1^*$. ∎

Sketch of Proof of Theorem 1. Consider a tuple of initial states and parameters $(R_1, U_1, \Theta) \in \mathbb{R}_{\geqslant 0}^6$ and an input sequence $\Psi : \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ for (2)–(5) such that the assumptions in (7) hold. Let $x_1 := (\bar{S}_1, \bar{I}_1, \gamma, \beta, \bar{\phi}, \bar{\omega}) \in \mathbb{R}_{\geqslant 0}^6$ and $u : \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ be the corresponding initial state and input sequence for the augmented system (9).

Suppose $\dim d\Xi_{u,3}(x_1) < 6$. Then there exists a vector $v \neq 0 \in \mathbb{R}^6$ such that

$$Dh_C(x_1)v = 0, \tag{11a}$$
$$Dh_D(x_1)v = 0, \tag{11b}$$
$$Dh_C(x_2)Df(x_1)v = 0, \tag{11c}$$
$$Dh_D(x_2)Df(x_1)v = 0, \tag{11d}$$
$$Dh_C(x_3)Df(x_2)Df(x_1)v = 0, \tag{11e}$$
$$Dh_D(x_3)Df(x_2)Df(x_1)v = 0, \tag{11f}$$

where $h(x) =: (h_C(x), h_D(x))$ and, for brevity, we denote by x_2 and x_3 the state of (9) at time $t = 2$ and $t = 3$, respectively, with the initial state x_1 and input sequence u. First, the conditions (11a)–(11d) and (11f), together with the assumptions (7a), (7b), and (7c) for $t = 1, 2$, imply that

$$v = c \left[\frac{\beta \bar{I}_1 \bar{S}_1}{\beta \bar{I}_1 \bar{S}_1 + u(1)} \bar{S}_1, \bar{I}_1, 0, -\frac{\beta \bar{I}_1 \bar{S}_1}{\beta \bar{I}_1 \bar{S}_1 + u(1)} \beta, -\bar{\phi}, -\bar{\omega} \right]' \tag{12}$$

for some scalar $c \neq 0$. Second, (12) together with the assumptions (7c) for $t = 3$ and (7d), implies that (11e) cannot hold for a nonzero vector v. Consequently, we have $\dim d\Xi_{u,3}(x_1) = 6$. Then Lemma 2 implies that x_1 is locally weakly observable on the interval $\{1, 2, 3\}$ for the input sequence u, and thus (R_1, U_1, Θ) is locally identifiable on the interval $\{1, 2, 3\}$ for the input sequence Ψ. ∎

B. Estimating the States and Parameters

Using the assumption that the disturbances and measurement noises are independent with standard deviations $V = (\sigma_\nu, \sigma_\rho, \sigma_C, \sigma_D, \sigma_\beta, \sigma_\phi)$, the joint probability density function of the state-parameters vector

$$Z := (U(1), \ldots, U(T), R(1), \ldots, R(T),$$
$$\beta(1), \ldots, \beta(T), \phi(1), \ldots, \phi(T), \omega, \gamma)$$

and the measurement vector

$$Y := (y_C(1), \ldots, y_C(T), y_D(1), \ldots, y_D(T)),$$

is given by

$$p_{Z,Y}(Z, Y; V) = \prod_{t=1}^{T-1} p_D\left(y_D(t) - \omega\, I(t) \right)$$
$$p_\nu\left(U(t+1) - U(t) - \bar{\nu}(t) - \Psi(t) \right) p_\beta\left(\beta(t+1) - \beta(t) \right)$$
$$p_\rho\left(R(t+1) - R(t) - \bar{\rho}(t) - \Psi(t) \right) p_\phi\left(\phi(t+1) - \phi(t) \right)$$
$$p_C\left(y_C(t) - \phi(t)\left(U(t+1) - U(t) - \Psi(t) \right) \right) \tag{13}$$

where $\bar{\nu}(t) := \beta(t)\,(U(t) - R(t))\,(N_0 - U(t))/N_0$, $\bar{\rho}(t) := \gamma\,(U(t) - R(t))$ and $p_\nu(\cdot)$, $p_\rho(\cdot)$, $p_C(\cdot)$, $p_D(\cdot)$, $p_\beta(\cdot)$ and $p_\phi(\cdot)$ are the probability density functions of d_ν, d_ρ, w_C, w_D, d_β and d_ϕ, respectively.

It was shown in [10] that if the conditional distribution $p_{Z|Y}(\cdot)$ is a multivariable Gaussian, then the maximum likelihood estimator for V can be obtained using

$$\hat{V} = \arg\max_{V} \left[-\frac{1}{2} \log\det\left(\frac{d^2}{dZ^2} \log p_{Z,Y}(\hat{Z}, Y; V) \right) \right.$$
$$\left. + \max_{Z} \log p_{Z,Y}(Z, Y; V) \right]$$

and the associated minimum variance estimator for Z is given by $\hat{Z} := \max_{Z} \log p_{Z,Y}(Z, Y; \hat{V})$. This motivates the use of Algorithm 1 which alternates between computing the minimum variance estimate for the state Z and the maximum likelihood estimate for the standard deviations on V, to obtain the estimate $\hat{Z}$.

Algorithm 1 Alternating descent

Require: Initial estimates $\hat{Z}^{(0)}$, $\hat{V}^{(0)}$, a tolerance ϵ
1: **while** $\left\| \hat{V}^{(k)} - \hat{V}^{(k-1)} \right\| \leqslant \epsilon$ **do**
2: $\quad \hat{Z}^{(k+1)} = \arg\max_{Z} \log p_{Z,Y}(Z, Y; \hat{V}^{(k)})$
3: $\quad \hat{V}^{(k+1)} = \arg\max_{V} \left[\log p_{Z,Y}(\hat{Z}^{(k+1)}, Y; V) \right.$
4: $\quad\quad \left. -\frac{1}{2}\log\det\left(\frac{d^2}{dZ^2} \log p_{Z,Y}(\hat{Z}^{(k+1)}, Y; V) \right) \right]$
5: **end while**

IV. OPTIMAL VACCINATION

Our goal is to determine optimal vaccination strategies to minimize a cost of the form

$$J := \sum_{t=T+1}^{T+P} I(t), \tag{14}$$

where $\{1, \ldots, T\}$ is a past time interval for which we have measurements available statisfying the identifiability assumptions in Theorem 1, and $\{T+1, \ldots, T+P\}$ is a future time interval over which we want to optimize the vaccination input $\Psi(t)$. The criterion (14) is evaluated along solutions to

(1)–(4), with future values for the zero-mean disturbances replaced by their most likely values of zero.

Minimizing J is equivalent to minimizing ωJ which is the total number of people who will die of COVID-19. Denoting by μ the average economic cost associated with one infected individual not being able to work for one day, minimizing J also minimizes the total economic cost of the pandemic which is given by μJ.

The selection of the vaccination input Ψ to (1)–(4) to minimize (14) must respect several constraints, including daily constraints on vaccine availability of the form

$$\sum_{s=1}^{t}\sum_{j=1}^{m_i}\psi_{ij}(s) \leqslant \sum_{s=1}^{t}A_i(s), \quad \forall i \in \{1,\dots,M\}, \, t \geqslant T+1,$$
(15)

where $A_i(t)$ denotes the number of vaccines of type i newly available on day t and m_i is the number of doses of vaccine i; and daily constraints on vaccine administration of the form

$$\sum_{i=1}^{M}\sum_{j=1}^{m_i}\psi_{ij}(t) \leqslant B(t), \quad \forall t \geqslant T+1;$$
(16)

where $B(t)$ denotes the maximum number of vaccinations supported by the healthcare system on day t. For multi-dose vaccines that require τ_{ij} days between the administration of doses j and $j+1$, we also need to enforce that

$$\sum_{s=1}^{t}\psi_{i(j+1)}(s) \leqslant \sum_{s=1}^{t-\tau_{ij}}\psi_{ij}(s),$$
$$\forall i \in \{1,\dots,M\}, j \in \{1,\dots,m_i-1\}, t \geqslant T+1. \quad (17)$$

Assuming that vaccines are only given to people that have not been infected, we need to maintain estimates of how many people that have been vaccinated remain uninfected. We do this by splitting $S(t)$ into compartments $S_{ij}(t)$, $i \in \{1,\dots,M\}, j \in \{1,\dots,m_i\}$, where $S_{ij}(t)$ denotes the number of people that received the jth dose of vaccine i but remain susceptible. The dynamics of each compartment are given by

$$S_0(t+1) = S_0(t) - \frac{\beta(t)I(t)}{N_0}S_0(t) - \sum_{i=1}^{M}\psi_{i1}(t) \quad (18a)$$

$$S_{ij}(t+1) = S_{ij}(t) - \frac{\beta(t)I(t)}{N_0}S_{ij}(t) \quad (18b)$$
$$+ (1-\theta_{i1}-\dots-\theta_{ij})\Big(\psi_{ij}(t)-\psi_{i(j+1)}(t)\Big)$$

For a feasible vaccination policy, the daily vaccination number must be chosen to keep $S_0(t)$ and all the $S_{ij}(t)$ in the interval $[0, N_0]$.

V. NUMERICAL EXPERIMENTS

A. Identification of synthetic data

To verify our identification results, we have tested Algorithm 1 in several synthetic scenarios. In this section we present the results to a particular numerical example.

In Figure 1a we present simulated data that incorporates some of the key challenges present in real world data: β was
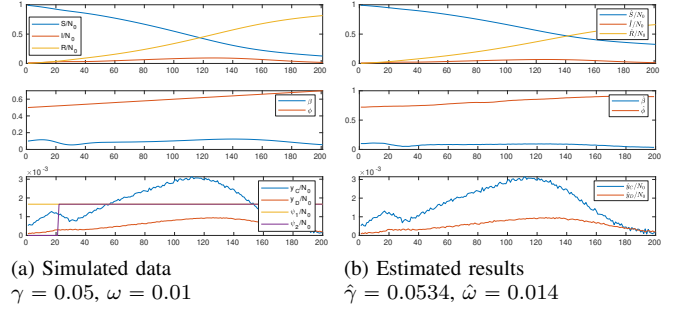

(a) Simulated data
$\gamma = 0.05$, $\omega = 0.01$

(b) Estimated results
$\hat{\gamma} = 0.0534$, $\hat{\omega} = 0.014$

Fig. 1.   Identification of synthetic data

(a) Data for cases, deaths and vaccination

(b) Estimated results
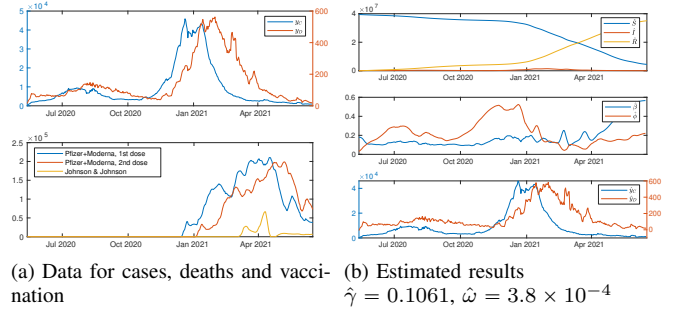$\hat{\gamma} = 0.1061$, $\hat{\omega} = 3.8 \times 10^{-4}$

Fig. 2.   Identification from California data of May 2020 to June 2021

chosen to reflect two waves of infection, with the infection peak separated of about 100 days. We also chose ϕ to be time varying and generally increasing which would be consistent with reporting becoming more efficient as the pandemic evolves. We use only one type of vaccine with two doses and parameters $\theta_{11} = 0.50$ and $\theta_{12} = 0.30$.

In Figure 1b we see that Algorithm 1 is able to accurately estimate the states of the SIR model and to capture the variations of β and ϕ. The latter is slightly over estimated but it still captures the general trends.

B. Identification with real data

We used our algorithm to identify the parameters for the US state of California. The data on cases, deaths and vaccination was obtained from the California Department of Public Health (CDPH). The efficacy of the Johnson and Johnson (Jansen) vaccine was taken to be 66% based on the Centers for Disease Control and Prevention (CDC) Grading of Recommendation [24]. CDPH only provides the combined number of Moderna and Pfizer doses. As their total efficacy is similar [24], we assumed both have the same per dose efficacy of 72% for the first dose and 14% for the second dose [5]. As the early pandemic data tends to be less reliable, our estimation window starts in mid May. The estimation window ends late June because, as the Delta variant started to dominate the pandemic in the United States, our model vaccination is less accurate.

The result of the estimation can be seen in Figure 2 which seem to indicate following conclusions. First, from the values of $\hat{\phi}$, it seems that the majority of cases were either not reported or asymptomatic. Second, the increased
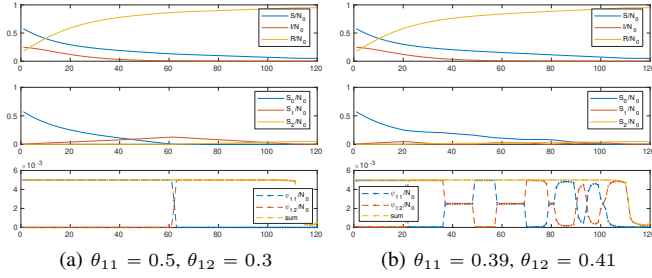
Fig. 3. When to apply the first and the second dose for two values θ_{11} and θ_{12} when $A_1(t) = B(t) = 5 \times 10^{-3}$

value of $\hat{\beta}$ in the end of the estimation window indicated some early signs of concern with regard to the new variants. June 15th was also California reopening, this might also have influenced on the increase in cases.

C. When to give the second dose

Different country/state/regions have followed different vaccination strategies. Some, as the United Kingdom, prioritized giving the first dose to most people before giving the second dose [25]. Others, such as France, started by prioritizing giving both doses to most people but later modified their approach [26]. Determining an optimal allocation of first and second dose for a single type of vaccine can be obtained by minimizing (14) subject to the constraints (15)–(18) with $M = 1$ and $m_1 = 2$.

We have solved this problem numerically for a variety of parameters and have observed essentially two vaccination patterns, illustrated in Figure 3.

The first case, shown in Figure 3a, corresponds to a scenario in which the first dose of the vaccine is more efficient than the second dose (*i.e.,* $\theta_{11} > \theta_{12}$), which is the case for most vaccines currently in the market [5]. In this case, the optimal strategy is to give the maximum number of first dose to most people, and only start to give the second dose after running out of non-vaccinated susceptible individuals. This results is consistent with what has been done in the United Kingdom which has seen a dramatically decrease in new cases [27].

The second case, shown in Figure 3b, corresponds to a scenario in which the second dose is more efficient (*i.e.,* $\theta_{12} > \theta_{11}$). The optimal strategy is now to give the first and the second dose to as many people as possible. The switching observed in the graph has to do with having to give the first dose and then waiting for some time before giving the second.

D. Balacing between availability and efficacy

As more vaccines become available, countries have to decide which one to administer. The decision is particularly difficult with two types of vaccines: one that is less efficient but more available, while the other is less available but more efficient.We will focus on a particular scenario where decision makers can chose between a vaccine that requires two doses and another that requires only one dose.

Which corresponds to the minimization in (14) subject to the constraints (15)–(18) with $M = 2$, $m_1 = 2$ and $m_2 = 1$.

Similar to the previous scenario, we have solved this optimization numerically for a variety of parameters and initial conditions and have observed the four patterns illustrated in Figure 4.

The first pattern is obtained when $\theta_{11} > \max\{\theta_{12}, \theta_{21}\}$, and is shown in Figure 4a. In this case, similar to what happened in Figure 3a, the whole supply of vaccines of type $i = 1$ is used to give the first dose. Because its supply $A_1(t)$ is smaller than the total healthcare capacity $B(t)$, the rest of this capacity is used to give the vaccine of type $i = 2$. This is consistent with what most medical experts have been stating: it is better to give a less effective vaccine (in this case of type $i = 2$) rather than not giving a vaccine while waiting for larger supplies of the better vaccine [28].

The second and third types of behavior are obtained when $\theta_{21} > \max\{\theta_{11}, \theta_{12}\}$, and are shown in Figures 4b and 4d. In this case, the optimal strategy is to give as much vaccine of type $i = 2$ as possible and distribute the rest of the daily healthcare capacity $B(t)$ as in the single vaccine scenario. Notice that we obtain this result even though $\theta_{21} < \theta_{11} + \theta_{12}$, meaning that the efficacy of the single-dose vaccine is smaller than the efficacy of the two-dose vaccine with both doses. This is one of the most surprising results, but it is consistent with the exponential behavior of the SIR system: it is better to act now and remove just a few susceptible individuals rather waiting τ_1 days (in this case, 21 days) before being able to remove a larger fraction of an ever shrinking number of susceptible people.

Finally, the last type of behavior is obtained when $\theta_{12} > \max\{\theta_{11}, \theta_{21}\}$, and are shown in Figure 4c (a slightly different behavior is obtained if $\theta_{21} > \theta_{11}$, but it is qualitatively the same.). This case is also very similar to the single vaccine type of behavior in the sense that the optimal strategy is to give priority to the vaccine of type $i = 1$ and use the rest of the total daily healthcare capacity $B(t)$ for the vaccine of type $i = 2$.

VI. Conclusion and Future Works

The colossal vaccination effort in the beginning of 2021 has been the light at the end of the tunnel for the pandemic. The results in this paper suggest that the vaccination strategy adopted by some countries such as the United Kingdom is the most effective at minimizing the pandemic cost. These results also suggest that the reasoning "a bad (but still safe) vaccine is better than no vaccine" is correct and that we should not restrain ourselves to give the maximum number of vaccines as soon as possible while waiting for potentially better vaccines.

Future work needs to address some of the limitations of our model. The first one is that we assume that only the susceptible individuals are going to be vaccinated. As of now, most countries are generally giving priority to people who have not contracted the disease but as the vaccination efforts progress, the number of removed people who get vaccinated will eventually increase. The second limitation is
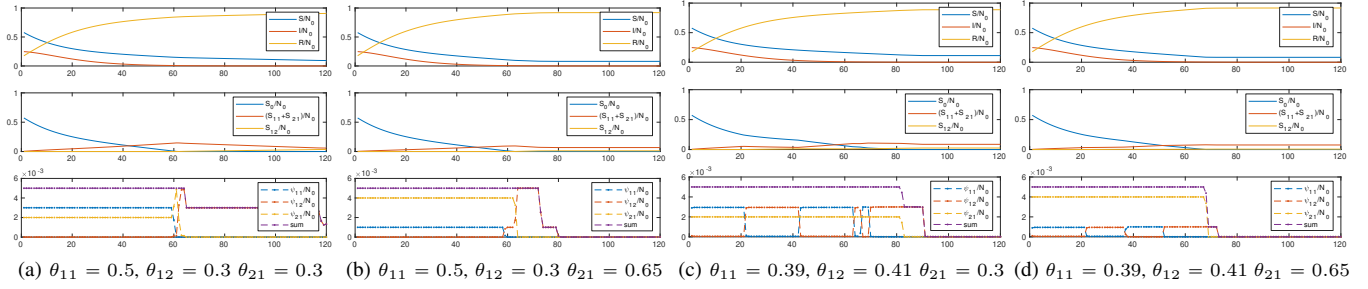
(a) $\theta_{11} = 0.5$, $\theta_{12} = 0.3$ $\theta_{21} = 0.3$ (b) $\theta_{11} = 0.5$, $\theta_{12} = 0.3$ $\theta_{21} = 0.65$ (c) $\theta_{11} = 0.39$, $\theta_{12} = 0.41$ $\theta_{21} = 0.3$ (d) $\theta_{11} = 0.39$, $\theta_{12} = 0.41$ $\theta_{21} = 0.65$

Fig. 4. When to apply the first and the second dose or the one dose vaccine for different values θ when $B(t) = 5 \times 10^{-3}$, $A_1(t) = 3 \times 10^{-3}$ and $A_2(t) = 4 \times 10^{-3}$.

that people who get infected after the first or the second dose have lighter symptoms and may be cured faster. This could be addressed by subdividing I into multiple compartments, each associated to a specific vaccination status, with different recovery rates.

Ending of A Tale of Two Cities by Charles Dickens: I see a beautiful city and a brilliant people rising from this abyss, and, in their struggles to be truly free, in their triumphs and defeats, through long years to come, I see the evil of this time and of the previous time of which this is the natural birth, gradually making expiation for itself and wearing out.

References

[1] T. T. Le, Z. Andreadakis, A. Kumar, R. G. Román, S. Tollefsen *et al.*, "The COVID-19 vaccine development landscape," *Nature Reviews Drug Discovery*, vol. 19, no. 5, pp. 305–306, 2020.

[2] S. E. Oliver, J. W. Gargano, M. Marin, M. Wallace, K. G. Curran *et al.*, "The advisory committee on immunization practices' interim recommendation for use of Pfizer-BioNTech COVID-19 vaccine — United States, December 2020," *MMWR Morbidity and Mortality Weekly Report*, vol. 69, no. 50, pp. 1922–1924, 2020.

[3] ——, "The advisory committee on immunization practices' interim recommendation for use of Moderna COVID-19 vaccine — United States, December 2020," *MMWR Morbidity and Mortality Weekly Report*, vol. 69, no. 51–52, pp. 1653–1656, 2021.

[4] S. E. Oliver, J. W. Gargano, H. Scobie, M. Wallace, S. C. Hadler *et al.*, "The advisory committee on immunization practices' interim recommendation for use of Janssen COVID-19 vaccine — United States, February 2021," *MMWR Morbidity and Mortality Weekly Report*, vol. 70, no. 9, pp. 329–332, 2021.

[5] V. J. Hall, S. Foulkes, A. Saei, N. Andrews, B. Oguti *et al.*, "Effectiveness of BNT162b2 mRNA vaccine against infection and COVID-19 vaccine coverage in healthcare workers in England, multicentre prospective cohort study (the SIREN study)," https://ssrn.com/abstract=3790399, accessed: Mar. 23, 2021.

[6] World Health Organization, "Draft landscape and tracker of COVID-19 candidate vaccines," https://www.who.int/publications/m/item/draft-landscape-of-covid-19-candidate-vaccines, accessed: Mar. 23, 2021.

[7] C. Zimmer, J. Corum, and S.-L. Wee, "The New York Times coronavirus vaccine tracker," https://www.nytimes.com/interactive/2020/science/coronavirus-vaccine-tracker.html, accessed: Mar. 23, 2021.

[8] H. Pettersson, B. Manley, S. Hernandez, and D. McPhillips, "Tracking Covid-19 vaccinations worldwide," https://www.cnn.com/interactive/2021/health/global-covid-vaccinations/.

[9] K. Kupferschmidt, "Despite obstacles, WHO unveils plan to distribute vaccine," *Science*, vol. 369, no. 6511, p. 1553, 2020.

[10] J. P. Hespanha, R. Chinchilla, R. R. Costa, M. K. Erdal, and G. Yang, "Forecasting COVID-19 cases based on a parameter-varying stochastic SIR model," *Annual Reviews in Control*, 2021.

[11] A. Comunian, R. Gaburro, and M. Giudici, "Inversion of a SIR-based model: a critical analysis about the application to COVID-19 epidemic," *Physica D: Nonlinear Phenomena*, vol. 413, p. 132674, 2020.

[12] N. D. Evans, M. J. Chapman, M. J. Chappell, and K. R. Godfrey, "Identifiability of uncontrolled nonlinear rational systems," *Automatica*, vol. 38, no. 10, pp. 1799–1805, 2002.

[13] "The COVID-19 forecast hub," https://covid19forecasthub.org/, accessed: Mar. 23, 2021.

[14] IHME COVID-19 Forecasting Team, R. C. Reiner *et al.*, "Modeling COVID-19 scenarios for the United States," *Nature Medicine*, 2020.

[15] A. Srivastava, T. Xu, and V. K. Prasanna, "Fast and accurate forecasting of COVID-19 deaths using the SIkJα model," *arXiv preprint arXiv:2007.05180*, 2020.

[16] A. Srivastava and V. K. Prasanna, "Learning to forecast and forecasting to learn from the COVID-19 pandemic," *arXiv preprint arXiv:2004.11372*, 2020.

[17] D. Zou, L. Wang, P. Xu, J. Chen, W. Zhang, and Q. Gu, "Epidemic model guided machine learning for COVID-19 forecasts in the United States," *medRxiv*, 2020.

[18] G. Zaman, Y. H. Kang, and I. H. Jung, "Stability analysis and optimal vaccination of an SIR epidemic model," *BioSystems*, vol. 93, no. 3, pp. 240–249, 2008.

[19] T. K. Kar and A. Batabyal, "Stability analysis and optimal control of an SIR epidemic model with vaccination," *BioSystems*, vol. 104, no. 2-3, pp. 127–135, 2011.

[20] J. D. Chapman and N. D. Evans, "The structural identifiability of susceptible–infective–recovered type epidemic models with incomplete immunity and birth targeted vaccination," *Biomedical Signal Processing and Control*, vol. 4, no. 4, pp. 278–284, 2009.

[21] F. Albertini and D. D'Alessandro, "Observability and forward–backward observability of discrete-time nonlinear systems," *Mathematics of Control, Signals and Systems*, vol. 15, no. 4, pp. 275–290, 2002.

[22] A. Isidori, *Nonlinear Control Systems*, 3rd ed. Springer, 1995.

[23] L. W. Tu, *An Introduction to Manifolds*, 2nd ed. Springer Science & Business Media, 2010.

[24] "GRADE Evidence Tables - Recommendations in MMWR," https://www.cdc.gov/vaccines/acip/recs/grade/table-refs.html, accessed: Sep. 20, 2021.

[25] Joint Committee on Vaccination and Immunisation, "Priority groups for coronavirus (COVID-19) vaccination, 30 December 2020," https://www.gov.uk/government/publications/priority-groups-for-coronavirus-covid-19-vaccination-advice-from-the-jcvi-30-december-2020.

[26] Haute Autorité de Santé, "Modification du schéma vaccinal contre le SARS-CoV-2 dans le nouveau contexte épidémique, 23 janvier 2021," https://www.has-sante.fr/jcms/p_3234097/fr/modification-du-schema-vaccinal-contre-le-sars-cov-2-dans-le-nouveau-contexte-epidemique, in French.

[27] European Center for Disease Prevention and Control, "Data on the daily number of new reported COVID-19 cases and deaths by EU/EEA country," https://www.ecdc.europa.eu/en/publications-data/data-daily-new-cases-covid-19-eueea-country, accessed: Mar. 23, 2021.

[28] "Fauci urges Americans to take any of three Covid vaccines available to them," https://www.theguardian.com/us-news/2021/feb/28/covid-vaccines-us-anthony-fauci, the Guardian, Accessed: Mar. 23, 2021.