

Dynamic planning of a two-dose vaccination campaign with uncertain supplies

Original

Dynamic planning of a two-dose vaccination campaign with uncertain supplies / Calafiore, Giuseppe Carlo; Parino, Francesco; Zino, Lorenzo; Rizzo, Alessandro. - In: EUROPEAN JOURNAL OF OPERATIONAL RESEARCH. - ISSN 0377-2217. - STAMPA. - 304:3(2023), pp. 1269-1278. [10.1016/j.ejor.2022.05.009]

Availability:

This version is available at: 11583/2970343 since: 2022-07-28T11:12:55Z

Publisher:

Elsevier

Published

DOI:10.1016/j.ejor.2022.05.009

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Dynamic planning of a two-dose vaccination campaign with uncertain supplies

Giuseppe Carlo Calafiore^{a,d}, Francesco Parino^a, Lorenzo Zino^{b,*}, Alessandro Rizzo^{a,c}^a Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino 10129, Italy^b Faculty of Science and Engineering, University of Groningen, Nijenborgh 4, Groningen 9747 AG, the Netherlands^c Institute for Invention, Innovation, and Entrepreneurship, New York University Tandon School of Engineering, 6 Metrotech Center, Brooklyn, New York 11201, USA^d Institute of Electronics, Computer and Telecommunication Engineering (IEIIT), National Research Council of Italy, Corso Duca degli Abruzzi 24, 10129 Turin, Italy

ARTICLE INFO

Article history:

Received 4 October 2021

Accepted 6 May 2022

Available online 13 May 2022

Keywords:

OR in medicine

Conic programming and interior point methods

Uncertainty modeling

Epidemics

Vaccination

ABSTRACT

The ongoing COVID-19 pandemic has led public health authorities to face the unprecedented challenge of planning a global vaccination campaign, which for most protocols entails the administration of two doses, separated by a bounded but flexible time interval. The partial immunity already offered by the first dose and the high levels of uncertainty in the vaccine supplies have been characteristic of most of the vaccination campaigns implemented worldwide and made the planning of such interventions extremely complex. Motivated by this compelling challenge, we propose a stochastic optimization framework for optimally scheduling a two-dose vaccination campaign in the presence of uncertain supplies, taking into account constraints on the interval between the two doses and on the capacity of the healthcare system. The proposed framework seeks to maximize the vaccination coverage, considering the different levels of immunization obtained with partial (one dose only) and complete vaccination (two doses). We cast the optimization problem as a convex second-order cone program, which can be efficiently solved through numerical techniques. We demonstrate the potential of our framework on a case study calibrated on the COVID-19 vaccination campaign in Italy. The proposed method shows good performance when unrolled in a sliding-horizon fashion, thereby offering a powerful tool to help public health authorities calibrate the vaccination campaign, pursuing a trade-off between efficacy and the risk associated with shortages in supply.

© 2022 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

In response to the COVID-19 health crisis, we witnessed an unprecedented mobilization of the scientific community. Within this joint effort, mathematical models and operational research (OR) have emerged as valuable tools to help assist public health authorities in their complex decisions during epidemic outbreaks (Nowzari, Preciado, & Pappas, 2016; Zino & Cao, 2021). In particular, during the first stages of the COVID-19 pandemic, they have been supportive in forecasting the epidemic spread (Calafiore, Novara, & Possieri, 2020; Parino, Zino, Porfiri, & Rizzo, 2021b; Taylor & Taylor, 2021; Truszkowska, 2021), including local outbreaks (Chang & Kaplan, 2021), understanding how to allocate patients to facilities (Hosseini-Motlagh, Samani, & Homaei, 2021), and

helping to devise nonpharmaceutical interventions (Eryarsoy, Shahmanzari, & Tanrisever, 2022; Fischetti, Fischetti, & Stoustrup, 2021; Tsilianni, Tsiliannidis, & Tsiliannidis, 2022). With the availability of vaccines, new challenges have emerged concerning the planning of an unprecedented mass vaccination campaign. Also in this phase, mathematical models and OR have been key to help assist public health authorities (Choi, 2021).

Most of the OR literature regarding vaccination has focused on optimizing the vaccine supply chain (Duijzer, van Jaarsveld, & Dekker, 2018). For instance, in Lin, Zhao, & Lev (2022), the authors focus on matching uncertainties on the supply-side and on the demand-side; in Arifoğlu & Tang (2022), an incentive program for coordinating the supply chain is proposed for the seasonal Flu vaccine; and in De Boeck, Decouttere, Jónasson, & Vandaele (2021), supply chains in resource-limited settings are optimized using a spatial model. Also, from an epidemic modeling perspective, the design of optimal vaccination strategies is a largely debated topic

* Corresponding author.

E-mail address: lorenzo.zino@rug.nl (L. Zino).

and a whole body of literature has been developed toward encapsulating vaccines and their effect in epidemic models. In these models, vaccinations are incorporated by means of additional compartments in population or network models, or via game-theoretic mechanisms to capture individual-level decisions on whether to take the vaccine shot or not (Wang et al., 2016).

Despite these efforts, the 2021 COVID-19 vaccination campaign has laid bare unprecedented difficulties in planning a massive vaccine rollout that entails the administration of two vaccine doses in a short time, facing limited stocks and uncertain supplies, especially in the early phases of the campaign. Many of the works modeling the COVID-19 vaccination campaign mimic the effect of immunization by encapsulating a simplistic implementation of the vaccination process into the disease progression model. In Giordano (2021); Sinha, Kumar, & Chandra (2021), one-dose vaccination processes are used to investigate the prioritization of the recipients of the vaccine. Leveraging a coordinate descent algorithm, Bertsimas, Digalakis, Jacquillat, Li, & Previero (2022) study how to optimally locate mass vaccination centers. The problem of optimally scheduling vaccination appointments given a set of spatially-located vaccination sites is treated in Zhang, Li, Cao, & Wen (2022) by formulating a mixed-integer linear program and approximating the solution via heuristics.

Another critical issue concerns devising inventory policies to optimize the completion of the recommended two-dose administration. In Parino, Zino, Calafiore, & Rizzo (2021a), a model predictive control approach is used to optimally design two-dose vaccination rollouts in a deterministic framework. A similar problem is considered in Mak, Dai, & Tang (2021) using a deterministic compartmental model. Other studies assess the effectiveness of delaying the second vaccine dose (Moghadas, 2021; Tuite, Fisman, Zhu, & Salomon, 2021). While these works provide interesting insight, they rely on the limiting assumption that vaccine supplies are deterministic. To the best of our knowledge, the problem of a two-dose vaccine allocation with a stochastic supply is treated only in Shumsky, Smith, Hoen, & Gilbert (2021). Therein, the problem is formulated as a stochastic dynamic program, for which some insight into the structure of optimal solutions can be gained. However, the size of the state space and the complexity of the stochastic dynamic program make the computation of optimal solutions unfeasible, limiting the findings to heuristics.

To fill in this gap, here we present an effectively-solvable optimization framework for devising a two-dose vaccination campaign in the presence of uncertain supplies. First, we introduce a linear programming (LP) approach for dealing with the ideal, deterministic scenario in which the time profile of the vaccine supplies is assumed to be exactly known in advance. Such a simplified approach allows us to discuss the constraints, including the bounds on the interval between the administration of the first and the second dose of the vaccine, and to formulate an objective function that accounts for the effectiveness of the partial immunization achieved with a single dose of the vaccine (Polack, 2020; Voysey, 2021). Then, we discuss a more realistic scenario in which random uncertainty affects the vaccine supply, toward the definition of a stochastic optimization problem with chance constraints and affine decision variables to address such uncertainty in supply. This problem is then shown to be representable in the form of a convex second-order cone program (SOCP), and hence amenable to an efficient global solution (Lobo, Vandenberghe, Boyd, & Lebret, 1998).

We use the proposed optimization framework to investigate the optimal vaccination policies for the 2021 COVID-19 vaccination campaign. In particular, we calibrate the model on the Italian vaccination campaign, utilizing publicly available data to fit the uncertainty model for the supplies (Vaccini, 2021), and we consider the two most commonly used vaccines, which have different characteristics in terms of effectiveness and administration timing: Comir-

naty by Pfizer-BioNTech and Vaxzevria by AstraZeneca (Polack, 2020; Voysey, 2021). Besides demonstrating the use of the proposed optimization framework in a real-world scenario, our findings allow us to contribute to the fast-growing field of vaccine distribution in OR by discussing some important properties of the optimal solution. In particular, we focus on the effectiveness of implementing sliding-horizon schemes, whereby the optimization problem is solved over a shorter time-window than the entire optimization horizon, the optimal solution is implemented, and then re-evaluated at fixed time-steps. Furthermore, we analyze the sensitivity of the solution to the uncertainty in the supply. Finally, we discuss how the proposed optimization framework can be utilized by public health authorities to trade-off efficacy of the vaccination campaign and the risk of setbacks due to shortages in the supply.

Our contribution is threefold. First, we propose a novel optimization framework to address the so far overlooked challenge of optimally devising a two-dose vaccination campaign in the presence of uncertain supplies. The flexibility of our optimization framework allows for its applicability to many real-world scenarios. Second, we provide a mathematical formulation of the optimization problem in the form of a SOCP. Such a formulation enables the use of efficient and computationally-effective solvers and algorithms to derive the optimal solution of the optimization problem. Third, we discuss a relevant application of the proposed framework to the 2021 COVID-19 vaccination campaign in Italy, illustrating how the framework allows for a straightforward calibration of the model parameters to real-world data and discussing some important properties of the optimal solution and of the real-world implementation of our OR method.

2. Problem setting

2.1. Notation

We denote the set of real, nonnegative real, nonnegative integer, and strictly positive integer numbers as \mathbb{R} , $\mathbb{R}_{\geq 0}$, $\mathbb{Z}_{\geq 0}$, and $\mathbb{Z}_{> 0}$, respectively. A vector \mathbf{x} is denoted with bold lowercase font, with i th entry x_i . A matrix \mathbf{A} is denoted with bold capital font, with A_{ij} denoting the j th entry of the i th row. Vector \mathbf{e}_i is a all-0 vector, except for a 1 in the i th entry. The all-1 and all-0 matrices are denoted by $\mathbf{1}$ and $\mathbf{0}$, respectively, and the identity matrix by \mathbf{I} . When needed, dimensions are denoted as one positive integer index (for vectors) and two indices (for matrices), which are the number of rows and columns, respectively. The superscript \top denotes the transpose operator. Inequalities for vectors or matrices are intended entrywise (i.e., $\mathbf{x} \geq \mathbf{0}$ means $x_i \geq 0, \forall i$). We denote by $\mathbb{P}[\cdot]$ the probability of an event and by $\mathbb{E}[\cdot]$ the expectation of a random variable (RV). We denote by Ψ the cumulative distribution function (CDF) of a Gaussian RV with zero mean and unit variance, and by Ψ^{-1} its inverse (quantile function).

2.2. Problem setting

Due to the practical nature of the problem, which considers the scheduling of vaccination for large populations (e.g., entire countries), we approximate the population as a continuum of individuals, which allows us to avoid all the issues and limitations related to the optimization of integer variables. Let

- $t \in \mathbb{Z}_{\geq 0}$ be a discrete-time index;
- $T \in \mathbb{Z}_{> 0}$ denote the *time-horizon* of the vaccination campaign. The time-horizon is measured in time units of constant (fixed) duration. In the rest of this paper, we will consider one week as a time unit;
- $k_{\min} \in \mathbb{Z}_{> 0}$ and $k_{\max} \in \mathbb{Z}_{> 0}$ denote the *minimum* and *maximum interval* between the two doses of the vaccine (here, in weeks),

respectively (i.e., the second dose should be administered between the beginning of the k_{\min} th week and the beginning of the k_{\max} th week after the first one); $\delta := k_{\max} - k_{\min}$ is the duration of the *time-window* for the second dose;

- $x^e(t) \geq 0$ be the number of individuals in the population who are *eligible* for receiving the vaccine but have not received the first dose at the beginning of week t ;
- $\mathbf{x}^f(t)$ be a k_{\max} -dimensional vector where the component $x_i^f(t) \geq 0$ denotes the number of individuals at the beginning of week t who have received the *first* dose i weeks before, but have not received the second dose yet;
- $x^d(t) \geq 0$ denote the number of individuals who, at the beginning of week t , have received the first dose k_{\max} weeks before or more, but have not received the second dose yet. For these individuals, the second dose is *delayed* by more than the maximum amount of time admitted;
- $x^s(t) \geq 0$ be the number of individuals who have already received the second dose of the vaccine at the beginning of week t ;
- $u^f(t) \geq 0$ be the number of *first doses* administered to individuals counted by the $x^e(t)$ variable during week t ;
- $\mathbf{u}^s(t)$ be a $(\delta + 1)$ -dimensional vector, where the i th component $u_i^s(t) \geq 0$, for $i = 1, \dots, \delta$, denotes the number of *second doses* administered during week t to individuals who have received the first dose $k_{\min} - 1 + i$ weeks before (i.e., those counted by $x_{k_{\min}-1+i}^f(t)$); the entry $u_{\delta+1}^s(t) \geq 0$ denotes the number of second doses administered to individuals who have received the first dose since more than k_{\max} weeks (i.e., those counted by $x^d(t)$);
- $S(t) \geq 0$ be the (cumulative) vaccine *supply* received up to and including the beginning of week t ;
- $c(t) \geq 0$ be the total *capacity* of vaccines that the healthcare system is able to administer during week t .

The dynamics of the vaccination procedure is described by the following $(k_{\max} + 3)$ -dimensional system of linear recurrence equations:

$$\begin{aligned}
 x^e(t+1) &= x^e(t) - u^f(t) \\
 x_1^f(t+1) &= u^f(t) \\
 x_i^f(t+1) &= x_{i-1}^f(t) \quad i = 2, \dots, k_{\min} \\
 x_i^f(t+1) &= x_{i-1}^f(t) - u_{i-k_{\min}}^s(t) \quad i = k_{\min} + 1, \dots, k_{\max} \\
 x^d(t+1) &= x_{k_{\max}}^f(t) - u_{\delta+1}^s(t) + x^d(t) \\
 x^s(t+1) &= x^s(t) + \sum_{i=1}^{\delta+1} u_i^s(t),
 \end{aligned} \tag{1}$$

which are initialized at the beginning of the vaccination campaign with $x^e(1) = N$, where N is the total number of individuals in the population that should be vaccinated, $x_i^f(1) = 0$ for $i = 1, \dots, k_{\max}$, $x^d(1) = 0$, and $x^s(1) = 0$.

The linear dynamics in Eq. (1) can be conveniently cast as

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{b}^f u^f(t) + \mathbf{B}^s \mathbf{u}^s(t), \tag{2}$$

where the *state variable* $\mathbf{x}(t)$ is the $(k_{\max} + 3)$ -dimensional vector $\mathbf{x}(t) := [x^e(t) \ \mathbf{x}^f(t)^\top \ x^d(t) \ x^s(t)]^\top$, the *decision variables* are the scalar $u^f(t)$ and the $(\delta + 1)$ -dimensional vector $\mathbf{u}^s(t)$, and $\mathbf{A} \in \mathbb{R}^{k_{\max}+3, k_{\max}+3}$, $\mathbf{b}^f \in \mathbb{R}^{k_{\max}+3}$, and $\mathbf{B}^s \in \mathbb{R}^{k_{\max}+3, \delta+1}$ are defined as

$$\mathbf{A} := \begin{bmatrix} 1 & \mathbf{0}_{1, k_{\max}-1} & 0 & 0 & 0 \\ 0 & \mathbf{0}_{1, k_{\max}-1} & 0 & 0 & 0 \\ \mathbf{0}_{k_{\max}-1, 1} & \mathbf{I}_{k_{\max}-1, k_{\max}-1} & 0 & 0 & 0 \\ 0 & \mathbf{0}_{1, k_{\max}-1} & 1 & 1 & 0 \\ 0 & \mathbf{0}_{1, k_{\max}-1} & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{b}^f := \begin{bmatrix} -1 \\ 1 \\ \mathbf{0}_{k_{\max}+1, 1} \end{bmatrix}, \quad \mathbf{B}^s := \begin{bmatrix} \mathbf{0}_{k_{\min}+1, \delta+1} \\ -\mathbf{I}_{\delta+1, \delta+1} \\ \mathbf{1}_{1, \delta+1} \end{bmatrix}.$$

3. Optimization problem

3.1. An LP approach for the deterministic scenario

Before introducing the stochastic framework for the uncertain supply case, we first discuss a “nominal” scenario in which supplies are deterministic.

Assumption 1. We assume that the cumulative supply $S(1), \dots, S(T)$ is a deterministic sequence, entirely known at time $t = 1$.

Under this assumption, the problem of designing an optimal two-dose vaccination strategy can be cast as an LP, hence its global optimal solution can be computed with extreme numerical efficiency (Boyd & Vandenberghe, 2004). Our purpose is to devise an optimal vaccination strategy in terms of first and second doses, i.e., designing $u^f(t)$ and $\mathbf{u}^s(t)$, over a given time horizon T , so that an appropriate objective function involving the number of vaccinated individuals over the considered time horizon is maximized, while satisfying suitable constraints on the vaccination process. We start by discussing the constraints.

Nonnegativity. The state variable and the optimization variables should be nonnegative (entrywise), i.e.,

$$\mathbf{x}(t) \geq 0, \quad t = 1, \dots, T + 1; \tag{3}$$

$$u^f(t) \geq 0, \quad t = 1, \dots, T; \tag{4}$$

$$\mathbf{u}^s(t) \geq 0, \quad t = 1, \dots, T. \tag{5}$$

Note that in Eq. (3), nonnegativity is required also at time $T + 1$, to ensure that the control implemented at time T yields a feasible state.

Stock. The cumulative number of doses administered until week t cannot exceed the cumulative vaccine supply received up to the beginning of that week, which can be conveniently formulated as

$$\sum_{k=1}^t u^f(k) + \sum_{k=1}^t \sum_{i=1}^{\delta+1} u_i^s(k) \leq S(t), \quad t = 1, \dots, T. \tag{6}$$

Capacity. The total number of doses administered during week $t = 1, \dots, T$, cannot exceed the weekly capacity, i.e.,

$$u^f(t) + \sum_{i=1}^{\delta+1} u_i^s(t) \leq c(t), \quad t = 1, \dots, T. \tag{7}$$

Delay. Second doses should be performed in the interval between k_{\min} weeks and k_{\max} weeks after the first dose. We may allow for a fraction $\alpha \in [0, 1]$ of the population who received the first dose to wait past the limit k_{\max} weeks for receiving the second one. If excess wait is not allowed, then $\alpha = 0$ is set in the model. This yields the following constraints:

$$x^d(t) \leq \alpha \left(x^d(t) + \sum_{i=1}^{k_{\max}} x_i^f(t) \right), \quad t = 1, \dots, T + 1. \tag{8}$$

For the sake of readability, we define three vectors $\mathbf{x} := [\mathbf{x}(1)^\top \dots \mathbf{x}(T)^\top]^\top$, $\mathbf{u}^f := [u^f(1) \dots u^f(T)]^\top$, and $\mathbf{u}^s := [\mathbf{u}^s(1)^\top \dots \mathbf{u}^s(T)^\top]^\top$ that gather the state and decision variables over the entire duration of the vaccination campaign, respectively. We define the following linear objective function to be maximized:

$$f(\mathbf{x}, \mathbf{u}^f, \mathbf{u}^s) := \sum_{t=1}^T x^s(t) - \gamma \sum_{t=1}^T x^e(t). \tag{9}$$

The first summand of Eq. (9) carries a positive contribution and accounts for the number of fully vaccinated individuals during the entire duration of the vaccination campaign. The second summand, on the other hand, entails a negative contribution and accounts for the population that is still waiting to start the vaccination process. The parameter $\gamma \geq 0$ weights the relative contribution of the two

terms: for $\gamma = 0$, the optimizer maximizes the number of fully vaccinated individuals; for $\gamma \gg 1$, instead, the key objective tends to be the minimization of the population that has still to receive the first dose. Note that, even though the expression in Eq. (9) does not directly depend on the decision variables, such a dependence is indirectly shaped by the evolution of the state \mathbf{x} , which is determined by Eq. (2) and thus depends on the decision variables \mathbf{u}^f and \mathbf{u}^s . To sum up, under Assumption 1, the vaccination problem can be formulated as an LP, in which the linear objective function $f(\mathbf{x}, \mathbf{u}^f, \mathbf{u}^s)$ is maximized subject to the linear constraints in Eqs. (2)–(8).

Remark 1. Eq. (9) can be re-written in order to explicitly consider the number of individuals with one and two vaccine doses. By rewriting the number of unvaccinated individuals $x^s(t)$ as the difference between the total population N and the number of individuals vaccinated with at least one dose, we obtain

$$\begin{aligned} f &= \sum_{t=1}^T x^s(t) - \gamma \sum_{t=1}^T x^e(t) \\ &= \sum_{t=1}^T x^s(t) - \gamma \sum_{t=1}^T \left(N - x^s(t) - \sum_{i=1}^{k_{\max}} x_i^f(t) - x^d(t) \right) \\ &= \sum_{t=1}^T \left[(1 + \gamma)x^s(t) + \gamma \left(\sum_{i=1}^{k_{\max}} x_i^f(t) + x^d(t) \right) \right] - \gamma NT \\ &= (1 + \gamma) \sum_{t=1}^T x^s(t) + \gamma \sum_{t=1}^T \left[\sum_{i=1}^{k_{\max}} x_i^f(t) + x^d(t) \right] - \gamma NT. \end{aligned} \tag{10}$$

The first term accounts for the number of fully vaccinated individuals; the second one for the number of individuals with a single dose; while the last one (i.e., γNT) is a constant that can be obviously dropped in the process of maximization. Hence, if we conveniently set the ratio $\gamma/(1 + \gamma)$ equal to the ratio between the effectiveness of a single dose of vaccine and the one of two doses, then maximizing the function f implies the maximization of the vaccination coverage over time, in terms of the sum of the total number of individuals vaccinated with a single dose multiplied by the effectiveness of a single dose and the total number of individuals with two doses multiplied by the effectiveness of two doses, up to some constants that play no role in the maximization.

3.2. Dealing with supply uncertainty

In the approach discussed in the previous section, we computed the entire optimal vaccination schedule over the horizon $t = 1, \dots, T$, based on an assumed perfect knowledge at $t = 1$ of the entire supply function over $t = 1, \dots, T$ (Assumption 1). However, this assumption may be quite unrealistic in practice, since on the one hand there is uncertainty on the future delivery of vaccines, and on the other hand, there is no necessity of devising a fully “here-and-now” schedule at $t = 1$ for the whole foreseeable future, since as time unrolls we may want to exploit the opportunity of suitably adjusting the schedule on the basis of the actual supply. In the remaining of this section, we discuss how to tackle both these issues, and we eventually formalize the problem of designing an optimal vaccination strategy in the scenario of uncertain supplies as a second-order cone programming (SOCP) problem (Lobo et al., 1998).

In the following, we assume that the (cumulative) supply $S(t)$ at the generic time-instant t is known up to a given level of uncertainty, and we assume that the level of uncertainty (i.e., its standard deviation) may vary with time in a fashion known a priori. This setting is summarized in the following assumption.

Assumption 2. We assume that the cumulative supply is in the form

$$S(t) = \hat{S}(t) + \sigma(t)\delta(t), \quad t = 1, \dots, T \tag{11}$$

where $\hat{S}(t)$ is the given nominal (or expected) time profile of the cumulative supply, $\sigma(t)$ is the given time profile of the standard deviation of the uncertainty, and $\delta(t)$, $t = 1, \dots, T$, is a sequence of independent and identically distributed (IID) RVs with zero mean and unit variance, representing the uncertainty on the cumulative supply at time t .

The uncertainty model in Assumption 2 is able to capture the presence of delays (or earliness) in the actual supply with respect to the scheduled one $\hat{S}(t)$ and relies on the assumption that suppliers, when not able to comply with an order, will try to accommodate it in the next delivery. Hence, a delay (or earliness) at time t does not affect the expected future supplies.

At this stage, we let the decision variables related to the first and second dose to react to changes in the cumulative supply. Specifically, we assume *vaccination policies* to take an affine form, i.e.,

$$\begin{aligned} u^f(t) &= \hat{u}^f(t) + g^f(t)(S(t) - \hat{S}(t)), \\ u^s(t) &= \hat{u}^s(t) + g^s(t)(S(t) - \hat{S}(t)), \end{aligned} \tag{12}$$

where $\hat{u}^f(t) \in \mathbb{R}$ and $\hat{u}^s(t) \in \mathbb{R}^{\delta+1}$ represent the nominal vaccination schedules, and $g^f(t) \in \mathbb{R}$ and $g^s(t) \in \mathbb{R}^{\delta+1}$ represent the adjustment gains, whose purpose is to correct the nominal decisions $\hat{u}^f(t)$ and $\hat{u}^s(t)$ as a consequence to deviations of the cumulative supply from its nominal (or expected) value. Note that the policies in Eq. (12) are fully *causal*, since the decision variables at t depend only on the uncertain supply realization up to t , while the future supplies remain uncertain and are suitably treated via probabilistic constraints, as discussed next. The idea of linear policies, or affine-adjustable variables, has been introduced in the context of robust optimization in Ben-Tal, Goryashko, Guslitzer, & Nemirovski (2004) and has been adopted in several uncertain decision problems, see, e.g., Calafiore (2009).

The decision variables of the optimization problem will now be $\hat{u}^f(t)$, $\hat{u}^s(t)$ and $g^f(t)$, $g^s(t)$, $t = 1, \dots, T$, rather than $u^f(t)$, $u^s(t)$. The optimization problem will therefore provide us with the parameters to plug into Eq. (12) and, from these equations at each t , based on the observation of the actual outcome of $S(t)$, we shall decide for the actual number of first and second dose vaccinations to be delivered during day t .

Under Assumption 2, substituting Eqs. (11) and (12) into the dynamics in Eq. (2), we can determine the recursions that govern the evolution of the stochastic system. Specifically, the state variable $\mathbf{x}(t)$ of the system evolves according to the stochastic recursion

$$\mathbf{x}(t + 1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}^f \hat{u}^f(t) + \mathbf{B}^s \hat{u}^s(t) + \sigma(t) (\mathbf{B}^f g^f(t) + \mathbf{B}^s g^s(t)) \delta(t). \tag{13}$$

Since the state $\mathbf{x}(t)$ is an RV, it might be convenient and useful for the derivations in the following to explicitly write two different expressions for the temporal evolution of its expected value and for the shift of the RV with respect to it. Specifically, if we define the expected value $\hat{\mathbf{x}}(t) := \mathbb{E}[\mathbf{x}(t)]$ and the shift $\tilde{\mathbf{x}}(t) := \mathbf{x}(t) - \hat{\mathbf{x}}(t)$, we have that

$$\hat{\mathbf{x}}(t + 1) = \mathbf{A}\hat{\mathbf{x}}(t) + \mathbf{B}^f \hat{u}^f(t) + \mathbf{B}^s \hat{u}^s(t), \tag{14}$$

and

$$\tilde{\mathbf{x}}(t + 1) = \mathbf{A}\tilde{\mathbf{x}}(t) + \mathbf{b}(t)\delta(t), \tag{15}$$

where the vector $\mathbf{b}(t) \in \mathbb{R}^{k_{\max}+3}$ is defined as

$$\mathbf{b}(t) := \sigma(t) (\mathbf{B}^f g^f(t) + \mathbf{B}^s g^s(t)). \tag{16}$$

Finally, we derive a recursion equation for the covariance matrix of $\mathbf{x}(t)$, defined as $\Sigma(t) := \mathbb{E}[\tilde{\mathbf{x}}(t)\tilde{\mathbf{x}}^\top(t)]$, for which it holds

$$\Sigma(t + 1) = \mathbf{A}\Sigma(t)\mathbf{A}^\top + \mathbf{b}(t)\mathbf{b}^\top(t). \tag{17}$$

Due to the stochasticity of the system, the constraints should be expressed in terms of probabilities (Calafiore & Ghaoui, 2006; Prékopa, 1995). To this aim, we fix the risk parameter $\varepsilon \in (0, 1/2)$, which represents the probability that a constraint is violated, and we cast the constraints as follows.

Nonnegativity. Since the state and the number of doses performed are stochastic, we re-formulate Eqs. (3)–(5) in a probabilistic fashion as

$$\mathbb{P}[x_i(t) \geq 0] \geq 1 - \varepsilon, \quad t = 1, \dots, T + 1, \quad i = 1, \dots, k_{\max} + 3, \quad (18)$$

$$\mathbb{P}[u^f(t) \geq 0] \geq 1 - \varepsilon, \quad t = 1, \dots, T, \quad (19)$$

$$\mathbb{P}[u_i^s(t) \geq 0] \geq 1 - \varepsilon, \quad t = 1, \dots, T, \quad i = 1, \dots, \delta + 1, \quad (20)$$

Stock. Similarly, we impose that

$$\mathbb{P}\left[\sum_{k=1}^t u^f(k) + \sum_{k=1}^t \sum_{i=1}^{\delta+1} u_i^s(k) \leq \hat{S}(t) + \sigma(t)\delta(t)\right] \geq 1 - \varepsilon, \quad t = 1, \dots, T. \quad (21)$$

Capacity. We set

$$\mathbb{P}\left[u^f(t) + \sum_{i=1}^{\delta+1} u_i^s(t) \leq c(t)\right] \geq 1 - \varepsilon, \quad t = 1, \dots, T. \quad (22)$$

Delay. The constraints can be stated on the expected states, i.e.,

$$\hat{x}^d(t) \leq \alpha \left[x^d(t) + \sum_{i=1}^{k_{\max}} \hat{x}_i^f(t) \right], \quad t = 1, \dots, T + 1. \quad (23)$$

Finally, the objective function f in Eq. (9) is now an RV, since it depends on the realization of the stochastic process $\mathbf{x}(t)$. The approach we take now is to maximize its expected value, yielding the following expectation objective

$$J := \mathbb{E}[f] = \mathbb{E}\left[\sum_{t=1}^T x^s(t) - \gamma \sum_{t=1}^T x^e(t)\right] = \sum_{t=1}^T \hat{x}^s(t) - \gamma \sum_{t=1}^T \hat{x}^e(t), \quad (24)$$

where $\hat{\mathbf{x}}(t)$ must satisfy the recursion equation in Eq. (14), and the constraints discussed above. In general, maximizing Eq. (24) might be a difficult task, due to the stochastic nature of the system and, ultimately, the nature of the uncertainties $\delta(t)$. In the following, we will show that, under some reasonable assumptions on the nature of the uncertainties $\delta(t)$, the optimization problem can be cast as a SOCP, and thus easily solved.

3.3. Convex second-order cone problem formulation

From now on, we make the following assumptions about the uncertainties.

Assumption 3. We assume that Assumption 2 holds, and that $\delta(t)$, for $t = 1, \dots, T$ is a sequence of IID Gaussian RVs with zero mean and unit variance.

Under Assumption 3, Eq. (13) yields a sequence of Gaussian random vectors with expected value $\hat{\mathbf{x}}(t)$ and covariance matrix $\Sigma(t)$ (Ross, 2019). Using this property, we derive the following explicit expressions for the constraints.

Nonnegativity (state variable). From Eq. (18), we observe that the generic i th component of the state variable $x_i(t)$ is a Gaussian-distributed RV with expected value $\hat{x}_i(t)$ and variance $\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i$. We have therefore that $\mathbb{P}[x_i(t) \geq 0] = 1 - \mathbb{P}[x_i(t) \leq 0] = 1 - \mathbb{P}\left[\frac{x_i(t) - \hat{x}_i(t)}{\sqrt{\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i}} \leq -\frac{\hat{x}_i(t)}{\sqrt{\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i}}\right] = 1 - \Psi\left(-\frac{\hat{x}_i(t)}{\sqrt{\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i}}\right)$. Thus, $\mathbb{P}[x_i(t) \geq 0] \geq 1 - \varepsilon$ if and only if (iff) $\Psi\left(-\hat{x}_i(t)/\sqrt{\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i}\right) \leq \varepsilon$, which in

turn holds iff $-\hat{x}_i(t)/\sqrt{\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i} \leq \Psi^{-1}(\varepsilon)$. Hence, being $\varepsilon < 1/2$, the generic i th constraint in Eq. (18) is equivalent to

$$|\Psi^{-1}(\varepsilon)|\sqrt{\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i} \leq \hat{x}_i(t). \quad (25)$$

Notice that we have from Eq. (17) that $\Sigma(1) = 0$, and that we can write the following recursion $\Sigma(t) = \sum_{k=1}^{t-1} \mathbf{A}^{t-k-1} \mathbf{b}(k) \mathbf{b}(k)^T (\mathbf{A}^{t-k-1})^T$, $t \geq 2$. Hence, we can write $\mathbf{e}_i^T \Sigma(t) \mathbf{e}_i = \|\mathbf{f}^i\|_2^2$, where $\mathbf{f}^i(t) := [\mathbf{e}_i^T \mathbf{A}^{t-2} \mathbf{b}(1), \mathbf{e}_i^T \mathbf{A}^{t-3} \mathbf{b}(2), \dots, \mathbf{e}_i^T \mathbf{A} \mathbf{b}(t-2), \mathbf{e}_i^T \mathbf{b}(t-1)]^T$. It follows that Eq. (25) is expressed as a second-order cone constraint $|\Psi^{-1}(\varepsilon)|\|\mathbf{f}^i(t)\|_2 \leq \hat{x}_i(t)$, where, by Eq. (16), $\mathbf{f}^i(t)$ is an affine function of the variables $\mathbf{g}^f(k), \mathbf{g}^s(k)$, $k = 1, \dots, t-1$. Hence, the (probabilistic) nonnegativity constraint on the state variables in Eq. (18) can be written as

$$|\Psi^{-1}(\varepsilon)|\|\mathbf{f}^i(t)\|_2 \leq \hat{x}_i(t), \quad t = 2, \dots, T + 1, \quad i = 1, \dots, k_{\max} + 3. \quad (26)$$

Nonnegativity (optimization variables). Eq. (19) yields $\mathbb{P}[u^f(t) \geq 0] = 1 - \mathbb{P}[u^f(t) < 0] = 1 - \mathbb{P}[\hat{u}^f(t) + \sigma(t)|g^f(t)|\delta(t) < 0] = 1 - \mathbb{P}[\delta(t) < -\hat{u}^f(t)/\sigma(t)|g^f(t)|] = 1 - \Psi(-\hat{u}^f(t)/\sigma(t)|g^f(t)|)$. Hence $\mathbb{P}[u^f(t) \geq 0] \geq 1 - \varepsilon$ iff

$$\hat{u}^f(t) \geq v_\varepsilon \sigma(t) |g^f(t)|, \quad t = 1, \dots, T. \quad (27)$$

Similar, from Eq. (20) we derive the explicit constraint

$$\hat{u}_2(t) \geq v_\varepsilon \sigma(t) |g_i^s(t)|, \quad t = 1, \dots, T, \quad i = 1, \dots, \delta + 1. \quad (28)$$

Stock. We rewrite the stock constraints in Eq. (21) as

$$\mathbb{P}\left[\sum_{k=1}^t \left(\hat{u}^f(k) + \sum_{i=1}^{\delta+1} \hat{u}_i^s(k)\right) + \xi(t) \leq \hat{S}(t)\right] \geq 1 - \varepsilon, \quad t = 1, \dots, T, \quad (29)$$

where $\xi(t) := \sum_{k=1}^t \sigma(k) (g^f(k) + \sum_{i=1}^{\delta+1} g_i^s(k)) \delta(k) - \sigma(t) \delta(t)$. We observe that $\xi(t)$ is the linear combination of t independent Gaussian RVs. Hence, also $\xi(t)$ is a Gaussian RV with $\mathbb{E}[\xi(t)] = 0$ and $\mathbb{E}[\xi^2(t)] = \|\mathbf{v}(t)\|_2^2$, where $\mathbf{v}(t) := [\sigma(1)(g^f(1) + \sum_{i=1}^{\delta+1} g_i^s(1)), \sigma(2)(g^f(2) + \sum_{i=1}^{\delta+1} g_i^s(2)), \dots, \sigma(t)(g^f(t) + \sum_{i=1}^{\delta+1} g_i^s(t) - 1)]^T$. The constraints in Eq. (29) are thus rewritten as $\mathbb{P}[\xi(t)/\|\mathbf{v}(t)\|_2 \leq (\hat{S}(t) - \hat{U}(t))/\|\mathbf{v}(t)\|_2] = \Psi((\hat{S}(t) - \hat{U}(t))/\|\mathbf{v}(t)\|_2)$, where we set for brevity $\hat{U}(t) := \sum_{k=1}^t (\hat{u}^f(k) + \sum_{i=1}^{\delta+1} \hat{u}_i^s(k))$. The above probability is greater than or equal to $1 - \varepsilon$ iff $(\hat{S}(t) - \hat{U}(t))/\|\mathbf{v}(t)\|_2 \geq \Psi^{-1}(1 - \varepsilon) = -\Psi^{-1}(\varepsilon) = v_\varepsilon$, whence we obtain the explicit constraint:

$$\hat{U}(t) + v_\varepsilon \|\mathbf{v}(t)\|_2 \leq \hat{S}(t), \quad t = 1, \dots, T. \quad (30)$$

Capacity. The capacity constraints in Eq. (22) are expressed, for all $t = 1, \dots, T$, as $\mathbb{P}[\hat{u}^f(t) + \sum_{i=1}^{\delta+1} \hat{u}_i^s(t) + \sigma(t)\delta(t) \leq c(t)] \geq 1 - \varepsilon$, which, proceeding as in the previous cases, is equivalent to

$$\hat{u}^f(t) + \sum_{i=1}^{\delta+1} \hat{u}_i^s(t) + \sigma(t) v_\varepsilon \left| g^f(t) + \sum_{i=1}^{\delta+1} g_i^s(t) \right| \leq c(t), \quad t = 1, \dots, T. \quad (31)$$

Delay. Finally, we note that the expression in Eq. (23) is already formulated in terms of expected states.

If we stack all the $\hat{u}^f(t), \hat{\mathbf{u}}^s(t), g^f(t)$, and $\mathbf{g}^s(t)$ variables in four vectors $\hat{\mathbf{u}}^f = [\hat{u}^f(1), \dots, \hat{u}^f(T)]^T$, $\hat{\mathbf{u}}^s = [\hat{\mathbf{u}}^s(1)^T, \dots, \hat{\mathbf{u}}^s(T)^T]^T$, $\mathbf{g}^f = [g^f(1), \dots, g^f(T)]^T$, and $\mathbf{g}^s = [\mathbf{g}^s(1)^T, \dots, \mathbf{g}^s(T)^T]^T$, we can formulate the optimization problem as:

$$\max_{\hat{\mathbf{u}}^f, \hat{\mathbf{u}}^s, \mathbf{g}^f, \mathbf{g}^s} J = \sum_{t=1}^T \hat{x}^s(t) - \gamma \sum_{t=1}^T \hat{x}^e(t) \quad (32)$$

s.t.: Eqs. (14), (26), (27), (28), (30), (31), (23).

Remark 2. Problem (32) has $(k_{\max} + \delta + 5) \times T + k_{\max} + 3$ variables $-(k_{\max} + 3) \times (T + 1)$ state variables and $(\delta + 2) \times T$ decision variables— and $(k_{\max} + \delta + 8) \times T + k_{\max} + 4$ constraints, and involves the optimization of a linear objective function subject to linear constraints and the intersection of second-order cones. Hence, it is a SOCP, for which the global optimal solution can be computed in polynomial time via interior point methods (Lobo et al., 1998), and standard and efficient algorithms have been developed, see, e.g., Mosek (Andersen, Roos, & Terlaky, 2003).

Remark 3. Despite the fact that the optimal policy in Eq. (12) is reactive and takes into account the uncertainty, one should consider that, as time unrolls and vaccines are supplied, part of the uncertainty is revealed. Hence, in practical applications, it makes sense to deploy Problem (32) within a sliding-horizon framework.

Fixing a simulation horizon $T_s < T$ and a re-evaluation step $T_r < T_s$, every T_r time steps starting from $t = 1$, one could utilize Problem (32) to compute the optimal vaccination policy for the following T_s time steps and apply it for the first T_r steps. Then, after T_r time-steps, one can re-evaluate the policy by utilizing again Problem (32), with updated initial conditions and uncertainty model (informed by the time-series of supplied observed in the real world), and apply it for the following T_r time-steps. This process can be repeated until the end of the time-horizon of the vaccination campaign. Similar sliding-horizon schemes have been successfully used in the literature of optimization under uncertainty, sometimes under the name of folding horizon. See, e.g., Ben-Tal, Golany, Nemirovski, & Vial (2005); Cohen, Golany, & Shtub (2016); Dabadghao, Marandi, & Roy (2021); De Ruyter, Brekelmans, & den Hertog (2016); Gorissen, Yanıkoğlu, & den Hertog (2015).

Remark 4. Particular combinations of the parameters and input values may result in unfeasible optimization problems. Specifically, the nonnegativity of the optimization variables in Eqs. (19) and (20) and the stock constraint in Eq. (21) may not be simultaneously verified for some values of t . This can typically happen in the first weeks, in the presence of small supplies with high uncertainty. A simple way to address this issue in practice is to fix $u^f(t) = 0$ and $u^s(t) = \mathbf{0}$ for all $t = 1, \dots, k$ and solve the Problem (32) starting from week $t = k + 1$, where $k = \min\{s: \text{Problem (32) is feasible for all } t \in [s + 1, T]\}$.

4. Application to 2021 COVID-19 vaccination campaign in Italy

To demonstrate the proposed optimization framework against real-world data, we use it to investigate the 2021 COVID-19 vaccination campaign in Italy.

4.1. Vaccine parameters

We consider Comirnaty and Vaxzevria, which have been the two mostly used vaccines in the first phases of the vaccination campaign.

Comirnaty. In Italy, the two doses are administered at least $k_{\min} = 3$ weeks apart and at most $k_{\max} = 6$ weeks apart. The efficacy during the interval between the two doses was estimated in clinical trials as 52%, while both doses yield a 95% protection (Polack, 2020). Hence, following Remark 1, we set $\gamma = 1.21$.

Vaxzevria. The Italian Medicines Agency indicates that the administration of the second dose should be ideally performed during the twelfth week and, in any case, at least ten weeks after the first dose. Consequently, we set $k_{\min} = 10$ and $k_{\max} = 12$. Clinical data confirm a vaccine efficacy of 76% after the first dose and 82% after two doses (Voysey, 2021), yielding $\gamma = 12.66$.

4.2. Expected supply and uncertainty model

We use the official data of the daily supply of COVID-19 vaccines in Italy from January 4, 2021 to July 19, 2021 (Vaccini, 2021). We set the discrete time unit equal to one week and we aggregate the data accordingly.

Based on these data, we construct a model for the expected cumulative supply $\hat{S}(t)$. The unfolding of the supply delivered (see Fig. 1A, B) suggests that the size of weekly supplies increases linearly with time. Hence, we suppose a quadratic trend for the expected cumulative supply $\hat{S}(t)$. We model the data by using a weighted least-square technique (Chatterjee & Simonoff, 2013). In particular, since the t th entry of the cumulative supply vector is the sum of t vaccine deliveries, each one characterized by its level of uncertainty, in fitting the weighted least-square model, we associate with the t th data-point of the Italian cumulative supply the weight $1/t$. Using the weighted least-square estimator we obtain $\hat{S}(t) = 65,160t^2 - 2,459t + 814,000$ for Comirnaty and $\hat{S}(t) = 7,309t^2 + 398,300t + 440,800$ for Vaxzevria, where Comirnaty is administrated starting from the week of January 19, while the administration of Vaxzevria starts from the week of February 22. Also, we use the fitted model to estimate the uncertainty associated with the predictions (i.e., the model prediction errors or forecast error), whose values are reported in Table S1 in the Supplementary Material. We use these estimations as $\sigma(t)$, and we finally obtain the supply in the form discussed in Eq. (11), under Assumption 3, i.e., assuming that $\sigma(t)$, for $t = 1, \dots, T$ is a sequence of IID Gaussian RV with zero mean and unit variance. The resulting fitting with the correspondent 95% prediction intervals for the two vaccine supplies are reported in Fig. 1.

Finally, we set the maximum weekly capacity $c(t) = 1.2(\hat{S}(t) - \hat{S}(t - 1))$, i.e., we assume that the healthcare system can accommodate a number of vaccine doses that is slightly greater than the expected weekly supply.

4.3. Analysis of the optimal solutions

We compute the optimal vaccination rollout using the optimization framework described in Section 3.3, with the expected supplies and uncertainty model derived for the 2021 COVID-19 vaccination campaign in Italy. We use the actual time-series of the Italian supply from Vaccini (2021) as the realization of the process. Unless stated otherwise, we set $\varepsilon = 0.05$.

First, we compare the performance of the optimal solution with the baseline performance of the solution of the LP proposed in Section 3.1, computed assuming that the exact time-series of the supplies are known a priori. This solution, even though unrealistic since it ignores the uncertainty, serves as a benchmark for evaluating the performance of the other approaches and the effect of uncertainty on the vaccination campaign. With respect to the LP solution, with both vaccines we observe that the presence of uncertainty may slow down a vaccination campaign. In particular, we observe a 3.60% reduction of the total number of doses administrated for the Comirnaty vaccination campaign. For Vaxzevria, which is characterized by higher levels of uncertainty, the performance are even worse, as we record a 8.61% drop in the number of administered doses.

We next discuss a sliding-horizon implementation of the proposed approach, following Remark 3. We compare different solutions, obtained using different sliding-horizon policies (all with $T_s = 20$): (a) *no re-evaluation*, which is the numerical solution of SOCP computed over the entire horizon; (b) *weekly re-evaluation*, in which the sliding horizon moves on a weekly basis ($T_r = 1$): every week the uncertainty model is re-estimated following the procedure described in Sec. 4.2, appropriately shifting the time index; then, the optimal solution is re-computed by solving the SOCP over




Fig. 1. Italian cumulative supply $S(t)$ for the two types of vaccines (blue) and the fitted model $\hat{S}(t)$, with the 95% confidence intervals (orange and gray, respectively).




Fig. 2. Optimal solution with no re-evaluation (blue), weekly (orange), and monthly re-evaluation (green) for the two vaccines. Panels A–B show the time-series of first doses; C–D of second doses; and E–F the evolution of the number of fully vaccinated individuals.

a time horizon $T_s = 20$ and it is applied for $T_r = 1$ week, before the following re-evaluation; (c) *monthly re-evaluation*, in which the uncertainty model is updated and the strategy is re-evaluated on a 4-week basis ($T_r = 4$); and (d) *bimonthly re-evaluation*, in which the uncertainty model is updated and the strategy is re-evaluated every $T_r = 8$ weeks. In other words, once an optimal solution is found, it will be applied for the next T_r weeks before being re-evaluated. Together with the re-evaluation of the optimal solution based on the increased data on the time-series of the supplies, a key advantage of the sliding-horizon technique resides in the possibility of utilizing such data to also re-calibrate the uncertainty model at each iteration, thereby reducing the level of uncertainty in the supply as the vaccination campaign proceeds.

In Fig. 2A–D, for both vaccines we report the optimal solutions for the cases of no re-evaluation, weekly, and monthly

re-evaluation. We notice that the solutions for the entire horizon result in a smooth trend of first and second doses, while the optimal solutions with the sliding horizon have more discontinuous trends, as they are able to re-evaluate the campaign in response to changes in the supply. Fig. 2E, F shows the temporal evolution of the number of vaccinated population under the different policies, illustrating how the implementation of a sliding horizon could be beneficial. We analyze in detail the different cases in Table 1. We compare the strategy that uses the sliding horizon with the case with no re-evaluation. The comparison is presented as percentage variation of i) the average utility function during the optimization J , which is related to the average coverage of the population during the vaccination campaign, as discussed in Remark 1, and ii) the total number of doses of vaccine administered V during the entire campaign. The results show the beneficial effect of

Table 1
Performance with sliding horizon, with respect to the scenario with no re-evaluation.

	Comirnaty		Vaxzevria	
	ΔJ	ΔV	ΔJ	ΔV
re-evaluation $T_r = 1$ week	2.95%	6.30%	0.26%	7.74%
re-evaluation $T_r = 4$ weeks	1.71%	6.64%	0.15%	5.54%
re-evaluation $T_r = 8$ weeks	1.52%	6.68%	0.13%	3.52%

implementing a sliding-horizon technique and highlight a non-trivial trade-off between performance and re-evaluation periods, whereby more frequent re-evaluations may result in a performance increase.

In particular, the use of a sliding horizon yields an increase in the total number of doses administered. For Comirnaty—which is characterized by lower levels of uncertainty—such an increase is approximately by 6–7% and it seems not to be strongly influenced by the re-evaluation frequency. However, we observe that shorter re-evaluation periods can be beneficial toward increasing the vaccination coverage during the campaign, with increases from 1.52% for bimonthly re-evaluation to 2.95% for weekly re-evaluation. On the other hand, the higher levels of uncertainty in the Vaxzevria supply seem to further increase the benefits of frequent re-evaluations, as also the total number of doses administered seems to be strongly affected by the re-evaluation policy. In fact, the implementation of a sliding-horizon technique increases the doses administered by 7.74% in the case of 1 week re-evaluation.

Summarizing, being able to set up flexible vaccination rollouts that can be re-evaluated in order to accommodate potential changes in the expected supply by utilizing the proposed optimization framework with a sliding horizon can bring to a considerable improvement of the performance of the vaccination campaign. In addition, the frequency of re-evaluation seems key in determining the increase in the performance, especially when the supply is highly uncertain.

Remark 5. We solved the problem using Mosek (Andersen et al., 2003) implemented with CVX in MatLab R2021a. The most computationally demanding problem (Comirnaty) involves 667 variables and 809 constraints. It was solved in approximately 24 seconds on a 2.9 Gigahertz 6-Core Intel Core i9.

4.4. Sensitivity with respect to the uncertainty

In the previous analysis, we considered the level of uncertainty estimated from Italian data (Vaccini, 2021). Here, we extend our analysis by considering the effect of different levels of uncertainty on the performance of our optimization framework. We use as a baseline the level of uncertainty $\sigma(t)$ estimated from the Italian data, and, for the two sliding-horizon policies described in the above (weekly and monthly), we consider different scenarios in which the level of uncertainty is increased by two, four or ten times, or decreased by the same amounts—for Vaxzevria, we omit to report the results with tenfold uncertainty because, due to the high baseline uncertainty, the problem is often infeasible. Table 2 reports the variation of the average utility function J and total number of doses V administered with respect to the baseline scenario.

Predictably, increasing the uncertainty negatively affects the performance, while a decrease has a positive impact. Interestingly, the effect of such an impact is highly nonlinear and is affected by the characteristics of the vaccine and the re-evaluation policy implemented. In particular, it seems that our optimization framework is able to deal with increased levels of uncertainty without big performance losses up to a critical threshold. Once this crit-

Table 2
Effect of the level of uncertainty on the performance of the optimization framework.

	Comirnaty				Vaxzevria			
	weekly		monthly		weekly		monthly	
	ΔJ	ΔV	ΔJ	ΔV	ΔJ	ΔV	ΔJ	ΔV
10σ	-10.32%	-21.04%	-8.05%	-8.26%	-	-	-	-
4σ	-3.17%	-7.37%	-2.89%	-0.97%	-1.66%	-20.27%	-1.70%	-21.90%
2σ	-0.95%	-2.50%	-1.03%	-0.16%	-0.44%	-5.89%	-0.55%	-5.95%
0.5σ	0.53%	1.28%	0.64%	2.10%	0.26%	2.12%	0.31%	1.63%
0.25σ	0.82%	1.87%	0.96%	3.21%	0.40%	3.25%	0.45%	1.27%
0.1σ	0.99%	2.19%	1.14%	3.79%	0.49%	4.01%	0.54%	1.25%

ical threshold is reached, the performance quickly drops (as can be observed with a tenfold increase in the uncertainty for Comirnaty and with a fourfold increase for Vaxzevria, which has already a higher baseline level of uncertainty). Besides depending on the characteristics of the vaccine, the critical point seems also to be influenced by the re-evaluation policy. Specifically, policies with less frequent re-evaluations seem to be more resilient than policies with frequent re-evaluations. These observations should be carefully considered in view of our findings in Section 4.3. In fact, in the previous section, we showed that frequent re-evaluations of the optimal policy may be beneficial in the presence of higher levels of uncertainty. However, it seems that above a critical threshold for the uncertainty—fortunately, several times larger than the one observed in real data—the optimization framework fails in dealing with the uncertainty. A possible explanation could be that, at each re-evaluation, the uncertainty model for the expected supply $\hat{S}(t)$ is re-computed. In the presence of an extremely high uncertainty, this may lead to overfitting the data, and the advantage of having frequent re-evaluation is therefore lost, thus decreasing the performance. To sum up, our findings suggest that frequent re-evaluation of the vaccination policy may be a double-edged sword when the uncertainty level is very high and that the accurate estimation of the uncertainty model (based on historical data and on information given by the suppliers) is key toward optimally design the sliding horizon policy. However, our estimation of the uncertainty model for the COVID-19 vaccination campaign in Italy suggests that real-world scenarios have levels of uncertainty that are below such a critical threshold, and frequent re-evaluations are thus often preferable.

4.5. Sensitivity with respect to the risk parameter

One of the key parameters in our framework is the risk parameter ε , i.e., the probability of violating the constraints in the stochastic optimization framework. Practically, setting the value of ε means to determine the probability with which we allow the computed optimal solution to be unfeasible due to the stochasticity of the supply. Thus, large values of ε correspond to allowing higher risks of unfeasible plans that could result in problematic online steerings of the vaccination rollout. On the other hand, smaller values of ε would provide more conservative solutions, which are often safer but may yield unsatisfactory performance. Intuitively, policy makers should set the value of ε to trade off between potentially high performance and the risk of implementing problematic steerings. In the simulations above, we proposed optimal solutions using $\varepsilon = 0.05$. Here, we explore the effect of different values of ε , spanning from very safe implementations with $\varepsilon = 0.005$, to extremely risky scenarios, in which we set $\varepsilon = 0.4$.

The results are reported in Table 3, where we compare the performances with the one obtained in the safest scenario, ($\varepsilon = 0.005$). Predictably, we observe decreasing performance as long as we approach safer scenarios. Interestingly, while increasing

Table 3
Effect of the risk parameter on the performance of the optimization framework.

	Comirnaty				Vaxzevria			
	weekly		monthly		weekly		monthly	
	ΔJ	ΔV	ΔJ	ΔV	ΔJ	ΔV	ΔJ	ΔV
$\varepsilon = 0.4$	1.41%	3.55%	1.77%	4.04%	0.72%	7.41%	0.85%	5.84%
$\varepsilon = 0.3$	1.23%	3.18%	1.58%	3.39%	0.63%	6.47%	0.75%	5.89%
$\varepsilon = 0.2$	1.01%	2.75%	1.33%	2.45%	0.52%	5.39%	0.65%	6.22%
$\varepsilon = 0.1$	0.71%	2.05%	1.00%	1.21%	0.36%	3.89%	0.49%	5.84%
$\varepsilon = 0.05$	0.49%	1.46%	0.71%	0.46%	0.25%	2.77%	0.34%	4.53%
$\varepsilon = 0.01$	0.14%	0.39%	0.10%	-0.19%	0.07%	0.77%	0.09%	1.34%

ε toward more risky strategies has a strong initial impact on the performance, such an effect seems to saturate, becoming smaller and smaller as we move to more risky strategies. In fact, while increasing ε from 0.005 to 0.1 increases the number of doses administered for Comirnaty with a weekly evaluation by 2.05%, a further increase of the parameter by 0.1 would increase the doses by just a further 0.7%. This suggests that policy makers can determine an optimal trade-off value for such a parameter, corresponding to the point at which the increase in the performance starts becoming not sufficiently large to justify the corresponding increases in the risk. Another interesting observation can be found when comparing the solutions with different sliding-horizon policies. In fact, the optimal solutions with monthly re-evaluation seem to be slightly more affected by changes in ε than those revised on a weekly basis. Hence, opting for vaccination strategies that are more risky could be more beneficial when policy makers cannot re-evaluate the strategy with high frequency. However, in the presence of high levels of uncertainty (i.e., with Vaxzevria) and with less frequent re-evaluations, the increase in the performance associated with the implementation of risky strategies seems to quickly saturate, suggesting that conservative scenarios may be preferred when the uncertainty in the supply is large. From the policy makers, point of view, this could constitute an interesting approach to evaluate the trade-off between the maximization of the vaccination and the risk of an infeasible vaccination plan.

5. Conclusion

We proposed a novel optimization framework for dealing with the optimal design of a two-dose vaccination rollout in the presence of uncertain vaccine supply. We cast the optimization problem as a numerically-solvable SOCP, with an objective function that takes into account the vaccine efficacy and the partial immunity gained after a single dose. The framework incorporates realistic constraints on the administration timing of the vaccine—minimum and maximum period that should elapse between the doses, limited capacity of the healthcare system, and tunable parameters that model the uncertainty in the supply.

We demonstrated the proposed optimization framework on a case study calibrated on the COVID-19 vaccination campaign in Italy. The results illustrated the applicability of the proposed framework to design vaccination rollout, in real-world scenarios, thanks to its formulation as a SOCP. In particular, we showed that a sliding-horizon implementation could improve the performance, even in the presence of high levels of uncertainty. This case study also showed how the framework might be used by public health authorities to accurately design a vaccination rollout, by setting the re-evaluation frequency of the rollout and trading off performances and risk due to uncertainties in the supplies.

The flexibility of our framework paves the way for several extensions. First, further features may be incorporated (e.g., additional booster doses, or the combination of different types of vaccine) through additional state variables. Second, we considered a

uniform population, where the vaccination of each individual provides the same contribution to the objective function. Partitioning the population in classes, depending on the individuals' risk—and thus on the benefit of vaccinating them—may be a suitable strategy to use our framework for devising a vaccination rollout not only in terms of scheduling first and second doses, but also in terms of their prioritization. All these extensions could be implemented at the cost of slightly increasing the computational complexity, without changing the very fabric of the optimization framework. Finally, we modeled uncertainty by means of Gaussian noise in the cumulative supply, which captures delays (or earliness) of committed weekly supplies. Further sources of uncertainty, such as the cancellation of supply may be modeled by including additional features to the uncertainty model, such as a Bernoulli process for missing deliveries. These extensions, however, would call for a different formulation of the stochastic optimization problem and are envisaged for our future research.

Acknowledgments

The work by C.C., F.P., and A.R. is partially supported by the Italian Ministry of Foreign Affairs and International Cooperation ("Mac2Mic").

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.ejor.2022.05.009](https://doi.org/10.1016/j.ejor.2022.05.009)

References

- Andersen, E. D., Roos, C., & Terlaky, T. (2003). On implementing a primal-dual interior-point method for conic quadratic optimization. *Mathematical Programming*, 95, 249–277. <https://doi.org/10.1007/s10107-002-0349-3>.
- Arifoğlu, K., & Tang, C. S. (2022). A two-sided incentive program for coordinating the influenza vaccine supply chain. *Manufacturing Service Operations Management*, 24, 235–255. <https://doi.org/10.1287/msom.2020.0938>.
- Ben-Tal, A., Golany, B., Nemirovski, A., & Vial, J. P. (2005). Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing Service Operations Management*, 7, 248–271. <https://doi.org/10.1287/msom.1050.0081>.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., & Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99, 351–376. <https://doi.org/10.1007/s10107-003-0454-y>.
- Bertsimas, D., Dugalakis, V., Jr, Jacquillat, A., Li, M. L., & Previero, A. (2022). Where to locate COVID-19 mass vaccination facilities? *Naval Research Logistics*, 69, 179–200. <https://doi.org/10.1002/nav.22007>.
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization* (1st). Cambridge University Press.
- Calafiore, G. C. (2009). An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM Journal on Control and Optimization*, 48, 2254–2274. <https://doi.org/10.1137/080723776>.
- Calafiore, G. C., & Ghaoui, L. E. (2006). On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, 130, 1–22. <https://doi.org/10.1007/s10957-006-9084-x>.
- Calafiore, G. C., Novara, C., & Possieri, C. (2020). A time-varying SIRD model for the COVID-19 contagion in Italy. *Annual Reviews in Control*, 50, 361–372. <https://doi.org/10.1016/j.arcontrol.2020.10.005>.
- Chang, J. T., & Kaplan, E. H. (2021). Modeling local corona virus outbreaks. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.07.049>.
- Chatterjee, S., & Simonoff, J. S. (2013). *Handbook of regression analysis, volume 5*. John Wiley & Sons. <https://doi.org/10.1002/9781118532843>.
- Choi, T. M. (2021). Fighting against COVID-19: What operations research can help and the sense-and-respond framework. *Annals of Operations Research*. <https://doi.org/10.1007/s10479-021-03973-w>.
- Cohen, I., Golany, B., & Shtub, A. (2016). The stochastic time-cost tradeoff problem: A robust optimization approach. *Networks*, 49, 175–188. <https://doi.org/10.1002/net.20153>.
- Dabadghao, S., Marandi, A., & Roy, A. Optimal Interventions in Robust Optimization with Time-Dependent Uncertainty (December 8, 2021). Available at SSRN: <https://ssrn.com/abstract=3624421> or <http://dx.doi.org/10.2139/ssrn.3624421>.
- De Boeck, K., Decouttere, C., Jónasson, J. O., & Vandaele, N. (2021). Vaccine supply chains in resource-limited settings: Mitigating the impact of rainy season disruptions. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.10.040>.
- De Ruiter, F. J. C. T., Brekelmans, R. C. M., & den Hertog, D. (2016). The impact of the existence of multiple adjustable robust solutions. *Mathematical Programming*, 160, 531–545. <https://doi.org/10.1007/s10107-016-0978-6>.

- Duijzer, L. E., van Jaarsveld, W., & Dekker, R. (2018). Literature review: The vaccine supply chain. *European Journal of Operational Research*, 268, 174–192. <https://doi.org/10.1016/j.ejor.2018.01.015>.
- Eryarsoy, E., Shahmanzari, M., & Tanrisever, F. (2022). Models for government intervention during a pandemic. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.12.036>.
- Fischetti, M., Fischetti, M., & Stoustrup, J. (2021). Safe distancing in the time of COVID-19. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.07.010>.
- Giordano, G., et al., (2021). Modeling vaccination rollouts, SARS-CoV-2 variants and the requirement for non-pharmaceutical interventions in Italy. *Nature Medicine*. <https://doi.org/10.1038/s41591-021-01334-5>.
- Gorissen, B. L., Yanıkoğlu, I., & den Hertog, D. (2015). A practical guide to robust optimization. *Omega*, 53, 124–137. <https://doi.org/10.1016/j.omega.2014.12.006>.
- Hosseini-Motlagh, S. M., Samani, M. R. G., & Homaei, S. (2021). Design of control strategies to help prevent the spread of COVID-19 pandemic. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.11.016>.
- Lin, Q., Zhao, Q., & Lev, B. (2022). Influenza vaccine supply chain coordination under uncertain supply and demand. *European Journal of Operational Research*, 297, 930–948. <https://doi.org/10.1016/j.ejor.2021.05.025>.
- Lobo, M. S., Vandenbergh, L., Boyd, S., & Lebret, H. (1998). Applications of second-order cone programming. *Linear Algebra and Its Applications*, 284, 193–228. [https://doi.org/10.1016/S0024-3795\(98\)10032-0](https://doi.org/10.1016/S0024-3795(98)10032-0).
- Mak, H.-Y., Dai, T., & Tang, C. S. Managing Two-Dose COVID-19 Vaccine Rollouts with Limited Supply: Operations Strategies for Distributing Time-Sensitive Resources (March 31, 2022). Production and Operations Management, Forthcoming, Available at SSRN: <https://ssrn.com/abstract=3790836> or <http://dx.doi.org/10.2139/ssrn.3790836>.
- Moghadas, S. M., et al., (2021). Evaluation of COVID-19 vaccination strategies with a delayed second dose. *PLOS Biology*, 19, 1–13. <https://doi.org/10.1371/journal.pbio.3001211>.
- Nowzari, C., Preciado, V. M., & Pappas, G. J. (2016). Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems Magazine*, 36, 26–46. <https://doi.org/10.1109/MCS.2015.2495000>.
- Parino, F., Zino, L., Calafiore, G. C., & Rizzo, A. (2021a). A model predictive control approach to optimally devise a two-dose vaccination rollout: A case study on COVID-19 in Italy. *International Journal of Robust and Nonlinear Control*. <https://doi.org/10.1002/rnc.5728>.
- Parino, F., Zino, L., Porfiri, M., & Rizzo, A. (2021b). Modelling and predicting the effect of social distancing and travel restrictions on COVID-19 spreading. *Journal of the Royal Society Interface*, 18, 20200875. <https://doi.org/10.1098/rsif.2020.0875>.
- Polack, F. P., et al., (2020). Safety and efficacy of the BNT162b2 mRNA COVID-19 vaccine. *The New England Journal of Medicine*, 383, 2603–2615. <https://doi.org/10.1056/NEJMoa2034577>.
- Prékopa, A. (1995). *Stochastic programming* (1st). Springer Netherlands.
- Ross, S. M. (2019). *Introduction to probability models* (12th). Academic Press.
- Shumsky, R. A., Smith, J., Hoen, A., & Gilbert, M. Allocating COVID-19 Vaccines: Save One for the Second Dose? (August 10, 2021). Tuck School of Business Working Paper No. 3816200, Available at SSRN: <https://ssrn.com/abstract=3816200> or <http://dx.doi.org/10.2139/ssrn.3816200>
- Sinha, P., Kumar, S., & Chandra, C. (2021). Strategies for ensuring required service level for COVID-19 herd immunity in indian vaccine supply chain. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.03.030>.
- Taylor, J. W., & Taylor, K. S. (2021). Combining probabilistic forecasts of COVID-19 mortality in the united states. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.06.044>.
- Truszkowska, A., et al., (2021). High-resolution agent-based modeling of COVID-19 spreading in a small town. *Advanced Theory and Simulations*, 4, 2170005. <https://doi.org/10.1002/adts.202000277>.
- Tsiligianni, C., Tsiligiannis, A., & Tsiliyannis, C. (2022). A stochastic inventory model of COVID-19 and robust, real-time identification of carriers at large and infection rate via asymptotic laws. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2021.12.037>.
- Tuite, A. R., Fisman, D. N., Zhu, L., & Salomon, J. A. (2021). Alternative dose allocation strategies to increase benefits from constrained COVID-19 vaccine supply. *Annals of Internal Medicine*, 174, 570–572. <https://doi.org/10.7326/M20-8137>.
- covid19-opendata-vaccini. (2021). Commissario straordinario per l'emergenza COVID-19. <https://github.com/italia/covid19-opendata-vaccini>.
- Voysey, M., et al., (2021). Single-dose administration and the influence of the timing of the booster dose on immunogenicity and efficacy of chadox1 nCoV-19 (AZD1222) vaccine: a pooled analysis of four randomised trials. *The Lancet*, 397, 881–891. [https://doi.org/10.1016/S0140-6736\(21\)00432-3](https://doi.org/10.1016/S0140-6736(21)00432-3).
- Wang, Z., Bauch, C. T., Bhattacharyya, S., d'Onofrio, A., Manfredi, P., Perc, M., ... Zhao, D. (2016). Statistical physics of vaccination. *Physics Reports*, 664, 1–113. <https://doi.org/10.1016/j.physrep.2016.10.006>.
- Zhang, C., Li, Y., Cao, J., & Wen, X. (2022). On the mass COVID-19 vaccination scheduling problem. *Computers Operations Research*, 141, 105704. <https://doi.org/10.1016/j.cor.2022.105704>.
- Zino, L., & Cao, M. (2021). Analysis, prediction, and control of epidemics: A survey from scalar to dynamic network models. *IEEE Circuits and Systems Magazine*, 21, 4–23. <https://doi.org/10.1109/MCAS.2021.3118100>.