IISE Transactions

Taylor & Francis
Taylor & Francis Group

IISE Transactions

ISSN: 2472-5854 (Print) 2472-5862 (Online) Journal homepage: www.tandfonline.com/journals/uiie21

Release immediately or sequentially? Strategies for allocating scarce therapeutic resources during disease outbreaks

Bismark Singh & Steffen Rebennack

To cite this article: Bismark Singh & Steffen Rebennack (25 Jul 2025): Release immediately or sequentially? Strategies for allocating scarce therapeutic resources during disease outbreaks, IISE Transactions, DOI: 10.1080/24725854.2025.2525918

To link to this article: https://doi.org/10.1080/24725854.2025.2525918

IISE TRANSACTIONS
https://doi.org/10.1080/24725854.2025.2525918

Release immediately or sequentially? Strategies for allocating scarce therapeutic resources during disease outbreaks

Bismark Singh[a] and Steffen Rebennack[b]

[a]School of Mathematical Sciences, University of Southampton, Southampton, UK; [b]Institute for Operations Research, Karlsruhe Institute of Technology, Karlsruhe, Germany

ABSTRACT

During infectious disease outbreaks, public health authorities face the challenge of fairly distributing limited medical therapeutics, such as antiviral drugs and respirators, to save lives. Such resources are not always widely available at the start of the outbreak when the progression of the outbreak is uncertain. Then, resources may be allocated from a centralized inventory—such as a federal stockpile—to regional authorities for further distribution on a periodic basis. Using mathematical optimization models for spatiotemporal resource allocations, we address a fundamental question of whether such stockpiles should be fully released at the outset or gradually distributed throughout the threat. Immediately releasing the therapeutics risks its consumption by non-symptomatic populations who are concerned about the disease; this is particularly so in the absence of a diagnostic test that distinguishes such individuals from those that are genuinely sick. We find that if the benefit a resource provides to an individual, whether genuinely sick or not, decreases as the outbreak progresses, then an immediate release is optimal. However, if some population groups are expected to benefit more than others, then sequential releases that conserve resources and match peak demand are advisable. Sequential release policies can pose challenging ethical and political dilemmas for policymakers under public pressure to provide immediate relief. To this end, using simulated scenarios sampled from six historic pandemics including the COVID-19 pandemic, we provide computational experiments for the state of Texas, with over 25 million people, that empirically demonstrate a significant number of lives can be saved by following optimal release policies.

ARTICLE HISTORY
Received 4 May 2024
Accepted 11 June 2025

KEYWORDS
Resource allocation; stochastic programs with recourse; healthcare policy; epidemiology; pandemics

1. Background

Prophylactic vaccines are often the most effective countermeasure against the spread of an infectious disease; however, their timeline for research, development and distribution against a novel threat can be quite long (Felter, 2021). In contrast to prophylactic countermeasures that work by preventing a disease, another line of defense is provided by therapeutic countermeasures that work by treating existing infections. Examples of such resources include antivirals, monoclonal antibodies, and mechanical ventilators—all of which help save lives of populations infected by a focal pathogen. Although typically less effective than prophylactic vaccines, such therapeutic resources are often stockpiled by federal and state agencies for limited release during infectious disease threats. For example, US authorities released 11 million courses of antivirals, such as oseltamivir and zanamivir, from the Strategic National Stockpile to states during the 2009 H1N1 influenza pandemic (Dimitrov et al., 2009), and several other therapeutic drugs were authorized to reduce the severity of COVID-19 infections early in the pandemic (Nature, 2021). Further, some states (e.g., Texas) reserved a small cache of vaccines for discretionary allocation to under-served populations in the 2009 H1N1 pandemic (Huang et al., 2017).

Although exact numbers, and even locations, of such therapeutic resources are classified, their available quantities are far fewer than potential demand (Facher, 2020). Thus, policymakers face the challenge of distributing them *fairly* and *effectively* from the central stockpile to a large and diverse population. These two goals are well-known in the healthcare (and, more generally, in the public policy) literature to be conflicting, see, e.g., (Emanuel et al., 2020; Huang et al., 2017; Lee et al., 2021; Persad et al., 2009). From an ethical standpoint, a fair allocation should provide a resource to any resource-seeking individual (or, any individual identified eligible for a resource) (Emanuel et al., 2020). From an efficiency standpoint, a rationing is required as resources consumed by different individuals at different times provide a different societal benefit (Emanuel &

CONTACT Bismark Singh b.singh@southampton.ac.uk

Supplemental data for this article can be accessed online at https://doi.org/10.1080/24725854.2025.2525918.

Wertheimer, 2006). Any decision of denying a resource-seeking individual a resource that is available (even if it is scarce) poses challenges for a policymaker via allegations of bias and discrimination. Indeed, during the COVID-19 pandemic several such racial and demographic disparities became clear as resource allocation decisions were made from a healthcare professional's judgment without clear guidelines (Laventhal et al., 2020).

During this early phase of novel disease outbreaks, it is especially difficult for a healthcare professional to discern infected and uninfected individuals. This is particularly because pathogen-detection testing is either not readily available, not reliable, or too expensive. Thus, a publicly available resource risks immediate consumption by the so-called "worried-well" population; i.e., individuals who are not infected by the pathogen but are severely concerned by its spread. Ethically, in the absence of a clear distinguishing mechanism a policymaker cannot regulate or restrict consumption by the worried-well population; see, e.g., (Goodwin et al., 2009) and (Walensky & Paltiel, 2006) for concerns of unmonitored usage of antiviral drugs for H1N1 and detection tests for HIV, respectively. It is with this background that the stored therapeutic resources acquire a particularly critical nature—awaiting their release—compounding policymaking challenges of rationing for consumption by the infected individuals who would provide the largest societal benefit. Such priority-based resource conservation policies were both advocated and found application during the early phases of the COVID-19 pandemic; see, e.g., (Laventhal et al., 2020). Further, in the early stages of the COVID-19 pandemic, individuals with symptoms were prioritized for testing by the US Centers for Disease Control and Prevention (Ward et al., 2020).

Resource allocation decisions are also impeded in the early phase of the disease by the uncertainty surrounding its spatiotemporal spread (Rahmandad & Sterman, 2008); i.e., the population affected by the disease in a given geography. Once resources are allocated to a specific jurisdiction or provider—such as a clinic or pharmacy—redeployment elsewhere is often costly or even impossible (Singh et al., 2015). A lack of situational awareness, coupled with the challenges of redeployment, suggests that policymakers should consider *sequentially* adaptive distributions, in which portions of the stockpile are held in reserve and released optimally as conditions change. However, due to reasons we highlight above, such policies of holding on to available resources can be legislatively difficult to enforce particularly when significant fear or worry of the disease is prevalent in addition to the underlying pathogen. As such, policymakers are often under immense pressure to *immediately* provide relief (Squazzoni et al., 2020).

With this background, we develop a mathematical model addressing such concerns for the spatiotemporal allocation and rationing of scarce therapeutic resources at the onset of disease outbreaks (henceforth, pandemics). Resource allocation strategies during pandemic, particularly for medical drugs, have been extensively studied using epidemiological susceptible-exposed-infected-recovered (SEIR) models (Arinaminpathy & McLean, 2008; Kleczkowski et al., 2019; Longini, 2004). To derive optimal control policies, researchers often compare a relatively small set of candidate strategies via simulation (Matrajt & Longini, 2010; Rahmandad & Sterman, 2008). For example, Longini et al., find that administering drugs prophylactically to 80% of the exposed population for eight weeks is almost as effective as vaccinating 80% of the entire population (Longini et al., 2004). Another study compares the merits of aggressive distribution strategies with higher infection peaks against a conservative plan with a longer pandemic duration (Arinaminpathy & McLean, 2008). To evaluate larger policy spaces—beyond the obvious candidates—researchers have turned to mathematical optimization methods (Dimitrov et al., 2009; Long et al., 2018; Singh et al., 2015). Optimization-based approaches can search large combinatoric spaces to determine distribution policies that are expected to save the most lives, across a wide range of epidemiological scenarios sampled via simulation of SEIR models (Duque et al., 2020).

In this work, we use a two-stage stochastic programming approach to determine resource allocation policies. The distribution decisions of policymakers are governed by maximizing an expected societal *benefit*—a notion well-studied in healthcare policy making (Emanuel et al., 2020; Matrajt & Longini, 2010; Persad et al., 2009); e.g., the quality-adjusted life years or number of lives saved. In Section 3, we provide a precise definition of our benefit. In addition to the uncertain benefit, the demand for resources is also uncertain. We simulate scenarios for demand and benefit using a compartmental SEIR-styled model (similar to the studies cited above) to construct discrete representations of the uncertainty. Our scenarios are seeded via a wide variety of epidemiological parameters (see, Section 3 and Appendix B) from Texas. During early stages of a pandemic, estimates for such epidemiological parameters vary widely across different scientific studies; our optimization model inputs these different simulations to output a *single* release schedule that is robust across all projections. In this sense, we seek to combine the advantages of both optimization and simulation-based mathematical approaches. Although the model we propose belongs to the family of general mathematical resource allocation problems as originally studied by Dantzig and Wolfe (Dantzig & Wolfe, 1961), there are at least three distinguishing features.

F.1 First, we assume that allocated resources cannot be pulled back or redeployed. Such overhauling and/or reallocation decisions are often virtually impossible due to prohibitive costs (Feldman & Sakhartov, 2022) or legal challenges (Hodge and O'Connell, 2006).

F.2 Second, releasing resources at certain times can have negative consequences; e.g., due to excessive consumption by the worried-well populations as was observed during the 2009 H1N1 pandemic (Ungchusak et al., 2012). The effects of this behavior may be compounded if the resources are available without prescription (Liang & Mackey, 2012). In addition to such worried-well behavior, this

feature finds application in a few other classes of resource allocation problems with negative externalities (Nadiminti et al., 2002).

F.3 Third, no person can be denied a resource at any time if a resource is available; this feature is designed to ensure fairness in the sense of individuals having equal rights to life-saving medical countermeasures (Lee et al., 2021; National Academies of Sciences et al., 2020). This feature disallows pharmacies from selectively dispensing critical resources during scarcity, but also provides no control over the worried-well populations seeking resources. Further, this feature also prevents providers from hoarding critical resources during times of scarcity (World Health Organization, 2016).

These three features result in several interesting mathematical properties of our model that further lead to important policy decisions; we present these in Section 2. For an introduction to the mathematics of resource allocation problems, we refer to Erlebacher (2000) and Singh (2020). For an introduction to fairness in critical resource allocation problems, we refer to Emanuel et al. (2020) and Persad et al. (2009); see, also the recent work on fairness measured via simulations (D'Amour et al., 2020). We summarize the main contributions and findings of this article as follows.

i. We propose a framework for optimizing an apriori release schedule for limited stockpiles of therapeutic resources to different regions at the very onset of a pandemic when its progression is uncertain (Section 2).
ii. We formally distinguish between sequential and immediate release policies available for a policymaker (Definition 1). We provide sufficient conditions guaranteeing the optimality of either, or both, release policies as well as bounds on the suboptimality of the other policy (Theorems 2–5).
iii. We prove that in the absence of age-differential response to a resource, immediate release of scarce resources as soon as they are available provides the maximum overall benefit (Proposition 5). However, if the therapeutic response of receiving a resource differs by age, we demonstrate that sequential release rather than immediate release is optimal, and can save significantly more lives, across a wide range of epidemiological scenarios grounded in data from six historical pandemics (Section 5).
iv. Finally, we study the lost opportunity costs that occur through worried-well uptake, and derive optimal allocation strategies under various levels of wastage (Section 4).

The structure of the rest of the article is as follows. In Section 2, we provide the mathematical optimization models that form the basis of our work. In Section 3, we provide formulae to compute the parameters of the optimization models, as well as an estimation procedure for these parameters. In Section 4, we analyze the effect of these parameters on the optimal solution; while, in Section 5 we use the estimates in a computational case study. We reserve some of the details for the Appendix. We conclude with limitations of our work and a summary in Section 6.

2. Mathematical formulations

2.1. Notation

Indices/ Sets	
$c \in C$	set of locations of infected individuals [counties]
$t \in T$	set of times; $T = \{1, 2, \ldots, \lvert T\rvert\}$ [months]
$\omega \in \Omega$	set of scenarios [-]
Parameters	
$A_t \in \mathbb{N}_+$	number of new resources available for release at time t [doses]
$S^{\omega}_{c,t} \in \mathbb{N}_+$	demand for resources in county c under scenario ω at time t; [doses]
$B^{\omega}_{c,t} \in \mathbb{R}$	average per-person benefit on receiving one dose in county c under scenario ω at time t; [benefit/dose]
$M_t \in \mathbb{R}_+$	Big-$\mathcal{M}$ value given by upper bound of $q^{\omega}_{c,t}$ [doses]
$P^{\omega} \in \mathbb{R}_+$	probability of scenario ω; $\sum_{\omega\in\Omega} P^{\omega} = 1$ [-]
Decision Variables	
$r_{c,t}$	(continuous, first stage) number of resources released in county c at time t [doses]
$f^{\omega}_{c,t}$	(continuous, second stage) fraction of satisfied demand in county c under scenario ω at time t; $0 \le f^{\omega}_{c,t} \le 1$ [-]
$q^{\omega}_{c,t}$	(continuous, second stage) surplus resources in county c under scenario ω at time t [doses]
$x^{\omega}_{c,t}$	(binary, second stage) =1 if entire demand for resources in county c under scenario ω at time t is satisfied; else 0 [-]

2.2. Optimization model

Our stochastic mixed-integer linear optimization model for the resource allocation problem reads as follows:

$$z^* = \max \sum_{c\in C,\, t\in T} \mathbb{E}_{\omega}\left[B^{\omega}_{c,t} S^{\omega}_{c,t} f^{\omega}_{c,t}\right] \tag{1a}$$

$$s.t. q^{\omega}_{c,t} = q^{\omega}_{c,t-1} - f^{\omega}_{c,t} S^{\omega}_{c,t} + r_{c,t} \quad \forall t \in T, \omega \in \Omega, c \in C \tag{1b}$$

$$\sum_{c\in C,\, t'\in T: t'\le t} r_{c,t'} \le \sum_{t'\in T: t'\le t} A_{t'} \forall t \in T \tag{1c}$$

$$x^{\omega}_{c,t} \le f^{\omega}_{c,t} \quad \forall t \in T, \omega \in \Omega, c \in C \tag{1d}$$

$$q^{\omega}_{c,t} \le M_t \; x^{\omega}_{c,t} \quad \forall t \in T, \omega \in \Omega, c \in C \tag{1e}$$

$$x^{\omega}_{c,t} \in \{0, 1\} \quad \forall t \in T, \omega \in \Omega, c \in C \tag{1f}$$

$$0 \le f^{\omega}_{c,t} \le 1 \quad \forall t \in T, \omega \in \Omega, c \in C \tag{1g}$$

$$q^{\omega}_{c,t}, r_{c,t} \ge 0 \quad \forall t \in T, \omega \in \Omega, c \in C. \tag{1h}$$

$$\begin{aligned} &\textit{Boundary conditions}: \\ &q^{\omega}_{c,0} = 0 \qquad\qquad \forall \omega \in \Omega, c \in C \end{aligned} \tag{1i}$$

$$f^{\omega}_{c,t} = 1, \; x^{\omega}_{c,t} = 1 \quad \forall (t, \omega, c) : S^{\omega}_{c,t} = 0. \tag{1j}$$

Model (1) is a two-stage stochastic program with recourse having multiple time periods (Shapiro et al., 2009). In the

first-stage, resources are allocated across all counties $c \in C$ and all times $t \in T$; *i.e.*, the entire release schedule, $r_{c,t}$, is decided. Then, the uncertainty is realized through scenario ω revealing the precise demand, $S^{\omega}_{c,t}$, and benefit, $B^{\omega}_{c,t}$. If the available resources exceed the demand, the surplus resources $q^{\omega}_{c,t}$ "roll-forward" to the next time; the binary variable $x^{\omega}_{c,t}$ indicates this situation. On the other hand, if the available resources are short of the demand they are completely consumed. The variable $f^{\omega}_{c,t}$ denotes the fraction of the demand that is satisfied; then, the quantities $S^{\omega}_{c,t}f^{\omega}_{c,t}$ and $B^{\omega}_{c,t}S^{\omega}_{c,t}f^{\omega}_{c,t}$ define the satisfied demand and the correspondingly realized value of satisfying the demand in a (c,t,ω) triplet, respectively.

With this background, the objective function (1a) maximizes the sum of the expected benefit from resource pickups across all counties and times. In Section 3, we provide a detailed discussion of such benefits and quantify $B^{\omega}_{c,t}$ in terms of the lives saved due to resource pickups; then, the objective function computes the expected number of lives saved relative to the situation if there were no resources available. Importantly, $B^{\omega}_{c,t}$ can be negative due to feature F.2 of Section 1. Constraint (1b) is an inventory balancing constraint and relates the surplus resources at time t with those at time $t-1$; the quantity $q^{\omega}_{c,t-1}+r_{c,t}$ is the total number of resources available for the triplet (c,t,ω). We enforce the boundary condition $q^{\omega}_{c,0}=0, \forall c \in C,\ \omega \in \Omega$ via equation (1i). Constraint (1c) restricts the number of resources that can be released at any time by those that are available; we consider resources are available for distribution on a periodic basis. Linear constraints (1d)-(1g) enforce the nonlinear complementary restriction:

$$q^{\omega}_{c,t}(1-f^{\omega}_{c,t})=0 \qquad \forall t \in T, \omega \in \Omega, c \in C \tag{2}$$

i.e., resources are never in surplus if the current demand is not completely satisfied. This important property of the model is due to feature F.3 of Section 1; we enforce resources cannot be held back if available and someone is in need. Constraints (1e) and (1f) ensure that $x^{\omega}_{c,t}=1$ if $q^{\omega}_{c,t}>0$; then, constraints (1d) and (1g) yield $f^{\omega}_{c,t}=1$.

In model (1), we do not constrain the variables $q^{\omega}_{c,t}$ and $r_{c,t}$—that denote doses of resources—to be integers. As we show later in this article, even in the absence of these integer restrictions, model (1) requires significant computational effort for its solution. In an effort toward tractability, the following proposition provides a sufficiently large value for the Big-$\mathcal{M}$ terms M_t, required in constraints (1e) to linearize the nonlinear complementary condition (2).

Proposition 1. *A sufficiently large value for M_t for model (1) is given by $\sum_{t' \in T: t' \le t} A_{t'}$.*

Proof. Adding constraints (1b) sequentially, we obtain

$$q^{\omega}_{c,t}=\sum_{t' \in T: t' \le t}(r_{c,t'}-f^{\omega}_{c,t'}S^{\omega}_{c,t'}) \le \sum_{t' \in T: t' \le t} r_{c,t'} \le \sum_{t' \in T: t' \le t} A_{t'}, \quad \forall t \in T, \omega \in \Omega, c \in C.$$

From constraint (1e) the result follows. ■

A policymaker is mainly interested in the release schedule, $r_{c,t}$, of the resources. Indeed, the following theorem shows that the second-stage decisions are given by the first-stage decisions; this is a very special structure. As such, the second-stage decision variables evaluate the release schedule in the stochastic environment, determining the objective function value associated with the deterministic decisions $r_{c,t}$.

Theorem 1.

i. *In an optimal solution for model (1), $f^{\omega}_{c,t}=\min\{1, \frac{q^{\omega}_{c,t-1}+r_{c,t}}{S^{\omega}_{c,t}}\}$, $\forall t \in T, \omega \in \Omega, c \in C$.*

ii. *Given values of $r_{c,t}, \forall c \in C, t \in T$, the values of the variables $q^{\omega}_{c,t}, f^{\omega}_{c,t}$ and $x^{\omega}_{c,t}$ are completely determined $\forall t \in T, \omega \in \Omega$, $c \in C$ as follows: $q^{\omega}_{c,t}=q^{\omega}_{c,t-1}-f^{\omega}_{c,t}S^{\omega}_{c,t}+r_{c,t}$; $x^{\omega}_{c,t}=1$ if $f^{\omega}_{c,t}=1$ else 0, where the $f^{\omega}_{c,t}$ values are given in (i).*

Proof.

i. If $S^{\omega}_{c,t}=0$, then the statement holds by definition. Hence, consider (c,t,ω) s.t. $S^{\omega}_{c,t}>0$. We distinguish two cases.
   a. For (c,t,ω) s.t. $x^{\omega}_{c,t}=1$ in an optimal solution, we have $f^{\omega}_{c,t}=1$ by constraint (1d). From constraint (1b) we obtain $\frac{q^{\omega}_{c,t-1}+r_{c,t}}{S^{\omega}_{c,t}}=1+\frac{q^{\omega}_{c,t}}{S^{\omega}_{c,t}} \ge 1$.
   b. For (c,t,ω) s.t. $x^{\omega}_{c,t}=0$ in an optimal solution, we have $q^{\omega}_{c,t}=0$ by constraint (1e); then from constraints (1b) and (1g) it follows that $f^{\omega}_{c,t}=\frac{q^{\omega}_{c,t-1}+r_{c,t}}{S^{\omega}_{c,t}} \le 1$.

ii. We prove by induction on t. Let the base case be $t=1$. Given $r_{c,1}$, $f^{\omega}_{c,1}$ is given by (i). If $f^{\omega}_{c,1}=1$, then $x^{\omega}_{c,1}=1$; similarly, if $f^{\omega}_{c,1}<1$, then $x^{\omega}_{c,1}=0$. By constraint (1b), $q^{\omega}_{c,1}$ is given by $r_{c,1}-f^{\omega}_{c,1}S^{\omega}_{c,1}$ and the base case holds. Next, as the inductive hypothesis, assume (ii) holds for time $t-1$. We show that it also holds for time t. Again, via (i) we determine $f^{\omega}_{c,t}$ since $r_{c,t}$ is known by the hypothesis and $q^{\omega}_{c,t-1}$ is known by the inductive hypothesis. If $f^{\omega}_{c,t}=1$, then $x^{\omega}_{c,t}=1$; similarly, if $f^{\omega}_{c,t}<1$, then $x^{\omega}_{c,t}=0$. By constraint (1b), $q^{\omega}_{c,t}$ is given by $q^{\omega}_{c,t-1}-f^{\omega}_{c,t}S^{\omega}_{c,t}+r_{c,t}$. ■

Corollary 1. *Model (1) has relatively complete recourse.*

Proof. It follows from Theorem 1 that a feasible second-stage solution exists for any $r_{c,t} \ge 0$. ■

From Theorem 1 it follows that a policymaker only needs to determine a release schedule for distributing resources. Next, we distinguish between two exclusive release schedules available at the policymaker's disposal.

2.3. Types of release schedules

In this section, we define two types of allocation policies that determine the release schedules. Further, we derive sufficient conditions for the optimality of either policy.

Definition 1. *Release Policies*

Immediate Release Policy (IRP): A set of values for $r_{c,t}$ that is feasible for model (1) and satisfies $\sum_{c \in C} r_{c,t}=A_t$ for all $t \in T$.

Sequential Release Policy (SRP): A set of values for $r_{c,t}$ that is feasible for model (1) with $\sum_{c\in C} r_{c,t} \neq A_t$ for at least one $t \in T$. ■

It follows from Definition 1 that an optimal release schedule, $r^*_{c,t}$, for model (1) is either an IRP or a SRP. In the following, we distinguish between an "optimal" policy—optimal here is with respect to model (1)—and, a "best" IRP or "best" SRP which is a policy with the largest value of the objective function of model (1) among all IRPs or SRPs, respectively, if they exist.

The IRP provides a policymaker an easy to implement policy—release all available resources within counties as soon as they are available (although the policymaker still needs to decide the distribution of the resources among the different counties $c \in C$). However, within the IRP, released resources are immediately consumed even if there is a lower benefit; thus, an optimal policy for model (1) might not always be an IRP. In contrast, the SRP allows a matching of the release schedule with the peak demand, and any optimal policy is a SRP when the best IRP is provably sub-optimal. Hence, counter-intuitively, it can be beneficial overall to hold on to resources even when there is a demand for them; we provide a simple bound to quantify the sub-optimality of the IRP in Theorem 4 and computational experiments in Section 5. We begin with the following theorem that provides a sufficient condition to guarantee optimality of a best IRP.

Theorem 2. Let $B^{\omega}_{c,t} \geq B^{\omega}_{c,t+1} \geq 0 \quad \forall t \in T, \omega \in \Omega, c \in C$, $t < |T|$. Then, there exists an IRP which is also an optimal policy for model (1).

Proof. Because of the decreasing per-person benefits $B^{\omega}_{c,t}$, we no longer need to enforce the special inventory constraints in equation (1b). (Otherwise, consider a policy not satisfying the inventory constraint. Then we can re-arrange the resources to satisfy the inventory constraints by using the resources from future times; because of the decreasing per-person benefit, the objective function does not decrease.) Thus, model (1) simplifies to

$$z^* = \max \sum_{c\in C,\, t\in T} \sum_{\omega\in\Omega} P^{\omega} B^{\omega}_{c,t} y^{\omega}_{c,t} \tag{3a}$$

$$\sum_{t'\leq t} y^{\omega}_{c,t'} \leq \sum_{t'\leq t} r_{c,t} \;\; \forall c \in C, t \in T, \omega \in \Omega \tag{3b}$$

$$\sum_{c\in C,\, t'\in T:t'\leq t} r_{c,t'} \leq \sum_{t'\in T:t'\leq t} A_{t'} \;\; \forall t \in T \tag{3c}$$

$$0 \leq y^{\omega}_{c,t} \leq S^{\omega}_{c,t} \;\; \forall t \in T, \omega \in \Omega, c \in C, \tag{3d}$$

where decision variable $y^{\omega}_{c,t}$ is the number of individuals who successfully receive resources in triplet (c, t, ω). We note that in constraint (3b) the decision variables $y^{\omega}_{c,t}$ are limited by the variables $r_{c,t}$ which are in turn limited by the parameter A_t in constraint (3c). Because of this special structure of equations (3b)-(3d), it is optimal to choose $r_{c,t}$ as large as possible as early as possible to allow the most flexibility for the $y^{\omega}_{c,t}$ variables in equation (3b). Thus, $\sum_{c\in C} r_{c,t} = A_t, \forall t \in T$ yields an optimal release schedule. ■

We note that Theorem 2 only provides a set of sufficient conditions to guarantee optimality of a best IRP, but these conditions are not necessary. Analogously, the next result provides a sufficient condition for the optimality of a best SRP.

Theorem 3. Let $0 \leq B^{\omega}_{c,t} \leq B^{\omega}_{c,t+1} \;\; \forall t \in T, \omega \in \Omega, c \in C, t < |T|$ and $\exists \bar{t} \in T : \sum_{t'\in T:t'\geq \bar{t}} S^{\omega}_{c,t'} > \sum_{t'\in T:t'\geq \bar{t}} A_{t'} \;\; \forall c \in C, \;\; \omega \in \Omega$ and $\sum_{t'\in T:t'<\bar{t}} A_{t'} > 0$. Then, there exists a SRP which is an optimal policy for model (1).

Proof. This proof is by construction. Assume there exists an IRP which is an optimal policy for model (1). (If no such IRP exists, then from Definition 1 an optimal policy is a SRP and there is nothing to show.) Given such an IRP, we construct a SRP with a no worse objective function value, impyling that this newly constructed SRP is also an optimal policy for model (1).

Let $y^{\omega,*}_{s,t}$ be the number of individuals who successfully receive a resource and $r^*_{c,t}$ be an associated solution for such a IRP. From the assumptions, there exists a time $\bar{t} \in T$ with $y^{\omega,*}_{c,\bar{t}} < S^{\omega}_{c,\bar{t}}$ for all $c \in C, \;\; \omega \in \Omega$ and $\sum_{t'\in T:t'<\bar{t}} A_{t'} > 0$. This implies that there exists some $\tau \in T$ with $\tau < \bar{t}$ and some $c \in C$ with $r^*_{c,\tau} > 0$. Now, construct the following policy for model (1) by copying all $r^*_{c,t}$ but reducing $r^*_{c,\tau}$ by $\ell = \min\{r^*_{c,t}, 1\} > 0$ and increasing $r^*_{c,\bar{t}}$ by ℓ; and, choose the $y^{\omega,*}_{s,t}$ accordingly. The technical assumption "$\exists \bar{t} \in T : \sum_{t'\in T:t'\geq \bar{t}} S^{\omega}_{c,t'} > \sum_{t'\in T:t'\geq \bar{t}} A_{t'} \;\; \forall c \in C, \;\; \omega \in \Omega$ and $\sum_{t'\in T:t'<\bar{t}} A_{t'} > 0$" ensures the feasibility of this policy for model (1) by ensuring there is additional capacity to distribute available resources beyond the IRP. Then, this policy is a SRP. The objective function value cannot decrease; further, it improves by at most $\sum_{c\in C,\, \omega\in\Omega} P^{\omega}(B^{\omega}_{c,\bar{t}} - B^{\omega}_{c,\tau})\ell \geq 0$. This shows that the newly constructed SRP is (also) an optimal policy. ■

The sufficient conditions in Theorem 2 and Theorem 3 relate to the uncertain parameters, $B^{\omega}_{c,t}$ and $S^{\omega}_{c,t}$, which can be verified before the allocation of resources. Thus, they allow a policymaker to assuredly make the entire set of release decisions without risking suboptimality. A natural follow-up question from this discussion is what happens when the sufficient conditions of Theorem 2 or Theorem 3 are not met; i.e., how suboptimal is the best IRP and the best SRP from an optimal policy. To this end, we first note that any feasible (and, hence, optimal) policy of model (1) is easily verifiable as an IRP or a SRP. However, if an optimal policy is an IRP, there might not even exist a best SRP since we may choose $\sum_{c\in C} r_{c,t}$ arbitrarily close to A_t for some $t \in T$. This is because the feasible region of the optimization problem defining the SRPs is no longer a closed set. For example, consider a simple instance with $|C| = |T| = |\Omega| = 1$ and, correspondingly, $B^{\omega}_{c,t} = S^{\omega}_{c,t} = A_t = 1$. The unique optimal solution is $r^*_{c,t} = f^{\omega,*}_{c,t} = 1$ with $z^* = 1$. This optimal policy is an IRP. Any SRP requires $r^*_{c,1} < 1$. However, there exists no finite optimal solution for model (1) with the

additional constraint $r^*_{c,1} < 1$. This mathematical subtlety needs to be kept in mind when comparing the objective function values of the IRP and the SRP. An easy workaround to avoid this issue is to enforce $r_{c,t}$ as integer-valued for all $c \in C$ and $t \in T$. However, doing so introduces additional integer variables for our model. With this background, the following theorem provides a simple bound on the difference of the objective function values of the IRP and SRP.

Theorem 4. Let z^*_{IRP} denote the objective function value of model (1) for the best IRP. Define

$$\underline{B} = \min_{c\in C,\, t\in T,\, \omega\in\Omega} B^{\omega}_{c,t} \quad \text{and} \quad \bar{B} = \max_{c\in C,\, t\in T,\, \omega\in\Omega} \{B^{\omega}_{c,t}, 0\}$$
$$\text{and} \qquad A = \sum_{t\in T} A_t.$$

Then

$$z^* - z^*_{\text{IRP}} \le (\bar{B} - \underline{B})A.$$

Proof. We distinguish two cases.

i. Consider $\underline{B} \le 0$. Then, a trivial upper bound for model (1) for any policy is given by $\bar{B}A$; i.e., $z^* \le \bar{B}A$. A lower bound on model (1) for any IRP is given by $\underline{B}A$; i.e., $z^*_{\text{IRP}} \ge \underline{B}A$. The result follows.
ii. Consider $\underline{B} > 0$. Then, $\bar{B} > 0$. Hence, for one unit of an allocated resource, the maximum difference in the objective function values between an optimal policy and any IRP for model (1) is $\bar{B} - \underline{B}$. Since at most A resources may be allocated, the result follows. ■

For a policymaker, Theorem 4 provides an easy to compute bound on the potential suboptimality of a best IRP from a best IRP. Theorem 4 also provides the insight that the objective function value of the two policies, IRP and SRP, are bound by the largest per-person benefit differences among all countries $c \in C$, times $t \in T$ and scenarios $\omega \in \Omega$. Large values of this bound serve as a warning to a policymaker for the potential loss which may result by not following the optimal policy. Although this bound is expected to be rather loose for most instances (i.e., the quantified loss is highly conservative), the following simple example shows that it may be tight.

Corollary 2. *The bound of Theorem 4 may hold with an equality.*

Proof. Consider an instance with $|C| = |\Omega| = 1, |T| = 2$ $B^{\omega}_{c,1} = -1$, $B^{\omega}_{c,2} = 1$, $S^{\omega}_{c,t} = 1$, and available resources $A_1 = 1$ and $A_2 = 0$. The (unique) best IRP is $r^*_{c,1} = f^{\omega,*}_{c,1} = 1$ and $r^*_{c,2} = f^{\omega,*}_{c,2} = 0$ with $z^*_{\text{IRP}} = -1$. The (unique) optimal policy is $r^*_{c,1} = f^{\omega,*}_{c,1} = 0$ and $r^*_{c,2} = f^{\omega,*}_{c,2} = 1$ with $z^* = 1$; this policy is a (best) SRP. From Theorem 4, we have

$$z^* - z^*_{\text{IRP}} = 1 - (-1) = 2 \le 2 = (1 - (-1))1 = (\bar{B} - \underline{B})A.$$

■

We conclude this section by mentioning that both an IRP and a SRP might be optimal policies for model (1); i.e., the model may have multiple optimal solutions. The following theorem provides a sufficient condition in this regard.

Theorem 5. Let $B^{\omega}_{c,t} = B^{\omega}_{c} \ge 0 \;\; \forall t \in T, \omega \in \Omega, c \in C$ and $\exists \bar{t} \in T : \sum_{t'\in T: t' \ge \bar{t}} S^{\omega}_{c,t'} > \sum_{t'\in T: t' \ge \bar{t}} A_{t'} \;\; \forall c \in C, \;\; \omega \in \Omega$ $\sum_{t'\in T: t' < \bar{t}} A_{t'} > 0$. Then there exists an IRP and a SRP which are both optimal policies for model (1).

Proof. Consider that there exists an optimal policy for model (1). Then, there are three cases which we distinguish.

i. Consider that the given optimal policy is an IRP. Then, we employ the same construction as in Theorem 3 to construct a SRP with an objective function value at least that of the optimal IRP. Further, from the hypothesis, the objective function value remains the same; thus, the constructed SRP is also an optimal policy.
ii. Consider that the given optimal policy is a SRP and that $\sum_{c\in C, t\in T} r_{c,t} = \sum_{t\in T} A_t$. Then, $\exists t, \tau \in T : \tau < t$ and $\bar{c} \in C$ with $\sum_{c\in C} r_{c,t} > A_t$, $r_{\bar{c},t} > 0$ and $\sum_{c\in C} r_{c,\tau} < A_\tau$. We assign the quantity $\min\{r_{\bar{c},t}, \sum_{c\in C} r_{c,t} - A_t, A_\tau - \sum_{c\in C} r_{c,\tau}\} > 0$ to $r_{\bar{c},\tau}$. Since, $B^{\omega}_{c} \ge 0$ is independent of t, this reallocation does not change the objective function value. Then, we continue this process until the policy becomes an IRP. Note that there are at most $(|C| \cdot |T|)^2$ such reallocations.
iii. Consider that the given optimal policy is a SRP and that $\sum_{c\in C, t\in T} r_{c,t} < \sum_{t\in T} A_t$. In this case, since $B^{\omega}_{c} \ge 0$, we just "fill up" the remaining resources to obtain an IRP; the objective function value remains unchanged due to the optimality of the SRP. ■

Theorem 5 allows a policymaker to construct a different optimal policy without losing a given optimal policy's efficiency. Such a change of policy finds value in the face of criticism and pressure from political opponents, see, e.g., (Singer et al., 2020) for a discussion of the bipartisan politics during the Zika pandemic in the US. To conclude, we note that if the conditions of Theorem 5 are satisfied, then Theorem 2, Theorem 3 and Theorem 4 also hold; specifically, $z^* - z^*_{\text{IRP}} = 0$. However, we note further that in Theorem 5, the quantity $(\bar{B} - \underline{B})A$ might be > 0 since the benefits might differ for different counties $c \in C$ and scenarios $\omega \in \Omega$. In the next section, we study methods to estimate or determine the parameters of model (1).

3. Estimation of the demands and benefits of resources

Our work employs realistic scenarios used by public health officials in Texas for pandemic preparedness programs and exercises. To generate these scenarios, we use the state-of-the-art simulation-based Texas Pandemic Flu Exercise (TPFE) decision-support tool designed for real-time decision-making by the Texas Department of State Health Services (Texas Advanced Computing Center & The University of Texas at Austin, 2020). This toolkit employs a high-fidelity compartmental model to determine the trajectories of the pandemic's progression. We provide only a

summary of our scenario generation procedure from the TPFE toolkit here, see Appendix B for details.

Two input parameters define a scenario in model (1): $B^{\omega}_{c,t}$ and $S^{\omega}_{c,t}$. By running the TPFE tool, we obtain estimates for the number of individuals infected with the focal pathogen in county c at time t in scenario ω. The TPFE toolkit's output of infected individuals in a (c,t,ω) triplet is further stratified by five age-groups and two risk levels. We let $g \in G$ denote this set of infected population groups; correspondingly, $U^{\omega}_{c,t,g}$ denotes the demand for resources by individuals of group g in a (c,t,ω) triplet. Then, the total number of genuinely infected people seeking resources is given by $\tilde{S}^{\omega}_{c,t} = \sum_{g\in G} U^{\omega}_{c,t,g}$; we refer to $\tilde{S}^{\omega}_{c,t}$ as the sick population. Since one individual seeks one dose of a resource, $\tilde{S}^{\omega}_{c,t}$ is the demand for resources by the sick population alone.

The quantity $B^{\omega}_{c,t}$ denotes the therapeutic societal benefit of providing resources that varies across the three dimensions (c,t,ω). Urban areas often have better access to resource providing healthcare infrastructure than rural areas, see, e.g, (Huang et al., 2017); this demonstrates the dependence of the benefit on the geography, c. Further, the benefit of allocating the same resource may vary with time, t, as new research on the novel disease outbreak is conducted, thereby updating the pandemic's epidemiological parameters (van der Weijden et al., 2013). Naturally, the benefit is uncertain across different scenarios, $\omega \in \Omega$, in the absence of precise information of its spread. For a detailed discussion on several metrics to determine benefits, see, e.g., (Matrajt et al., 2020). Three metrics to quantify $B^{\omega}_{c,t}$ include: (i) the number of lives saved, (ii) the number of life-years saved, and (iii) the number of hospitalizations averted (Singh, 2016). As novel disease parameters are often hard to estimate, our choice of the benefit is guided by the use of parameters that are more easily available, particularly in the early stages of an outbreak. Thus, we consider (i) the number of lives saved by the use of therapeutic resources in a (c,t,ω) triplet as our measure of the benefit.

The efficacy of therapeutic resources may vary by age and preexisting chronic conditions of the infected recipient. For example, the drug aspirin affects children and adults differently (Hall, 1986; Santos, 2015). Similarly, famciclovir—a therapeutic drug against herpes—is known to preferentially affect populations aged above 50 years (Perry & Wagstaff, 1995). We define the efficacy of a single dose of a resource in averting the death of a genuinely infected member of a population group, $g \in G$, as

$$\beta_g = \mathbb{P}(\text{death}|\text{without resources})_g - \mathbb{P}(\text{death}|\text{with resources})_g. \tag{4}$$

By definition, $-1 \le \beta_g \le 1$. In Appendix A, we provide estimates of β_g that we use for the computational experiments in this work. Then, the therapeutic benefit of a resource to genuinely infected individuals in a (c,t,ω) triplet is given by $\tilde{B}^{\omega}_{c,t} = \frac{\sum_{g\in G}\beta_g U^{\omega}_{c,t,g}}{\sum_{g\in G} U_{c,t,g}}$.

Next, we seek to estimate the demands and benefits of resources consumed by the worried-well population that is not actually infected by the focal pathogen. A precise estimate of such demand is outside the scope of this work; indeed, very few studies address this important epidemiological and healthcare policy question (Lipsitch et al., 2011; Siddiqui & Edmunds, 2008). We assume an additional constant fraction, $0 \le \alpha < 1$, of the population is worried-well; i.e., the demand of resources by the worried-well is $\alpha \sum_{g\in G} U^{\omega}_{c,t,g}$. The condition $\alpha < 1$ ensures that the population of the worried-well is no more than the population of the sick, throughout the pandemic. Scarce resources received by the worried-well populations are detrimental overall by depriving lifesaving opportunities for infected individuals. We capture this *lost opportunity* fraction, γ_{ww}, as the negative mean of the resource's efficacy across the infected groups i.e., $\gamma_{\text{ww}} = -\sum_{g\in G}\frac{\beta_g}{|G|}$. Then, the *overall* demand for a (c,t,ω) triplet is the sum of the sick and the worried-well populations, while the overall per-person benefit is the weighted average of the benefits of the sick and worried-well demands. These are given, respectively, by

$$S^{\omega}_{c,t} = (1+\alpha)\sum_{g\in G} U^{\omega}_{c,t,g} \quad \text{and} \quad B^{\omega}_{c,t} = \frac{1}{1+\alpha}(\tilde{B}^{\omega}_{c,t} + \alpha\gamma_{\text{ww}}). \tag{5}$$

We conclude this section with a simple observation that we make use of in our computational experiments of Section 5. For triplets $(c,t,\omega) : \tilde{S}^{\omega}_{c,t} = 0$, we have $S^{\omega}_{c,t} = 0, \forall \alpha \in [0,1)$. From the boundary condition in Equation (1j), these triplets are excluded from any contributions to the objective function of model (1); thus, we fix the corresponding second-stage variables as $f^{\omega}_{c,t} = x^{\omega}_{c,t} = 1$. Optimization over instances of model (1) with a sparse S matrix benefits from this reduction in the number of decision variables. Such a sparse structure arises especially when demand is forecast using so-called intermittent demand forecasting methods (Nikolopoulos, 2021). Further, when demand is sporadic or intermittent, the distribution of S over time might have a number of zeros before finally a positive demand appears. If the string of zero demands occurs toward the end of the time horizon—such as for distributions characterized by long and narrow tails—the following proposition further helps reduce the number of decision variables a priori to the optimization. Demand for critical medical drugs during the course of an epidemic frequently follow such narrow tailed distributions (Flahault et al., 1988).

Proposition 2. *Define* $\bar{t}_c := \min_{t\in T}\{t \mid \sum_{t'\in T: t'\ge t}\sum_{\omega\in\Omega} S^{\omega}_{c,t'} = 0\}, \forall c \in C$. *An optimal solution to model (1) exists with* $r_{c,t} = 0, \forall t \ge \bar{t}_c, c \in C$ *and* $f^{\omega}_{c,t} = x^{\omega}_{c,t} = 1, q^{\omega}_{c,t} = q^{\omega}_{c,\bar{t}_c-1}$, $\forall t \ge \bar{t}_c, c \in C, \omega \in \Omega$.

Proof. Since $B^{\omega}_{c,t}S^{\omega}_{c,t} = 0$, $\forall t \ge \bar{t}_c, c \in C, \omega \in \Omega$, the objective function value does not change for triplets $(c,t,\omega) : t' \ge \bar{t}_c, c \in C, \omega \in \Omega$. Then the feasible solution $r_{c,t} = 0$, $\forall t \ge \bar{t}_c, c \in C$ is also optimal, and the hypothesized values of the

second-stage variables $(f^{\omega}_{c,t}, q^{\omega}_{c,t}, x^{\omega}_{c,t})$ follow directly from constraints (1b)-(1f). ■

In the next section, we analyze the detrimental impact caused by the worried-well populations.

4. Impact of benefits on release policies

4.1. Resources with negative efficacies

It follows from equation (5) that $B^{\omega}_{c,t}$ can be negative for some (c,t,ω) triplets depending on particular values of β_g and $U^{\omega}_{c,t,g}$. Then, the worried well individuals provide an overall detrimental effect by consuming scarce resources meant for at-risk populations; see, e.g., (Watkins, 2005). We provide two situations where this may happen.

i. First, certain drugs can have high cardiotoxicity levels, or cause toxic reactions in particular population groups; see, e.g., (Boelaert et al., 2002) for toxic reactions of drugs for visceral leishmaniasis, and (Peters et al., 2007) for possible adverse reactions of sulfadoxine/pyrimethamine use in pregnant women.
ii. Second, by consuming scarce resources meant for sick individuals, the worried-well population cause society to lose opportunities in the future stages of the pandemic; e.g., worried-well populations can lead to resource shortages (Watkins, 2005) and overcrowded hospitals (Ungchusak et al., 2012).

The framework in Section 3 handles the first of the above-mentioned two situations with $\beta_g < 0$ for the corresponding population-group g. We model the second situation by assigning resource pickups by the worried-well populations with negative values of β; hence, in Section 3 we consider $\gamma_{\text{ww}} < 0$. From our definition of γ_{ww}, this condition is equivalent to $\sum_{g\in G}\beta_g > 0$; or, a dose of a critical medical resource positively benefits at least one sick population-group g. Intuitively, if a resource is not harmful to any of the population-groups, then as long as there are not "too many" worried-well individuals there is always a corresponding non-negative benefit of releasing the resource. Proposition 3 provides a precise definition of "too many".

Proposition 3. *Let $\beta_g \geq 0, \forall g \in G$. For sufficiently small values of $\alpha \in [0,1)$, we have $B^{\omega}_{c,t} \geq 0, \forall t \in T, \omega \in \Omega, c \in C$.*

Proof. From Equation (4) and the definition of γ_{ww}, it follows that for any triplet (c,t,ω), $B^{\omega}_{c,t} \geq 0$ if and only if $\sum_{g\in G} U^{\omega}_{c,t,g}[\beta_g - \frac{\alpha}{|G|}\sum_{g\in G}\beta_g] \geq 0$. Under the hypothesis, the result holds for all $\alpha \in \left[0, \frac{|G|\min_{g\in G}\beta_g}{\sum_{g\in G}\beta_g}\right]$. ■

Proposition 3 provides a sufficient condition—a small enough value of α—for all the per-person benefits to be non-negative *independent* of the corresponding demand $S^{\omega}_{c,t}$. In contrast, even for large values of $\alpha < 1$, the condition $\beta_g - \frac{\alpha}{|G|}\sum_{g\in G}\beta_g < 0, \forall g \in G$ cannot hold true, since $\sum_{g\in G}\beta_g > 0$. In other words, the presence of a large fraction of worried-well population alone does not guarantee a negative benefit of the triplet (c,t,ω). However, if there exists at least one group $g' \in G$ with $\beta_{g'} < 0$ and its population $U^{\omega}_{c,t,g'}$ is sufficiently large, then resources are wasted on this group; this wastage can be larger than the benefit throughout the pandemic, and it is more beneficial to not release resources at all. This situation finds application for non-prescription drugs, particularly high-risk therapeutic drugs such as opioids and laxatives, that are liable to misuse as they are easily available over the counter; consequently, pharmacists devise policies such as keeping them away from sight (Hughes et al., 1999). Such preferential policies, although not completely fair, allow a user-differentiated allocation mechanism in favor of societal efficiency based on the resource's potential for misuse. This can be modeled within our framework by a risk-based adjustment to the efficacy of the resource, $\beta_{g'}$, for the group, g', most likely to misuse the resource. The following proposition summarizes this discussion.

Proposition 4. *Let $B^{\omega}_{c,t} \leq 0, \forall t \in T, \omega \in \Omega, c \in C$. Then, $r_{c,t} = 0, \forall t \in T, c \in C$ is optimal for model (1).*

Proof. Since $B^{\omega}_{c,t} \leq 0, \forall t \in T, c \in C$, we have $z^* \leq 0$. Consider the feasible solution $r_{c,t} = 0$. From Theorem 1 (ii) we have $f^{\omega}_{c,t} = 0, \forall t \in T, \omega \in \Omega, c \in C$; this solution provides $z^* = 0$. ■

4.2. Resources with equal efficacies

Next, consider a resource with the same efficacy for all infected population groups; i.e., $\beta_g = \bar{\beta}, \forall g \in G$. Although some prophylactic drugs have shown to demonstrate almost age-independent efficacy—such as the BNT162b2 mRNA vaccine developed by Pfizer and BioNtech against COVID-19 (Polack et al., 2020)—there are only a few therapeutic resources, to the best of our knowledge, that have the potential to induce such response. Examples include therapeutic vaccines for cancer (Melief et al., 2015) and HIV (Graziani & Angel, 2015). Then, the per-person benefit, $B^{\omega}_{c,t} = \frac{1-\alpha}{1+\alpha}\bar{\beta}$, varies neither across locations nor across time. The following proposition follows from Theorem 2, and says that such resources provide the largest benefit when they are immediately released.

Proposition 5. *Let $\beta_g = \bar{\beta} > 0 \quad \forall g \in G, \quad \alpha \in [0,1)$; i.e., $B^{\omega}_{c,t} = \frac{1-\alpha}{1+\alpha}\bar{\beta} > 0, \forall t \in T, \omega \in \Omega, c \in C$. Then, the IRP is optimal.*

Proof. Under the hypothesis, $B^{\omega}_{c,t} = B^{\omega}_{c,t+1} > 0, \forall t \in T, \omega \in \Omega, c \in C$. The result immediately follows from Theorem 2. ■

Proposition 5 has important policy implications for the early stages of a pandemic. It suggests that if a newly approved resource is provably effective and in an unbiased manner across the infected groups, then the optimal policy is to release the resource immediately. In other words, despite the resource's scarcity, it is not advisable to hold it for subsequent release. In line with this reasoning, several

countries followed the IRP by approving regulatory or limited use of drugs against the SARS-COV-2 virus despite limited testing (Kyriakidis et al., 2021). Currently, there is limited epidemiological research to determine age-specific efficacy of critical therapeutic resources; e.g., therapeutic vaccines against major chronic infectious diseases are still not widely available (Boukhebza et al., 2012). In the next section, we provide computational experiments under a variety of input parameters for a resource that does not satisfy the conditions of Proposition 5.

5. A Numerical study

5.1. Setup

We run all computational experiments using CPLEX 12.10 with GAMS on the Frontera supercomputer at the Texas Advanced Computing Center (Stanzione et al., 2020). We use a time limit of 42,000 seconds; an optimality gap of 0.5% and 1% for the SRP and IRP instances, respectively; and, set `lpmethod` to 4 to use the barrier method. We ran instances with a larger time limit as well, the results do not differ significantly. We use a time horizon of 15 months, i.e., $|T| = 15$. This number provides a conservative estimate of the length of a pandemic's progression envisioned by epidemiologists and policymakers at the onset of the disease, see, e.g., (Katella, 2021); in this time frame all our demand scenarios from the TPFE rise to a peak and then fall to zero (see Appendix B for details). We randomly sample $|\Omega| = 50$ for each of the six pandemics using the methodology described in Section 3 and Appendix B, and use $|C| = 254$ for all the counties of Texas.

Figure 1 provides an example of the scenarios we consider for two counties of Texas. Here, we plot the 50 scenarios that we consider later in Table 1 for the 2020 pandemic. Freestone County is a county in Texas with a population just over 19,000; i.e., at the peak almost 12% of the population was infected. Harris County is the most populated county of Texas—with a population of over 4.7 million. The scenarios indicate a greater diversity in Freestone county.

We compare outputs using three availability regimes with a million doses of resources each: (I) $A_1 = 10^6$, $A_t = 0, \forall t \neq 1$, (II) $A_3 = 10^6$, $A_t = 0, \forall t \neq 3$, and (III) $A_1 = A_3 = 0.5 \cdot 10^6$, $A_t = 0, \forall t \neq 1, 3$; i.e., in the first setting we assume a million resource doses are immediately available at the start of the pandemic, in the second setting we assume the infected in the first two months of the pandemic are without any resources, while in the third setting the same quantity of resources are released equally over two periods. These availability regimes ensure the resource is sufficiently scarce; past studies estimate a likelihood of 0.036 per person for receiving a critical resource in a national emergency (Rule, 2021) and the total population of Texas is about 29 million. We further assume four regimes for the worried-well population described in Section 3: $\alpha = 0, 0.10, 0.15, 0.20$. Next, we compare the IRP and SRP for these parameters.

5.2. Results

Table 1 summarizes our results. The "Gain" columns measure the extent to which the reported IRP solution lags behind the reported SRP solution; this quantity is an estimate of the number of lives that are potentially lost if the IRP is followed. We note that these results are conservative if either of the models is solved suboptimally (as indicated by *). As it turns out, the IRP instances are computationally more challenging than the SRP instances; 46 and 42 of the 72 instances in Table 1 were not solved to optimality for the SRP and IRP, respectively. For these particular instances, optimality gaps range from 0.5% to 1.02% for the SRP, but from 1.1% to 33.0% for the IRP; see, Table S4 in Appendix C for additional detailed results. We draw several conclusions from our computational experiments.

First, we note that an increase in α decreases the expected benefit (lives saved) for both the SRP and IRP, as resources consumed by worried-wells have a negative effect for all the pandemics and availability regimes. For example, compare

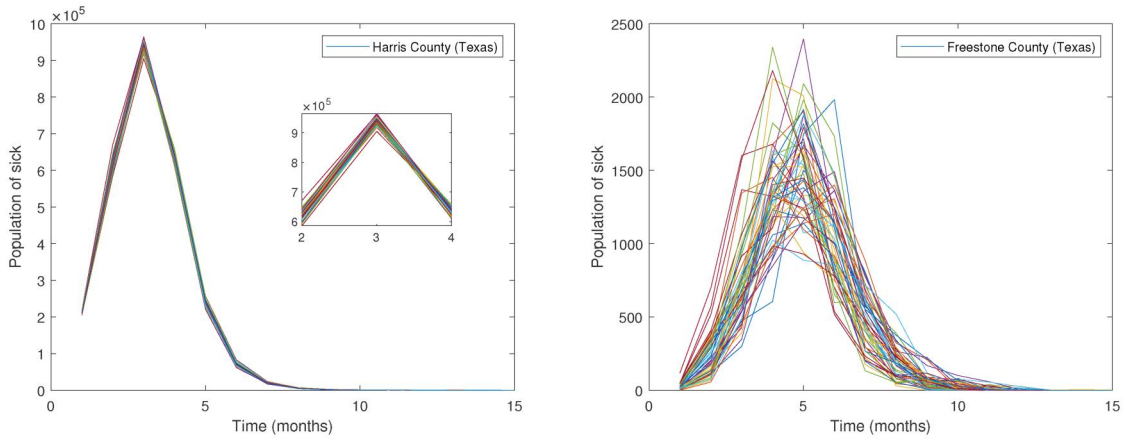

(a) Harris County (b) Freestone County

Figure 1. Spaghetti plot of the number of infectious populations in (a) Harris County—the most populated county of Texas—with a population of over 4.7 million, and (b) Freestone county—with a population of over 19,000. The inlet for harris county zooms in around the peak for $t = 2, 3, 4$. for details, see section 5 and appendix B.

Table 1. All SRP and IRP instances are solved to a MIP gap of 0.5% and 1% with a maximum time limit of 42,000 seconds; those marked with a * did not converge in the time limit. For both the SRP and the IRP we report the best available solution. The "Gain" column denotes the conservative gap between the best bound of the SRP and the best available solution of the IRP. For details, see section 5.2

Instance			Expected Benefit			Instance			Expected Benefit		
Pandemic	Availability	α	SRP	IRP	Gain (%)	Pandemic	Availability	α	SRP	IRP	Gain (%)
1918	I	–	27,619.4*	24,175.3	13.0%	1928	I	–	27,635.1	24,100.5	13.2%
(Spanish		0.10	20,822.2*	17,358.4*	17.2%	(Influenza		0.10	20,872.7	17,359.8*	17.2%
flu)		0.15	17,887.4*	14,478.0*	19.7%	epidemic)		0.15	17,877.3*	14,651.5*	18.7%
		0.20	15,157.1*	12,143.3	20.7%			0.20	15,170.4*	11,919.3*	22.1%
	II	–	27,376.8	26,226.2	4.7%		II	–	27,383.0	26,207.0	4.8%
		0.10	20,607.0	19,623.0	5.3%			0.10	20,622.0	19,531.1	5.8%
		0.15	17,624.9*	16,633.7	6.3%			0.15	17,678.7	16,595.6	6.6%
		0.20	14,946.0*	13,880.3	7.7%			0.20	14,937.2*	13,875.9	7.8%
	III	–	27,640.0	25,189.9*	9.3%		III	–	27,663.7	20,091.0*	27.7%
		0.10	20,840.1*	14,173.1*	32.4%			0.10	20,882.4	18,740.1*	10.7%
		0.15	17,890.1*	15,286.2*	15.2%			0.15	17,922.4*	15,523.6*	13.9%
		0.20	15,156.7*	9,079.3*	40.7%			0.20	15,163.3*	12,832.5*	16.2%
1957	I	–	27,650.9*	23,886.0*	14.1%	1968	I	–	27,680.2	24,207.6	13.0%
(Asian		0.10	20,864.9*	17,504.6*	16.7%	(Hong Kong		0.10	20,870.0*	17,393.7*	17.1%
flu)		0.15	17,888.6*	14,436.0*	20.0%	flu)		0.15	17,932.1*	14,608.6*	19.0%
		0.20	15,232.6*	11,781.4*	23.1%			0.20	15,170.9*	11,907.5*	22.2%
	II	–	27,416.4	26,245.0	4.8%		II	–	27,392.1	26,256.1	4.6%
		0.10	20,646.8	19,650.3	5.2%			0.10	20,619.2	19,595.7	5.4%
		0.15	17,696.6	16,646.4	6.4%			0.15	17,661.1*	16,642.9	6.3%
		0.20	14,973.5*	13,919.7	7.6%			0.20	14,948.7*	13,916.3	7.5%
	III	–	27,650.8*	25,271.7*	9.1%		III	–	27,679.9	23,637.3*	15.0%
		0.10	20,861.3*	18,344.7*	12.7%			0.10	20,840.4*	18,643.9*	11.1%
		0.15	17,919.9*	15,722.3*	12.9%			0.15	17,901.0*	16,010.3*	11.2%
		0.20	15,236.5*	9,143.0*	40.4%			0.20	15,216.6*	12,474.0*	18.5%
2009	I	–	27,693.6	24,235.9	12.9%	2020	I	–	27,601.2	24,126.8	13.0%
(swine		0.10	20,867.2*	17,454.3*	17.0%	(COVID-19)		0.10	20,842.2	17,443.5*	16.7%
flu)		0.15	17,912.6*	14,691.2*	18.7%			0.15	17,882.2*	14,515.7*	19.3%
		0.20	15,262.0*	12,123.9*	21.1%			0.20	15,145.8*	12,073.1*	21.0%
	II	–	27,418.7	26,260.6	4.7%		II	–	27,367.5	26,186.3	4.8%
		0.10	20,627.2*	19,574.6	5.7%			0.10	20,581.6	19,583.3	5.3%
		0.15	17,681.2*	16,678.5	6.3%			0.15	17,602.0*	16,626.5	6.2%
		0.20	15,007.7*	13,905.8	7.9%			0.20	14,936.9	13,854.4	7.7%
	III	–	27,702.6	25,123.4*	9.8%		III	–	27,595.4	20,069.7*	27.6%
		0.10	20,917.9*	18,294.9*	13.0%			0.10	20,799.2*	17,335.3*	17.2%
		0.15	17,962.2*	15,540.6*	14.0%			0.15	17,843.1*	15,841.9*	11.9%
		0.20	15,226.3*	11,726.8*	23.7%			0.20	15,133.5*	12,116.1*	20.6%

the expected benefit for the SRP from the first and fourth rows of the 2020 pandemic (i.e., availability regime (I)) in Table 1. Here, at least 12,053 lives are lost due to resources consumed by the worried-wells if they represent 20% of the sick population. The situation is similar for other instances. Resources released too early are consumed by the worried-wells leading to a decrease in the overall benefit; gaps for $\alpha = 0.2$ are as high as 41% for the 1918 pandemic with availability regime (III). If a laboratory diagnosis is mandatory for the availability of resources, the number of worried-wells can be significantly reduced rendering more resources available for the genuinely sick; such allocation measures that prioritize "evidenced-based medicine" over "emotional biases" have been advocated before (Johnson & Levine, 2013).

Second, gaps between the SRP and IRP are largest for the availability regime (III) (average gap = 18.1%), and smallest for availability regime (II) (average gap = 6.1%). The variability across the six outbreaks is small. The peak of the total sick population, $\sum_{c\in C} S^{\omega}_{c,t}$, occurs at $t = 3$ for all of the 50 scenarios. Thus, we empirically observe that if the availability of resources times with the peak demand (such as, availability regime (II)) there are relatively fewer lives lost by not following the optimal SRP. Further, this loss is nearly independent of the fraction of worried-well; the standard deviation of the loss for availability regime (II) is only 1.1%, while that of availability regime (I) and (III) are 3.2% and 9.3%, respectively. This suggests another intuitive result—when therapeutic resources become available only during the course of the outbreak, releasing them immediately is more justifiable than if resources were available from the start of the outbreak.

Our results in the preceding discussion include scenarios sampled from a given pandemic. However, predicting the course of a pandemic is hard (Taubenberger et al., 2007). Thus, to further validate our results, we additionally sample a new (independent to the previously-considered scenarios) set of 50 mixed scenarios *across* the six pandemics. We observe similar trends to those we previously report (see Table S5 in Appendix C for details). Although this does not provide an out-of-sample performance validation in the traditional sense for stochastic optimization models (see, e.g., (Mak et al., 1999)), similar values of the optimal objective function values to those in Table 1 suggest the qualitative release policies does not differ significantly with the disease parameters.

To provide further support of our results, consider the following null and alternate hypotheses:

$$H_0 : \mu_{\text{SRP}} = \mu_{\text{IRP}} \qquad \text{and} \qquad H_A : \mu_{\text{SRP}} \neq \mu_{\text{IRP}}. \tag{6}$$

Here, μ_{SRP} and μ_{IRP} denote the means of either (i) the sick population that receives resources, $\sum_{c\in C} f^{\omega}_{c,t} \tilde{S}^{\omega}_{c,t}$, (ii) the

sick population that does not receive resources, $\sum_{c\in C}(1-f^{\omega}_{c,t})\tilde{S}^{\omega}_{c,t}$, or (iii) the worried-well population that receives resources, $\sum_{c\in C}\alpha f^{\omega}_{c,t}\tilde{S}^{\omega}_{c,t}$, for the SRP and IRP, respectively. Then, for the mixed scenarios with availability regimes (I) and (III), for both the sick and the worried-well, we reject H_0 at a 97.5% level for $t = 1,\ldots,10$; for the remaining times the population is significantly smaller. Since $\alpha > 0$, trends for the resources wasted mirror those for the sick population. Specifically, for $t = 1,2,3$, $\mu_{\text{SRP}} < \mu_{\text{IRP}}$ while $\mu_{\text{SRP}} > \mu_{\text{IRP}}$ for $t = 4,\ldots,10$. Further, since peak demand occurs at $t = 3$ for all scenarios and the overall expected benefit of the SRP is larger than IRP, we conclude that the IRP exhausts a majority of its resources before the peak is realized. These depleted resources are consumed not only by the worried-wells, but also by population subgroups with smaller values of β_g.

Trends for the sick that do not receive resources are, naturally, the reverse. For $t = 1,2,3$ $\mu_{\text{SRP}} > \mu_{\text{IRP}}$ while $\mu_{\text{SRP}} < \mu_{\text{IRP}}$ for $t = 4,\ldots,7$ (for $t = 8,\ldots,10$ we cannot reject H_0). Similar results hold for instance II; the only difference is that since no resources are available for $t = 1,2$ there is no test at these times. Table S6 in the Appendix summarizes the results for the hypotheses testing. Although our results are naturally dependent on the input parameters, the rejection of the null hypotheses provides policy recommendations. If resources are available at the start of the outbreak—but are scarce—then, policies that advise holding on to resources and releasing in a sequential manner are preferable.

The above results show that following the IRP provides a myopic benefit that is rapidly observed by the public, but it can be precarious in the long run. We conclude by acknowledging, as we mention in Section 1, that holding on to resources when there is a demand for them is a challenging question from an ethical standpoint (Emanuel et al., 2020). Public pressure to immediately respond, puts an increased burden on both the scientific and political communities to defend farsighted choices (Squazzoni et al., 2020). Our results help provide data-driven guidance that a policymaker may use to support particular release policies.

6. Conclusions

We present a stochastic optimization model to determine allocation decisions of scarce therapeutic resources at the onset of a pandemic. Optimal policies, under the stochastic programming framework we pursue, are governed by maximizing an expected benefit. We distinguish two types of resource release policies available to a policymaker. A strategy of releasing resources as soon as they become available is intuitive, and helps a policymaker relieve potential pressure arising from the public who is seeking this resource. Surprisingly this greedy strategy could be optimal and we provide sufficient conditions under which this is indeed so. However, as our computational experiments demonstrate, following this policy when the sufficient conditions are not met can lead to poor long-term benefits; e.g., significantly lesser number of lives saved. Then, a sequential release policy of holding on to resources and gradually releasing them provides the maximum benefit. In our computational study across six different pandemics in Texas, the sequential release policy saves between 4.6% and 40.7% more lives.

Problems with this background also find application in a number of other settings, where resources are allocated from a centralized inventory. Following emergency disasters, several types of resources that provide relief are both valuable and scarce lending them a critical nature (Timbie et al., 2013). Such disasters further produce an immediate demand for relief-providing resources that existing supply cannot handle (Brotcorne et al., 2003; Doan & Shaw, 2019). Further, this setting finds value in allocation of high-value spare parts that are immediately required by the customer, e.g., aircraft components or inventory for the defense industry (Simao & Powell, 2009). We note that the unifying feature of both these settings, and that of the healthcare planner, is that once resources are allocated to a specific class of users—such as pharmacy chains or geographic regions—redeployment elsewhere is not possible; indeed, this was observed for the case of distributing antivirals during the 2009 H1N1 pandemic (Singh et al., 2015).

As in any modeling framework, there are several limitations in our work that we present next. We begin with limitations in our optimization framework. First, we consider the entire release schedule, $r_{c,t}$, is made before the scenarios are realized. A higher fidelity model would consider dynamic decision-making where resource pickups in the current time influence resource pickups in the following times. Multi-stage stochastic programs, as opposed to two-stage stochastic programs, allow such decision-making (Rebennack, 2016; Shapiro et al., 2009). Second, while our decision-making process explicitly depends on the underlying uncertainty, the uncertainty is not determined by the decisions. Mathematically, we do not consider so-called endogenous uncertainty in our model (Jonsbråten et al., 1998; Krasko & Rebennack, 2017). Finally, we mention that model (1) may be considered as defined with respect to a *true* distribution of the uncertainty, with a sample space Ω', that could be discrete or continuous. Then, we consider finitely many scenarios, $|\Omega|$, of this true distribution obtained via a Monte Carlo sampling procedure to obtain a so-called sample average approximation (SAA) of the true model (Mak et al., 1999). Our work views these scenarios as equally likely and considers $\Omega' = \Omega$.

Next, we present the epidemiological limitations of our work. First, we do not consider a reduction in transmission of the disease after populations receive resources; i.e., our model assumes the parameters of the pandemic do not change over time. Including a time-dependent R_0 could help with this limitation. Second, our simulations do not consider a reduction in the number of future infections following an intervention in the current time step. Simulating such scenarios allows an adjustment of the demand after observing a release schedule. The TPFE tool is capable of handling vaccine interventions if the timing of the vaccine availability is known. However, this assumption is relatively mild for our work as we consider therapeutic (as opposed to

prophylactic) resources; intake of prophylactic drugs notably reduces transmission of a virus, thereby resulting in lower demands in future time periods. Third, the production schedule of resources, A_t, could itself be uncertain. Further, our model aggregates demand over all the infected individuals thereby losing characteristics of age and risk status; future work could consider these subgroups separately. Then, in addition to ensuring equal access to resources for the infected and uninfected populations, measures to ensure equity among the infected populations (e.g., dispersion or the so-called Gini index) could be established. Such a segregated allocation allows consideration of misuse of resources (see, Section 4) by certain groups; then, constraints that limit uptake of their resources could be imposed. Further, the fraction of population that is worried-well is assumed to be a constant; however, this fraction might change with time as the perception of the pandemic evolves. Currently, to the best of our knowledge, there is no state-of-the-art way to estimate worried-well populations. Finally, we assume only a single dose of a resource is required per individual and we do not allow multiple resource types (e.g., those with a low-risk of misuse, such as antivirals, and those with a high-risk of misuse, such as opioids) to be distributed simultaneously.

Third, our study has computational limitations as well. The two-stage stochastic optimization model we consider is computationally expensive—even on the fastest academic supercomputer in the US with a time limit of half a day and a modest number of scenarios, we were unable to solve all the instances considered to optimality. Future work could examine tailored algorithms that allow faster solutions for a larger sample size. Future work could also include necessary conditions for optimality of the IRP (as opposed to sufficient conditions that we provide), and/or a quantification of the optimality gaps.

The data that support the findings of this study are available on the following repository: https://github.com/bissi1/ResourceAllocation.

Acknowledgments

We thank Nedialko Dimitrov for helpful discussions and earlier mathematical formulations of the model presented in this article. We remain grateful to Paul Navratil for his assistance with the Frontera supercomputer and keep him in our memory. We thank Oliver Rehberg for the computational experiments conducted as well as the data organization.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

Both authors were partially supported by the EOSCsecretariat.eu project "Optimal spatiotemporal antiviral release under uncertainty". EOSCsecretariat.eu has received funding from the European Union's Horizon Programme call H2020-INFRAEOSC-05-2018-2019, grant Agreement number 831644. Bismark Singh was further supported by the Deutsche Forschungsgemeinschaft project "Optimal decision making during pandemics" under grant number 490948220.

ORCID

Bismark Singh http://orcid.org/0000-0002-6943-657X
Steffen Rebennack http://orcid.org/0000-0002-8501-2785

References

Arinaminpathy, N., & McLean, A. R. (2008). Antiviral treatment for the control of pandemic influenza: Some logistical constraints. *Journal of the Royal Society, Interface*, *5*(22), 545–553. https://doi.org/10.1098/rsif.2007.1152

Boelaert, M., Ray, D. L., & der Stuyft, P. V. (2002). How better drugs could change kala-azar control. Lessons from a cost-effectiveness analysis. *Tropical Medicine & International Health: TM & IH*, *7*(11), 955–959. https://doi.org/10.1046/j.1365-3156.2002.00959.x

Boukhebza, H., Bellon, N., Limacher, J. M., & Inchauspé, G. (2012). Therapeutic vaccination to treat chronic infectious diseases. *Human Vaccines & Immunotherapeutics*, *8*(12), 1746–1757. https://doi.org/10.4161/hv.21689

Brotcorne, L., Laporte, G., & Semet, F. (2003). Ambulance location and relocation models. *European Journal of Operational Research*, *147*(3), 451–463. https://doi.org/10.1016/S0377-2217(02)00364-8

D'Amour, A., Srinivasan, H., Atwood, J., Baljekar, P., Sculley, D., & Halpern, Y. (2020, January). Fairness is not static: deeper understanding of long term fairness via simulation studies. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20* (pp. 525–534). Acm.

Dantzig, G. B., & Wolfe, P. (1961). The decomposition algorithm for linear programs. *Econometrica*, *29*(4), 767. https://doi.org/10.2307/1911818

Dimitrov, N., Goll, S., Meyers, L. A., Pourbohloul, B., & Hupert, N. (2009). Optimizing tactics for use of the U.S. antiviral Strategic National Stockpile for pandemic (H1N1) influenza, 2009. *PLoS Currents*, *1*, RRN1127. https://doi.org/10.1371/currents.rrn1127

Doan, X. V., & Shaw, D. (2019). Resource allocation when planning for simultaneous disasters. *European Journal of Operational Research*, *274*(2), 687–709. https://doi.org/10.1016/j.ejor.2018.10.015

Duque, D., Morton, D. P., Singh, B., Du, Z., Pasco, R., & Meyers, L. A. (2020). Timing social distancing to avert unmanageable COVID-19 hospital surges. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(33), 19873–19878. https://doi.org/10.1073/pnas.2009033117

Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., & Phillips, J. P. (2020). Fair allocation of scarce medical resources in the time of Covid-19. *The New England Journal of Medicine*, *382*(21), 2049–2055. https://doi.org/10.1056/NEJMsb2005114

Emanuel, E. J., & Wertheimer, A. (2006). Who should get influenza vaccine when not all can? *Science (New York, N.Y.)*, *312*(5775), 854–855. https://doi.org/10.1126/science.1125347

Erlebacher, S. J. (2000). Optimal and heuristic solutions for the multi-item newsvendor problem with a single capacity constraint. *Production and Operations Management*, *9*(3), 303–318. https://doi.org/10.1111/j.1937-5956.2000.tb00139.x

Facher, L. (2020, March). The coronavirus outbreak has left medical supplies in short supply. is the nation's emergency stockpile ready to help?. techreport, STAT.

Feldman, E. R., & Sakhartov, A. V. (2022). Resource redeployment and divestiture as strategic alternatives. *Organization Science*, *33*(3), 926–945. https://doi.org/10.1287/orsc.2021.1474

Felter, C. (2021, March). A guide to global COVID-19 vaccine efforts.

Flahault, A., Letrait, S., Blin, P., Hazout, S., Ménarés, J., & Valleron, A.-J. (1988). Modelling the 1985 influenza epidemic in France. *Statistics in Medicine*, *7*(11), 1147–1155. https://doi.org/10.1002/sim.4780071107

Goodwin, R., Haque, S., Neto, F., & Myers, L. B. (2009). Initial psychological responses to Influenza A, H1N1 ("swine flu"). *BMC Infectious Diseases*, *9*, 166. https://doi.org/10.1186/1471-2334-9-166

Graziani, G. M., & Angel, J. B. (2015). Evaluating the efficacy of therapeutic HIV vaccines through analytical treatment interruptions. *Journal of the International AIDS Society*, *18*(1), 20497. https://doi.org/10.7448/IAS.18.1.20497

Hall, S. M. (1986). Reye's syndrome and aspirin: A review. *Journal of the Royal Society of Medicine*, *79*(10), 596–598. https://doi.org/10.1177/014107688607901014

Hodge, J. G., & O'Connell, J. P. (2006). The legal environment underlying influenza vaccine allocation and distribution strategies. *Journal of Public Health Management and Practice: JPHMP*, *12*(4), 340–348. https://doi.org/10.1097/00124784-200607000-00007

Huang, H.-C., Singh, B., Morton, D. P., Johnson, G. P., Clements, B., & Meyers, L. A. (2017). Equalizing access to pandemic influenza vaccines through optimal allocation to public health distribution points. *PloS One*, *12*(8), e0182720. https://doi.org/10.1371/journal.pone.0182720

Hughes, G. F., McElnay, J. C., Hughes, C. M., & McKenna, P. (1999). Abuse/misuse of non-prescription drugs. *Pharmacy World & Science: PWS*, *21*(6), 251–255. https://doi.org/10.1023/a:1008788726842

Johnson, L.-M., & Levine, D. (2013). Allocation of drugs in short supply. *AMA Journal of Ethics*, *15*(8), 645–648.

Jonsbråten, T. W., Wets, R. J.-B., & Woodruff, D. L. (1998). A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, *82*(0), 83–106. https://doi.org/10.1023/A:1018943626786

Katella, K. (2021, March). Our pandemic year—a COVID-19 timeline. techreport. Yale Medicine.

Kleczkowski, A., Hoyle, A., & McMenemy, P. (2019). One model to rule them all? Modelling approaches across OneHealth for human, animal and plant epidemics. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *374*(1775), 20180255. https://doi.org/10.1098/rstb.2018.0255

Krasko, V., & Rebennack, S. (2017). Two-stage stochastic mixed-integer nonlinear programming model for post-wildfire debris flow hazard management: Mitigation and emergency evacuation. *European Journal of Operational Research*, *263*(1), 265–282. https://doi.org/10.1016/j.ejor.2017.05.004

Kyriakidis, N. C., López-Cortés, A., González, E. V., Grimaldos, A. B., & Prado, E. O. (2021). SARS-CoV-2 vaccines strategies: A comprehensive review of phase 3 candidates. *NPJ Vaccines*, *6*(1), 28. https://doi.org/10.1038/s41541-021-00292-w

Laventhal, N., Basak, R., Dell, M. L., Diekema, D., Elster, N., Geis, G., Mercurio, M., Opel, D., Shalowitz, D., Statter, M., & Macauley, R. (2020). The ethics of creating a resource allocation strategy during the COVID-19 pandemic. *Pediatrics*, *146*(1), e20201243. https://doi.org/10.1542/peds.2020-1243

Lee, J., Lee, H. S. H., Shin, H., & Krishnan, V. (2021). Alleviating drug shortages: The role of mandated reporting induced operational transparency. *Management Science*, *67*(4), 2326–2339. January) https://doi.org/10.1287/mnsc.2020.3857

Liang, B. A., & Mackey, T. K. (2012). Vaccine shortages and suspect online pharmacy sellers. *Vaccine*, *30*(2), 105–108. https://doi.org/10.1016/j.vaccine.2011.11.016

Lipsitch, M., Finelli, L., Heffernan, R. T., Leung, G. M., & Redd, S. C., for the 2009 H1N1 Surveillance Group. (2011). Improving the evidence base for decision making during a pandemic: The example of 2009 Influenza A/H1N1. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, *9*(2), 89–115.

Long, E. F., Nohdurft, E., & Spinler, S. (2018). Spatial resource allocation for emerging epidemics: A comparison of greedy, myopic, and dynamic policies. *Manufacturing & Service Operations Management*, *20*(2), 181–198. https://doi.org/10.1287/msom.2017.0681

Longini, I. M., Halloran, M. E., Nizam, A., & Yang, Y. (2004). Containing pandemic influenza with antiviral agents. *American Journal of Epidemiology*, *159*(7), 623–633. https://doi.org/10.1093/aje/kwh092

Mak, W.-K., Morton, D. P., & Wood, R. K. (1999). Monte Carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters*, *24*(1-2), 47–56. https://doi.org/10.1016/S0167-6377(98)00054-6

Matrajt, L., Eaton, J., Leung, T., & Brown, E. R. (2020). Vaccine optimization for COVID-19: Who to vaccinate first? *Science Advances*, *7*(6), eabf1374. https://doi.org/10.1126/sciadv.abf1374

Matrajt, L., & Longini, I. M. (2010). Optimizing vaccine allocation at different points in time during an epidemic. *PloS One*, *5*(11), e13767. https://doi.org/10.1371/journal.pone.0013767

Melief, C. J., van Hall, T., Arens, R., Ossendorp, F., & van der Burg, S. H. (2015). Therapeutic cancer vaccines. *The Journal of Clinical Investigation*, *125*(9), 3401–3412. https://doi.org/10.1172/JCI80009

Nadiminti, R., Mukhopadhyay, T., & Kriebel, C. H. (2002). Research Report: Intrafirm resource allocation with asymmetric information and negative externalities. *Information Systems Research*, *13*(4), 428–434. https://doi.org/10.1287/isre.13.4.428.70

National Academies of Sciences, Engineering, and Medicine. (2020). *Framework for Equitable Allocation of COVID-19 Vaccine*. The National Academies Press.

Nature. (2021, January). *COVID research updates: COVID vaccines might lose potency against new viral variants*. Nature.

Nikolopoulos, K. (2021). We need to talk about intermittent demand forecasting. *European Journal of Operational Research*, *291*(2), 549–559. https://doi.org/10.1016/j.ejor.2019.12.046

Perry, C. M., & Wagstaff, A. J. (1995). Famciclovir: A review of its pharmacological properties and therapeutic efficacy in herpesvirus infections. *Drugs*, *50*(2), 396–415. https://doi.org/10.2165/00003495-199550020-00011

Persad, G., Wertheimer, A., & Emanuel, E. J. (2009). Principles for allocation of scarce medical interventions. *Lancet (London, England)*, *373*(9661), 423–431. https://doi.org/10.1016/S0140-6736(09)60137-9

Peters, P. J., Thigpen, M. C., Parise, M. E., & Newman, R. D. (2007). Safety and toxicity of sulfadoxine/pyrimethamine. *Drug Safety*, *30*(6), 481–501. https://doi.org/10.2165/00002018-200730060-00003

Polack, F. P., Thomas, S. J., Kitchin, N., Absalon, J., Gurtman, A., Lockhart, S., Perez, J. L., Marc, G. P., Moreira, E. D., Zerbini, C., Bailey, R., Swanson, K. A., Roychoudhury, S., Koury, K., Li, P., Kalina, W. V., Cooper, D., Frenck, R. W., Hammitt, L. L., … Gruber, W. C., C4591001 Clinical Trial Group. (2020). Safety and efficacy of the BNT162b2 mRNA covid-19 vaccine. *The New England Journal of Medicine*, *383*(27), 2603–2615. https://doi.org/10.1056/NEJMoa2034577

Rahmandad, H., & Sterman, J. (2008). Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. *Management Science*, *54*(5), 998–1014. https://doi.org/10.1287/mnsc.1070.0787

Rebennack, S. (2016). Combining sampling-based and scenario-based nested Benders decomposition methods: Application to stochastic dual dynamic programming. *Mathematical Programming*, *156*(1-2), 343–389. https://doi.org/10.1007/s10107-015-0884-3

Rule, T. (2021). Toward a more Strategic National Stockpile. *Texas A&M Law Review*, *9*(1), 49–108. https://doi.org/10.37419/LR.V9.I1.2

Santos, R. M. (2015). *Differential drug response as a function of age* [Ph. D. thesis]. University of Cambridge.

Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2009, January). *Lectures on stochastic programming*. Society for Industrial and Applied Mathematics.

Siddiqui, M. R., & Edmunds, W. J. (2008). Cost-effectiveness of antiviral stockpiling and near-patient testing for potential influenza pandemic. *Emerging Infectious Diseases*, *14*(2), 267–274. https://doi.org/10.3201/eid1402.070478

Simao, H., & Powell, W. (2009). Approximate dynamic programming for management of high-value spare parts. *Journal of Manufacturing Technology Management*, *20*(2), 147–160. https://doi.org/10.1108/17410380910929592

Singer, P. M., Willison, C. E., & Greer, S. L. (2020). Infectious disease, public health, and politics: United States response to Ebola and Zika. *Journal of Public Health Policy*, *41*(4), 399–409. https://doi.org/10.1057/s41271-020-00243-0

Singh, B. (2016). *Optimal spatiotemporal resource allocation in public health and renewable energy* [Ph. D. thesis]. The University of Texas at Austin.

Singh, B. (2020). Fairness criteria for allocating scarce resources. *Optimization Letters*, *14*(6), 1533–1541. https://doi.org/10.1007/s11590-020-01568-1

Singh, B., Huang, H.-C., Morton, D. P., Johnson, G. P., Gutfraind, A., Galvani, A. P., Clements, B., & Meyers, L. A. (2015). Optimizing distribution of pandemic influenza antiviral drugs. *Emerging Infectious Diseases*, *21*(2), 251–258. https://doi.org/10.3201/eid2102.141024

Squazzoni, F., Polhill, J. G., Edmonds, B., Ahrweiler, P., Antosz, P., Scholz, G., Chappin, É., Borit, M., Verhagen, H., Giardini, F., & Gilbert, N. (2020). Computational models that matter during a global pandemic outbreak: A call to action. *Journal of Artificial Societies and Social Simulation*, *23*(2), 10. https://doi.org/10.18564/jasss.4298

Stanzione, D., West, J., Evans, R. T., Minyard, T., Ghattas, O., & Panda, D. K. (2020, July). Frontera: The evolution of leadership computing at the National Science Foundation. In *Practice and Experience in Advanced Research Computing*. ACM.

Taubenberger, J. K., Morens, D. M., & Fauci, A. S. (2007). The next influenza pandemic: Can it be predicted? *Jama*, *297*(18), 2025–2027. https://doi.org/10.1001/jama.297.18.2025

Texas Advanced Computing Center, The University of Texas at Austin. (2020). Texas Pandemic Flu Toolkit.

Timbie, J. W., Ringel, J. S., Fox, D. S., Pillemer, F., Waxman, D. A., Moore, M., Hansen, C. K., Knebel, A. R., Ricciardi, R., & Kellermann, A. L. (2013). Systematic review of strategies to manage and allocate scarce resources during mass casualty events. *Annals of Emergency Medicine*, *61*(6), 677–689.e101. https://doi.org/10.1016/j.annemergmed.2013.02.005

Ungchusak, K., Sawanpanyalert, P., Hanchoworakul, W., Sawanpanyalert, N., Maloney, S. A., Brown, R. C., Birmingham, M. E., & Chusuttiwat, S. (2012). Lessons learned from influenza A (H1N1) pdm09 pandemic response in Thailand. *Emerging Infectious Diseases*, *18*(7), 1058–1064. https://doi.org/10.3201/eid1807.110976

van der Weijden, C. P., Stein, M. L., Jacobi, A. J., Kretzschmar, M. E. E., Reintjes, R., van Steenbergen, J. E., & Timen, A. (2013). Choosing pandemic parameters for pandemic preparedness planning: A comparison of pandemic scenarios prior to and following the influenza A(H1N1) 2009 pandemic. *Health Policy (Amsterdam, Netherlands)*, *109*(1), 52–62. https://doi.org/10.1016/j.healthpol.2012.05.007

Walensky, R. P., & Paltiel, A. D. (2006). Rapid HIV testing at home: Does it solve a problem or create one? *Annals of Internal Medicine*, *145*(6), 459–462. https://doi.org/10.7326/0003-4819-145-6-200609190-00010

Ward, S., Lindsley, A., Courter, J., & Assa'ad, A. (2020). Clinical testing for COVID-19. *The Journal of Allergy and Clinical Immunology*, *146*(1), 23–34. https://doi.org/10.1016/j.jaci.2020.05.012

Watkins, J. (2005). The UK flu vaccine shortage—Who is at fault? *Bmj*, *331*(7529), 1413. https://doi.org/10.1136/bmj.331.7529.1413

World Health Organization. (2016). Medicines shortages: Global approaches to addressing shortages of essential medicines in health systems. *WHO Drug Information*, *30*(2), 180–185.