
A STOCHASTIC PROGRAMMING APPROACH FOR DYNAMIC ALLOCATION OF BED CAPACITY AND ASSIGNMENT OF PATIENTS TO COLLABORATING HOSPITALS DURING PANDEMIC OUTBREAKS

A PREPRINT

Stef Baas

Department of Stochastic Operations Research
University of Twente
7522 NH Enschede, the Netherlands

Sander Dijkstra

Department of Stochastic Operations Research
University of Twente
7522 NH Enschede, the Netherlands

Richard J. Boucherie

Department of Stochastic Operations Research
University of Twente
7522 NH Enschede, the Netherlands

Anne Zander*

Department of Stochastic Operations Research
University of Twente
7522 NH Enschede, the Netherlands

November 28, 2023

ABSTRACT

We consider a region containing several hospitals collaborating to treat patients during an infectious outbreak. Collaboration occurs through the dynamic allocation of hospital bed capacity to infectious patients and by assigning infectious patients to specific hospitals. Scaling up capacity for infectious patients means that beds are removed from regular care, which can only be done simultaneously for all beds in a room. Moreover, as opening rooms for infectious patients takes some lead time, we have to decide on preparing to open rooms for infectious patients ahead of time. We apply a stochastic direct lookahead approach. Each day, we make decisions on room allocation and patient assignment based on the solutions of two stochastic programs with scenarios using short-time forecasts of the number of infectious hospitalizations in the region and the bed occupancy in each collaborating hospital. We aim to balance costs for bed shortages and unutilized beds for infectious patients and opening and closing rooms. We demonstrate that a stochastic lookahead is superior to a deterministic one. We furthermore compare our solution approach with two heuristic strategies in a simulation study based on historical COVID-19 data of a region with three hospitals in the Netherlands. In one strategy, hospitals decide on their capacity allocation individually. In another strategy, we assume a pandemic unit, where one hospital is designated to take all regional infectious patients until full. The numerical results show that our stochastic direct lookahead approach considerably outperforms these heuristics.

Keywords OR in health services · Stochastic programming · Simulation · Pandemic response management · COVID-19

*Corresponding author

1 Introduction

Access to healthcare is a fundamental human right to be maintained during a crisis. Pandemics such as the COVID-19 pandemic or larger outbreaks of infectious diseases such as seasonal influenza, result in seriously ill patients, straining healthcare delivery from a clinical perspective as well as from the perspective of available hospital resources such as staffed beds in the ward and ICU, operating rooms (ORs) and capacity of diagnostics (Bernstein et al., 2022). During infectious disease outbreaks, adequate allocation of healthcare resources to infectious patients is of utmost importance. However, sustaining regular care for non-infectious patients is just as essential. Especially non-urgent care is at risk of being postponed. For example, during the COVID-19 pandemic, regular patients' surgeries and treatments were postponed, leading to a significant loss of healthy life years (de Graaff et al., 2022; Rovers et al., 2022).

Although the COVID-19 pandemic is perceived by many as a rare event, we should consider this as a wake-up call to be ready for future pandemics. The likelihood of COVID-19-like pandemics is rising due to the increasing population, agricultural intensification and expansion, climate change, and increased mobility (Bernstein et al., 2022). It is to be expected that within the coming decades, pandemics will again stress our healthcare system like COVID-19 did, especially with the capacity in healthcare tightening due to both an increase in demand, e.g., due to the aging population and a decrease in available resources, such as a stable decrease of available hospital beds over the last decades (StatLine, 2023). Pandemics and large outbreaks with new or already known pathogens, however, are not the only threat: seasonal influenza poses dynamic operational stress across healthcare systems, causing major impacts outside of declared pandemics every year, with a larger influenza outbreak every few years, causing major problems for the entire chain of care such as during the 2017/2018 influenza outbreak (Shearer et al., 2020). Whatever the type of future pandemic or larger infectious outbreak, we must decide how to adequately allocate limited resources to different patient groups (with infectious patients on one side and regular, non-infectious patients on the other). We must aim to serve the infectious patient population while avoiding over-allocation of capacity to this group to ensure the continuation of regular care to the largest extent possible.

This paper focuses on dynamic bed capacity allocation and patient assignment at a regional level, assuming a set of collaborating hospitals in a region. We propose decision rules to centrally and dynamically allocate bed capacity to infectious patients (which means that we take these beds out of the regular care bed capacity) in the collaborating hospitals and to assign regional infectious patients to those hospitals. The decision rules are based on a stochastic direct lookahead given by solving two stochastic programs with scenarios in sequence each day. There, we take into account that bed capacity can only be allocated in terms of rooms and that it takes a certain number of days until a room can be opened for infectious patients. This is because, for example, regular patients who are currently accommodated in this room must be discharged first. The first stochastic program decides on opening or closing regular care rooms to accommodate infectious patients over the upcoming days. The second stochastic program assigns arriving infectious patients to the collaborating hospitals, given previous capacity allocation decisions. The input scenarios given to the stochastic programs are based on a forecasting method and contain the number of regional patients in need of infectious care and the forecast bed occupancy at individual hospitals (excluding patients in need of assignment) in the region for the coming days. The decision rules are compared to two simpler heuristic decision rules in a simulation study that simulates the evolution of demand for COVID-19 care in hospital wards in the ROAZ (Regional Consultation Body on Acute Care) region *Acute Zorg Euregio* in the Netherlands at the end of 2021. We demonstrate that a stochastic lookahead is superior to a deterministic one. Furthermore, the numerical results show that our stochastic direct lookahead approach considerably outperforms these simpler heuristics. The research leading to this paper had a direct practical impact, as the comparison of one of these heuristics, the pandemic unit, with regional collaboration led to the conclusion for Acute Zorg Euregio that dynamic regional collaboration would be preferred over a pandemic unit in case of a new pandemic outbreak, see, e.g., NOS nieuws (2023).

The remainder of this paper is organized as follows. In the next section, a comprehensive literature review is presented. Section 3 introduces the sequential decision modeling framework used for capacity allocation and patient assignment and also presents the model to describe the evolution of the infectious bed occupancy in the considered region that is

used to generate scenarios. Section 4 presents the stochastic programs and resulting decision rules for opening rooms and assignment of infectious patients. Section 5 describes the setup and results of a simulation study where the decision rules are compared to two simpler heuristic decision rules with an application to the COVID-19 pandemic. Section 6 concludes the paper and gives ideas for further research.

2 Literature review

On the one hand, this paper focuses on hospital bed capacity allocation to different patient groups. Within a single hospital, non-pandemic outbreak setting, this was studied in the healthcare operations research literature before, under various circumstances using various modeling approaches, see, e.g., Seung-Chul et al. (2000), Andersen et al. (2017), or Gong et al. (2022). Also, in the non-pandemic setting, joint bed capacity reservation within a region was studied, see, e.g., Litvak et al. (2008) or Marquinez et al. (2021). During the early weeks of the COVID-19 pandemic, Kaplan (2020) was one of the first to apply a model within an individual hospital to determine the necessary COVID-19 bed capacity. The model used the Erlang loss formula to determine whether existing COVID-19 capacity is sufficient. Similar to our research, but still within the setting of an individual hospital, both Lay et al. (2023) and Ma et al. (2022) developed models, resp. a discrete-event simulation model and a dynamic programming model for allocating bed capacity to infectious patients by taking these beds out of the regular care processes. At a later point in time, this action can be reversed so that beds used for infectious care can be used for regular care again. Gökalp et al. (2023) also considered dynamic capacity management of the most crucial resources for COVID-19 care, which are ICU beds, ventilators, and normal ward beds, but in their developed Markov decision process, these resources are not (temporarily) taken away from regular care, but brought into the hospital from an external source. A similarity that Lay et al. (2023) and Gökalp et al. (2023) share with our approach, but Ma et al. (2022) does not, is that it takes some lead time to make resources available for the infectious patient group.

On the other hand, this paper focuses on assigning patients to different hospitals. During the COVID-19 pandemic, the assignment of patients to different hospitals was also studied, e.g., Sarkar et al. (2021) proposed a data-driven decision-making tool to optimize the assignment of infectious patients, Aydin and Cetinkale (2022) presented a linear programming model for patient assignment and AbdelAziz et al. (2021) a bi-objective optimization model. Unlike our direct lookahead approach, these papers did not make patient assignment decisions based on bed occupancy predictions but assigned patients myopically. Therefore, these papers also did not evaluate their patient assignment decisions in a simulation study, as we do. Sarkar et al. (2021) solved a problem instance where the number of patients to be assigned at that certain point in time follows from a compartment model, whereas Eriskin et al. (2022) and AbdelAziz et al. (2021) knew for every patient to be assigned the exact time-stamp from which on they need hospitalization. In a previous paper, Dijkstra et al. (2023), we have also developed decision rules for assigning patients to hospitals, aiming at fair balancing of the burden of COVID-19 care on the bed occupancy levels within and across regions. The dynamic fair balancing model for patient assignment within a region was a load balancing model incorporating a forecast of the bed occupancy, while across regions, it was an integer linear program taking into account forecasts of future bed surpluses or shortages. The work of Ye et al. (2022) also presented a patient assignment method based on a defined principle of fairness to balance the occupancy levels in various affected areas but did not incorporate up and down scaling of capacity.

Several papers in the literature consider both allocation of hospital resources and patient assignment. Parker et al. (2020) and Eriskin et al. (2022) proposed optimization models resulting in centralized decisions on resource re-allocations (nurses and other critical resources) as well as daily patient-hospital assignments. Parker et al. (2020) minimized the total bed shortage in a network of collaborating hospitals over a pre-specified time period via a linear programming relaxation. Here, it was assumed that the (deterministic) number of newly admitted patients is known for every day, together with a length of stay (LoS) distribution, which allows the total expected bed shortage to be calculated as a real number. Eriskin et al. (2022) developed a mixed integer linear program (MILP) that determines the periodic resource re-allocations and daily patient-hospital assignments. This MILP was extended

to a p -robust program afterward by including a stochastic bed occupancy and adding robustness constraints such that the relative regret for each scenario is kept under a threshold. Barbato et al. (2023) proposed a mathematical programming approach for rebalancing the health resources among hospitals, which included, in addition to resource re-allocations, the re-purposing of hospital wards (i.e., isolation beds and ordinary beds) and, as a last resort, the selective discharge of less severely ill patients. In comparison to our approach, the model in Barbato et al. (2023) is not evaluated in a dynamic way, and parameters, such as the number of patients to be assigned, were assumed to be known. Fattahi et al. (2023) developed a multi-stage stochastic program to make three types of decisions: redistribution of patients between hospitals; resource allocation, i.e., providing more external resources for hospitals such as creating new suitable beds; and resource relocation, that is, sharing of capacity between hospitals, such as healthcare personnel and ventilators. Their objective was to minimize patient refusals, the added external resources, and the number of both patient and resource transfers. The policy resulting from the multi-stage stochastic program prescribed decisions for a fixed-length horizon and was evaluated in a rolling horizon way via a simulation model.

Prior to the COVID-19 pandemic, a limited number of studies considered the assignment decision of infectious patients among multiple healthcare facilities in the case of a pandemic outbreak. Sun et al. (2014) was, to the best of our knowledge, the only one to develop multi-objective optimization models for both patient and resource allocation among hospitals in a healthcare network during an influenza pandemic outbreak. Opposed to our approach, where we make capacity allocation and patient assignment every day, they develop an optimization model for a fixed planning horizon (which they fix at seven days, so after every seventh day, the model is solved and prescribes allocation decisions for the upcoming seven days). Furthermore, their objectives focus on total and maximum patient traveling distance instead of capacity utilization, and their model does not include stochastic parameters, i.e., the LoS is known for every patient (and assumed to be shorter than the planning horizon), which also holds for the number of patients to be assigned (i.e., before solving the linear optimization model, the number of new patients to be assigned is known for each day of the seven-days planning horizon).

In this paper, we generalize the existing models in the literature as follows. First, we centrally allocate bed capacity to infectious patients in a region of collaborating hospitals by removing the beds from regular care room by room, and we do that prospectively by taking into account the lead time it takes to make a room available for infectious care. Previously, we only saw these types of decisions within an individual hospital (Lay et al., 2023; Ma et al., 2022). Also, we base these decisions on scenarios of bed occupancy resulting from inflow scenarios and an estimated LoS for each hospital. Second, patient assignment in times of a pandemic outbreak has been studied before, but not as an integrated decision together with scaling up or down bed capacity as proposed in this paper. Third, the degree to which we include stochasticity of the underlying bed occupancy in each hospital was not observed in previous work.

3 Model

In this section, we present collaborative bed allocation and patient assignment as a sequential decision problem and motivate our stochastic direct lookahead solution approach to that problem. We further explain our modeling and forecasting approach for bed occupancy, which will then be used to generate scenarios for the stochastic programs presented in Section 4.

3.1 Region of collaborating hospitals

We consider a region of H collaborating hospitals, indexed $h = 1, \dots, H$, that centrally decide on infectious bed capacity and assignment of infectious patients to those hospitals during each day of an infectious outbreak. This collaboration aims to ensure enough capacity for infectious patients while maintaining regular care as well as possible.

Each hospital has a dedicated ward for infectious patients with bed capacity c_h . In addition, there are n_h regular care rooms at hospital h that may be opened for infectious care. Hence, infectious bed capacity can only be made available in terms of rooms and not in terms of individual beds for health safety reasons. These rooms in hospital h are labeled $1, \dots, n_h$ and can only be opened and closed in sequence: room n can be opened only if room $n - 1$

is open, and room n can be closed only if room $n + 1$ is closed. The reason for this is that there must be a clear separation between infectious care rooms and regular care rooms to lower the risk of disease spread. Let room n in hospital h contain $b_{h,n}$ beds. In accordance with practice, we assume that it takes a fixed amount of two days to empty a regular care room and open it for infectious patients, e.g., by relocating regular patients to other regular care rooms and/or discharging regular patients. If a room allocated to infectious patients is empty, then it may be closed and will be immediately available for regular patients. Further, over the course of a day, infectious patients arrive either autonomously at individual hospitals or in the region, in which case they have to be assigned to individual hospitals.

3.2 Sequential decision-making problem

We model the capacity allocation and patient assignment as a sequential decision problem according to the unified framework proposed in Powell (2022). The *state* variables indicate the rooms that are currently open (in the morning) on day t and that are in preparation for opening on day $t + 1$ and $t + 2$ in hospital h , as well as the current number of infectious patients per hospital. The *decisions* are which rooms for each hospital h will actually be opened in hospital h on day t (which had to be prepared on day $t - 1$ to be opened on day t), which rooms will be prepared to be opened on day $t + 1$ (which had to be prepared on day $t - 1$ to open on day $t + 1$), which rooms will be prepared to be opened on day $t + 2$ and which rooms to close. Further, the hospitals decide online which regional patient to assign to which hospital. In between decision epochs, the *exogenous information* becomes known, i.e., the number of autonomously arriving infectious patients and the number of discharged patients. The *transition* to the next state happens through implementing the decisions and taking the exogenous information into account. Our *objective* is to minimize the cumulative daily costs, which consist of costs for opening and closing rooms (reflecting cleaning and set-up costs), costs for overbeds (infectious patients that cannot be accommodated), costs for the number of reserved infectious beds (beds in rooms open for infectious patients) and costs for the number of infectious beds that are in the opening process (beds in rooms from which regular patients are transferred or discharged).

3.3 Modeling bed occupancy

We assume that the number of infectious patients demanding care follows an inhomogeneous Poisson process with rate λ_τ at time τ . The assumption of Poisson arrivals was shown to be justified for arrivals to Emergency Departments (Whitt and Zhang, 2017) and is also verified for arrivals of COVID-19 patients (Dijkstra et al., 2023). A fraction f_h of the regional patients arrives autonomously at hospital h , leading to an autonomous inflow with Poisson rate $\lambda_{h,\tau} = f_h \lambda_\tau$ of infectious patients to hospital h . This autonomous inflow contains, e.g., patients who are diagnosed positive upon arrival at the hospital's emergency department. The remaining inflow of regional patients follows a Poisson process with rate $(\lambda_\tau - \sum_h \lambda_{h,\tau})$.

Following Baas et al. (2021), we model the infectious ward of each hospital h as an infinite server queue that records the number of hospitalized infectious patients in the ward. Let L_h denote the random variable of the Length of Stay (LoS) of the patients in the ward. Hence, the hospital is described by an $M_\tau/G/\infty$ queue. Following Massey and Whitt (1993, Theorem 1.2), starting from an empty system at time 0, the occupancy by infectious patients $N'_{h,\tau}$ in hospital h at time τ has a time-dependent Poisson distribution, i.e.,

$$\mathbb{P}[N'_{h,\tau} = n] = \frac{(\rho_{h,\tau})^n}{n!} e^{-\rho_{h,\tau}}, \text{ where } \rho_{h,\tau} = \mathbb{E} \left[\int_{u=\max\{0,\tau-L_h\}}^{\tau} \lambda_{h,u} du \right]. \quad (1)$$

The Poisson distribution for the number of infectious patients in the ward (1) allows us to evaluate various performance measures. Let $\mathbf{L}_{h,\tau}$ denote tuples of the attained LoSs (up to time τ) of patients residing in the ward of hospital h at time τ . The expected occupancy in the ward at time $\tau + \sigma$ given the LoSs of the residing patients at time τ is determined by the patients present at time τ who are still present at time $\tau + \sigma$ and the patients arriving between time τ

and time $\tau + \sigma$ in the $M_\tau/G/\infty$ queue that starts empty at time τ :

$$\mathbb{E}[N'_{h,\tau+\sigma} \mid \mathbf{L}_{h,\tau} = \ell] = \sum_{i=1}^{|\ell|} \frac{1 - F_h(\ell_i + \sigma)}{1 - F_h(\ell_i)} + \mathbb{E} \left[\int_{u=\max\{\tau, \tau+\sigma-L_h\}}^{\tau+\sigma} \lambda_{h,u} du \right].$$

Other measures, such as the variance and quantiles of the occupancy at time $\tau + \sigma$ may be calculated from similar principles.

3.4 Solution approach

The two time-scale sequential decision problem of opening and closing rooms on a daily basis and assigning patients to hospitals during the day corresponds to an infinite horizon MDP with large discrete state and action spaces and partially unknown dynamics. In particular, it is assumed that the regional arrival rate λ_τ and the fractions f_h of autonomous regionally arriving infectious patients to each hospital are unknown, while the distribution of L_h is known. However, we do assume that we can reasonably predict the regional arrival rate over a short time horizon of a few days based on the arrival history. Therefore, it cannot directly be solved to optimality for real-life instances. Instead, we aim to find a good approximate policy, i.e., a mapping from states to actions. An approximate policy can be constructed based on one or a mix of the four meta-policies for solving sequential decision-making problems (Powell, 2022). Since our focus is on making good decisions for the current decision epoch, we use a (stochastic) direct lookahead approximation (DLA) approach, which estimates the impact of current decisions on current and future costs. Through modeling relationships between succeeding decision epochs, we can avoid the introduction of parameters that are required for the other meta-policies. This means that we can directly apply our DLA method to make decisions without prior computation to tune parameters. In reality, when experiencing an outbreak, tuning would have to be done via simulation or at the beginning of the outbreak, making the application of policies that need tuning more difficult.

We develop two (related) stochastic DLA models for the two types of decisions (on opening and closing of rooms and assigning regional infectious patients) that we need to make. To build tractable lookahead models, we truncate the time horizon (looking at least two days ahead), we sample scenarios with respect to the exogenous information, and we apply stage aggregation, in this case working with two stages. Because we are working with discrete and constraint decisions and assume linear costs, we develop two 2-stage stochastic programs (SPs), which we apply in a rolling horizon fashion.

Before explaining the stochastic programs in more detail in Section 4, we will first explain how we model bed occupancy and how we consequently generate scenarios.

3.5 Scenario generation

The stochastic programs that we present in Section 4 use scenarios of daily occupancy and infectious patients arriving in the region to be assigned to the hospitals. These scenarios are generated based on an estimate of the parameters of the model described in Section 3.3. Following the approach developed in Baas et al. (2021) and Dijkstra et al. (2023), using historical data up to day t , we use a prediction $\hat{\lambda}_\tau$ for $\tau \in [t, t+s)$ of the regional arrival rate in the period $[t, t+s)$, an estimate \hat{f}_h of the fraction of regional patients arriving at each hospital h , and an estimate of the LoS distribution. Here, s is the prediction horizon, which is usually not more than a week. Combining $\hat{\lambda}_{h,\tau}$ for $\tau \in [t, t+s)$ with the estimated LoS distribution, a sample path $\tilde{N}_{h,t+1}, \tilde{N}_{h,t+2}, \dots, \tilde{N}_{h,t+s}$ of the occupancy at days $t+1, \dots, t+s$ by autonomously arriving infectious patients at hospital h is generated. This method was seen to yield accurate forecasts of daily bed occupancy for the COVID-19 pandemic (Baas et al., 2021; Dijkstra et al., 2023).

Following Section 3.3, letting $\partial \hat{\Lambda}_{t+u} = \int_{t+u-1}^{t+u} \hat{\lambda}_\tau d\tau$, the autonomous arrival rate of regional infectious patients to hospital h in $[t+u-1, t+u)$ can be predicted as $\partial \hat{\Lambda}_{h,t+u} = \hat{f}_h \partial \hat{\Lambda}_{t+u}$ for all $u \in \{1, \dots, s\}$. Using this predictor, scenarios A_{t+u} of daily regionally arriving infectious patients to be assigned to hospitals in $[t+u-1, t+u)$ can be determined for $u = 1, \dots, s$ by generating sample paths from the inhomogeneous Poisson process with intensity $((1 - \sum_h \hat{f}_h) \partial \hat{\Lambda}_{t+u})_{u=1}^s$.

Therefore, a scenario i is defined as an independently sampled path $((\tilde{N}_{h,t+u}^{(i)})_{h=1}^H, A_{t+u}^{(i)})_{u=1}^s$ of the occupancy and regionally arriving infectious patients.

4 Decision rules for opening rooms and patient assignment

In this section, we will explain the two stochastic programs and the resulting decision rules. The program SP1, which decides on the number of regular care rooms that are made available and planned to be made available for infectious patients in the coming days, is described in Section 4.1, while the program SP2, which allocates patients to hospitals during the current day, conditional on the previous decision to open rooms for that day made by SP1, is given in Section 4.2.

4.1 Stochastic program for room allocation

This section introduces the stochastic program SP1 for opening and closing rooms for infectious care patients. In the first stage, SP1 decides on the number of regular care rooms that are actually opened on day t and prepared to be made available on days $t + 1, t + 2$ for infectious patients. In the second stage of SP1, we assign future regional infectious patients to the hospitals, which are, hence, wait-and-see variables. The aim is to minimize a weighted sum of the number of overbeds, available and reserved regular care beds and opened and closed regular care rooms.

4.1.1 Decision variables and scenarios

The primary decision variables for SP1 are indicator variables $(z_{h,n,t+u})_{u=0}^{s-1}$, each denoting whether room n in hospital h is opened at a day $t + u$. As secondary decision variables, we determine indicator variables $(v_{h,n,t+u}^{(d)})_{u=0}^{s-1}$, each denoting whether room n in hospital h is prepared to be open on day $t + u + d$ at a day $t + u$. As it was assumed in Section 3 that a regular care room can only be opened for infectious care after 2 days, given that it is currently in use for regular care, we must have that $z_{h,n,t+u} \leq v_{h,n,t+u-1}^{(1)} \leq v_{h,n,t+u-2}^{(2)}$, i.e., rooms opened at a current day were scheduled to be opened yesterday and the day before yesterday. A third set of decision variables are non-negative integers $(x_{h,t+u}^{(i)})_{u=1}^s$, each denoting the number of regionally arriving infectious patients assigned to hospital h in a period $[t + u - 1, t + u)$, depending on scenario i . Based on these three sets of decision variables and the scenarios, we can determine the offered capacity at days $t, \dots, t + s - 1$, and corresponding occupancy at (the beginning) of the days $t + 1, \dots, t + s$. In this calculation, it is assumed that regional patients assigned to a hospital stay at the hospital for each day in the time interval $[t, t + s)$. The difference between occupancy and capacity, if it is positive, results in the number of overbeds for that day (also see Dijkstra et al. (2023)).

4.1.2 Objective

SP1 minimizes the average daily costs over the days until the end of the considered planning horizon, where the average is taken over scenarios. The costs are determined as a linear combination of the number of opened and closed rooms, regular care beds used and scheduled to be used, as well as the number of overbeds, with coefficients $\alpha, \beta, \gamma, \delta, \epsilon > 0$, respectively. Note that all costs besides those for overbeds can be incurred at the time of decision-making. The costs for overbeds can only be incurred the next day when the occupancies become known. At the end of the considered lookahead horizon, the remaining number of overbeds and opened rooms are multiplied by costs $(s - 1) \cdot \epsilon$ and $(s - 1) \cdot \gamma$, respectively, and added to the objective as a terminal cost so that the final amount of overbeds and opened rooms are counted for s days in total. The choice of the terminal cost was made to incentivize consistency in the decisions made by the stochastic program. Suppose the terminal cost term is not added. In that case, the program can choose to allow for a large number of overbeds on day $t + s$ or opened rooms on day $t + s - 1$ when running the program at day t , but when the program is actually used to decide on open rooms on day $t + s - 1$ these overbeds or used beds induce a higher cost and will be avoided as a result.

4.1.3 Implementation and reuse of first-stage decisions

The first-stage decision variables $(z_{h,n,t+u})_{u=0}^{s-1}$, $v_{h,n,t}^{(1)}$, and $v_{h,n,t}^{(2)}$ are stored after determining the optimal solution. We implement the decisions for $u = 0$ and take the decision variables $(z_{h,n,t+u})_{u=0}^{s-1}$ into account in the second stochastic program, i.e., when assigning patients to hospitals. Further, the decision variables $z_{h,n,t}$, $v_{h,n,t}^{(1)}$ and $v_{h,n,t}^{(2)}$ are used in the first stochastic program when running it to determine rooms to open on the next day.

4.1.4 Stochastic program

Table 1 describes the indices, parameters and optimization variables used in SP1, where we define $\mathcal{U}_- = \{0, \dots, s-1\}$ and $\mathcal{U}_+ = \{1, \dots, s\}$.

Table 1: Indices, parameters, and decision variables used in the stochastic program SP1 for opening and closing regular care rooms. The function $\mathbb{I}(\cdot)$ denotes the indicator function.

Symbol	Description
Indices	
$h \in \{1, \dots, H\}$	Hospital
$n \in \{1, \dots, n_h\}$	Room in hospital h
$u \in \{0, \dots, s\}$	Days ahead
$i \in \{1, \dots, I\}$	Scenario (scen.)
$d \in \{1, 2\}$	Day ahead in schedule
Parameters	
$b_{h,n} \in \mathbb{N}$	Number of beds in room n , hospital h
$\tilde{N}_{h,t+u}^{(i)} \in \mathbb{N}$	Autonomous occupancy, hospital h , on day $t+u$, $u \in \mathcal{U}_+$, scen. i
$A_{t+u}^{(i)} \in \mathbb{N}$	Regional patients arriving in $[t+u-1, t+u)$, $u \in \mathcal{U}_+$, scen. i
$z_{h,n,t-1} \in \{0, 1\}$	$\mathbb{I}(\text{room } n \text{ in hospital } h \text{ is open on day } t-1)$
$v_{h,n,t-1}^{(d)} \in \{0, 1\}$	$\mathbb{I}(\text{room } n, \text{ hospital } h \text{ scheduled open in } d \text{ days on day } t-1)$
$c_h \in \mathbb{N}$	Standard capacity infectious ward in hospital h
$\alpha, \beta, \gamma, \delta, \epsilon \in \mathbb{R}_{>0}$	Weights used in objective function
$M \in \mathbb{R}_{>0}$	Big number
Variables	
$z_{h,n,t+u} \in \{0, 1\}$	$\mathbb{I}(\text{room } n \text{ is open at hospital } h \text{ on day } t+u)$, $u \in \mathcal{U}_-$
$v_{h,n,t+u}^{(d)} \in \{0, 1\}$	$\mathbb{I}(\text{room } n, \text{ hospital } h \text{ scheduled open in } d \text{ days on day } t+u)$, $u \in \mathcal{U}_-$
$x_{h,t+u}^{(i)} \in \mathbb{N}_0$	Patients assigned to hospital h in $[t+u-1, t+u)$, $u \in \mathcal{U}_+$, scen. i
$N_{h,t+u}^{(i)} \in \mathbb{N}_0$	Total occupancy in hospital h on day $t+u$, $u \in \mathcal{U}_+$, scen. i
$o_{h,t+u}^{(i)} \in \mathbb{N}_0$	Overbeds in hospital h on day $t+u$, $u \in \mathcal{U}_+$, scen. i
$y_{h,t+u}^{(+)} \in \mathbb{N}_0$	Total infectious rooms opened in hospital h on day $t+u$, $u \in \mathcal{U}_-$
$y_{h,t+u}^{(-)} \in \mathbb{N}_0$	Total infectious rooms closed in hospital h on day $t+u$, $u \in \mathcal{U}_-$

The stochastic program SP1 for opening and closing rooms is defined below. Unless indicated otherwise, the index variables are assumed to lie in the domains indicated in Table 1.

$$\begin{aligned}
 \text{(SP1)} \quad \min \quad & \sum_{u=0}^{s-1} \sum_{h=1}^H [\alpha y_{h,t+u}^{(+)} + \beta y_{h,t+u}^{(-)} + \sum_{n=1}^{n_h} (\gamma z_{h,n,t+u} + \delta (v_{h,n,t+u}^{(1)} + v_{h,n,t+u}^{(2)})) b_{h,n}] \\
 & + \frac{1}{I} \sum_{i=1}^I \sum_{h=1}^H [(\epsilon \sum_{u=1}^s o_{h,t+u}^{(i)} + (s-1)(\epsilon o_{h,t+s}^{(i)} + \gamma z_{h,n,t+s-1})] \tag{2}
 \end{aligned}$$

s.t.

$$\sum_{h=1}^H x_{h,t+u}^{(i)} = A_{t+u}^{(i)} \quad \forall u \in \mathcal{U}_+, i, \tag{3}$$

$$N_{h,t+u}^{(i)} = \tilde{N}_{h,t+u}^{(i)} + \sum_{\ell=1}^u x_{h,t+\ell}^{(i)} \quad \forall u \in \mathcal{U}_+, h, i, \tag{4}$$

$$z_{h,n,t+u} \leq z_{h,n-1,t+u} \quad \forall u \in \mathcal{U}_-, n \geq 2, h, \tag{5}$$

$$N_{h,t+u}^{(i)} - \sum_{n=1}^{n_h} z_{h,n,t+u-1} b_{h,n} - c_h \leq o_{h,t+u}^{(i)} \quad \forall u \in \mathcal{U}_+, h, i, \tag{6}$$

$$\sum_{n=1}^{n_h} (z_{h,n,t+u} - z_{h,n,t+u-1}) = y_{h,t+u}^{(+)} - y_{h,t+u}^{(-)} \quad \forall u \in \mathcal{U}_-, h, \tag{7}$$

$$z_{h,n,t+u} \leq v_{h,n,t+u-1}^{(1)} + z_{h,n,t+u-1} \quad \forall u \in \mathcal{U}_-, h, n, \tag{8}$$

$$v_{h,n,t+u}^{(1)} \leq v_{h,n,t+u-1}^{(2)} + v_{h,n,t+u-1}^{(1)} \quad \forall u \in \mathcal{U}_-, h, n, \tag{9}$$

$$z_{h,n,t+u+1} - z_{h,n,t+u} - 1 \leq z_{h,n,t+u} - z_{h,n,t+u-1} \quad \forall u \in \mathcal{U}_-, h, n, \tag{10}$$

$$z_{h,n,t+u}, v_{h,n,t+u}^{(1)}, v_{h,n,t+u}^{(2)} \in \{0, 1\} \quad \forall u \in \mathcal{U}_-, h, n, \tag{11}$$

$$x_{h,t+u}^{(i)}, N_{h,t+u}^{(i)}, o_{h,t+u}^{(i)} \in \mathbb{N}_0 \quad \forall u \in \mathcal{U}_+, h, n, i, \tag{12}$$

$$y_{h,t+u}^{(+)}, y_{h,t+u}^{(-)} \in \mathbb{N}_0 \quad \forall u \in \mathcal{U}_-, h, n, i. \tag{13}$$

In the above problem formulation, Constraint set (3) ensures that all regionally arriving infectious patients in need of assignment are assigned to a hospital for all scenarios and considered days. Constraint set (4) determines the total occupancy at a hospital based on the assigned regional patients and the autonomous occupancy under the scenario on a given day. Constraint set (5) enforces that rooms can only be opened sequentially. Constraint set (6) determines the amount of overbeds at a hospital under a scenario on a given day. Constraint set (7) determines the number of regular care rooms opened or closed at a hospital on a given day. Constraint set (8) ensures that rooms can only be open on a given day if that room was open the day before or if the room was scheduled to be open on that day the day before. Constraint set (9) ensures that a room can only be scheduled to open tomorrow if it was already scheduled to open in two days the day before, i.e., ensures that opening a room takes two days. Constraint set (10) ensures that rooms that are closed on a given day cannot be opened the next day. Constraint sets (11)-(13) are domain constraints.

4.2 Stochastic program for patient assignment

This section introduces the stochastic program SP2 for the assignment of patients to hospitals during the current day. It takes the first stage decisions on rooms of SP1 as input and, in the first stage, decides on the assignment of the j -th regional patient of today for $j = 1, 2, \dots$ simultaneously and, in the second stage, on the assignment of future regional patients. Hence, we solve SP2 once directly after SP1 to decide on the assignment of all future regional infectious patients of the day. The aim of the program is to minimize the number of overbeds for infectious patients in the coming days $t + 1, \dots, t + s$.

4.2.1 Decision variables and scenarios

The probability $p^{(j)}$ that j regional patients need to be assigned today coming from the Poisson distribution with rate $\partial \tilde{\Lambda}_{t+1} = (1 - \sum_h \hat{f}_h) \partial \tilde{\Lambda}_{t+1}$ is given as an input parameter to SP2. These probabilities are truncated at a value J , set to the 97.5% quantile of the Poisson distribution with rate $\partial \tilde{\Lambda}_{t+1}$. Similar to the scenarios for SP1, a scenario (i, j) for SP2 is defined as an independent sample path $((\tilde{N}_{h,t+u}^{(i)})_{h=1}^H)_{u=1}^s, (A_{t+u}^{(i)})_{u=2}^s, j)$ of the autonomous occupancy at

the hospitals at days $t + 1, \dots, t + s$, regionally arriving infectious patients in need of assignment in $[t + 1, t + s)$ and where j is the number of infectious patients in need of assignment in $[t, t + 1)$.

The primary decision variables of the stochastic program for patient assignment come in the form of an indicator variable $w_{h,t+1}^{(j)}$ of the event that the j -th arriving patient in the period $[t, t + 1)$ will be assigned to hospital h . A second set of decision variables $(x_{h,t+u}^{(i,j)})_{u=2}^s$ represents the number of regionally arriving infectious patients assigned to hospital h in $[t + u - 1, t + u)$ for scenario (i, j) . Given the patient assignments and the autonomous occupancy scenario, the total occupancy can be determined. The capacity $C_{h,t+u}$ at the hospitals at time $t + u$, determined as $C_{h,t+u} = c_h + \sum_{n=1}^{n_h} z_{h,n,t+u} b_{h,n}$, is determined by the solution of SP1. The capacity over time and occupancy scenarios now determine the number of overbeds $(o_{h,t+u}^{(i,j)})_{u=1}^s$ for each scenario (i, j) , hospital h , and time $t + u$.

4.2.2 Objective

The stochastic program minimizes the average number of overbeds over the days $t + 1, \dots, t + s$ that are related to decisions made on the days $t, \dots, t + s - 1$, where the average is taken over the chosen scenarios.

4.2.3 Implementation of first-stage decisions

The first-stage stage decision variables $w_{h,t}^{(j)}$ are stored after solving SP2 to optimality. When new regional infectious patients arrive throughout the day, we assign them to hospitals according to the values of those first-stage variables. If the actual number of patients that arrive during a given day exceeds the upper bound J , the choice is made to run the stochastic program again with limit $2 \cdot J$ setting $p^{(j)} = 0$ for $j \leq J$, in order to obtain decision variables for a larger amount of arriving patients.

4.2.4 Stochastic program

Table 2 describes the indices, parameters, and optimization variables used in the stochastic program SP2 for patient assignment, where $\mathcal{U}_{\geq 2} = \{2, \dots, s\}$.

Table 2: Indices, parameters, and decision variables used in the stochastic program SP2 for patient assignment.

Symbol	Description
Indices	
$j \in \{0, \dots, J\}$	Index of infectious patient to be assigned on day t
$h \in \{1, \dots, H\}$	Hospital
$u \in \{0, \dots, s\}$	Days ahead
$i \in \{1, \dots, I\}$	Scenario
Parameters	
$C_{h,t+u} \in \mathbb{N}$	Number of infectious beds in hospital h on day $t + u$, $u \in \mathcal{U}_-$
$p^{(j)} \in [0, 1]$	Probability that j regional infectious patients arrive today
$\tilde{N}_{h,t+u}^{(i)} \in \mathbb{N}$	Autonomous occupancy in hospital h , at time $t + u$, $u \in \mathcal{U}_+$, scenario i
$A_{t+u}^{(i)} \in \mathbb{N}$	Patients in need of assignment in $[t + u - 1, t + u)$, $u \in \mathcal{U}_+$, scenario i
Variables	
$w_{h,t}^{(j)} \in \{0, 1\}$	Indicator that the j -th patient at day t is assigned to hospital h
$x_{h,t+u}^{(i,j)} \in \mathbb{N}_0$	Patients assigned to hospital h in $[t + u - 1, t + u)$, scenario (i, j) , $u \in \mathcal{U}_{\geq 2}$
$N_{h,t+u}^{(i,j)} \in \mathbb{N}_0$	Infectious occupancy in hospital h on day $t + u$, scenario (i, j) , $u \in \mathcal{U}_+$
$o_{h,t+u}^{(i,j)} \in \mathbb{N}_0$	Overbeds in hospital h on day $t + u$ under scenario (i, j) , $u \in \mathcal{U}_+$

The stochastic program for patient assignment is defined in SP2. Again, unless indicated otherwise, the index variables are assumed to lie in the domains indicated in Table 2.

$$(SP2) \quad \min \frac{1}{I} \sum_{u=1}^s \sum_{h=1}^H \sum_{i=1}^I \sum_{j=0}^J p^{(j)} o_{h,t+u}^{(i,j)} \quad (14)$$

s.t.

$$\sum_{h=1}^H w_{h,t}^{(j)} = 1 \quad \forall j, \quad (15)$$

$$\sum_{h=1}^H x_{h,t+u}^{(i,j)} = A_{t+u}^{(i)} \quad \forall u \in \mathcal{U}_{\geq 2}, i, j, \quad (16)$$

$$N_{h,t+u}^{(i,j)} = \tilde{N}_{h,t+u}^{(i)} + \sum_{\ell=2}^u x_{h,t+\ell}^{(i,j)} + \sum_{k=1}^j w_{h,k,t}^{(j)} \quad \forall u \in \mathcal{U}_+, h, i, j, \quad (17)$$

$$N_{h,t+u}^{(i,j)} - C_{h,t+u-1} \leq o_{h,t+u}^{(i,j)} \quad \forall u \in \mathcal{U}_+, h, i, j, \quad (18)$$

$$w_{h,t}^{(j)} \in \{0, 1\} \quad \forall h, j, \quad (19)$$

$$N_{h,t+u}^{(i,j)}, o_{h,t+u}^{(i,j)} \in \mathbb{N}_0 \quad \forall u \in \mathcal{U}_+, h, i, j, \quad (20)$$

$$x_{h,t+u}^{(i,j)} \in \mathbb{N}_0 \quad \forall u \in \mathcal{U}_{\geq 2}, h, i, j. \quad (21)$$

In the above, Constraint set (15) ensures that all regionally arriving infectious patients in need of assignment today are assigned to a hospital. Constraint set (16) ensures that all regionally arriving infectious patients in need of assignment are assigned to a hospital for all scenarios and future days. Constraint set (17) determines the total occupancy at a hospital based on the assigned regional patients and the autonomous occupancy under the scenario on a given day. Constraint set (18) determines the amount of overbeds at a hospital under a scenario on a given day. Constraint sets (19)-(21) are domain constraints.

4.2.5 Relationship between the two stochastic programs

Every day, we run SP2 directly after SP1 without knowledge of the actual number of arriving regional infectious patients that day. This construction is especially suited to be used in a simulation to test and evaluate the policies resulting from the two SPs. In reality, one could instead run an SP every time a new regional infectious patient arrives and decide on the placement for this patient alone. This would also allow us to use the information on already-arrived autonomous patients until then. However, it would also substantially increase computation time.

Note that in the second stage of SP1, we decide on the assignment of patients to hospitals for today and the upcoming days. In SP2, the assignment decision for today's patient is moved to the first stage, while the second stage still considers the assignment of patients in the upcoming days. Therefore, the quality of the room decisions made by SP1 is dependent on using SP2 for the patient assignment decisions afterward. However, given any room decisions for the next few days, SP2 can be applied on its own independently of SP1 to provide suggestions for patient-hospital assignments.

5 Case study: collaboration of hospitals in Acute Zorg Euregio during the COVID-19 pandemic

This section evaluates the performance of using SP1 and SP2 to determine the number of rooms to open, reserve to open, and assign patients. The performance of this decision rule, denoted SP, is compared to the performance of two myopic decision rules in a simulation study. The section ends with a sensitivity analysis.

The simulation study considers occupancy and regional arrivals by COVID-19 patients in the period 1 October 2021 until 31 December 2021 (91 days) at three hospitals in ROAZ region Acute Zorg Euregio (Euregio): Medisch Spectrum Twente (MST), Ziekenhuisgroep Twente (ZGT) and Streekiekenhuis Koningin Beatrix (SKB). During COVID-19 outbreaks in the Netherlands, the ROAZ consultation bodies controlled the assignment of infectious patients to hospitals within a region. At each decision epoch, the method described in Section 3.5 generates scenarios of daily occupancy and arrivals in the period $[t, t + s)$, used by the considered decision rules. As in Dijkstra et al. (2023), the decision epochs are set to 10 AM for each day, which was determined as the time that hospitals report their occupancy to the

region. At the start of each simulation run, there are no COVID-19 patients at the hospitals. The standard capacity and the capacity of regular care rooms that can be opened for COVID-19 care per hospital can be found in Table 3. In the following paragraphs, we will first describe each parameter used in the simulation study, after which we will describe the (common) estimation method for this parameter used by the decision rules.

Table 3: Standard (COVID-19) capacity and regular care room capacity per hospital.

Hospital h	Standard capacity c_h	Reg. capacity $b_{h,n}$				
MST	20	4	12	6	2	4
ZGT	8	8	5	5	6	-
SKB	8	5	7	-	-	-

Following Dijkstra et al. (2023), an exponentially weighted moving average (weight 0.1) of the sum of the historic daily COVID-19 arrivals to the three hospitals during the considered period is used to determine the Poisson arrival rate λ_τ of regional infected patients, which is used to sample regional arrivals of COVID-19 patients in the simulation study. Figure 1 shows the resulting arrival intensity. At each decision epoch t , the decision rules use a 5-parameter Richards' curve predictor as described in (Baas et al., 2021) and improved in (Dijkstra et al., 2023) to determine $\hat{\lambda}_\tau$ for $\tau \in [t, t + s)$. The choice was made to use the Richards' curve predictor and not the daily infections predictor proposed in (Dijkstra et al., 2023) as there is no guarantee that in future pandemics, carefully recorded daily infection data will always be publicly available. To stabilize arrival rate predictions at the start of the simulation study, the prediction method for the regional arrival rate uses 2 months of historical daily arrival data prior to the start of the simulation study.




Figure 1: Daily regional arrival rates (amount of patients per day) of regional COVID-19 patients in the region for the simulation study of Section 5 (1/10/2021 - 31/12/2021).

The fractions f_h of COVID-19 patients arriving autonomously at the hospitals are not available from the data and are set to 15%, 4%, and 4% for MST, ZGT, and SKB, respectively. These values were determined by pre-computation as they yielded a similar probability of overbeds (around 1%) for all hospitals, given that the hospitals do not scale up (i.e., when only the standard capacity, given in Table 3, is utilized). At each decision epoch t , letting $t = 0$ denote the start of the simulation study, the decision rules use an estimate $\hat{f}_{h,t}$ of f_h , determined using the formula

$$\hat{f}_{h,t} = (\hat{f}_{h,0} + \sum_{u=1}^t \tilde{A}_{h,u}) / (1 + \sum_{u=1}^t [A_u + \sum_{h=1}^H \tilde{A}_{h,u}]), \quad (22)$$

where $\tilde{A}_{h,u}$ is the number of patients arriving autonomously to hospital h in $[u - 1, u)$ of the simulation and $\hat{f}_{h,0}$ is a prior fraction to stabilize the estimation procedure at the start of the simulation period, set to 20%, 5%, and 5% for MST, ZGT, and SKB. Note that $\hat{f}_{h,t}$ provides a consistent estimator of f_h under the model in Section 3.3. Hence,

although the (arbitrary) choice 20%, 5%, and 5% for the prior fractions is different from the actual values 15%, 4%, and 4%, the fractions will, in practice, quickly converge to the truth.

A shared LoS distribution in the ward is determined based on pooled data collected from the hospital data warehouses of ZGT and SKB, where the weight of the LoS data from ZGT is set to 2.0 as, after consulting with representatives of MST, it was determined to have the best correspondence with the LoS distribution at MST during this period, where this data was missing. The LoS distribution used in the simulation study is determined using the Kaplan-Meier estimator in order to account for right censoring due to patient transfers to other hospitals and patients still residing at the moment of estimation. The LoS distribution is assumed to be known to the decision rules.

We now describe the decision rules considered in the simulation study:

- **Individual Hospitals (IH):** This decision rule mimics the situation in which all regional patients arrive autonomously to the hospitals and hospitals open rooms according to individual forecasts of the infectious occupancy in the coming days. The fractions of regional patients arriving autonomously at the hospitals are set to 48%, 32%, and 20% for MST, ZGT, and SKB, respectively, in agreement with the ratio between maximum capacity available for infectious patients for these hospitals, as given in Table 3.

A forecast of the occupancy for the coming days is generated by taking the maximum of the current occupancy at the hospital and a 90% quantile over scenarios of the occupancy at the hospital on day $t + 2$, as rooms can only be opened in 2 days. The number of rooms to be opened at the hospital is determined as the smallest number of rooms such that the capacity exceeds this occupancy forecast when these rooms are opened. The number of rooms opened and scheduled to open is set to the minimum of the previous decision made by the decision rule and the currently determined amount of rooms to be opened. In order to dampen the fluctuation in decisions over time, the number of opened rooms cannot be decreased by IH when the occupancy forecast for the hospital plus a safety margin exceeds the current hospital capacity. This safety margin is set to 3 beds for MST and 2 beds for ZGT and SKB.

- **Pandemic Unit (PU):** This decision rule mimics the situation in which all patients are first sent to the hospital with the largest capacity in the region, which was MST in this case (see Table 3). When assigning patients, PU sends all patients to MST until the hospital is at capacity minus a safety margin. The safety margin equals the occupancy forecast (maximum of occupancy today and 90% quantile of occupancy in 2 days) coming from autonomous arrivals. The remainder of the patients to be assigned are assigned at random with probabilities 61% vs. 39% for ZGT and SKB, respectively, in agreement with the ratio between maximum capacity available for infectious patients for these hospitals, as given in Table 3.

When determining which rooms to open or close, all arrivals are assumed to go to MST when this hospital has not opened all its rooms. Otherwise, the predicted arrival rate of patients to be assigned to MST is set such that MST has no capacity left when the expected amount of arriving patients stay at MST during the interval $[t, t + s)$. The occupancy forecast, determined in the same manner as for IH, is included in the calculation of the remaining capacity. It is assumed that the remainder of the regional arrival rate of patients in need of assignment is divided over ZGT and SKB in the ratio 61% vs. 39%. As in IH, the number of rooms to be opened at each hospital is determined as the smallest number of rooms such that the capacity exceeds the corresponding occupancy forecast when these rooms are opened. The number of rooms opened and scheduled to open is set to the minimum of the previous choice and the currently determined amount of rooms to be opened. For PU, no rooms at ZGT and SKB can be opened or scheduled to be open if MST has not opened all its rooms. The safety margins for scaling down, used in the same manner as for IH, are set to 3 beds for MST and 2 beds for ZGT and SKB.

- Stochastic Program (SP):** Three cost settings for SP are considered, where the cost vector corresponding to the first, second, and third row of Table 4 is denoted by SP-O, SP-B, and SP-R, respectively. First, the cost parameters of SP-O are chosen such that it outperforms PU (in terms of all cost components given in (2)) and has a low amount of regional overbeds. Second, the cost parameters of SP-B are chosen such that it outperforms IH and allocates a low amount of regular beds to infectious care. Third, the cost parameters of SP-R are chosen such that it outperforms PU and opens/closes a low amount of regular rooms for infectious care. As we will mainly focus on opening or closing a room in the evaluation, we set $\alpha = \beta$ and furthermore set $\gamma = 1$ to fix the scale. The parameter choice is found by hand and was completed when changing any of the parameters significantly either led to the regional overbeds increasing or a worse performance in the other KPIs than the compared decision rule (IH or PU) in any of the KPIs. In practice, the cost parameters can be based on expert opinion, reflect actual real-life costs, or can be tuned based on scenario analyses.

Table 4: Scenarios for cost vectors considered for SP.

Setting	α	β	γ	δ	ϵ
Overbeds SP (SP-O)	15	15	1	1.5	40
Reg. beds SP (SP-B)	6	6	1	1	13
Open/close rooms (SP-R)	60	60	1	1.25	25

For each decision rule, 250 independent simulation runs of 91 days are performed, and 100 scenarios of regional arrivals and daily occupancy are used by the decision rules for each day in a simulation run. On average, a simulation run for the IH and PU decision rules took around 550 seconds, while a simulation run with SP took around 700-1300 seconds.

5.1 Numerical comparison of decision rules

Table 5 presents the main results of the simulation study. Decision rules are evaluated based on several KPIs, namely, the number of overbeds, underbeds, regular care beds used, regular care beds reserved, and the number of rooms added/removed on average per day during the simulation period. The underbeds KPI equals the number of regular care beds opened for infectious care that is unoccupied on a given day, averaged over days in a simulation run. This KPI is not directly minimized in the objective function of SP1. To quantify the significance of the differences seen in the table, 95% confidence intervals (CI) for the means, based on Student's t-distribution, are also shown in the table.

The first column of Table 5, corresponding to IH, shows a similar number of overbeds for MST, ZGT, and SKB, while the number of underbeds is slightly higher for MST and ZGT than for SKB. The ordering in the number of beds reserved and rooms added/removed roughly follows the ordering in the total capacities for the hospitals, while ZGT has the highest need for extra capacity under IH, with around 14 extra beds used on average per day.

In comparison to IH, the second column of Table 5, corresponding to PU, shows that the match between occupancy and capacity is worse for MST, as can be seen from the higher amount of underbeds per day, while this match is better for ZGT and SKB. For ZGT and SKB, most KPIs decrease, often significantly (as the CIs do not overlap), when comparing IH to PU. For MST, the number of beds reserved and rooms added/removed decreases. Looking at the regional level, it is seen that the match between occupancy and capacity becomes better than under IH, as both the average number of over and underbeds decrease at the cost of a higher number of regular care beds used over time. In agreement with the KPIs for MST, the number of beds reserved and rooms added/removed also decreases, indicating less fluctuation in the decisions for scaling up/down COVID-19 capacity over time. The reason for this can be that under PU, there is less variability in the occupancy at ZGT and SKB, as most COVID-19 patients will first go to MST. Due to this, the forecast quality of future occupancy improves and, with it, the performance of PU.

Table 5: Average KPIs for each hospital and the total region, along with 95% confidence interval, for the decision rules compared in the simulation study. The average costs for IH and PU shown in columns 4, 5, and 6 are determined using the same cost parameters as those used for the SP decision rules.

Hospital	KPI	IH	PU	SP-O	SP-B	SP-R
MST	Overbeds	0.230 ± 0.025	0.102 ± 0.012	0.140 ± 0.018	0.249 ± 0.022	0.167 ± 0.019
	Underbeds	4.654 ± 0.135	6.570 ± 0.073	3.729 ± 0.104	2.747 ± 0.083	3.586 ± 0.110
	Reg. beds used	13.056 ± 0.178	22.903 ± 0.115	12.187 ± 0.188	10.946 ± 0.179	11.588 ± 0.206
	Beds reserved	1.163 ± 0.034	0.835 ± 0.024	0.828 ± 0.024	0.926 ± 0.026	0.711 ± 0.022
	Rooms added/removed	0.119 ± 0.004	0.078 ± 0.003	0.085 ± 0.003	0.100 ± 0.003	0.057 ± 0.001
ZGT	Overbeds	0.269 ± 0.033	0.181 ± 0.025	0.096 ± 0.012	0.165 ± 0.013	0.106 ± 0.011
	Underbeds	4.779 ± 0.093	3.359 ± 0.084	3.397 ± 0.079	2.573 ± 0.060	3.660 ± 0.077
	Reg. beds used	14.309 ± 0.128	9.126 ± 0.118	12.310 ± 0.171	11.141 ± 0.168	12.877 ± 0.164
	Beds reserved	1.001 ± 0.029	0.879 ± 0.026	0.846 ± 0.024	0.906 ± 0.026	0.779 ± 0.022
	Rooms added/removed	0.085 ± 0.003	0.089 ± 0.003	0.084 ± 0.003	0.088 ± 0.003	0.047 ± 0.001
SKB	Overbeds	0.207 ± 0.029	0.268 ± 0.036	0.090 ± 0.010	0.159 ± 0.014	0.110 ± 0.012
	Underbeds	3.388 ± 0.107	1.986 ± 0.078	2.509 ± 0.071	1.906 ± 0.053	2.426 ± 0.068
	Reg. beds used	7.174 ± 0.101	4.836 ± 0.069	7.221 ± 0.144	6.816 ± 0.130	7.381 ± 0.135
	Beds reserved	0.679 ± 0.023	0.437 ± 0.019	0.469 ± 0.018	0.486 ± 0.019	0.417 ± 0.016
	Rooms added/removed	0.060 ± 0.003	0.040 ± 0.002	0.042 ± 0.002	0.042 ± 0.002	0.032 ± 0.001
Region	Overbeds	0.706 ± 0.051	0.552 ± 0.046	0.326 ± 0.035	0.573 ± 0.041	0.383 ± 0.037
	Underbeds	12.821 ± 0.197	11.915 ± 0.137	9.635 ± 0.181	7.226 ± 0.146	9.672 ± 0.193
	Reg. beds used	34.539 ± 0.251	36.865 ± 0.208	31.717 ± 0.240	28.902 ± 0.242	31.846 ± 0.233
	Beds reserved	2.843 ± 0.052	2.151 ± 0.039	2.143 ± 0.042	2.318 ± 0.043	1.907 ± 0.034
	Rooms added/removed	0.263 ± 0.006	0.207 ± 0.005	0.211 ± 0.004	0.230 ± 0.005	0.136 ± 0.001
	Avg. cost IH	-	-	70.990 ± 2.077	48.139 ± 0.726	71.547 ± 1.329
	Avg. cost PU	-	-	65.291 ± 1.870	47.439 ± 0.654	65.797 ± 1.202
	Avg. cost SP	-	-	51.129 ± 1.407	40.055 ± 0.598	51.959 ± 0.932
	Forecast cost SP (hor. 1)	-	-	48.990 ± 1.238	38.469 ± 0.508	50.248 ± 0.784
	Forecast cost SP (hor. 2)	-	-	50.908 ± 1.325	39.255 ± 0.501	51.855 ± 0.831
	Forecast cost SP (hor. 3)	-	-	58.380 ± 1.797	41.768 ± 0.630	59.388 ± 1.127

The third column of Table 5, corresponding to SP-O, shows that in comparison to PU, the number of overbeds at ZGT and SKB is lower, while the amount of overbeds for MST is higher. The number of regular care beds used and rooms added/removed is more similar for MST and ZGT in comparison to PU. SP-O outperforms IH for all KPIs and hospitals. The ordering seen in the over and underbeds corresponds to the order of the sizes of the hospitals. The similarity of the amount of opened/closed beds for MST and ZGT can be explained by the similarity in total regular care capacity available for COVID-19 for these hospitals, 28 in comparison to 24 respectively (see Table 3).

The bottom rows of the third column Table 5 show the simulated average daily cost for IH, PU, and SP-O under the SP-O cost vector, as well as the forecast cost for SP-O (1, 2, and 3 days ahead) using the scenarios generated under SP-O. The daily costs are realizations of the cost during one simulation run, while the forecasts are made by averaging several scenarios used by SP-O during the simulation run. It is seen that IH has the highest cost, then PU and the smallest cost is attained for SP-O, which could be a result of PU and SP-O being geared towards the reduction of regional overbeds, while IH is geared towards the reduction of overbeds at each hospital separately. It is seen that the forecast cost increases with the forecast horizon, overshooting the realized cost around 7 for three days ahead. This could be due to the assumption that patients assigned to the hospitals stay there during the period $[t, t + s)$, while in practice, the LoS might be shorter.

The fourth column of Table 5, corresponding to SP-B, shows that in comparison to IH, the amount of overbeds is higher at MST, while this amount is lower at ZGT and SKB. Most other KPIs significantly decrease in comparison to IH. In comparison to PU, the most notable change is the number of regular beds used at MST, which decreases by roughly 50%, while the number of regular beds used at the other hospitals increases. The regional KPIs show that the amount of regional overbeds is close to that of IH, while all other KPIs are significantly lower. In comparison to PU, the number of underbeds and regular beds used is lower. As the regular beds used are the main component of the cost and this KPI involves less uncertainty, the cost forecast stays closer to the realization for SP-B than for SP-O.

The fifth column of Table 5, corresponding to SP-R, shows that in comparison to PU, the number of overbeds increases at MST while all other KPIs decrease. For ZGT and SKB, the number of underbeds and regular beds used per day increases in comparison to PU, while overbeds, reserved beds, and opened/closed rooms decrease. Regionally, SP-R outperforms PU and, hence also, IH in all KPIs. As the cost given to overbeds is higher for SP-R than for SP-B, it is again seen that the forecast becomes higher than the realization for longer horizons.

The results show that SP is a flexible decision rule, where the cost parameters can be tuned to make a trade-off between relevant KPIs, while SP still outperforms rule-of-thumb heuristics.

5.2 Sensitivity analyses

In this section, we analyze the sensitivity of the decision rules IH, PU, and SP to several parameters. In order to analyze SP, we chose to analyze only SP-O as the quality of the solution depends most on the scenarios given as input to SP. First, we analyze the performance of the IH and PU decision rules when changing the quantile used for determining the occupancy forecast to investigate the robustness of the cost vectors in Table 4 to the chosen quantile. Second, we analyze the sensitivity of SP-O to the lookahead horizon, the number of scenarios, as well as the arrival rate predictor used in generating the scenarios. Third, we compare SP-O with deterministic versions of SP-O based on the median and a quantile over scenarios to investigate whether results similar to those for SP can be reached with a less computationally intensive approach.

5.2.1 Sensitivity of IH and PU to choice of quantile

Table 6 shows the results for IH and PU when the 80%, 85%, 90%, and 95% quantile is taken over scenarios to determine the forecasts. The number of overbeds decreases for IH when the quantile over scenarios increases. For IH, the number of underbeds and regular beds used increases, while the number of beds reserved and rooms added/removed slightly decreases. PU seems less sensitive to the choice of quantile, with most differences not being significant. The main difference between IH and PU is that the number of beds reserved increases slightly for PU in the value of the

quantile, while it decreases for IH. Tables 5 and 6 show that all SP decision rules still outperform the respective heuristic (IH or PU) even though they were designed to outperform it for other values of the quantile.

Table 6: Regional KPIs for decision rules IH and PU using the 80, 85, 90, 95% quantile over scenarios to forecast the occupancy in 2 days.

DR	KPI	80%	85%	90%	95%
IH	Overbeds	0.971 ± 0.053	0.831 ± 0.052	0.706 ± 0.051	0.575 ± 0.051
	Underbeds	10.398 ± 0.168	11.489 ± 0.178	12.821 ± 0.197	14.690 ± 0.220
	Reg. beds used	31.954 ± 0.247	33.134 ± 0.248	34.539 ± 0.251	36.476 ± 0.240
	Beds reserved	2.908 ± 0.052	2.875 ± 0.054	2.843 ± 0.052	2.841 ± 0.054
	Rooms add./rem.	0.272 ± 0.006	0.267 ± 0.007	0.263 ± 0.006	0.256 ± 0.006
PU	Overbeds	0.564 ± 0.045	0.559 ± 0.045	0.552 ± 0.046	0.526 ± 0.044
	Underbeds	11.802 ± 0.132	11.841 ± 0.134	11.915 ± 0.137	12.126 ± 0.139
	Reg. beds used	36.743 ± 0.207	36.784 ± 0.207	36.865 ± 0.208	37.091 ± 0.208
	Beds reserved	2.128 ± 0.039	2.140 ± 0.039	2.151 ± 0.039	2.158 ± 0.039
	Rooms add./rem.	0.206 ± 0.005	0.207 ± 0.005	0.207 ± 0.005	0.204 ± 0.005

5.2.2 Comparison of lookahead horizons, number of scenarios, and arrival predictors

Table 7 shows the KPIs for SP-O, as well as the resulting KPIs for using SP with a different arrival rate predictor and different lookahead horizons. The second column shows the KPIs when taking the 90% upper bound of the confidence interval for the predicted arrival rate, following from ordinary least squares, to generate scenarios of regional arrivals and occupancy in the stochastic program. For this arrival rate predictor, the table shows that SP-O results in a lower amount of overbeds, beds reserved, and rooms added/removed, while it results in a higher amount of regular beds used, which makes sense as a higher arrival rate leads to a higher forecast occupancy and required capacity. Table 7 and 5 show that SP-O with the UB arrival rate predictor still outperforms PU, while the average cost for SP-O under the UB arrival rate predictor is lower, although not significantly lower based on the CIs, indicating that this arrival rate could be used as an alternative to the arrival rate used in Table 5. The lower cost under the UB arrival rate could possibly be explained by the negative bias for the Richards' curve predictor observed in Baas et al. (2021) and Dijkstra et al. (2023). The table shows that the error between realized and forecast cost grows with the horizon of the scenarios, which is expected, as the predictor used for generating arrival scenarios lies above that used for SP-O.

The last two columns of Table 7 show results for different lookahead horizons $s = 4, 3$. The table shows that, on average, the realized daily average amount of overbeds increases, and the realized average daily amount of underbeds and regular beds used decreases when decreasing the horizon for the scenarios. For shorter scenario horizons, the stochastic program often underestimates the total cost, where for a scenario horizon of 3 days, the forecast cost stays below the realized cost for all forecast horizons. As the difference in average cost between SP-O with a horizon of 4 or 5 days is not significant, it is concluded that taking a scenario horizon of 4 days would also have been sufficient. Comparing the average costs in columns 1 and 2 and columns 1, 3, and 4 in Table 7, it is concluded that the program is more sensitive to the scenario horizon than the arrival rate predictor used.

Table 7: Regional KPIs for SP-O for lookahead horizons 5, 4, and 3 days and when taking the 90% CI upper bound for the predicted arrival intensity (UB arrival rate).

KPI	SP-O	UB arrival rate	Horizon 4 days	Horizon 3 days
Overbeds	0.326 ± 0.035	0.253 ± 0.036	0.380 ± 0.038	0.559 ± 0.041
Underbeds	9.635 ± 0.181	11.212 ± 0.210	8.832 ± 0.172	7.610 ± 0.148
Reg. beds used	31.717 ± 0.240	33.464 ± 0.240	30.814 ± 0.239	29.183 ± 0.242
Beds reserved	2.143 ± 0.042	2.023 ± 0.040	2.355 ± 0.043	2.198 ± 0.039
Rooms added/removed	0.211 ± 0.004	0.187 ± 0.004	0.228 ± 0.005	0.218 ± 0.004
Avg. cost SP	51.129 ± 1.407	49.429 ± 1.459	52.956 ± 1.533	58.120 ± 1.676
Forecast cost SP (hor. 1)	48.990 ± 1.238	50.002 ± 1.423	50.323 ± 1.293	52.878 ± 1.328
Forecast cost SP (hor. 2)	50.908 ± 1.325	55.907 ± 1.644	50.311 ± 1.356	51.105 ± 1.307
Forecast cost SP (hor. 3)	58.380 ± 1.797	68.829 ± 2.325	57.086 ± 1.867	56.332 ± 1.738

Another setting of SP is the number of scenarios used. Figure 2 shows the average cost of SP-O, as well as 95% CI for the expected cost vs. the number of scenarios. It is seen that the average cost stabilizes roughly after using 50 scenarios, indicating that taking 50 scenarios would have also been sufficient for SP-O.




Figure 2: Average total cost (and 95% CI) for SP-O vs. the number of included scenarios.

5.2.3 Comparison to deterministic programs

Table 8 shows the KPIs for SP-O, as well as the resulting KPIs when using deterministic versions of SP-O. The second and third columns of Table 8 show the KPIs for SP-O when taking the median and 85% quantile over scenarios of arrivals and occupancy and using this as the only scenario. The 85% quantile is chosen as it resulted in the smallest average cost when compared to deterministic programs with quantiles 60%, 70%, 80%, 85%, 90%, and 95%.

The results show that taking the median over scenarios leads to a more optimistic forecast than under SP-O, resulting in a realized average cost that is more than 50% higher than the 1-day ahead forecast. The long-term forecasts show an even more extreme behavior, where the realized cost is more than three times higher. Regionally, this leads to a higher amount of realized overbeds and rooms added/removed, while there is a lower amount of regular beds used compared to SP-O.

The program that results when taking the 85% quantiles shows another extreme, where the forecast cost lies above the realized cost, being more than three times higher than the realized cost for a forecast horizon of four days. For this

program, the amount of rooms added/removed is slightly lower than those under SP-O, while the other KPIs are higher. The realized average cost is around 12% higher.

Even though the deterministic decision rule with a well-chosen quantile percentage comes close to the performance of SP-O, however this can only be achieved after tuning the quantile parameter based on the outcomes of the simulation study. In the real world, the cost parameters of SP might correspond to actual costs and can be determined without using simulations, while this is not the case for the choice of quantile over scenarios. Hence, for such situations, we likely overestimate the performance of the deterministic decision rule in our results.

Table 8: Regional KPIs for SP-O, as well as the deterministic version of the stochastic programs, when taking the median and 85% quantile over the arrivals and occupancy scenarios for each day.

KPI	SP-O	Median Scen.	Quantile Scen.
Overbeds	0.326 ± 0.035	2.455 ± 0.061	0.419 ± 0.047
Underbeds	9.635 ± 0.181	4.084 ± 0.083	12.029 ± 0.199
Reg. beds used	31.717 ± 0.240	23.744 ± 0.245	34.079 ± 0.240
Beds reserved	2.143 ± 0.042	2.209 ± 0.034	2.340 ± 0.044
Rooms added/removed	0.211 ± 0.004	0.239 ± 0.004	0.196 ± 0.005
Avg. cost SP	51.129 ± 1.407	128.850 ± 2.542	57.299 ± 1.883
Forecast cost SP (hor. 1)	48.990 ± 1.238	85.416 ± 1.812	72.173 ± 2.988
Forecast cost SP (hor. 2)	50.908 ± 1.325	56.003 ± 1.257	105.605 ± 4.429
Forecast cost SP (hor. 3)	58.380 ± 1.797	41.251 ± 1.445	156.446 ± 5.864

Comparing the results shown in Table 8 with those shown in Table 7, the results for SP-O change more drastically when making the program deterministic, in comparison to when changing certain settings of the stochastic program such as the arrival rate predictor, horizon and number of scenarios.

6 Discussion

In this paper, we developed decision rules for a region with collaborating hospitals to allocate hospital bed capacity to regular or infectious care and to assign infectious patients to hospitals in a region during an infectious outbreak, aiming to serve all infectious demand while maintaining regular care. The presented decision rules result from a stochastic lookahead approach by solving two stochastic programs with scenarios in sequence on each day.

In our numerical experiments, we evaluated the performance of the developed decision rules considering the use case of COVID-19 in a region with three hospitals in the Netherlands at the end of 2021. The results indicate that our approach considerably outperforms simpler decision rules for non-collaborating hospitals and for having a pandemic unit, showing a clear benefit of using our solution approach and of hospital collaboration in a region during infectious outbreaks. The comparison of one of these heuristics, the pandemic unit, with regional collaboration led to the conclusion for the considered region that a strategy involving regional collaboration with dynamic capacity allocation and patient assignment to all hospitals would be preferred over a pandemic unit, indicating the practical relevance of this paper, see, e.g., NOS nieuws (2023). The numerical results also demonstrate that our stochastic lookahead approach is superior to deterministic lookaheads. For our use case, we show that only a few scenarios in the stochastic programs, i.e., approximately 50-100, and a short lookahead horizon, i.e., 4-5 days, are needed, ensuring a fast runtime of the stochastic programs. Further, a computational advantage of our decision rules is that they do not rely on tunable parameters compared to simpler ones. Therefore, we can directly apply our rules instead of having to tune parameters in a simulation upfront.

Our generic solution approach allows the application of decision rules during future infectious outbreaks to guide bed capacity and patient assignment. For that, the weights for the objective function must be determined by the decision-makers, and scenarios for the stochastic programs must be generated. To create those input scenarios, any suitable method to predict infectious arrival rates can be applied together with our method or alternative methods to produce bed occupancy scenarios. Further, the stochastic programs can easily be adjusted to model different numbers of days it takes to empty a room. The constraints encoding the order of opening and closing rooms in a hospital can also easily be removed or adjusted.

For future research, solution quality could potentially be increased further by revisiting some of our simplifying assumptions. First, one could investigate including a stochastic number of days to empty a room, reconsider the assumption that assigned patients stay until the end of the lookahead horizon, that the LoS distribution is known, or that the fractions of autonomously arriving patients are time-invariant. Second, one could also possibly consider multiple patient types. Third, multi-stage stochastic programs can be employed instead of two-stage programs. Fourth, one could solve a stochastic program for each arriving regional infectious patient instead of one single program at the beginning of the day. This way, one can include the realizations of the autonomous inflow of patients until the arrival of that patient. However, while it might be possible to work with multi-stage programs and stochastic programs for individual patients at decision times, evaluating those programs via simulation could become intractable.

In this paper, we considered three choices of the weights of the objective function of the stochastic program, chosen such that one of the considered heuristics is outperformed and the stochastic program results in either a low amount of overbeds, regular beds used, or rooms opened and/or closed. These parameters were found by hand, and in future research, it could be interesting to consider optimizing the weight parameters such that the stochastic program satisfies the aforementioned properties. In practice, at the start of a pandemic, such an optimization method can optimize the weights based on scenario analysis, or the weights might first reflect real-life costs, while the weight optimization method can force the behavior to have desired properties if, at interim analyses, the previously set weights seem insufficient to do so.

To support even more planning decisions during infectious outbreaks, ICU capacity and patient transfers between the ward and ICU, as well as between hospitals, could be modeled. We could further model bed capacity that can be added without reducing regular care bed capacity. Furthermore, occupancy by high-priority regular care patients who have to be admitted to a hospital (i.e., regular patients arriving autonomously) could also be modeled according to the proposed queueing model and can be taken into account in our proposed decision rule. For instance, one could decide to allocate less infectious patients to hospitals where the strain on (future) capacity by high-priority (i.e., more severely ill) regular patients will be higher. Finally, it would be interesting to extend our modeling approach from the regional to the national level, also considering travel distances.

7 Acknowledgments

The authors thank Manon Bruens (Acute Zorg Euregio), Bart Veltman, Sophie Ligtenstein, and Franka Schellekens (Rhythm) for their collaboration in obtaining the input data for the case study in this paper.

References

- Amr Mohamed AbdelAziz, Louai Alarabi, Saleh Basalamah, and Abdeltawab Hendawi. A Multi-Objective Optimization Method for Hospital Admission Problem—A Case Study on Covid-19 Patients. *Algorithms*, 14, 38(2):1–14, 2021. doi:<https://doi.org/10.3390/a14020038>.
- Anders Reenberg Andersen, Bo Friis Nielsen, and Line Blander Reinhardt. Optimization of hospital ward resources with patient relocation using Markov chain modeling. *European Journal of Operational Research*, 260(3):1152–1163, 2017. ISSN 0377-2217. doi:<https://doi.org/10.1016/j.ejor.2017.01.026>.

- Nezir Aydin and Zeynep Cetinkale. Analyses on ICU and non-ICU capacity of government hospitals during the COVID-19 outbreak via multi-objective linear programming: An evidence from Istanbul. *Computers in Biology and Medicine*, 146:1–21, 2022. ISSN 0010-4825. doi:<https://doi.org/10.1016/j.compbiomed.2022.105562>.
- S. Baas, S. Dijkstra, A. Braaksma, P. van Rooij, F. J. Snijders, L. Tiemessen, and R. J. Boucherie. Real-time forecasting of COVID-19 bed occupancy in wards and Intensive Care Units. *Health Care Management Science*, 24:402–419, 2021. doi:<https://doi.org/10.1007/s10729-021-09553-5>.
- Michele Barbato, Alberto Ceselli, and Marco Premoli. On the impact of resource relocation in facing health emergencies. *European Journal of Operational Research*, 308(1):422–435, 2023. ISSN 0377-2217. doi:<https://doi.org/10.1016/j.ejor.2022.11.024>.
- Aaron S. Bernstein, Amy W. Ando, Ted Loch-Temzelides, Mariana M. Vale, Binbin V. Li, Hongying Li, Jonah Busch, Colin A. Chapman, Margaret Kinnaird, Katarzyna Nowak, Marcia C. Castro, Carlos Zambrana-Torrel, Jorge A. Ahumada, Lingyun Xiao, Patrick Roehrdanz, Les Kaufman, Lee Hannah, Peter Daszak, Stuart L. Pimm, and Andrew P. Dobson. The costs and benefits of primary prevention of zoonotic pandemics. *Science Advances*, 8:1–13, 2 2022. ISSN 2375-2548. doi:<https://doi.org/10.1126/sciadv.abl4183>.
- Michelle R de Graaff, Rianne N M Hogenbirk, Yester F Janssen, Arthur K E Elfrink, Ronald S L Liem, and et al. Impact of the COVID-19 pandemic on surgical care in the Netherlands. *British Journal of Surgery*, 109:1282–1292, 11 2022. ISSN 0007-1323. doi:<https://doi.org/10.1093/bjs/znac301>.
- Sander Dijkstra, Stef Baas, Aleida Braaksma, and Richard J Boucherie. Dynamic fair balancing of COVID-19 patients over hospitals based on forecasts of bed occupancy. *Omega*, 116:1–21, 2023. doi:<https://doi.org/10.1016/j.omega.2022.102801>.
- L Eriskin, M Karatas, and Y J Zheng. A robust multi-objective model for healthcare resource management and location planning during pandemics. *Ann Oper Res*, pages 1–48, 2022. . doi: <https://doi.org/10.1007/s10479-022-04760-x>.
- Mohammad Fattahi, Esmaeil Keyvanshokoo, Devika Kannan, and Kannan Govindan. Resource planning strategies for healthcare systems during a pandemic. *European journal of operational research*, 304(1):192–206, 2023. doi:<https://doi.org/10.1016/j.ejor.2022.01.023>.
- Elvan Gökcalp, M Selim Cakir, and Hasan Satis. Dynamic capacity planning of hospital resources under COVID-19 uncertainty using approximate dynamic programming. *Journal of the Operational Research Society*;1–13, 2023. doi:<https://doi.org/10.1080/01605682.2023.2168570>.
- Xuran Gong, Xiuxian Wang, Liping Zhou, and Na Geng. Managing hospital inpatient beds under clustered overflow configuration. *Computers & Operations Research*, 148:1–26, 2022. ISSN 0305-0548. doi:<https://doi.org/10.1016/j.cor.2022.106021>.
- Edward H. Kaplan. OM Forum-COVID-19 scratch models to support local decisions. *Manufacturing & Service Operations Management*, 22(4):645–655, 2020. doi:<https://doi.org/10.1287/msom.2020.0891>.
- Jules Le Lay, Vincent Augusto, Edgar Alfonso-Lizarazo, Malek Masmoudi, Baptiste Gramont, Xiaolan Xie, Bi-venu Bongue, and Thomas Celarier. COVID-19 Bed Management Using a Two-Step Process Mining and Discrete-Event Simulation Approach. *IEEE Transactions on Automation Science and Engineering*;;1–12, 2023. doi:<https://doi.org/10.1109/TASE.2023.3274847>.
- N. Litvak, M. van Rijsbergen, R.J. Boucherie, and M. van Houdenhoven. Managing the overflow of intensive care patients. *European Journal of Operational Research*, 185(3):998–1010, 2008. doi:<https://doi.org/10.1016/j.ejor.2006.08.021>.
- Xin Ma, Xue Zhao, and Pengfei Guo. Cope with the COVID-19 pandemic: Dynamic bed allocation and patient subsidization in a public healthcare system. *International Journal of Production Economics*, 243:1–14, 2022. ISSN 0925-5273. doi:<https://doi.org/10.1016/j.ijpe.2021.108320>.

- José Tomás Marquinez, Antoine Sauré, Alejandro Cataldo, and Juan-Carlos Ferrer. Identifying proactive ICU patient admission, transfer and diversion policies in a public-private hospital network. *European Journal of Operational Research*, 295(1):306–320, 2021. ISSN 0377-2217. doi:<https://doi.org/10.1016/j.ejor.2021.02.045>.
- W.A. Massey and W. Whitt. Networks of infinite-server queues with nonstationary Poisson input. *Queueing Systems*, 13(1-3):183–250, 1993.
- NOS nieuws. Onderzoek in Twente: speciale coronavleugel ziekenhuis heeft geen zin (Research in Twente: special Coronavirus department in hospital is of no additional use), January 2023. URL <https://nos.nl/artikel/2458753-onderzoek-in-twente-speciale-coronavleugel-ziekenhuis-heeft-geen-zin>. Accessed 9-11-2023.
- Felix Parker, Hamilton Sawczuk, Fardin Ganjkanloo, Farzin Ahmadi, and Kimia Ghobadi. Optimal Resource and Demand Redistribution for Healthcare Systems Under Stress from COVID-19, 2020. URL <https://arxiv.org/abs/2011.03528>. Accessed 19-10-2023.
- Warren B. Powell. *Reinforcement Learning and Stochastic Optimization*. Wiley, Hoboken, 1st edition, 4 2022. ISBN 9781119815037. doi:<https://doi.org/10.1002/9781119815068>. URL <https://onlinelibrary.wiley.com/doi/book/10.1002/9781119815068>. Accessed 19-10-2023.
- Maroeska M. Rovers, Stan R.W. Wijn, Janneke P.C. Grutters, Sanne J.J.P.M. Metsemakers, Robin J. Vermeulen, Ron Van Der Pennen, Bart J.J.M. Berden, Hein G. Gooszen, Mirre Scholte, and Tim M. Govers. Development of a decision analytical framework to prioritise operating room capacity: lessons learnt from an empirical example on delayed elective surgeries during the COVID-19 pandemic in a hospital in the Netherlands. *BMJ Open*, 12:1–10, 4 2022. ISSN 20446055. doi:<https://doi.org/10.1136/bmjopen-2021-054110>.
- Sobhan Sarkar, Anima Pramanik, J. Maiti, and Genserik Reniers. Covid-19 outbreak: A data-driven optimization model for allocation of patients. *Computers & Industrial Engineering*, 161:1–23, 2021. doi:<https://doi.org/10.1016/j.cie.2021.107675>.
- Kim Seung-Chul, Horowitz Ira, Young Karl K, and Buckley Thomas A. Flexible bed allocation and performance in the intensive care unit. *Journal of Operations Management*, 18(4):427–443, 2000. ISSN 0272-6963. doi:[https://doi.org/10.1016/S0272-6963\(00\)00027-9](https://doi.org/10.1016/S0272-6963(00)00027-9).
- Matthew P. Shearer, Diane Meyer, Divya Hosangadi, Michael R. Snyder, Marc Trotochaud, Syra Madad, and Jennifer B. Nuzzo. Operational stresses on New York City Health+Hospitals Health System frontline hospitals during the 2017-18 influenza season. *American Journal of Disaster Medicine*, 15:99–111, 4 2020. ISSN 1932-149X. doi:<https://doi.org/10.5055/ajdm.2020.0360>.
- StatLine. Gezondheid, leefstijl, zorggebruik en -aanbod, doodsoorzaken; vanaf 1900 (health, lifestyle, utilization and supply of healthcare, and causes of death since 1900), 7 2023. URL <https://opendata.cbs.nl/statline/?dl=37883#/CBS/nl/dataset/37852/table>. Accessed 19-10-2023.
- Li Sun, Gail W. DePuy, and Gerald W. Evans. Multi-objective optimization models for patient allocation during a pandemic influenza outbreak. *Computers and Operations Research*, 51:350–359, 2014. ISSN 03050548. doi:<https://doi.org/10.1016/j.cor.2013.12.001>.
- W. Whitt and X. Zhang. A data-driven model of an emergency department. *Operations Research for Health Care*, 12:1–15, 2017. doi:<https://doi.org/10.1016/j.orhc.2016.11.001>.
- Y Ye, L Huang, and J Wang et al. Patient allocation method in major epidemics under the situation of hierarchical diagnosis and treatment. *BMC Medical Informatics and Decision Making*, 22:1–18, 2022. doi:<https://doi.org/10.1186/s12911-022-02074-3>.

A Full tables for Section 5.2

Table 9: KPIs per hospital for decision rules (decision rule) IH and PU using the 80, 85, 90, 95% quantile over scenarios to forecast the occupancy in 2 days.

Decision rule	Hospital	KPI	80%	85%	90%	95%
IH	MST	Overbeds	0.350 ± 0.028	0.285 ± 0.026	0.230 ± 0.025	0.168 ± 0.024
		Underbeds	3.641 ± 0.117	4.092 ± 0.124	4.654 ± 0.135	5.439 ± 0.148
		Reg. beds used	11.970 ± 0.173	12.463 ± 0.179	13.056 ± 0.178	13.875 ± 0.175
		Beds reserved	1.154 ± 0.033	1.154 ± 0.033	1.163 ± 0.034	1.190 ± 0.036
		Rooms added/removed	0.119 ± 0.004	0.119 ± 0.004	0.119 ± 0.004	0.120 ± 0.004
	ZGT	Overbeds	0.366 ± 0.034	0.315 ± 0.034	0.269 ± 0.033	0.221 ± 0.033
		Underbeds	3.952 ± 0.081	4.329 ± 0.088	4.779 ± 0.093	5.414 ± 0.101
		Reg. beds used	13.423 ± 0.135	13.833 ± 0.132	14.309 ± 0.128	14.969 ± 0.119
		Beds reserved	1.070 ± 0.030	1.036 ± 0.030	1.001 ± 0.029	0.976 ± 0.025
		Rooms added/removed	0.092 ± 0.003	0.088 ± 0.003	0.085 ± 0.003	0.080 ± 0.003
	SKB	Overbeds	0.255 ± 0.029	0.231 ± 0.029	0.207 ± 0.029	0.186 ± 0.028
		Underbeds	2.805 ± 0.096	3.067 ± 0.101	3.388 ± 0.107	3.837 ± 0.119
		Reg. beds used	6.561 ± 0.103	6.838 ± 0.101	7.174 ± 0.101	7.632 ± 0.103
		Beds reserved	0.684 ± 0.025	0.685 ± 0.025	0.679 ± 0.023	0.675 ± 0.025
		Rooms added/removed	0.061 ± 0.003	0.060 ± 0.003	0.060 ± 0.003	0.056 ± 0.003
PU	MST	Overbeds	0.100 ± 0.011	0.099 ± 0.011	0.102 ± 0.012	0.096 ± 0.011
		Underbeds	6.581 ± 0.073	6.582 ± 0.073	6.570 ± 0.073	6.614 ± 0.074
		Reg. beds used	22.916 ± 0.114	22.917 ± 0.114	22.903 ± 0.115	22.943 ± 0.113
		Beds reserved	0.825 ± 0.022	0.826 ± 0.022	0.835 ± 0.024	0.829 ± 0.023
		Rooms added/removed	0.077 ± 0.003	0.077 ± 0.003	0.078 ± 0.003	0.077 ± 0.003
	ZGT	Overbeds	0.187 ± 0.025	0.184 ± 0.025	0.181 ± 0.025	0.173 ± 0.024
		Underbeds	3.313 ± 0.082	3.328 ± 0.082	3.359 ± 0.084	3.415 ± 0.086
		Reg. beds used	9.072 ± 0.119	9.088 ± 0.119	9.126 ± 0.118	9.199 ± 0.118
		Beds reserved	0.876 ± 0.026	0.882 ± 0.027	0.879 ± 0.026	0.895 ± 0.028
		Rooms added/removed	0.089 ± 0.003	0.089 ± 0.003	0.089 ± 0.003	0.088 ± 0.003
	SKB	Overbeds	0.277 ± 0.036	0.275 ± 0.036	0.268 ± 0.036	0.257 ± 0.035
		Underbeds	1.908 ± 0.075	1.931 ± 0.076	1.986 ± 0.078	2.097 ± 0.080
		Reg. beds used	4.754 ± 0.071	4.778 ± 0.070	4.836 ± 0.069	4.950 ± 0.066
		Beds reserved	0.427 ± 0.017	0.432 ± 0.017	0.437 ± 0.019	0.434 ± 0.017
		Rooms added/removed	0.040 ± 0.002	0.041 ± 0.002	0.040 ± 0.002	0.040 ± 0.002

Table 10: KPIs for SP-O for lookahead horizons 5, 4, and 3 days and when taking the 90% CI upper bound for the predicted arrival intensity (UB arrival rate).

Hospital	KPI	Overbeds	UB arrival rate	Horizon 4 days	Horizon 3 days
MST	Overbeds	0.140 ± 0.018	0.105 ± 0.017	0.161 ± 0.020	0.240 ± 0.020
	Underbeds	3.729 ± 0.104	4.403 ± 0.135	3.535 ± 0.106	3.132 ± 0.086
	Reg. beds used	12.187 ± 0.188	13.189 ± 0.209	11.822 ± 0.204	11.260 ± 0.199
	Beds reserved	0.828 ± 0.024	0.807 ± 0.024	0.916 ± 0.026	0.863 ± 0.023
	Rooms added/removed	0.085 ± 0.003	0.073 ± 0.002	0.092 ± 0.003	0.088 ± 0.003
ZGT	Overbeds	0.096 ± 0.012	0.072 ± 0.012	0.112 ± 0.011	0.165 ± 0.015
	Underbeds	3.397 ± 0.079	4.032 ± 0.086	2.952 ± 0.072	2.485 ± 0.068
	Reg. beds used	12.310 ± 0.171	13.044 ± 0.190	11.773 ± 0.179	10.983 ± 0.187
	Beds reserved	0.846 ± 0.024	0.774 ± 0.021	0.917 ± 0.026	0.865 ± 0.026
	Rooms added/removed	0.084 ± 0.003	0.074 ± 0.002	0.091 ± 0.003	0.089 ± 0.003
SKB	Overbeds	0.090 ± 0.010	0.076 ± 0.011	0.107 ± 0.012	0.154 ± 0.012
	Underbeds	2.509 ± 0.071	2.777 ± 0.084	2.345 ± 0.063	1.993 ± 0.052
	Reg. beds used	7.221 ± 0.144	7.231 ± 0.149	7.220 ± 0.137	6.940 ± 0.139
	Beds reserved	0.469 ± 0.018	0.442 ± 0.019	0.521 ± 0.022	0.470 ± 0.020
	Rooms added/removed	0.042 ± 0.002	0.040 ± 0.002	0.045 ± 0.002	0.041 ± 0.002
Region	Overbeds	0.326 ± 0.035	0.253 ± 0.036	0.380 ± 0.038	0.559 ± 0.041
	Underbeds	9.635 ± 0.181	11.212 ± 0.210	8.832 ± 0.172	7.610 ± 0.148
	Reg. beds used	31.717 ± 0.240	33.464 ± 0.240	30.814 ± 0.239	29.183 ± 0.242
	Beds reserved	2.143 ± 0.042	2.023 ± 0.040	2.355 ± 0.043	2.198 ± 0.039
	Rooms added/removed	0.211 ± 0.004	0.187 ± 0.004	0.228 ± 0.005	0.218 ± 0.004
Avg. cost SP	Avg. cost SP	51.129 ± 1.407	49.429 ± 1.459	52.956 ± 1.533	58.120 ± 1.676
	Forecast cost SP (hor. 1)	48.990 ± 1.238	50.002 ± 1.423	50.323 ± 1.293	52.878 ± 1.328
	Forecast cost SP (hor. 2)	50.908 ± 1.325	55.907 ± 1.644	50.311 ± 1.356	51.105 ± 1.307
	Forecast cost SP (hor. 3)	58.380 ± 1.797	68.829 ± 2.325	57.086 ± 1.867	56.332 ± 1.738

Table 11: KPIs for SP-O, as well as the deterministic version of the stochastic programs, when taking the median and 85% quantile over the arrivals and occupancy scenarios for each day.

Hospital	KPI	Overbeds	Median Scen.	Quantile Scen.
MST	Overbeds	0.140 ± 0.018	0.831 ± 0.039	0.107 ± 0.018
	Underbeds	3.729 ± 0.104	1.718 ± 0.061	5.512 ± 0.172
	Reg. beds used	12.187 ± 0.188	8.811 ± 0.179	14.439 ± 0.268
	Beds reserved	0.828 ± 0.024	0.885 ± 0.019	0.987 ± 0.034
	Rooms added/removed	0.085 ± 0.003	0.104 ± 0.003	0.086 ± 0.003
ZGT	Overbeds	0.096 ± 0.012	0.551 ± 0.028	0.070 ± 0.012
	Underbeds	3.397 ± 0.079	1.602 ± 0.054	4.389 ± 0.120
	Reg. beds used	12.310 ± 0.171	9.000 ± 0.181	12.779 ± 0.235
	Beds reserved	0.846 ± 0.024	0.880 ± 0.025	0.892 ± 0.024
	Rooms added/removed	0.084 ± 0.003	0.095 ± 0.003	0.076 ± 0.003
SKB	Overbeds	0.090 ± 0.010	1.073 ± 0.047	0.243 ± 0.034
	Underbeds	2.509 ± 0.071	0.764 ± 0.039	2.128 ± 0.080
	Reg. beds used	7.221 ± 0.144	5.933 ± 0.148	6.860 ± 0.168
	Beds reserved	0.469 ± 0.018	0.444 ± 0.018	0.460 ± 0.019
	Rooms added/removed	0.042 ± 0.002	0.040 ± 0.002	0.034 ± 0.002
Region	Overbeds	0.326 ± 0.035	2.455 ± 0.061	0.419 ± 0.047
	Underbeds	9.635 ± 0.181	4.084 ± 0.083	12.029 ± 0.199
	Reg. beds used	31.717 ± 0.240	23.744 ± 0.245	34.079 ± 0.240
	Beds reserved	2.143 ± 0.042	2.209 ± 0.034	2.340 ± 0.044
	Rooms added/removed	0.211 ± 0.004	0.239 ± 0.004	0.196 ± 0.005
Avg. cost SP	Avg. cost SP	51.129 ± 1.407	128.850 ± 2.542	57.299 ± 1.883
	Forecast cost SP (hor. 1)	48.990 ± 1.238	85.416 ± 1.812	72.173 ± 2.988
	Forecast cost SP (hor. 2)	50.908 ± 1.325	56.003 ± 1.257	105.605 ± 4.429
	Forecast cost SP (hor. 3)	58.380 ± 1.797	41.251 ± 1.445	156.446 ± 5.864