European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Capacity planning for effective cohorting of hemodialysis patients during the coronavirus pandemic: A case study

Cem D. C. Bozkir[a], Cagri Ozmemis[a], Ali Kaan Kurbanzade[a], Burcu Balcik[a,*], Evrim D. Gunes[b], Serhan Tuglular[c]

[a] Industrial Engineering Department, Ozyegin University, Istanbul, Turkey
[b] Business Administration, College of Administrative Sciences and Economics, Koc University, Sariyer, Istanbul, Turkey
[c] Medical Faculty, Department of Internal Medicine, Marmara University, Istanbul, Turkey

A R T I C L E I N F O

Article history:
Received 15 February 2021
Accepted 20 October 2021
Available online xxx

Keywords:
OR in health services
COVID-19 pandemic
Hemodialysis
Patient cohorting
Stochastic programming

A B S T R A C T

Planning treatments of different types of patients have become challenging in hemodialysis clinics during the COVID-19 pandemic due to increased demands and uncertainties. In this study, we address capacity planning decisions of a hemodialysis clinic, located within a major public hospital in Istanbul, which serves both infected and uninfected patients during the COVID-19 pandemic with limited resources (i.e., dialysis machines). The clinic currently applies a 3-unit cohorting strategy to treat different types of patients (i.e., uninfected, infected, suspected) in separate units and at different times to mitigate the risk of infection spread risk. Accordingly, at the beginning of each week, the clinic needs to allocate the available dialysis machines to each unit that serves different patient cohorts. However, given the uncertainties in the number of different types of patients that will need dialysis each day, it is a challenge to determine which capacity configuration would minimize the overlapping treatment sessions of different cohorts over a week. We represent the uncertainties in the number of patients by a set of scenarios and present a stochastic programming approach to support capacity allocation decisions of the clinic. We present a case study based on the real-world patient data obtained from the hemodialysis clinic to illustrate the effectiveness of the proposed model. We also compare the performance of different cohorting strategies with three and two patient cohorts.

1. Introduction

The coronavirus disease (COVID-19) has caused a global pandemic, as declared by the World Health Organization in March 2020 (Sohrabi et al., 2020). It has created a significant burden on the capacity of healthcare systems all around the world due to a surge in demand for outpatient clinics, inpatient beds, and intensive care units. While many people could postpone their nonurgent hospital visits to protect from infection risk (Hafner, 2020), some patients with chronic health conditions (such as hemodialysis, chemotherapy, physical therapy) do not have the luxury to avoid or postpone visits to health care facilities since they are obliged to obtain medical care regularly. It has also been a challenge for health care providers to serve such chronic patients without exposing them to infection risk with ever-tighter resources. This paper addresses the capacity planning decisions of a hospital's hemodialysis clinic that provides care to infected and uninfected patients during the COVID-19 pandemic by implementing cohorting strategies to mitigate the risk of infection spread among patients.

Hemodialysis (HD) is a treatment to remove waste products and excess fluid from the blood with the help of a machine when kidneys fail to work (The National Health Service, 2021). Most chronic HD patients must receive dialysis treatment three times a week, during which patients stay connected to a dialysis machine for about four hours. During their frequent and long hospital visits to get treatments, chronic HD patients can become highly susceptible to COVID-19 transmission risk (Lano et al., 2020). Indeed, their infection risk is found as nearly two-fold greater than those receiving dialysis at home (Hsu & Weiner, 2020). Moreover, chronic kidney disease is a prevalent risk factor for severe COVID-19 cases and deaths (ERACODA Working Group, 2021; Ozturk et al., 2020). Therefore, it is imperative for health care providers to take precautions to ensure that the uninfected chronic HD patients are not in contact with the COVID-19 infected and suspected HD patients during their clinic visit (Basile et al., 2020). Given that the global prevalence rate of chronic kidney disease (all stages) is 9.1%

* Corresponding author.
E-mail address: burcu.balcik@ozyegin.edu.tr (B. Balcik).

https://doi.org/10.1016/j.ejor.2021.10.039
0377-2217/© 2021 Elsevier B.V. All rights reserved.

(Bikbov et al., 2020), effective planning of HD treatments during the pandemics by accounting for infection spread risk can protect the lives of millions of people in vulnerable conditions. However, even without a pandemic disturbance, scheduling treatments in a single HD unit is a difficult problem (e.g., Fleming, Gartner, Padman, & James, 2019; Holland, 1994; Liu et al., 2019). When resources become tighter, such as the aftermath of an earthquake (Sever, Lameire, & Vanholder, 2009) or during a pandemic (Corbett et al., 2020), the increased demands and uncertainties make it even more challenging to plan for HD services.

In this study, motivated by a real case, we focus on HD services provided in a hospital's HD clinic during the COVID-19 pandemic where multiple patient types are treated. Specifically, besides the chronic HD patients, the clinic serves admitted patients that need dialysis due to acute kidney failure for various medical conditions. Moreover, the HD clinic is responsible for treating the suspected and infected COVID-19 patients that cannot be treated in small unequipped private HD centers. Managing HD resources effectively to threat these patients is challenging for the hospital during the pandemic due to several reasons. First, the effective HD capacity can be reduced due to the additional time required for cleaning the rooms in-between HD sessions, as well as arrangements required to cohort patients. Specifically, the uninfected (clean), infected (confirmed), and suspected (shows symptoms or had contact with a confirmed case) HD patients must be treated at different designated units; and if possible, during different (nonoverlapping) periods. Secondly, the demand for HD during a pandemic increases because dialysis sessions may be needed for some infected patients due to the adverse effects of COVID-19 on kidneys (Sperati, 2020); indeed, 35.2% of COVID-19 patients needed dialysis (Klein et al., 2020; Smith, 2020). Thirdly, while the number of chronic HD patients on a treatment schedule is known, the number of infected and clean acute patients that need HD treatments each day can be highly uncertain since the number of people infected with COVID-19 disease fluctuates due to various reasons such as the course of the lockdown measures and emergence of new variants. In this resource-constrained uncertain environment, employing Operations Research techniques can be valuable to make effective cohorting, capacity planning and treatment scheduling decisions.

Treating the infected, suspected, and uninfected HD patients at different times and in separate areas are common approaches followed in HD facilities to mitigate the risk of infection transmission among patients during the COVID-19 pandemic (Boushab et al., 2021; Kliger & Silberzweig, 2020). Different cohorting strategies, which involve defining patient types and assigning them to separate areas, have been implemented around the world (e.g., Collison et al., 2020; Whiteside, Kane, Aljohani, Alsamman, & Pourmand, 2020), and there is no well-established standard yet. Park et al. (2020) and Meijers, Messa, & Ronco (2020) summarize the classification of different hemodialysis patients that are either infected by COVID-19 or under the suspicion of infection and outline the operationalization of the HD clinics in Korea, Belgium, and Italy. While cohorting has been implemented in HD clinics around the world during the COVID-19 pandemic, to the best of our knowledge, no study has yet focused on modeling and evaluating alternative cohorting strategies where both rooms and treatment times of different patient cohorts are separated to mitigate infection spread, and the capacity configuration decisions are made to minimize overlapping HD sessions of different patient cohorts. We present a novel stochastic programming approach for implementing cohorting strategies in a hospital's HD clinic effectively.

Our case study focuses on the operations of an HD clinic in a large public hospital in Istanbul, Turkey, which is one of the designated "pandemic hospitals" during the COVID-19 pandemic. To treat different patient types (i.e., uninfected chronic and acute patients, and infected and suspected patients), the clinic implements a 3-unit cohorting strategy to mitigate infection spread risk. To facilitate this, the HD clinic is divided into units separated by drywalls, where each unit has several machines. Accordingly, uninfected, suspected and infected patients are treated at different units. Moreover, if possible, the HD sessions of different patient cohorts are scheduled at different times to further decrease physical interaction. An alternative 2-unit cohorting strategy could separate the clinic into two units (Whiteside et al., 2020), one for the uninfected patients, and the other for suspected and infected patients, where these patients would be treated sequentially in the unit. Given a cohorting strategy, we address the problem of assigning dialysis machines to the units in the HD clinic thereby setting their capacity to serve different patient cohorts during a week, whose demand is uncertain.

Our collaboration with the HD clinic has two main objectives: (1) To provide the clinic with a decision support tool based on mathematical modeling to support their capacity planning decisions and cohorting practices, (2) To evaluate the performance of alternative cohorting strategies with two and three units and provide insights for their potential use. To address the uncertainty in demands for HD, we use a stochastic programming approach, which is an optimization framework used to deal with problems that involve uncertain parameters (Birge & Louveaux, 2011). Specifically, we develop a two-stage stochastic programming model, which consists of decisions made before (i.e., first stage) and after (i.e., second stage) the realization of uncertainty, represented by a set of scenarios (Shapiro, Dentcheva, & Ruszczyński, 2014; Wallace & Ziemba, 2005). Accordingly, we create a set of discrete probabilistic scenarios based on forecasts developed from the HD clinic's historical demand data. We then set the number of dialysis machines allocated to each unit at the beginning of a week (first stage), based on the expected outcomes of the second stage decisions that determine a daily dialysis treatment schedule in each scenario. The objective is to treat all patients while minimizing the expected number of patients from different cohorts having HD at the overlapping sessions. Following such a stochastic optimization approach is essential in our context since neglecting the significant uncertainties in HD demands of different patient types may lead to suboptimal solutions with higher overlapping sessions and increased infection risk. Moreover, a solution that assumes deterministic demand may even be infeasible in some demand realizations. We test the proposed approach on real-world data obtained from our collaborating hospital and present results that illustrate the benefits of using the proposed models and the comparative performance of different cohorting strategies.

The remainder of this paper is structured as follows. In Section 2, we review the related literature. In Section 3, we describe the system, define the problem and present our models. In Section 4, we present a case study and discuss the results. Finally, we present our conclusions in Section 5.

2. Literature review

We review the literature related to relevant health care planning problems under uncertainty, cohorting practices and planning of hemodialysis services during pandemics.

2.1. Health care planning under uncertainty

There exists a rich literature that utilizes operations research (OR) methods to address various health care planning problems. Existing studies address problems in different settings ranging from planning the operations of a single hospital or a unit (e.g., intensive care units, operating rooms) to a network of health providers (see reviews by Ahmadi-Javid, Jalali, & Klassen, 2017; Hulshof, Kortbeek, Boucherie, Hans, & Bakker, 2012; Rais & Viana, 2011). Stochastic optimization techniques are widely used in healthcare systems given the significant uncertainties in various parameters such as demand, arrival times, and service durations. It is common to represent uncertain parameters by a set of scenarios and then develop two-stage stochastic programming (SP) models, in which first-stage decisions are made under uncertainty by considering the possible outcomes of the second stage decisions that are made when scenarios are realized. For instance, Marchesi, Hamacher, & Fleck (2020) address a physician staffing and scheduling problem in an emergency department and account for the uncertainties in patient arrival times and service durations to minimize the expected number of patients in the queue. Punnakitikashem, Rosenberber, & Buckley-Behan (2013) address a nurse staffing and assignment problem and consider uncertain patient care times in minimizing the expected excess workload of nurses. Gul (2018) set surgery appointment schedules under the uncertainties in surgery duration and room turnover times to optimize waiting and idleness objectives. Chalabi, Epstein, McKenna, & Claxton (2008) focus on allocating resources between different healthcare programs with uncertain outcomes to maximize expected benefits.

A growing number of studies address problems for allocating and managing valuable health care resources during epidemics and pandemics, where significant uncertainty exists in demands. Klein et al. (2020) present a review of studies that present models to estimate demand surge for hospital capacity and resources (such as hospital beds, ventilators) during the pandemics. Parker, Sawczuk, Ganjkhanloo, Ahmadi, & Ghobadi (2020) study a demand and resource redistribution problem among multiple facilities and develop alternative models to minimize the required surge capacity and resource shortages by considering demand uncertainty. Mehrotra, Rahimian, Barah, Luo, & Schantz (2020) present a two-stage SP model to allocate and share ventilators among risk-averse states by considering demand uncertainty. Finally, Yarmand, Ivy, Denton, & Lloyd (2014) and Tanner, Sattenspiel, & Ntaimo (2008) propose two-stage SP models for optimizing vaccine allocation during epidemics under uncertain demand and epidemiological parameters.

In aforementioned studies that present two-stage SP models for health care planning problems, it is common to generate scenarios for uncertain parameters by sampling realizations from a fitted probability distribution based on the available data (e.g., Punnakitikashem et al., 2013), from the confidence interval of the forecasts (e.g., Mehrotra et al., 2020), or from the specific distributions based on epidemiological studies (e.g., Tanner et al., 2008), as well as by using other analytical techniques such as Monte Carlo or Latin Hypercube Sampling (e.g., Chalabi et al., 2008). In this study, we focus on weekly capacity planning decisions in an HD clinic to treat multiple patient types in a pandemic setting and propose a two-stage SP model, where the uncertainty in patient demands are represented by scenarios generated from the prediction intervals of forecasts.

2.2. Cohorting practices

Several studies provide examples of cohorting strategies implemented in health care facilities throughout the world during the COVID-19 pandemic. For example, Alharbi, Alharbi, & Alqaidi (2020) and George et al. (2020) describe how dental emergencies and cholesteatoma surgeries are ordered based on patients' infection status in HD clinics in Saudi Arabia and Switzerland, respectively. Moreover, Park et al. (2020), Meijers et al. (2020), Rubin (2020) and Boushab et al. (2021) describe cohorting practices implemented in HD clinics in South Korea, Italy & Belgium, United States and Mauritania, respectively. As described in these examples, patient classification can be different in different cohorting applications.

Although cohorting practices are widely used to avoid infection spread during pandemics, a limited number of OR studies address cohorting related problems. Most studies focus on allocating beds to different hospital wards that serve different types of patients. Pinker & Tezcan (2013) address capacity configuration decisions in a hospital with the existence of patients needing isolation due to infectious diseases. They present a stochastic optimization model with a revenue maximization objective to reserve beds to the isolation and non-isolation patients, where the two groups have different requirements for spacing and admission criteria. Chia & Lin (2015) use simulation to analyze the effectiveness of different cohorting strategies in a pediatric hospital for allocating beds among different patient types. Wang, Gong, Geng, Jiang, & Zhou (2020) focus on allocating beds among different departments in a hospital under stochastic patient arrival and service times and propose an approach that combines optimization and simulation models. Melman, Parlikad, & Cameron (2020) focus on the operations of a hospital during the COVID-19 pandemic and develop a discrete event simulation model to evaluate alternative strategies for allocating limited resources effectively in a hospital to schedule the surgeries of infected and uninfected patients. In this paper, we introduce a new capacity allocation problem for implementing 2- and 3-unit cohorting strategies in an HD clinic during a pandemic. While the proposed models are tailored to represent the operational rules of the case clinic, the proposed approach can be easily adapted to different cohorting applications and operational settings around the world.

2.3. Hemodialysis planning during pandemics

The literature that addresses HD planning during pandemics mostly focuses on providing practical medical guidelines to manage this process effectively (Naicker et al., 2020). For instance, Ikizler & Kliger (2020), Kliger & Silberzweig (2020) and Rubin (2020) highlight the importance of patient screening and environmental disinfection in HD clinics. CMS (2020) suggests that uninfected HD patients should wait for their sessions outside the hospital and there should be at least two meters among patients during their treatments. Alberici et al. (2020) describe the layout changes made in a hospital to provide treatments to uninfected and infected kidney transplant patients. Several studies highlight the importance of cohorting HD patients based on their infection status (e.g., Ikizler & Kliger, 2020; Kliger & Silberzweig, 2020; Roper et al., 2020). As reviewed previously, several studies describe the implementation of cohorting in HD clinics during the COVID-19 pandemic (e.g., Boushab et al., 2021; Meijers et al., 2020; Park et al., 2020; Rubin, 2020).

Although there exist a few OR studies that address problems related to planning HD treatments, no study has yet addressed the particular challenges faced during a pandemic. The existing studies usually consider deterministic settings. Liu et al. (2019) study a dialysis scheduling problem with different types of patients (conventional, hepatitis B, and hepatitis C) and present an integer programming model to minimize night shifts and meet patient preferences. Fleming et al. (2019) also address a scheduling problem by considering each dialysis machine as a workstation and present an integer programming model to minimize the waiting time of dialysis patients at the clinic for treatment and the scheduled finish time of treatments each day. Yu, Kulkarni, & Deshpande (2020) present appointment scheduling policies for patients that need a series of appointments (e.g., for chemotherapy and HD treatments) and present a Markov decision process model that considers revenue per service per patient, and the costs of staffing, overtime, overbooking, and delay.

Table 1
Patient types treated in the HD clinic during the pandemic.

Patient Types	Description
Uninfected Acute (Type 1)	Uninfected patients that are admitted to the hospital and need hemodialysis.
Uninfected Chronic (Type 2)	Uninfected chronic patients that receive regular HD treatment
Infected COVID-19 (Type 3)	COVID-19 infected HD patients.
Suspected COVID-19 (Type 4)	Suspected HD patients with the possibility of having COVID-19 infection.

Currie et al. (2020) express the need for scientific studies to prevent the infection rate of COVID-19 in HD units and highlight that this area is lacking in the literature. Corbett et al. (2020) discuss the challenges in planning the HD treatments under a high level of uncertainty and stress the need for an analytical method. Ikizler & Kliger (2020) highlight that the prolonged pandemic overextends the healthcare capacities and creates shortages, which necessitates allocating resources efficiently in cohorting symptomatic and clean patients. Although several studies draw attention to the need for effective management of dialysis resources during a pandemic, no work in the literature provides analytical methods for supporting cohorting and the capacity planning decisions in an HD unit to serve different types of patients by avoiding infection spread. We contribute to the literature by presenting a stochastic programming approach to support effective decision making for configuring an HD clinic's capacity to allocate resources and plan HD sessions during the pandemic to avoid infection spread. While we consider four types of patients and cohorting strategies with two and three units, as implementable in our case hospital, the proposed models can be extended to other settings that consider cohorting strategies with a different number of cohorts and patient types.

3. Problem description & mathematical model

In this section, we first describe the system based on the operations of our collaborating hemodialysis (HD) clinic located in a major hospital in Turkey (Section 3.1). We then define the problem (Section 3.2) and present our formulations developed for two alternative cohorting strategies (Section 3.3).

3.1. System description and objectives of the study

We focus on the operations of an HD clinic that treats different types of patients during the COVID-19 pandemic. The clinic is located within Marmara University's Pendik Training and Research Hospital, which is a major public hospital located in Istanbul. In Turkey, all major public hospitals have been designated as "pandemic hospitals" to treat COVID-19 infected or suspected patients. COVID-19 patients could be treated only at the pandemic hospitals for several months after the official declaration of the pandemic in March 2020. During this period, the pandemic hospitals continued providing their regular services as well, which necessitated using resources efficiently more than ever.

In Turkey, chronic HD patients can get their regular treatments in HD clinics located in public hospitals or private HD centers, covered by national health insurance. Chronic patients are served by their registered clinics three days a week, according to a Monday–Wednesday–Friday (MWF) or a Tuesday–Thursday–Saturday (TTS) regime. Different from private HD centers, major public hospitals also have inpatients (i.e., admitted to various wards or intensive care units) who may need temporary HD treatment during their stay at the hospital (e.g., after an operation). However, the COVID-19 pandemic required making significant changes in this routine in our case hospital, which has additionally become responsible for serving the COVID-19 patients with the same level of resources (dialysis machines, rooms). Specifically, since the private dialysis facilities are not fully equipped for managing possible complications of dialysis patients with COVID-19, their chronic HD patients are referred to pandemic hospitals if they become infected with COVID-19. Additionally, some COVID-19 patients who do not have a priori chronic kidney disease may need HD due to COVID-19 related complications. These patients are also treated at pandemic hospitals. Our case hospital's HD clinic has to admit all confirmed and suspected COVID-19 patients for dialysis treatment until they fully recover. To summarize, four types of patients with different characteristics are served in the HD clinic during the COVID-19 pandemic (see Table 1).

Dialysis operations and the current (3-unit) cohorting strategy: The HD clinic in our case hospital has 14 dialysis machines. Starting from 8:00 am, four HD sessions can be performed on each machine each day (Monday–Saturday). On Sundays, the clinic serves only emergency patients. In HD clinics throughout the world, it is common to run sessions in parallel; that is, HD sessions start and end at the same time throughout the day. One machine serves one patient in each session. In our case hospital, an HD session for a patient takes four hours including the preparation times. Although there may be some variability in treatment start times and durations due to unexpected events such as patient lateness or technical problems, the clinic constructs daily treatment schedules by considering a fixed four-hour duration per session. We, therefore, assume that each HD session has a fixed duration for all patients.

Since the beginning of the pandemic, our case hospital separates the HD patients into three cohorts (i.e., uninfected, infected, suspected), where each cohort is treated in a separate unit (standard, isolated, and quarantine, respectively). All machines and units are cleaned intensively between HD sessions. A nurse can handle the treatment of multiple patients simultaneously in the same unit in our case. The clinic arranges the configuration of the units at the beginning of each week and makes weekly plans such as nurse shifts. For instance, only two machines were allocated to infected patients during the first months of the pandemic, while there were five machines in the isolated unit during the data collection period (November–December 2020). A schematic of the HD clinic during this period is given in Fig. 1. As shown in the figure, seven, five and two machines are currently allocated to standard, isolated and quarantine units, respectively.

Scheduling treatments: In addition to treating different patient cohorts in separate units, the clinic schedules the treatments of different cohorts at nonoverlapping sessions. Since all units are along the same corridor, overlapping sessions may increase the risk of infection spread among patients and their companions. Hence, as long as the unit capacities permit, the uninfected patients are treated in morning sessions, followed by the sessions of suspected, and then the infected patients. Therefore, the number of machines allocated to each unit at the beginning of a week directly affects the number of overlapping HD sessions that will be incurred each day. Currently, the hospital does not use any analytical tools that can support decisions for determining the size of units. Given the uncertainties in the number of patients that will need dialysis treatments each day, it is challenging to make weekly capacity planning decisions by considering their implications on daily treatment schedules and overlapping sessions. Our first objective in this study is to develop an optimization model that can assist the hos-

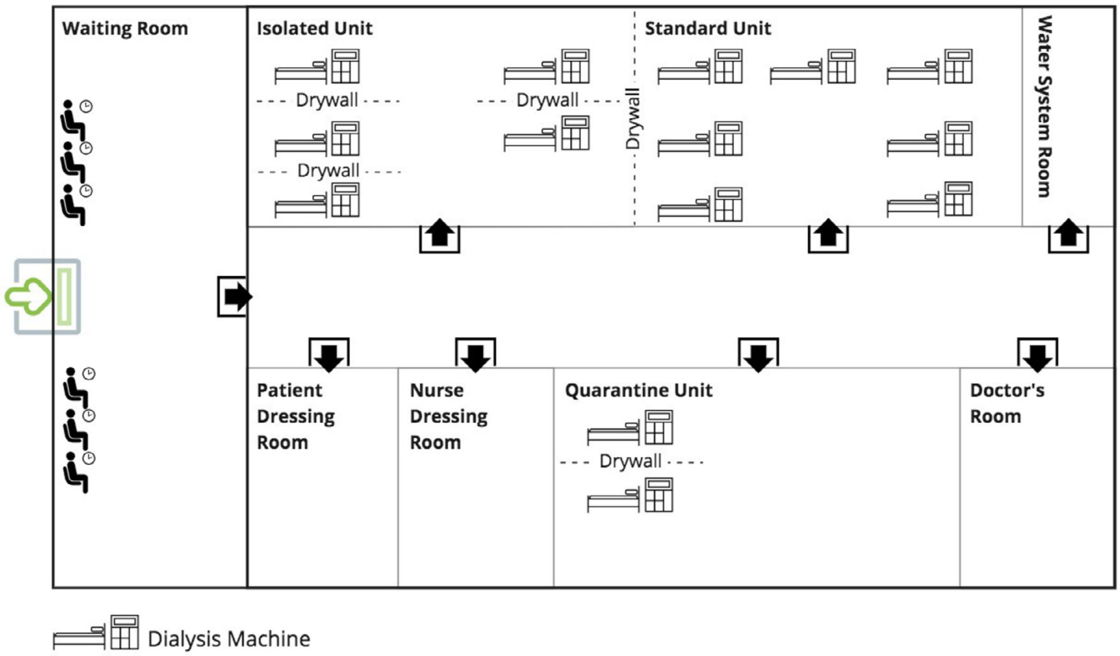

Fig. 1. HD clinic floor plan and units.

pital to determine unit capacities effectively under demand uncertainty.

Alternative (2-unit) cohorting strategy: While our case hospital has separated the HD clinic into three units, an alternative cohorting strategy could be to divide the clinic into two units, where the uninfected patients are treated in the first unit, and the suspected and infected patients can be sequentially treated in the second unit, as practiced in some cases (Whiteside et al., 2020). We investigate whether this alternative cohorting strategy might work better for our HD clinic to avoid overlapping sessions, which will highly depend on patient demands over a week. Hence, our second objective in this paper is to evaluate the performance of different cohorting strategies and make recommendations to the hospital.

Next, we define our problem reflecting the described characteristics of the HD operations in the case hospital during the COVID-19 pandemic.

3.2. Problem definition

We introduce the capacity planning problem of an HD clinic, which must allocate dialysis machines to different units that serve different patient cohorts during a pandemic. Each HD unit has a capacity restriction in terms of number of machines imposed by physical constraints. The number of uninfected acute patients (Type 1) that need dialysis each day is uncertain. The daily demand for chronic HD patients (Type 2) is known as these patients receive regular treatments. The daily demand by suspected and infected COVID-19 patients (Type 3 and 4) are also uncertain, which depend on uncontrollable factors including the current infection rate in the population and referrals from other HD clinics. We represent the uncertainties in the number of Type 1, 3 and 4 patients that will need HD each day by a set of discrete scenarios, which can be generated by making predictions from the past demand data (see Section 4).

We consider two alternative cohorting strategies: dividing the HD clinic into three units, which is the current practice, or two units. We assume that the hospital sets the cohorting strategy in advance as a policy. Table 2 shows the patient types that are assigned to clinic units in each cohorting strategy. Besides treating

Table 2
Unit assignments of patient types in 3-unit and 2-unit cohorting strategies.

Units	3-unit cohorting	2-unit cohorting
Standard	Uninfected Acute (Type 1) Uninfected Chronic (Type 2)	Uninfected Acute (Type 1) Uninfected Chronic (Type 2)
Isolated	Infected COVID-19 (Type 3)	Infected COVID-19 (Type 3) Suspected COVID-19 (Type 4)
Quarantine	Suspected COVID-19 (Type 4)	

different patient cohorts in different units, it is desirable to treat different patient types at different HD sessions (i.e. time slots) to avoid further physical interaction in the HD clinic. Only Type 1 and 2 patients can be served in the same unit during an HD session since these patients are uninfected. Otherwise, an HD session must be composed of patients of the same type. When Type i and Type j patients from different cohorts have to be treated in the same HD session, an "overlap $i \times j$", denoted by $O_{i\times j}$, occurs. Avoiding the overlaps of uninfected patients with infected or suspected patients (i.e., overlaps 1×3, 1×4, 2×3, 2×4) is of primary importance. Moreover, in the 3-unit cohorting strategy where the infected and suspected patients are treated in different rooms, their overlap (overlap 3×4) should be avoided if possible; however, this overlap is considered to be less serious compared to overlaps that involve clean patients. Furthermore, in the 2-unit cohorting strategy, Type 3 and 4 patients must be sequentially treated in the same unit, and Type 4 patients are always treated earlier than Type 3 patients.

Fig. 2 illustrates example daily treatment schedules and the resulting overlaps for both cohorting strategies on a small instance with seven machines and 15 patients. The number of patients of Type 1–4 are three, five, four and three, respectively. As shown in Fig. 2(a), when three, two, and two dialysis machines are allocated to standard, isolated and quarantine units, respectively, some patients from each type have to be served in overlapping sessions. In Fig. 2(b), three and four machines are allocated to standard and isolated units, respectively, and the HD sessions of some uninfected (Type 1 and 2) patients and suspected (Type 4) patients overlap. Therefore, both the cohorting policy and the capacity configuration

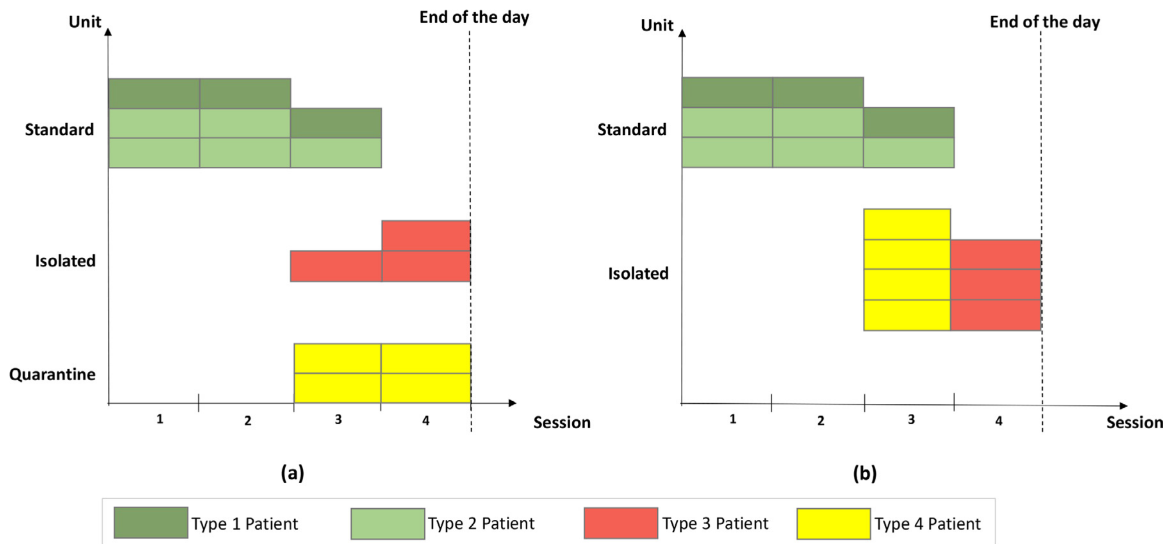

Fig. 2. Example daily treatment schedules under (a) 3-unit cohorting and (b) 2-unit cohorting.

decisions can have a significant effect on the number of patients who are exposed to infection risk each day.

In our problem, we consider a planning horizon of one week for the capacity planning decisions, which involve determining the number of dialysis machines to be allocated to each unit. We follow a two-stage approach to formulate the HD clinic's problem, where capacity planning is done in the first stage, accounting for the second stage treatment scheduling decisions made under different demand scenarios. Treatment scheduling decisions involve determining the timing of HD sessions allocated to each patient group each day given the predetermined first-stage decisions and the realized patient demands.

The main objective of our capacity planning problem is to allocate machines to units by minimizing the total expected penalties that will occur over a week due to overlapping sessions of different cohorts. Specifically, a penalty is incurred for each patient that is scheduled in an overlapping session with other patients treated in different units. We impose an overlapping penalty charge per patient since there might be a higher infection risk if a larger number of conflicting patients are in the clinic at the same time. Additionally, if there are patients that cannot be served due to capacity insufficiency, we impose a very large penalty per patient. Finally, to avoid alternative solutions that spread patients unnecessarily over multiple HD sessions, which may arise when capacity is abundant, we impose a small penalty for each HD session performed over a day.

While we capture all main features of the hospital's capacity planning problem during the COVID-19 pandemic, we make two simplifying assumptions in formulating this problem. First, we assume that each patient requires one HD session of fixed duration, which may not be true for some acute patients that need longer dialysis as required by their varying medical conditions. However, such cases can be easily handled by the proposed model by allocating multiple HD sessions to such patients by creating dummy patients. Secondly, we assume that all patients can be treated according to the constructed schedules, and ignore exceptions due to emergency needs. While this is a reasonable assumption for our hospital since emergency arrivals constitute a small percentage of patients in our case, extending the proposed models to accommodate random dialysis requests (e.g., by enhancing the current scenario definition to include random arrival times) is possible.

We next present two stochastic optimization models developed for solving the described capacity planning problems associated with alternative cohorting strategies.

3.3. Mathematical modeling

We present the notation and the two-stage stochastic programming models for 3- and 2-unit models in Sections 3.3.1 and 3.3.2, respectively.

Sets

K: set of scenarios; $k \in K$

I: set of patient types; $i \in I$; $i = 1, \ldots, 4$

J: set of HD units; $j \in J$; $j = 1, \ldots, 3$ (1: Standard unit, 2: Isolated unit, 3: Quarantine unit)

T: set of days in the planning horizon; $t \in T$; $t = 1, \ldots, 6$

S: set of HD sessions on each day; $s \in S$; $s = 1, \ldots, 4$

Parameters

p_k: probability of scenario k

H_{it}^k: number of type i patients who need to receive HD on day t of scenario k

C_j: maximum number of dialysis machines that can be allocated to unit j

$\hat{C}$: total number of available dialysis machines

α_1: penalty (per patient) for uninfected patients and infected COVID-19 patients that are treated in overlapping HD sessions

α_2: penalty (per patient) for uninfected (Type 1, 2) patients and suspected COVID-19 (Type 4) patients that are treated in overlapping HD sessions

α_3: penalty (per patient) for suspected (Type 4) and infected (Type 3) patients that are treated in overlapping HD sessions

Π_i: penalty (per patient) for infeasibility if a Type i patient cannot be treated due to capacity insufficiency.

ϵ: small penalty coefficient for starting an HD session (to avoid unnecessary sessions)

First stage decision variables

R_j: number of dialysis machines allocated to HD unit j

Second stage decision variables

X_{its}^k: number of Type i patients scheduled to receive HD treatment in session s on day t of scenario k

F_{it}^k: number of Type i patients that cannot be served on day t of scenario k

N^k_{jts}: $\begin{cases} 1, & \text{if HD unit type } j \text{ is used for session } s \text{ on day } t \text{ of scenario } k, \\ 0, & \text{otherwise.} \end{cases}$

Q^k_{ts}: number of uninfected (Type 1, 2) and infected (Type 3) patients treated in session s on day t of scenario k

G^k_{ts}: number of uninfected (Type 1, 2) and suspected (Type 4) patients treated in session s on day t of scenario k

W^k_{ts}: number of infected (Type 3) and suspected (Type 4) patients treated in session s on day t of scenario k

U^k_{ts}: $\begin{cases} 1, & \text{if an overlap exists in session } s \text{ on day } t \text{ of scenario } k \text{ for uninfected (Type 1, 2) and infected (Type 3) patients,} \\ 0, & \text{otherwise.} \end{cases}$

D^k_{ts}: $\begin{cases} 1, & \text{if an overlap exists in session } s \text{ on day } t \text{ of scenario } k \text{ for uninfected (Type 1, 2) and suspected (Type 4) patients} \\ 0, & \text{otherwise.} \end{cases}$

V^k_{ts}: $\begin{cases} 1, & \text{if an overlap exists in session } s \text{ on day } t \text{ of scenario } k \text{ for infected (Type 3) and suspected (Type 4) patients,} \\ 0, & \text{otherwise.} \end{cases}$

Z_{3u}: auxiliary variable for the objective function

3.3.1. 3-unit cohorting model

The formulation for the capacity planning problem to implement the 3-unit cohorting strategy is as follows.

$$\min \quad Z_{3u} = \sum_{k\in K} p_k \Bigg[\alpha_1 \sum_{t\in T, s\in S} Q^k_{ts} + \alpha_2 \sum_{t\in T, s\in S} G^k_{ts} + \alpha_3 \sum_{t\in T, s\in S} W^k_{ts} + \Pi_i \sum_{i\in I, t\in T} F^k_{it} + \epsilon \sum_{j\in J, t\in T, s\in S} N^k_{jts} \Bigg] \tag{1}$$

subject to

$$R_j \le C_j \qquad \forall j \in J \tag{2}$$

$$\sum_j R_j \le \hat{C} \tag{3}$$

$$X^k_{1ts} + X^k_{2ts} \le R_1 \qquad \forall t \in T, s \in S, k \in K \tag{4}$$

$$X^k_{3ts} \le R_2 \qquad \forall t \in T, s \in S, k \in K \tag{5}$$

$$X^k_{4ts} \le R_3 \qquad \forall t \in T, s \in S, k \in K \tag{6}$$

$$\sum_s X^k_{its} = H^k_{it} - F^k_{it} \qquad \forall i \in I, t \in T, k \in K \tag{7}$$

$$X^k_{1ts} + X^k_{2ts} \le M_1 \times N^k_{1ts} \qquad \forall t \in T, s \in S, k \in K \tag{8}$$

$$N^k_{1ts} \ge N^k_{1t(s+1)} \qquad \forall s \in S : s \le 3, t \in T, k \in K \tag{9}$$

$$X^k_{3ts} \le M_1 \times N^k_{2ts} \qquad \forall t \in T, s \in S, k \in K \tag{10}$$

$$X^k_{4ts} \le M_1 \times N^k_{3ts} \qquad \forall t \in T, s \in S, k \in K \tag{11}$$

$$\sum_{s'=s+2}^{S} N^k_{jts'} \le |S|(1 - N^k_{jts} + N^k_{jt(s+1)}) \qquad \forall s \in S : s \le 2, t \in T, j \in J : j \ne 1, k \in K \tag{12}$$

$$1 + U^k_{ts} \ge N^k_{1ts} + N^k_{2ts} \qquad \forall t \in T, s \in S, k \in K \tag{13}$$

$$Q^k_{ts} \ge X^k_{1ts} + X^k_{2ts} + X^k_{3ts} - M_2(1 - U^k_{ts}) \qquad \forall t \in T, s \in S, k \in K \tag{14}$$

$$1 + V^k_{ts} \ge N^k_{2ts} + N^k_{3ts} \qquad \forall t \in T, s \in S, k \in K \tag{15}$$

$$W^k_{ts} \ge X^k_{3ts} + X^k_{4ts} - M_2(1 - V^k_{ts}) \qquad \forall t \in T, s \in S, k \in K \tag{16}$$

$$1 + D^k_{ts} \ge N^k_{1ts} + N^k_{3ts} \qquad \forall t \in T, s \in S, k \in K \tag{17}$$

$$G^k_{ts} \ge X^k_{1ts} + X^k_{2ts} + X^k_{4ts} - M_2(1 - D^k_{ts}) \qquad \forall t \in T, s \in S, k \in K \tag{18}$$

$$R_j \in \mathbb{Z}^+ \qquad \forall j \in J \tag{19}$$

$$X^k_{its} \in \mathbb{Z}^+ \qquad \forall i \in I, t \in T, s \in S, k \in K \tag{20}$$

$$F^k_{it} \in \mathbb{Z}^+ \qquad \forall i \in I, t \in T, k \in K \tag{21}$$

$$Q^k_{ts}, W^k_{ts}, G^k_{ts} \in \mathbb{Z}^+ \qquad \forall t \in T, s \in S, k \in K \tag{22}$$

$$N^k_{jts} \in \{0, 1\} \qquad \forall j \in J, t \in T, s \in S, k \in K \tag{23}$$

$$U^k_{ts}, V^k_{ts}, D^k_{ts} \in \{0, 1\} \qquad \forall t \in T, s \in S, k \in K \tag{24}$$

The first three terms of the objective function (1) minimizes the weighted sum of the penalties due to overlapping sessions. The fourth and fifth terms in (1) are for penalizing patients that cannot be treated in the clinic due to capacity insufficiency. The last term in the objective function minimizes the number of HD sessions provided during a week to facilitate constructing compact schedules. Constraints (2) ensure that the number of machines allocated to each unit does not exceed its capacity. Constraint (3) ensures that the total number of dialysis machines assigned to units does not exceed the number of machines available in the clinic. Constraints (4)–(6) guarantee that for any unit, the number of patients assigned to an HD session cannot exceed the capacity reserved for that unit. Constraints (7) balance the daily number of HD sessions required and provided. Constraints (8) and (9) ensure that HD sessions of uninfected (Type 1 and 2) patients start from the beginning of the day, and multiple HD sessions are conducted consecutively in the standard unit. Constraints (10) and (11) are for determining used HD sessions of suspected (Type 4) and infected (Type 3) patients anytime during a day; additionally, constraints (12) ensure that the sessions in their corresponding HD units must be conducted consecutively for the compactness of the schedule. In other words, by keeping track of the two consecutive sessions, (12) prevents leaving an idle session between any sessions in which suspected (Type 4) and infected (Type 3) patients are treated. Constraints (13) determine the overlapping HD sessions among standard and isolated units. Accordingly, the number of uninfected (Type 1 and 2) and infected (Type 3) patients that receive HD in the same session is determined by constraints (14). Similarly, the number of infected (Type 3) and suspected (Type 4) patients that receive HD in the same session is determined by constraints (15) and (16); further, the number of uninfected (Type 1 and 2) and suspected (Type 4) patients that receive HD in the same session is determined by constraints (17) and (18). Finally, the integer variables are defined by constraints (19)–(22), and binary variables are defined by constraints (23) and (24). In implementing this model, we set the value of $M_1 = \hat{C}$ in constraints (8), (10), and (11) and $M_2 = |S| \times \hat{C}$ in constraints (14), (16), and (18).

3.3.2. 2-Unit cohorting model

The 2-unit model has a few differences compared to the 3-unit model presented above. First, there exist two units, that is, $j = 1, 2 \in J$, where $j = 1$ denotes the standard unit where the uninfected patients are treated, and $j = 2$ denotes the isolated unit where suspected and infected patients are treated sequentially. To sequence the treatments of infected patients after the suspected patients in the isolated unit, we introduce a new variable Y^k_{ts}, which takes the value of 1 when all suspected patients that can be treated in a day are assigned to an HD session, thereby specifying the earliest session that the infected patients can be scheduled in the unit. Moreover, a dummy HD session is needed to schedule consecutive HD sessions of suspected and infected patients in the same cohort by using variables Y^k_{ts}. We denote the dummy session by $\{0\}$ and let $S_0 = S \cup \{0\}$. The W^k_{ts} and V^k_{ts} variables that are used before in the 3-unit cohorting model are not relevant in the 2-unit model, since the suspected and infected patients are treated in the same unit in separate HD sessions. Finally, we define Z_{2u} as an auxiliary variable to keep the objective function value attained for the 2-unit model.

We next present the model for the capacity planning problem to implement a 2-unit cohorting strategy.

$$\min \quad Z_{2u} = \sum_{k\in K} p_k \Big[\alpha_1 \sum_{t\in T, s\in S} Q^k_{ts} + \alpha_2 \sum_{t\in T, s\in S} G^k_{ts} + \Pi_i \sum_{i\in I, t\in T} F^k_{it} + \epsilon \sum_{j\in J, t\in T, s\in S} N^k_{jts}\Big] \tag{25}$$

subject to

(2), (3), (4), (7), (8), (9), (13), (14)
(17), (18), (19), (20), (21), (23)

$$X^k_{3ts} + X^k_{4ts} \le R_2 \qquad \forall t \in T, s \in S, k \in K \tag{26}$$

$$X^k_{3ts} + X^k_{4ts} \le M_3 \times N^k_{2ts} \qquad \forall t \in T, s \in S, k \in K \tag{27}$$

$$\sum_{s'=s+2}^{4} N^k_{2ts'} \le |S|(1 - N^k_{2ts} + N^k_{2t(s+1)}) \qquad \forall s \in S : s \le 2, t \in T, k \in K \tag{28}$$

$$H^k_{4t} - \sum_{s'=0}^{s} X^k_{4ts'} - F^k_{4t} \le M_4 \times Y^k_{ts} \qquad \forall t \in T, s \in S_0, k \in K \tag{29}$$

$$X^k_{3ts} \le M_4 \times (1 - Y^k_{t(s-1)}) \qquad \forall t \in T, s \in S, k \in K \tag{30}$$

$$X^k_{it0} = 0 \qquad \forall i \in I, t \in T, k \in K \tag{31}$$

$$Q^k_{ts}, G^k_{ts} \in \mathbb{Z}^+ \qquad \forall t \in T, s \in S, k \in K \tag{32}$$

$$U^k_{ts}, D^k_{ts}, Y^k_{ts} \in \{0, 1\} \qquad \forall t \in T, s \in S, k \in K \tag{33}$$

The objective function (25) is similar to (1), except that there is no penalty now associated with overlapping Type 3 and 4 patients, which are sequentially treated in the same unit in the 2-unit cohorting strategy. Moreover, constraints (5) and (6) in the 3-unit model are consolidated in a single constraint (26) in the 2-unit model since Type 3 and 4 patients are assigned to the same unit. Similar to constraints (10) and (11), constraints (28) ensure that HD sessions for Type 3 and 4 infected patients can start anytime during a day, but their sessions must be conducted consecutively. Constraints (29) and (30) specify the earliest HD session for Type 3 patients after Type 4 patients are scheduled. Constraints (33) define the domains of variables. The remaining constraints are identical to those of the 3-unit model. In implementing this model, we set $M_3 = \hat{C}$ in constraints (27). M_4 values in constraints (29) and (30) are set to the largest number of Type 4 and Type 3 demand occurrences among $|K|$ scenarios, respectively.

4. Case study

This section presents a case study to implement the proposed models based on data obtained from our collaborating HD clinic. Fig. 3 illustrates our roadmap for the case study, which involves two main phases. The first phase focuses on **Data Analysis and Instance Generation** (Section 4.1), which consists of two main steps: data collection and processing, followed by demand prediction and scenario generation. In the **Model Implementation and Results** phase (Section 4.2), we analyze the capacity allocation problem of the HD clinic in three steps. Specifically, we present results to evaluate the performance of the hospital's current solution, analyze the capacity allocation solutions of the stochastic optimization models, and compare the solutions obtained by alternative cohorting strategies. Finally, we discuss the practical implication of results (Section 4.3).

4.1. Data analysis and instance generation

4.1.1. Data collection and processing

As described before, the hemodialysis clinic of our collaborating hospital has been applying a 3-unit cohorting strategy since the beginning of the pandemic to treat four types of patients (see Tables 1 and 2). During our collaboration, a nurse working in the hemodialysis clinic of the hospital has recorded patient data (i.e., the number of patients of each type treated in the clinic each day) for eight weeks in November and December 2020, which corresponds to a period around the second peak of the pandemic.

While our study focuses on solving the problem of this specific HD clinic based on the provided patient data, the proposed approach is general and can be followed similarly in other settings, where there are several units to be used for cohorting purposes. To implement our approach, it is sufficient to have information on patient demands, the unit's physical capacities, and the total number of dialysis machines available. Here we describe our observations based on the clinic's data to give more information on our case.

The analysis of patient data shows that the clinic served an average of 25 patients per day and provided a total of 1049 HD sessions over the two months. 47 of the 1049 HD sessions, mainly conducted for COVID-19 infected patients, lasted between 2 and 4 hours, whereas the remaining sessions lasted 4 hours. Moreover, the number of Type 3 and 4 patients treated in the clinic was 46 using 210 HD sessions. 32 of these 210 sessions, 6 of which were for infected patients, could be scheduled in earlier slots (between 8 AM - 4 PM). In other words, while the clinic scheduled most of the sessions of infected and suspected patients at the end of the day, in some cases, infected patients received dialysis in overlapping HD sessions with uninfected patients. Furthermore, Type 4

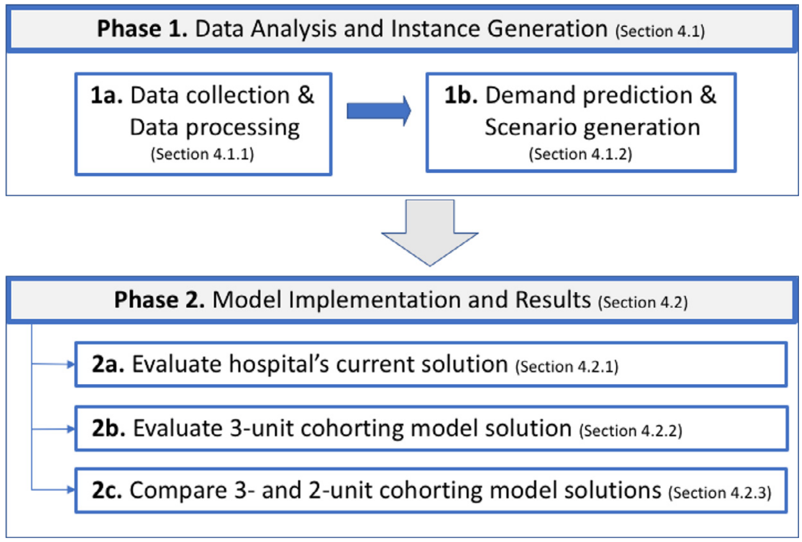

Fig. 3. Case study design.

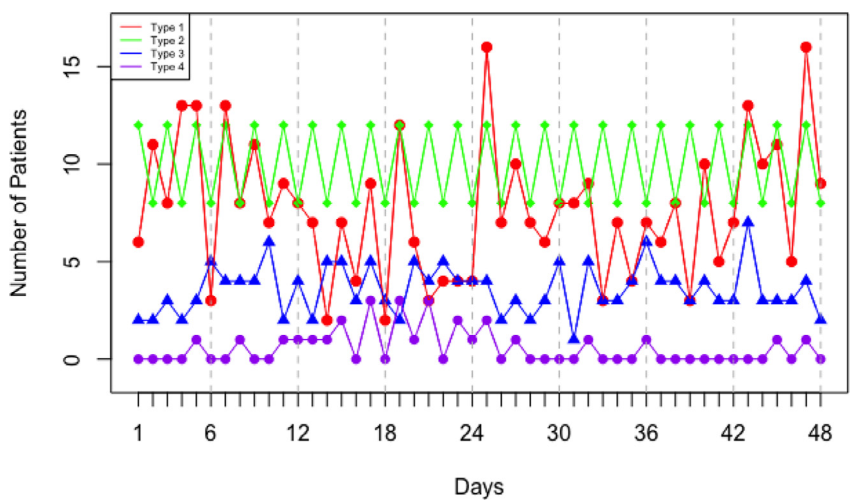

Fig. 4. Daily number of HD patients by type.

patients received only one HD session before their status was confirmed as COVID-19 infected, and they continued treatments as an infected patient. Patients with a confirmed COVID-19 case (Type 3) receives 4.3 HD sessions on average, whereas a maximum of 16 HD sessions were observed for such a patient. This wide range shows that the number of HD sessions provided per patient, especially for COVID-19 infected cases, is highly variable, increasing the demand uncertainty.

After processing the patient data to create inputs for implementing the proposed approach, we observe that only a few Type 1 patients are treated on Sundays, so cohorting or capacity planning does not have any effects. Therefore, we exclude Sundays from analysis, and thus we obtain a data set for 48 days in total. Fig. 4 presents a plot of the daily number of patients for each type where the grey dashed lines indicate weekly intervals (see Appendix A for the associated raw data). We observe that the daily number of uninfected acute patients (Type 1) is highly uncertain, as shown by the fluctuations in the plot. In contrast, the number of uninfected chronic patients (Type 2) shows a periodic pattern. Furthermore, the infected COVID-19 (Type 3) and suspected COVID-19 (Type 4) patients constitute only 19% of the overall demand, while the number of infected COVID-19 (Type 3) patients also varies considerably. Based on these patient data, we generate demand predictions and demand scenarios, explained in the next subsection.

4.1.2. Demand prediction and scenario generation

We consider a weekly (six days) planning horizon since the unit configuration can only be changed at the beginning of the week. Therefore, the number of patients to be treated over a week is a key input to our model. One possible approach to predict demands could be based on the general infection rate in the country (or state) and the percentage of dialysis patients in the population (for example Mehrotra et al., 2020 uses IHME forecast for ventilator demand which is based on epidemiological models). Fig. 5 shows the number of infected (Type 3) and suspected (Type 4) COVID-19 patients treated in our case hospital and the daily reported COVID-19 infected patients in Turkey between November 2 and December 26 (excluding Sundays). We observe that the number of Type 3 and 4 patients treated in the pandemic hospital does not closely follow the same pattern with the number of cases in the country. When the COVID-19 cases peaked in the middle of November, the numbers of dialysis patients with COVID-19 increased and peaked as well. However, after a curfew was announced at the beginning of December, indicated with the purple dashed line in the figure, the number of nationwide cases started to diminish, while the num-

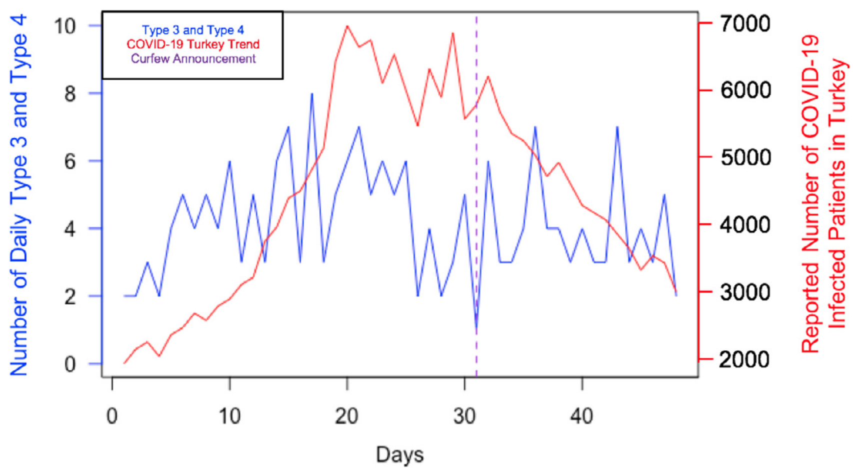

Fig. 5. Total number of Type 3 and 4 patients treated in the hemodialysis clinic and the reported national COVID-19 cases (November–December 2020).

ber of dialysis patients with COVID-19 seems stationary. There is also a higher variability in demand for HD over time, which is at a smaller scale compared to the aggregate number of cases at the national level. SIR-based disease progression dynamic models have been used to estimate the demand for hospital services (cf. Weissman et al., 2020). However, we need to predict the demand for a single hospital, for which aggregate predictions may not work well. Moreover, time-series methods are expected to work well for short-term forecasting during the pandemic (Doornik, Castle, & Hendry, 2020). Therefore, we chose to use time series forecasting methods to estimate the demand for the upcoming week based on the hospital's historical data. We next explain the steps followed to apply forecasting methods and generate scenarios.

As shown by hospital's data (Fig. 4), the number of Type 2 patients follows a periodic pattern; that is, 12 and 8 patients are enrolled in Monday-Wednesday-Friday (MWF) and Tuesday-Thursday-Saturday (TTS) regimes, respectively. We set the future daily demands for Type 2 patients based on these values. For the other demand types, we generate demand predictions and scenarios for three weeks (i.e., Weeks 6, 7, and 8), using the data from the previous weeks. We implement the following procedure to generate patient demand scenarios for each week: (1) Split the historical data into training and testing sets, (2) Choose a forecasting method, (3) Compute prediction intervals for the next week's demand, (4) Generate scenarios by randomly sampling instances from the estimated demand distribution, explained below.

Data splitting. To choose a forecasting method, the last week of the data are used as a testing set, whereas the earlier data points are used as the training data set. The accuracy of the estimators is verified by comparing the estimators with the testing set containing six data points. More specifically, our training sets consist of 30, 36, and 42 data points for Weeks 6, 7, and 8, respectively (as we consider six-day weeks). However, we keep the testing set as six days for all of our experiments because we aim to make an accurate decision for the study clinic over a planning horizon of one week. In other words, our training-test split ratio varies between datasets, the smallest being 12.5% (for 8 weeks) and the largest being 16.7% (for 6 weeks), which are within widely used 10% and 20% range for split ratios (Kuhn & Johnson, 2013).

Forecasting. To predict the demands of Type 1, 3, and 4 patients, we evaluate simple exponential smoothing and double exponential smoothing methods. We choose the simple exponential smoothing method for Weeks 6, 7, and 8 since this method minimizes the Root Mean Square Errors (RMSE) for the test data. We then find the point estimate for the next week's daily demand for patient type i, $\hat{Y}_i$.

Computing prediction intervals. The 80% and 90% prediction intervals (PIs) are calculated as $\hat{Y}_i \pm 1.28\sigma_i$ and $\hat{Y}_i \pm 1.64\sigma_i$ respectively, where the standard deviation σ_i is approximated by the RMSE. These intervals can be considered as representing a less and more conservative approach in accounting for uncertainty, respectively. The PIs for each patient type for each week are presented in Table 3. Fig. 6 depicts the data and the prediction intervals (PI) for Type 1, 3, and 4 patients, respectively, using seven weeks of past data to predict week eight. The blue line represents the point estimate, the green and red dashed lines represent the 80% and 90% PI, respectively.

Scenario generation. For each instance, we generate 30 equiprobable scenarios to test our models. Each scenario includes demand realizations for each type of patient for each day. Assuming a uniform distribution over the range of each PI (Table 3), we generate demand realizations by randomly sampling from the PI, and rounding negative values to 0 and real numbers to the nearest integer.

Finally, other parameters used in our case instance are as follows. The total number of dialysis machines available in the clinic, $\hat{C}$, is 14. The unit capacities are set as $C_1 = 11$, $C_2 = 8$ and $C_3 = 5$ for the 3-unit cohorting model, and as $C_1 = 11$ and $C_2 = 8$ for the 2-unit cohorting model. In both models, the objective function weights are set as $\alpha_1 = \alpha_2 = 1000$, $\alpha_3 = 100$, $\Pi_i = 100,000$ and $\epsilon = 2$. The proposed models are solved by Gurobi 9.0.3 on a computer with an Intel(R) Core (TM) i7-9750H CPU @ 2.60 GHz processor and a 16 GB RAM. We next present the results.

4.2. Model implementation and results

This section presents our numerical results (as outlined in Fig. 3). First, in Section 4.2.1, we evaluate the effectiveness of the hospital's current capacity configuration based on the actual (realized) demand. We also evaluate the effect of using optimization in making capacity allocation decisions. In Section 4.2.2, we evaluate the solutions of the stochastic optimization model for the 3-unit cohorting strategy. Finally, in Section 4.2.3, we compare the solutions of alternative cohorting policies and evaluate the solutions of the 2-unit cohorting model.

4.2.1. Evaluation of the hospital's current capacity configuration

This section evaluates how the hospital's current solution performs in terms of the number of overlapping patients. As described in Section 3.1, our case hospital applies a 3-unit cohorting strategy to minimize the spread risk of COVID-19 among its patients. During November–December 2020, the number of machines allocated

Table 3
Demand prediction intervals for Weeks 6, 7 and 8.

	Week 6		Week 7		Week 8	
Patient Type	80% PI	90% PI	80% PI	90% PI	80% PI	90% PI
Type 1	(2.31, 11.94)	(0.96, 13.29)	(2.30, 11.40)	(1.01, 12.68)	(2.40, 11.12)	(1.18, 12.65)
Type 3	(1.87, 5.18)	(1.41, 5.64)	(2.20, 5.71)	(1.70, 6.21)	(1.87, 5.17)	(1.40, 5.64)
Type 4	(−0.73, 1.77)	(−1.08, 2.21)	(−0.76, 1.59)	(−1.09, 1.92)	(−1.01, 1.17)	(−1.32, 1.48)

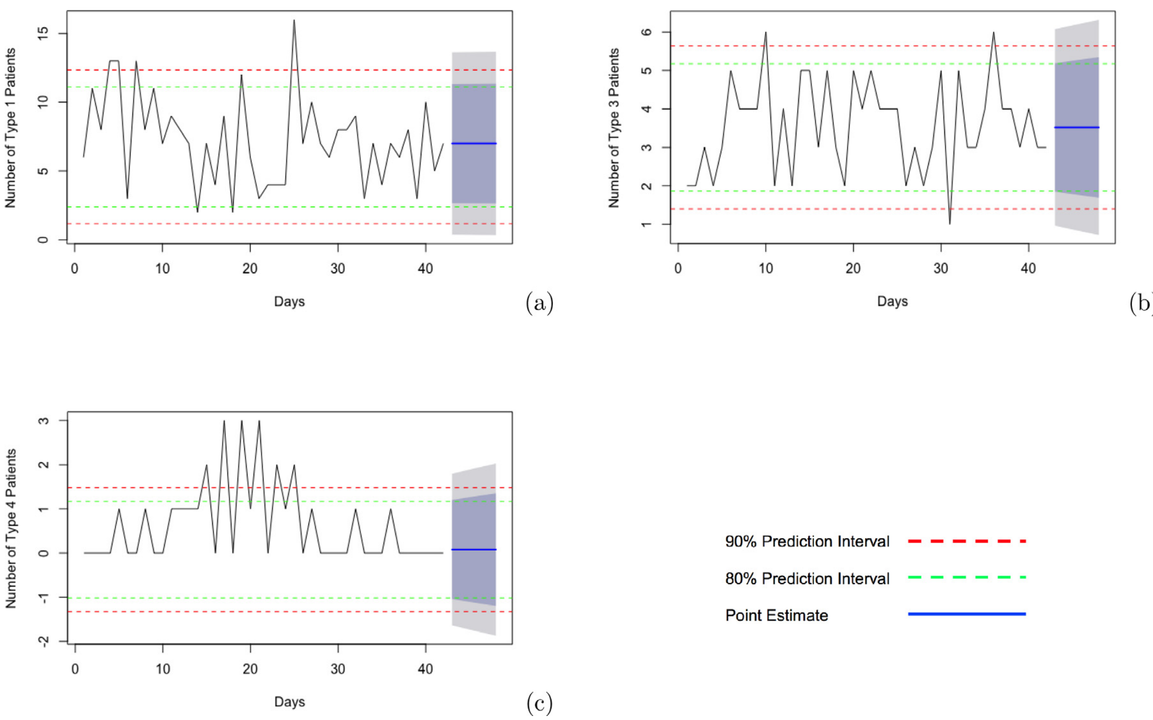

Fig. 6. Prediction intervals for Type 1 (a), Type 3 (b) and Type 4 (c) patients for Week 8.

Table 4
Comparison of the overlaps under hospital's current capacity allocation ($R_1 = 7, R_2 = 5, R_3 = 2$) with optimal allocation under perfect demand information.

	Overlaps under realized demand (hospital's allocation)				Deterministic solution			Overlaps under realized demand (deterministic solution)				
Week	$O_{(1,2)\times 3}$	$O_{(1,2)\times 4}$	$O_{3\times 4}$	Z_{3u}	R_1	R_2	R_3	$O_{(1,2)\times 3}$	$O_{(1,2)\times 4}$	$O_{3\times 4}$	Z'_{3u}	$\frac{Z_{3u}-Z'_{3u}}{Z_{3u}}$
1	7	5	4	12,450	10	3	1	0	0	4	444	96%
2	16	0	13	17,360	9	4	1	0	0	3	350	98%
3	0	8	17	9752	7	5	2	0	8	17	9752	0%
4	5	15	15	21,558	8	3	3	0	0	17	1758	92%
5	15	11	4	26,456	8	4	2	8	6	10	15,050	43%
6	2	0	12	3254	9	4	1	0	0	5	546	83%
7	0	0	0	48	10	4	0	0	0	0	36	25%
8	34	11	4	45,458	9	4	1	15	2	19	17,950	61%
Total	79	50	69	136,336	–	–	–	23	16	75	45,886	66%

to standard, isolated and quarantine units were $R_1 = 7, R_2 = 5, R_3 = 2$, respectively.

We compare the overlaps that occur by using hospital's current solution and an optimized solution under perfect information. First, to determine the number of weekly overlaps that would occur given the hospital's capacity configuration, we solve the second stage problem of the 3-unit cohorting model for each week by using realized demands. Second, we find the optimal capacity allocation by solving the 3-unit cohorting model with the observed demand as the only scenario in the second stage, which is labelled as the deterministic solution. Table 4 presents the number of overlaps for each type of $O_{i\times j}$ for these two solutions. For example, $O_{(1,2)\times 3}$ is the number of overlaps among uninfected (Type 1 and 2) patients with infected (Type 3) patients.

As shown in Table 4, except for Week 3, there is a difference in capacity allocation decisions. Moreover, allocation decisions can significantly impact the overlaps of uninfected and infected patients. We explain the differences in the following. For all weeks except Week 3, the hospital reserves fewer dialysis machines for the uninfected patients than they should (i.e., optimal R_1 is greater than 7). As a result, since the capacity of earlier HD sessions is not enough, some uninfected patients receive HD at later sessions of the day, during which suspected and infected patients also receive their treatments. As a result, under the hospital's allocation policy, the number of overlaps in uninfected and suspected patients' treatments is larger. Furthermore, the hospital's allocation is optimal only in Week 3, in which demands from uninfected acute patients and suspected patients are lower than aver-

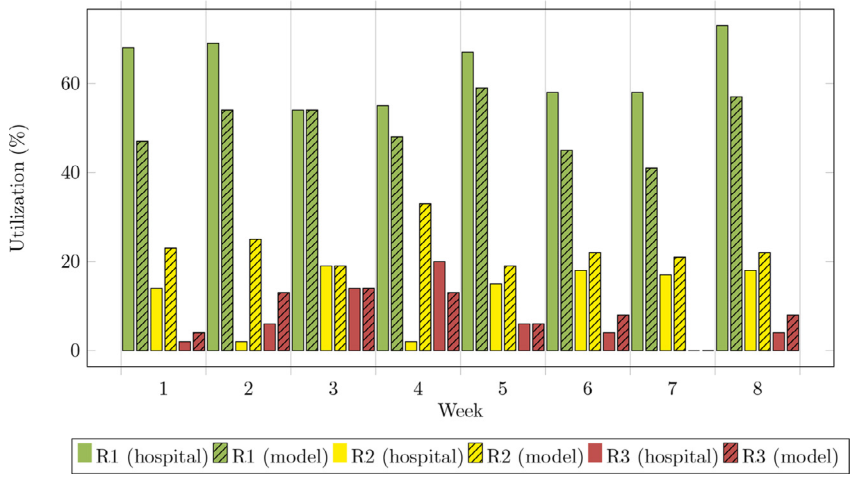

Fig. 7. Utilization of units in hospital's capacity allocation policy versus optimal allocation policy.

Table 5
Performance of the stochastic optimization model for 3-unit cohorting.

		Optimal allocation for 3-unit cohorting (stochastic solution)							Overlaps under realized demand		
Week	PI (%)	R_1	R_2	R_3	$E[O_{(1,2)\times 3}]$	$E[O_{(1,2)\times 4}]$	$E[O_{3\times 4}]$	Z_{3u}	$O_{(1,2)\times 3}$	$O_{(1,2)\times 4}$	$O_{3\times 4}$
6	80	8	5	1	0.0	0.9	9.3	1883	0	0	8
6	90	8	5	1	0.8	1.3	7.6	2979	0	0	8
7	80	8	5	1	1.4	0.5	10.1	2998	0	0	0
7	90	10	3	1	3.4	0.4	6.4	4561	0	0	0
8	80	8	5	1	0.0	0.0	5.1	561	24	5	9
8	90	8	5	1	1.0	0.0	4.2	1497	24	5	9

age, whereas a larger number of HD sessions is needed for infected patients.

The average daily utilization of each unit provides more insights into the HD clinic's current allocation policy. Fig. 7 presents the utilization of units in model's and hospital's allocations. We observe that for the hospital's current allocation policy, other than the third week, the standard unit's utilization is higher than the optimal allocation's utilization. In contrast, the average utilization of both isolated and quarantine units is lower than the optimal. This is because the optimal deterministic solution recommends allocating lower capacity for treating infected and suspected patients.

The results here, which are based on actual demand, show that there might be an opportunity to make better capacity allocation decisions and reduce the number of overlapping patients in the HD clinic. That is, an optimization approach can be helpful to reach more effective solutions. However, in reality, weekly capacity allocation decisions must be made under significant demand uncertainty. Therefore, in the following subsection, we evaluate the performance of the proposed stochastic optimization model in making capacity allocation decisions by accounting for demand uncertainty.

4.2.2. Evaluation of 3-unit cohorting model's solutions

We now present the solutions of the stochastic optimization model for 3-unit cohorting. As described in Section 4.1, we generate scenarios by predicting demands for Weeks 6, 7 and 8 based on the data from the HD clinic. Table 5 presents the solutions achieved by the stochastic optimization model for these weeks under two scenario sets generated by assuming two Prediction Intervals (PI). Further, the number of overlaps that would result from these capacity allocation solutions based on the realized patient demands is presented.

According to the results, the stochastic optimization model tends to allocate a larger number of dialysis machines to the standard unit (R_1) that treats uninfected (Type 1 and 2) patients compared to the hospital's current solution. Moreover, the capacity configuration achieved by the stochastic model leads to a lower number of overlaps for these three weeks compared to those incurred under the hospital's allocation (Table 4). Specifically, in Week 6, the hospital's allocation causes an overlap of uninfected and infected patients, which are avoided in stochastic model's solution. Moreover, in Week 8, the model decreases the overlaps of infected and suspected patients with uninfected patients. While this comes at a cost of some increase in the overlaps of suspected and infected patients, since these overlaps are considered less critical, the model solution is an improvement over the hospital's allocation. We also observe that while Week 6 and 8 capacity allocation solutions are robust to the different PIs used to develop scenarios, Week 7 solutions are affected by the different set of scenarios; specifically, in Week 7, the model allocates more machines to treat uninfected patients when more conservative demand realizations are accounted for.

4.2.3. Evaluation of different cohorting strategies

We next investigate the performance effects of different cohorting strategies. First, we solve the 3-unit and 2-unit cohorting models with realized demand for eight weeks, and the cost of overlap values of these deterministic solutions are compared. In addition, we evaluate the performance of the stochastic optimization model for the 2-unit cohorting case under the realized demand. Table 6 presents the solutions obtained by 2- and 3-unit cohorting models by considering the realized demands of the hospital for eight weeks. The last column of Table 6 presents the percentage difference in the objective functions of 2- and 3-unit models.

As observed in Table 6, the solutions of the 3-unit model dominate those of the 2-unit model in seven weeks out of eight. Since

Table 6
Solutions of 2-unit and 3-unit cohorting models under deterministic demand.

	3-unit cohorting							2-unit cohorting					
	Deterministic solution			Overlaps under realized demand (deterministic solution)				Deterministic solution		Overlaps under realized demand (deterministic solution)			
Week	R_1	R_2	R_3	$O_{(1,2)\times 3}$	$O_{(1,2)\times 4}$	$O_{3\times 4}$	Z_{3u}	R_1	R_2	$O_{(1,2)\times 3}$	$O_{(1,2)\times 4}$	Z_{2u}	$\frac{Z_{3u}-Z_{2u}}{Z_{3u}}$
1	10	3	1	0	0	4	444	11	3	0	4	4042	−810%
2	9	4	1	0	0	3	350	10	4	0	2	2050	−486%
3	7	5	2	0	8	17	9752	9	5	0	11	11,050	−13%
4	8	3	3	0	0	17	1758	9	5	0	9	9048	−415%
5	8	4	2	8	6	10	15,050	10	4	0	13	13,046	13%
6	9	4	1	0	0	5	546	8	6	0	2	2044	−274%
7	10	4	0	0	0	0	36	10	4	0	0	36	0%
8	9	4	1	15	2	19	17,950	10	4	8	13	21,048	−13%
Total	–	–	–	23	16	75	45,886	–	–	8	54	62,364	–

Type 3 and Type 4 patients can be treated in the same session in different units in the 3-unit model, the model assigns all of these patients to the last (i.e., the fourth) HD session of a day if possible. If the capacity of a single HD session is not sufficient, some patients are assigned to the third session, during which uninfected (Type 1 and 2) patients are usually treated as well. Therefore, by allowing overlaps in Type 3 and 4 patients' treatments, the 3-unit model avoid more serious overlaps of these patients' treatments with uninfected (Type 1 and 2) patients, thus incurring smaller penalties in our instances. In the 2-unit case, Type 3 and 4 patients receive dialysis in the same unit, but suspected patients' HD sessions precede those of the infected ones. Consequently, since the whole unit is allocated to a single patient group when the demand from one group of patients is low, the unit capacity cannot be utilized fully. Moreover, assigning suspected patients to HD sessions earlier in the day may cause increased overlaps with uninfected patients and result in worse objective values, as observed from Table 6.

According to Table 6, the 2-unit model slightly outperforms the 3-unit model only in Week 5, which is an exceptional case for the considered time horizon. When we examine the solutions for this week, we observe that the 2-unit model allocates more machines to the standard unit than the 3-unit model. Specifically, the optimal allocations are $R_1 = 10$, $R_2 = 4$ for the 2-unit model, and $R_1 = 8$, $R_2 = 4$, $R_3 = 2$ for the 3-unit model. In Fig. 8, we present the resulting treatment schedules for two days from Week 5, which include overlaps. In Week 5, the number of uninfected, infected and suspicious patient demands are 28, 4 and 2 on Monday (Day 1), and 22, 1 and 3 on Wednesday (Day 3). The number of patients assigned to each HD session is indicated on the figure. As shown in Fig. 8(b), on Day 1, since a larger number of machines are allocated to the standard unit in the 2-unit model, uninfected patients' dialysis treatments could be completed on the third HD session, and they only overlap with two Type 3 patients on that session. On the other hand, as shown in Fig. 8(a), due to the fewer dialysis machines allocated to the standard unit in the 3-unit case, Type 1 and 2 patients' treatments continue until the last HD session of the day, and a larger penalty is incurred due to several overlaps. In contrast, on Day 3, the allocation made by the 3-unit model results in a lower objective value (Fig. 8(c)). Since Type 3 and 4 patients are treated in parallel in the last HD session, they do not overlap with any uninfected patients in this case. However, as shown in Fig. 8(d), overlapping treatments had to be scheduled in the 2-unit solution.

We also evaluate the performance of the 2-unit cohorting strategy by solving the 2-unit scenario-based stochastic model, in which the uncertainties in inpatient demands are accounted for while deciding on the number of machines for an upcoming week. Table 7 presents the results of the 2-unit model, which are compared with those obtained by the 3-unit model for the same instances (Table 5). We observe that the 3-unit cohorting strategy results in lower expected penalties than those of the 2-unit cohorting strategy.

4.3. Discussion

Our case study results show that determining the best cohorting policy and capacity configuration can be challenging for HD clinics during a pandemic, even under the availability of perfect demand information. Furthermore, the uncertainties in demands for dialysis during a pandemic can significantly affect the performance of cohorting strategies. We show that developing demand forecasts based on historical data and solving the proposed scenario-based stochastic optimization models can help reduce the overlapping sessions of uninfected and infected patients thereby improving the effectiveness of cohorting. Using the methods proposed in this paper, HD clinics can evaluate the performance of alternative cohorting policies and obtain the best allocation decision based on the demand forecasts for the upcoming time horizon.

Our results also highlight the significant impact of allocation decisions on the number of overlaps between infected and uninfected patients. Evaluation of the case hospital's allocation policy showed that dialysis machines could be better allocated in the past, leading to fewer overlapping HD sessions if the allocation was optimized. Thus, we can infer that the hospital was overly conservative and reserved more buffer capacity for the infected and suspected patients to ensure a high service level to these patients. One can argue that this makes the system more resilient to surges in suspected and infected patient numbers. However, this cautious approach comes at the cost of increasing overlaps among patient cohorts. Our results show that stochastic optimization can be a valuable approach to effectively manage this trade-off.

Different cohorting strategies, in terms of the number of cohorts and the patient types in each cohort, can be modeled using the proposed optimization framework. Comparing 2-unit and 3-unit cohorting strategies for the case hospital showed that the 3-unit strategy was more effective overall by decreasing the more serious overlaps of the uninfected patients with suspected and infected patients. However, in general, the performance of 2- and 3-unit models depends on the number of patients throughout the planning horizon. For example, the 3-unit model performs well when there are not many uninfected patients, leading to the possibility of finishing their treatment earlier in the day. However, if the number of uninfected patients is also very high, this advantage of the 3-unit model will diminish. On the contrary, the 2-unit model can allocate more capacity to uninfected patients and manage to end the treatment of uninfected patients relatively earlier.

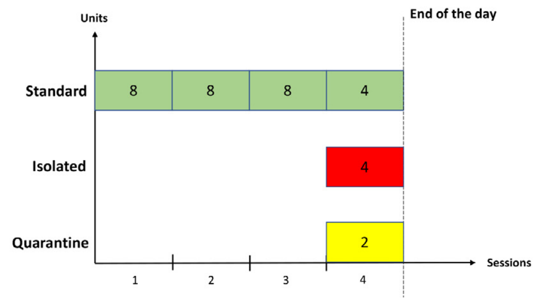

(a) Day 1, 3-unit model's schedule
(Daily penalty = 14,612)

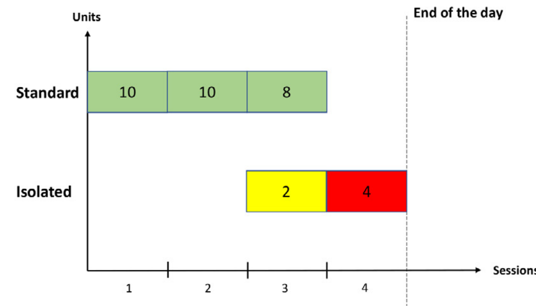

(b) Day 1, 2-unit model's schedule
(Daily penalty = 10,010)

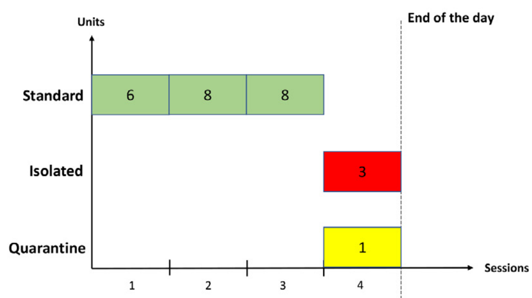

(c) Day 3, 3-unit model's schedule
(Daily penalty = 410)

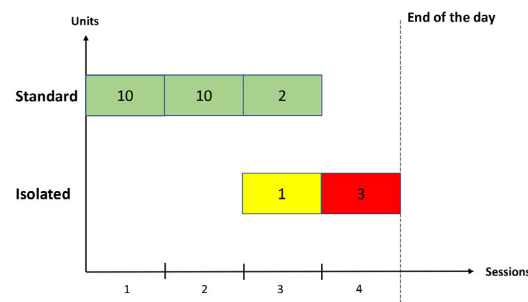

(d) Day 3, 2-unit model's schedule
(Daily penalty = 3,010)

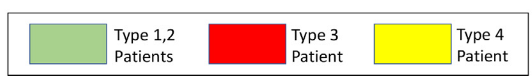

Fig. 8. Illustration of treatment schedules for Day 1 and Day 3 of Week 5.

Table 7
Performance of the stochastic optimization model for 2-unit cohorting.

		Optimal allocation for 2-unit cohorting (stochastic solution)					Overlaps with realized demand	
Week	PI (%)	R_1	R_2	$E[O_{(1,2)\times 3}]$	$E[O_{(1,2)\times 4}]$	Z_{2u}	$O_{(1,2)\times 3}$	$O_{(1,2)\times 4}$
6	80	9	5	0.0	4.7	4713	0	7
6	90	9	5	0.0	4.1	4112	0	7
7	80	9	5	0.9	5.9	6815	0	0
7	90	8	6	0.0	7.2	7214	0	0
8	80	9	5	0.0	2.2	2211	16	14
8	90	9	5	0.7	3.1	3811	16	14

The case hospital was able to treat all patients with the available dialysis resources since the beginning of the pandemic, albeit incurring some overlapping HD sessions. However, as the pandemic waves grow, it might be impossible to treat all patients. The proposed models can also be used to assess when such insufficiency of capacity might happen. In such a case, the proposed models can be used by policymakers to evaluate which patients would be better to transfer to other HD clinics, which is not currently allowed.

5. Conclusion and future research

Treating vulnerable groups such as people with chronic diseases needs special attention to prevent infection spread in health facilities during a pandemic. In this paper, we have focused on the operations of a hemodialysis (HD) clinic in a major pandemic hospital in Turkey during the COVID-19 pandemic, which serves multiple patient groups with different infection statuses by applying cohorting strategies. While similar HD facilities worldwide have also taken adequate precautions to manage different types of patients through cohorting, analytical tools have not been used yet to make cohorting plans. We present a two-stage stochastic programming modeling approach to make effective capacity planning and treatment scheduling decisions. We show with actual data collected during the COVID-19 pandemic that the clinic can use such an analytical tool to mitigate the infection transmission risk at the hospital by decreasing overlapping HD sessions among infected and uninfected patient groups. Moreover, we present results that compare two alternative cohorting strategies based on the hospital's data. We show that while 3-unit cohorting is generally more effective than 2-unit cohorting to reduce overlaps among patient cohorts in our case study, the performance of each strategy can be highly dependent on the relative number of each patient type. As clinical trials may not be allowed to test different practices, our proposed models can be valuable in testing different cohorting, capacity allocation, and scheduling approaches without harming patients, and identifying the best strategy to mitigate infection risk by using scarce resources efficiently in other countries and settings.

Given the scarce research on planning the treatments of patients with different infection statuses during the epidemics and pandemics, several future research directions exist. The non-medical operational solution proposed in this paper can be adapted to support cohorting, capacity planning, and treatment scheduling decisions in other clinics. Chemotherapy, radiotherapy, and physiotherapy patients may present different aspects and constraints. For instance, while dialysis treatment session durations are fixed, chemotherapy or physical therapy treatment durations may differ for each patient, which brings additional complexity for the scheduling problem. Another future research direction can be to incorporate uncertainty in patient arrivals during the day. We assumed that the demands from all patient types are known at the beginning of the day while scheduling treatments to minimize overlaps. Arrivals of patients with urgent treatment needs can be accommodated by reserving capacity for these patients, which is a common practice in appointment systems. Our stochastic optimization models can be extended in future research to account for such dynamics (e.g., by including such uncertainty in scenarios). Alternatively, different cohorting strategies can be analyzed with queueing models to obtain insights on their effectiveness in different environments.

Acknowledgments

This research has been supported by AXA Award Grant from AXA Research Fund. We sincerely thank hemodialysis nurse Ayşegül Işık for collecting data in the clinic to assist this study.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2021.10.039

References

Ahmadi-Javid, A., Jalali, Z., & Klassen, K. J. (2017). Outpatient appointment systems in healthcare: A review of optimization studies. *European Journal of Operational Research, 258*(1), 3–34.
Alberici, F., Delbarba, E., Manenti, C., Econimo, L., Valerio, F., Pola, A., et al. (2020). Management of patients on dialysis and with kidney transplant during SARS–COV-2 (COVID-19) pandemic in Brescia, Italy. *Kidney International Reports, 5*(5), 580–585.
Alharbi, A., Alharbi, S., & Alqaidi, S. (2020). Guidelines for dental care provision during the COVID-19 pandemic. *The Saudi Dental Journal, 32*(4), 181–186.
Basile, C., Combe, C., Pizzarelli, F., Covic, A., Davenport, A., Kanbay, M., et al. (2020). Recommendations for the prevention, mitigation and containment of the emerging SARS-CoV-2 (COVID-19) pandemic in haemodialysis centres. *Nephrology Dialysis Transplantation, 35*(5), 737–741.
Bikbov, B., Purcell, C. A., Levey, A. S., Smith, M., Abdoli, A., Abebe, M., et al. (2020). Global, regional, and national burden of chronic kidney disease, 1990–2017: A systematic analysis for the global burden of disease study 2017. *The Lancet, 395*(10225), 709–733.
Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming*. Springer.
Boushab, B. M., Koné, N., Baba, S. E.-W. O., El Mokhtar, M. J. O. M., Ahmed, M. M. O., Habiboullah, H. O., et al. (2021). Management of COVID-19 infection in patients undergoing hemodialysis. *The Pan African Medical Journal, 38*(13).
Chalabi, Z., Epstein, D., McKenna, C., & Claxton, K. (2008). Uncertainty and value of information when allocating resources within and between healthcare programmes. *European Journal of Operational Research, 191*(2), 530–539.
Chia, L., & Lin, W. D. (2015). Analytical framework to establish requisite bed capacity when cohorting paediatric patients. In *2015 IEEE international conference on industrial engineering and engineering management (IEEM)* (pp. 726–731). IEEE.
CMS (2020). Guidance for infection control and prevention of coronavirus disease 2019 (COVID-19) in dialysis facilities. *Technical Report*. Centers for Medicare & Medicaid Services.
Collison, M., Beiting, K. J., Walker, J., Huisingh-Scheetz, M., Pisano, J., Chia, S., et al. (2020). Three-tiered COVID-19 cohorting strategy and implications for memory-care. *Journal of the American Medical Directors Association, 21*(11), 1560–1562.
Corbett, R. W., Blakey, S., Nitsch, D., Loucaidou, M., McLean, A., Duncan, N., et al. (2020). Epidemiology of COVID-19 in an urban dialysis center. *Journal of the American Society of Nephrology, 31*(8), 1815–1823.
Currie, C. S. M., Fowler, J. W., Kotiadis, K., Monks, T., Onggo, B. S., Robertson, D. A., et al. (2020). How simulation modelling can help reduce the impact of COVID-19. *Journal of Simulation, 14*(2), 83–97.
Doornik, J. A., Castle, J. L., & Hendry, D. F. (2020). Short-term forecasting of the coronavirus pandemic. *International Journal of Forecasting*. In press
ERACODA Working Group (2021). Chronic kidney disease is a key risk factor for severe COVID-19: A call to action by the ERA-EDTA. *Nephrology Dialysis Transplantation, 36*(1), 87–94.
Fleming, R., Gartner, D., Padman, R., & James, D. (2019). An analytical approach for improving patient-centric delivery of dialysis services. In *Amia annual symposium proceedings: vol. 2019* (p. 418). American Medical Informatics Association.
George, M., Alexander, A., Mathew, J., Iyer, A., Waterval, J., Simon, C., et al. (2020). Proposal of a timing strategy for cholesteatoma surgery during the COVID-19 pandemic. *European Archives of Oto-Rhino-Laryngology, 277*, 2619–2623.
Gul, S. (2018). A stochastic programming approach for appointment scheduling under limited availability of surgery turnover teams. *Service Science, 10*(3), 277–288.
Hafner, K. (2020). Fear of COVID-19 leads other patients to decline critical treatment. Available at https://www.nytimes.com/2020/05/25/health/coronavirus-cancer-heart-treatment.htmlMay 25, 2020; Last Accessed: February 1, 2021.
Holland, J. (1994). Scheduling patients in hemodialysis centers. *Production and Inventory Management Journal, 35*(2), 76.
Hsu, C. M., & Weiner, D. E. (2020). COVID-19 in dialysis patients: Outlasting and outsmarting a pandemic. *Kidney International, 98*(6), 1402–1404.
Hulshof, P. J. H., Kortbeek, N., Boucherie, R. J., Hans, E. W., & Bakker, P. J. M. (2012). Taxonomic classification of planning decisions in health care: A structured review of the state of the art in OR/MS. *Health Systems, 1*(2), 129–175.
Ikizler, T. A., & Kliger, A. S. (2020). Minimizing the risk of COVID-19 among patients on dialysis. *Nature Reviews Nephrology, 16*(6), 311–313.
Klein, M. G., Cheng, C. J., Lii, E., Mao, K., Mesbahi, H., Zhu, T., et al. (2020). COVID-19 models for hospital surge capacity planning: A systematic review. *Disaster Medicine and Public Health Preparedness*, 1–17.
Kliger, A. S., & Silberzweig, J. (2020). Mitigating risk of COVID-19 in dialysis facilities. *Clinical Journal of the American Society of Nephrology, 15*(5), 707–709.
Kuhn, M., & Johnson, K. (2013). *Applied predictive modeling*: vol. 26. Springer.
Lano, G., Braconnier, A., Bataille, S., Cavaille, G., Moussi-Frances, J., Gondouin, B., et al. (2020). Risk factors for severity of COVID-19 in chronic dialysis patients from a multicentre french cohort. *Clinical Kidney Journal, 13*(5), 878–888.
Liu, Z., Lu, J., Liu, Z., Liao, G., Zhang, H. H., & Dong, J. (2019). Patient scheduling in hemodialysis service. *Journal of Combinatorial Optimization, 37*(1), 337–362.
Marchesi, J. F., Hamacher, S., & Fleck, J. L. (2020). A stochastic programming approach to the physician staffing and scheduling problem. *Computers and Industrial Engineering, 142*, 106281.
Mehrotra, S., Rahimian, H., Barah, M., Luo, F., & Schantz, K. (2020). A model of supply-chain decisions for resource sharing with an application to ventilator allocation to combat COVID-19. *Naval Research Logistics, 67*(5), 303–320.
Meijers, B., Messa, P., & Ronco, C. (2020). Safeguarding the maintenance hemodialysis patient population during the coronavirus disease 19 pandemic. *Blood Purification, 49*(3), 259–264.
Melman, G. J., Parlikad, A. K., & Cameron, E. A. B. (2020). Balancing scarce hospital resources during a pandemic: A case-study of COVID-19 at Addenbrooke's Hospital (UK). Working paper.
Naicker, S., Yang, C.-W., Hwang, S.-J., Liu, B.-C., Chen, J.-H., & Jha, V. (2020). The novel coronavirus 2019 epidemic and kidneys. *Kidney International, 97*(5), 824–828.
Ozturk, S., Turgutalp, K., Arici, M., Odabas, A. R., Altiparmak, M. R., Aydin, Z., et al. (2020). Mortality analysis of COVID-19 infection in chronic kidney disease, hemodialysis and renal transplant patients compared with patients without kidney disease: A nationwide analysis from turkey. *Nephrology Dialysis Transplantation, 35*(12), 2083–2095.
Park, H. C., Do Hyoung Kim, K. D. Y., Kim, Y.-G., Lee, S.-H., Yoon, H. E., Kim, D. K., et al. (2020). Korean clinical practice guidelines for preventing transmission of coronavirus disease 2019 (COVID-19) in hemodialysis facilities. *Kidney Research and Clinical Practice, 39*(2), 145.
Parker, F., Sawczuk, H., Ganjkhanloo, F., Ahmadi, F., & Ghobadi, K. (2020). Optimal resource and demand redistribution for healthcare systems under stress from COVID-19. Working paper.
Pinker, E., & Tezcan, T. (2013). Determining the optimal configuration of hospital inpatient rooms in the presence of isolation patients. *Operations Research, 61*(6), 1259–1276.
Punnakitikashem, P., Rosenberber, J. M., & Buckley-Behan, D. F. (2013). A stochastic programming approach for integrated nurse staffing and assignment. *IIE Transactions, 45*(10), 1059–1076.
Rais, A., & Viana, A. (2011). Operations research in healthcare: A survey. *International Transactions in Operational Research, 18*(1), 1–31.
Roper, T., Kumar, N., Lewis-Morris, T., Moxham, V., Kassimatis, T., Game, D., et al. (2020). Delivering dialysis during the COVID-19 outbreak: Strategies and outcomes. *Kidney International Reports, 5*(7), 1090–1094.
Rubin, R. (2020). Finding ways to reduce coronavirus exposure during dialysis. *Jama, 323*(20), 1993–1995.
Sever, M. S., Lameire, N., & Vanholder, R. (2009). Renal disaster relief: From theory to practice. *Nephrology Dialysis Transplantation, 24*(6), 1730–1735.
Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2014). *Lectures on stochastic programming: Modeling and theory*. SIAM Series on Optimization.
Smith, V. K. (2020). New York hospitals face another COVID-19 equipment shortage: Dialysis machines. Available at https://www.marketwatch.com/story/new-york-hospitals-face-another-covid-19-equipment-shortage-dialysis-machines-and-nurses-2020-04-16Last Accessed: April 16, 2020.

Sohrabi, C., Alsafi, Z., O'Neill, N., Khan, M., Kerwan, A., Al-Jabir, A., et al. (2020). World health organization declares global emergency: A review of the 2019 novel coronavirus (COVID-19). *International Journal of Surgery, 76*, 71–76.

Sperati, J. J. (2020). Coronavirus: Kidney damage caused by COVID-19. Available at https://www.hopkinsmedicine.org/health/conditions-and-diseases/coronavirus/coronaviruskidney-damage-caused-by-covid19May 14, 2020; Last Accessed: February 1, 2021.

Tanner, M. W., Sattenspiel, L., & Ntaimo, L. (2008). Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical Biosciences, 215*(2), 144–151.

The National Health Service (2021). Overview - dialysis. Available at https://www.nhs.uk/conditions/dialysis/Last Accessed: 1 February 2021.

Wallace, S. W., & Ziemba, W. T. (2005). *Applications of stochastic programming*. SIAM Series on Optimization.

Wang, X., Gong, X., Geng, N., Jiang, Z., & Zhou, L. (2020). Metamodel-based simulation optimisation for bed allocation. *International Journal of Production Research, 58*(20), 6315–6335.

Weissman, G. E., Crane-Droesch, A., Chivers, C., Luong, T., Hanish, A., Levy, M. Z., et al. (2020). Locally informed simulation to predict hospital capacity needs during the COVID-19 pandemic. *Annals of Internal Medicine, 173*(1), 21–28.

Whiteside, T., Kane, E., Aljohani, B., Alsamman, M., & Pourmand, A. (2020). Redesigning emergency department operations amidst a viral pandemic. *The American Journal of Emergency Medicine, 38*(7), 1448–1453.

Yarmand, H., Ivy, J. S., Denton, B., & Lloyd, A. L. (2014). Optimal two-phase vaccine allocation to geographically different regions under uncertainty. *European Journal of Operational Research, 233*(1), 208–219.

Yu, S., Kulkarni, V. G., & Deshpande, V. (2020). Appointment scheduling for a health care facility with series patients. *Production and Operations Management, 29*(2), 388–409.