INTEGRATED ONCOLOGIST CONSULTATION AND CHEMOTHERAPY APPOINTMENT COORDINATION UNDER UNCERTAINTY IN OUTPATIENT CHEMOTHERAPY CLINICS

by
EGE DİLAN ÖZYÜKSEL

Submitted to the Graduate School of Engineering and Natural Sciences
in partial fulfillment of
the requirements for the degree of Master of Science

Sabancı University
July 2023

INTEGRATED ONCOLOGIST CONSULTATION AND CHEMOTHERAPY APPOINTMENT SCHEDULING UNDER UNCERTAINTY IN OUTPATIENT CHEMOTHERAPY CLINICS

Approved by:

Prof. XXX YYY .....
(Thesis Supervisor)

Assoc. Prof. AAA AAA .....

Prof. BBB BBB .....

Date of Approval: July, 2023

ABSTRACT

INTEGRATED ONCOLOGIST CONSULTATION AND CHEMOTHERAPY APPOINTMENT SCHEDULING UNDER UNCERTAINTY IN OUTPATIENT CHEMOTHERAPY CLINICS

EGE DİLAN ÖZYÜKSEL

INDUSTRIAL ENGINEERING MSc. THESIS, JULY 2023

Thesis Supervisor: Prof. Dr. Tonguç ÜNLÜYURT
Thesis Co-Supervisor: Asst. Prof. Serhat GÜL

Keywords: Chemotherapy Scheduling, Oncologist Consultation Appointments, Care Coordination, Stochastic Programming, Scenario Reduction

The scheduling of chemotherapy treatments in outpatient chemotherapy clinics poses significant challenges due to limited resources, uncertainty in infusion durations, and the critical nature of cancer treatment. This study addresses these challenges through integrating oncologist consultation and chemotherapy scheduling by coordinating oncologist appointment times and chemotherapy treatment start times for a daily list of patients. A two-stage stochastic mixed-integer programming model is developed, considering continuous time frames for appointments and stochastic factors such as infusion times and the status of chemotherapy treatment approval. The first stage of the model arranges patients in a sequence based on their designated oncologists, while the second stage assigns patients to chairs and nurses. The objective function penalizes the expected weighted sum of the closing time of the chemotherapy clinic and patient waiting times. To reduce problem complexity, a scenario reduction algorithm is applied to the original scenario set prior to optimization. The proposed method, a Wasserstein Distance-Based Local Search Algorithm (WDB-LSA), is tested using real data obtained from a major academic oncology hospital in Turkey. The algorithm is compared with several practical heuristics from the literature using a commercial solver. The results demonstrate the effectiveness and computational efficiency of WDB-LSA in optimizing chemotherapy scheduling in outpatient clinics, taking into account multiple uncertainties and limitations. The

impact of varying model parameters was assessed under sensitivity analysis, and the solution methodology was tested against the mean value solution in order to estimate the value of the stochastic solution.

ACKNOWLEDGEMENTS

I would like to express my profound appreciation to my thesis co-supervisor Asst. Prof. Serhat Gül for his willingness and care to ensure an exceptional learning experience for me throughout this journey. I would also like to extend my genuine gratitude to my thesis supervisor Prof. Dr. Tonguç Ünlüyurt for his navigation when I needed it, and for sharing his precious time and intelligence with me. I could not have commenced this endeavor without their detailed and perpetual guidance with immense patience. I consider myself extremely lucky to have worked with such dedicated academics who have been actively supportive of this work, which is going to keep inspiring me endlessly in my personal life as well as in my professional life.

I am truly thankful to my mother for providing all the care and resources leading up to today. I always appreciate her unique ways of responding to my ways of expressing my distress. She somehow manages to still love me, to be the listening ear to my struggles, and to morally support me through my instabilities. I would also like to acknowledge my father's vibrant attention and his honest curiosity about my thesis study and many other undertakings of mine, which has encouraged me a lot to aim further after every conversation. I would also like to recognize my sister as an essential character in my life, who effortlessly alleviates any burden by only existing.

Lastly, I'd like to mention the support and faith of my friends who have been always there for me. I cannot thank these people enough for making me feel way smarter than I am and keeping me moving along the way.

This study is dedicated to my family and my favorite human, Brandon.
I must have done a world of good to receive such a tremendous amount of support, patience, care, and love from you along this rugged journey.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1. INTRODUCTION

Cancer is a disease that might affect a variety of organs caused by uncontrollable growth and the spread of abnormal cells, which commonly leads to death if proper treatment is not provided. Unfortunately, the risks of cancer apply to any age group and are expedited if the individual develops unhealthy habits such as smoking, alcohol abuse, and poor diet, and exposes themselves to many other environmental risk factors. According to the American Cancer Society, more than 16.9 million Americans were having a history of invasive cancer on January 1, 2019. Furthermore, approximately 1.9 million additional cancer cases were foreseen to be recorded in 2023 with 609,820 anticipated deaths in the United States, implying around 1,670 deaths per day (ACS, 2023). Cancer is a major public health concern in Turkey as well. The most recent estimates of the Global Cancer Observatory (GCO) database are that the age-standardized incidence rate of cancer in Turkey in 2020 was supposed to be 205.4 cases per 100,000 people.

Due to the avalanche of cancer diagnosis rates, the demand for outpatient chemotherapy clinics (OCC) from cancer patients is growing drastically (Haghi, Hashemi Doulabi, Contreras & Bhuiyan, 2022). An OCC is a healthcare institution where chemotherapy patients can undergo treatment without needing to stay overnight. While ongoing improvements in cancer treatments reduce death rates, maintaining and supporting survival is challenging but possible with effective management of chemotherapy processes in OCCs. Such institutions function with multiple resources and uncertainties to be coordinated effectively while constituting the appointment schedules for the chemotherapy patients (Corsini, Costa, Fichera & Parrinello, 2022; Gul, 2022) .

Since determining the days of treatment and appointment hours in an integrated manner is a very complicated problem (Benzaid, Lahrichi & Rousseau, 2020), this task is typically divided into two isolated problems (Cataldo, Sufan, Lorca, Andresen, Sánchez & Sauré, 2023; Lyon, Cataldo, Angulo, Rey & Sauré, 2023;

Ramos, Cataldo & Ferrer, 2020). These two main subproblems for chemotherapy treatment are the planning and scheduling phases. In the *Planning* phase, an inter-day scheduling problem is solved to determine the days of treatment to minimize the delay before the initiation of chemotherapy treatment throughout the long-term treatment plan. Once the days of treatment over a time horizon are determined, in the (intra-day) *scheduling* phase the daily patient sequences are determined including appointment start times or drug preparation and delivery times throughout a working day in an OCC. Since the number of patients assigned to a day is pre-determined, such problems do not have to consider patients' frequency or arrival rates. An additional division that is often combined with the scheduling phase is the *assignment* perspective. The assignment aspect considers the fairness and capacity issues for limited resources in the clinic by taking care of optimal resource-patient assignments (Hadid, Elomri, El Mekkawy, Jouini, Kerbache & Hamad, 2022).

As a patient is diagnosed with cancer, it is a critical matter that the chemotherapy starts within the maximum allowed delay period so that the treatment serves the purpose and stress-oriented adverse effects are prevented (Alexander, Blum, Burbury, Coutsouvelis, Dooley, Fazil, Griffiths, Ismail, Joshi, Love & others, 2017; Khorana, Tullio, Elson, Pennell, Grobmyer, Kalady, Raymond, Abraham, Klein, Walsh & others, 2019). This type of delay in chemotherapy treatment is related to the *chemotherapy planning* phase of the literature.

Outpatient chemotherapy clinics experience an extensive demand variety resulting in complex treatment flow schemes that are customized to each patient with varying needs by cancer types and disease backgrounds. Throughout chemotherapy treatment pathways, process durations, availability of limited resources, patient punctuality or unexpected fluctuations in health conditions commonly create an overwhelming uncertainty in OCCs (Hadid et al., 2022). For the *chemotherapy scheduling* phase, patient waiting time is also a crucial factor that affects the patient's chemotherapy outcomes.

The daily flow of a patient in an OCC follows a pathway that includes registration, lab testing, oncologist consultation, drug preparation, premedication, and infusion (see Figure 1.1). In practice, some OCCs conduct blood tests and vitals measurements on the day before the scheduled treatment day. According to this plan, the patients are urged to visit the clinic on two consecutive days with potentially more

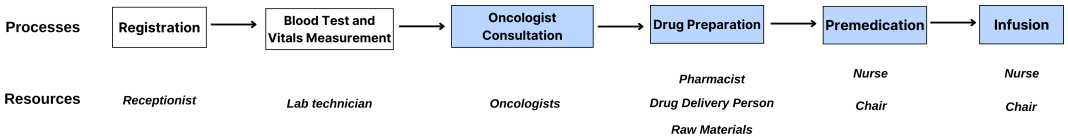

Figure 1.1 Processes and associated resources along a patient pathway in an OCC

distress. Noting this, Liang, Turkcan, Ceyhan & Stuart (2015) cared to arrange these procedures on the day of treatment in their studies. In a two-day visit scheme, oncologist consultation, drug preparation, premedication, and infusion constitute the critical steps of the second visit day. Oncologist consultation may be a bottleneck process depending on the duration of blood testing and vitals examination, which are rarely considered in chemotherapy planning and scheduling studies (Hadid et al., 2022). Patients wait for their chemotherapy treatment during the drug preparation phase. Premedication and infusion are the two fundamental consecutive steps of chemotherapy treatment. Patients utilize a nurse and a chair simultaneously during the premedication phase where they are prepared to prevent the side effects of chemotherapy drugs. Then, the patient receives the infusion drugs through a catheter and an IV bag under the surveillance of a nurse (Karakaya, Gul & Çelik, 2023). After the infusion step, the patient is discharged from the OCC.

No matter how many components a patient pathway is designed to have, the existence of multiple consecutive processes inevitably results in prolonged patient waiting times due to bottlenecks and prominent challenges in resource management (Hadid et al., 2022).

Changing the number of existing resources is not a preferred decision due to high recruitment costs for management. For that reason, effective management of healthcare resources in a limited capacity, in the long run, helps to serve more patients with reduced waiting times. Therefore the service quality and chances of survival enhance (Haghi et al., 2022). Commonly emphasized resources in a daily chemotherapy flow can be considered nurses and chairs. While oncologists and pharmacists are also incorporated occasionally, resources like receptionists, lab technicians, and drug delivery staff are usually overlooked in the literature. Raw materials for drug preparation are also considered as resources that are out of the scope of OCC resource management (Hadid et al., 2022).

The limited availability of nurses and chairs paves the way for the interdependence of resources (Vidal-Carreras, Garcia-Sabater & Marin-Garcia, 2022). Nurses are an

essential type of resource that can be simultaneously utilized by multiple patients for simple tasks such as monitoring multiple infusions. The complication arises from the fact that each patient utilizes both a nurse and a chair simultaneously throughout the entire treatment (Gul, 2022). This means that an outpatient clinic can simultaneously serve several patients even when the number of patients exceeds the number of available nurses, only if there is a satisfying number of chairs. However, a single nurse can perform premedications for only one patient at a time since absolute attention is required through the process (Lyon et al., 2023).

Based on the experience at Sir Charles Gairdner Hospital (a major tertiary academic hospital in Australia), for patients receiving treatment on the same day, approximately 20% of pre-arranged chemotherapy and infusion drugs were wasted due to cancellations or postponement of treatments (Lau, Watson & Hasani, 2014). Therefore, it is important that drugs are prepared on the day of treatment to prevent a remarkable amount of labor and drug-waste cost for such institutions. Although it may increase the waiting time for patients within the day of treatment, this procedure is widely implemented in many clinics. Fortunately, the advancement in automated chemotherapy drug preparation devices is expected to make this approach even more fetching (Hesaraki, Dellaert & de Kok, 2019). These systems are already adopted in different countries including Denmark, Germany, Italy, Japan, Spain, Turkey, and the United States (Masini, Nanni, Antaridi, Gallegati, Marri, Paolucci, Minguzzi & Altini, 2014).

The primary objective of an OCC is often associated with delay minimization. Typical performance measures consist of delay minimization attached to waiting times, clinic overtime, or chair idle times. Workload balance and overtime management for nurses are also key issues that promote employee satisfaction that are mostly related to cost minimization which is the secondary objective in an OCC. Patient waiting time is directly related to the satisfaction levels of patients (Gul, 2022). The total working time of a clinic (makespan) is also an alternative to nursing overtime, which is a frequently studied criterion in the related literature. The total working time of a clinic can be calculated as the final discharge time of the last patient on the daily appointment list. Although makespan minimization does not restrict the closing time of the care facility, it is helpful to optimize the utilization of resources (Heshmat, Nakata & Eltawil, 2018).

Other than the need for more than one resource at a time, patient journey and resource management in outpatient chemotherapy units are burdensome on account of uncertain parameters as well (Karakaya et al., 2023). Acknowledging treatment

durations as stochastic parameters renders imitating real-life instances easy. However, Haghi et al. (2022) remarks that literature still lacks stochastic approaches toward treatment times.

The importance of developing stochastic models for appointment scheduling is indisputable due to legitimate reasons. Assuming an expected value for infusion times is a venture due to unexpected circumstances. The duration of the chemotherapy treatment of a patient might take a few minutes on the lower extreme or a few hours on the upper extreme depending on the specified drug types, dosages, or the method of treatment. A patient might have to terminate the treatment due to their inability to tolerate the drugs resulting in an abnormally short infusion time. On the other end, any kind of complication caused by adverse effects may require additional time for revisions in medication content (Gul, 2022).

Inherently, the distribution for the infusion durations has a large variance. If a single point estimation was made for each patient, the dispersion of the data set results in either overestimated or underestimated values for a large portion of patients. While overestimating infusion times favors patient waiting times, it might also increase the total working time of the clinic. Conversely, underestimated infusion times are likely to prolong patient waiting times while the total working time is potentially shortened. Therefore, the decision maker should acknowledge the trade-off between patient waiting times and the total working time (Gul, 2022).

An overlooked source of uncertainty in an OCC is the status of treatment approval for patients depending on the results of blood tests and the progress in their course of treatment. Prior to chemotherapy administration, oncologists examine lab reports for the vitals and conditions of patients to decide if they should proceed with the treatment on the same day. If the lab results signal a potential health hazard that might arise due to treatment, the infusion must be postponed for about a week (Hesaraki et al., 2019). Since patient waiting times and the total working time of the clinic might be affected drastically due to deferrals on the day of treatment, consideration of the treatment approval uncertainty also helps to build reliable schedules.

The presence of uncertainty in both infusion times and treatment approval status arouses the need for an integrated appointment scheduling of treatment and consultation stages. The integration of both scheduling procedures has the potential to reduce patient waiting times and the total working time in OCCs. Due to the doses and types of medication prescribed and the treatment approval decision of oncologists, the outcomes of the consultation process have a remarkable impact on the patient flow through an OCC. Therefore, considering the treatment scheduling independent from the consultation scheduling facilitates the negligence of the in-

terdependency and leads to a myopic approach to treatment appointment schedule (Haghi et al., 2022).

This thesis study addresses the optimal patient scheduling and resource assignment challenges in OCCs with an aim to enhance patient satisfaction by reducing waiting times on treatment days. The patient pathway in this problem includes the oncologist consultation, drug preparation, premedication, and infusion processes. The limited resources are nurses and chairs that are seized simultaneously during premedication and infusion durations. Beyond the effort in patient waiting time minimization, the total working time of the clinic is also taken as a performance measure that is in conflict with the patient waiting time. An essential aspect of this research emphasizes the inherent variability in infusion durations. Furthermore, for patients who are not considered ready for an infusion treatment on the appointment day, the uncertain possibility of patient deferral is incorporated into the problem framework. Thus, the contributions of this thesis work include assessing the value of considering uncertainty in infusion durations and the status of treatment approval while integrating the scheduling of daily appointments for oncologist consultation and chemotherapy infusion processes in an OCC. The first stage of the stochastic TSMIP model determines precedence amongst patients within their oncologist groups. It sets appointment times, while the second stage assigns patients to a limited number of nurses and chairs. The uncertainty incorporated in the model is due to the possibility of treatment cancellations and stochastic infusion durations. The objective function minimizes the weighted sum of patient waiting times and the total working time (makespan) of the clinic. Next, a scenario reduction algorithm is implemented to represent the original scenario set by a smaller and more manageable representative scenario set. The results of our solution methodology are compared to the optimal value of the original model using a commercial solver. Furthermore, our approach is compared with several practical scheduling heuristics from the literature. Moreover, sensitivity analysis on several model parameters is conducted to generate managerial insights. Finally, the value of the stochastic solution (VSS) is estimated.

The remainder of this thesis is organized as follows. A review regarding similar models and solution methodologies is provided from the relevant literature in Section 2. Next, the stochastic TSMIP model details and assumptions are elaborated in Section 3. The scenario reduction approach as our solution methodology is introduced in Section 4, and computational results are discussed in Section 5. We conclude and discuss possible extensions of this work in Section 6.

2. LITERATURE REVIEW

In this section, we first review the related literature on deterministic chemotherapy appointment scheduling in Section 2.1. Then, we review the recent results for stochastic chemotherapy appointment scheduling in Section 2.2, which is the essential intent of this study. Finally, an extensive review of alternative scenario reduction methods in the literature is provided in Section 2.3. The comparison table in Figure 2.1 demonstrates a descriptive summary of various deterministic and stochastic two-stage models for chemotherapy scheduling problems.

2.1 Deterministic Chemotherapy Appointment Scheduling

Sadki, Xie & Chauvin (2011) concentrate on the two-stage outpatient appointment scheduling problem with a flow that follows oncologist consultation, drug preparation, and injection processes. In this work, nurses are not considered as limited resources, while the limitations are defined on oncologists with no idle time allowance and beds required for the infusion process. Makespan and patient waiting times are minimized using a Lagrangian relaxation-based heuristic. Turkcan, Zeng & Lawley (2012) combine this patient scheduling aspect of OCC studies with resource-patient assignments and treatment planning aspects. In a rolling horizon approach, treatment delays for patients and clinic overtime are minimized within a two-stage deterministic frame. Drug preparation and infusion processes are focal points while acuity level and treatment day tolerance complexities are resolved. Heshmat et al. (2018) approach the same problem with the intent to reduce the problem size using clustering algorithms to construct optimal patient groups for the first stage. The clusters are produced based on patient similarities such as duration of treatment, cancer type, or acuity levels. The latter stage is an improved extension of the mathematical model of Turkcan et al. (2012), in which every nurse is assigned to these

optimal patient clusters and chairs at optimum time slots while minimizing the total completion time of all treatments.

Alvarado & Ntaimo (2018) propose a variation on next-day chemotherapy scheduling in which they handle patient-nurse assignments and treatment day identification within the first stage. The researchers also incorporate nurse acuity levels into their model to minimize the surplus acuity and maximum allowed acuity level, along with the nurse overtime and deviation in treatment start time in the second stage. Apart from the drug infusion scheduling aspect in which the most emphasis is devoted; patient planning and patient-nurse assignment approaches are also considered simultaneously in their three-mean-risk stochastic integer programming model. Throughout the planning horizon, every day is divided into equal time slots.

2.2 Stochastic Chemotherapy Appointment Scheduling

An instance of two-stage stochastic integer programming models that focuses on the chemotherapy infusion process in an OCC is studied by Castaing, Cohn, Denton & Weizer (2016). The problem framework consists of a single nurse and multiple patients to coordinate next-day patient appointment schedules under stochastic infusion and preparation durations. A weighted sum of patient waiting times and expected total time spent throughout a day of treatment is minimized in the multi-objective function. The nurse and chair assignments are taken care of in the second stage, after the first stage decisions are made. On the other end, Demir, Gul & Çelik (2021) develop a two-stage functional care delivery model that schedules patients' appointments on the same day, under the limited availability of chairs and nurses with identical skills. Uncertainty in premedication and infusion duration of patients introduced stochasticity to their model and the patient sequence remains the same upon arrival at the clinic. The second stage assigns patients to nurses and chairs using the first stage outputs by penalizing chair idle times, nurse overtime, and patient waiting time in the clinic throughout the day. To maintain the workload balance among nurses, Gul (2021) makes subtle structural alterations to the outline of this same model by ensuring nurse-patient assignments in the first stage of their stochastic programming model. Corsini et al. (2022) develop an efficient Hybrid Harmony Search meta-heuristic to reduce the total flow time and patient waiting time within a similar problem setting with the idle time factor included as well. As a

	Paper	Processes				Decision		Limited Resources			Performance Measures	Uncertainty	Model	Method
		Consultation	Premedication	Drug Prep	Infusion	Patient Sequencing	Appointment Time Setting	Oncologists	Nurse	Chair				
Deterministic	Sadki et al. (2011)	Yes		Yes	Yes	Yes	Yes	Yes		Yes	Makespan, Waiting time		MIP	Lagrangean Relaxation Heuristic
	Turkcan et al. (2012)			Yes	Yes	Yes	Yes		Yes	Yes	Treatment Delay, Makespan, Overtime, Idle Time		IP	Heuristic
	Heshmat et al. (2018)				Yes	Yes	Yes		Yes	Yes	Total Working Time		IP	Clustering
Stochastic	Garaix et al. (2020)	Yes		Yes	Yes	Yes				Yes	Makespan, Overtime	Status of Treatment Approval	TSMIP	Greedy Heuristic
	Castaing et al. (2016)				Yes		Yes			Yes	Makespan, Waiting Time	Infusion Duration, Patient Preparation Time	TSMIP	Heuristic
	Alvarado and Ntaimo (2018)				Yes		Yes		Yes	Yes	Makespan, Overtime, Excess Nurse Acuity	Infusion Duration, Nurse Availability	Mean-Risk IP	Heuristic
	Demir et al. (2021)		Yes		Yes	Yes	Yes		Yes	Yes	Overtime, Chair Idle Time, Waiting Time	Infusion Duration, Premedication Duration	TSMIP	Progressive Hedging Algorithm
	Gul (2021)		Yes		Yes		Yes		Yes	Yes	Overtime, Waiting Time	Infusion Duration, Premedication Duration	TSMIP	Solver
	Corsini et al. (2022)	Yes		Yes	Yes	Yes		Yes	Yes	Yes	Flow Time	**Status of Treatment Approval, Consultation,** Drug Preparation, Drug Transportation, Setup, Treatment Times		Hybrid Harmony Search Metaheuristic
	Haghi et al. (2022)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Overtime, Waiting Time	Injection Duration	TSMIP, multi-TSP	Sample Average Approximation
	Karakaya et al. (2023)		Yes		Yes	Yes	Yes		Yes	Yes	Waiting Time, Excess Acuity, Nurse Overtime	Infusion Durations	TSMIP	Scenario Bundling Based Decomposition
	This Study	**Yes**		**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Total Working Time, Waiting Time**	**Infusion Duration, Status of Treatment Approval**	**TSMIP**	**Scenario Reduction**

Figure 2.1 Literature Review Comparison Table

distinction from the previously mentioned studies, all the stages of a chemotherapy process are modeled in a stochastic fashion. Deferrals and consultation times are incorporated as sources of uncertainty.

Garaix, Rostami & Xie (2020) study the same deferral concept, namely the permission status for the injection process of patients following the oncologist consultation. Their model constructs daily global patient sequences as in Demir et al. (2021)'s work while minimizing the makespan of the clinic. Only bed-patient assignments and patient waiting times are considered while the oncologist consultation duration is identical for all patients. Deferral times are obtained from a stochastic distribution with variable deferral probabilities for each scenario. The required number of time slots with identical lengths for drug preparation and infusion processes are deterministic parameters.

As the pioneer of studies that integrate consultation and treatment appointment scheduling, Haghi et al. (2022) present two separate two-stage stochastic programming models; one of which employs machine scheduling constraints for resource/patient assignments, and the other uses a multi-TSP formulation to minimize overtime and patient waiting times. Non-identical patient types are taken into consideration, and sample average approximation is used as the solution methodology.

Karakaya et al. (2023) penalizes the excess workload of nurses by assigning patients to nurses in the first stage of their TSMIP model for daily scheduling. Nurses in this framework are equipped with varying (non-identical) skills while patient acuity levels are considered as well. Appointment times are determined by assigning patients to time slots. The expected weighted sum of excess acuity levels, nurse overtime, and patient waiting time are minimized. Scenario bundles in adequate sizes are created to obtain near-optimal schedules using a scenario bundle-based decomposition algorithm. Gul (2022) approaches the same OCC TSMIP patient scheduling model by introducing flexibility to the nursing care delivery system. The alternatives for the system are broken down into fully flexible, partially flexible, and inflexible policies. Even though patients have their primary nurses to whom they are assigned, they may be paired with any nurse with adequate skills for them. However, alternative nurses may be paired with a restrained number of patients with upper and lower limits.

2.3 Scenario Reduction

The discrete scenario reduction concept is introduced by Dupačová, Gröwe-Kuska & Römisch (2003) as an alternative approach to approximate a discrete distribution with a smaller/subset distribution where the components of the sub-distribution must be selected arbitrarily from the set of components in the initial distribution. The authors present a discrete scenario reduction approach to diminish the computational complexity by extracting a subset of scenarios from the main distribution to compare the discrepancy between those two distributions later.

Sample average approximation (SAA) is one of the most prevalent approaches that function by generating random samples of scenarios that are equally likely to occur (Kleywegt, Shapiro & Homem-de Mello, 2002). Then, the sample average function approximates the expected value function. Despite the asymptotical optimality of the results and its significance in effectiveness, this approach may hinder the likelihood of obtaining absolute optimality for insufficiently small sample sizes (Zhang, Wang, Jacquillat & Wang, 2023). As an instance of this concept, Fei, Gülpınar & Branke (2019) work with an SAA framework and develop a two-stage heuristic solution methodology to their computational resource allocation model. A Wasserstein distance-based screening approach is used to measure the discrepancy between sampling measures and rank solutions to detect and promote potentially superior solutions for the simulation. The combination of the optimal computing budget allocation technique and the Wasserstein screening approach can produce the optimal number of replications for any potential solution and penalizes the selection of undesired solutions.

As an alternative to probability metrics for distance-oriented approaches, Wasserstein distance can be an eligible selection. Based on this idea, Rujeerapaiboon, Schindler, Kuhn & Wiesemann (2022) ground their scenario reduction discussions on this metric to approximate an original scenario distribution with a reduced scenario distribution for both discrete and continuous scenario reduction practices. Incorporating the Wasserstein distance, a prominent Local Search algorithm is used as a basis for a polynomial-time constant-factor approximation algorithm. Additionally, an exact mixed-integer programming (MIP) reformulation is developed for the scenario reduction stage. Within a similar frame, Bertsimas & Mundru (2022) develop a convex optimization-oriented alternating-minimization algorithm with an emphasis on the cost structure of the problem. The quality of decisions is also based on the Wasserstein distance between two discrete distributions. Authors construct

their ideas regarding their worst-case error-bound results based partially on Rujeerapaiboon et al. (2022)'s work. However, Bertsimas & Mundru (2022) differentiate their study from theirs by using a novel divergence technique rather than the typical Euclidean norm-based distance, and their study significantly outperforms the other most recent scenario-reduction-related methods in the literature. Their algorithm is influenced by Lloyd (1982)'s k-means clustering algorithm.

Out of the scope of healthcare applications, Abouelrous, Gabor & Zhang (2022) implement a methodology resembling Bertsimas & Mundru (2022) study for their two-stage stochastic inventory optimization problem to derive a cost-effective fulfillment policy of online and in-store demand. This problem has a sophisticated nature due to demand uncertainty and the computational complexity of the fulfillment procedure. Therefore, to overcome this combinatorial complexity, a novel proximity measure with their novel technique that combines Good-Turing sampling and linear programming is used to determine clusters of scenarios that represent randomly generated instances (Good, 1953). The fact that it does not require a pre-specified number of scenarios is the primary benefit of this framework and one of the two major distinctions from Bertsimas & Mundru (2022)'s work. The latter distinction is that scenarios can be grouped in the same cluster only if the proximity measure with the centroid of the cluster is less than the given threshold. Otherwise, if scenarios seem to be distant, the algorithm creates new clusters for them.

As opposed to scenario generation practices focused on distance and probability metrics with poor convergence, Prochazka & Wallace (2020) introduce a novel scenario-tree construction method for scenario generation in the problem-oriented realm. The heuristic compares and minimizes the incoherence of performances between in-tree and out-tree problem solutions, having this distance measure embedded in a loss function. What distinguishes their work from an accustomed version of this implementation is that a subset of feasible solutions is used rather than having the problem solved to optimality. Narum, Fairbrother & Wallace (2022) introduce a problem-oriented approach called singular value decomposition to observe the impact of various potential decisions on output distributions for stochastic programming. This is a generalized scenario reduction approach that is suitable even for significantly challenging distributions in two-stage stochastic models. Usually within the scope of scenario reduction practices, the number of scenarios required for reliable results cannot be known beforehand. Therefore, the novelty of Narum et al. (2022)'s approach is that their method proposes a suitable number of scenarios for any accuracy level provided. Fairbrother, Turner & Wallace (2022) propose a different extension to problem-oriented scenario generation approaches by considering problems with tail-risk measures, where the distribution of the stochastic input parameter has some

regions that do not have an impact on the cost function. Scenarios that fall into these regions are considered unnecessary to be included in the reduced scenario sets while their "aggregation sampling algorithm" that collapses other scenarios outside of the risk region into a single scenario point. Henrion & Römisch (2022) study the problem-oriented method in which a semi-infinite optimization problem is solved to obtain the best approximation of the main distribution by utilizing stability estimates based on problem-specific input. They eventually conclude that the method does not provide decent tractability since it is a generalized problem.

The scenario-grouping approach is often combined with common decomposition algorithms such as the progressive hedging algorithm (Crainic, Hewitt & Rei (2014), Escudero, Garín, Pérez & Unzueta (2013), Gade, Hackebeil, Ryan, Watson, Wets & Woodruff (2016), Jiang, Bai, Wallace, Kendall & Landa-Silva (2021)), Lagrangian decomposition algorithm (Escudero et al. (2013)) and L-shaped (Oliveira, Sagastizábal & Scheimberg (2011)) algorithm. In one of the most recent studies carried out by Karakaya et al. (2023), a two-step scenario bundling-based decomposition algorithm is implemented to create near-optimal schedules for their TSMIP model that investigates the impact of the flexibility in nurse care delivery system in an OCC framework. Gul (2022) also pursued the effort for the assessment of the flexibility factor within a similar TSMIP framework. As the solution methodology for their model, they propose variations of a scenario grouping-based decomposition (SGBD) algorithm in which each of the four variants has a different procedure to partition the main scenario set into smaller clusters of scenarios. Progressive-SGBD groups scenarios according to the proximity of single-scenario subproblem solutions instead of the proximity among stochastic parameters. Inspired by k-means clustering, input-based-SGBD determines a reference scenario at each iteration and adds scenarios to the cluster that are closest or furthest away to the centroid. Random-SGBD, on the other hand, groups scenarios randomly while the rest of the procedure is the same. Eventually, TSMIP gets to be solved for each scenario group, and the initial original problem is solved to examine first-stage solutions. Keutchayan, Ortmann & Rei (2023) also employ a similar approach that clusters scenarios into plausible-sized groups by replicating a smaller portion of the original problem as a reduced scenario cluster in a problem-oriented fashion. These smaller clusters are constructed in such a way that their objective values approximate the objective value provided by the original set of scenarios. Their approach considers the proximity between scenarios in terms of their cost function value instead of directly measuring the proximity between scenarios to find the best subset of scenarios that provides a decent approximation of the original problem. Despite having both studies used clustering-originated approaches, it is evident that Keutchayan et al. (2023)'s work

falls into the scenario reduction category whereas Gul (2022)'s contribution is an instance of scenario grouping-based decomposition algorithms, yielding even stronger results in terms of computational efficiency, even for problems with larger instance sizes.

Within the scope of problem-oriented scenario reduction, Zhang et al. (2023) formulate another novel scenario decomposition methodology to overcome limitations that exist in conventional SAA and distribution-based scenario reduction contexts. A scenario subset selection model is constructed as an MIP to generate high-quality solutions by optimizing stochastic recourse function approximations among possible first-stage solutions. To obtain tighter stochastic lower bounds, a scenario assortment optimization (SAO) is utilized so that smaller-scale stochastic models could be solved within smaller scenario bundles. As a proprietary contribution of this study, a column-evaluation and generation approach is proposed to obstruct the complexities that emerged from the SAO formulation by solving optimization problems with small confidence in the predictability of objective parameters.

3. THE CHEMOTHERAPY APPOINTMENT SCHEDULING PROBLEM

3.1 A TSMIP Model for the Chemotherapy Appointment Scheduling Problem

This section provides a TSMIP formulation for the integrated chemotherapy and consultation scheduling problem. This model attempts to coordinate oncologist consultations and chemotherapy treatments by determining daily patient sequences and appointment start times under the limited availability of nurses and infusion chairs. The stochastic nature of the problem stems from uncertainties in infusion durations and the status of treatment approvals after oncologist consultations. Decisions made in this model can be listed in four steps:

- Determining a sequence for a daily patient list
- Setting appointment times for patients
- Patient-chair assignments
- Patient-nurse assignments

In the first stage of the TSMIP, the model ensures that the patients are sequenced within their designated oncologist groups while the second stage is designed to assign patients to chairs and nurses. The objective function minimizes the weighted expected total patient waiting time and total working time over a sufficiently large number of scenarios. Scenario instances are created by sampling from real distributions of infusion durations and treatment approval probabilities. Our solution methodology to obtain near-optimal patient schedules ensures improvement in terms of computational performance. To reduce the problem size, a Wasserstein Distance-Based Local Search Algorithm (WDB-LSA) is introduced as a distribution-based discrete scenario reduction technique. Next, the algorithm is tested against the op-

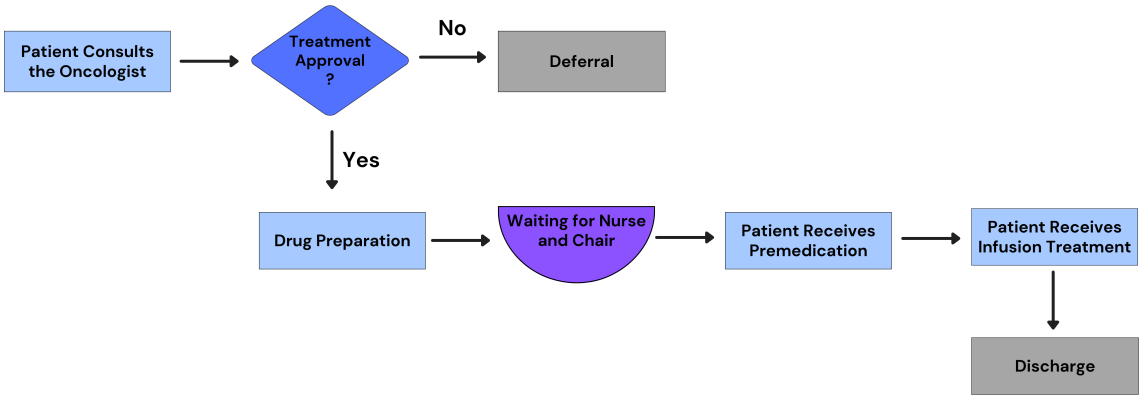

Figure 3.1 OCC Patient Flow considered in our stochastic TSMIP formulation

timal solution, practically relevant sequencing, and appointment-setting heuristics to advocate performance.

After the arrival at the OCC, a patient follows the illustrated pathway in Figure 3.1. Since the blood draw process is completed one day in advance, patients start their journey with an oncologist consultation without waiting in a queue. Patients are pre-assigned to existing oncologists before their arrival. This means that patient-oncologist assignment is not one of the intended efforts of the stochastic TSMIP model described in section 3.1.1.

Next, the oncologist examines the vitals, the condition of organs, and laboratory test results associated with the patient's blood sample to decide whether the patient is ready for the pre-planned treatment on this particular day. If the patient is not granted to proceed with the chemotherapy treatment, the patient leaves the system. Those patients are rescheduled to a future day by the head nurse, but this task is out of the scope of this study. If the treatment is approved, a prescription is delivered to the pharmacy for drug preparation. As soon as the custom-made drug mix is ready, the patient waits for an available chair and nurse. Finally, the patient seizes both resources at the same time throughout the premedication and infusion processes. Once the infusion treatment is completed, the patient is discharged from the OCC.

3.1.1 The Stochastic TSMIP Formulation

This section elaborates on the stochastic TSMIP model for the chemotherapy appointment scheduling problem. The assumptions of the model are as follows:

- Patient allocations to different days over a broader period in the realm of chemotherapy scheduling are already handled, meaning that the study only tackles identifying the daily sequence and appointment times of patients who are assigned to the same day for consultation and treatment.
- This study considers the processes and waiting times shown on the patient flow in Figure 3.1
- Waiting time calculations are assumed to be the total non-value added time. Therefore, drug preparation time is not included in the waiting time calculations since it is considered necessary (value-added) waiting time.
- Blood tests are taken care of before the day of appointment for every patient.
- The model does not consider time slots since it is a restrictive approach. In this model, patients are scheduled at any minute within the day to prevent excessive idle times.
- Functional care delivery system is assumed to be utilized in the OCC. This scheme allows the patients to be treated by any nurse available, meaning that nurses have identical skills.
- A nurse can administer only one premedication at a time.
- While conducting premedication, the same nurse can monitor multiple infusions simultaneously.
- All patients are assumed to be punctual.
- All patients have the same level of acuity/urgency.
- Regardless of the initial appointment sequence of the patients, patients who are done with the oncologist consultation can move on to the treatment stage, meaning that the sequence of the patients may change in the chemotherapy treatment phase depending on their oncologist discharge times.
- Premedication duration is assumed to be deterministic and equal to 15 minutes for each patient.
- Infusion duration is assumed to be stochastic.

Table 3.1 Notation used throughout the two-stage SMIP model

	Sets
I	Set of patients
C	Set of chairs
Θ	Set of oncologists
$I\Theta_\theta$	Set of patients of oncologist $\theta \in \Theta$, where $I\Theta_\theta \in I$
Ω	Set of scenarios
N	Set of nurses
	Parameters
s	Premedication duration
λ	Trade-off parameters in the objective function in the interval $[0, 1]$
M_1, M_2, M_3	A large value
k_i	Drug preparation duration for patient $i \in I$
o_i	Oncologist consultation duration for patient $i \in I$
t_i^ω	Infusion time of patient $i \in I$ in scenario $\omega \in \Omega$
β_i^ω	$= \begin{cases} 1, & \text{if patient } i \in I \text{ proceeds with their treatment after the oncologist consultation in scenario } \omega \in \Omega \\ 0, & \text{otherwise} \end{cases}$
	First-Stage Decision Variables
b_{ij}	$= \begin{cases} 1, & \text{if patient } i \in I \text{ precedes patient } j \in I \text{ in the daily oncologist appointment sequence.} \\ 0, & \text{otherwise} \end{cases}$
e_{ij}	$= \begin{cases} 1, & \text{if the oncologist discharge time of patient } i \in I \text{ is earlier than or equal to the oncologist discharge time of patient } j \in I \\ 0, & \text{otherwise} \end{cases}$
a_i	Appointment start time of patient $i \in I$
	Second-Stage Decision Variables
x_{in}^ω	$= \begin{cases} 1, & \text{if patient } i \in I \text{ is assigned to nurse } n \in N \text{ in scenario } \omega \in \Omega \\ 0, & \text{otherwise} \end{cases}$
y_{ic}^ω	$= \begin{cases} 1, & \text{if patient } i \in I \text{ is assigned to chair } c \in C \text{ in scenario } \omega \in \Omega \\ 0, & \text{otherwise} \end{cases}$
w_i^ω	Waiting time of patient $i \in I$ in scenario $\omega \in \Omega$
d_i^ω	Final discharge time of patient $i \in I$ in scenario $\omega \in \Omega$
f_i^ω	Treatment start time for patient $i \in I$ under scenario $\omega \in \Omega$ if the patient's treatment is approved
μ^ω	Total working time in scenario $\omega \in \Omega$

- Oncologist consultation duration for each patient is deterministic.
- Drug preparation duration is deterministic.
- Status of treatment approval is a binary stochastic parameter.
- There are sufficient number of pharmacists in our model, meaning that pharmacists are not considered a bottleneck resource. Therefore, the drug preparation process can be simultaneously carried out for multiple patients.
- Prior to their visits, patients have their own designated oncologists, and patients are allowed to consult only their oncologists.

$$(3.1)\quad min \quad \mathcal{Q}(\mathbf{a},\mathbf{b})$$

$$(3.2)\quad b_{ij}+b_{ji}=1 \qquad \forall i,j\in I\Theta_\theta, j>i, \forall\theta\in\Theta$$

$$(3.3)\quad a_j \geq a_i+o_i-M_1(1-b_{ij}) \qquad \forall i,j\in I\Theta_\theta, j\neq i, \forall\theta\in\Theta$$

$$(3.4)\quad e_{ij}+e_{ji}=1 \qquad \forall i,j\in I, j>i$$

$$(3.5)\quad a_j+o_j \geq a_i+o_i-M_2(1-e_{ij}) \qquad \forall i,j\in I, j\neq i$$

$$(3.6)\quad b_{ij}\in\{0,1\} \qquad \forall i,j\in I\Theta_\theta, j\neq i, \forall\theta\in\Theta$$

$$(3.7)\quad e_{ij}\in\{0,1\} \qquad \forall i,j\in I, j\neq i$$

$$(3.8)\quad a_i : integer \qquad \forall i\in I$$

where

$$\mathcal{Q}(\mathbf{a},\mathbf{b})=E_\xi[\mathrm{Q}(\mathbf{a},\mathbf{b},\xi(\omega))]$$

is the expected recourse function, and is given for each scenario $\omega\in\Omega$ by

$$\mathrm{Q}(\mathbf{a},\mathbf{b},\xi(\omega))=\min\left\{\lambda\sum_{i\in I}w_i^\omega+(1-\lambda)\mu^\omega\right\}$$

$$(3.9)\quad \sum_{n\in N}x_{in}^\omega=1 \qquad \forall i\in I$$

$$(3.10)\quad \sum_{c\in C}y_{ic}^\omega=1 \qquad \forall i\in I$$

$$(3.11)\quad f_i^\omega=a_i+o_i+\beta_i^\omega(k_i+w_i^\omega) \qquad \forall i\in I$$

$$(3.12)\quad d_i^\omega=f_i^\omega+\beta_i^\omega(s+t_i^\omega) \qquad \forall i\in I$$

$$(3.13)\quad f_j^\omega \geq f_i^\omega+s-M_3(5-e_{ij}-x_{in}^\omega-x_{jn}^\omega-\beta_i^\omega-\beta_j^\omega) \qquad \forall i,j\in I, j\neq i, \forall n\in N$$

$$(3.14)\quad f_j^\omega \geq d_i^\omega-M_3(5-e_{ij}-y_{ic}^\omega-y_{jc}^\omega-\beta_i^\omega-\beta_j^\omega) \qquad \forall i,j\in I, j\neq i, \forall c\in C$$

$$(3.15)\quad f_j^\omega \geq f_i^\omega-M_3(3-e_{ij}-\beta_i^\omega-\beta_j^\omega) \qquad \forall i,j\in I, j\neq i$$

$$(3.16)\quad \mu^\omega \geq d_i^\omega \qquad \forall i\in I$$

$$(3.17)\quad x_{in}^\omega\in\{0,1\} \qquad \forall i\in I, \forall n\in N$$

$$(3.18)\quad y_{ic}^\omega\in\{0,1\} \qquad \forall i\in I, \forall c\in C$$

$$(3.19)\quad d_i^\omega, w_i^\omega, f_i^\omega \geq 0 \qquad \forall i\in I$$

3.1.2 Explanations for the Stochastic TSMIP Formulation

The first-stage objective function (3.1) contains nothing but the expected second-stage objective function since none of the variables in the first stage has an immediate impact on the objective function value. Given the first-stage decisions, the expected second-stage recourse function minimizes the scenario average of the weighted sum of patient waiting times and total working time. Here, the parameter λ imposes a trade-off between the two performance measures in the objective function.

Constraints for the first stage are illustrated by the expressions (3.2) - (3.7). Constraints (3.2) provide the order of precedence between patients concerning the oncologist groups they are assigned to. For each patient subgroup associated with the oncologist θ, if patient $i \in I\Theta_\theta$ precedes patient $j \in I\Theta_\theta$, the corresponding b_{ij} value is equal to 1, and 0 otherwise. Constraints (3.3) coordinate the relationship between the appointment start times and the sequence of patients. Amongst the patients of the oncologist θ, if patient $i \in I\Theta_\theta$ precedes patient $j \in I\Theta_\theta$, then the appointment time of patient j should not be earlier than the summation of appointment time of patient i and their consultation duration since they are consulting the same oncologist. Since the chemotherapy treatment sequence is determined based on the order of oncologist discharge times of patients in our TSMIP model, we define a binary variable e_{ij} to represent that order. Constraints (3.4) ensure that either patient i precedes j in the oncologist discharge time sequence, or vice versa. Considering the complete list of patients, constraints (3.5) deal with coordinating the relationship between oncologist discharge times of patients. If e_{ij} is 1, it is expected that $a_j + o_j$ (the oncologist discharge time of patient j) is greater than or equal to $a_i + o_i$ (the oncologist discharge time of patient i). Otherwise, if the oncologist discharge time of patient i is not earlier than that of patient j ($e_{ij} = 0$), the constraint inherently becomes redundant due to the big M value. Constraints (3.6) - (3.7) indicate the binary restriction on the first-stage variables b_{ij} and e_{ij} while constraints (3.8) represent the integrality restriction on the first-stage variable a_i.

In order to obtain the assignment decisions through the second stage formulation, a subproblem must be solved for each scenario representing a different set of chemotherapy durations and treatment approval statuses. The expressions (3.9) – (3.19) stand for these second-stage constraints. Constraints (3.9) handle the assignment of each patient to an available nurse for every scenario while ensuring that a patient is assigned to only one nurse. Similarly, constraints (3.10) assign each patient to only one of the available chairs in every scenario. Since our model considers the binary state of treatment approval after the oncologist consultation, constraints

(3.11) rely on this approval status while calculating the treatment start time for each patient under each scenario. The binary stochastic second-stage parameter β_i^ω stands for the state of this approval. As long as patient $i \in I$ is approved to go under treatment in scenario $\omega \in \Omega$ (meaning that $\beta_i^\omega = 1$), the treatment start time of patient $i \in I$ equals the summation of their oncologist discharge time, drug preparation time and waiting time. A similar logic applies to constraints (3.12) in which the final discharge time of each patient is calculated. Likewise, if β_i^ω equals 1, the discharge time of patient $i \in I$ should be the summation of the premedication duration, infusion duration, and the treatment start time that was calculated in constraint (3.11). Under the opposite circumstance in which the beta parameter is 0, patient $i \in I$ would be discharged at their potential treatment start time since the treatment of the patient is not approved. If the oncologist discharge time of patient $i \in I$ is not later than that of patient $j \in I$, and patients $i, j \in I$ are assigned to the same nurse $n \in N$, and both patients are confirmed to go under treatment after the oncologist consultation, then constraints (3.13) ensure that the treatment time of patient $j \in I$ should not start earlier than the premedication finish time of patient $i \in I$. Similarly, constraints (3.14) ensure that if the oncologist discharge time of patient $i \in I$ is not later than that of patient $j \in I$, and patient $i, j \in I$ are assigned to the same chair $c \in C$, and both patients are confirmed to go under treatment after the oncologist consultation, then the treatment start time of patient $j \in I$ should not be earlier than the final discharge time of patient $i \in I$. Constraints (3.15) is formulated to determine the treatment start times according to the order of oncologist discharge times of patients. The constraints secure that the treatment start time of patient $i \in I$ is earlier than or equal to that of patient $j \in I$ as long as the oncologist discharge time of patient $i \in I$ is earlier than or equal to that of patient $j \in I$ when both patients are confirmed to go under chemotherapy treatment. For constraints (3.13), (3.14) and (3.15), failure to sustain any of the binary requirements inside the parentheses leads to redundancy due to the existence of big M values. Constraints (3.16) guarantee that the daily total working time in the OCC is later than or equal to the final discharge time of the last patient in the treatment sequence. Constraints (3.17)-(3.18) are binary restrictions over the assignment variables x_{in}^ω and y_{ic}^ω, and (3.19) ensure non-negativity for the rest of the second-stage variables.

Patient assignments to nurses and chairs are handled in the second stage after the sequencing and appointment time decisions are obtained since the opposite approach might be too restrictive for the appointment times and result in inefficient schedules. In cases of real-life applications, patient-to-resource assignment decisions are likely to be made dynamically right as patients arrive prior to the full realization of the uncertainty on the infusion durations (Demir et al., 2021). In the corresponding

multi-stage stochastic programming (SP) model, the decision stages would represent the time points at which the nurses and chairs become available. However, a particular nurse and chair assignment policy that we impose in our model renders the formulation of a multi-stage SP model unnecessary. In particular, we enforce a myopic resource assignment policy through constraints (3.15). The policy ensures that the assignment of patients to an available chair and nurse, and hence the treatment start times of patients only depend on the order of oncologist discharge times, which is determined through the values of first-stage variables. In other words, among the patients waiting in the OCC, an available nurse and chair are assigned to the patient who is done with their oncologist consultation at the earliest time. Therefore, realizations of uncertain infusion durations of the patients who are not treated yet are irrelevant when patient-to-chair-nurse assignment decisions are made. This implies that the patient-to-resource assignment decisions can be made independently for each scenario by only considering the values of e_{ij} variables. Consequently, our TSMIP model would provide the same solution as the multi-stage SP formulation of our problem.

According to a likely inaccurate inference to be made from the model at first glance, infusion times seem to be revealed before patient-resource assignments. However, the uncertain infusion times are in fact realized randomly after the assignment is completed, which is in line with practice and the discrete event simulation concept. It is essential to note that this simple assumption is realistic and applicable in real life since a head nurse cannot optimize the treatment sequence instantaneously at each time the resources become available by taking uncertain infusion durations into account.

To provide an example patient flow through the oncologist clinic and OCC, we provide an instance of a realization of a schedule for a case having 2 oncologists, 6 patients, 3 chairs, and 2 nurses. The Gantt chart 3.2 for our TSMIP model, reveals the patient appointment sequence for oncologist consultation as $P_3 = P_4 < P_5 < P_2 < P_6 < P_1$. Unlike the case in Garaix et al. (2020), since there is no globally fixed sequence assumption in our model, it is clearly visible that the patient sequences change after the consultation as $P_4 < P_5 < P_2 < P_1 < P_6$. The chart also shows that the sequence of patients from different oncologist groups is independent of each other.

In this particular scenario presented in Figure 3.2, the only type of waiting time under consideration is between the end of drug preparation and the beginning of premedication. The chart depicts the premedication process (PM) (15 minutes) and the infusion process for each patient P_i along the same bar without an interruption,

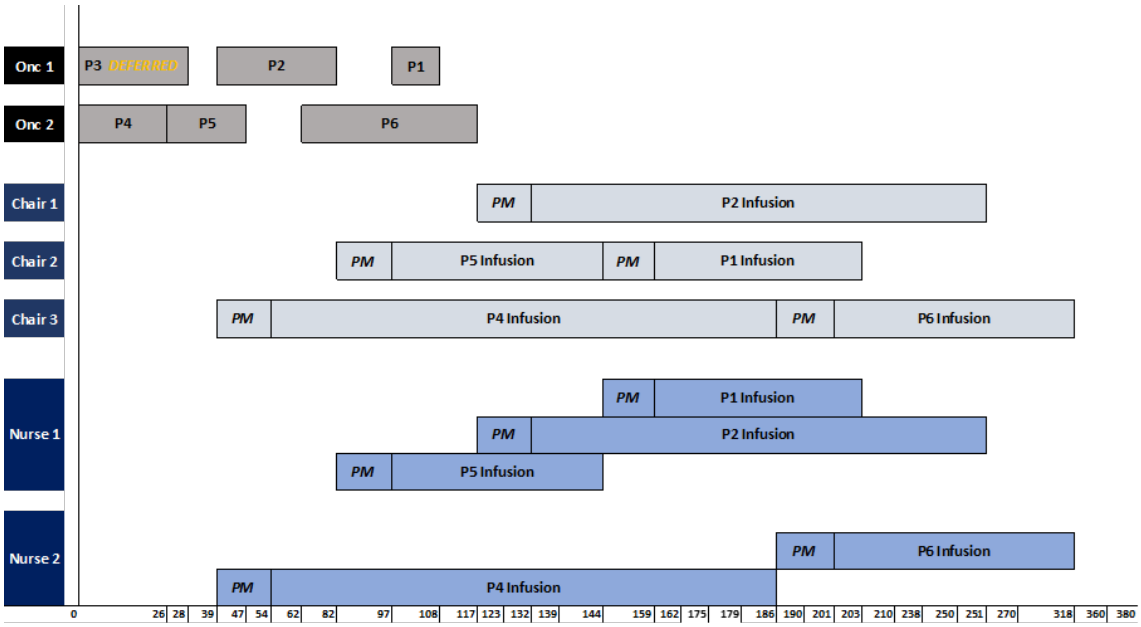

Figure 3.2 Gantt Chart for the OCC Patient Flow with a 2 oncologist, 6 patient, 3 chair, 2 nurse setup for $a_1 = 97$, $a_2 = 39$, $a_3 = 0$, $a_4 = 0$, $a_5 = 26$, $a_6 = 62$ (in minutes).

since patients seize both a chair and a nurse simultaneously throughout the two processes. The drug preparation process is not shown here since there are a sufficient amount of pharmacists in the OCC. Therefore, drug preparation can be carried out simultaneously for all patients if required. It was stated that nurses might be used by several patients simultaneously under convenient circumstances. As can be seen in Figure 3.2, a particular nurse can start administering premedication drugs to another patient only after the premedication of the previously assigned patient is completed. It can also be seen that Nurse 1 can start the premedication of Patient 2 while still monitoring Patient 5. Then, nurse 1 starts administering premedication for Patient 1 while monitoring Patient 2, and simultaneously monitors both until discharge.

The Gantt chart shown in 3.2 reveals that the oncologist discharge time of Patient 6 is at time 117, and the treatment starts for the same patient at time 186. This results in a 69-minute time gap in between for Patient 6, which also comprises the 13-minute drug preparation process. Therefore, since drug preparation is considered a value-added process, the actual waiting time for Patient 6 can be reported as 56 minutes. Similarly, Patient 1 leaves the oncologist consultation at time 123 and the treatment for this patient begins at time 144, leading to a 21-minute waiting time. However, since the drug preparation for Patient 1 is 15 minutes, the actual waiting time for this patient can be reported as 6 minutes. Therefore, the total waiting time in this particular scenario realization is recorded as 62 minutes. Meanwhile,

the total working time (makespan) in the OCC is realized as 318 minutes, which corresponds to Patient 6's treatment completion time.

3.1.3 Big M Calculations

Selection of a sufficiently large value for *big M* parameters is a delicate matter. The big M parameter should be as small as possible to remain conservative. A small enough big M value restricts the feasible solution so much so that the computational complexity can be reduced. Nevertheless, parameters M_1, M_2, and M_3 are required to be large enough since it is desired to help violate constraints they are used in, if desired requirements for those specific constraints to be valid are not satisfied. A large enough parameter ensures that the feasible region is not cut down further than necessary, which may result in eliminating a part of the feasible region that might have a potentially optimal solution. In other words, approximating the original optimal value of the problem is not ensured if big M is set to a too-small value. On the other hand, too-large M values may decrease both the precision and the numerical stability of the solution (Cococcioni & Fiaschi, 2021).

In order to obtain the most conservative big M calculations for our constraints in the stochastic TSMIP model, we assume a single chair, single oncologist system. Nurses are not included as a resource in calculations and all β_i^ω parameters are assumed to be 1. We use the maximum values of each duration parameter. Chairs are bottleneck resources since they are utilized during both premedication and infusion, which is the longest process in the clinic.

Constraints (3.3) concern the appointment start times of patients in a sequence. The maximum infusion duration for a patient is known to be 217 minutes relying on the data collected. We do not consider drug preparation duration in the calculation since this process does not utilize a physical resource. Premedication duration is deterministic and constant for each patient and equal to 15 minutes. Big M must be equal to the appointment start time of the last patient and we assume there are 9 patients assigned to a single day in all our experiments. In the worst-case situation, a patient would occupy a chair as long as the summation of premedication and infusion durations, which is $15 + 217 = 232$ minutes per patient. We can accept this duration as the upper bound on the time when the next patient on the list can start their appointment. The reason for this logic is that, since the chair is the bottleneck resource, no matter how long the waiting, drug preparation, and oncologist consultation times are, patients will be ready to occupy the chair after

the previously assigned patient leaves the chair. 8 x 232 will give us the latest appointment of the first 8 patients, which is also the appointment start time of patient 9, the last patient assigned on the same day. Hence, the big M value for constraints (3.3) is calculated as $M_1 = 1856$.

Constraints (3.5) concern the relation between oncologist discharge times of patients based on their sequencing. The maximum oncologist consultation duration possible is 30 minutes for a single patient. Since we have already calculated the latest appointment time of the 8th patient as 1856, the maximum oncologist discharge time for the last patient is going to be 1856 + 30 = 1886. Hence, the big M value for constraint (3.5) is set as $M_2 = 1886$.

Constraints (3.13), (3.14), and (3.15) are related to treatment start and final discharge times of patients. The only parameter excluded from the previous big M calculations is the drug preparation duration with a maximum value of 35 for a patient. Adding this value to the maximum oncologist discharge time of the last patient results in the treatment start time of the last patient in the system. Therefore, the common big M value calculated for constraints (3.13), (3.14), and (3.15) is $M_3 = 1921$.

4. SOLUTION METHODOLOGY

Future uncertainty is formulated and introduced in TSMIP models as stochastic second-stage constraints to obtain strategic decisions in the first stage of the problem (Bertsimas & Mundru, 2022). Thus, expected cost calculations in such stochastic models are often a numerically non-practical matter since a large number of scenarios is required to realistically simulate the long-run performance of any system under future uncertainty. Additionally, the inclusion of numerous parameters and variables is a hindrance to computational complexity. This reveals the need for a smaller subset of scenarios to represent the entire scenario set. In this study, it is not possible to solve the TSMIP model to optimality for a sufficient number of scenarios in a reasonable amount of time due to the computational burden. Therefore, we implement a scenario reduction methodology to solve our large instance problem by combining a distance measure called the Wasserstein metric and a typical Local Search Algorithm for k-medians clustering from the literature.

Section 4.1 details the goals and branching concepts in the scenario reduction approach from other studies. Section 4.2 provides the structure of the aggregation of the Wasserstein distance metric and the Local Search Algorithm (WDB-LSA) as the solution methodology.

4.1 Scenario Reduction

The scenario reduction realm of the literature, pioneered by Dupačová et al. (2003), is a fundamental step for solving stochastic programs for many reasons, one of which being the fact that the size of the scenario space has an impact on the instance size of a problem. One other goal of scenario reduction is to determine the scenarios and their probabilities that produce the tightest bounds possible while the gap is approximately zero (Keutchayan et al., 2023) and produce optimal or near-optimal

decisions with a reduced scenario space to increase tractability and interpretability of continuous distributions by discretizing them, or by reducing the size of an already discrete distribution (Narum et al., 2022).

Stochastic programming models are commonly solved by approximating the underlying probability distribution $\mathbb{P}$ by a discrete probability distribution $\mathbb{R}$ with fewer number of scenarios, where each scenario appears with a certain probability. This reduced distribution R must resemble the original distribution $\mathbb{P}$ so much so that the solution of stochastic mathematical models does not deviate significantly. Approximating a distribution by another distribution often means that a distance calculation between two distributions is required. Although the numerical calculation of such a distance is not always straightforward, obtaining estimates of $D_l(\mathbb{P}_{\ltimes},\mathbb{R}_{\gtrdot})$ by other distances of simpler distributions is possible. Laying this knowledge as the groundwork, optimal scenario reduction can be described as the determination of the best approximation of $\mathbb{P}$ by distribution $\mathbb{R}$, according to the distance of probability measures Römisch (2009).

The discrete scenario reduction concept is introduced by Dupačová et al. (2003) as an alternative approach to approximate a discrete distribution with another discrete distribution where the components of the sub-distribution must be selected arbitrarily from the set of components in the initial distribution. Dupačová et al. (2003) present a discrete scenario reduction approach to diminish the computational complexity by extracting a subset of scenarios from the main distribution to compare the discrepancy between those two distributions. While forming a reduced scenario set under discrete scenario reduction, any scenario to be selected from the initial distribution must be exactly one of the existing scenario points. This restriction is relaxed in continuous scenario reduction, meaning that new scenarios can be selected without being limited to only discrete points in the initial distribution, yielding more flexible and superior approximations. This also means that the scenarios derived from the initial distribution are mostly unforeseeable, especially for a distribution with non-convex or multi-partite support. Therefore, the discrete scenario reduction occupies a remarkable place in literature for stochastic programming (Rujeerapaiboon et al., 2022). Continuous scenario reduction might also be referred to as scenario generation in related literature (Löhndorf, 2016).

Discrete scenario reduction can be formulated as below, where n is the cardinality of $\mathbb{P}$ and m is the cardinality of $\mathbb{R}$.

$$\mathbb{D}_W(\mathbb{P}_n, m) = min_{\mathbb{R}}\left\{ D_W(\mathbb{P}_n, \mathbb{R}) : \quad \mathbb{R} \in P(\{\xi^1 \dots \xi^n\}, m) \right\}$$

For such computationally challenging scenario reduction methods, two main approaches can be classified as distribution-oriented and problem-oriented. While distribution-oriented methods make use of probabilistic distance measures, the solution to the original problem itself can be observed as a metric for problem-oriented methods. The problem-oriented scenario reduction method holistically treats the problem by comparing the ultimate cost function and parameters of the problem for different scenarios and works with scenario-based output data instead of input data of a probability distribution. This may mean that problems with the same fundamental distribution and different objective function values are not formulated using the same set of scenarios (Keutchayan et al., 2023; Narum et al., 2022). Keutchayan et al. (2023) suggest that the problem-oriented method is more useful and powerful for problems with larger distribution sizes and thus higher complexity, to obtain tractable and stable results. However, Keutchayan et al. (2023) also argues that problem-oriented approaches are not as common as distribution-oriented approaches since their results are not practical enough to systematically generate scenarios for problems. Furthermore, the distribution-oriented methods are compatible with a vast variety of algorithms such as Monte Carlo sampling or clustering algorithms. Other than its simplicity, our study also steers towards distribution-oriented methods taking their practicality into account.

Since this study adopts a distribution-oriented approach, only the distribution for the scenario space is considered and the cost function and constraints are neglected while generating the best reduced discrete distribution $\mathbb{R}$. Distribution-oriented methods often make use of probabilistic proximity measures such as Wasserstein (Bertsimas & Mundru, 2022; Ketkov, 2023; Pflug & Pichler, 2015; Rujeerapaiboon et al., 2022) Lévy–Prokhorov or Hellinger (Vidyashankar & Xu, 2015) distance metrics. Our solution method adopts the Wasserstein metric for the distribution-oriented discrete scenario reduction concept. This methodology approximates the initial distribution $\mathbb{P}$ by another distribution $\mathbb{R}$ in a reduced size. Each scenario point in $\mathbb{P}$ has a constant realization probability p_i, where $i \in 1, \ldots, n$ and $n > 0$. This approach intends to produce a proximate discrete distribution $\mathbb{R}$ that contains scenarios with constant realization probabilities q_j, where $j \in 1, \ldots, m$ and $n > m > 0$. As the proximity between $\mathbb{P}$ and $\mathbb{R}$ increases, it is expected that the solutions associated with those scenario sets get closer to each other in terms of quality. The trade-off, in this case, is that the computational burden significantly diminishes as the n-point distribution $\mathbb{P}$ is replaced with the new m-point distribution $\mathbb{R}$.

4.2 Wasserstein Distance Based Local Search Algorithm (WDB-LSA)

The remainder of this section provides the Wasserstein Distance formulation in Section 4.1, and the Local Search Algorithm in Section 4.2.2.

4.2.1 A Proximity Measure Between Distributions: Wasserstein Distance

- $\mathbb{P}$: A discrete probability distribution with scenario points $\left\{\xi_1, \ldots, \xi_n\right\}$
- $\mathbb{R}$: A reduced discrete probability distribution with scenario points $\left\{\zeta_1, \ldots, \zeta_m\right\}$
- p_i : Probability of scenario ξ_i, where $i \in \{1, \ldots, n\}$ and $p_i \in [0, 1]$
- q_j : Probability of scenario ζ_j, where $j \in \{1, \ldots, m\}$ and $q_j \in [0, 1]$
- π_{ij} : The amount of probability transferred from scenario ξ_i to scenario ζ_j, *where* $\sum_{i=1}^{n} \sum_{j=1}^{m} \pi_{ij} = \sum_{i=1}^{n} p_i = \sum_{j=1}^{m} q_j = 1$

(4.1)

$$D_l(\mathbb{P}, \mathbb{R}) = \left[\min \left\{ \sum_{i=1}^{n} \sum_{j=1}^{m} \pi_{ij} \|\xi_i - \zeta_j\|^l : \sum_{j=1}^{m} \pi_{ij} = p_i, \forall i = \{1, \ldots, n\} \right. \right.$$
$$\left. \left. \sum_{i=1}^{n} \pi_{ij} = q_j, \forall j = \{1, \ldots, m\}, \quad \pi_{ij} \geq 0 \right\} \right]$$

The joint probability distribution for Wasserstein distance (Rujeerapaiboon et al., 2022) can be modeled and optimized as a transportation problem that aims to transfer probabilities from an n-point origin $\mathbb{P}$ to an m-point destination $\mathbb{R}$ with the minimum transportation cost possible. The unit transportation cost between two scenarios of different distributions can be calculated as $\|\xi_i - \zeta_j\|^l$.

The expression (4.1) minimizes the Wasserstein difference between two distributions where $\|\xi_i - \zeta_j\|^l$ represents the l-norm of scenario $\zeta_i \in \mathbb{R}$ and scenario $\xi_i \in \mathbb{P} \setminus \mathbb{R}$. In

this study, we consider the $l2$-norm for the distance vector, and we assume that the realization probabilities of all scenarios are equal.

4.2.2 Local Search Algorithm

Both Rujeerapaiboon et al. (2022) and Bertsimas & Mundru (2022) have previously modified Arya, Garg, Khandekar, Meyerson, Munagala & Pandit (2001)'s local search algorithm for k-median clustering. The authors of both studies performed and discussed several computational experiments on various branches of scenario reduction approaches. The structure of the same algorithm is tailored to this study's requirements and the Wasserstein metric is considered for distance calculations.

For the initialization step of WDB-LSA, an initial reduced scenario set, $\mathbb{R}$, is defined with size m, where $m < n$. Besides, the initial reduced scenario set is marked as the best reduced set, $\mathbb{R}'$. Next, using the distance matrix, the algorithm calculates the Wasserstein distance between $\mathbb{R}$ and $\mathbb{P}$. The distance matrix is built to store the Euclidean distance values between every $\xi_i \in \mathbb{P}$ and $\zeta_j \in \mathbb{R}$ pair in a two-dimensional array.

Subsequently, the algorithm creates a candidate reduced scenario set by selecting a scenario $\zeta_j \in \mathbb{R}$ and swapping that with a scenario $\xi_i \in \mathbb{P} \setminus \mathbb{R}$. Then, the Wasserstein distance between the candidate reduced set and $\mathbb{P}$ is calculated and compared to the distance between $\mathbb{R}$ and $\mathbb{P}$. If the candidate reduced set results in a smaller distance value from $\mathbb{P}$, then the best reduced set is updated as the candidate reduced set. Otherwise, the best reduced set remains the same. The algorithm checks all possible candidate reduced sets by swapping $\zeta_j \in \mathbb{R}$ with all $\xi_i \in \mathbb{P} \setminus \mathbb{R}$. When a new swap is not possible, the algorithm chooses a new scenario, ζ_k, from $\mathbb{R}$ instead of ζ_j. Next, similar swap operations are conducted by replacing ζ_k with all possible scenarios in $\mathbb{P} \setminus \mathbb{R}$ to generate candidate reduced scenario sets. The algorithm terminates after all such candidate reduced scenario sets are checked, and the one with the smallest distance from $\mathbb{P}$ becomes the final reduced scenario set (i.e., best reduced set, $\mathbb{R}'$) representing the original scenario set.

The initial true probability distribution dominates the uncertain problem parameters. Therefore, when comparing the distance between scenarios (approximating the true distribution), the values to be compared are stochastic model parameters since they vary in each scenario. In our TSMIP model, the diversity and uncertainty are mainly originated from infusion duration parameter t_i^{ω}. However, our model also

consists of a binary variable β_i^ω for the status of treatment approval of each patient in each scenario. Since the revealed binary value of β determines whether a patient receives treatment or not, it also determines if the corresponding infusion duration is in the picture. Therefore, these two essential parameter arrays are multiplied with each other to form a new consolidated parameter for distance calculations. The resulting parameter inherently contains either the value 0 for patients without treatment approval or contains the corresponding infusion durations for patients with approval in each scenario.

Algorithm 1 Local Search Algorithm

Initialize: $\mathbb{R} = \{1, \ldots, m\}$
$\delta_w =$ Wasserstein distance between $\mathbb{P}$ and $\mathbb{R}$.
    $\delta_{bestfit} = \delta_w$
if $\gamma \in \mathbb{R}$ **and** $\kappa \notin \mathbb{R}$ **then**
    $\delta_{candidate} = \min_{\gamma \in \mathbb{R}, \kappa \in \mathbb{P} \setminus \mathbb{R}}$ Wasserstein distance between $\mathbb{P}$ and $\mathbb{R} \cup \{\kappa\} \setminus \{\gamma\}$
    $\gamma, \kappa = \arg \min_{\gamma \in \mathbb{R}, \kappa \in \mathbb{P} \setminus \mathbb{R}}$ Wasserstein distance between $\mathbb{P}$ and $\mathbb{R} \cup \{\kappa\} \setminus \{\gamma\}$
    if $\delta_{candidate} < \delta_{bestfit}$ **then**
        $\delta_{bestfit} = \delta_{candidate}$
        update $\mathbb{R}' = \mathbb{R} \cup \{\kappa\} \setminus \{\gamma\}$
    end if
end if
Repeat until no further improvement on best-fit

This algorithm may outline that the discrete scenario reduction problem accompanied by the Wasserstein metric the same as the k-median clustering problem (Rujeerapaiboon et al., 2022).

While exchanging single scenarios across sets, $\mathbb{R}$ and $\mathbb{P} \setminus \mathbb{R}$, we search for the swap that results in the best improvement in terms of distance from the original scenario set $\mathbb{P}$. This search strategy is called as the the *best-fit* strategy. An alternative strategy, *first-fit*, stops swapping scenarios right after an improvement in distance value is obtained. Even though the latter strategy would be inferior to former strategy in terms of approximating the original scenario set, it is obviously associated with lower computational time.

5. COMPUTATIONAL EXPERIMENTS

This section outlines the course of computational experiments and provides inferences of the comparisons for the integrated oncologist and chemotherapy scheduling problem.

The experiments are performed on a problem instance set that comprises 10 problem instances that are generated by sampling infusion durations in our data set gathered from the OCC at Hacettepe University Oncology Hospital between November 2017 and March 2018.

In the following set of experiments, models, and algorithms are coded and run in C++ using Microsoft Visual Studio 2022 along with CPLEX Studio IDE 20.1.0. The computations are performed with Intel Core i7- 1165G7 CPU @2.80 GHz and 16GB RAM.

The descriptive statistics regarding the data set and the problem instance generation method are explained in Section 5.1. WDB-LSA performance is evaluated in Section 5.2. Section 5.3 demonstrates the results of sensitivity analysis on various model parameters, and managerial insights are presented according to inferences. Finally, the value of a stochastic solution (VSS) is estimated in Section 5.4.

5.1 Data Description and Instance Generation

Our data set, which is detailed in Demir et al. (2021), comprises the estimated treatment durations along with the actual treatment durations of 204 patients. The data of these patients are collected over 11 different days. Infusion durations vary between 16 and 217 minutes; the average is 112.5 minutes while the 95% CI for infusion durations is [104.56, 120.35] minutes.

Table 5.1 Infusion duration intervals (in minutes) and their frequencies in the data set

Class	Predicted Interval	Realization Probability	True Interval for Infusion
1	(20, 45]	26.96%	[16, 44]
2	(45, 100]	7.85%	[29, 80]
3	(100, 150]	33.33%	[74, 132]
4	(150, 240]	31.86%	[125, 217]

For problem instance generation, β (treatment approval status) and infusion duration parameters are considered components. Both the status of treatment approval and infusion duration samples are uniquely varied for all instances in an instance set. Infusion durations are sampled from the data set of Hacettepe Oncology Hospital. While generating treatment approval status parameters, we benefit from the values reported in Garaix et al. (2020) and our own observations. Garaix et al. (2020) argue that a deferral probability ≤ 0.2 sounds realistic and plausible based on the setting they study. However, we choose a slightly larger value (0.25) than the limit they suggest for deferral probability to be less conservative. Therefore, 75% percent of patients are assumed to be approved to undergo their chemotherapy treatments following their oncologist consultations, while the remaining 25% are assumed to have their treatments disapproved (i.e. deferred).

Patients are grouped into classes based on the feedback from the head nurse and predicted infusion durations. Table 5.1 displays the realized infusion duration interval for classes along with the probability of observing patients from each class. Note that the sampling of binary beta parameters is independent of patient classes.

The initial step while generating instances is to determine each patient's class by drawing a random number between 0 and 1 and using the associated class probabilities in Table 5.1. In each instance, the number of patients in each group is the same and the patient categories do not change. In the latter step for deriving infusion durations for each patient, a value is sampled from the actual data set for the corresponding patient class. A different infusion duration is generated while also considering patient classes for each patient in each scenario. This procedure is repeated for every single patient in each problem instance. Likewise, the binary β parameter is simulated for each patient in each instance by drawing a random number between 0 and 1. If the random number drawn is smaller than the deferral probability (0.25), the β value in the sample data set is realized as 0 for the corresponding patient, and 1 otherwise.

Besides the stochastic input parameters, deterministic drug preparation durations are randomly generated between [13, 35] in minutes; deterministic oncologist consul-

tation durations are randomly generated between [15, 30] in minutes for each patient, but remaining constant across different scenarios. The data ranges are determined based on our observations in Hacettepe University Oncology Hospital. Premedication durations are constant at 15 minutes for all patients and across scenarios (Karakaya et al., 2023).

Since the assignment of patients to oncologists is not a decision made by the TSMIP problem, the pre-assignment of patients to oncologists is handled according to patient indices on the daily list prior to sequencing. Explicitly, the number of patients is divided by the number of available oncologists in the clinic, and patients are evenly distributed among oncologists. If the numbers are indivisible, the last oncologist gets fewer patients to preserve the integrality. For instance, if there are 8 patients and 2 oncologists, both oncologists would be assigned to 4 patients each. If the number of patients is 9 for 2 oncologists, the first oncologist is assigned to 5 patients while the other is assigned to 4.

5.2 Assessment of WDB-LSA Performance

In the remainder of this section, results of the performance assessment of WDB-LSA methodology against the optimal solutions of the original TSMIP model are discussed in Section 5.2.1. Next, Section 5.2.2 compares WDB-LSA to the well-known sequencing and job hedging heuristics from the relevant literature.

5.2.1 Comparison with the Optimal Solutions

This section intends to verify that the WDB-LSA provides solutions close to the optimal solutions for several instances generated for experimentation purposes. In this section, a computational time-wise comparison is intended while ensuring that the optimality gap is as small as possible, proving that the solution methodology provides an improvement in computational time without deteriorating the optimal value significantly.

For a credible comparison, it is ensured that the TSMIP model is solved to optimality by CPLEX before exceeding a reasonable time limit that is set as 3 hours in this

study. The parameter settings are set as 9 patients, 2 nurses, 4 chairs, and 3 oncologists. In the objective function, the weight of the cost for patient waiting time λ is set as 0.3. We set the number of scenarios as 35 in these experiments, because CPLEX requires more than 3 hours to find optimal solutions for larger scenario size instances. The objective, average makespan, and waiting time values are recorded as well as the computational time for each instance.

The value of $|\mathbb{R}|$ which stands for the length of the reduced set for the WDB-LSA was determined by comparing the computational efficiency associated with various values. Considering a too-small size for the reduced set would mean that the entire scenario set is approximated by only a small number of scenarios despite computational gains. On the contrary, having the $|\mathbb{R}|$ value as close to the size of the original set of scenarios as possible would produce more realistic and powerfully representative reduced sets. However, this approach is in contradiction with the essential goal of scenario reduction. Therefore, this trade-off is investigated by testing the model time with different $|\mathbb{R}|$ values, and a sufficiently large value for the reduced set size is decided as $|\mathbb{R}| = 10$.

CPLEX spends 1479 seconds, on average, to find the optimal solutions across 10 instances. The WDB-LSA spends only 306 seconds, on average, to provide a solution. Whereas, Table 5.2 shows that the average optimality gap associated with WDB-LSA solutions is 3%. The largest gap between TSMIP and WDB-LSA is detected as 6% while there is also an instance with 0% gap, indicating that it is possible to obtain the optimal solution in some instances using the WDB-LSA methodology. This implies that since the average computational time for CPLEX is 79% worse than that for WDB-LSA, a satisfactory near-optimal solution can be obtained within a significantly shorter time.

5.2.2 Comparison with Scheduling Heuristics

We compare the solutions found by WDB-LSA against those found by practically relevant scheduling heuristics from the literature. The heuristics are the combinations of well-known sequencing heuristics with appointment start time setting heuristics that can be easily used in practice by the head nurse. Four basic sequencing heuristics taken into consideration are *shortest (increasing) mean infusion duration (SPT)*, *longest (decreasing) mean infusion duration (LPT)*, *increasing variance of infusion duration (VAR)*, and *increasing coefficient of variance of infusion duration (CoV)*. While sequencing patients, we do not consider premedication or consultation dura-

Table 5.2 CPLEX optimal solution and WDB-LSA comparison in terms of objective value and model time (in seconds)

	CPLEX		**WDB-LSA**		
Instance #	**Objective**	**Time**	**Objective**	**Time**	**Gap%**
1	238.0	609.8	253.1	336.0	5.9%
2	261.6	751.2	261.7	214.2	0.0%
3	266.9	821.2	276.6	122.7	3.5%
4	267.4	165.8	275.5	253.7	2.9%
5	276.2	1665.6	284.8	631.7	3.0%
6	273.1	4716.3	287.3	290.1	4.9%
7	263.0	126.7	268.6	566.6	2.1%
8	246.6	2712.5	251.2	404.3	1.8%
9	255.8	159.0	264.1	37.3	3.1%
10	252.2	3063.3	259.5	203.5	2.8%
Average	*260.1*	*1479.1*	*268.2*	*306.0*	*3.0%*

tions, because mean or variance of these durations do not change from one patient to another. After sequencing patients based on their infusion times, estimated patient appointment times are calculated using a job hedging heuristic.

It is demonstrated in our TSMIP model formulation that the binary first stage decision variable b_{ij} depicts the precedence relationship between patients $i \in I$ and $j \in I$ while the first stage decision variable a_i stands for the appointment time for patient $i \in I$. Note that e_{ij} is also a first-stage binary decision variable whose value is dependent directly on a_i and indirectly on b_{ij}. Since the derivation of e_{ij} values would be fairly burdensome for the head nurse, the job hedging heuristic does not tackle determining e_{ij} values. This task is rather handed over to the two of the original first stage constraints (3.4)-(3.5) that contain the e_{ij} variables. As a result, we determine the values of b_{ij} and a_i using the heuristics, and then solve the remaining part of the TSMIP model (i.e., we exclude constraints which include only these fixed variables).

5.2.2.1 Implementation of Heuristics

Based on the classes to which certain patients belong, the mean infusion durations within those particular classes (using the original infusion duration data) are calculated. Similarly, variance and coefficient of variation values of original infusion durations are calculated for each patient's class and are assigned to patients of those classes. Eventually, patients could be sorted based on which class they belong to,

Table 5.3 Average gap percentages of various sequencing rules from the WDB-LSA solutions under different job hedging levels

Job Hedging Level	**LPT&CoV Gap%**	**SPT Gap%**	**VAR Gap%**
40%	25.6	24.4	23.4
45%	26.4	25.3	23.6
50%	27.1	26.0	24.4
55%	27.7	26.5	25.1
60%	28.1	26.8	25.5
65%	29.3	28.2	26.7
70%	30.4	29.0	27.9
75%	30.9	29.6	28.5
80%	32.0	30.2	29.6
Average	*28.6*	*27.4*	*26.1*

which oncologist they are assigned to, and which sequencing heuristic is under consideration. It is of great importance to note that throughout the sequencing process, it is ensured that patients are sequenced within their oncologist groups so that we could make an eligible comparison to WDB-LSA.

After sorting the patient indices and determining b_{ij} values with the assistance of scheduling heuristics, we notice that the LPT rule and the CoV rule resulted in the same sequence of patients. Therefore, the following steps of the algorithm has the same results for both cases in which LPT and CoV are considered as sequencing rules. In order to apply job hedging to obtain the appointment start times of patients, different percentile values of the treatment durations are calculated in each patient class. Next, we estimate the treatment duration as the summation of infusion duration and the constant premedication duration (15 min) for each patient. Then, to calculate the appointment time of a patient, we add the treatment time of the preceding patient of the same oncologist to the appointment time of the preceding patient. However, appointment times for patients that precede all others in their oncologist group are set as 0. For each patient class, percentile values are varied between 40% and 80% in increments of 5% to allow and assess a wider range of values coming from the original distribution.

For the experiments, runs are performed with Algorithm 2 for 100 scenarios. While appointment times (a_i) between the 40th and 80th percentiles are set as parameters, the increase in percentile values results in increased total working time (makespan). Indicating a trade-off here, the waiting time either keeps decreasing or remains fixed to a certain value since the patients are assigned later appointment times as the percentile value is increased. Scheduling a patient at a larger appointment time means that the patients arrive at the clinic later. Therefore, the expected waiting

times in the clinic change in a non-increasing fashion.

Table 5.3 demonstrates the average objective gap percentages for the comparison of WDB-LSA to varying combinations of job hedging levels and sequencing heuristics. Ideally, it is expected to have these gap percentage values as high as possible in favor of WDB-LSA to ensure the methodology we suggest outperforms the conventional one. Table 5.3 shows that the WDB-LSA significantly outperforms practical heuristics. Although SPT and VAR sequencing rules perform slightly better than LPT and CoV, the performance gap against these rules is fairly similar. Since the Gap% keeps increasing with respect to job hedging levels, the best solution performance for the sequencing rules is obtained on the 40% level for each.

Algorithm 2 Job Hedging Heuristic

1: Calculate the average infusion duration for each patient.

$$avg_i = \frac{\sum_{\omega \in \Omega} t_i^{\omega}}{\|\Omega\|}, \quad \forall i \in I$$

2 :Assign index numbers to patients and determine the classes for each patient. Sort avg_i values within oncologist groups with respect to one of the sequencing heuristics (LTP, SPT, VAR or CoV). Sequence the index numbers of the patients.
3 :Assign b_{ij} values according to the patient sequence found in Step 2.
4 :Determine a_i values (appointment times) for all patients according to a range of varying percentile levels of job hedging heuristic. In each oncologist group, if a patient is preceding all other patients, the appointment time of that patient is assigned as zero. The appointment times of other patients in each oncologist group are determined in sequence with respect to their precedence by taking account of the first average available time of nurse and chair simultaneously.
5 : Call CPLEX to solve TSMIP by fixing b_{ij} and a_i values and the optimal objective function value.

5.3 Sensitivity Analysis on Model Parameters

In this section, sensitivity analysis results for objective function coefficient λ are discussed in subsection 5.3.1, and sensitivity analysis results for varying numbers of nurses and oncologists are provided in subsection 5.3.2.

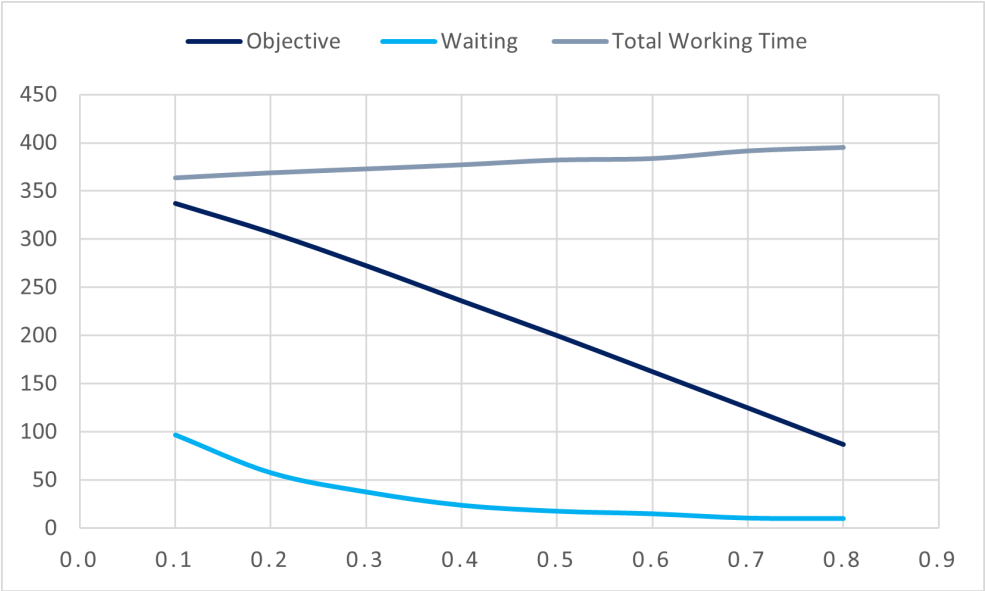

Figure 5.1 The change in objective value, waiting times, and the makespan with respect to varying λ values.

5.3.1 Impact of the λ Value

The $\lambda \in [0,1]$ value is the objective function coefficient for the patient waiting time while $(1-\lambda)$ is the coefficient for the total working time in the TSMIP formulation. Different values chosen for λ indicate the importance attached to key performance measures in the model. In this section, the trade-off between these performance measures is assessed by varying the λ value between 0.1 and 0.8 with increments of 0.1. Extreme values for λ, such as 0 or 1, are excluded since that would mean ejecting one of the key measures from the picture. Throughout the experiments, the number of patients, nurses, oncologists, chairs, and scenarios are kept constant at 9, 2, 3, 4, and 35 respectively.

The numerical results of the experiment are provided in Table 5.4. As the coefficient of patient waiting time λ increases by 0.1, the objective function decreases substantially. In fact, a gradually more dramatic percentage difference between two consecutive objective function values is observed for every 0.1 increase in the λ value. Waiting time naturally decreases as lambda increases, eventually with a significant reduction of 90%. Total working time, on the other hand, constantly increases, with a maximum increase rate of 10%. This indicates that waiting time and total working time are two conflicting components of the objective function. This significant reduction rate in waiting time results in a 74% improvement in the objective function from $\lambda = 0.1$ to $\lambda = 0.8$. Thus, favoring the waiting time in terms of relative importance (against the total working time) leads to substantially superior results

Table 5.4 Sensitivity of average objective value, patient waiting time (in minutes), and clinic closing time (in minutes) to λ

	Objective Value	**Waiting Time**	**Makespan**
λ=0.1	366.8	96.5	363.5
λ=0.2	306.5	57.3	368.8
λ=0.3	272.2	37.3	372.8
λ=0.4	235.7	23.5	377.1
λ=0.5	199.7	17.4	382.0
λ=0.6	162.2	14.6	383.6
λ=0.7	127.7	10.2	391.5
λ=0.8	86.8	9.7	395.1

for the expected objective function value.

5.3.2 Impact of the Number of Nurses and Oncologists

The sensitivity of the model outputs against different combinations of varying numbers of nurses and oncologists is assessed in this section. Keeping all other parameters the same, the number of oncologists ranges from 2 to 4, and the number of nurses ranges from 2 to 3. Since this problem considers nurse assignments in the second stage, the number of nurses must be at least 2. The sensitivity analysis in this part reveals peculiar results at first appearance, which can be explained by the nature of the solution approach. The numerical values of nurse-oncologist variations are shown in Tables 5.5, 5.6, and Figure 5.2.

Patterns observed:

For two nurses, as the number of oncologists increases from 2 to 3, waiting time slightly increases and then remains constant from 3 oncologists to 4 oncologists. For three nurses, the waiting time initially decreases, then remains constant. For two oncologists, as the number of nurses increases, the average waiting time increases. The same peculiarity is not observed for O = 3 or O = 4.

The objective function value (see Table 5.6) automatically provides similar unexpected patterns since the formulation is solely based on the patient waiting time and the total working time (makespan) of the clinic. However, the objective function tends to decrease as the number of oncologists increases. Computational time does not exhibit an expected pattern with respect to the increase of problem parameters as well.

Figure 5.2 The change in the makespan, waiting times, and objective function values based on varying numbers of nurses and oncologists in the respective order.

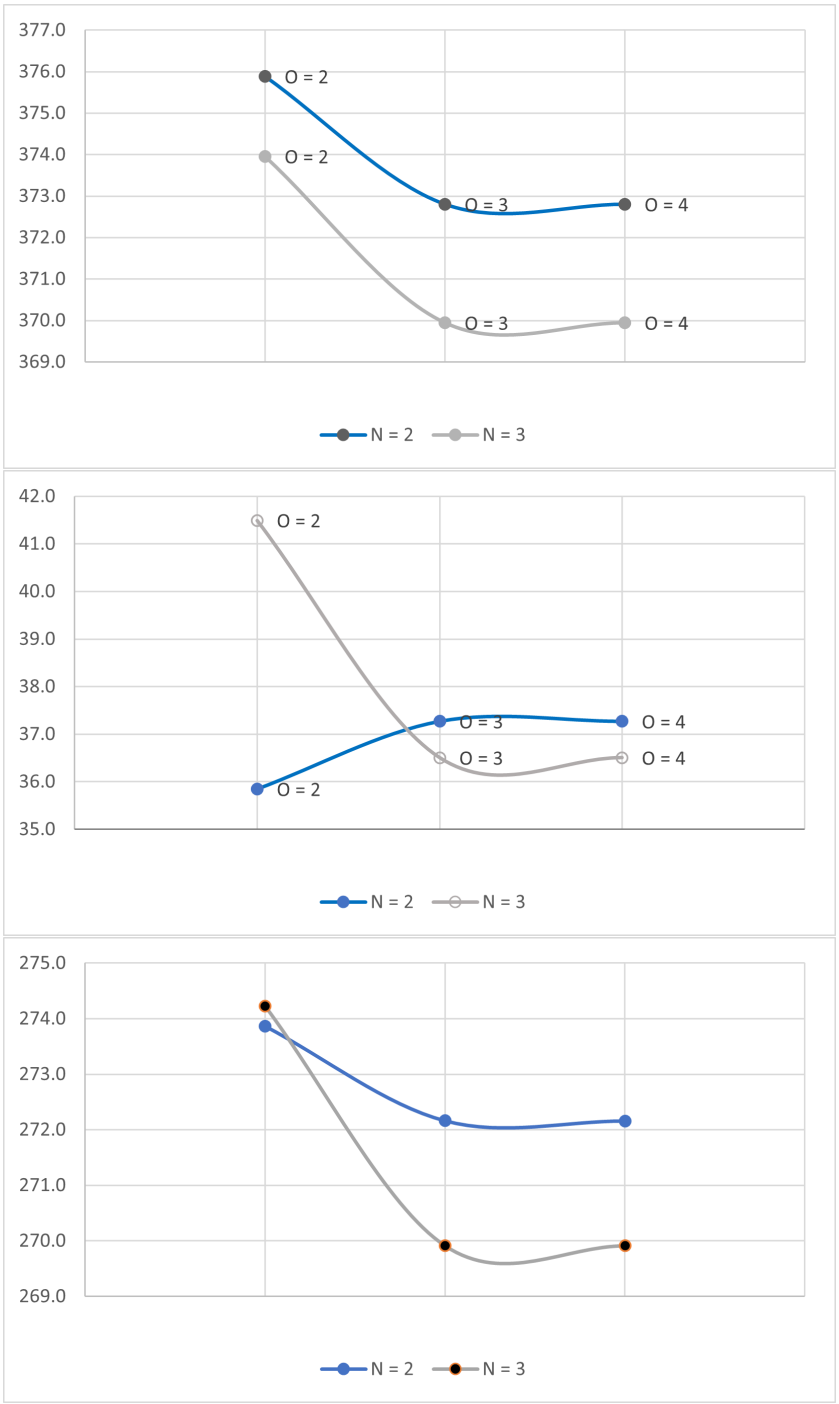

Table 5.5 Average patient waiting time and total working time (makespan) values (in minutes) varying with respect to different numbers of oncologists and nurses.

	Waiting Time			**Total Working Time**		
	$\|O\| = 2$	$\|O\| = 3$	$\|O\| = 4$	$\|O\| = 2$	$\|O\| = 3$	$O\| = 4$
$\|N\| = 2$	35.8	37.3	37.3	375.9	372.8	372.8
$\|N\| = 3$	41.5	36.5	36.5	374.0	369.9	369.9

Table 5.6 Average objective function (in minutes) and computational run time values (in seconds) varying with respect to different numbers of oncologists and nurses.

	Objective			**Computational Time**		
	$\|O\| = 2$	$\|O\| = 3$	$\|O\| = 4$	$\|O\| = 2$	$\|O\| = 3$	$O\| = 4$
$\|N\| = 2$	273.9	272.2	272.2	643.0	690.8	671.8
$\|N\| = 3$	274.2	269.9	269.9	791.9	685.4	723.1

As the number of nurses in the OCC increases, the total working time of the clinic conceivably exhibits a slight reduction (see Table 5.5). Similarly, as the number of oncologists increases, the total working time of the clinic initially slightly decreases and then remains the same.

Unexpected pattern observations are likely to occur in our experiments considering the existence of potential alternative optimal solutions when the TSMIP is solved based on the reduced scenario set (10 scenarios). This model outputs the first-stage variables to be used as input parameters for the second-stage problem which is solved using 100 scenarios. Two alternative optimal solutions for the TSMIP model with the reduced scenario set may yield different objective values when the second-stage problem is solved using the original scenario set. In other words, the evaluation process of the solution methodology alters the average values and results in peculiar outputs since the scenario reduction method produces imperfect approximations.

5.4 VSS Estimation

The value of stochastic solution assessment is the concept of solving the original stochastic model with a deterministic setup for a single scenario. When the stochastic values are substituted by mean/deterministic values for the input parameters, VSS allows us to assess the quality of the expected solution value (Escudero, Garín, Merino & Pérez, 2007). In other words, VSS estimates the predicted advantage from solving a stochastic model rather than its deterministic version. While a low value

Table 5.7 Percent improvement acquired on average mean value (MV) solutions based on varying λ values using the WDB-LSA.

λ	**Average VSS%**	**Min VSS%**	**Max VSS%**
0.3	10%	0%	16%
0.5	12%	4%	24%
0.7	19%	6%	38%
Avg	*14%*	*3%*	*26%*

for VSS nearly always indicates a flaw in the modeling itself, obtaining satisfactory VSS values proves that stochastic programming models are necessary despite the computational challenges (Maggioni & Wallace, 2012).

The VSS is estimated by calculating the difference between the WDB-LSA solutions and the expected objective value of the mean value (MV) problem. The assessment is made considering different trade-off parameter values (λ), and average improvement percentages WDB-LSA makes on the deterministic MV solution are reported in Table 5.7 along with the maximum and minimum improvement rate that can be obtained from individual instances. Evidently, the improvement percentage increases as the objective function coefficient of the waiting time increases. The average VSS over all instances is found as 14%, while the VSS can reach up to 38%.

6. CONCLUSION

Cancer is a global health concern and a leading cause of death in Turkey. The risk factors that contribute to cancer development include genetic predisposition, unhealthy habits, or other unavoidable external factors. Outpatient Chemotherapy Clinics (OCC) are prominent healthcare institutions that facilitate oncology patients' intra-day follow-up chemotherapy treatments. Given the limited resources, the uncertainty surrounding infusion durations, and the critical nature of cancer treatment, scheduling chemotherapy in outpatient clinics poses significant challenges and constraints. To address this, our study integrates oncologist consultation and chemotherapy scheduling, ensuring coordination of daily sequences and appointment times for patients scheduled on the same day through the medium of a two-stage stochastic mixed integer programming model. The model considers integer appointment times and incorporates stochastic elements such as infusion times and the approval status of chemotherapy treatments based on the outcomes of consultations with oncologists. In the first stage of the model, after the patient sequence is arranged, appointment times for consultation are set taking into account patients' designated oncologists. The second stage involves assigning patients to chairs and nurses. The objective function incorporates a weighted sum of the total clinic working time and patient waiting times. To solve the model efficiently, we employ a scenario reduction algorithm to reduce problem complexity. The suggested algorithm, referred to as the Wasserstein Distance-Based Local Search Algorithm (WDB-LSA), has been examined using authentic data gathered from a prominent oncology hospital in Turkey. The algorithm is benchmarked against various practical heuristics drawn from scholarly sources. It is shown that WDB-LSA is both effective and computationally efficient at improving chemotherapy scheduling in outpatient departments, even when considering a range of uncertainties and constraints. The influence of different model parameters has been evaluated through sensitivity analysis. Additionally, the solution methodology has been examined against the mean value solution to assess the value of the stochastic solution.

6.1 Future improvement opportunities

Limitations of the model include a lack of consideration for other processes and limited resources in the patient pathway. In future studies, blood testing can become a part of the patient flow considered for the scheduling problem in an OCC, considering uncertainty in blood draw durations and limitations on responsible lab technicians.

In addition, the model can be revised by the assumption that the number of pharmacy technicians is limited. In such a case, drug preparation is likely to become a bottleneck process, evoking a new space for additional patient waiting time.

Another revision of the model can be the consideration of fairness issues in the objective function.

In this thesis study, oncologists are not considered resources to assign patients to, and rather entities that are already assumed to be designated to certain patients in a pre-determined list. It is also possible to consider these entities as limited resources of an OCC.

In order to improve and expand the TSMIP model, varying treatment deferral probabilities can be assessed. Additionally, oncologist consultation times could also be introduced as uncertain parameters to the problem by collecting realistic data from healthcare institutions. These approaches may play an important role in enhancing patient satisfaction and better time management in future studies.

For the implementation of the scenario reduction, integrating different distance metrics into the LSA could be considered another experiment opportunity. It is possible to alter the randomization method for the initial reduced set in scenario reduction. Additionally, rather than randomly initializing the reduced set, a different systematic approach can also be explored. After the randomization, the scenario swapping operation between the reduced set and the entire scenario set can also be diversified. Other than the distribution-based discrete scenario reduction approach adopted in this study, problem-based reduction methods can be tested in terms of practicality, computational performance, and solution quality.

BIBLIOGRAPHY

Abouelrous, A., Gabor, A. F., & Zhang, Y. (2022). Optimizing the inventory and fulfillment of an omnichannel retailer: a stochastic approach with scenario clustering. *Computers & Industrial Engineering*, *173*, 108723.

ACS, A. (2023). Cancer facts & figures 2023. *American Cancer Society, National Home Office, Atlanta*.

Alexander, M., Blum, R., Burbury, K., Coutsouvelis, J., Dooley, M., Fazil, O., Griffiths, T., Ismail, H., Joshi, S., Love, N., et al. (2017). Timely initiation of chemotherapy: a systematic literature review of six priority cancers–results and recommendations for clinical practice. *Internal Medicine Journal*, *47*(1), 16–34.

Alvarado, M. & Ntaimo, L. (2018). Chemotherapy appointment scheduling under uncertainty using mean-risk stochastic integer programming. *Health care management science*, *21*, 87–104.

Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., & Pandit, V. (2001). Local search heuristic for k-median and facility location problems. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, (pp. 21–29).

Benzaid, M., Lahrichi, N., & Rousseau, L.-M. (2020). Chemotherapy appointment scheduling and daily outpatient–nurse assignment. *Health care management science*, *23*, 34–50.

Bertsimas, D. & Mundru, N. (2022). Optimization-based scenario reduction for data-driven two-stage stochastic optimization. *Operations Research*.

Castaing, J., Cohn, A., Denton, B. T., & Weizer, A. (2016). A stochastic programming approach to reduce patient wait times and overtime in an outpatient infusion center. *IIE Transactions on Healthcare Systems Engineering*, *6*(3), 111–125.

Cataldo, A., Sufan, S., Lorca, Á., Andresen, M., Sánchez, C., & Sauré, A. (2023). Multi–stage process for chemotherapy scheduling and effective capacity determination. *International Transactions in Operational Research*, *30*(1), 151–180.

Cococcioni, M. & Fiaschi, L. (2021). The big-m method with the numerical infinite m. *Optimization Letters*, *15*(7), 2455–2468.

Corsini, R. R., Costa, A., Fichera, S., & Parrinello, V. (2022). Hybrid harmony search for stochastic scheduling of chemotherapy outpatient appointments. *Algorithms*, *15*(11), 424.

Crainic, T. G., Hewitt, M., & Rei, W. (2014). Scenario grouping in a progressive hedging-based meta-heuristic for stochastic network design. *Computers & Operations Research*, *43*, 90–99.

Demir, N. B., Gul, S., & Çelik, M. (2021). A stochastic programming approach for chemotherapy appointment scheduling. *Naval Research Logistics (NRL)*, *68*(1), 112–133.

Dupačová, J., Gröwe-Kuska, N., & Römisch, W. (2003). Scenario reduction in stochastic programming. *Mathematical programming*, *95*, 493–511.

Escudero, L. F., Garín, A., Merino, M., & Pérez, G. (2007). The value of the stochastic solution in multistage problems. *Top*, *15*(1), 48–64.

Escudero, L. F., Garín, M. A., Pérez, G., & Unzueta, A. (2013). Scenario cluster decomposition of the lagrangian dual in two-stage stochastic mixed 0–1 optimization. *Computers & Operations Research*, *40*(1), 362–377.

Fairbrother, J., Turner, A., & Wallace, S. W. (2022). Problem-driven scenario generation: an analytical approach for stochastic programs with tail risk measure. *Mathematical Programming*, 1–42.

Fei, X., Gülpınar, N., & Branke, J. (2019). Efficient solution selection for two-stage stochastic programs. *European Journal of Operational Research*, *277*(3), 918–929.

Gade, D., Hackebeil, G., Ryan, S. M., Watson, J.-P., Wets, R. J.-B., & Woodruff, D. L. (2016). Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs. *Mathematical Programming*, *157*, 47–67.

Garaix, T., Rostami, S., & Xie, X. (2020). Daily outpatient chemotherapy appointment scheduling with random deferrals. *Flexible Services and Manufacturing Journal*, *32*, 129–153.

Good, I. J. (1953). The population frequencies of species and the estimation of population parameters. *Biometrika*, *40*(3-4), 237–264.

Gul, S. (2021). Chemotherapy appointment scheduling under uncertainty by considering workload balance among nurses. *Pamukkale Üniversitesi Mühendislik Bilimleri Dergisi*, *27*(4), 570–578.

Gul, S. (2022). Scheduling chemotherapy appointments under uncertainty by considering different nursing care delivery schemes. *arXiv preprint arXiv:2211.10258*.

Hadid, M., Elomri, A., El Mekkawy, T., Jouini, O., Kerbache, L., & Hamad, A. (2022). Operations management of outpatient chemotherapy process: An optimization-oriented comprehensive review. *Operations Research Perspectives*, *9*, 100214.

Haghi, M., Hashemi Doulabi, H., Contreras, I., & Bhuiyan, N. (2022). Integrated consultation and chemotherapy scheduling with stochastic treatment times. *Journal of the Operational Research Society*, 1–16.

Henrion, R. & Römisch, W. (2022). Problem-based optimal scenario generation and reduction in stochastic programming. *Mathematical Programming*, *191*(1), 183–205.

Hesaraki, A. F., Dellaert, N. P., & de Kok, T. (2019). Generating outpatient chemotherapy appointment templates with balanced flowtime and makespan. *European Journal of Operational Research*, *275*(1), 304–318.

Heshmat, M., Nakata, K., & Eltawil, A. (2018). Solving the patient appointment scheduling problem in outpatient chemotherapy clinics using clustering and mathematical programming. *Computers & Industrial Engineering*, *124*, 347–358.

Jiang, X., Bai, R., Wallace, S. W., Kendall, G., & Landa-Silva, D. (2021). Soft clustering-based scenario bundling for a progressive hedging heuristic in stochastic service network design. *Computers & Operations Research*, *128*, 105182.

Karakaya, S., Gul, S., & Çelik, M. (2023). Stochastic scheduling of chemotherapy appointments considering patient acuity levels. *European Journal of Operational Research*, *305*(2), 902–916.

Ketkov, S. S. (2023). Distributionally robust mixed-integer programming with

wasserstein metric: on the value of uncertain data. *arXiv preprint arXiv:2304.01343*.

Keutchayan, J., Ortmann, J., & Rei, W. (2023). Problem-driven scenario clustering in stochastic optimization. *Computational Management Science*, *20*(1), 13.

Khorana, A. A., Tullio, K., Elson, P., Pennell, N. A., Grobmyer, S. R., Kalady, M. F., Raymond, D., Abraham, J., Klein, E. A., Walsh, R. M., et al. (2019). Time to initial cancer treatment in the united states and association with survival over time: an observational study. *PloS one*, *14*(3), e0213209.

Kleywegt, A. J., Shapiro, A., & Homem-de Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on optimization*, *12*(2), 479–502.

Lau, P. K., Watson, M. J., & Hasani, A. (2014). Patients prefer chemotherapy on the same day as their medical oncology outpatient appointment. *Journal of oncology practice*, *10*(6), e380–e384.

Liang, B., Turkcan, A., Ceyhan, M. E., & Stuart, K. (2015). Improvement of chemotherapy patient flow and scheduling in an outpatient oncology clinic. *International Journal of Production Research*, *53*(24), 7177–7190.

Lloyd, S. (1982). Least squares quantization in pcm. *IEEE transactions on information theory*, *28*(2), 129–137.

Löhndorf, N. (2016). An empirical analysis of scenario generation methods for stochastic optimization. *European Journal of Operational Research*, *255*(1), 121–132.

Lyon, G., Cataldo, A., Angulo, G., Rey, P. A., & Sauré, A. (2023). A column generation approach to intraday scheduling of chemotherapy patients. *International Journal of Production Research*, *61*(7), 2231–2249.

Maggioni, F. & Wallace, S. W. (2012). Analyzing the quality of the expected value solution in stochastic programming. *Annals of Operations Research*, *200*, 37–54.

Masini, C., Nanni, O., Antaridi, S., Gallegati, D., Marri, M., Paolucci, D., Minguzzi, M., & Altini, M. (2014). Automated preparation of chemotherapy: quality improvement and economic sustainability. *American Journal of Health-System Pharmacy*, *71*(7), 579–585.

Narum, B. S., Fairbrother, S. W., & Wallace, J. (2022). Problem-based scenario generation by decomposing output distributions.

Oliveira, W., Sagastizábal, C., & Scheimberg, S. (2011). Inexact bundle methods for two-stage stochastic programming. *SIAM Journal on Optimization*, *21*(2), 517–544.

Pflug, G. C. & Pichler, A. (2015). Dynamic generation of scenario trees. *Computational Optimization and Applications*, *62*, 641–668.

Prochazka, V. & Wallace, S. W. (2020). Scenario tree construction driven by heuristic solutions of the optimization problem. *Computational Management Science*, *17*(2), 277–307.

Ramos, C., Cataldo, A., & Ferrer, J.-C. (2020). Appointment and patient scheduling in chemotherapy: a case study in chilean hospitals. *Annals of Operations Research*, *286*, 411–439.

Römisch, W. (2009). Scenario reduction techniques in stochastic programming. In *Stochastic Algorithms: Foundations and Applications: 5th International Symposium, SAGA 2009, Sapporo, Japan, October 26-28, 2009. Proceedings*

5, (pp. 1–14). Springer.

Rujeerapaiboon, N., Schindler, K., Kuhn, D., & Wiesemann, W. (2022). Scenario reduction revisited: Fundamental limits and guarantees. *Mathematical Programming*, *191*(1), 207–242.

Sadki, A., Xie, X., & Chauvin, F. (2011). Appointment scheduling of oncology outpatients. In *2011 IEEE International Conference on Automation Science and Engineering*, (pp. 513–518). IEEE.

Turkcan, A., Zeng, B., & Lawley, M. (2012). Chemotherapy operations planning and scheduling. *IIE Transactions on Healthcare Systems Engineering*, *2*(1), 31–49.

Vidal-Carreras, P. I., Garcia-Sabater, J. J., & Marin-Garcia, J. A. (2022). Applying value stream mapping to improve the delivery of patient care in the oncology day hospital. *International Journal of Environmental Research and Public Health*, *19*(7), 4265.

Vidyashankar, A. N. & Xu, J. (2015). Stochastic optimization using hellinger distance. In *2015 Winter Simulation Conference (WSC)*, (pp. 3702–3713). IEEE.

Zhang, W., Wang, K., Jacquillat, A., & Wang, S. (2023). Optimized scenario reduction: Solving large-scale stochastic programs with quality guarantees. *INFORMS Journal on Computing*.