TESIS DE LA UNIVERSIDAD DE ZARAGOZA

2024 281

Daniel Julio Calcinaro

Clinical Trials Supply Chain Optimization

Director/es
Gürbüz, Mustafa

ISSN 2254-7606

Prensas de la Universidad
Universidad Zaragoza

Universidad de Zaragoza
Servicio de Publicaciones

ISSN 2254-7606

Universidad Zaragoza
1542

Tesis Doctoral

CLINICAL TRIALS SUPPLY CHAIN OPTIMIZATION

Autor

Daniel Julio Calcinaro

Director/es

Gürbüz, Mustafa

UNIVERSIDAD DE ZARAGOZA
Escuela de Doctorado

Programa de Doctorado en Logística y Gestión de la Cadena de Suministro

2024

Clinical Trials Supply Chain Optimization

by

Daniel J. Calcinaro

Bachelor of Industrial Organization, Universidad Tecnológica Nacional (2013)

Submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in the

MIT-ZARAGOZA INTERNATIONAL LOGISTICS PROGRAM

at the

ZARAGOZA LOGISTICS CENTER,

a research institute affiliated with the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
and the
UNIVERSITY OF ZARAGOZA

December, 2023

Author .....
Daniel Julio Calcinaro
Ph.D. Candidate, MIT-Zaragoza International Logistics Program

Certified by .....
Mustafa Çagri Gürbuz, Ph.D.
Professor MIT-Zaragoza International Logistics Program
Research Affiliate, MIT Center for Transportation and Logistics
Thesis Supervisor

Certified by .....
Yasel Costa Salas, Ph.D.
Professor and Ph.D. Program Director, MIT-Zaragoza International Logistics Program
Research Affiliate, MIT Center for Transportation and Logistics
PhD Program Director

Clinical Trials Supply Chain Optimization: A Multi-Stage Stochastic Programming Approach

by

Daniel Julio Calcinaro

Submitted to the MIT-Zaragoza International Logistics Program in partial fulfillment of the requirements for the degree of Doctor of Philosophy at the Zaragoza Logistics Center, a research institute affiliated with the Massachusetts Institute of Technology and the University of Zaragoza.

ABSTRACT

The growth of development of new pharmaceutical products, combined with a rigid regulation have saturated the capacity of nations that traditionally used to lead the execution of clinical trials, leading to a sustained an increasing globalization of such experiments, increasing as well their logistic complexity, costs and completion time. Clinical trials costs have gone from marginal to represent a significant portion of R&D expenditures, attracting attention from the industry and the interest on finding methodologies to efficiently manage production and distribution of pharmaceutical stock.

The main challenge that clinical trials present is the fact that we never know when exactly, or even where patients will register to receive their first dose. Furthermore, it is unknown how long it will take until the trial ends, which makes it more complicated than traditional inventory problems. This is a finite horizon problem, but without a specified completion time. Still, the design is expected to grant that they are able to receive treatment wherever and whenever they show up. If patients could wait indefinitely, in a trial that enrols H patients to receive m doses each in one of the N enrolment sites, we would just place $H \times m$ units in the Central Warehouse (CW), and we would wait to see where do patients enrol in order to make precise downstream deliveries; if, on the contrary, production cost was not an issue or patients could not wait at all, we would then place inventory downstream, exactly $H \times m$ in each one of the enrolment sites, which means we would be producing N times as many units as we will finally use. In practice, reality is not described by these extreme cases: cost will be one of the main objectives, and the need of filling demand soon one of the main constraints, with several factors to be accounted for, like production, transportation and holding costs, as well as lead times. The

main goal in this research is the determination of optimal production and distribution decisions in a realistic representation able to capture the complexities that characterize most clinical trials.

The proposed model is a Multi-Stage Mixed Integer Linear Stochastic Program which, as expected, requires a prohibitively large number of scenarios to cover all the possible progressions of the system. For this reason, we propose a clustering-based scenario reduction methodology, which groups similar scenarios together according to a novel dissimilarity metric that we introduce, and which can be adjusted to any multi-stage multi-echelon problem.

Finally, a set of case studies is used to test the proposed methodologies, with highly satisfactory results both in terms of answering our research questions and also bringing new questions for future research.

RESUMEN

El crecimiento del desarrollo de nuevos productos farmacéuticos, en combinación con un rígido sistema regulatorio han saturado la infraestructura de los países que tradicionalmente lideraron los ensayos clínicos, resultando ello en la creciente globalización de tales experimentos, elevando también la complejidad logística de éstos, su costo y el tiempo necesario para completarlos. El costo de estos ensayos ha pasado de ser marginal ha representar una porción importante del I+D, atrayendo la atención de la industria y el interés por encontrar metodologías para gestionar eficientemente la producción y distribución del inventario farmacéutico.

La principal dificultad de los ensayos clínicos radica en que no se sabe con exactitud dónde o cuándo los pacientes se registrarán para recibir su primera dosis. Más aún, se desconoce qué tiempo tomará completar el experimento, lo cual lo hace más complicado que los problemas de inventario tradicionales. Es éste un problema en el que el horizonte de tiempo es finito, pero no especificado. Sin embargo, se espera que el diseño garantice que los pacientes puedan recibir tratamiento allí donde decidan registrarse, sin someterlos a una espera excesiva. Si los pacientes pudieran esperar indefinidamente, en un ensayo donde H pacientes deben recibir m dosis cada uno en alguno de los N sitios habilitados, colocaríamos simplemente $H \times m$ unidades en el nivel superior, y esperaríamos a ver dónde los pacientes se presentan para

entonces realizar la distribución; si, por el contrario, el costo de producción no fuera un problema y/o los pacientes no pudieran esperar en absoluto, colocaríamos $H \times m$ unidades en cada uno de los sitios habilitados, produciendo una cantidad N veces mayor a la que finalmente sera administrada. En la práctica, la realidad está alejada de estos casos extremos: el costo es un factor de decisión crucial, los pacientes no esperarán indefinidamente, y una serie de factores influyen en la decisión óptima, incluyendo los costos de mantenimiento de inventario, los costos de producción y distribución y los tiempos de entrega. El objetivo que perseguimos en este trabajo es la determinación de las decisiones adecuadas en un escenario capaz de capturar tales particularidades, adaptándose a una porción considerable de ensayos clínicos.

El modelo propuesto es un Programa Stocástico Lineal Multi-Etapa que, como es esperable, requiere una cantidad prohibitivamente grande de escenarios para cubrir todas las posibles evoluciones del sistema. Por ello, proponemos una metodología de reducción de escenarios basada en clustering, la cual agrupa escenarios similares de acuerdo a una original métrica que presentamos, la cual se ajusta a cualquier problema multi-etapa y multi-escalón.

Finalmente, una serie de casos de estudio proporcionarán el marco para poner a prueba las metodologías propuestas, con resultados altamente satisfactorios en cuanto a las respuestas a nuestras preguntas de investigación, así como también abrir camino a futuras investigaciones.

Thesis Supervisor:
Mustafa Çagri Gürbuz, Ph.D.
Professor MIT-Zaragoza International Logistics Program
Research Affiliate, MIT Center for Transportation and Logistics

Dedication

To my beloved parents: I know that you wanted me to be a different type of Doctor, but I really hope the findings in this thesis will contribute to save some lives.

Acknowledgements

I especially thank my advisor, Professor Çagri Gürbuz: I am out of words to describe his dedication and guidance not only during the research and thesis, but even also before it even started: this thesis would just not exist without the generosity and commitment that he demonstrated to help me achieve my dream. I want to also thank to my former advisor, Professor Yasel Costa Salas for his support during the earlier stages of my PhD. I express my gratitude to Dr. Christopher Caplice and Dr. Christopher Mejía Argueta who, during my time at MIT, first encouraged me in pursuing a Ph.D. and endorsed my candidacy with their recommendation letters. To Roche Pharmaceuticals, my thanks for fully opening the company's doors to me and to providing me with operational data and tacit knowledge, fully engaging in this endeavor. I gratefully acknowledge the support of the MIT-Zaragoza International Logistics Program, including all faculty and staff. To my classmate and friend Luis Eduardo Wilbert, thank you for the mutual encouragement along those intense years of our doctoral studies: he demonstrated gratitude when he had my help, and offered his support and friendship when I was facing difficult times. I am also grateful to my friend Luiggi Laporte for his indirect contribution to shaping my research ideas. To Anastasiia, I am thankful for the emotional care during the final stage of this thesis. To my parents, Mirta and Julio, for their patience and love.

Contents

List of tables **11**

List of figures **12**

List of acronyms **14**

1 Introduction **16**

1.1 Context and Motivation 16

1.2 Objectives and Contributions 19

1.3 Methodological Approach 20

1.4 Research Questions 21

1.5 Outline of the thesis 21

2 Literature Review **23**

3 Statistical Analysis **28**

3.1 Industry Data and Descriptive Statistics 28

3.2 New patients arrival pattern 29

3.2.1 Linearity 29

3.2.2 Relationship between Poisson behaviour and linearity 29

3.3 Memory-less inter-arrival times hypothesis testing 33

3.3.1 Drivers for non-Poisson behavior 33

4 Mathematical Model **37**

4.1 Problem Statement 37

4.2 Assumptions 38

4.3 Stochastic characterization 39

4.4 Decision variables, parameters and sets 43

4.4.1 Sets and indexes . . . 43
4.4.2 Parameters . . . 44
4.4.3 Decision variables . . . 45
4.5 A multi-objective mathematical model . . . 46
4.5.1 Objectives . . . 47
4.5.2 Constraints . . . 48
4.5.3 A note on production batches . . . 55
4.5.4 A note on Completion Time minimization as main objective . . . 57
4.6 Conditional sampling and sequential optimization . . . 58

5 Scenario Reduction **60**
5.1 Scenario Reduction . . . 60
5.2 Scenario reduction by clustering . . . 61
5.2.1 K-medoids++ clustering . . . 61
5.2.2 A new metric for multi-stage scenario clustering in a multi-echelon setting 64

6 Case Studies **69**
6.1 WC28325: A phase III CT for *diabetes mellitus* treatment . . . 70
6.1.1 Stability Test . . . 73
6.1.2 Conditional Sampling and Sequential Optimization . . . 75
6.1.3 Scenario Reduction . . . 79
6.1.4 Sequential Optimization with $S' = 60$ and $K = 9$ . . . 80
6.1.5 Comparison with other policies . . . 81
6.1.6 Sensitivity over arrival rates . . . 84
6.1.7 Consideration of a fixed production cost . . . 89
6.2 BO29389: A phase II CT for malignant neoplasm of branchus and lung . . . 92
6.2.1 Modeling considerations . . . 92
6.2.2 Numerical Results . . . 93
6.2.3 Comparison with other policies . . . 96
6.2.4 A note on the consequences of using weekly stages . . . 98
6.3 BP30002: A phase I CT for glaucoma treatment . . . 99
6.3.1 Data, Assumptions and Modeling Considerations . . . 99
6.3.2 Numerical Results . . . 101

7 Extensions to the model **105**

7.1 Enforcement of a simpler policy (Periodic Review s-S) 105

7.2 Multi-product CT 110

7.2.1 Changes to the model described in Chapter 4 110

7.2.2 Periodic Review s-S and can-order policy 113

8 Conclusion **116**

List of tables

2.1 Clinical Trial Supply Chain Management related papers 27
2.2 Stochastic Programming Scenario Reduction related papers 27

5.1 Enrolment sites data for illustrative example in 5.2.2 66
5.2 Country depots data for illustrative example in 5.2.2 66
5.3 Scenarios generated for illustrative example in 5.2.2 66
5.4 Cost for each scenario in the illustrative example in 5.2.2 67

6.1 summary of cases to analyze 70
6.2 Enrolment sites data for CT WC28325 72
6.3 Country depots data for CT CW28325 72
6.4 Distribution of final inventory across the WC28325 network 78
6.5 Changes applied to the arrival rates site by site in section 6.1.6 88
6.6 Enrolment sites data for CT BO29389 94
6.7 Country depots data for CT BO29389 94
6.8 Country depots data for CT BP30002 100
6.9 Enrolment sites data for CT BP30002 100

List of figures

3.1 Example of enrolment over time on trial 30
3.2 Distribution of coefficient of determination on enrolment sites 31
3.3 Arrival rates distribution 33
3.4 Distribution for trial length 34
3.5 Distribution of number of patients for a clinical trial 35
3.6 Proportion of trials with exponentially distributted inter-arrival times, by country 36

4.1 General structure of a clinical trial 38
4.2 Scheme of a Scenario Tree 39

6.1 Structure of the SC Network in CT WC28325 71
6.2 Stability Test (WC28325) 74
6.3 Spread Stability Test (WC28325) 75
6.4 Average solving time by stage for each sample size 75
6.5 Evolution of the estimation made by the model for trial completion time as the enrolment evolved over time 76
6.6 Progress of enrolment in trial WC28325 77
6.7 Evolution of the estimations made by the model as the enrolment evolved over time 77
6.8 Relative error density plots 80
6.9 The model that optimizes after clustering scenarios perfectly reproduces estimations made by the original mode 81
6.10 The model that optimizes after clustering scenarios follows pretty much the same pattern both in cost and overage inventory estimation 82
6.11 Comparison of our two models against two benchmarks 83
6.12 Impact of overestimated arrival rates on completion time and cost estimation . . 86

6.13 Comparison of results with right arrival rates and '×3' overestimated arrival rates 87
6.14 Change in initial inventory country by country, under arrival rates changes simulated in section 6.1.6 . . . 87
6.15 Performance comparison of model with right arrival rates and model described in 6.1.6 . . . 89
6.16 Comparison of evolution of cost and overage inventory estimation when rates deviate as described in 6.1.6 . . . 90
6.17 Comparison of number of shipments from CW to each DC, with right arrival rates and with wrong arrival rates as described in 6.1.6 . . . 90
6.18 Evolution of total cost and overage inventory estimation for clinical trial WC28325 with a fixed production cost of €10,000 . . . 91
6.19 Comparison WC28325 with and without fixed production cost . . . 91
6.20 Clinical Trial BO29389 . . . 93
6.21 Stability Test (BO29389) . . . 93
6.22 Spread Stability Test (BO29389) . . . 95
6.23 Estimation made of Completion Time for trial BO29389 . . . 96
6.24 Evolution of the estimation made for cost and ending inventory in clinical trial BO29389 . . . 97
6.25 Density plots to determine an adequate number of clusters for CT BO29389 . . . 97
6.26 Evolution of estimations for key indicators of CT BO29389 using the S40K10 model . . . 98
6.27 BO29389: Comparison of our models against 2 benchmarks . . . 99
6.28 Clinical Trial BP30002 . . . 100
6.29 Evolution of estimations for key indicators of CT BP30002 using the S40K10 model 102
6.30 BP30002: Comparison against 2 benchmark models . . . 102
6.31 Evolution of estimation of cost and overage inventory for CT BP30002 with shelf life of 150 days . . . 104
6.32 Distribution of cost for units produced in each production batch in clinical trial BP30002 . . . 104

List of acronyms

ALT Actual Lead Time

BD Benders Decomposition

CD Country Depot

CT Clinical Trial

CW Central Warehouse

CTSC Clinical Trial Supply Chain

DND Do Not Dispense

DNS Do Not Ship

DNT Do Not Transfer

FIFO First In, First Out

IMP Investigational Medical Product

LHS Left Hand Side

MILP Mixed Integer Linear Programming

MINLP Mixed Integer Non-Linear Programming

MISLP Mixed Integer Stochastic Linear Programming

MLT Model Lead Time

MSSP Multi-Stage Stochastic Program

MOO Multi-Objective Optimization

NAC Non Anticipative Constraints

NNLS Non-Negative Least Squares

OLS Ordinary Least Squares

PAM Partition Around Medoids

PHA Progressive Hedging Algorithm

RHS Right Hand Side

SP Stochastic Programming

SC Supply Chain

Chapter 1

Introduction

1.1. Context and Motivation

According to World Health Organization (WHO), clinical trials (CT) are *'a type of research that studies new tests and treatments and evaluates their effects on human health outcomes'*, where *'people volunteer to take part in clinical trials to test medical interventions including drugs, cells and other biological products, surgical procedures, radiological procedures, devices, behavioural treatments and preventive care'* (WHO 2023).

As far as is known, the first time that a clinical trial guided a public health decision is as antique as 562 BCE, and was conducted by King Nebuchadnezzar of Babylon. More recently, in XVIII century, Dr. James Lind conducted the first controlled clinical trial (Mitchel et al. 2022). Since then, clinical trials have been tightly regulated and have been increasingly going global over the past 4 decades, given that a multi-country setting efficiently pools resources to provide faster recruitment and more generalisable results across patient populations, ethnicities and disease management paradigms (IM and I 2015).

This dissertation investigates production, inventory and distribution decisions made within clinical trial supply chains (CTSC) in order to reduce both time to market and drug supply costs.

Before a novel drug can be widely used in patients, clinical trials are mandatory. In these studies, of scientific nature, the effect of a new investigational medical product (IMP) is examined in a very controlled setting. These patients typically constitute a representative sample of the population to be treated (Thiers 2006) and, since it is a small sample of volunteers, a combination of ethical commitment and practical sense demand the service level required to

be extremely high, in the sense that drugs must be available on time for the patient to receive treatment immediately upon arrival. As stated by A. Fleischhacker and Y. Zhao 2011, *'Since patient recruitment is the typical bottleneck in conducting clinical trials, a shortage of clinical drug is considered unacceptable'*. In summary, clinical trials face increasing challenges as a result of the drive to bring new drugs to market faster (Peterson et al. 2004).

From lab to shelf, drugs traverse a long journey that lasts around 10-15 years and costs €1 bn to 5 bn (Tohidi 2020). Before reaching the market, three phases of clinical trials are usually required and in each phase, the drug candidate is tested against either a placebo and/or an already commercialized medication: i) First, a few patients (typically 50 to 100 healthy individuals) are tested in order to determine the toxicity of the intervention, ii) If phase I is successful, a greater sample (a few hundreds of potential patients) is tested to determine the optimal way to administer the new drug, iii) safety and efficacy of the optimized treatment are tested in a greater sample (typically a few thousands patients) only if phases I and II are successful. Phase IV takes place after approval, in order to collect information about effects en specific populations.

The most challenging issues in managing today's clinical trials are patient recruitment and securing drug supply. Patients are enrolled into trials through both doctors and hospitals that agree to aid the study.

As a result of slow patient recruitment and increasing demand for clinical development, clinical trials are increasingly going global, and this globalization implies that the needs of multiple international regulatory bodies are now to be taken into account, with the consequent complexities associated.

The globalization of clinical trials has also meant that this organizational effort needs to be taken not only in traditionally wealthy nations, but progressively more and more in emergency economies, where the infrastructure presents numerous challenges, leading also to an increase in costs (Thiers 2006), longer lead times and a considerable higher organizational effort.

As of May 2023, the number of registered clinical trials globally was 24% greater than the total

reported in 2021 (Statista 2023). The number of drugs in pipeline has also been increasing at an average rate of 6.2% for the past 2 decades, i.e. it doubles every 11.5 years.

While in the past pharmaceutical companies considered drug supplying costs to be negligible, it is now recognized that these costs can absorb 20% of the final value of a drug. The potential savings in this area are now seen by pharmaceutical companies as valuable sources for financing R&D.

A key fact about clinical trials is that one never knows where or when patients will show up. This, combined with the stringent level of service required, lives at stake and complex importing procedures puts a lot of pressure on CT Management, forcing them to produce a considerable excess of drug units, as big as twice as much as necessary (A. Fleischhacker, Ninh, and Y. Zhao 2015), which sets an essential trade-off not only because of the financial impact, but mainly because of the increasing pressure over companies to embrace sustainability.

During the past decade, clinical supply managers are being asked to cut costs, as clinical supply costs can potentially account for a up to 40% of total clinical trial spending (A. J. Fleischhacker and Y. Zhao 2013). With budgetary pressures due to eroding pharmaceutical profits, overages reduction is a logical initiative, and CT Managers are expected to do more with less (A. Fleischhacker, Ninh, and Y. Zhao 2015).

This thesis is motivated by the real life problem of a major global pharmaceutical firm, with a cabal interest on developing new drugs and testing them in the most agile possible way while meeting every regulatory, technical, clinical and ethical requirement, but also expressing interest on cutting time to market while reducing the leftover inventory and saving costs. As stated by (Newton 2023), *'the environmental goals of Environmental, Social, Governance (ESG) have shifted from a lofty ideal to an economic imperative if pharma is to maintain its current pace of innovation'*, and *'Sustainability means less costs, less drug waste, less transport, less CO2 emissions, and, ultimately, more drugs to patients'*. In this sense, this collaboration intends to save lives in first place, at the minimum possible cost (in order to feed the virtuous circle of R&D), and with the minimum possible environmental impact.

1.2. Objectives and Contributions

There are three main goals:

- Provide a mathematical background to faithfully represent Clinical Trial Supply Chains (CTSC).
- Highlight the adequateness of Stochastic Programming as a valid approach to optimize decisions under uncertainty in a multi-stage multi-echelon setting.
- Provide a solid procedure for scenario reduction in multi-stage multi-echelon systems, with minimal accuracy loss and considerable efficiency gain in terms of required time to find a solution.

The literature encountered reveals that previous efforts in approaching the CTSC problem have, with little exceptions, focused in cost reduction, neglecting both the objective of completing enrolment soon and the minimization of environmental impact. More importantly, existing literature tends to oversimplify the complex CTSC problem by at least one of the following:

- Neglecting perishability, which is a distinctive trait of pharmaceutical products.
- Either assuming that patients who enrol will complete the trial, or modeling dropouts like a fixed proportion independent of the supply chain decisions.
- Assuming an inflexible production process where all the necessary inventory is produced in a single run and before the CT starts; we have found this to be an usual practice rather than a hard constraint.
- Assuming a stationary inventory policy, even when demand is non-stationary.
- Assuming a one dose treatment, which is the exception, rather than the most frequent case.

We present a model where cost reduction, overage inventory reduction and the need of finishing the clinical trial as soon as possible are jointly considered. Furthermore, we consider the limited shelf life of pharmaceutical products, the limited patience of patients who wait for

treatment, multiple production runs and a treatment consisting of multiple doses. As expected, all these considerations come at expenses of a considerable computational effort: we have overcome such a challenge by presenting a clustering based scenario reduction method, which relies on an original problem-based dissimilarity metric.

1.3. Methodological Approach

The main tools that we use in this research are:

i. Statistical Analysis

ii. Mathematical Programming & Multi-Objective Decision Making

iii. Clustering / Scenario Reduction

iv. Supervised Machine Learning (Regression)

v. Simulation

As the main goal is the optimization of production and distribution decisions in a Clinical Trial Supply Chain (CTSC), mathematical models will be presented, being the main one a multi-objective mixed-integer linear stochastic program (MILSP). It considers the network structure, lead-times and patient arrival rates as inputs. Based on the expected delays and associated costs, it determines how much should be produced, and when, as well as the right quantities and moments to deliver both from Central Warehouse (CW) to country depots (or enrolment sites, if they are served directly from the CW), and from country depots to enrolment sites.

Due to dimensionality issues, as is usual in stochastic programming (SP), a stability test is conducted in order to determine the right sample size to preserve tractability with minimal accuracy loss. Also, the model that we will present in Chapter 4 is multi-stage, in the sense that decisions are enriched as new information is incorporated while the clinical trial evolves over time. Each decision step is called a stage. Since we have to solve as many stochastic programs as stages there are in the problem, we need to perform additional scenario reduction:

we apply a *k-medoids++* technique with a problem based dissimilarity metric with 3 terms whose relative contribution we must first determine using linear regression. With the relative weights of the 3 terms already determined, the dissimilarity metric is defined and we apply scenario reduction by clustering over the set of scenarios obtained by conditional sampling, stage by stage, until the clinical trial (CT) is complete.

Statistical analysis was used to gain insights on the arrival patterns, in order to perform scenario generation when applying our model in Chapter 6, where we also applied simulation to test the performance of the model and compare with a benchmark.

1.4. Research Questions

The experimentation phase is then proposed, guided by a set of research questions defined as follows:

1. What is the impact of a flexible production process in the efficiency of the CT in terms of time, cost and overage inventory?
2. How much does the enforcement of a stationary policy impact on the CT in terms of time, cost and overage inventory?
3. What is the most adequate clustering technique and how should the dissimilarity metric be defined for performing scenario reduction with minimal accuracy loss?
4. What is the value generated by the optimization model when compared to industry benchmarks and other policies found in the literature?

1.5. Outline of the thesis

This thesis is organized as follows: in chapter 2 we systematically review the relevant literature related to Clinical Trial Supply Chains, Stochastic Programming methods and Scenario Reduction techniques; in chapter 3 we perform statistical analysis in order to extract insights from clinical trials data; in chapter 4 we present a multi-objective multi-stage stochastic programming model for the CTSC problem; in chapter 5 we present a scenario reduction procedure to accelerate the solution while preserving as much accuracy as possible, including a novel

problem-based metric approach suitable for any multi-echelon multi-stage stochastic program; in chapter 6 we apply our model and scenario reduction techniques to different settings inspired in actual clinical trials, and compare the results with other existing models; in chapter 7 we present a couple extensions to the model: one for multi-product and one for policy enforcement; and in chapter 8 we conclude by summarizing our findings and discussing the overall results, tracing back to the CTSC and Stochastic Programming literature. At the end of each chapter, we provide intermediate conclusions, address limitations, and provide directions for future research.

Chapter 2

Literature Review

Inventory Management in multi-echelon and multi-period settings with stochastic demand has attracted considerable attention during the past decade, and produced a variety of creative solutions, with different approaches like non-linear optimization (Shu, Li, and Huang 2013), robust optimization (Aharon, Boaz, and Shimrit 2009) (Wang and Yang n.d.), genetic algorithm, (Zhou, L. Chen, and Ge 2013), miopic policies (Goh and Porteus 2016) (Zhang, Chai, and Ma 2021), and others.

The problem of inventory management in CTSC, on the contrary, has attracted very limited attention from both the industry and academia. More generally, perishable goods have been receiving little attention in comparison with their counterpart, the non-perishable goods, although there have been some interesting efforts that contributed to understand the convenience of explicitly including the perishability property in the formulation, rather than using the classical inventory management techniques that work well for non-perishable: (Haijema 2013), for instance, have used Dynamic Programming (DP) to derive an interesting time-dependent (s, S, q, Q) periodic review policy. As indicated by (Nahmias 1982), products have either a fixed shelf life or a random one. When dealing with non-fixed lifetimes, things can get significantly complex, as deterioration rate could depend on age or on inventory level (Bakker, Riezebos, and Teunter 2012). Luckily for us, a fixed lifetime is suitable for products highly regulated like the ones in pharmaceutical industry, where a product gets a best-before or use-by date on the package, reflecting this fixed maximum shelf life (Pauls-Worm 2016).

While a lot is known and has been written about hypothesis testing to asses the goodness of fit of empirical data (Horn 1977; D'Agostino and Stephens 1986), the literature consulted has dedicated little or none effort to challenge the usual assumption of patient arrivals following a

Poisson process. Some authors (H. Zhao et al. 2019; A. Fleischhacker, Ninh, and Y. Zhao 2015) have assumed a non-stationary Poisson process, and others have preferred the use of empirical distributions (Y. Chen, Mockus, et al. 2012).

The complexity of interactions in this type of system has led to simulation techniques (Peterson et al. 2004; Parke 2011; Ninh, Lefew, and Anisimov 2019; Pereira 2006) being the most popular tool to assist an inventory policy selection. In contrast, a first analytical model was proposed by (A. Fleischhacker, Ninh, and Y. Zhao 2015), who presented a deterministic MINLP and its equivalent MILP model to minimize total cost assuming a single production run and echelon-dependent policies, finding an elegant solution for the limited subset of clinical trials meeting the rather simplistic assumptions: in particular, the premise of 1 single dose means that nearly 90% the clinical trials do not meet the criteria, while the assumption of a single production run to take place before the enrolment starts was found to be contrary to the preferences of our sponsoring company. We find also a few examples where both simulation and optimization are combined (Y. Chen, Mockus, et al. 2012; Y. Chen, Pekny, and Reklaitis 2012).

The consideration of cost minimization as only objective is shared by the vast majority of past research in this area, although it may not only lead to ethical issues, but also economical inefficiencies, as noted by (Clemento 1999). More recently, the cost of delaying one day the launch of a blockbuster medicine was estimated to reach an average $ 8 millions (Clariness 2023). Exception of this inconsideration of completion time as a crucial objective can be found in (H. Zhao et al. 2019), who have proposed a multi-objective stochastic model primarily focused on minimization of completion time. This model also considers dropouts, although they are modeled as a fixed proportion, therefore considered as exogenous and not related to decisions made in the system. The consideration of the link between the decisions taken and dropouts has been found to be a knowledge gap, to the extent of our knowledge.

As noted by (Peterson et al. 2004), traditionally the supply of medication to many centres is achieved by issuing sufficient medication to each centre to account for the requirements of all recruited patients without any resupply from a central depot. However if this is done, it is inevitable that either demand will be unmet or a considerable excess of units will be produced

only to be disposed at the end of the trial. (A. Fleischhacker and Y. Zhao 2011) have brought attention to multiple production runs as an element for controlling for demand failure, in contrast with the usual single production run approach (Y. Chen, Mockus, et al. 2012) that the same authors apply in (A. Fleischhacker, Ninh, and Y. Zhao 2015).

Other important aspects of CTs have been receiving attention and approached thru the lens of mathematical optimization, like the adequate number of clinical sites (Rubio-Herrero, Ninh, and Lefew 2023), outcome analysis (Colvin and Maravelias 2010) and pipeline planning (Miller 1996, Tohidi 2020).

A good introduction on Stochastic Programming can be found on (Birge and Louveaux 1997) and (Ruszczyński and Shapiro 2003).

SP application in Supply Chain has been rather limited: different works depict the application of SP in Supply Chain problems as conceptually enriching, although applied with limited results due to the essential challenges to preserve tractability as a certain number of scenarios is used in order to model reality with an acceptable level of detail. Fahimnia and Jabbarzadeh 2016 and Manopiniwes and Irohara 2017 have highlighted the power of SP to address problems where anticipated decisions are to be taken in the presence of uncertainty while dealing with essential trade-offs. (Fattahi, Govindan, and Keyvanshokooh 2018), (Azizi and Hu 2021), (Lima, Relvas, and Barbosa-Póvoa 2018) and (Fattahi and Govindan 2018) have shown its applicability in multi-stage problems and multi-echelon problems.

The level of complex interactions present in the CTSC has been approached with a 2-stage SP approach in (H. Zhao et al. 2019), (W.-A. Chen 2019) and (Zheng et al. 2021); in the first of them, the authors deal with the course of dimensionality using PHA, but their numerical results are limited to rough bounds with heavy dependence on the penalty parameter.

The problem of sampling and stability is well addressed by (Heitsch, W. Römisch, and Strugarek 2006) and (Shapiro 2003), with key contributions regarding the importance of conditional sampling in multi-stage problems. The problem of the essential trade-off between stability and tractability is extensively treated in (Linderoth, Alexander Shapiro, and Wright

2006). (Gangammanavar and Sen 2021) provides an algorithm for conditional sampling in order to perform sequential optimization to get a consistent estimator for the multi-stage SP.

The generalities of scenario reduction methods have been extensively addressed in (Werner Römisch 2009) and (Heitsch and W. Römisch 2009). In particular, (Beraldi and Bruni 2014) have successfully applied clustering methods in a systematic approach to prune scenario trees with minimal information loss. (Henrion and Roemisch 2018) and (Keutchayan, Ortmann, and Rei 2021) are good examples of problem-based scenario reduction methods, in contrast to general reduction methods that do not rely on the structure of the problem at the time of pruning the scenario tree. However, to the extent of our knowledge, no previous research has used or suggested a problem-based dissimilarity metric, like we do in this thesis.

Scenario reduction methods have been little applied in SC problems, being (Heitsch and W. Römisch 2009), (Dillon, Oliveira, and Abbasi 2017) and (Paulo et al. 2017) rare exceptions.

K-medoids algorithm is perfectly explained in (Kaufman and Peter J Rousseeuw 2009). Its convenience for scenario clustering has been addressed by (Dupačová, Groewe-Kuska, and Roemisch 2003) and (Keutchayan, Ortmann, and Rei 2021). Different python implementations have been sufficiently explained in (Schubert and Lenssen 2022) and (Schubert and Peter J. Rousseeuw 2021).

Tables 2.1 and 2.2 summarise some of the important papers reviewed and the main characteristics collected from close related publications, with the first table containing Clinical Trial Supply Chain related papers, and the second one covering Stochastic Programming Scenario Reduction related research. These summaries contribute to reinforce that there is no knowledge up to now of other papers that consider the CTSC problem with the level of detail that we will present in this thesis, neither about previous efforts towards defining an adequate scenario reduction technique for Multi-Stage Stochastic Programming problems.

Title	Authors	Journal	Year	Stochastic Arrivals	Multi-Objective	Considers Perishability	Allows Multiple Production Runs	Multi-Dose	Considers Dropouts	Links Dropouts to Decisions	Multi-Product
Clinical Trials Supply Chain Optimization	Calcinaro, D.	ZLC (PhD Thesis)	2023	✓	✓	✓	✓	✓	✓	✓	✓
Simulation-optimization approach to clinical trial supply chain management with demand scenario forecast	Chen, Y. Mockus, L. Orcun, S. Reklaitis, G.	Computers & Chemical Engineering	2012	✓		✓	✓	✓	✓		✓
Positioning Inventory in Clinical Trial Supply Chains	Fleischhacker, A. Ninh, A. Zhao, Y.	POMS	2015	✓							
A multi-objective production planning problem with the consideration of time and cost in clinical trials	Zhao, H. Huang, E. Dou, R. Wu, K.	Expert Systems with Applications	2019	✓	✓			✓	✓		
A Study on the Optimal Inventory Allocation for Clinical Trial Supply Chains	Zheng, M. Du, N. Zhao, H. Huang, E. Wu, K.	Applied Mathematical Modeling	2009	✓				✓			
Drug Supply Chain Optimization for Adaptive Clinical Trials	Chen, W.	Purdue University (PhD Thesis)	2019	✓			✓	✓	✓		✓

Table 2.1: *Clinical Trial Supply Chain Management related papers*

Title	Authors	Journal	Year	Problem Based	Clustering Based	*Ad-hoc* Distance Metric	Suitable for MSSP
Clinical Trials Supply Chain Optimization	Calcinaro, D.	ZLC (PhD Thesis)	2023	✓	✓	✓	✓
A clustering approach for scenario tree reduction: an application to a stochastic programming portfolio optimization problem	Beraldi, P. Bruni, M.E.	SEIO	2013	✓	✓		
Decision-based scenario clustering for decision making under undertainty	Hewitt, M. Ortmann, J. Rei, W.	Annals of Operations Research	2022	✓	✓		
Problem-based optimal scenario generation and reduction in stochastic programming	Henrion, R. Römisch, W.	Mathematical Programming	2022	✓			
Scenario Reduction Algorithms in Stochastic Programming	Heitsch, H. Römisch, W.	Computational Optimization and Applications	2003	✓			✓
Problem-Driven Scenario Clustering in Stochastic Optimization	Keutchayan, J. Ortmann, J. Rei, W.	Computational: Management Science	2023	✓	✓		
Scenario Reduction Techniques in Stochastic Programming	Römisch, W.	Stochastic Algorithms: Foundations and Applications	2009	✓			

Table 2.2: *Stochastic Programming Scenario Reduction related papers*

Chapter 3

Statistical Analysis

3.1. Industry Data and Descriptive Statistics

Our data set, provided by the sponsoring company, contains 178,816 records and 53 columns with data for 2852 clinical trials, collected over 7 years (January 2013 to October 2019), involving 93 countries and 45,593 enrolment sites. Relevant data provided includes:

- Clinical Trial Code
- Country
- Enrolment site code
- Date
- Number of people enrolled
- Disease area
- Indication
- Phase
- Type of control

After cleaning data removing incomplete or unreliable records we were left with 21,662 sites over 572 different trials of phases I, II and III, covering from 20 to 2155 enrolled patients, from 1 to 41 country depots, from 1 to 394 enrolment sites, and up to 2,249 days completion time. If we use the median to summarize these attributes, a typical clinical trial in this sample enrolls 95 patients in 113 days, over 16 sites distributed in 3 countries.

Further cleaning was necessary to perform statistical analysis that could lead to significant findings. For these purposes, we focused on 2,132 sites over 405 clinical trials. This sample is the result of cleaning the data by:

- Removing trials of phase IV or unknown: This is because we would expect trials in phase IV to follow a much more irregular arrival process, since it does not require enrolment of patients in a clinical site, and therefore they do not represent the type of process that is object of this research.
- Removing trials where there is an observation about uncertain data
- Removing trials taking less than 2 months
- Removing trials where the number of records is less than 10.

3.2. New patients arrival pattern

3.2.1. Linearity

Our sponsoring company has expressed concerns about the linearity of the number of enrolled patients over time, in order to be able to extrapolate. We refer here to the new patients arrivals, i.e. patients coming for a first dose of the treatment. We have found that, at clinical trial level, 90% of trials exhibit $R^2 \geq 0.8$, and it is $R^2 \geq 0.9$ for the best 73%.

Figure 3.1 shows an example of a double blind phase II trial for breast cancer treatment that enrolled 151 patients over 2 years.

Even dis-aggregating to site level, we have 90% with $R^2 > 0.745$, and 54% with $R^2 \geq 0.9$ (Fig 3.2).

3.2.2. Relationship between Poisson behaviour and linearity

It is important to remark that the model proposed in this research, to be described in Chapter 4, does not rely on the patients arrival process following any specific distribution. With this being said, it will come a point where we will have to assume some distribution in order to generate scenarios to asses the performance of our model in Chapter 6. We have here dedicated some effort to asses the suitability of Poisson distribution mainly for two reasons:

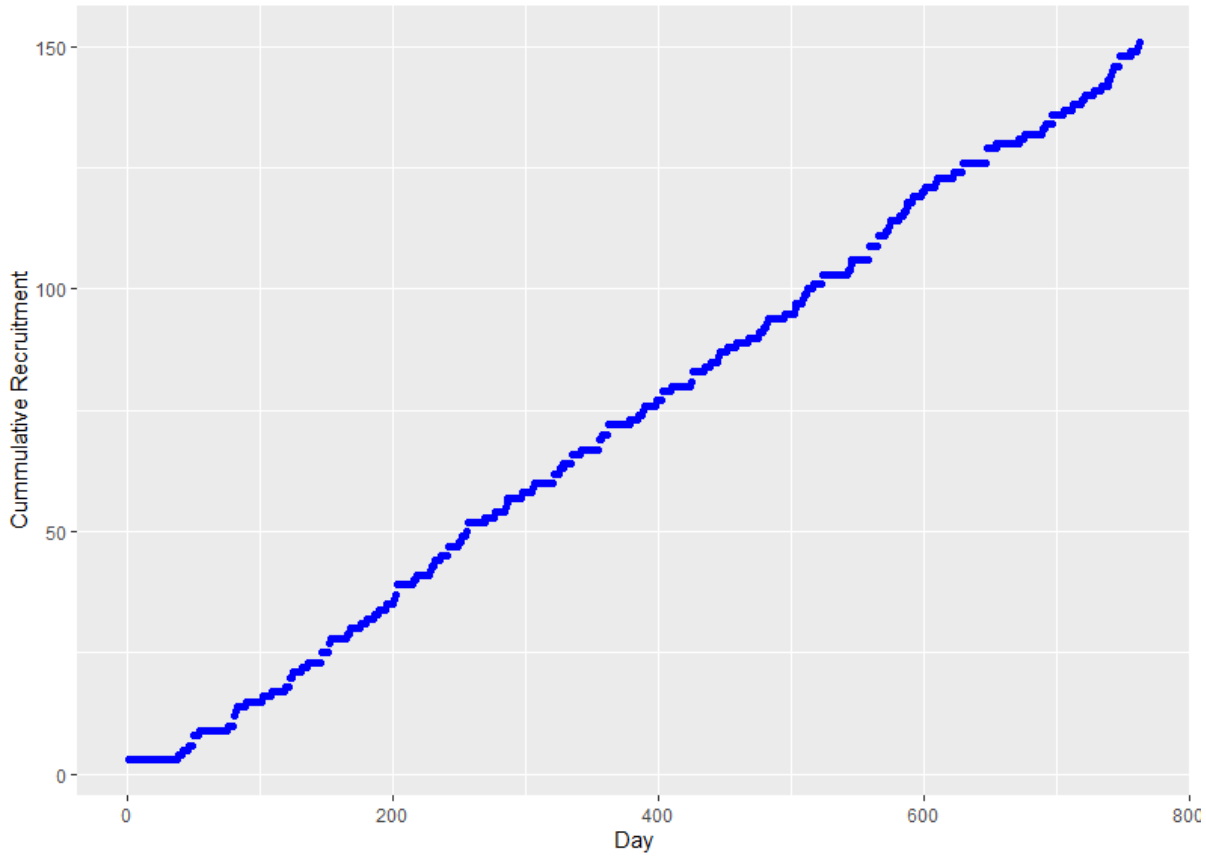

Figure 3.1: *Example of enrolment over time on trial*

1. The vast majority of the literature on CTSC assumes that patients arrive to enrolment sites following a Poisson distribution, i.e. that inter-arrival times follow an exponential distribution.

2. The exponential distribution is the only continuous distribution that holds the memory-less property, and therefore it is the expectable distribution of inter-arrival times of events that take place over time with independence of each other. Clinical sites are chosen based on their ability to enrol patients who meet the eligibility requirements, and they are typically geographically dispersed, drawing patients from distinct populations, and therefore the assumption of recruitment process in a site being independent from the one in another site is in general well justified

The positive findings of about a reasonably high coefficient of determination for cumulative enrolment over time intuitively suggests that there exists a pretty stable arrival rate over time, which is what we would expect from a Poisson process.

In this section we will derive a formula for the coefficient of determination that results

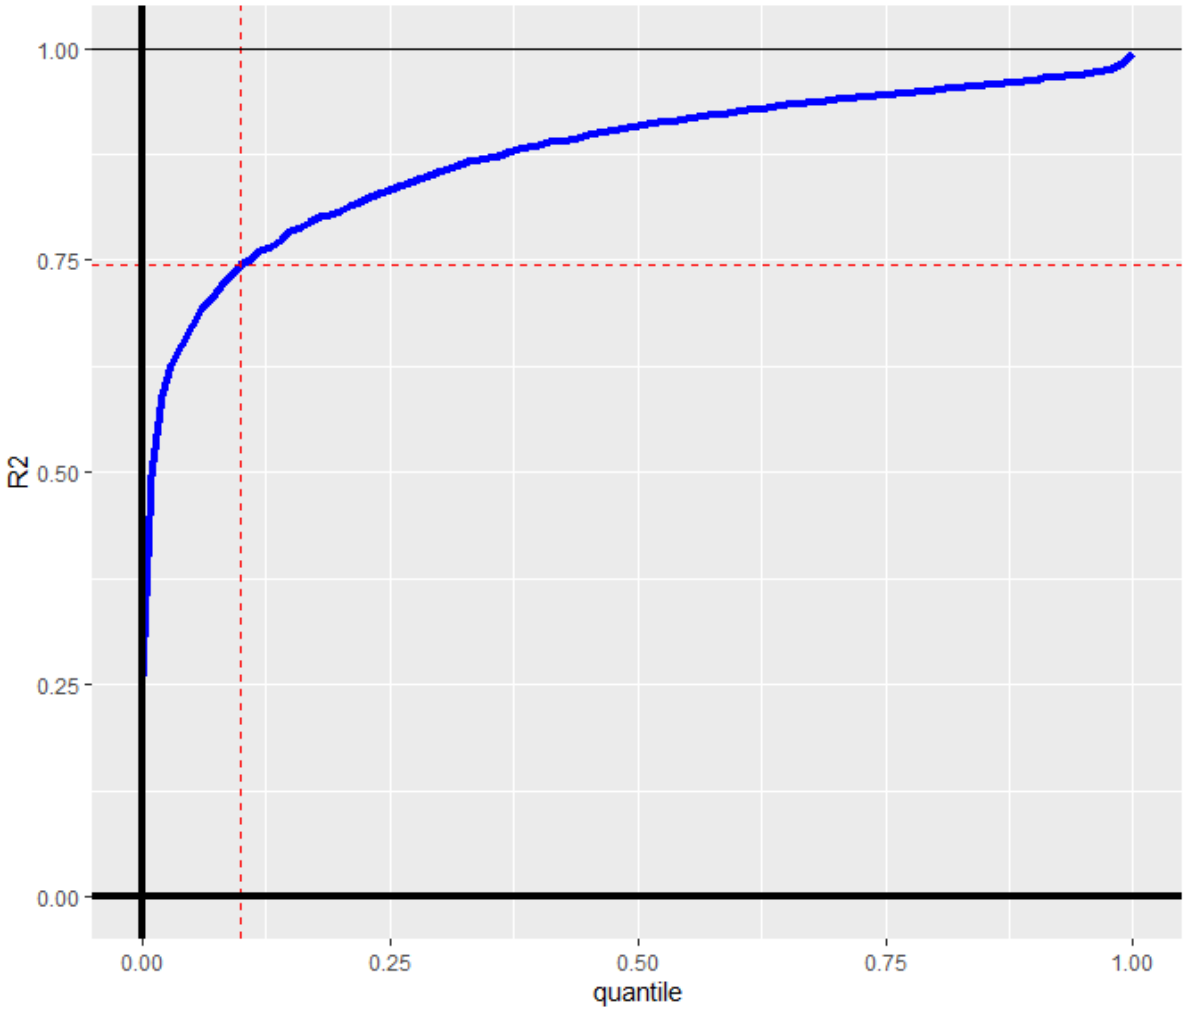

Figure 3.2: *Distribution of coefficient of determination on enrolment sites*

from enrolments over time assuming a Poisson process with rate λ taking place over a discrete period of time that ranges from $t = 1$ to $t = L$. We then have, for the horizontal coordinate $x \in \{1, 2, ..., L\}$, while the vertical coordinate follows a Poisson distribution with rate $\lambda \cdot x$, component-wise. By definition, we have:

$$R^2 = \left(\frac{cov(x, y)}{\sigma_x \sigma_y} \right)^2 \tag{3.1}$$

From the definition of mean and variance of a discrete uniform variable over a range [1,L], it follows that $E[x] = \frac{L+1}{2}$ and $\sigma_x^2 = \frac{L^2-1}{12}$. To take the first two momentum for y, we need to use Law of total expectation and Law of total variance, respectively:

$$E[y] = E\left(E[y|x]\right) = E[\lambda \cdot x] = \lambda \cdot E[x] = \lambda \cdot \frac{L+1}{2} \tag{3.2}$$

$$\begin{aligned}\sigma_y^2 = E[Var(y|x)] + Var[E(y|x)] = E[\lambda \cdot x] + Var(\lambda \cdot x) = \\ = \lambda \cdot \frac{L+1}{2} + \lambda^2 \cdot Var(x) = \lambda \cdot \frac{L+1}{2} + \lambda^2 \cdot \frac{L^2-1}{12} = \\ = \frac{\lambda \cdot (L+1)}{12} \cdot [6 + \lambda \cdot (L-1)]\end{aligned} \tag{3.3}$$

In order to bring clarity to the previous derivation, let us take into account that the variable $y|x$ follows a Poisson distribution with rate $\lambda \cdot x$.

For the covariance, we will use the well known formula $cov(x, y) = E[x \cdot y] - E[x] \cdot E[y]$. For the first term, let us use again the law of total expectation:

$$\begin{aligned}E[x \cdot y] = E\left[E[(xy)|x]\right] = E\left[\sum_{k=1}^{L} \left(k \cdot y|_{y=k} \cdot \Pr[x = k]\right)\right] = \\ E\left[\sum_{k=1}^{L} \left(k \cdot y|_{y=k} \cdot \frac{1}{L}\right)\right] = \frac{1}{L} \cdot E\left[\sum_{k=1}^{L} \left(k \cdot y|_{y=k}\right)\right] = \\ \frac{1}{L} \sum_{k=1}^{L} \left[k \left(E[y|_{x=k}]\right)\right] = \frac{1}{L} \sum_{k=1}^{L} [k \cdot (\lambda k)] = \frac{\lambda}{L} \cdot \sum_{k=1}^{L} k^2 = \\ \frac{\lambda}{L} \cdot \frac{L \cdot (L+1) \cdot (2L+1)}{6} = \frac{\lambda \cdot (L+1) \cdot (2L+1)}{6}\end{aligned} \tag{3.4}$$

Using equations 3.1, 3.2, 3.3 and 3.4, and rearranging a bit, we get:

$$R^2 = \frac{1}{1 + \frac{6}{\lambda \cdot (L-1)}} \tag{3.5}$$

In line with what common sense suggests, a higher arrival rate, longer recruitment or number of patients (essentially $\lambda \cdot L$) are expected to lead to a higher coefficient of linear determination.

In our data sample, the mean for time to complete enrolment is 372 days, with an average arrival rate of 0.05 patients/day. Using these values in Eq. 3.5, should it follow a Poison distribution, a typical trial would have $R^2 \approx 0.908$, which is indeed very close to the observed average in this sample ($R^2 = 0.875$).

3.3. Memory-less inter-arrival times hypothesis testing

As explained in 3.2.2, while our model does not require the new patients arrival process to follow any specific distribution, it is of our interest to asses the validity of the popular assumption of this one being a Poisson process.

We applied Chi-Squared distribution to test the goodness of fit (Horn 1977, and D'Agostino and Stephens 1986) of the Poisson distribution for the arrivals patterns in 21662 sites, failing to reject the null hypothesis (with an average p-value of 0.693) in 19478 of them (89.9%), which indicates that it is fairly accurate to assume Poisson arrivals for the purpose of scenarios generation..

3.3.1. Drivers for non-Poisson behavior

The exploration of reasons for some trials following a non-Poisson distribution has not shown any clear driver across numerical variables, as can be seen in figures 3.3, 3.4 and 3.5.

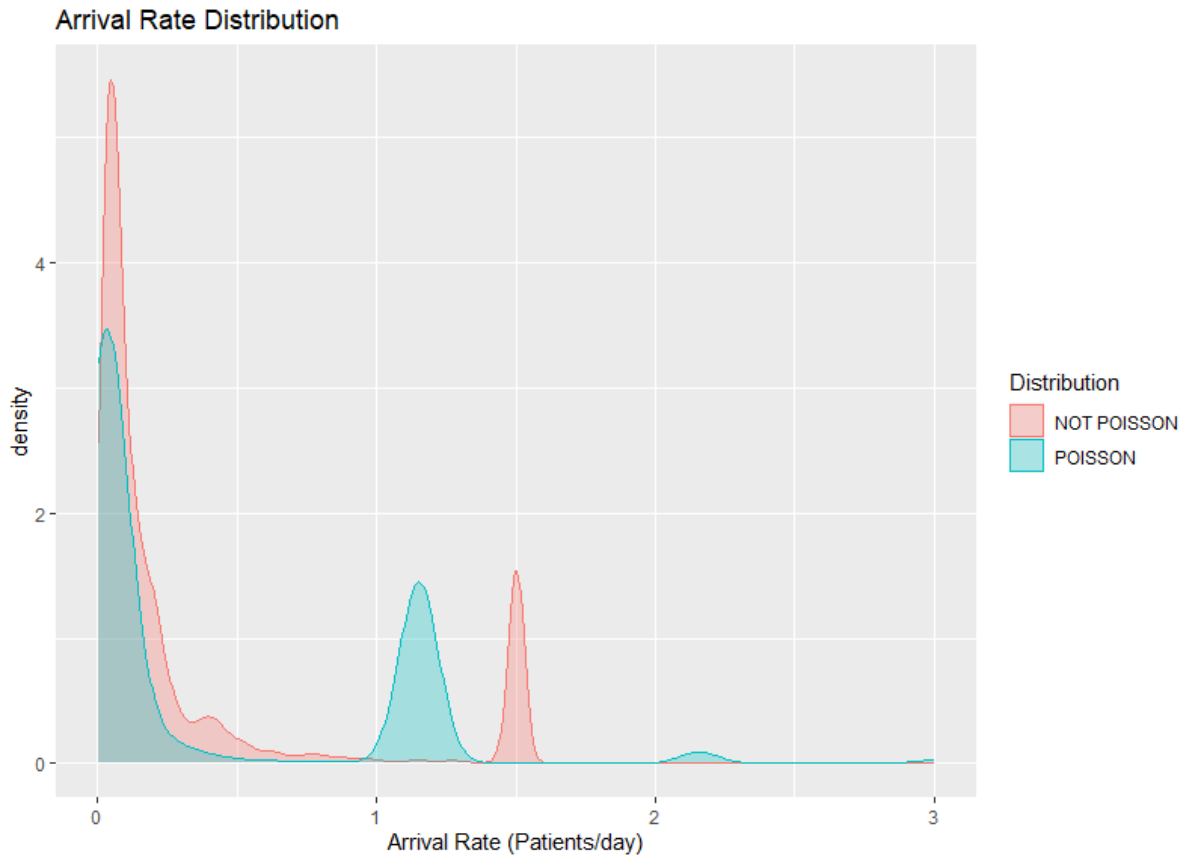

Figure 3.3: *Arrival rates distribution*

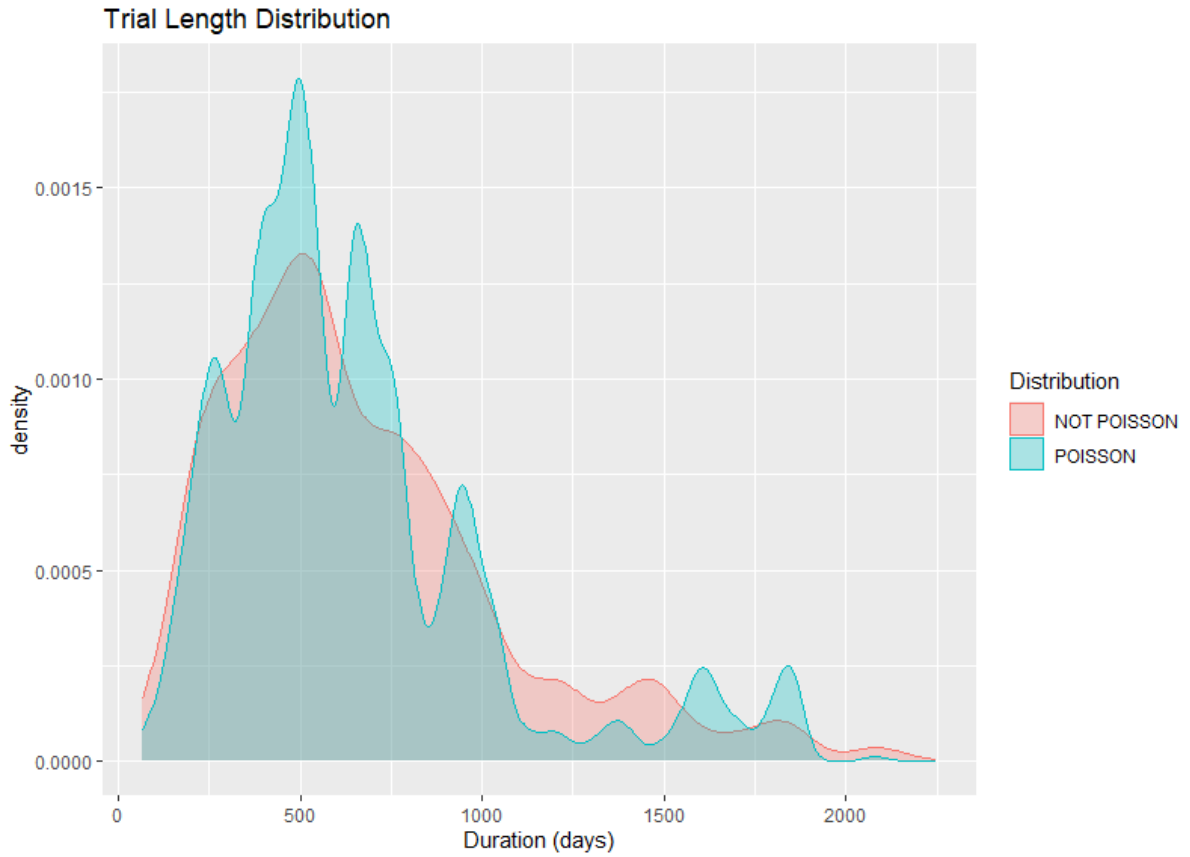

Figure 3.4: *Distribution for trial length*

However, some categorical variables appear to be considerably associated to a higher rate of deviations from Poisson behavior.

Disease area

Out of 10 different disease areas, MPS (*Mucopolysaccharidosis*) contributes to our data sample with 161 clinical studies, where the Poisson hypothesis is rejected in 44% of them, while for the remaining 9 disease areas rejection rate ranges from 6% to 20%.

Country

Among the countries that held at least 20 clinical trials in our data sample, some have consistently recruited following a pattern consistent with a Poisson distribution: we have not rejected the Poisson hypothesis for Latvia or Estonia in any of their 28 and 34 clinical sites respectively; for other Asian countries like Japan, India and Malaysia, the rejection rate has ranged between 2% and 3%. On the other side, Netherlands and Poland recruit following a non-Poisson distribution at least in 28% and 20% of their clinical trials respectively (Figure

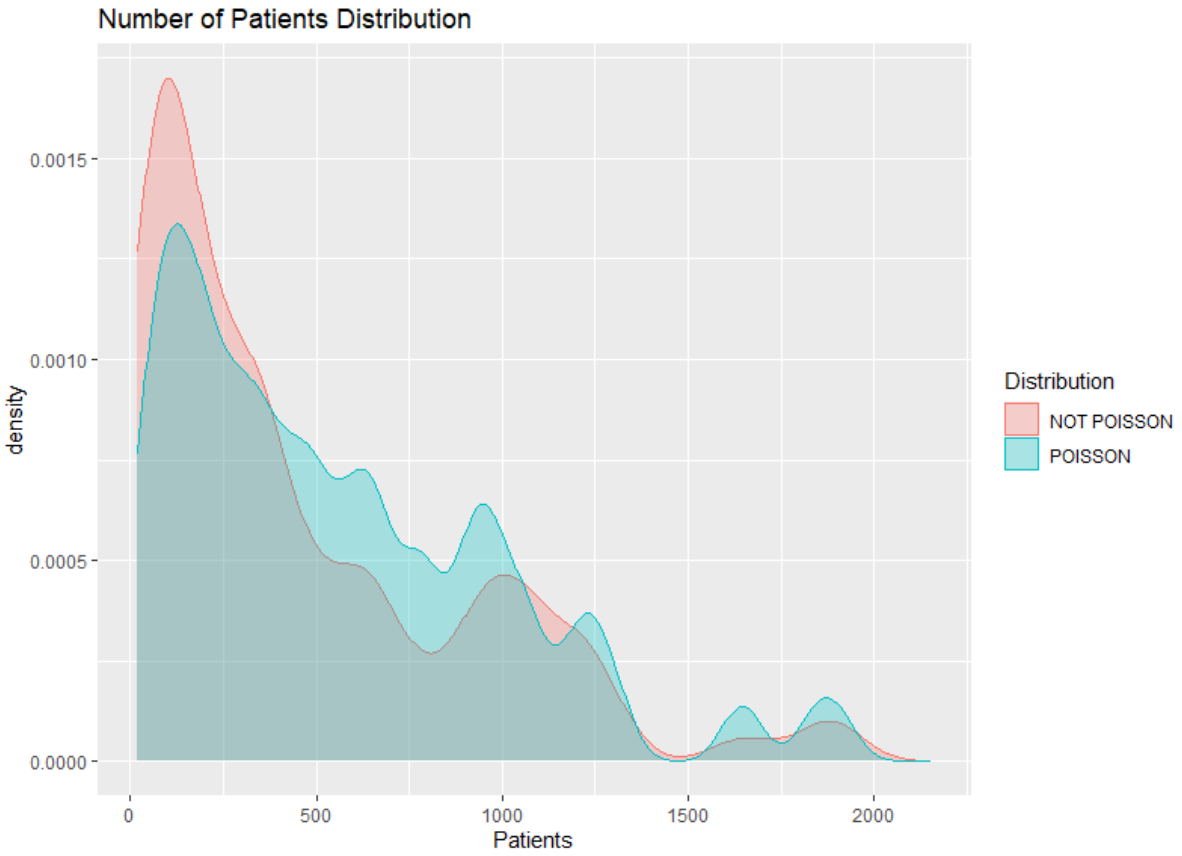

Figure 3.5: *Distribution of number of patients for a clinical trial*

3.6).

Although it is outside the scope of our research, it is our belief that some legal regulation or adaptive behavior (Chan and Green 2013) might be affecting the recruitment in certain countries.

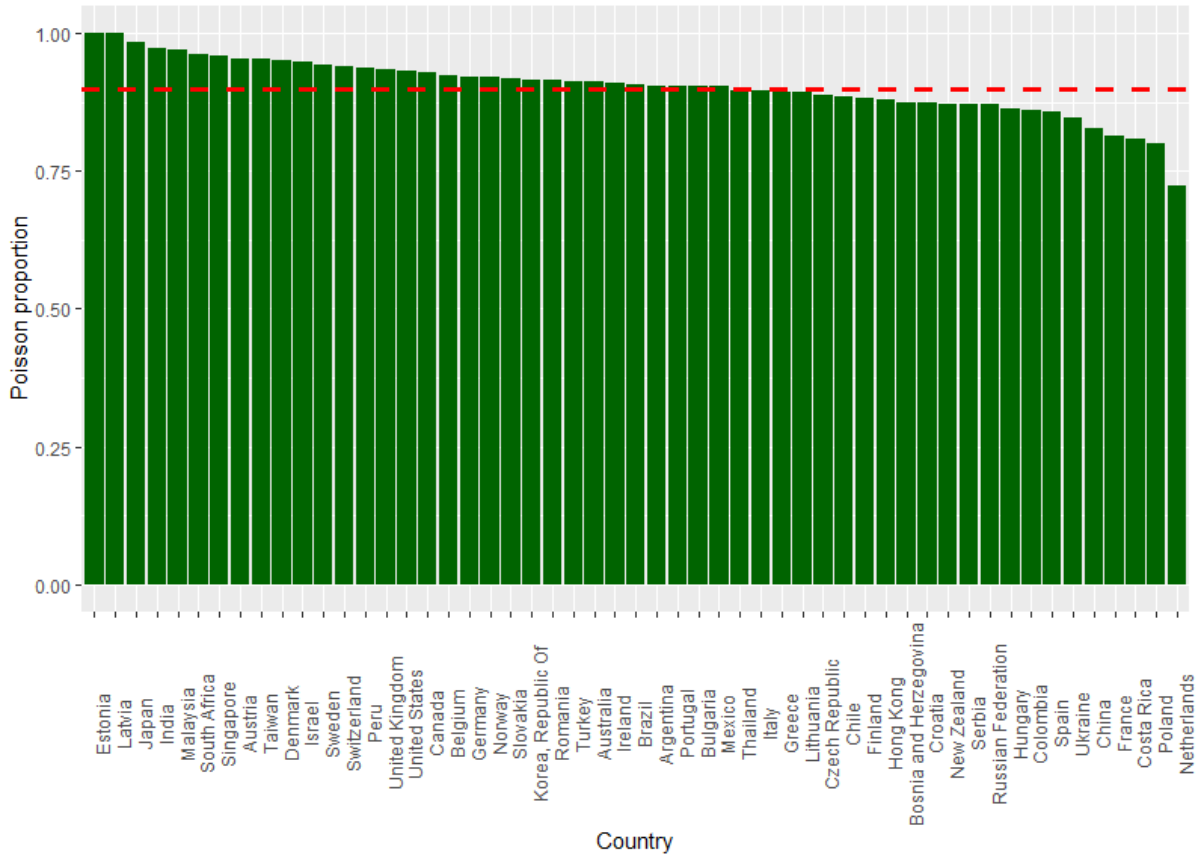

Figure 3.6: *Proportion of trials with exponentially distributted inter-arrival times, by country*

Chapter 4

Mathematical Model

4.1. Problem Statement

The problem, at a glance, is the determination of optimal production and distribution decisions in a Clinical Trial Supply Chain (CTSC), at each period of time.

By optimal decisions, we mean the ones that will lead to lower completion time, cost and overage inventory, in such a way that we perform multi-objective optimization with objectives of social, economical and sustainability nature.

As is usual with Multi-Objective Optimization (MOO), different approaches are possible, including:

- Create a single objective function that weighs the importance of each objective
- Find the Pareto-optimality boundary
- No preference methods
- Preemptive optimization (Lexicographic approach)

In the latter one, all objectives are ranked in order of importance and the optimization is performed by considering one objective at a time, based on priorities. We interpret that the priorities structure is clear enough as for exploiting the simplicity of this method.

While we understand the importance of reducing overage inventory for environmental reasons, and of course the economical objective of reducing the clinical trial cost, the existence of lives at stake combined with the potential economical benefits of reducing the time to market set

a clear priority in favour of minimizing completion time as main objective. As explained in (H. Zhao et al. 2019): *'On average, the trial delay of a single day can cause a loss of more than $1 million revenue for a typical drug. Hence, if a clinical trial can be completed earlier, the drug's commercialization time under patent protection may in- crease and the profit may increase accordingly. As a result, the clinical trial duration should be minimized to gain the most profit'*.

The preference between 2nd order and 3rd order objective is more nuanced. We have decided to rank Cost minimization as 2nd objective because pharmaceutical companies actively use a portion of their savings for environmental initiatives, which makes this objective more comprehensive.

4.2. Assumptions

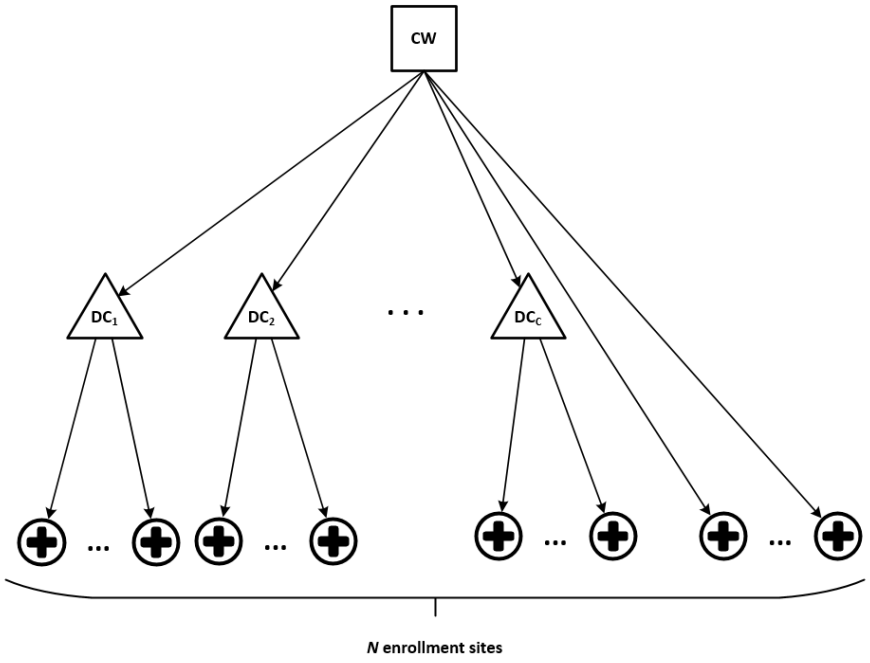

Figure 4.1: *General structure of a clinical trial*

Let us consider a clinical trial, where:

- There are N enrolment sites distributed in C countries
- An enrolment site is characterized by index j and may be served from a country depot or directly from the Central Warehouse (CW).

- Patient horizon is H: a total H patients will be enrolled to receive full treatment. If a patient drops out from the trial, then enrolment is extended to ensure that at least H patients receive full treatment.

- Treatment consists of m doses, separating each dose by τ periods.

- The duration of the trial is limited above by a sufficiently large number of periods, L.

- Patients will wait for inventory a maximum δ periods, and they will drop out if they are not filled by after waiting that time for a specific dose.

- We will assume one product for now, for simplicity. We will treat the extension to multi-product in a separate section.

- New production has a remaining shelf life SL, which is fixed and known in advance.

4.3. Stochastic characterization

A discrete stochastic process can be represented by a decision tree, where each node corresponds to a possible outcome.

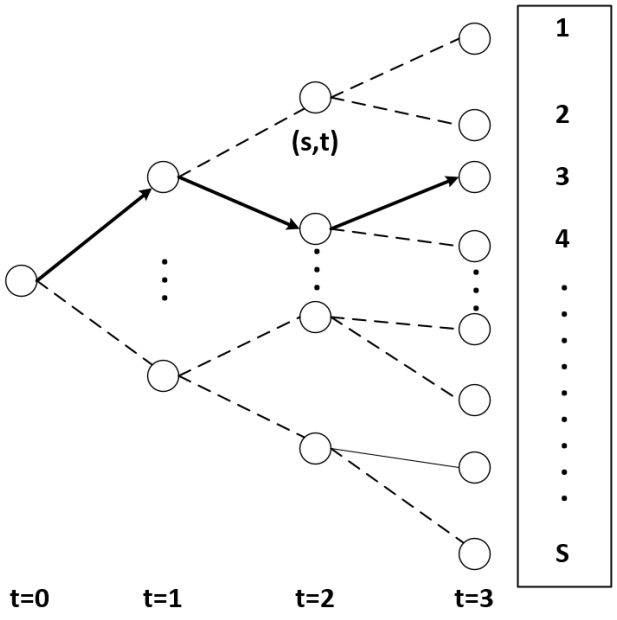

Figure 4.2: *Scheme of a Scenario Tree*

A path from the root node to the last stage L identifies a scenario, i.e. a joint realization of the uncertain parameters over all the stages (Beraldi and Bruni 2014). In Figure 4.2 we illustrate a scenario tree with 3 stages, where S different scenarios are possible, and the specific realization of scenario 3 takes place.

An important point to clarify is that stages do not necessarily relate to time periods: they correspond to steps in the decision process (Bertocchi, Moriggia, and Dupačová 2006).

Each scenario is fully characterized by the enrolment pattern, each period, at each enrolment site. Let us remark that the enrolment pattern is nothing but the new patients arrival pattern truncated when the patient horizon is reached. There is a finite number of possible realizations such that either the patient horizon H or the time limit L is reached. We will denote each of those scenarios s. If, for example, we have 4 enrolment sites, a patient horizon of 5, and a maximum of 7 periods to complete the trial, then some possible scenarios are:

	Period 1	Period 2	Period 3	Period 4	Period 5	Period 6	Period 7
Site 1	1	0	0	0	0	0	0
Site 2	0	0	1	0	0	0	0
Site 3	0	0	0	0	1	0	0
Site 4	0	0	2	0	0	0	0

indicating that 1 patient arrives to sites 1, 2 and 3 on periods 1, 3 and 5 respectively and 2 patients arrive to site 4 on day 3, meaning that the trial is completed (in case of a one-dose treatment) by the end of period 5. Another scenario, which also leads to a trial completion on day 5 could be:

	Period 1	Period 2	Period 3	Period 4	Period 5	Period 6	Period 7
Site 1	0	0	0	0	5	0	0
Site 2	0	0	0	0	0	0	0
Site 3	0	0	0	0	0	0	0
Site 4	0	0	0	0	0	0	0

representing that the 5 patients arrive on day 5 to site 1. Also, we could have:

	Period 1	Period 2	Period 3	Period 4	Period 5	Period 6	Period 7
Site 1	0	0	0	1	0	0	0
Site 2	0	0	0	0	0	1	0
Site 3	0	0	0	0	0	0	1
Site 4	0	0	0	0	0	0	0

where 3 patients in total arrived within the specified time. Let us assume that such case is excluded, because we rightly pick a sufficiently large L. We can then characterize each scenario as a $\mathbf{A}_{N\times L} = \{a_{jt}\}$ matrix, where $\sum_{j=1}^{N}\sum_{t=1}^{L} a_{jt} = H$. This means that the number of scenarios S is the number of weak compositions (see Page 2013) of H in $N \times L$ parts, so the number of scenarios is:

$$S = \binom{H + NL - 1}{H} \tag{4.1}$$

Let us note that the trial could also evolve in a way that more than H patients are initially recruited, if suddenly multiple arrivals happen, like for example:

	Period 1	Period 2	Period 3	Period 4	Period 5	Period 6	Period 7
Site 1	0	0	0	1	1	0	0
Site 2	0	0	1	0	1	0	0
Site 3	0	1	0	0	2	0	0
Site 4	1	0	0	0	3	0	0

The above example represents a case where 4 patients have enrolled over the first 4 periods, and 7 show up in the 5th period, resulting in an enrolment of 11 patients, even when the patient horizon had been set to 5. In general, the condition for $\mathbf{A}_{N\times L}$ is not $\sum_{j=1}^{N}\sum_{t=1}^{L} a_{jt} = H$ but the satisfaction of two simultaneous conditions:

$$\sum_{j=1}^{N}\sum_{t'=1}^{t-1} a_{jt} < H \tag{4.2}$$

$$\sum_{j=1}^{N}\sum_{t'=1}^{t} a_{jt} \geq H \tag{4.3}$$

for some t such that $1 \leq t \leq L$. while (4.2) and (4.3) represent the general case, using $\sum_{j=1}^{N}\sum_{t=1}^{L} a_{jt} = H$ as representation still provides an approximate representation of most arrival patterns, and (4.1) is a valid lower bound for the number of scenarios described by (4.2) and (4.3).

Each scenario is characterized by the sequence of new patient arrivals at each site in each period: $\left\{d_{\mathrm{ij,t}}^{(\mathrm{s})}\right\}$. Let us assume that we know the probability that, on a specific site, and a specific period, a certain number of patients will arrive, i.e. we know $p_{\mathrm{ij}}(d_{\mathrm{ij,t}})$. If, for instance, we assume that, at each site, new arrivals take place in each period independently of other periods, we can then calculate the probability of each of these S scenarios:

$$p_s = \prod_{ij,t} p_{\mathrm{ij}}(d_{\mathrm{ij,t}}^{(\mathrm{s})})$$

where $\sum_{t=1}^{L} \sum_{ij} d_{\mathrm{ij,t}} = H$.

Let us take into consideration that this is the probability of all the possible evolutions of the system when we run the model from the first time, at the beginning of the trial. The probability of each scenario must be updated as the time goes by, before each model run, in such a way that the probability of each scenario will indeed depend on the stage as well. Let us introduce $H_t^{(\mathrm{s})}$ as the remaining patient horizon at the beginning of period t', so the probability of each scenario will also depend on the stage where information has been updated:

$$p_{\mathrm{s,t'}} = \prod_{ij,t} p_{\mathrm{ij}}(d_{\mathrm{ij,t'}}^{(\mathrm{s})}, t')$$

where $\sum_{t=t'}^{L} \sum_{ij} d_{\mathrm{ij,t'}} = H_t^{(\mathrm{s})}$ and $H_1^{(\mathrm{s})} = H$.

4.4. Decision variables, parameters and sets

It is important to remark that, while demand is stochastic, remaining parameters are deterministic. Furthermore, demand behaves as deterministic within each scenario, and therefore, it will be treated as deterministic after the scenarios generation.

4.4.1. Sets and indexes

- $\mathcal{S}$: The set of scenarios, $s \in \mathcal{S}$
- $\mathcal{C}$: The set of countries, $i \in \mathcal{C}$
- $\mathcal{J}_i$: The set of sites served from country depot i, $j \in \mathcal{J}_i$
- $\mathcal{J}_0$: The set of sites served from CW, $j \in \mathcal{J}_0$

We will also refer to a period of time as t, where $0 \leq t \leq L$, and a production run as ρ, with $0 \leq \rho \leq \mu$, and where $\rho = 0$ is reserved for the inventory produced before the start of the enrolment. We will often refer to the whole set of enrolment sites, regardless of where they are served from: $\mathcal{U} = (\bigcup_{i \in C} \mathcal{J}_i) \cup \mathcal{J}_0$, and we will say that $j \in \mathcal{U}$.

4.4.2. Parameters

- v and f: variable and fixed production cost.
- v_i^{CD} and f_i^{CD}: variable and fixed delivery cost from CW to country depot i.
- v_j^h and f_j^h: variable and fixed delivery cost to enrolment site j.
- h^{CW}, h_i^{CD} and h_j^h: Unit holding cost per period at CW, country depot i and enrolment site j respectively.
- ζ: Level of service, understood as the probability that every enrolled patient will receive treatment (i.e. will not exceed the maximum waiting time and therefore dropout)
- ζ': Level of immediate service, understood as the probability that every patient will receive treatment without waiting, i.e. that every incoming patient will find enough inventory as for receiving treatment timely.
- μ: Number of production runs allowed.
- κ_c: Maximum number of shipments allowed from CW to CDs.
- κ_s: Maximum number of shipments allowed from CDs (or CW) to enrolment sites.
- LT_{prod}: The production lead time, measured in number of periods.
- LT_i^{CD}: Delivery lead time from CW to country depot i, measured in number of periods.
- LT_j^h: Delivery lead time to enrolment site j, measured in number of periods.
- DND: 'Do Not Dispense', if the remaining shelf life is less than this number of periods, the unit cannot be administered at the site.
- DNS: 'Do Not Ship', if the remaining shelf life is less than this number of periods, the unit cannot be delivered from a country depot to an enrolment site.
- DNT: 'Do Not Transfer', if the remaining shelf life is less than this number of periods, the unit cannot be delivered from the Central Warehouse to a country depot or site directly served from the Central Warehouse.

4.4.3. Decision variables

The only practical decision variables are:

- The initial inventory levels at each node of the network
- Production quantities and delivered quantities in each time period.

The model contains, however, other decision variables that are actually outcomes of these decision variables.

The only decisions to be made at time zero, before any random event takes place, are:

- IL_{CW}: Initial inventory level at the CW.
- IL_i^{CD}: Initial inventory level at country depot (CD) i.
- IL_{j}^{h}: Initial inventory level at enrolment site j.

Remaining decisions will be made with updated information at each period t.

- $IL_{t,\rho}^{\text{CW(s)}}$: Inventory level at CW at the end of period t, of units from production run ρ, under scenario s.
- $IL_{i,t,\rho}^{\text{CD(s)}}$: Inventory level at country depot i at the end of period t, of units from production run ρ, under scenario s.
- $IL_{j,t,\rho}^{\text{h(s)}}$: Inventory level at enrolment site j at the end of period t, of units from production run ρ, under scenario s.
- $\phi_t^{(\text{s})}$: 1 if, under scenario s, at the beginning of period t, the enrolment target has not been reached, 0 otherwise.
- $\psi_t^{(\text{s})}$: 1 if, under scenario s, at the beginning of period t, the clinical trial is not complete, 0 otherwise.
- $Q_{prod,t,\rho}^{(\text{s})}$: Quantity produced at period t, within production run ρ, under scenario s.
- $Q_{i,t,\rho}^{\text{CD(s)}}$: Quantity sent at time t, from CW to country depot i, from production run ρ, under scenario s. This means that, after LT_i^{CD} periods, this quantity will be added to the inventory level $IL_{i,t,\rho}^{\text{CD(s)}}$.

- $Q^{\text{h(s)}}_{j,t,\rho}$: Quantity sent at time t to enrolment site j, from production run ρ, under scenario s. This means that, after LT^h_j periods, this quantity will be added to the inventory level $IL^{\text{h(s)}}_{j,t,\rho}$.

- $y^{\text{(s)}}_{prod,t,\rho}, y^{\text{CD(s)}}_{i,t,\rho}, y^{\text{h(s)}}_{j,t,\rho}$: 1 if $Q^{\text{(s)}}_{prod,t,\rho} > 0$, $Q^{\text{CD(s)}}_{i,t,\rho} > 0$ and $Q^{\text{h(s)}}_{j,t,\rho} > 0$ respectively, 0 otherwise.

- y^{CW}_0: 1 if there is positive initial inventory (at time zero) in any node of the network, 0 otherwise.

- $y^{CD}_{i,0}$: 1 if there is not zero initial inventory, at time zero, in country depot i or any enrolment site served from it, 0 otherwise.

- $y^h_{j,0}$: 1 if the initial inventory, at time zero, level at enrolment site j is not zero, 0 otherwise.

- $E^{\text{(s)}}_{j,t}$: Number of patients enrolled at site j, during period t, under scenario s.

- $D^{\text{(s)}}_{j,t,k,r}$: Number of patients who will come at site j on period t for r$^{\text{th}}$ time in order to receive their k$^{\text{th}}$ dose, under scenario s.

- $X^{\text{(s)}}_{j,t,k,\rho,r}$: Number of patients who will receive their k$^{\text{th}}$ dose at site j on period t during their r$^{\text{th}}$ visit after having received the (k-1)$^{\text{th}}$ dose, with units from production run ρ , under scenario s.

- $A^{\text{(s)}}_{j,t,k}$: Number of patients who will drop out on period t after having received $k-1$ doses at site j and having visited the site for the δ^{th} time for a k$^{\text{th}}$ dose, under scenario s.

- $\alpha^{\text{(s)}}$: 1 if, under scenario s, at least one dropout happens during the whole trial, 0 otherwise.

- $\beta^{\text{(s)}}$: 1 if, under scenario s, at least one patient is not immediately filled for every dose in the treatment, during the whole trial, 0 otherwise.

4.5. A multi-objective mathematical model

In this section, we present a multi-stage, multi-objective, stochastic programming formulation.

4.5.1. Objectives

While we could rank our objectives in 24 different ways, the nature of these 4 objectives and because of the considerations discussed in Section 4.1, we follow a hierarchy of objectives as follows:

i. Minimization of expected necessary time to complete the trial

ii. Minimization of total cost

iii. Minimization of overage inventory

iv. Minimization of number of production runs

Then:

$$\min T = \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} p_s \cdot \psi_t^{(\mathrm{s})} \tag{4.4}$$

$$\begin{aligned}
\min C = v \cdot [IL_{CW} &+ \sum_{i\in\mathcal{C}} IL_i^{CD} + \sum_{j\in\mathcal{J}_i} IL_j^h + \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot Q_{prod,t,\rho}^{(\mathrm{s})}] \\
&+ \sum_{i\in\mathcal{C}} \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} (p_s \cdot v_i^{CD} \cdot Q_{i,t,\rho}^{\mathrm{CD(s)}}) + \sum_{j\in\mathcal{U}} \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} (p_s \cdot v_j^h \cdot Q_{j,t,\rho}^{\mathrm{h(s)}}) + \\
&+ f \cdot (y_0^{CW} + \sum_{s\in S} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot y_{prod,t,\rho}^{(\mathrm{s})}) + \sum_{i\in\mathcal{C}} [f_i \cdot (y_{i,0}^{CD} + \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot y_{i,t,\rho}^{\mathrm{CD(s)}})] \\
&+ \sum_{j\in\mathcal{U}} [f_j^h \cdot (y_{j,0}^h + \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot y_{j,t,\rho}^{\mathrm{h(s)}})] + \\
&+ \sum_{s\in\mathcal{S}} \sum_{t=0}^{L} \sum_{\rho=0}^{\mu} [p_s \cdot (h^{CW} \cdot IL_{t,\rho}^{\mathrm{CW(s)}} + \sum_{j\in\mathcal{U}} h_j^h \cdot IL_{j,t,\rho}^{\mathrm{h(s)}} + \sum_{i\in\mathcal{C}} h_i^{CD} \cdot IL_{i,t,\rho}^{\mathrm{CD(s)}})]
\end{aligned} \tag{4.5}$$

$$\min \Delta = IL_{CW} + \sum_{i\in\mathcal{C}} IL_i^{CD} + \sum_{j\in\mathcal{U}} IL_j^h + \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} (p_s \cdot Q_{prod,t,\rho}^{(\mathrm{s})}) - \sum_{s\in\mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} \sum_{k=1}^{m} \sum_{r=1}^{\delta} \sum_{j\in\mathcal{U}} p_s \cdot X_{j,t,k,\rho,r}^{(\mathrm{s})} \tag{4.6}$$

$$\min P = \sum_{s\in\mathcal{S}} p_s \cdot \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} y_{prod,t,\rho}^{(\mathrm{s})} + y_0 \tag{4.7}$$

4.5.2. Constraints

$$\sum_{z=1}^{t}\sum_{j\in\mathcal{U}}(E_{j,z}^{(s)}-\sum_{k=1}^{m}A_{j,z,k}^{(s)})\geq H\cdot(1-\phi_{t+1}^{(s)})\quad\forall s\in\mathcal{S},1\leq t\leq L \tag{4.8}$$

$$\sum_{z=1}^{t}\sum_{j\in\mathcal{U}}(E_{j,z}^{(s)}-\sum_{k=1}^{m}A_{j,z,k}^{(s)})\leq H-1+M\cdot(1-\phi_{t+1}^{(s)})\quad\forall s\in\mathcal{S},1\leq t\leq L \tag{4.9}$$

$$E_{j,t}^{(s)}\leq d_{j,t}^{(s)}\cdot\phi_{t}^{(s)}\quad\forall j\in\mathcal{U},s\in\mathcal{S},1\leq t\leq L \tag{4.10}$$

$$E_{j,t}^{(s)}\geq d_{j,t}^{(s)}-M\cdot(1-\phi_{t}^{(s)})\quad\forall j\in\mathcal{U},s\in\mathcal{S},1\leq t\leq L \tag{4.11}$$

$$\sum_{j\in\mathcal{U}}\sum_{k=1}^{m}\sum_{t'=t}^{L}\sum_{r=1}^{\delta}D_{j,t',k,r}^{(s)}\geq 1-M\cdot(1-\psi_{t+1}^{(s)})\quad\forall s\in\mathcal{S},1\leq t\leq L \tag{4.12}$$

$$\sum_{j\in\mathcal{U}}\sum_{k=1}^{m}\sum_{t'=t}^{L}\sum_{r=1}^{\delta}D_{j,t',k,r}^{(s)}\leq M\cdot\psi_{t+1}^{(s)}\quad\forall s\in\mathcal{S},1\leq t\leq L \tag{4.13}$$

$$D_{j,t,k,1}^{(s)}=\sum_{r=1}^{\delta}\sum_{\rho=0}^{\mu}X_{j,t-\tau,k-1,\rho,r}^{(s)}\quad\forall j\in\mathcal{U},\tau<t\leq L,1<k\leq m,s\in\mathcal{S} \tag{4.14}$$

$$D_{j,t,k,r}^{(s)}=D_{j,t-1,k,r-1}^{(s)}-\sum_{\rho=0}^{\mu}X_{j,t-1,k,\rho,r-1}^{(s)}\quad\forall j\in\mathcal{U},1<t\leq L,1\leq k\leq m,1<r\leq\delta,s\in\mathcal{S} \tag{4.15}$$

$$E_{j,t}^{(s)}=D_{j,t,1,1}^{(s)}\quad\forall j\in\mathcal{U},1\leq t\leq L,s\in\mathcal{S} \tag{4.16}$$

$$D_{j,t,k,\delta}^{(s)}=\sum_{\rho=0}^{\mu}X_{j,t,k,\rho,\delta}^{(s)}+A_{j,t,k}^{(s)}\quad\forall j\in\mathcal{U},\delta\leq t\leq L,1\leq k\leq m,s\in\mathcal{S} \tag{4.17}$$

$$\sum_{\rho=0}^{\mu}X_{j,t,k,\rho,r}^{(s)}\leq D_{j,t,k,r}^{(s)}\quad\forall j\in\mathcal{U},1\leq t\leq L,1\leq k\leq m,1\leq r<\delta,s\in\mathcal{S} \tag{4.18}$$

$$Q_{prod,t,\rho}^{(s)}\leq M\cdot y_{prod,t,\rho}^{(s)}\quad\forall 1\leq t\leq L,1\leq\rho\leq\mu,s\in\mathcal{S} \tag{4.19}$$

$$Q^{\text{CD(s)}}_{i,t,\rho} \leq M \cdot y^{\text{CD(s)}}_{i,t,\rho} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, 0 \leq \rho \leq \mu, s \in \mathcal{S} \tag{4.20}$$

$$Q^{\text{h(s)}}_{j,t,\rho} \leq M \cdot y^{\text{h(s)}}_{j,t,\rho} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, 0 \leq \rho \leq \mu, s \in \mathcal{S} \tag{4.21}$$

$$IL_{CW} + \sum_{i \in \mathcal{C}} IL^{CD}_i + \sum_{j \in \mathcal{U}} IL^h_j \leq M \cdot y_0 \tag{4.22}$$

$$IL^{CD}_i + \sum_{j \in \mathcal{J}_i} IL^h_j \leq M \cdot y^{CD}_{i,0} \quad \forall i \in \mathcal{C} \tag{4.23}$$

$$IL^h_j \leq M \cdot y^h_{j,0} \quad \forall j \in \mathcal{U} \tag{4.24}$$

$$IL^{\text{h(s)}}_{j,t,\rho} = IL^{\text{h(s)}}_{j,t-1,\rho} + Q^{\text{h(s)}}_{j,t-LT_j,\rho} \cdot \max(t - LT^h_j, 0) - \sum_{k=1}^{m} \sum_{r=1}^{\delta} X^{(s)}_{j,t,k,\rho,r} \quad \forall j \in \mathcal{U}, LT_j < t \leq L, 0 \leq \rho \leq \mu, s \in \mathcal{S} \tag{4.25}$$

$$IL^{\text{CD(s)}}_{i,t,\rho} = IL^{\text{CD(s)}}_{i,t-1,\rho} + Q^{\text{CD(s)}}_{i,t-LT_i,\rho} \cdot \max(t - LT^{CD}_i, 0) - \sum_{j \in \mathcal{J}_i} Q^{\text{h(s)}}_{j,t,\rho} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, 0 \leq \rho \leq \mu, s \in \mathcal{S} \tag{4.26}$$

$$IL^{(s)}_{CW,t,\rho} = IL^{(s)}_{CW,t-1,\rho} + Q^{(s)}_{t-LT_{prod},\rho} \cdot \max(t - LT_{prod}, 0) - \sum_{j \in \mathcal{J}_0} Q^{\text{h(s)}}_{j,t,\rho} - \sum_{i \in \mathcal{C}} Q^{\text{CD(s)}}_{i,t,\rho} \quad \forall 1 \leq t \leq L, 0 \leq \rho \leq \mu, s \in \mathcal{S} \tag{4.27}$$

$$IL^{(s)}_{CW,0,0} = IL_{CW} \quad \forall s \in \mathcal{S} \tag{4.28}$$

$$IL^{\text{CD(s)}}_{i,0,0} = IL^{CD}_i \quad \forall i \in \mathcal{C}, s \in \mathcal{S} \tag{4.29}$$

$$IL^{\text{h(s)}}_{j,0,0} = IL^h_j \quad \forall j \in \mathcal{U}, s \in \mathcal{S} \tag{4.30}$$

$$\sum_{i\in\mathcal{C}}\sum_{t=1}^{L}\sum_{\rho=0}^{\mu} y_{y,t,\rho}^{CD} \leq \kappa_c \quad \forall s \in \mathcal{S} \tag{4.31}$$

$$\sum_{j\in\mathcal{U}}\sum_{t=1}^{L}\sum_{\rho=0}^{\mu} y_{j,t,\rho}^{h(s)} \leq \kappa_s \quad \forall s \in \mathcal{S} \tag{4.32}$$

$$\sum_{t=1}^{L}\sum_{k=1}^{m}\sum_{j\in\mathcal{U}} A_{j,t,k}^{(s)} \leq M \cdot \alpha^{(s)} \quad \forall s \in \mathcal{S} \tag{4.33}$$

$$\alpha^{(s)} \leq \sum_{t=1}^{L}\sum_{k=1}^{m}\sum_{j\in\mathcal{U}} A_{j,t,k}^{(s)} \quad \forall s \in \mathcal{S} \tag{4.34}$$

$$\sum_{s\in\mathcal{S}}(p_s \cdot \alpha^{(s)}) \leq 1-\zeta \tag{4.35}$$

$$\sum_{t=1}^{L}\sum_{k=1}^{m}\sum_{j\in\mathcal{U}} D_{j,t,k,2}^{(s)} \leq M \cdot \beta^{(s)} \quad \forall s \in \mathcal{S} \tag{4.36}$$

$$\beta^{(s)} \leq \sum_{t=1}^{L}\sum_{k=1}^{m}\sum_{j\in\mathcal{U}} D_{j,t,k,2}^{(s)} \quad \forall s \in \mathcal{S} \tag{4.37}$$

$$\sum_{s\in\mathcal{S}}(p_s \cdot \beta^{(s)}) \leq 1-\zeta' \tag{4.38}$$

$$\sum_{t=1}^{L} y_{t,\rho}^{(s)} \leq 1 \quad \forall s \in \mathcal{S}, 1 \leq \rho \leq \mu \tag{4.39}$$

$$\sum_{k=1}^{m}\sum_{r=1}^{\delta}\sum_{j\in\mathcal{U}} X_{j,t,k,\rho,r}^{(s)} \leq M \cdot (1-y_{t',\rho}) \quad \forall s \in \mathcal{S}, 1 \leq \rho \leq \mu, 1 \leq t' < t+DND-SL < t \leq L \tag{4.40}$$

$$X_{j,t,k,0,r}^{(s)} = 0 \quad \forall s \in \mathcal{S}, 1 \leq SL-DND < t \leq L, 1 \leq k \leq m, 1 \leq r \leq \delta, j \in \mathcal{U} \tag{4.41}$$

$$\sum_{i\in\mathcal{C}}\sum_{j\in\mathcal{J}_i} Q^{(s)}_{j,t,\rho} \le M\cdot(1-y_{t',\rho}) \quad \forall s\in\mathcal{S}, 1\le\rho\le\mu, 1\le t'<t+DNS-SL<t\le L \tag{4.42}$$

$$Q^{(s)}_{j,t,0}=0 \quad \forall s\in\mathcal{S}, 1\le SL-DNS<t\le L, j\in\mathcal{U} \tag{4.43}$$

$$\sum_{i\in\mathcal{C}} Q^{(s)}_{i,t,\rho} \le M\cdot(1-y_{t',\rho}) \quad \forall s\in\mathcal{S}, 1\le\rho\le\mu, 1\le t'<t+DNT-SL<t\le L \tag{4.44}$$

$$Q^{(s)}_{i,t,0}=0 \quad \forall s\in\mathcal{S}, 1\le SL-DNT<t\le L, i\in\mathcal{C} \tag{4.45}$$

$$\sum_{j\in\mathcal{J}_0} Q^{h(s)}_{j,t,\rho} \le M\cdot(1-y_{t',\rho}) \quad \forall s\in\mathcal{S}, 1\le\rho\le\mu, 1\le t'<t+DNT-SL<t\le L \tag{4.46}$$

$$Q^{h(s)}_{j,t,0}=0 \quad \forall s\in\mathcal{S}, 1\le SL-DNT<t\le L, j\in\mathcal{J}_0 \tag{4.47}$$

$$\sum_{t'=1}^{t} y^{(s)}_{t',\rho} \ge \sum_{t'=1}^{t} y^{(s)}_{t',\rho+1} \quad \forall s\in\mathcal{S}, 1\le\rho<\mu, 1\le t\le L \tag{4.48}$$

$$IL_{CW}\ge 0 \tag{4.49}$$

$$IL^{CD}_{i}\ge 0 \quad \forall i\in\mathcal{C} \tag{4.50}$$

$$IL^{h}_{j}\ge 0 \quad \forall j\in\mathcal{U} \tag{4.51}$$

$$IL^{CW(s)}_{t,\rho}\ge 0 \quad \forall s\in\mathcal{S}, 1\le\rho<\mu, 1\le t\le L \tag{4.52}$$

$$IL_{i,t,\rho}^{CD(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L, i \in \mathcal{C} \tag{4.53}$$

$$IL_{j,t,\rho}^{h(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L, j \in \mathcal{U} \tag{4.54}$$

$$Q_{prod,t,\rho}^{(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L \tag{4.55}$$

$$Q_{i,t,\rho}^{CD(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L, i \in \mathcal{C} \tag{4.56}$$

$$Q_{j,t,\rho}^{h(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L, j \in \mathcal{U} \tag{4.57}$$

$$E_{j,t}^{(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq t \leq L, j \in \mathcal{U} \tag{4.58}$$

$$D_{j,t,k,r}^{(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq t \leq L, j \in \mathcal{U}, 1 \leq k \leq m, 1 \leq r \leq \delta \tag{4.59}$$

$$X_{j,t,k,\rho,r}^{(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq t \leq L, j \in \mathcal{U}, 1 \leq k \leq m, 1 \leq r \leq \delta, 1 \leq \rho \leq \mu \tag{4.60}$$

$$A_{j,t,k}^{(s)} \geq 0 \quad \forall s \in \mathcal{S}, 1 \leq t \leq L, j \in \mathcal{U}, 1 \leq k \leq m \tag{4.61}$$

$$\phi_t^{(s)} \in \{0,1\} \quad \forall s \in \mathcal{S}, 1 \leq t \leq L \tag{4.62}$$

$$\psi_t^{(s)} \in \{0,1\} \quad \forall s \in \mathcal{S}, 1 \leq t \leq L \tag{4.63}$$

$$y_{prod,t,\rho}^{(s)} \in \{0,1\} \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L \tag{4.64}$$

$$y_{i,t,\rho}^{CD(s)} \in \{0,1\} \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L, i \in \mathcal{C} \tag{4.65}$$

$$y_{j,t,\rho}^{h(s)} \in \{0,1\} \quad \forall s \in \mathcal{S}, 1 \leq \rho < \mu, 1 \leq t \leq L, j \in \mathcal{U} \tag{4.66}$$

$$y_0^{prod} \in \{0,1\} \tag{4.67}$$

$$y_{i,0}^{CD} \in \{0,1\} \quad \forall i \in \mathcal{C} \tag{4.68}$$

$$y_{j,0}^{h} \in \{0,1\} \quad \forall j \in \mathcal{U} \tag{4.69}$$

$$\alpha^{(s)} \in \{0,1\} \quad \forall s \in \mathcal{S} \tag{4.70}$$

$$\beta^{(s)} \in \{0,1\} \quad \forall s \in \mathcal{S} \tag{4.71}$$

Constraints represented in (4.8)-(4.11) indicate that enrolment will be paused when, by the starting of a period, the difference between enrolled patients and dropouts has reached the patient horizon. When enrolment is not paused, any new patient will be enrolled.

Constraints (4.12)-(4.13) indicate that treatment is finished when no patient is expected to return.

Eq.(4.14) indicates that people who received a certain dose in a certain period will return τ periods later for next dose.

Eq.(4.15) indicates that people who was not filled in a certain period, will return and demand that same dose on the next period.

Eq.(4.16) indicates that enrolled people will immediately demand their first dose

Eq.(4.17) indicates that people who are not filled for a specific dose after δ consecutive

periods will drop out from the clinical trial.

Eq.(4.18) indicates that the amount of people to receive a certain dose in a certain period and site cannot be greater than the nuber of people claiming such dose in such site during such period.

Eqs. (4.19)-(4.24) define linking constraints for binary variables that indicate drug production or deliveries.

Eqs. (4.25)-(4.27) represent inventory balance in enrolment sites, country depots and Central Warehouse respectively.

Eqs. (4.28)-(4.30) define initial conditions for inventory level variables.

Eqs. (4.31)-(4.32) express the potential limitation in number of shipments allowed to CDs and enrolment sites.

Eqs. (4.33)-(4.34) make sure that α takes value 1 if and only if some dropout happens, and 0 otherwise. Eq.(4.35) links α to the required level of service.

Eqs. (4.36)-(4.37) make sure that β takes value 0 if and only no waiting is required from any patient to be filled, and 1 otherwise. Eq.(4.38) links β to the required level of immediate service.

Eq.(4.39) indicates that each production run can happen once at most. Eq.(4.40)-(4.41) make explicit reference to the limited shelf life and define DND; in a similar way, Eq.(4.42)-(4.47) define DNS and DNT as the maximum remaining shelf lifes suitable for shipping from a country depot and from CW respectively.

Eq.(4.48) enforces that a production run number cannot take place unless previous ones have already happened.

Eq.(4.49)-(4.61) are non-negativity constranints, and (4.62)-(4.71) make explicit reference to the binary nature of some variables.

There is one more set of constraints that we will explicitly enforce, namely the non-anticipative constraints (NAC): if two scenarios share the exact same arrival pattern up to a certain period t, then all decisions made up to period t must be identical for both scenarios. This prevents from using future information:

$$\left\{d_{\mathrm{j,t}}^{(\mathrm{s})}\right\} = \left\{d_{\mathrm{j,t}}^{(\mathrm{s}')}\right\} \forall 1 \leq t \leq t' \implies \begin{cases} Q_{prod,t,\rho}^{(\mathrm{s})} = Q_{prod,t,\rho}^{(\mathrm{s}')} & \forall s \in \mathcal{S}, 1 \leq \rho \leq \mu, 1 \leq t \leq t' \\ Q_{i,t,\rho}^{\mathrm{CD(s)}} = Q_{i,t,\rho}^{\mathrm{CD(s')}} & \forall s \in \mathcal{S}, i \in \mathcal{C}, 0 \leq \rho \leq \mu, 1 \leq t \leq t' \\ Q_{j,t,\rho}^{\mathrm{h(s)}} = Q_{j,t,\rho}^{\mathrm{h(s')}} & \forall s \in \mathcal{S}, j \in \mathcal{U}, 0 \leq \rho \leq \mu, 1 \leq t \leq t' \end{cases} \tag{4.72}$$

To illustrate this, let's take the scenario tree represented in Figure 4.2. If we were at stage 1, we could not distinguish whether we are in scenario 1, 2, 3 or 4: therefore decisions made in that stage for those 4 scenarios should be identical. When we update information in stage 2, we know scenarios 1 and 2 are no longer possible, but we still cannot distinguish whether we are in scenario 3 or 4, and therefore decision made in stage 2 for scenarios 3 and 4 should be identical.

4.5.3. A NOTE ON PRODUCTION BATCHES

It is important to remark that in each node of the network, stock from different production runs may exist. The differentiation among them is not expected to be an issue, because we assume that they will be properly labeled. Furthermore, all the inventory from a certain production run is assumed to be produced in the same period, and with identical shelf life, therefore all the units from a same production run will have the exact same remaining shelf life period by period. This guarantees that inventory from a production run will either consist of usable units or all of them will be units not to be administered/shipped, making inventory management much simpler.

Also, it is worth to mention that while we are not enforcing a FIFO policy, the use of newer

units instead of older units would be a solution dominated by the use of older units when possible, therefore the optimization algorithm will enforce the right prioritization when two or more production batches coexist in any node.

Let us note that, while we have explicitly considered the expected amount of unused inventory at the end of the trial, we have not distinguished between units that have perished and units that are still usable by the end of the trial. The main reason for this is that, in practice, unused drugs will be disposed at the end of the trial regardless of whether they have perished or not, and therefore there is no point on considering such distinction in any objective or constraint. With this being said, it is still interesting to report the perished and non-perished inventory separately. Let us denote by $IL_{per}^{h(s)}$ and $IL_{good}^{h(s)}$ the number of perished and non-perished units respectively at the end of the trial under scenario s, and let us define $T^{(s)}$ as the clinical trial completion time under scenario s, in such a way that we can rewrite objective (4.4) like:

$$\min T = \sum_{s \in \mathcal{S}} p_s \cdot T^{(s)} \tag{4.73}$$

where:

$$T^{(s)} = \sum_{t=1}^{L} \psi_t^{(s)} \tag{4.74}$$

Then for the amount of perished units by the end of the trial under scenario s, we have:

$$IL_{per}^{h(s)} = \sum_{t=(SL-DND)}^{T^{(s)}} \sum_{\rho=0}^{\mu} IL_{j,t,\rho}^{h(s)} \cdot y_{prod,t+DND-SL,\rho}^{(s)} \tag{4.75}$$

Let us note that, although (4.75) is non-linear, this does not affect the model, since this is not an issue because this non-linearity takes no part in the model: we will only use (4.75) to calculate the amount of perished drugs under each scenario a posteriori, i.e. once the model has already been solved.

4.5.4. A note on Completion Time minimization as main objective

On section 4.5.1, we have explained the reasons for choosing a particular objectives hierarchy among the 24 possible sequences of 4 objectives. It is worth to highlight that the particular objective hierarchy that we have chosen (actually any of the 6 hierarchies where completion time is the principal objective) offers a considerable simplification opportunity for finding the minimum completion time without necessity of solving the model described in Section 4.5. For simplicity (the necessary conditions are much less restrictive), let us assume that there is no limitation in the allowed number of production runs and shipments, and that the shelf life is long enough as for neglecting perishability: we could simply place a sufficiently large initial inventory in each enrolment site ($IL_j^h \geq m \times H \quad \forall j \in \mathcal{U}$), and this would mean that no patient would have to wait for inventory. Then, the enrolment (1st dose) of the whole patient horizon for each scenario would, simply be the the time needed for the cumulative arrivals to reach the target number of patients. Completion time for each scenario is then:

$$T^{(s)} = (m-1) \cdot \tau + \min(t) | \sum_{t'=1}^{t} \sum_{j \in \mathcal{U}} d_{j,t'}^{(s)} \geq H \tag{4.76}$$

Eq. 4.76 indicates that, for each scenario, the completion time is the addition of two terms: 1) the enrolment time for the first H patients, and 2) the time to wait for the last enrolled patient to receive the last $m-1$ doses. This equation, combined with (Eq. 4.73) lead to the determination of the expected completion time, which is much simpler than solving the whole Mixed Integer Stochastic Linear Program.

Sufficient conditions for optimization procedure described in section 4.5.4

In order to build an intuition for the solution procedure under mild conditions when the expected completion time is the main objective, we have assumed conditions that might be too restrictive and that, fortunately, are not necessary.

The described solution will remain valid in any situation where the delay in the administration of doses to patients can be avoided by:

1. Allocation of enough initial inventory in every enrolment site, and
2. Production of a sufficiently large amount of units, delivered timely to enrolment sites

granting that they arrive before the expiration of the existing units in the enrolment sites.

It would be sufficient that the following condition is met:

$$\{SL - DND - \{LT_{prod} + \max[\max_{i \in \mathcal{C}}(LT_i^{CD} + \max_{j \in \mathcal{J}_i} LT_j^h), \max_{j \in \mathcal{J}_0} LT_j^h]\}\} \cdot \min(\mu, \kappa_c, \kappa_s) \geq L \quad (4.77)$$

4.6. Conditional sampling and sequential optimization

The output of the model described in section 4.5 is the set of all $Q^{(s)}_{prod,t,\rho}$, $Q^{CD(s)}_{i,t,\rho}$ and $Q^{h(s)}_{j,t,\rho}$, i.e. the prescription of:

1. Initial inventory at each node of the network
2. How much should be produced at the CW at each period, under each scenario
3. How much inventory should be shipped at each period from CW and from each country depot to each downstream node, under each scenario.

Let us note that, as the CT evolves over time, we will have updated information on where patients have been enrolling for treatment, which is critical to make the right distribution decisions. When we initially solve the model, it will give us a good knowledge of how much initial inventory we should locate at CW, at each country depot and at each enrolment site, but the recommendation on how much we should produce exactly 40 days later and how much of that production should be pushed to each downstream node is much less accurate, and is actually parameterized by scenario s. For this reason, sequential optimization is required, which follows the next steps:

1. Generate S scenarios.
2. Solve model and extract values of IL_{CW}, IL_i^{CD} for every country i and IL_{j}^h for every enrolment site j.
3. Create additional constraints enforcing IL_{CW}, IL_i^{CD} and IL_{j}^h to take the corresponding optimal values extracted in previous step.
4. Set $t' = 1$

5. Conditional sampling: Generate S new scenarios, this time enforcing the new patients arrivals from $t = 1$ to $t = t'$ to match the actual realizations.

6. Solve updated model and extract all values of $Q^{(s)}_{prod,t',\rho}$, $Q^{CD(s)}_{i,t',\rho}$ and $Q^{h(s)}_{j,t',\rho}$.

7. Create additional constraints enforcing $Q^{(s)}_{prod,t',\rho}$, $Q^{CD(s)}_{i,t',\rho}$ and $Q^{h(s)}_{j,t',\rho}$ to take the corresponding optimal values extracted in previous step.

8. Increase t' in one unit, i.e. $t' \leftarrow t' + 1$

9. While $t' < L$, return to step 5

We will find that, in general eq(4.1) will result in a prohibitively large number, so we will need to sample a certain number of scenarios that has been shown to be representative enough. It is important to remark that the accuracy of the sequential optimization will heavily depend on the conditional sampling, which basically depends on the sample size, which is likely to be limited by the computational power, considering that we need to go several times over the cycle *sample* $\rightarrow$ *optimize* $\rightarrow$ *add constraints* $\rightarrow$ *sample* $\rightarrow$...

The lower bound for our sample size should be determined *ad hoc* by a stability test. In occasions, this lower bound might exceed the computational power that we can use. In next chapter, we present a method to deal with the essential trade-off between stability/accuracy and tractability.

Chapter 5

Scenario Reduction

5.1. Scenario Reduction

As explained in the previous chapter, the pursuit of stable and accurate solutions requires conditional sampling with a sufficiently large sample size, whose lower bound is to be determined *ad hoc* by means of a stability test. This lower bound might exceed the computational power that we can use, hence emphasizing the importance of finding an adequate scenario reduction method that can deal with the problem taking care of both accuracy and tractability.

Given the scenario tree, we want to reduce its size while preserving as much information as possible.

As pointed by (Beraldi and Bruni 2014), two contrasting issues should be jointly taken into account:

1. The number of scenarios should be large enough as for guaranteeing an accurate representation of the underlying stochastic process.

2. The cardinality of the scenario set should be kept limited, to preserve the computational tractability.

Given the original scenario tree, the problem is the determination of a reduced tree that contains a 'good enough' subset of scenarios. In this paper, we address the scenario reduction problem by a clustering approach.

5.2. Scenario reduction by clustering

When applying clustering methods, we group scenarios that are close to each other in terms of optimal solution (Hewitt, Ortmann, and Rei 2022), obtaining a partition $C_1, ..., C_K$ of the space of scenarios $\mathcal{S}$. We choose one representative scenario $\sigma_j \in C_j$ for each cluster C_j, and assign it a probability π_j:

$$\pi_j = \sum_{s \in C_j} p_s \quad \forall j, 1 \leq j \leq K \tag{5.1}$$

It is worth to remark that the cardinality of each cluster $C_1, ..., C_K$ is in general different, and therefore, by (eq. 5.2), these clusters are assigned different probabilities. This is because not every scenario is equally representative of the possible evolution of the system.

As is known, the clustering procedure requires the definition of both:

- **A clustering method:** the convenience of k-medoids++ (Schubert and Lenssen 2022) is very clear in this case, since it guarantees that the selected representative scenarios (medoids) is an existing scenario, i.e. that the representative element of a cluster is an element of the original set, unlike it happens with other popular clustering methods like *k-means* or *k-medians*. Another advantage of *k-medoids++* is its robustness with respect to random initialization.

- **A distance definition:** The absence of an adequate metric to express dissimilarity between scenarios in a multi-stage setting for a multi-echelon network is one of the main knowledge gaps that we have detected during our extensive literature review, where most authors use Euclidean distance or Manhattan distance due to their popularity, with little or none reflection on its convenience.

5.2.1. K-MEDOIDS++ CLUSTERING

Clustering is a common unsupervised machine learning task, in which the data set has to be automatically partitioned into 'clusters', such that objects within the same cluster are more similar, while objects in different clusters are more different.

A classic method taught in textbooks is *k-means* (Bock 2007), where the data is modeled using k cluster means, that are iteratively refined by assigning all objects to the nearest mean, then recomputing the mean of each cluster. In *k-medoids*, the data is modeled similarly, using k representative objects called medoids, which are chosen from the dataset and that serve as prototypes for the clusters instead of means. While k-means tries to minimize the within cluster sum-of-squares, k-medoids tries to minimize the sum of distances between each point and the medoid of its cluster.

The fact that the representative element of each cluster is itself an element of the original set is a rare and highly desirable property at the time of partitioning scenarios.

The k-medoids algorithm has been shown to be NP-hard (Ushakov and Vasilyev 2019), but there exist very precise heuristics that run in polynomial time: the Partitions Around Medoids (PAM) algorithm used here has a complexity $O(N^2KT)$, where N is the sample size, K the number of clusters and T the number of iterations.

The use of a data point to represent each cluster's center allows the use of any distance metric for clustering, and actually in 5.2.2 we will present a specific metric for the problem that is object of this research.

PAM: Concept and definition

As explained by (Kaufman and Peter J Rousseeuw 2009), when partitioning a set of objects into k clusters the main objective is to find clusters, the objects of which show a high degree of similarity, while objects belonging to different clusters are as dissimilar as possible. The algorithm used in the program PAM is based on the search for k representative objects among the objects of the data set. These objects should represent various aspects of the data. In the cluster analysis literature such representative objects are often called *centroids*, but in the context of PAM the representative objects are the so-called *medoids* (Kaufman and Peter J. Rousseeuw 1987). After finding a set of k representative objects, the k clusters are constructed by assigning each object of the data set to the nearest representative object.

The algorithm consists of two.phases. In a first phase, called BUILD, an initial clustering is obtained by the successive selection of representative objects until k objects have been found. The first object is the one for which the sum of the dissimilarities to all other objects is as small as possible. This object is the most centrally located in the set of objects. Subsequently, at each step another object is selected. This object is the one which decreases the objective function as much as possible.

In the second phase of the algorithm (called SWAP), it is attempted to improve the set of representative objects and therefore also to improve the clustering yielded by this set. This is done by considering all pairs of objects (i, h) for which object i has been selected and object h has not. It is determined what effect is obtained on the value of the clustering when a swap is carried out (when object i is no longer selected as a representative object but object h is).

This algorithm ensures that a good partition is found, although it may be not optimal.

One of the undesired characteristics that *k-medoids* shares with *k-means* is that the resulting partition depends on the initialization. At this point, it is important to remark that granting the absolute minimum total dissimilarity is not of our interest: what we really need is to have the set of scenarios partitioned according to an adequate metric that we will define in 5.2.2, and that provides for each cluster a representative element that belongs to the original set. However, it is highly desirable that the selection of representative elements does not hide the existence of extreme scenarios, which we want to be represented in order for the model to make robust decisions taking into consideration the variety of uncertain scenarios that could occur.

For this reason, we have used the variant called *k-medoids*++ (Arthur and Vassilvitskii 2007), which uses a density based initialization method that favours the initial selection of elements that are distant from each other.

It is expected that, despite the reduced number of scenarios, better results are achieved due to the combination of two factors:

1. Robustness of *k-medoids*++ allows to properly weigh the relative likelihoods of different

scenarios, enabling for wiser decision making.

2. The density-based initialization of *k-medoids++* ensures the consideration of more extreme scenarios, considering a wider set of possibilities during the decision making process.

5.2.2. A new metric for multi-stage scenario clustering in a multi-echelon setting

We will define a distance $D(s_1, s_2)$ between scenarios s_1 and s_2 as:

$$D(s_1, s_2) = \sum_{t=1}^{L} \left[c_1 \cdot \sum_{j \in \mathcal{U}} \left| \sum_{t'=1}^{t} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right| + c_2 \cdot \left| \sum_{t'=1}^{t} \sum_{j \in \mathcal{U}} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right| + c_3 \cdot \sum_{i \in \mathcal{C}} \left| \sum_{t'=1}^{t} \sum_{j \in \mathcal{J}_i} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right| \right] \tag{5.2}$$

The first term in the right hand side allows for cumulative demand to compensate over periods of time, but it punishes any spatial difference between scenarios. This way, if two scenarios differ only in the permutation of the arrivals sequence in two different enrolment sites, that difference will be captured by this contribution.

Second term allows for cumulative demand to compensate not only over time, but also over the enrolment sites. If two scenarios differ only in the permutation of two nodes, this term will not contribute to the distance between these two scenarios. For example, if there is a scenario where 10 patients arrive on a given day to a hospital in Paris and 5 arrive to Brussels, and another scenario where 5 arrive to Paris and 10 to Brussels, that difference will not be captured by this term: this is something intended, because these two scenarios would have the same impact for the production decisions made at the CW, so we do not want them to be considered so different.

Third term is similar to second one, but it only allows for cumulative demand to compensate over nodes supplied by the same intermediate echelon node. This is because we want this term to not contribute when two scenarios differ in the permutation of demand between two sites served from the same country depot: this is because we expect this difference to have no impact not only in production decisions, but also in the shipments from CW to CD.

Coefficients c_1, c_2 and c_3 are to be determined. We want these coefficients to represent the importance of each of these three terms in the determination of the distance between two different scenarios. Next section deals with the determination of these coefficients.

Determination of c_1, c_2 and c_3

Let us denote the set of all decision variables by θ, and the objective function by $f(\theta, s)$. In general, we want to minimize the expected value $E_s[f(\theta, s)]$, and we will denote the set of optimal values for the decision variables as θ^*.

We want similar scenarios lead to similar results, so it is natural to pick $D(s_1, s_2)$ coefficients in such a way that the absolute difference between the objective function evaluated at two different scenarios is aligned with the distance between those scenarios:

$$|f(\theta^*, s_1) - f(\theta^*, s_2)| \sim D(s_1, s_2) \tag{5.3}$$

We will then proceed as follows:

1. Generate S' random scenarios
2. Compute all the pairwise distances as a function of c_1, c_2 and c_3 using (5.2)
3. Solve the program and evaluate the objective at each deterministic scenario.
4. Determine c_1, c_2 and c_3 by linear regression using (5.3)

Once c_1, c_2 and c_3 are determined, we have a valid metric to measure dissimilarity between scenarios, which we can safely use to perform K-medoids++ clustering in order to achieve the desired scenario reduction as for solving our problem with less computational effort.

An illustrative example

For the sake of clarity, let us hypothesize an extremely simple clinical trial with 4 enrolment sites only, each of them served from a country depot, where patents arrive following a Poisson Distribution.

Site	Country	λ(patients/day)	Lead Time (days)
Site 1	Country 1	1	1
Site 2	Country 1	4	1
Site 3	Country 2	7	1
Site 4	Country 2	10	1

Table 5.1: *Enrolment sites data for illustrative example in 5.2.2*

We will assume unit production cost of €100 and a patient horizon $H = 50$ for a treatment where a total of $m = 3$ doses are administered to each patient every $\tau = 14$ days. Tables 5.1 and 5.2 show data for enrolment sites, and country depots respectively.

CD	Fixed Delivery Cost	Variable Delivery Cost	Lead Time (days)
Country 1	€50	€10	1
Country 2	€50	€10	1

Table 5.2: *Country depots data for illustrative example in 5.2.2*

We have generated just 20 scenarios, for which we show a slice in Table 5.3. First 4 scenarios have been generated independently, using the arrival rates indicated in Table 5.1, while the remaining ones have been generated by permutation of rows and columns over the independent scenarios: this is intended to generate scenarios that would allow for compensation in some of the terms in the dissimilarity metric defined in Eq. 5.2.

Scenario	Site	t=1	t=2	t=3	t=4	t=5
1	1	1	1	0	0	0
1	2	0	3	5	0	0
1	3	3	8	7	0	0
1	4	8	5	13	0	0
2	1	2	1	3	0	0
2	2	2	3	6	0	0
2	3	5	9	8	0	0
2	4	6	8	11	0	0
...	...	...	...	...		
20	1	0	1	2	0	2
20	2	0	6	5	0	3
20	3	0	5	13	0	10
20	4	0	7	18	0	10

Table 5.3: *Scenarios generated for illustrative example in 5.2.2*

Optimizing for cost, the solution indicates a minimum expected cost of €9,511, but more importantly we can evaluate, at optimality, the cost in each specific scenario (Table 5.4).

Next, we take every pair of scenarios (s_1, s_2),(s_1, s_3),...,(s_{19}, s_{20}) and we compute for them

Scenario	Optimal Cost (€)
1	9,540
2	9,740
3	9,570
4	9,820
5	9,460
6	9,940
7	8,900
8	8,900
9	9,210
10	9,210
11	9,940
12	9,640
13	9,540
14	9,420
15	9,430
16	9,430
17	9,570
18	9,570
19	9,820
20	9,570

Table 5.4: *Cost for each scenario in the illustrative example in 5.2.2*

the left hand side (LHS) and right hand side (RHS) of Eq. 5.3.

LHS corresponding to scenarios s_1 and s_2 is simply $|9,540 - 9,740| = 200$. The corresponding RHS will depend on c_1, c_2 and c_3 as expressed in Eq. 5.2, resulting $50 \cdot c_1 + 40 \cdot c_2 + 40 \cdot c_3$. If we use Ordinary Least Squares (OLS), the repetition of this procedure for the $20 \times 19/2 = 190$ scenario pairs leads to a linear regression (with zero intercept) with $AdjR^2 = 0.956$, in which $c_1 = -2.3458$; $c_2 = 0$, and $c_3 = 5.2047$, with all p-values below 10^{-4} and $R^2 = 0.9496$. However, the fact that $c_1 < 0$ is not desirable, since we want Eq.5.2 to satisfy the properties to be considered a distance in a metric space (M. Deza and E. Deza 2009), i.e:

i. **Non-negativity:** $D(s_1, s_2) \geq 0 \quad \forall s_1, s_2$

ii. **Identity of indiscernibles:** $D(s_1, s_2) = 0 \Leftrightarrow s_1 = s_2 \quad \forall s_1, s_2$

iii. **Symmetry:** $D(s_1, s_2) = D(s_2, s_1) \quad \forall s_1, s_2$

iv. **Triangle inequality:** $D(s_1, s_2) \leq D(s_1, s_3) + D(s_1, s_3) \quad \forall s_1, s_2, s_3$

While these conditions are all met when coefficients c_1, c_2 and c_3 are all non-negative, this is not necessarily the case when any of them is negative. Actually, in this particular example,

we have 3 pairs where the estimated distance using Equation (5.2) with $c_1 = -2.3458$; $c_2 = 0$, and $c_3 = 5.2047$ would be negative. For this reason, instead of using OLS, we must apply Non-Negative Least Squares (NNLS) , which in this case leads to $c_1 = c_2 = 0$ and $c_3 = 2.738$, with a deviance just 10.36% greater than the one resulting from OLS, and a coefficient of determination $R^2 = 0.85675$. In this example, as we will do in Chapter 6, we have used Lawson and Hanson NNLS algorithm (Lawson and Hanson 1995).

For this specific clinical trial, we would then perform scenario reduction using *k-medoids++* with a dissimilarity metric:
$D(s_1, s_2) = 2.738 \cdot \sum_{t=1}^{L} \sum_{i \in \mathcal{C}} \left| \sum_{t'=1}^{t} \sum_{j \in \mathcal{J}_i} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right|$.

Note on the applicability of OLS for determination of c_1, c_2 and c_3

As shown in the previous section, the application of OLS may result in some of the coefficients being negative, with the possibility of violating some basic properties of a distance. It is important to remark that, while we can easily avoid this risk by applying NNLS regression instead of OLS, this is not always necessary. Let us note that, from the basic triangle inequality property of absolute value, it follows that:

$$\sum_{j \in \mathcal{U}} \left| \sum_{t'=1}^{t} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right| \geq \sum_{i \in \mathcal{C}} \left| \sum_{t'=1}^{t} \sum_{j \in \mathcal{J}_i} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right| \geq \left| \sum_{t'=1}^{t} \sum_{j \in \mathcal{U}} (d_{j,t'}^{(s_1)} - d_{j,t'}^{(s_2)}) \right| \tag{5.4}$$

Terms in Eq.5.2.2 are, from left to right, those multiplied by c_1, c_3 and c_2 respectively in Eq.5.2. This is why $c_1 < 0$ is a serious threat for the non-negativity property of distance. However, when $c_1 \geq 0$, $c_3 \geq 0$, $c_2 < 0$ and $c_1 + c_2 + c_3 \geq 0$, all distance properties are granted to be satisfied, and we can safely use OLS regression. Moreover, since this is a less restrictive regression, we would recommend OLS over NNLS in this particular case. We will, indeed, apply OLS later in Section 6.3.

Chapter 6

Case Studies

In order to test the model presented in Chapter 4, as well as the scenario reduction method presented in Chapter 5, we have chosen 3 case studies, informed by the trials completed by the sponsoring pharmaceutical company:

i. WC28325, a phase III clinical trial for *diabetes mellitus* treatment: we chose this one because we find interesting the relatively fast enrolment (it enrols an average 1.43 patients per day, compared to the mean of 0.12 for all the clinical trials in our database; it is situated in the upper 7% percentile in new patients arrival rate), which allowed for a fair comparison with (A. Fleischhacker, Ninh, and Y. Zhao 2015), as well as the pretty general structure, where not only 4 different depots serve 12 enrolment sites, but also there are 10 enrolment sites served directly from the CW.

ii. BO29389, A phase II clinical trial for malignant neoplasm of branchus and lung: we wanted to show this case because of the long enrolment time required and the relatively high number of doses per patient (15) in addition to a non-trivial structure where 18 sites are distributed among 6 different countries.

iii. BP3002, a phase I clinical trial for glaucoma treatment: It was interesting to include a phase I CT in order to cover all first 3 phases, and with a sufficiently large enrolment period as for shelf life being a critical factor in the production plan. Also, we wanted to include a single-dose treatment in order to make an even fairer comparison with (A. Fleischhacker, Ninh, and Y. Zhao 2015).

By choosing these 3 cases, we have covered an interesting range of characteristics:

- phases I, II and II are present
- enrolment period ranges from 132 days to 547 days

- number of doses ranges from 1 to 15
- enrolment sites range from 6 to 22
- very different disease areas are covered (diabetes, cancer and glaucoma)
- number of patients range from 45 to 190

Table 6.1 shows a summary of the main characteristics of each trial considered in this Chapter.

Case	# Countries	Sites	Total λ (pat/day)	# doses	τ	H	δ	Drug Cost
WC28325	5	22	1.44	3	7 days	190	3 days	€152
BO29389	6	18	0.257	15	3 weeks	49	2 weeks	€152
BP30002	2	6	0.239	1	-	45	3 days	€1500

Table 6.1: *summary of cases to analyze*

6.1. WC28325: A phase III CT for *diabetes mellitus* treatment

The proposed model was tested considering the clinical trial WC28325, in which 190 patients received treatment for *diabetes mellitus* every 7 days for a total of 3 doses per patient. This clinical trial was conducted during 147 days in year 2013 in 5 different countries, in a total 22 sites, with 4 country depots and 10 sites served directly from the Central Warehouse (located in USA), as is shown in Figure 6.1. The unit production cost of the drug is assumed to be €152. Although the model considers a finite production time as parameter of the model, for this specific case, the product is simple enough as for considering that the production made in a certain period is available to deliver in that same period. We will update information on a daily basis so, in the context of this case, when we speak about stages, we mean days.

Table 6.2 shows the lead times and arrival rates used in this optimization. Although our model does not have this limitation, we have decided to assume transportation costs to enrolment sites to be €0, and same for holding costs at every node of the network, mainly for 2 reasons:

1. We want to establish a fair comparison with the results derived from the model in A. Fleischhacker, Ninh, and Y. Zhao 2015, and

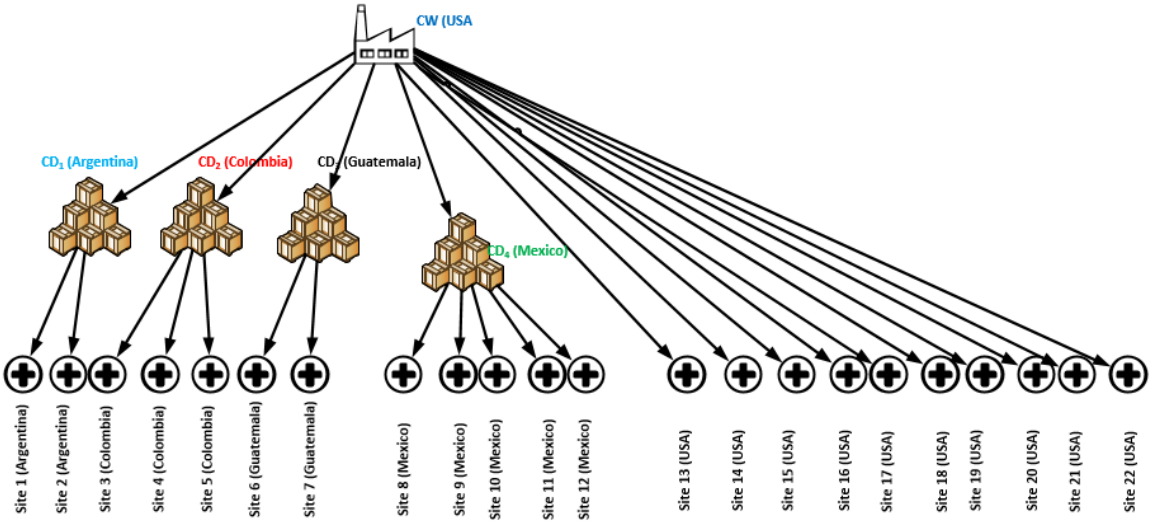

Figure 6.1: *Structure of the SC Network in CT WC28325*

2. The transportation costs to enrolment sites and holding costs at sites are negligible compared to the production costs and transportation costs from CW to the CDs.

Data relative to country depots is shown in Table 6.3. It is important to remark that, for confidentiality reasons, figures are not exact; with this being said, they represent reasonable estimates from our sponsoring company. Same applies for every cost estimate used in this section.

Also, we will initially assume no fixed production cost; later in this chapter we will include a fixed production cost in order to test how its consideration affects production and distribution decisions. Let us observe that such assumption will act in favour of having daily production runs of small size, maximizing the use of available information at the time of deciding the size of production batches, and taking advantage of the low commitment of producing small quantities. However, the existence of non-zero fixed shipment costs and considerable shipment lead times will encourage larger production quantities to be shipped from CW to CDs only a few times rather than every period. These competing forces configure a trade-off that justifies this elaborated optimization approach.

By Equation 4.1, the lower bound for the total number of scenarios in this problem is greater than 7×10^{313}. As is usual in Stochastic Programming, we need to sample a number of scenarios sufficiently large as for representing the multiple possible states of evolution of the system, while keeping the number of scenarios sufficiently small as for being able to solve the

j	i	Country	λ(patients/day)	Lead Time (days)
1	1	Argentina	0.114	1
2	1	Argentina	0.038	2
3	2	Colombia	0.121	1
4	2	Colombia	0.068	2
5	2	Colombia	0.008	3
6	3	Guatemala	0.098	1
7	3	Guatemala	0.053	2
8	4	Mexico	0.121	2
9	4	Mexico	0.114	2
10	4	Mexico	0.106	2
11	4	Mexico	0.091	2
12	4	Mexico	0.061	2
13	0	USA	0.106	1
14	0	USA	0.091	2
15	0	USA	0.076	2
16	0	USA	0.045	2
17	0	USA	0.038	2
18	0	USA	0.030	2
19	0	USA	0.023	2
20	0	USA	0.023	2
21	0	USA	0.008	2
22	0	USA	0.008	2

Table 6.2: *Enrolment sites data for CT WC28325*

CD	i	Fixed Delivery Cost	Variable Delivery Cost	Lead Time (days)
Argentina	1	€60	€180	3
Colombia	2	€40	€150	3
Guatemala	3	€25	€100	2
Mexico	4	€25	€100	2

Table 6.3: *Country depots data for CT CW28325*

problem with the available computational power within a reasonable amount of time.

The scenario generation assumes that new arrivals follow a Poisson distribution. It is important to highlight that, as the uncertainty is represented by a collection of discrete scenarios in the proposed model, any stochastic process could have been used to generate scenarios. However, the analysis on Section 3.3 has brought enough evidence in favour of considering Poisson arrivals. Furthermore, in this particular trial, we have failed to reject the null hypothesis of Poisson distribution in 16 out of 18 sites, where an 86.8 % of patients were enrolled.

As is usual in Stochastic Programming, the number of scenarios required to faithfully represent

the underlying uncertainty is a barrier to finding efficient numerical solutions. Monte Carlo sub-sampling simply consists in picking randomly a subset of S' scenarios in the original set and assigning them equal probability $1/S'$. The choice of S' is non-trivial decision to be made only after the appropriate stability test.

6.1.1. Stability Test

The procedure consists of the following steps:

1. Take M samples of size S', with some $M > 30$ and for different values of S'.
2. Solve each of the M stochastic programming problems for each of the values of S' under consideration.
3. Calculate the sample mean and sample standard deviation of the objective function(s), and plot them.
4. The recommended S' is one where both the sample mean and sample standard deviations seem to stabilize, i.e. where the expected gain in terms of standard deviation reduction is little.
5. Simultaneously, the computational time necessary to run the optimization for a sample of size S' is taken into consideration.

For this particular case study, we have taken 40 replications (i.e. $M = 40$) of each sample size between 10 and 60, in steps of 10, concluding that a sample size of 40 provides a good representation of the system, and a reasonable trade-off in terms of stability and computational effort.

In Figure 6.2, each bubble represents a stochastic optimization, and a larger size of the bubble indicates a smaller optimality gap, i.e. the difference between the LP relaxation bound and the best incumbent solution found (it is worth to mention that, for the sake of time, in this specific case, the program has been set to return a solution as soon as it gets an optimality gap below €25, which is the lowest associated cost in the trial, as shown in Table 6.3). The discontinuous line represents the sample mean. We see that expected completion time is pretty stable around 147 days, while the expected cost seems to stabilize right below €136,000, with the elbow method (Kleywegt, Alexander Shapiro, and Homem-de-Mello 2002) suggesting a

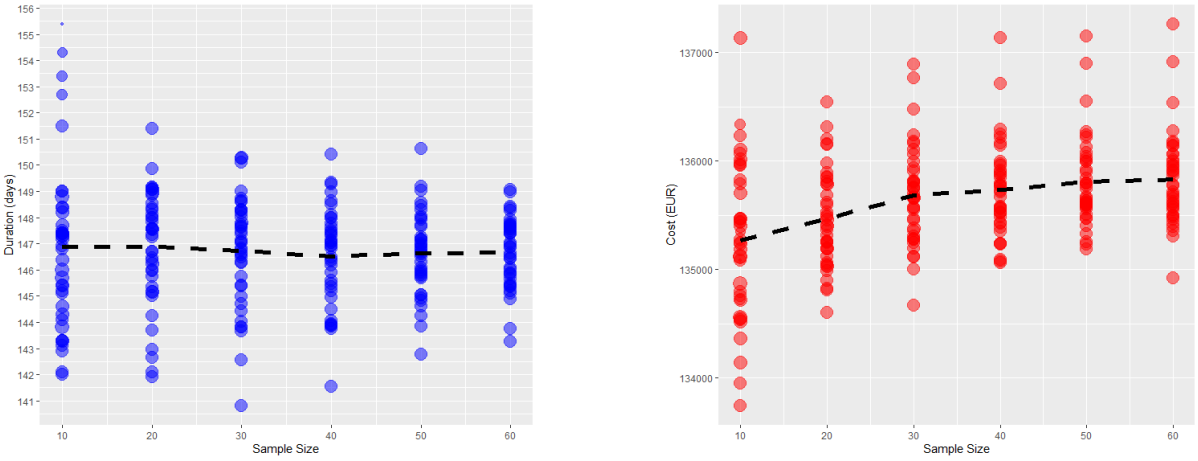

(a) *Completion Time* **(b)** *Cost*

Figure 6.2: *Stability Test (WC28325)*

sample size of 30. The elbow method criteria basically relies on the existence of a point where the indicator whose stability we intend to asses reduces its slope, as is the case at $S = 30$ for Figure 6.2b. Is is worth to mention that the consideration of a greater number of scenarios means that the (production and distribution) decisions made must be adequate for a bigger number of situations, which makes these decision more conservative and farther from the decisions that we would make under perfect information, resulting in a higher expected total cost. To illustrate this, let us consider a single scenario: we would optimize under perfect information, and make perfect decisions; If we add a second scenario, then chances are that the decisions that I will make are not optimal for any of the two scenarios, but there is an optimal 'compromised solution' tat considers the trade-offs and works well enough under both scenarios. Let us recall that, while we optimize an expected value, we have several constraints that must be satisfied for each existing scenario: the more scenarios we add, we will still find a solution that is optimal in expected value, but there will be more situations for which the solution has to be acceptable: basically we are adding constraints, so the total cost cannot get any better.

In addition to stability of the mean, we must consider stability for the standard deviation. Figures in 6.3 offer a better appreciation of the variability around the mean: the range and standard deviations are shown for Completion Time and Cost, respectively. The elbow method would advice a sample size of 20 based on these two plots, where standard deviation is close to 2 days and €500 respectively, i.e. coefficients of variation of 1.36 % and 0.37 % respectively,

which is fairly acceptable.

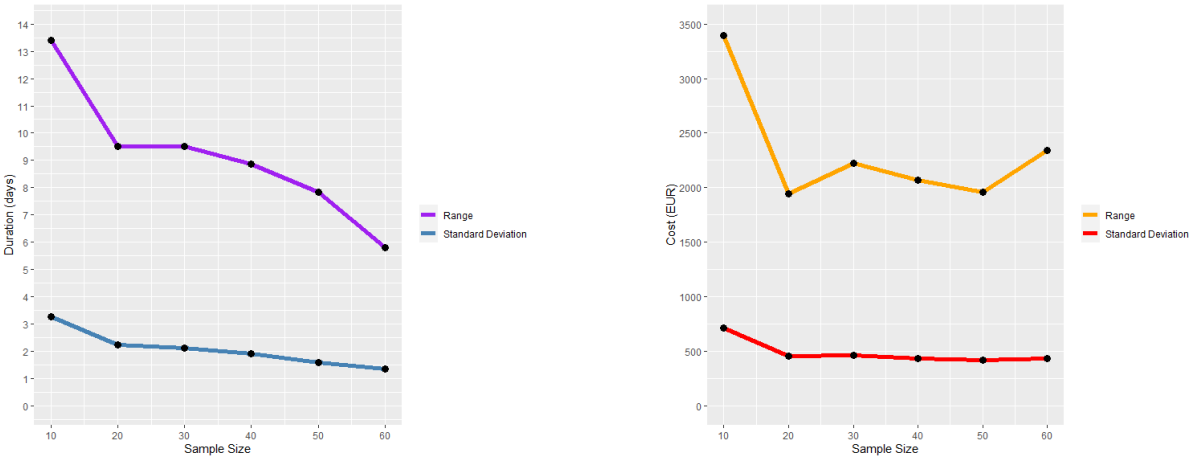

(a) *Completion Time* **(b)** *Cost*

Figure 6.3: *Spread Stability Test (WC28325)*

The joint consideration of Figures in 6.2 and 6.3 indicates that a sample size of 30 is good enough. We have decided to be a little more conservative, picking a sample size of 40, although this increases the computational effort by a 66% (see Figure 6.4).

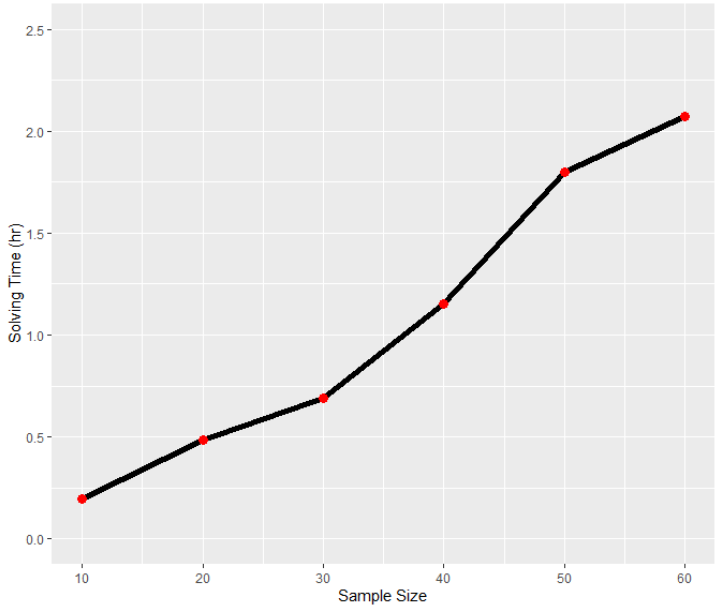

Figure 6.4: *Average solving time by stage for each sample size*

6.1.2. Conditional Sampling and Sequential Optimization

As explained in Section 4.6, in order to obtain a valid consistent estimator, we have performed conditional sampling (Shapiro 2003), limiting the randomness in future stages, and fixing the

past stages arrivals to the actual realizations. Also, we have performed sequential optimization, with one optimization per stage, using the previous stage optimal values as constraints for optimizing forward.

Numerical Results

Figure 6.5 shows the progression of the expected completion time during the optimization, with the corresponding upper and lower bounds, i.e. with the greatest and lowest completion times among the 40 different scenarios considered in each period of time (day). Please note that uncertainty disappears on day 133, which is when the last new patient was enrolled, and therefore the problem became deterministic until the trial completion.

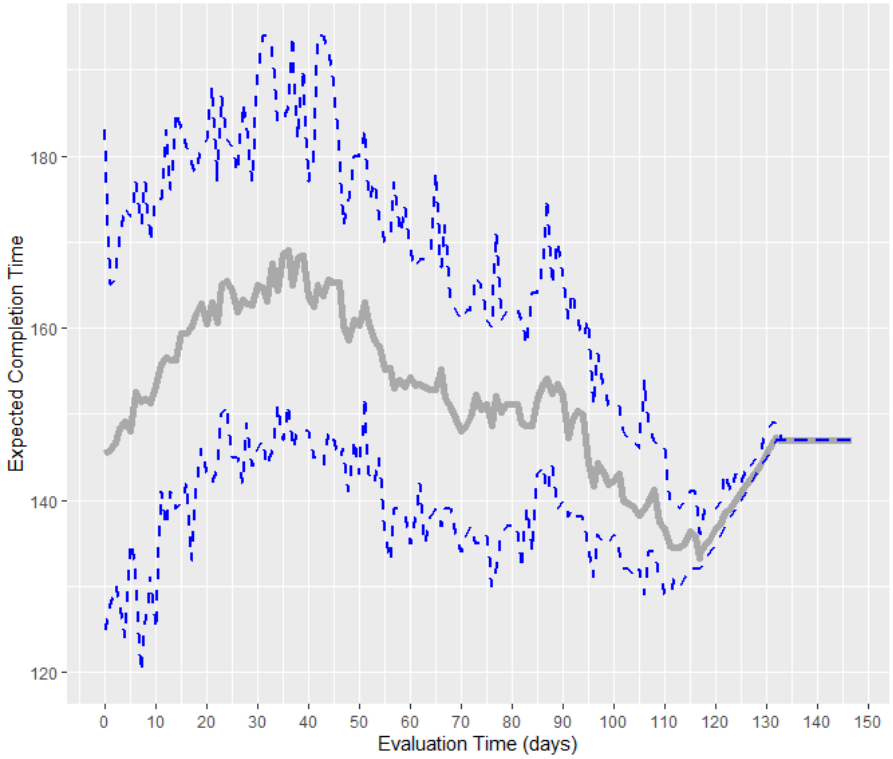

Figure 6.5: *Evolution of the estimation made by the model for trial completion time as the enrolment evolved over time*

Also from Figure 6.5 results interesting that the model initially overestimates the time that the trial is expected to take, growing in the first 36 days to an expected 169 days completion time. The estimation decreases progressively until day 117, when the expected completion time is estimated in 133 days. For the next 16 days, the estimated completion time grows to a definitive 147 days. The explanation for this lies on the particular pattern that the actual enrolment for this trial exhibited, which we can see in Figure 6.6.

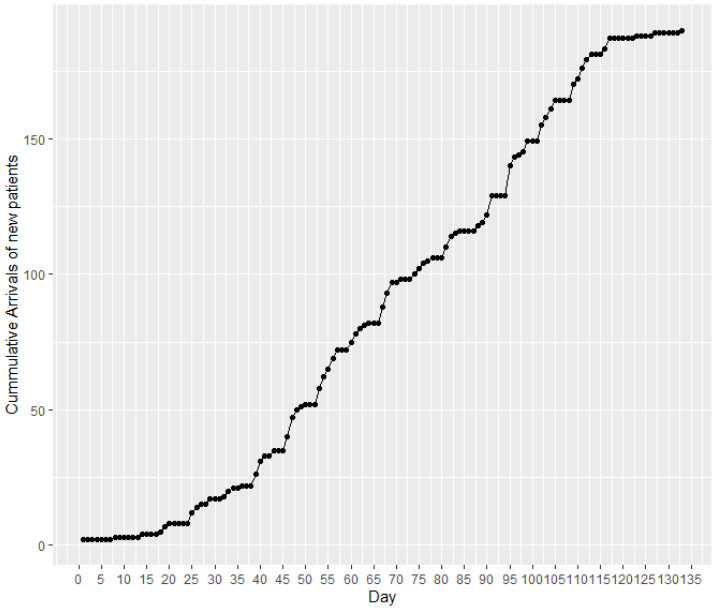

Figure 6.6: *Progress of enrolment in trial WC28325*

As we can see, the enrolment is extremely slow for the first 25 days, making the model estimate a larger completion time; then enrolment accelerates until day 117, when it is 3 patients away from reaching the patient horizon of 190 patients: at this point, with an expected 1.4 patients/day, it makes sense to estimate 2-3 more days for enrolling the last 3 patients, and 2 more weeks to administer doses 2 and 3 for these patients, making a total of 133 or 134 days. The actual enrolment, however, decelerated after day 117, taking 16 days to enrol the last 3 patients, causing the adjustments that the model successfully made as it received new information while the trial was progressing.

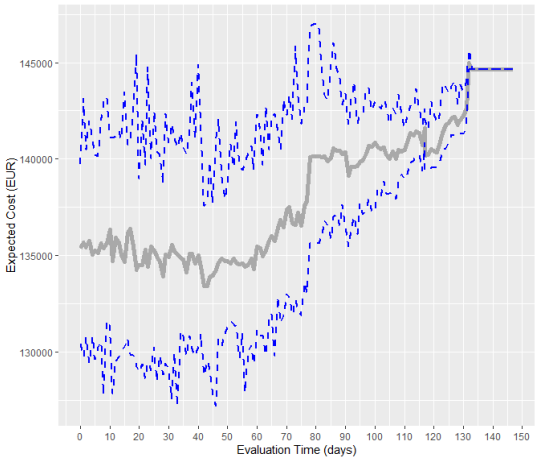

(a) *Total Cost*

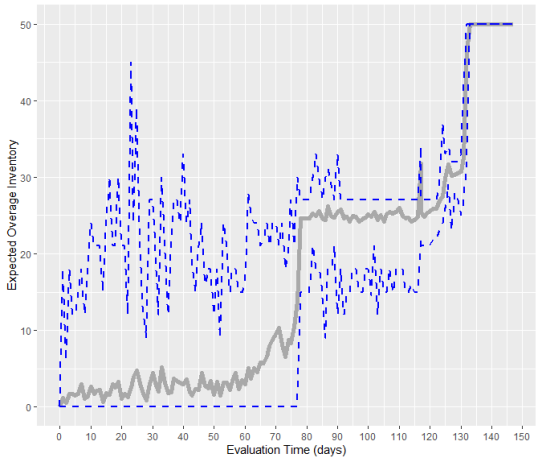

(b) *Overage Inventory*

Figure 6.7: *Evolution of the estimations made by the model as the enrolment evolved over time*

Figures 6.7a and 6.7b show the corresponding sequence for Cost and Overage Inventory: as

can be seen, this optimization estimates that the trial will be completed in 147 days, with a cost of €144,660, resulting of producing a total 620 units (394 initial stock, and 2 more production runs of 203 and 23 units), and 15 shipments (108 units) from CW to CDs. The overage inventory at the end of the trial is 50 units, distributed as shown in Table 6.4, where it is worth to observe the absence of final inventory in the middle echelon (Country Depots), which indicates that the model is effectively avoiding to incur in unnecessary costs by pushing inventory downstream; with that being said, it calls our attention the unexpectedly high inventory (9 units) in site 22: the reasons for this is that, in the random sampling made at $t = 112$, a very extreme and unlikely scenario was sampled, which led to the decision of sending 9 units from CW; however, later samplings did not reveal such demand on site 22, and therefore this inventory remained unused.

NODE	Final Inventory
CW	16
Site 2 (ARG)	1
Site 3 (COL)	1
Site 4 (COL)	1
Site 5 (COL)	1
Site 6 (GUA)	1
Site 7 (GUA)	2
Site 8 (MEX)	2
Site 9 (MEX)	2
Site 10 (MEX)	1
Site 11 (MEX)	1
Site 12 (MEX)	2
Site 14 (USA)	1
Site 15 (USA)	1
Site 16 (USA)	2
Site 17 (USA)	2
Site 18 (USA)	2
Site 20 (USA)	2
Site 22 (USA)	9

Table 6.4: *Distribution of final inventory across the WC28325 network*

For this case study shelf life was assumed long enough as for neglecting perishability, but still two additional production runs were necessary in order to minimize cost and overage inventory: 203 units on day 78, and 23 units on day 132, right before the problem became deterministic. This choice of having 2 additional production runs, far away from daily productions, but also from a single production is well aligned with the intuition built at the beginning of Section 6.1.

With respect to figures 6.7a and 6.7b, a shift in both estimated cost and overage inventory takes place exactly when production runs take place, on days 78 and 132. The explanation for this is that the model makes initial inventory allocation and posterior distribution according to the previous arrivals pattern and the assumed arrival rates λ_j; the scenarios considered before day 78 may suggest that the produced inventory (or the addition of a small production) would be enough for completing the trial, but the actual enrolment revealed the need of further production runs, which caused the observed adjustments in the estimation of overage inventory and cost.

6.1.3. Scenario Reduction

As shown in Fig. 6.4, the average time to run an optimization with 40 scenarios took 1.15 hr; taking into consideration that we estimated this trial to take about 147 days to complete, the expected running time was 169 hr (about a week). However, the addition of new constraints in each stage has resulted in the need of nearly 3 times the initial estimation to complete the simulation. This is why we developed the scenario reduction procedure in Chapter 5.

Determination of dissimilarity metric

As explained in Section 5.2.2, it is necessary to determine the weight of each term in (eq. 5.2) before performing K-medoids++.

We have solved the 1st stage of this SP with $S = 40$, using 10 different random seeds, creating a set of $10 \times (1 + 2 + ... + 39) = 10 \times 40 \times (40 - 1)/2 = 7800$ pairs of scenarios, for which we have calculated:

- Each term in the RHS of (eq. 5.2)
- The LHS in (eq. 5.3)

The OLS linear regression resulted in $AdjR^2 = 0.9954$ and the 3 coefficients being statistically significant with p-values in the order of 10^{-16}. However, since we had $c_1 = -0.27776 < 0$, we have used NNLS, resulting $c_1 = 0$, $c_2 = 2.1926$ and $c_3 = 0.3514$, with a deviance just 8.1% greater than the one resulting from OLS, and coefficient of determination $R^2 = 0.9871$.

Determination of adequate number of clusters

With the distance metric fully defined, we have performed K-medoids++ clustering on the samples of size $S' = 60$, the greatest one that we had considered for the stability test, trying number of clusters ranging from $K = 1$ (deterministic equivalent) to $K = 10$. The K-medoids algorithm is NP-hard, and therefore we have used the python FAST-PAM implementation by (Schubert and Lenssen 2022).

Figures in 6.8 show the distribution of the relative error for Completion Time and Cost (both calculated using K clusters with K ranging from 1 to 10), compared with the completion time obtained with the direct solution for $S' = 60$.

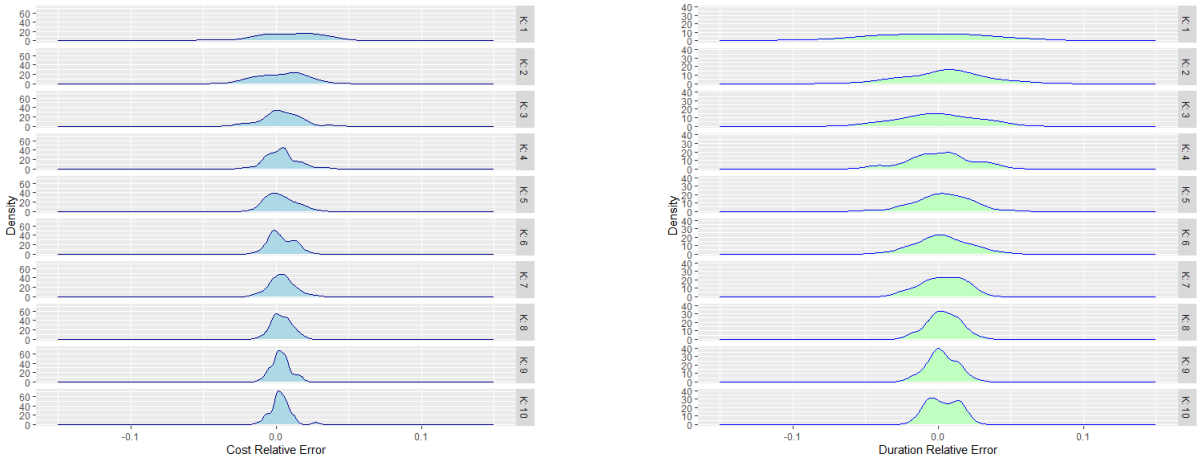

(a) *Cost* **(b)** *Completion Time*

Figure 6.8: *Relative error density plots*

As we can see, replacing by a sole most representative scenario ($K = 1$) would lead to huge relative error. A choice of $K = 8$ would result in a mean error below 0.4% in both objectives, with a standard deviation no greater than 1%. We will be conservative and pick $K = 9$.

6.1.4. Sequential Optimization with $S' = 60$ and $K = 9$

Just like in the sequential version with $S = 40$ and no clustering, the proposed solution required 2 additional production runs (days 91 and 130, instead of days 78 and 132; including the initial production run, the sequence according to this model is 420 + 178 + 22, which is reasonably close to the 394 + 203 + 23 resulting from the previous model) and led to virtually the same evolution of estimated completion time, as shown in Figure 6.9.

With regards to cost and overage inventory, final results are 0.7% off in cost and completely identical in overage inventory (figures in 6.10). Essentially, the difference is that the model with clustering (which we named S60K9) uses 16 deliveries from CW to CDs instead of 15 (123 units, instead of 108) as the one without clustering (which we named S40).

Lastly and no less importantly, model S60K9 took only 55 hr to run over the 147 periods, which represents a reduction of 88.4 % over the 473 hr necessary to sun model S40.

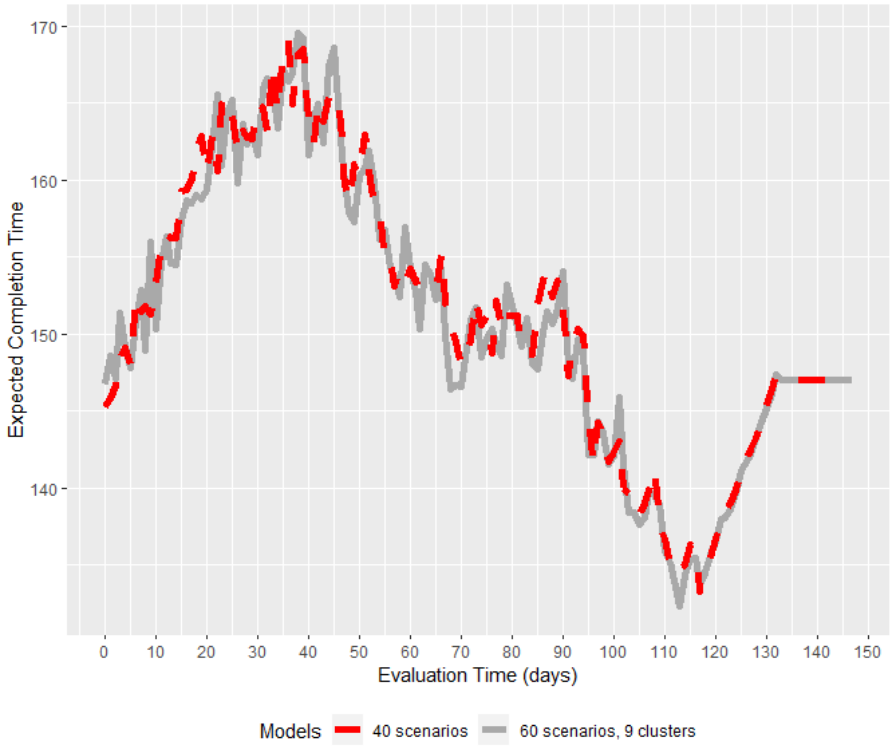

Figure 6.9: *The model that optimizes after clustering scenarios perfectly reproduces estimations made by the original mode*

6.1.5. Comparison with other policies

A. Fleischhacker, Ninh, and Y. Zhao 2015

We have coded the model presented in (A. Fleischhacker, Ninh, and Y. Zhao 2015), which led to:

- CW: Initial inventory of 569 units

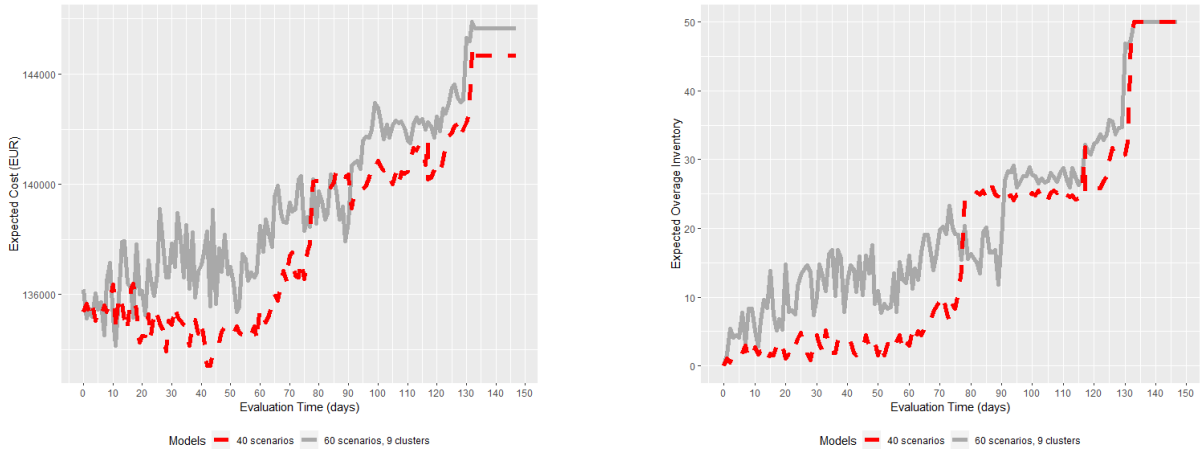

(a) *Cost* **(b)** *overage inventory*

Figure 6.10: *The model that optimizes after clustering scenarios follows pretty much the same pattern both in cost and overage inventory estimation*

- CDs: Continuous review-batch order policy, initial inventory of 2, reorder point of 1, and reorder quantity 1 in every CD.
- Enrolment sites: Continuous review base-stock policy with target level equal to 1 in every site.

The total production in this model is then 599 units. Replenishment policies used by the CW, CD, and sites described above are as such because the low enrolment rates, which leads to inventory being located in the most possible centralized way, especially taking into consideration that this model assumes patients to wait as much as necessary, and aims only for cost minimization, without any consideration on the completion time.

The objective function (total cost) in the optimization leads to an optimal €154,195. The comparison (see Figure 6.11) shows that our model leads to higher overage inventory, with consequent slightly higher production cost (+3.5%), but with a much lower distribution cost (-30%), thanks to the much more efficient deliveries (96 % less shipments from CW to CDs). The high level of match between the areas delimited with red and purple lines also indicates that our model obtained after clustering greatly approximates the results of our main model.

With regards to the maximum waiting time, our model satisfied the requirement of 4 days maximum waiting time, while the benchmark model required to relax that constraint, taking a maximum of 5 days to fill demand in certain cases: it would have lead to dropouts in one of the sites in Colombia if we had not allowed a maximum waiting time of 5 days.

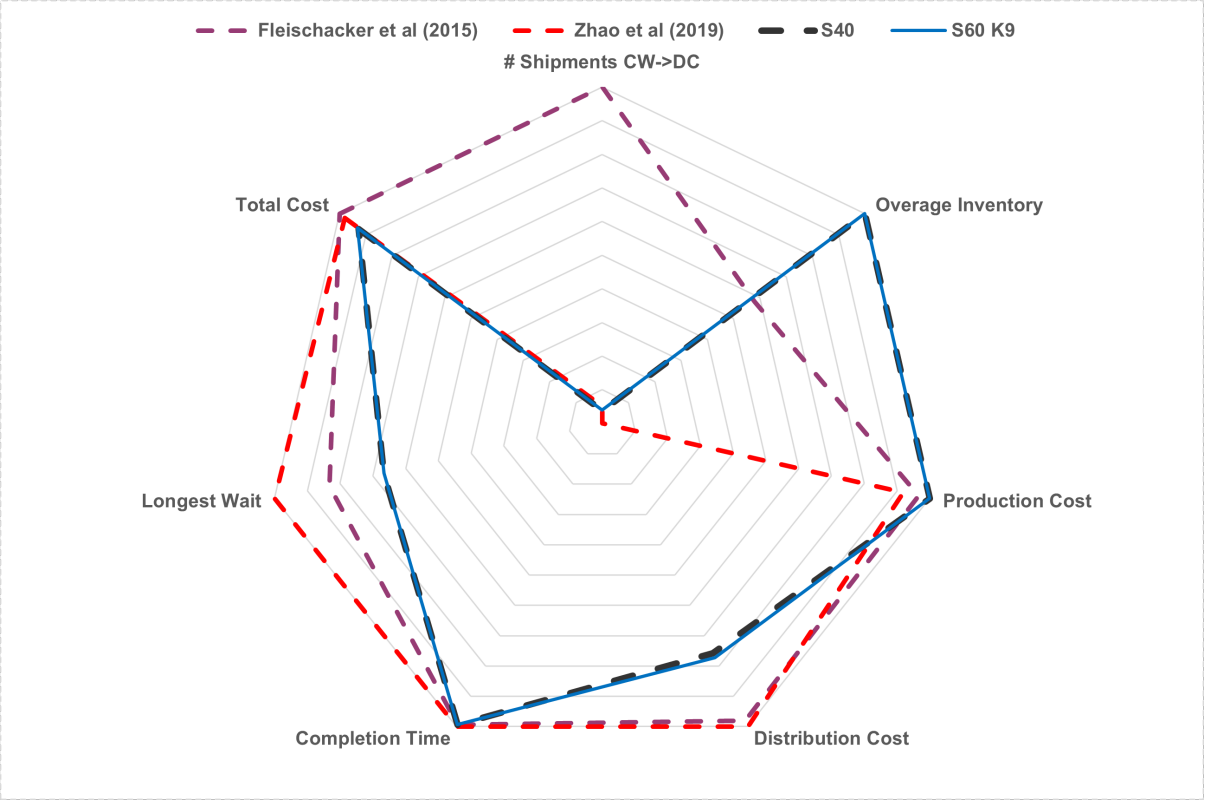

Figure 6.11: *Comparison of our two models against two benchmarks*

H. Zhao et al. 2019

We have coded the model presented in (H. Zhao et al. 2019), adapting it to our clinical trial setting by:

i Disabling the possibility of emergency deliveries: The authors of this paper enable the possibility of expensive deliveries with negligible lead time. We do not consider this possibility.

ii Inability of cross-shipping between enrolment sites: the authors enable sourcing from one enrolment site to another: we do not enable this possibility.

iii Inability of reverse sourcing: The authors of this paper enable reverse sourcing from enrolment sites to Distribution Centers: we do not enable this possibility.

iv We set the dropout rate to 0: The authors have assumed a fixed dropout rate in each period, not related to the supply chain decisions. We have assumed that these dropouts do not happen.

It is worth to mention that emergency deliveries, reverse sourcing and transshipment

would work in favour of granting feasibility, but they would not make the solution cheaper. Also a non-zero dropout proportion would create the need of more production, increasing the cost. In summary, all these adaptations tend to reduce the cost of the solution provided by the benchmark model, which anyway provides a more expensive solution than ours: the application of this model indicates for clinical trial WC28325 a completion time of 148 days and a total cost of €153,127.60 i.e. about 2% cheaper than A. Fleischhacker, Ninh, and Y. Zhao 2015, but 5% more expensive than our S60K9 model. The overage inventory is, on the contrary, 0 units as a result of:

- Assumption of patients waiting as much as necessary
- Non-commitment to a stationary policy

The waiting time, however, can be as long as 6 days, which violates the restriction of maximum waiting time of 4 days that we have imposed to our model. Results are shown in Figure 6.11, where the extremely big number of shipments of A. Fleischhacker, Ninh, and Y. Zhao 2015 does not allow to appreciate that H. Zhao et al. 2019 suggests an expected 20.53 shipments from CW to CDs, compared to the 16 necessary for S60K9 model.

6.1.6. Sensitivity over arrival rates

As can be seen in Figure 6.9, both the S40 and S60K9 models share not only a high match in the estimation made day by day for the completion time of the trial, but also it is noticeable that, in both of them, the initial estimation of the completion time is about the same estimation made at the end, i.e. about 147 days. The reason for this extremely good initial estimation of the expected completion time is that the arrival rates used for this model turned out to be the right ones. However, truth is that the arrival rates are hard to know in advance with good accuracy (Lefew, Ninh, and Anisimov 2021), and therefore it is worth to explore how this model would work when arrival rates are off.

Overestimation of arrival rates

We have applied S60K9 model with arrival rates 3 times bigger than the original ones, for every enrolment site. For the sequential optimization, however, the actual arrivals have been used (same as in previous sections). Figure 6.12a shows that the high hypothetical arrival rates

provoke a sub-estimation of the completion time, which grows consistently during nearly the whole trial. Some interesting facts about this simulation are:

1. The high arrival rates resulted in more inventory being initially situated downstream, since the system needed to be prepared for more frequent arrivals than they actually were.

2. Most of these enforced anticipated decisions were taken with wrong information about arrival rates; better decisions would have been taken with updated information about the arrivals, but the lead times and limited patience of the patients forced the model to take anticipated decisions that were affected by the wrong estimation of arrival rates: as a result of this, this setting needed 20 shipments from CW to DCs (25 % more than the one with right data), and 4 additional production runs (a total 5, including the initial inventory) instead of 2, to compensate the inventory wrongly pushed downstream at the beginning.

3. As Figure 6.12b indicates, the cost of arrival rates overestimation is at the end just 0.2%. We must take into account that this is partially due to the very low fixed cost that we had for this specific CT, where we also assumed that production cost was fully variable: the excess of 2 production runs and 4 international deliveries are likely to cause a greater cost excess under other settings. Figure 6.13 shows what we believe is a more comprehensive comparison. With this being said, the sequential optimization after updated information proves to be adequate to make the necessary adjustments and make adequate decisions with the available information. In this setting, where at least the arrival rate proportions are correctly estimated, the model performs considerable well with very little difference in cost and overage inventory.

Combination of over and under-estimation of arrival rates

While the overestimation of arrival rates shown in 6.1.6 causes anticipated and sub-optimal decisions, the fact that we have considered a fixed factor increasing the arrival rates in the same proportion still means that the proportion of arrivals in each site is essentially rightly estimated (note from Figure 6.13 that overage inventory is approximately the same as the one with accurate arrival rates).

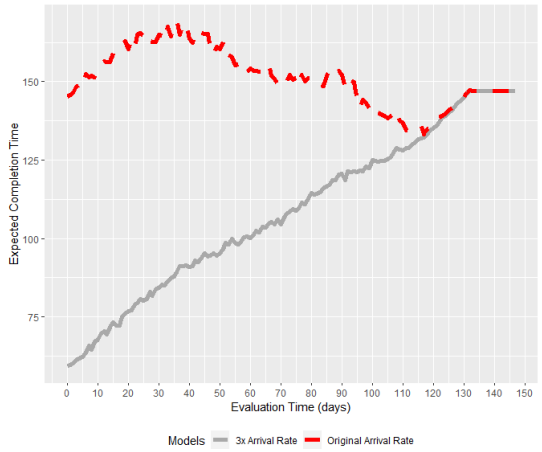

(a) *Impact of overestimated arrival rates on completion time estimation*

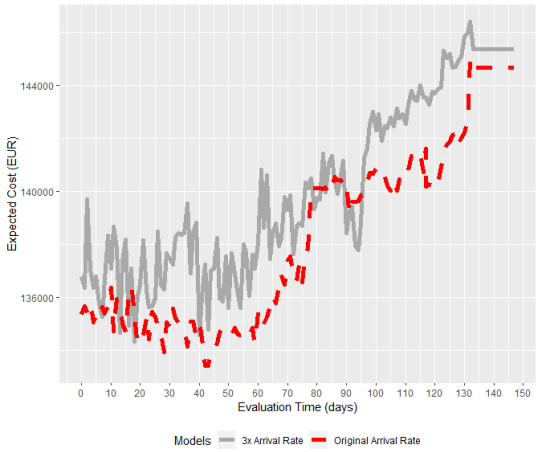

(b) *Impact of overestimated arrival rates on cost estimation*

Figure 6.12: *Impact of overestimated arrival rates on completion time and cost estimation*

Let us now consider the case where half the sites have their arrival rates overestimated, and half underestimated. Table 6.5 details the changes applied to original arrival rates, where rates for half the sites have been increased by a factor of 3, and have been reduced by the same factor for the other half. In this experimental design, we have:

- Sites in Argentina overestimate their arrival rates ($\times 3$)
- Sites in Guatemala underestimate their arrival rates ($\times \frac{1}{3}$)
- Sites in Colombia an Mexico have a mix of over and underestimation, with slightly more than half the arrival rates being overestimated.
- Sites in USA have a mix of over and underestimation, with slightly more than half the arrival rates being underestimated.

Figure 6.14 shows the effect of these changes in the initial inventory allocated in each country. Interactions here are far too complex as for going far with the interpretation, but we can clearly see that, while inventory allocated in Guatemala has decreased by a factor greater than the reduction applied to its arrival rates ($\times \frac{1}{5.5}$ vs. $\times \frac{1}{3}$), inventory allocated in Argentina increases in more modest proportion than the arrival rates overestimation ($\times 1.85$ vs $\times 3$).

In Colombia and Mexico, where the average arrival rates have increased by 36% and 73% respectively, the initial inventory have respectively decreased by 38% and increased by 24%.

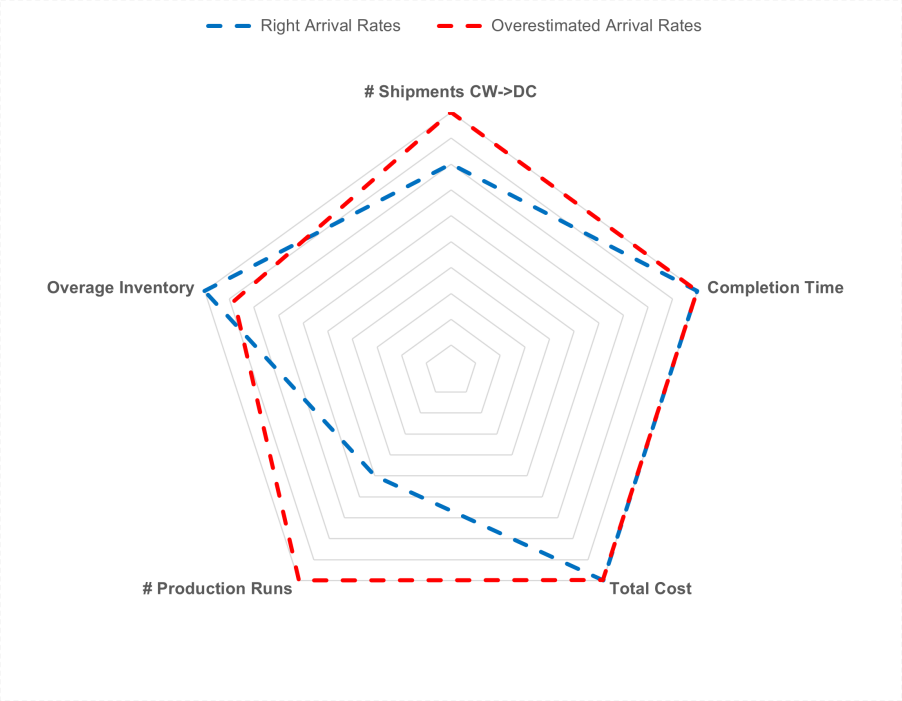

Figure 6.13: *Comparison of results with right arrival rates and '×3' overestimated arrival rates*

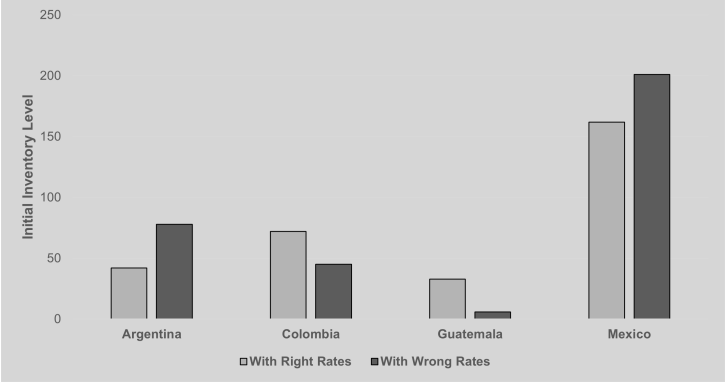

Figure 6.14: *Change in initial inventory country by country, under arrival rates changes simulated in section 6.1.6*

In any case, wrong *a priori* estimations of the arrival rates are likely to lead to allocation of too much inventory in some CDs (causing an excess of overage inventory), and too little inventory in others (causing extra production runs to compensate the inventory that had been wrongly moved to other CDs).

This is represented in Figure 6.15, where it is particularly noticeable the excess of overage inventory and number of shipments, both exceeding twice the baseline. Furthermore, deviation in cost and overage inventory get bigger over time: while there was an initial cost difference inherent to the dynamics expected for the estimated arrival rates, this difference

Enrolment Site	Change
Site 1 (ARG)	$\lambda_1 \leftarrow 3 \times \lambda_1$
Site 2 (ARG)	$\lambda_2 \leftarrow 3 \times \lambda_2$
Site 3 (COL)	$\lambda_3 \leftarrow \frac{1}{3} \times \lambda_3$
Site 4 (COL)	$\lambda_4 \leftarrow 3 \times \lambda_4$
Site 5 (COL)	$\lambda_5 \leftarrow 3 \times \lambda_5$
Site 6 (GUA)	$\lambda_6 \leftarrow \frac{1}{3} \times \lambda_6$
Site 7 (GUA)	$\lambda_7 \leftarrow \frac{1}{3} \times \lambda_7$
Site 8 (MEX)	$\lambda_8 \leftarrow \frac{1}{3} \times \lambda_8$
Site 9 (MEX)	$\lambda_9 \leftarrow \frac{1}{3} \times \lambda_9$
Site 10 (MEX)	$\lambda_{10} \leftarrow 3 \times \lambda_{10}$
Site 11 (MEX)	$\lambda_{11} \leftarrow 3 \times \lambda_{11}$
Site 12 (MEX)	$\lambda_{12} \leftarrow 3 \times \lambda_{12}$
Site 13 (USA)	$\lambda_{13} \leftarrow \frac{1}{3} \times \lambda_{13}$
Site 14 (USA)	$\lambda_{14} \leftarrow \frac{1}{3} \times \lambda_{14}$
Site 15 (USA)	$\lambda_{15} \leftarrow 3 \times \lambda_{15}$
Site 16 (USA)	$\lambda_{16} \leftarrow \frac{1}{3} \times \lambda_{16}$
Site 17 (USA)	$\lambda_{17} \leftarrow 3 \times \lambda_{17}$
Site 18 (USA)	$\lambda_{18} \leftarrow \frac{1}{3} \times \lambda_{18}$
Site 19 (USA)	$\lambda_{19} \leftarrow 3 \times \lambda_{19}$
Site 20 (USA)	$\lambda_{20} \leftarrow \frac{1}{3} \times \lambda_{20}$
Site 21 (USA)	$\lambda_{21} \leftarrow 3 \times \lambda_{21}$
Site 22 (USA)	$\lambda_{22} \leftarrow \frac{1}{3} \times \lambda_{22}$

Table 6.5: *Changes applied to the arrival rates site by site in section 6.1.6*

grows by a factor of $\times 2.75$ during the 147 days that the trial lasts (Fig. 6.16a); likewise, the difference in estimated overage inventory growth from nothing to a ratio of $\times 2.14$ (Fig. 6.16b).

With regards to the number of shipments from CW to CDs, which was as low as 16 with right rates and as high as 39 with the wrong ones, it is interesting to note that none of those 39 shipments was to Argentina, where the arrival rates had been overestimated, and therefore all the necessary inventory (plus 18 extra units) was initially allocated in this CD. The remaining 3 CDs, on the contrary, needed between 1.7 and 8 times as many shipments as they needed with right rates (Fig. 6.17).

As shown in this section, it is important to remark how much a good knowledge of the new patients arrival process helps in the decision making. Although it is outside of the scope of the present research, the development of methods to produce precise *a priori* estimations of arrival rates is highly encouraged in order to feed our model with appropriate inputs.

Figure 6.15: *Performance comparison of model with right arrival rates and model described in 6.1.6*

6.1.7. Consideration of a fixed production cost

In previous sections, within the context of clinical trial WC28325, we have shown that:

i. **Stability:** A sample size of 40 is good enough to represent the uncertainty in the new patients arrival process.

ii. **Scenario Reduction:** A sample of 60 scenarios, grouped in 9 clusters produces a result virtually identical to the one with 40 scenarios without clustering, with the practical advantage of a much shorter running time.

iii. **Sensitivity:** Analysis in Section 6.1.6 show the impact of a good estimation of arrival rates over the quality of the solution provided by our model.

iv. **Results quality compared with other policies:** The comparison made in Section 6.1.5 showed that our model performs better in general.

However, as mentioned before, we had made some significant simplifying assumptions in order to make a fair comparison. In particular, our model required a total 3 production runs, in contrast with the only production run assumed in (A. Fleischhacker, Ninh, and Y. Zhao 2015). This is, at least partially, due to the fact that we have assumed no fixed production cost, allowing the model to benefit from producing small batches and postponing further production runs.

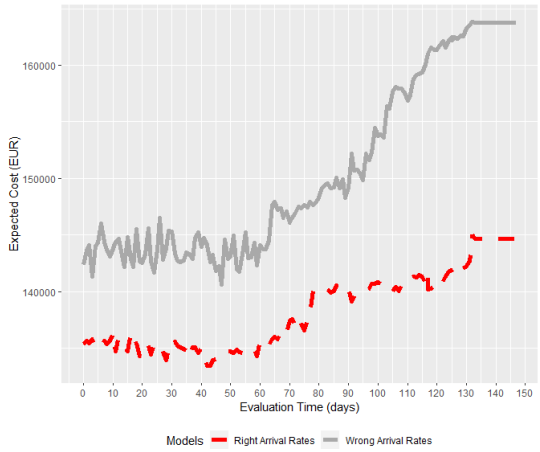

(a) *Cost*

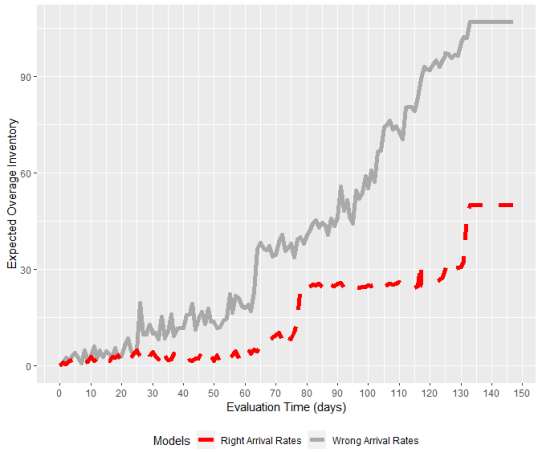

(b) *Overage Inventory*

Figure 6.16: *Comparison of evolution of cost and overage inventory estimation when rates deviate as described in 6.1.6*

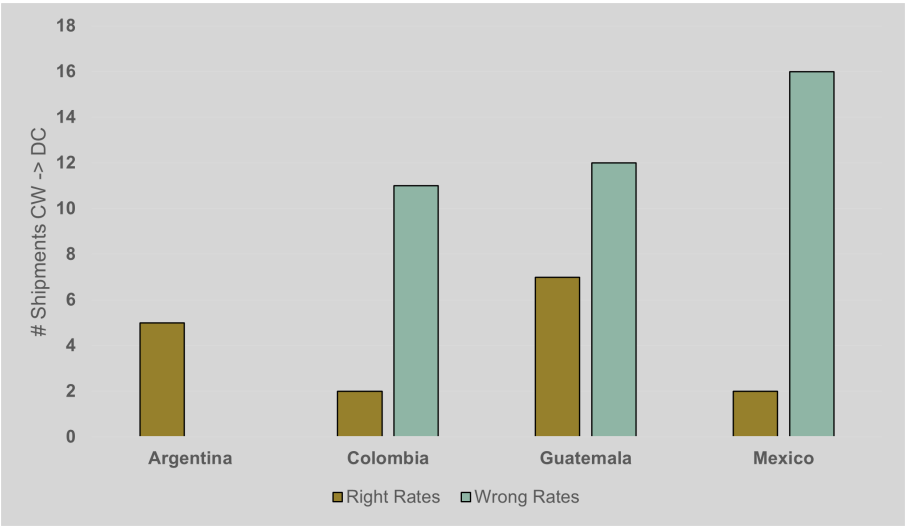

Figure 6.17: *Comparison of number of shipments from CW to each DC, with right arrival rates and with wrong arrival rates as described in 6.1.6*

We will now solve our S60K9 model with the only change of using a fixed production cost $f = €10{,}000$.

As expected, the initial inventory is now 570 units (instead of 420), i.e. just the bare minimum enough as for filling all 190 patients with 3 doses each. However, 2 additional production runs have still been necessary, although they have been postponed to days 124 (32 units) and 133 (4 units), which is the moment when the last new patient arrives.

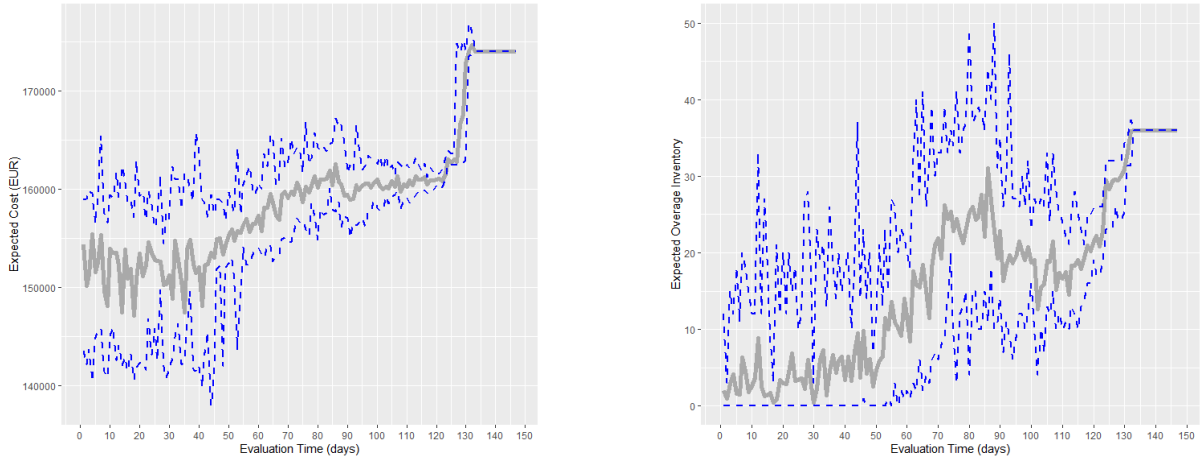

(a) *Cost* **(b)** *Overage Inventory*

Figure 6.18: *Evolution of total cost and overage inventory estimation for clinical trial WC28325 with a fixed production cost of €10,000*

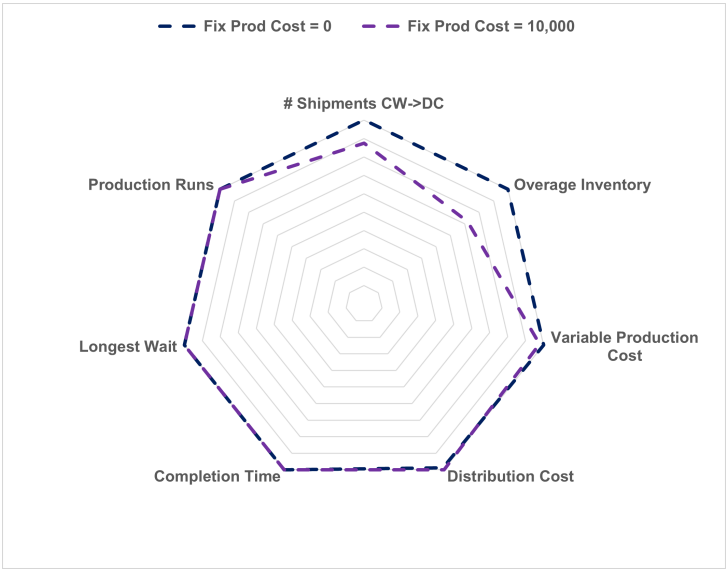

Figure 6.19: *Comparison WC28325 with and without fixed production cost*

Let us also observe that total cost (Fig. 6.18a) is now €174,067; this is €28,417 more expensive than the case without fixed cost, but subtracting the €30,000 corresponding to the fixed cost of the 3 production runs, there is a saving of €1,583 that we can associate to better production decisions due to this postponement, also evidenced in the reduction of the number of shipments to CDs from 16 to 14. In fact ending inventory (Fig. 6.18b) is now 36 units, i.e. a 28% less than when the non-consideration of fixed production cost encouraged production runs to take place earlier (Fig. 6.19).

Let us note in Figure 6.18a that, although there is a production run on day 124, there is no cost shift on that day: there is, however, a shift of over €10,000 on day 131. This is because, even several weeks before day 124, the model was counting on a second production run to be made in the future; the need of a third production run (that finally happened on day 133) was not revealed until day 131, when the actual demand took place quite differently from previous samplings, making this small production run necessary.

6.2. BO29389: A phase II CT for malignant neoplasm of branchus and lung

6.2.1. Modeling considerations

CT BO29389 enrolled only 49 patients between December 2013 and June 2016. Not only the recruitment period has been very long, but also this trial administered 15 doses to each patient in intervals of 3 weeks, in a network with 18 sites distributed over 6 countries. In terms of the parameters defined in Section 4.4.2, $m = 15$ and $N = 18$ are unavoidable, but we can certainly explore the possibility of using weekly stages instead of daily, in order to get the model running faster. We would then need $L \approx 80$ instead of $L \approx 560$, and $\tau = 3$ instead of $\tau = 21$.

However, this simplification forces us to make some approximations in enrolment sites lead-times and maximum waiting time δ: these quantities all range from 4 to 9 days, and we will be approximating using 1 week for all of them. Likewise, for the CD lead times of 23 and 28 days we use 3 and 4 weeks.

Figure 6.20 shows the structure of the trial, and Table 6.6 shows the arrival rates and both actual (ALT) and model (MLT) lead times for each enrolment site (Lerma and Lichtenstein 2020). Table 6.7 is the analogous for country depots. As we did in Section 6.1, transportation costs to enrolment sites are set to €0.

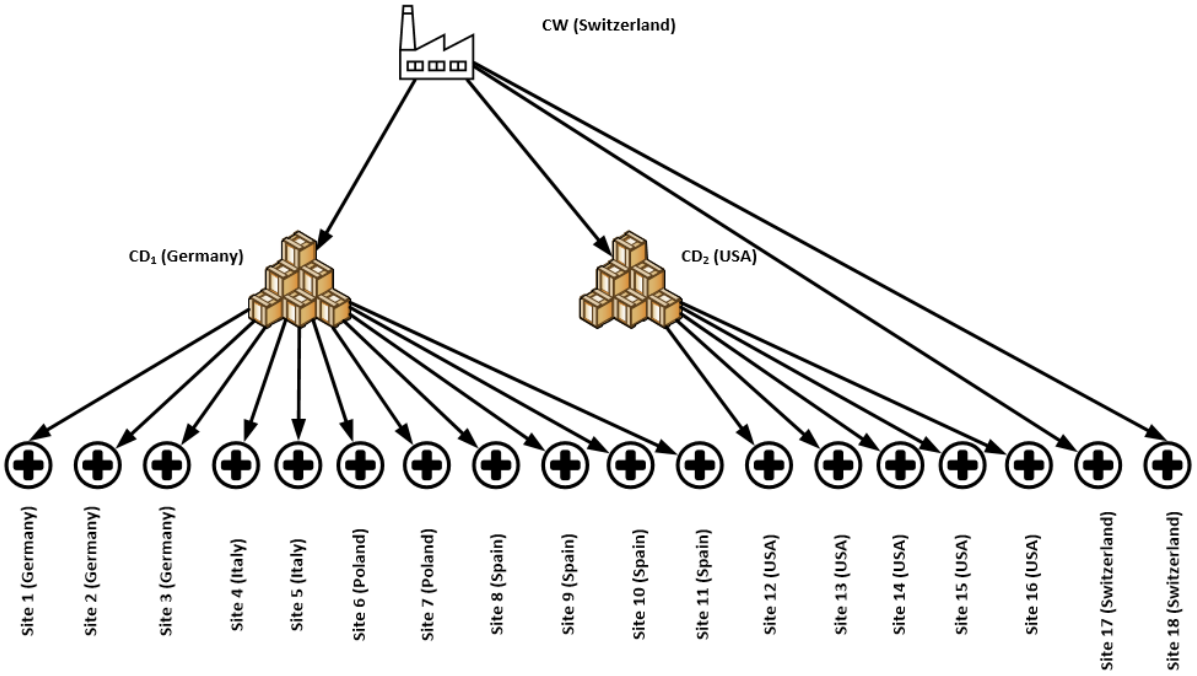

Figure 6.20: *Clinical Trial BO29389*

6.2.2. Numerical Results

Stability Test

Following the same procedure already described in 6.1.1, we have obtained, as figures 6.21a, 6.21b, 6.22a and 6.22b show, that a sample size of 40 is again sufficiently good as for representing the stochastic nature up to at least the 2nd momentum.

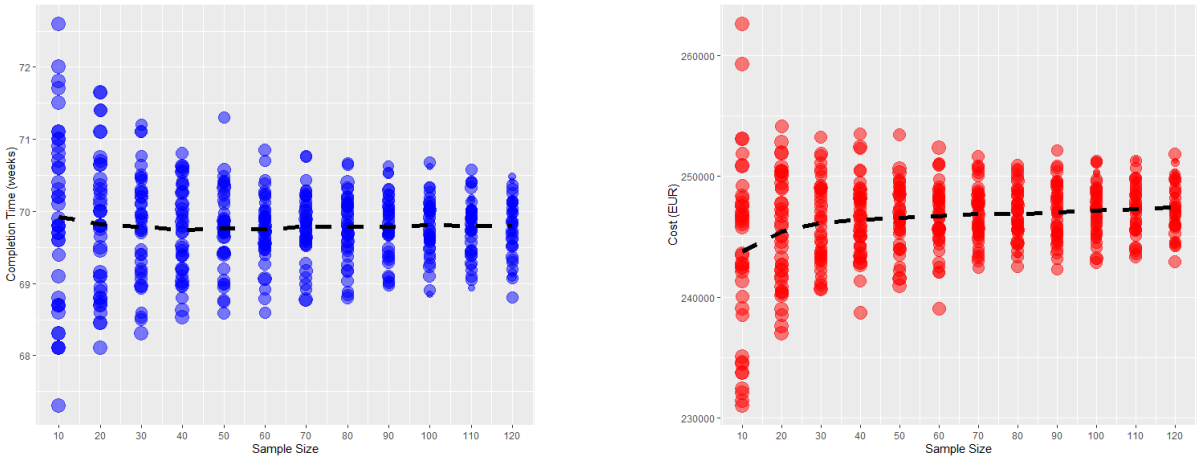

(a) *Completion Time* **(b)** *Cost*

Figure 6.21: *Stability Test (BO29389)*

j	i	Country	λ(patients/week)	ALT (days)	MLT (weeks)
1	1	Germany	0.093	9	1
2	1	Germany	0.140	9	1
3	1	Germany	0.070	7	1
4	2	Italy	0.047	7	1
5	2	Italy	0.140	7	1
6	2	Poland	0.093	7	1
7	2	Poland	0.047	7	1
8	2	Spain	0.070	7	1
9	2	Spain	0.093	7	1
10	2	Spain	0.093	7	1
11	2	Spain	0.117	7	1
12	2	USA	0.117	9	1
13	2	USA	0.070	9	1
14	2	USA	0.117	9	1
15	2	USA	0.187	9	1
16	2	USA	0.093	9	1
17	0	Switzerland	0.163	4	1
18	0	Switzerland	0.047	4	1

Table 6.6: *Enrolment sites data for CT BO29389*

CD	i	Fixed Delivery Cost	Variable Delivery Cost	ALT (days)	MLT (weeks)
Germany	1	€12.76	€180	23	3
USA	2	€111.88	€150	28	4

Table 6.7: *Country depots data for CT BO29389*

Sequential Optimization

As shown in Figure 6.23, the estimated arrival rates were too high, causing a too short initial estimation of trial completion time (around 70 weeks, well below the 121 weeks that the model estimated at the end, which happen to approximate very well the 841 days $\simeq$ 120.14 weeks that the trial actually lasted.

The model estimates a total cost of €238,210 and an ending inventory of 62 units (Figure 6.24), with only one additional production run (377 units, on week 50, in addition to the initial production of 420 units) and a total of 9 shipments (397 units) from CW to CDs.

Let us note that two major jumps exist in Figure 6.24b: the first of them, in week 50, is associated to the only additional production run: the uncertainty on the total number of recruited patients opens the possibility for an ending inventory ranging from 0 to 62, with

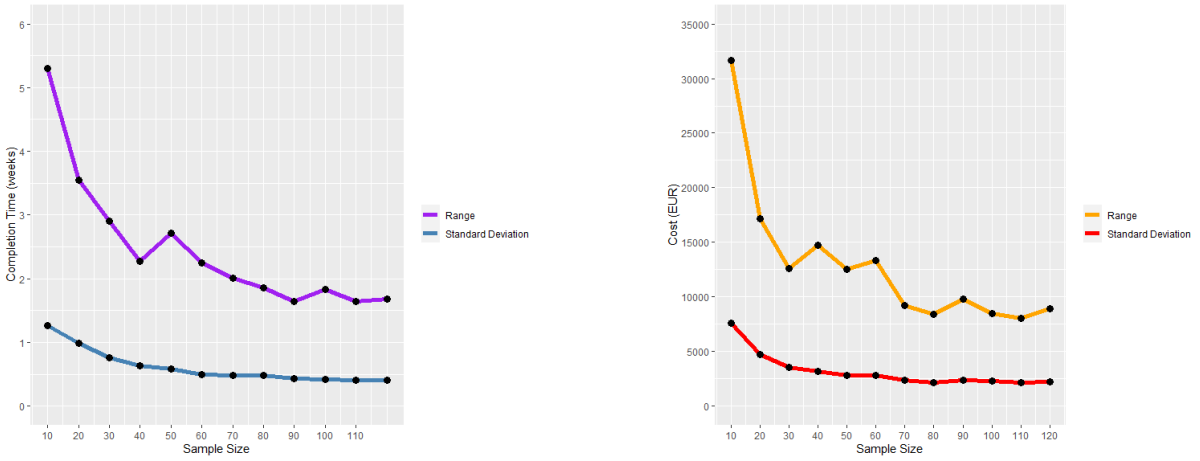

(a) *Completion Time* **(b)** *Cost*

Figure 6.22: *Spread Stability Test (BO29389)*

the expected value situated around 48; the second, in week 79, matches the recruitment of the last new patient: in this moment, the problem becomes deterministic and one of the extreme scenarios, that where a total 49 patients is enrolled, leads to a definitive ending inventory of 62 units. Let us take into account that, while the patient horizon had been set to 49 patients, had multiple patients shown up simultaneously it would have been possible that the cumulative recruited patients jumps from 48 to 50, for example, in which case a total 50 would have been enrolled, demanding the administration of $50 \times 15 = 750$ units instead of 735.

Scenario Reduction

We have followed the same procedure described in Section 6.1.3 in the context of trial WC 28325. Again, 10 different initializations with 40 scenarios each, have been used to build 7800 scenario pairs. The application of NNLS linear regression led to $c_1 = 0$, $c_2 = 11.6568$ and $c_3 = 18.5434$, with $R^2 = 0.992$.

As we did in Section 6.1.3 for CT BO29389, we have evaluated the relative errors in both completion time and total cost for a different number of clusters taken from samples of size 40.

Figure 6.25 shows that mean and standard deviation of relative errors in both objectives are well below 1% for $K \geq 8$. We will pick $K = 10$, which not only achieves a low relative error, but also is expected to match the best performance of PAM algorithm according to (Kaufman and Peter J. Rousseeuw 1987), where $S = 2K + 20$ is indicated as the most recommendable relationship between sample size and number of clusters.

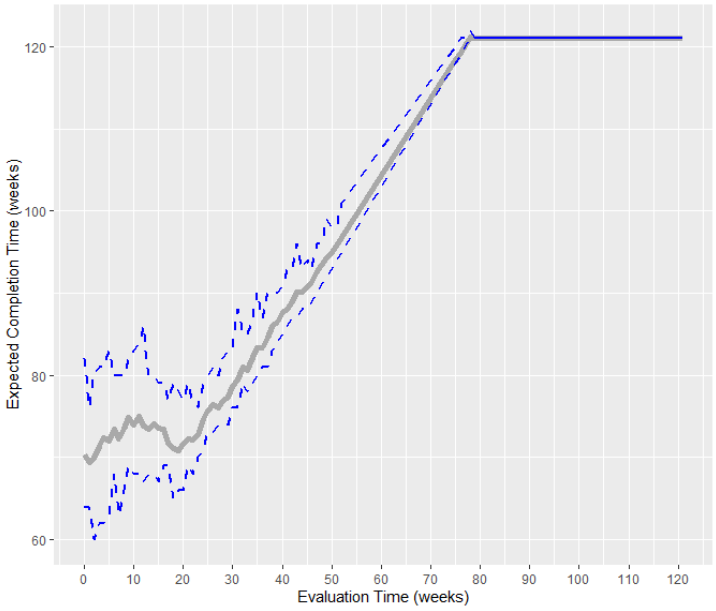

Figure 6.23: *Estimation made of Completion Time for trial BO29389*

Interestingly, the sequential optimization with sampling of size $S = 40$ and $K = 10$ clusters in each stage resulted in a single production run of 780 units, with total cost just 0.8% below the one estimated without clustering. It is also interesting that, despite using one less production run, this model estimates an ending inventory of 45 units, i.e. 17 units less than the model without clustering.

Results are shown in figures 6.26.

6.2.3. Comparison with other policies

A. Fleischhacker, Ninh, and Y. Zhao 2015

The benchmark model in (A. Fleischhacker, Ninh, and Y. Zhao 2015) was applied to this model in the Masters thesis by (Lerma and Lichtenstein 2020). We have also coded this implementation, finding a result of €338,514 and an ending inventory of 345 units, as a result of:

- Reorder points of 43 and 31 units in CD Germany and CD USA respectively
- Shipment quantities of 10 and 12 units in CD Germany and CD USA respectively
- Base level of 5 and 6 units in both sites served from CW; varying from a minimum of 4

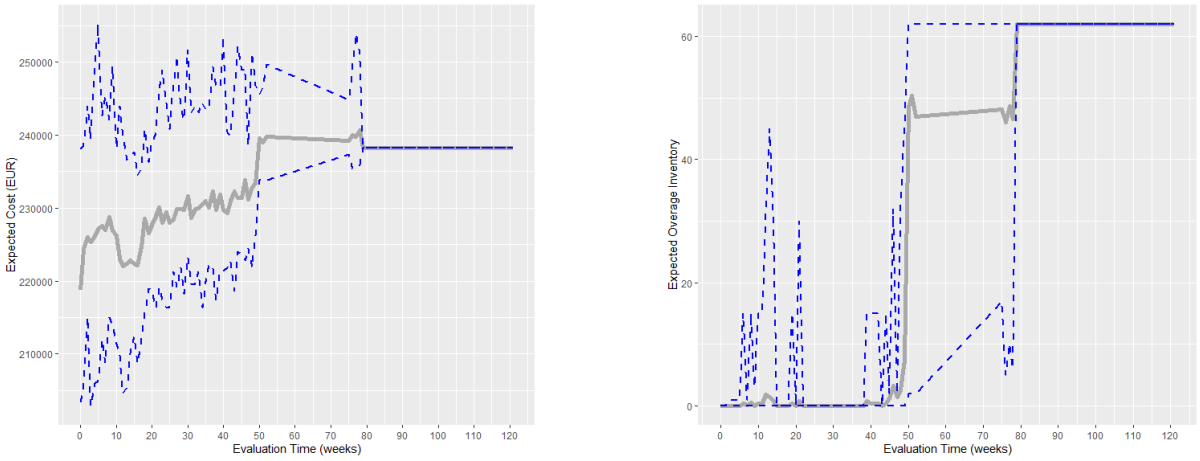

(a) *Cost* **(b)** *Overage Inventory*

Figure 6.24: *Evolution of the estimation made for cost and ending inventory in clinical trial BO29389*

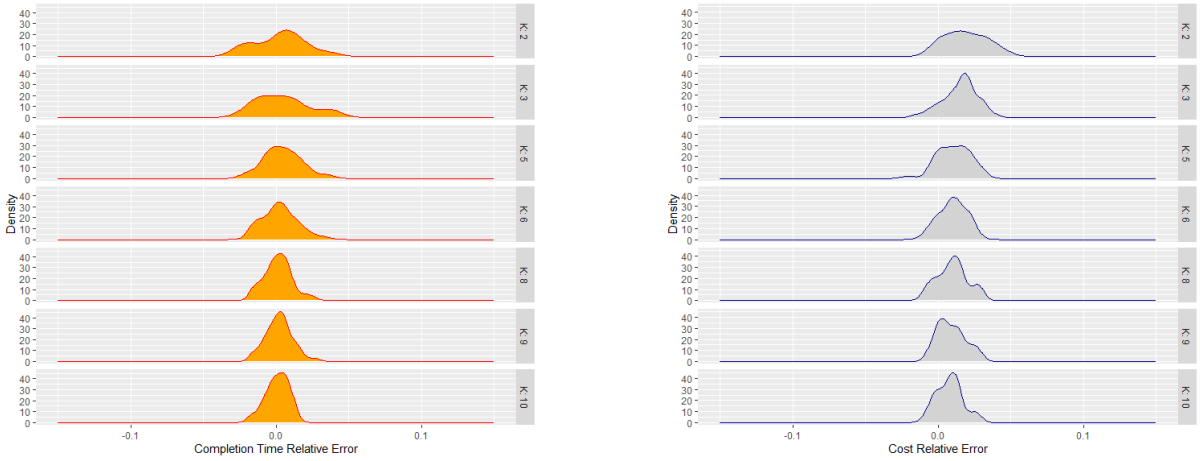

(a) *Completion Time* **(b)** *Overage Inventory*

Figure 6.25: *Density plots to determine an adequate number of clusters for CT BO29389*

and a maximum of 8 in sites served from CD Germany, and from 5 to 11 in sites located in USA.

which means that our model predicts a cost reduced by 30% and an ending inventory reduced by 85%. This huge improvement in terms of both cost and overage inventory is achieved by making distribution decisions tailored for the specific situation, instead of committing to a stationary policy. With this being said, we do not discard the possibility that a better parameters choice could result in a much better performing stationary policy able offer much greater simplicity at expenses of sacrificing a smaller portion of cost: in Chapter 7 we will present an extension to explore this possibility.

It is fair to point that A. Fleischhacker, Ninh, and Y. Zhao 2015 resulted in no patient having

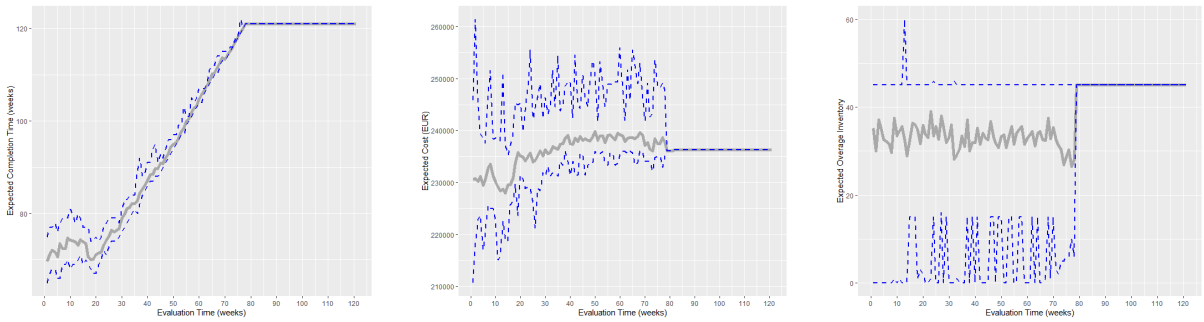

(a) *Completion Time* **(b)** *Cost* **(c)** *Overage Inventory*

Figure 6.26: *Evolution of estimations for key indicators of CT BO29389 using the S40K10 model*

to wait for inventory, while our mode resulted in 31 of the $15 \times 49 = 735$ units (4.2%) being administered late, which is within the 95% immediate service level required.

H. Zhao et al. 2019

We have applied this Stochastic Programming model, finding a a total cost of €238,302, which is virtually equal to the one obtained in our S40 model, and just 0.8 % more expensive than that of our S40K10 model. The overage inventory is just 23 units, which is significantly less than the one resulting from our models (45 units and 62 units in S40K10 and S40 models respectively), although it comes at a cost of a maximum waiting time that can be as long as 5 weeks, compared to the 16 days longest wait that resulted from our models. Figure 6.27 summarizes these results, along with models S40, S40K10 and the one from A. Fleischhacker, Ninh, and Y. Zhao 2015.

6.2.4. A NOTE ON THE CONSEQUENCES OF USING WEEKLY STAGES

As explained in Section 6.2.1, we have used weeks as stages in this case study, which has forced us to make some approximations in the lead times (Tables 6.6 and 6.7), which might result in actual results deviating from the predicted ones.

We have indeed run a simulation where we consider that decisions resulting from the model take place weekly (the first day of each week), and that demand occurs daily. Also in this model, we have considered the more realistic lead times referred as ALT in tables 6.6 and 6.7), instead of the approximated ones referred as MLT that we used in our optimization model.

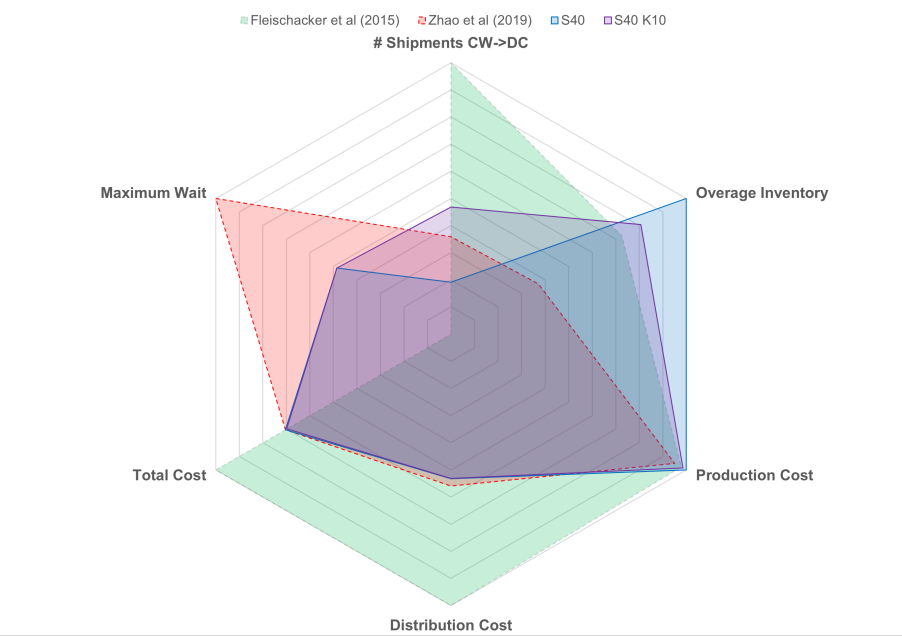

Figure 6.27: *BO29389: Comparison of our models against 2 benchmarks*

This simulation has shown that 2 patients would wait as much as 16 days for one of the doses, which slightly exceeds the maximum of 2 weeks allowed. This excess of 2 days is was due to two different simplifications:

- In one of the cases, it was the difference of 2 days between ALT and MLT (9 days vs 1 week) in the corresponding site (Site 13, in USA)
- In the other case, the difference of 2 days between ALT and MLT (23 days vs 3 weeks) in the corresponding country depot (CD Gernmany).

This means that the requirement of 2 weeks maximum waiting time has not been strictly satisfied and, if this is a hard constraint, it is advisable to either use daily stages or reduce the maximum waiting time to just 1 week.

6.3. BP30002: A phase I CT for glaucoma treatment

6.3.1. Data, Assumptions and Modeling Considerations

This clinical trial enrolled only 45 patients in USA and Singapore between December 2015 and June 2016 (Fig. 6.28). It is a treatment for glaucoma, consisting of a single dose of a quite expensive drug ($v_p = 1,500$). There is also a high fixed production cost ($f_p = 10,000$), and shipments costs from CW to CDs. Lead times from CDs to enrolment sites are 1 day in all cases.

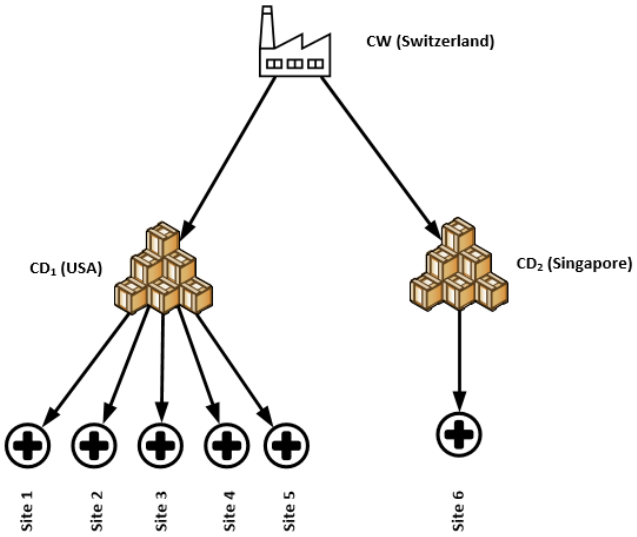

Figure 6.28: *Clinical Trial BP30002*

Data related to country depots is shown in Table 6.8, and estimated arrival rates are in Table 6.9.

CD	i	Fixed Shipment Cost	Variable Shipment Cost	Lead Time (days)
USA	1	€5,000	€500	10
Singapore	2	€2,000	€200	3

Table 6.8: *Country depots data for CT BP30002*

j	i	Country	λ(patients/day)
1	1	USA	0.0904
2	1	USA	0.0266
3	1	USA	0.0904
4	1	USA	0.0106
5	1	USA	0.0053
6	2	Singapore	0.016

Table 6.9: *Enrolment sites data for CT BP30002*

We will initially assume a very large shelf life as for neglecting perishability. Later, we will asses the impact of a shorter shelf life in the solution.

We will update information and re-optimize on a daily basis, so, in the context of this case study, when we speak about stages, we mean days.

While, in general, a stability test is recommended to determine an adequate sample size,

we will in this case assume that, just like in sections 6.1 and 6.2, a sample size of 40 is sufficiently good. Likewise, we will use $K = 10$ clusters to perform sequential optimization, basing this decision on:

i. 10 clusters has been more than enough in previous cases

ii. As shown in (Kaufman and Peter J. Rousseeuw 1987), the PAM algorithm runs at its best when $K = \frac{S-20}{2}$; therefore $S = 40$ leads to $K = 10$.

6.3.2. Numerical Results

Dissimilarity Metric

Following the procedure presented in Section 5.2.2 and used later on clinical trials WC28325 and BO29389, we have determined that the adequate coefficients for distance function (Eq. 5.2) are $c_1 = 0.025311$, $c_2 = -0.998159$ and $c_3 = 1.559996$, with all three being statistically significant at 10^{-6}.

Sequential Optimization and comparison with other policies

The one-dose clinical trial took 188 days to complete. Our model recommended a single production of 55 units, i.e. just 22% more than the patient horizon. The long lead time (10 days) from CW to USA, combined with the very low lead time from CDs to enrolment sites (1 day in all cases) resulted in most inventory being initially located in the country depots (45 in USA and 2 in Singapore), and a total just of 2 shipments from CW to country depots (excluding the shipments made for initial inventory allocation in such CDs), both of them to Singapore. This resulted in a total cost of €134,000 and an overage inventory of 10 units, as is shown in Figure 6.29.

In terms of total cost, these results are just slightly better than those that would result from applying the model proposed by A. Fleischhacker, Ninh, and Y. Zhao 2015 (€139,500) and H. Zhao et al. 2019 (€135,732), as can be seen in Figure 6.30. The cost savings with respect to (A. Fleischhacker, Ninh, and Y. Zhao 2015) is roughly 4%. This is due to the fact that this particular clinical trial satisfies much better the rather restrictive assumptions made by the authors, in particular the one related to the number of doses per patient being strictly equal to 1. However, in terms of overage inventory, our model does significantly better, with a

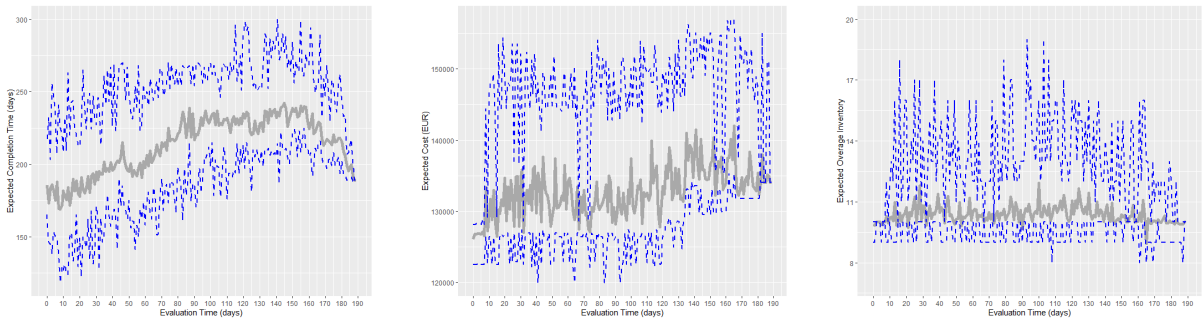

(a) *Completion Time* **(b)** *Cost* **(c)** *Overage Inventory*

Figure 6.29: *Evolution of estimations for key indicators of CT BP30002 using the S40K10 model*

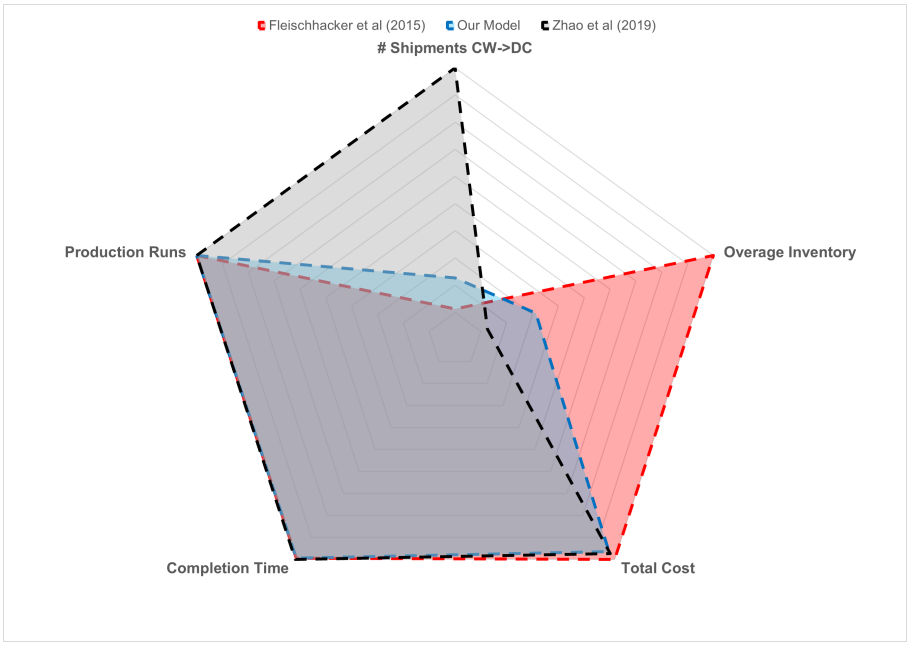

Figure 6.30: *BP30002: Comparison against 2 benchmark models*

reduction of about 70% with respect to the benchmark model (10 units instead of 32 units). With respect to (H. Zhao et al. 2019), our cost saving is just below 1%, due to the flexibility of both models in terms of not committing to a stationary policy. However, (H. Zhao et al. 2019) achieves an overage inventory of just 4 units, which is a considerable improvement with respect to the 10 units achieved with our model.

Consideration of Perishability

To illustrate the case of a limited shelf life and how the model captures this more realistic situation, let us consider a shelf life of 150 days. Since, in the absence of this limitation, 188 days were necessary to complete the trial, it is expected that one additional production run will be necessary. Indeed, the first stage solution indicates:

- An initial production run of 30 units (28 located in the USA CD and 2 in CD Singapore)
- According to the sampling we did, depending on the scenario, a second production run (between 15 and 18 units) would be necessary somewhere between day 17 and day 88.

Let us observe that, in the scenario where second production run takes place as early as day 17, the last patient enrols on day 165, so that remaining shelf life is still 2 days by the time that the last patient receives treatment.

Similarly, the second stage optimization (production and distribution decisions to be made at $t = 1$) considers a variety of scenarios with a second production run (between 15 and 25 units) to take place somewhere between days 8 and 93. The sequential optimization procedure results in a second production run finally happening on day 86, for 20 units, 18 of which are immediately pushed downstream to CD USA, where the remaining inventory from initial production run was 2 units (with a remaining shelf life of 64 days). This makes sense if we take into consideration that the expected number of arrivals to American sites during the 10 days lead time is $(10 days) \times 0.2234 patients/day = 2.23 patients$.

As we can see in Figure 6.31, there is a third production run on day 174 (7 units). Let us note that only 14 days after this third production run, the clinical trial finished with an ending inventory of 12 units (5 units from the second production run, with remaining shelf life of 48 days, and 7 units from the third and last production run, with remaining shelf life of 136 days), i.e. 5 units more than this last production batch. This production decision was indeed due to an unfortunate scenario sampling on stage 173. Figure 6.32 shows how the total cost of €160,800 is distributed for the 3 production batches, where we can observe that the second production batch carries 74% as much cost as the first one: they are both necessary and complementary, due to the limited shelf life. Cost associated to the third production batch, on the other hand, could have been avoided with a different sampling.

It is worth to mention that, units from the second production batch have been shipped from CD USA to enrolment sites in the USA only 7 days after the last unit from the first production batch had been pushed downstream. This confirms that, as we explained on Section 4.5.3, there is no need of enforcing FIFO.

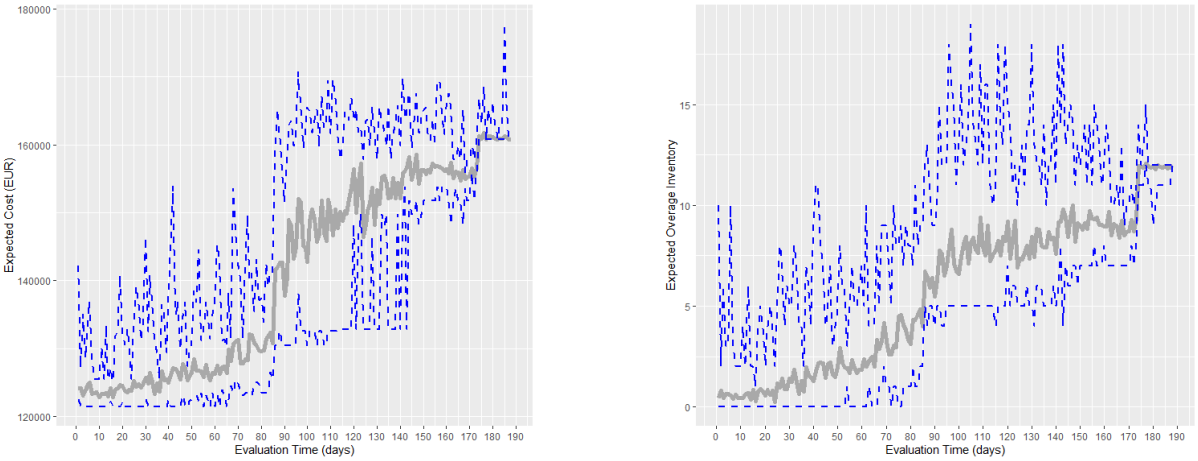

(a) *Cost* **(b)** *Overage Inventory*

Figure 6.31: *Evolution of estimation of cost and overage inventory for CT BP30002 with shelf life of 150 days*

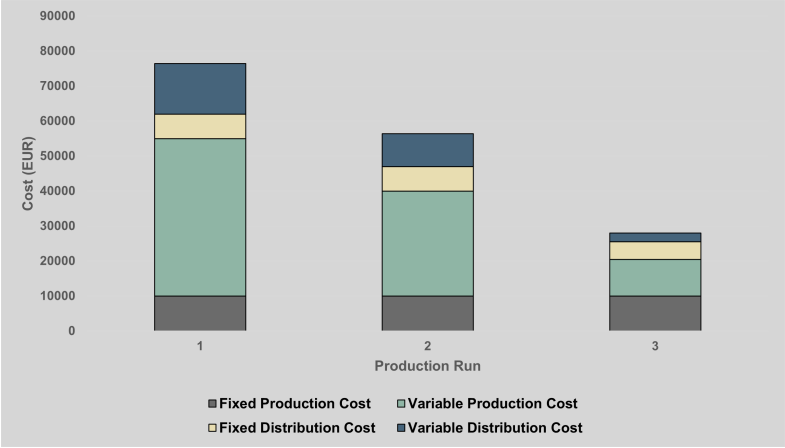

Figure 6.32: *Distribution of cost for units produced in each production batch in clinical trial BP30002*

Chapter 7

Extensions to the model

Although the model presented in Chapter 4 and techniques described in Chapter 5 answer the set of research questions that we listed in 1.4, we present here two extensions that illustrate how the model can be adapted to:

1. **Simpler policy enforcement:** Under certain circumstances (low drug cost and highly stable demand, for example), a simple inventory management policy might be preferred over an optimal one. This demands the addition of constraints to enforce a specific policy to be followed, making the model more complex, but the implementation much simpler.

2. **Multi-product CT**: We have noted that, while most CTs intend to test a specific new drug, sometimes the drug to be administrated is dependent on the genetic profile, which is only revealed after enrolment. This demands the need of managing inventory for more than one product simultaneously.

7.1. Enforcement of a simpler policy (Periodic Review s-S)

The output of the model shown in the previous section is a decision tree with a considerable amount of ramifications, and its application might be quite cumbersome. In this section, we aim to explore the trade-off between accuracy and complexity, by enforcing the solution to follow a policy of easy implementation. Although this will push the solution away from optimality, there exists the possibility that the huge gain in simplicity justifies the sacrifice of a portion of efficiency.

Variables

In addition to the variables defined in 4.4.3, we introduce:

- S_{CW}, S_i^{CD} and S_j^h: Target inventory level at CW, country depot i, and enrolment site j.
- s_{CW}, s_i^{CD} and s_j^h: Reorder point at CW, country depot i, and enrolment site j.
- $IP_{CW,t}^{(s)}$, $IP_{i,t}^{\text{CD(s)}}$ and $IP_{j,t}^{\text{h(s)}}$: Inventory position at the beginning of period t, under scenario s, on CW, country depot i and enrolment site j.
- $O_{i,t}^{\text{CD(s)}}$ and $O_{j,t}^{\text{h(s)}}$: Order quantity placed on period t, under scenario s, at country depot i and enrolment site j.
- $\chi_{CW,t}^{(s)}$, $\chi_{i,t}^{\text{CD(s)}}$ and $\chi_{j,t}^{(s)}$: 1 if inventory position is below the reorder point level (0 otherwise) at period t, under scenario s, at CW, country depot i and enrolment site j respectively.
- $\epsilon_{i,t}^{\text{CD(s)}}$ and $\epsilon_{j,t}^{\text{h(s)}}$: 1 if the delivery quantity is limited by the order quantity, 0 if it is limited by the inventory level upstream.

Constraints

In addition to the constraints presented in 4.5.2, we introduce:

$$IP_{j,t}^{(s)} = \sum_{\rho=0}^{\mu} IL_{j,t-1,\rho}^{(s)} + \sum_{t'=1}^{t-1} O_{j,t'}^{(s)} - \sum_{t'=1}^{t-1} \sum_{\rho=0}^{\mu} Q_{j,t',\rho}^{(s)} - \sum_{k=1}^{m} \sum_{r=2}^{\delta} D_{j,t,k,r}^{(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S} \quad (7.1)$$

$$O_{j,t}^{(s)} \leq S_j - IP_{j,t}^{(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \quad (7.2)$$

$$O_{j,t}^{(s)} \geq S_j - IP_{j,t}^{(s)} - M \cdot (1 - \chi_{j,t}^{(s)}) \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \quad (7.3)$$

$$IP_{j,t}^{(s)} \geq s_j - M \cdot \chi_{j,t}^{(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \quad (7.4)$$

$$O_{j,t}^{(s)} \leq M \cdot \chi_{j,t}^{(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \quad (7.5)$$

$$IP_{j,t}^{(s)} \leq s_j - 1 - M \cdot (1 - \chi_{j,t}^{(s)}) \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.6}$$

$$\sum_{\rho=0}^{\mu} Q_{j,t,\rho}^{(s)} \leq O_{j,t}^{(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.7}$$

$$\sum_{j \in \mathcal{J}_i} \sum_{\rho=0}^{\mu} Q_{j,t,\rho}^{(s)} \geq \sum_{\rho=0}^{\mu} IL_{i,t,\rho}^{(s)} - M \cdot \epsilon_{i,t}^{(s)} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.8}$$

$$\sum_{j \in \mathcal{J}_i} \sum_{\rho=0}^{\mu} Q_{j,t,\rho}^{(s)} \geq \sum_{j \in \mathcal{J}_i} O_{j,t}^{(s)} - M \cdot (1 - \epsilon_{i,t}^{(s)}) \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.9}$$

Eq. 7.1 provides a definition for inventory position at the enrolment site. Eqs. 7.2, 7.3 and 7.4 enforce that the necessary quantity to reach the target level is ordered in case of inventory position being below the reorder point. Eqs. 7.5 and 7.6 grant that no order is placed if inventory level is at or above the rorder point. Eqs. 7.7, 7.8 and 7.9 define the delivered quantity in such a way that it matches the maximum between the ordered quantity and the upstream inventory available.

Similarly, at country depot level:

$$IP_{i,t}^{(s)} = \sum_{\rho=0}^{\mu} IL_{i,t-1,\rho}^{(s)} + \sum_{t'=1}^{t-1} O_{i,t'}^{(s)} - \sum_{t'=1}^{t-1} \sum_{\rho=0}^{\mu} Q_{i,t',\rho}^{(s)} - \sum_{t'=1}^{t-1} \sum_{ij \in \mathcal{J}_i} O_{ij,t'}^{(s)} + \sum_{\rho=0}^{\mu} \sum_{t'=1}^{t-1} \sum_{ij \in \mathcal{J}_i} Q_{ij,t',\rho}^{(s)} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.10}$$

$$O_{i,t}^{(s)} \leq S_i - IP_{i,t}^{(s)} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.11}$$

$$O_{i,t}^{(s)} \leq M \cdot \chi_{i,t}^{(s)} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.12}$$

$$O_{i,t}^{(s)} \geq S_i - IP_{i,t}^{(s)} - M \cdot (1 - \chi_{i,t}^{(s)}) \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.13}$$

$$IP_{i,t}^{(s)} \leq s_i - 1 - M \cdot (1 - \chi_{i,t}^{(s)}) \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.14}$$

$$IP_{i,t}^{(s)} \geq s_i - M \cdot \chi_{i,t}^{(s)} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.15}$$

$$\sum_{\rho=0}^{\mu} Q_{i,t,\rho}^{(s)} \leq O_{i,t}^{(s)} \quad \forall i \in \mathcal{C}, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.16}$$

$$\sum_{i \in \mathcal{C}} \sum_{\rho=0}^{\mu} Q_{i,t,\rho}^{(s)} + \sum_{j \in \mathcal{J}_0} \sum_{\rho=0}^{\mu} Q_{j,t,\rho}^{(s)} \geq \sum_{\rho=0}^{\mu} IL_{CW,t,\rho}^{(s)} - M \cdot \epsilon_{CW,t}^{(s)} \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.17}$$

$$\sum_{i \in \mathcal{C}} \sum_{\rho=0}^{\mu} Q_{i,t,\rho}^{(s)} + \sum_{j \in \mathcal{J}_0} \sum_{\rho=0}^{\mu} Q_{j,t,\rho}^{(s)} \geq \sum_{i \in \mathcal{C}} O_{i,t}^{(s)} + \sum_{j \in \mathcal{J}_0} O_{j,t}^{(s)} - M \cdot (1 - \epsilon_{CW,t}^{(s)}) \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.18}$$

And same at CW:

$$IP_{CW,t}^{(s)} = \sum_{\rho=0}^{\mu} IL_{CW,t-1,\rho}^{(s)} + \sum_{\rho=0}^{\mu} Q_{t,\rho}^{(s)} - \sum_{t'=1}^{t-1} \sum_{i \in \mathcal{C}} (O_{i,t}^{(s)} - \sum_{\rho=0}^{\mu} Q_{i,t,\rho}^{(s)}) - \sum_{t'=1}^{t-1} \sum_{0 j \in \mathcal{J}_0} (O_{j,t}^{(s)} - \sum_{\rho=0}^{\mu} Q_{j,t,\rho}^{(s)}) \quad \forall 1 \leq t \leq L, s \in \mathcal{S} \tag{7.19}$$

$$Q_{t,\rho}^{(s)} \leq S_{CW} - IP_{CW,t}^{(s)} \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.20}$$

$$Q_{t,\rho}^{(s)} \leq M \cdot \chi_{CW,t}^{(s)} \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.21}$$

$$Q_{t,\rho}^{(s)} \geq S_{CW} - IP_{CW,t}^{(s)} - M \cdot (1 - \chi_{CW,t}^{(s)}) \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.22}$$

$$Q_{t,\rho}^{(s)} \leq s_{CW} - 1 + M \cdot (1 - \chi_{CW,t}^{(s)}) \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.23}$$

$$Q_{t,\rho}^{(s)} \geq s_{CW} - M \cdot \chi_{CW,t}^{(s)} \quad \forall 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N} \tag{7.24}$$

Plus non-negativity, integrality and binarity constraints:

$$S_{CW} \geq 0 \tag{7.25}$$

$$S_i^{CD} \geq 0 \quad \forall i \in \mathcal{C} \tag{7.26}$$

$$S_j^{CD} \geq 0 \quad \forall j \in \mathcal{J} \tag{7.27}$$

$$s_{CW} \geq 0 \tag{7.28}$$

$$s_i^{CD} \geq 0 \quad \forall i \in \mathcal{C} \tag{7.29}$$

$$s_j^{CD} \geq 0 \quad \forall j \in \mathcal{J} \tag{7.30}$$

$$O_{i,t}^{\mathrm{CD(s)}} \geq 0 \quad \forall i \in \mathcal{C}, \forall 1 \leq t \leq L \tag{7.31}$$

$$O_{j,t}^{\mathrm{CD(s)}} \geq 0 \quad \forall j \in \mathcal{J}, \forall 1 \leq t \leq L \tag{7.32}$$

$$\chi_{CW,t}^{(s)} \in \{0,1\} \quad \forall 1 \leq t \leq L \tag{7.33}$$

$$\chi_{i,t}^{\mathrm{CD(s)}} \in \{0,1\} \quad \forall i \in \mathcal{C}, \forall 1 \leq t \leq L \tag{7.34}$$

$$\chi_{j,t}^{\mathrm{h(s)}} \in \{0,1\} \quad \forall j \in \mathcal{J}, \forall 1 \leq t \leq L \tag{7.35}$$

$$\epsilon_{i,t}^{\text{CD(s)}} \in \{0,1\} \quad \forall i \in \mathcal{C}, \forall 1 \leq t \leq L \tag{7.36}$$

$$\epsilon_{j,t}^{\text{h(s)}} \in \{0,1\} \quad \forall j \in \mathcal{J}, \forall 1 \leq t \leq L \tag{7.37}$$

The application of this extended model is expected to provide reorder point and target level inventory for each node of the network, resulting in a much simpler decision rule to manage inventory, although at expenses of a higher cost and/or completion time.

7.2. Multi-product CT

In addition to the assumptions stated in 4.2, we assume that patients are tested upon enrolment, and only then it is revealed which drug type each patient will receive subject to her genetic profile. Let us assume that there is a fixed known probability γ_u for each one of the total ν genetic profiles Γ.

$$\Pr[\Gamma = \Gamma_u] = \gamma_u \quad \forall 1 \leq u \leq \nu \tag{7.38}$$

Let us denote by $\Psi_{u,j,t}$ the number of new patients with genetic profile u that arrive to site j on period t. The natural extension of the stochastic characterization explained in Section 4.3 results in the lower bound for the number of scenarios being:

$$S = \binom{H + NL\nu - 1}{H} \tag{7.39}$$

7.2.1. Changes to the model described in Chapter 4

A few variables will remain exactly the same as they were in the single-product formulation: $\alpha^{(s)}$, $\beta^{(s)}$, $\phi_t^{(s)}$ and $\psi_t^{(s)}$.

Most variables, however, are modified with the addition of an index u. For example:

- $Q_{i,t,\rho}^{\text{CD(s)}} \rightarrow Q_{i,t,\rho,u}^{\text{CD(s)}}$: Amount of drug u sent from CW to country depot i on day t, under scenario s.

- $D^{(s)}_{j,t,k,r} \rightarrow D^{(s)}_{j,t,k,r,u}$: Number of people who will come to site j on day t, for r^{th} consecutive day to ask for their k^{th} dose of drug u, under scenario s.

Remaining variables defined in 4.4.3 are also modified with the addition of index u and admit an equally intuitive interpretation.

Objectives

Just like in the single product formulation, objectives prioritization remains the same: main objective is expected completion time minimization, secondary objective is minimization of expected overage, a third objective is expected total cost minimization, and a fourth objective is minimization of expected number of production runs. While objectives 1 and 4 remain unchanged, and objective 3 is a natural extension, the 2nd objective considers the possibility of joint-ordering savings by introducing a major fixed cost per shipment besides minor fixed costs per shipment of each type of drug. That is why we introduce ξ^h_j, ξ^{CD}_i and ξ_{CW} variables to indicate whether there exists on day t any delivery or production, regardless of the drug type.

$$
\begin{aligned}
\min C = \sum_{u=1}^{v} \{ & v_u \cdot [IL_{CW,u} + \sum_{i \in \mathcal{C}} IL^{CD}_{i,u} + \sum_{j \in \mathcal{J}_i} IL^{h}_{j,u} + \sum_{s \in \mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot Q^{(s)}_{prod,t,\rho,u}] \\
& + \sum_{i \in \mathcal{C}} \sum_{s \in \mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} (p_s \cdot v^{CD}_{i,u} \cdot Q^{\mathrm{CD(s)}}_{i,t,\rho,u}) + \sum_{j \in \mathcal{U}} \sum_{s \in \mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} (p_s \cdot v^{h}_{j,u} \cdot Q^{\mathrm{h(s)}}_{j,t,\rho,u}) + \\
& + f_u \cdot (y^{CW}_{0,u} + \sum_{s \in S} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot y^{(s)}_{prod,t,\rho,u}) + \sum_{i \in \mathcal{C}} [f_{i,u} \cdot (y^{CD}_{i,0,u} + \sum_{s \in \mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot y^{\mathrm{CD(s)}}_{i,t,\rho,u})] \\
& + \sum_{j \in \mathcal{U}} [f^{h}_{j,u} \cdot (y^{h}_{j,0,u} + \sum_{s \in \mathcal{S}} \sum_{t=1}^{L} \sum_{\rho=1}^{\mu} p_s \cdot y^{\mathrm{h(s)}}_{j,t,\rho,u})] + \\
& + \sum_{s \in \mathcal{S}} \sum_{t=0}^{L} \sum_{\rho=0}^{\mu} [p_s \cdot (h^{CW,u} \cdot IL^{\mathrm{CW(s)}}_{t,\rho,u} + \sum_{j \in \mathcal{U}} h^{h}_{j,u} \cdot IL^{\mathrm{h(s)}}_{j,t,\rho,u} + \sum_{i \in \mathcal{C}} h^{CD}_{i,u} \cdot IL^{\mathrm{CD(s)}}_{i,t,\rho,u})]\} + \\
& + \sum_{s \in \mathcal{S}} p_s \cdot \{ \sum_{j \in \mathcal{J}} F^{(h)}_{j} \cdot \xi^{h(s)}_{j,t} + \sum_{i \in \mathcal{C}} F^{(CD)}_{i} \cdot \xi^{CD(s)}_{i,t} + F_{CW} \cdot \xi^{(s)}_{CW,t} \}
\end{aligned}
\tag{7.40}
$$

Only the last line in 7.40 is an addition to the simple addition of u index to the Cost Objective function defined in 4.5.1. In this last term added, we use $F^{(h)}_j$, $F^{(CD)}_i$ and F_{CW} as major fixed costs for delivery to site j, delivery to country depot i and production respectively.

Constraints

Most are natural extensions from the single product formulation, that require just the inclusion of index u. For example, (eq. 4.8) and (eq. 4.10) turn into (eq. 7.41) and (eq. 7.42):

$$\sum_{u=1}^{v}\sum_{z=1}^{t}\sum_{j\in\mathcal{U}}(E^{(s)}_{j,z,u}-\sum_{k=1}^{m}A^{(s)}_{j,z,k,u})\geq H\cdot(1-\phi^{(s)}_{t+1})\quad\forall s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.41}$$

$$E^{(s)}_{j,t,u}\leq d^{(s)}_{j,t,u}\cdot\phi^{(s)}_{t}\quad\forall j\in\mathcal{U},s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.42}$$

This simple adjustment is the rule rather than the exception, so we will elaborate on what must be changed in the single-product model in order to get a functional multi-product model either by addition, elimination or replacement (other than addition of an index) of constraints.

New constraints to be added for Multi-Product Model

The only additions to be made have the purpose of linking the new variables $\xi^{h(s)}_{j,t}$, $\xi^{CD(s)}_{i,t}$ and $\xi^{(s)}_{CW,t}$ to variables $y^{h(s)}_{j,t,u}$, $y^{CD(s)}_{i,t,u}$ and $y^{(s)}_{prod,t,u}$, in the sense that, if some delivery (or production) in general is taking place, then some shipment (or production) of a specific product is taking place.

$$\sum_{u=1}^{v}y^{h(s)}_{j,t,u}\leq\nu\cdot\xi^{h(s)}_{j,t}\quad\forall j\in\mathcal{U},s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.43}$$

$$\sum_{u=1}^{v}y^{h(s)}_{j,t,u}\geq\xi^{h(s)}_{j,t}\quad\forall j\in\mathcal{U},s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.44}$$

$$\sum_{u=1}^{v}y^{CD(s)}_{i,t,u}\leq\nu\cdot\xi^{CD(s)}_{i,t}\quad\forall i\in\mathcal{C},s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.45}$$

$$\sum_{u=1}^{v}y^{h(s)}_{i,t,u}\geq\xi^{CD(s)}_{i,t}\quad\forall i\in\mathcal{C},s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.46}$$

$$\sum_{u=1}^{v}y^{(s)}_{prod,t,u}\leq\nu\cdot\xi^{(s)}_{CW,t}\quad\forall s\in\mathcal{S},1\leq t\leq L,1\leq u\leq v \tag{7.47}$$

$$\sum_{u=1}^{\nu} y_{prod,t,u}^{(s)} \geq \xi_{CW,t}^{(s)} \quad \forall s \in \mathcal{S}, 1 \leq t \leq L, 1 \leq u \leq \nu \tag{7.48}$$

$$\xi_{j,t}^{h(s)} \in \{0,1\} \quad \forall j \in \mathcal{U}, s \in \mathcal{S}, 1 \leq t \leq L \tag{7.49}$$

$$\xi_{i,t}^{CD(s)} \in \{0,1\} \quad \forall i \in \mathcal{C}, s \in \mathcal{S}, 1 \leq t \leq L \tag{7.50}$$

$$\xi_{CW,t}^{(s)} \in \{0,1\} \quad \forall s \in \mathcal{S}, 1 \leq t \leq L \tag{7.51}$$

7.2.2. Periodic Review s-S and can-order policy

Combining sections 7.1 and 7.2, we can find a periodic review s-S policy in a multi-product setting. Since genetic profiles are far from being equally likely, it makes sense to define product-dependent values for reorder level s_u and target level S_u. With inventory for different products being managed independently, we would only need to add index u to the formulation described in 7.1. However, it is more interesting and challenging to embrace a policy where different products are managed in a not fully independent manner. In particular, we are interested on a can-order policy (Johansen and Melchiors 2003), in such a way that for any product u, it remains true that $IP_u < s_u \Rightarrow O_u = S_u - IP_u$; however, we could still place an order of product u even if $IP_u \geq s_u$, as long as two conditions are met:

1. $IP_u < s_{cu}$, where $s_{cu} > s_u$
2. $IP_{u'} < s_{u'}$, for some $u' \neq u$

This means that it is enough that one of the ν products satisfy the classical condition $IP < s_{u'}$ to raise every other reorder level from s_u to $s_{cu} > s_u$.

We then need to add a few decision variables.

- s_{site}^*, s_{CD}^* and s_{CW}^* : the joint reorder at site level point, country depot level and central warehouse respectively.
- $\eta_{j,t,u}^{h(s)}$, $\eta_{i,t,u}^{CD(s)}$ and $\eta_{CW,t,u}^{(s)}$: 1 if corresponding inventory position for product u is below $s_{j,u}^h$, $s_{i,u}^{CD}$ and $s_{CW,u}$ respectively, 0 otherwise.

- $\omega_{j,t}^{h(s)}$, $\omega_{i,t}^{CD(s)}$ and $\omega_{CW,t}^{(s)}$: 1 if, at time t for some product, inventory position is lower than corresponding specific reorder level, 0 otherwise.

Now we will formulate the corresponding constraints at enrolment site level:

We just need to add index u for (eq.7.1), (eq.7.2), (eq.7.3), (eq.7.5), (eq.7.7), (eq.7.8) and (eq.7.9).

For (eq.7.4) and (eq.7.4), however, we need a change and they will be substituted by (eq. 7.52), 7.53), 7.54) and 7.55):

$$IP_{j,t,u}^{\mathrm{h(s)}} \geq s_{j,u}^{h} - M \cdot \chi_{j,t,u}^{(\mathrm{s})} - M \cdot \omega_{j,t}^{h(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N}, 1 \leq u \leq \nu \tag{7.52}$$

$$IP_{j,t,u}^{\mathrm{h(s)}} \leq s_{j,u}^{h} - 1 + M \cdot (1 - \chi_{j,t,u}^{(\mathrm{s})}) - M \cdot \omega_{j,t}^{h(s)} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N}, 1 \leq u \leq \nu \tag{7.53}$$

$$IP_{j,t,u}^{\mathrm{h(s)}} \geq s_{site}^{*} - M \cdot \chi_{j,t,u}^{(\mathrm{s})} - M \cdot (1 - \omega_{j,t}^{h(s)}) \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N}, 1 \leq u \leq \nu \tag{7.54}$$

$$IP_{j,t,u}^{\mathrm{h(s)}} \leq s_{site}^{*} - 1 + M \cdot (1 - \chi_{j,t,u}^{(\mathrm{s})}) - M \cdot (1 - \omega_{j,t}^{h(s)}) \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N}, 1 \leq u \leq \nu \tag{7.55}$$

Please note that 7.52) and 7.53) are active if and only if $\omega_{j,t}^{h(s)} = 0$, while 7.54) and 7.55) are active if and only if $\omega_{j,t}^{h(s)} = 1$. This means that, if none of the ν products has its inventory position below its corresponding reorder level $s_{j,u}^{h}$, then no order is placed, but if some of the products has its inventory position below its reorder level, then it will be enough for others product to have inventory position below the joint reorder point in order to place an order.

We still need to create the constraints to track whether a product has its inventory level below its reorder level:

$$IP^{\mathrm{h(s)}}_{j,t,u} \geq s^{h}_{j,u} - M \cdot \eta^{(\mathrm{s})}_{j,t,u} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N}, 1 \leq u \leq \nu \tag{7.56}$$

$$IP^{\mathrm{h(s)}}_{j,t,u} \leq s^{h}_{j,u} - M \cdot (1 - \eta^{(\mathrm{s})}_{j,t,u}) \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S}, t = a \cdot R, a \in \mathbb{N}, 1 \leq u \leq \nu \tag{7.57}$$

$$\sum_{u=1}^{\nu} \eta^{(\mathrm{s})}_{j,t,u} \geq \omega^{h(s)}_{j,t} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.58}$$

$$\sum_{u=1}^{\nu} \eta^{(\mathrm{s})}_{j,t,u} \leq M \cdot \omega^{h(s)}_{j,t} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.59}$$

and finally:

$$\eta^{(\mathrm{s})}_{j,t,u} \in \{0,1\} \quad \forall j \in \mathcal{U}, 1 \leq u \leq \nu, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.60}$$

$$\omega^{(\mathrm{s})}_{j,t} \in \{0,1\} \quad \forall j \in \mathcal{U}, 1 \leq t \leq L, s \in \mathcal{S} \tag{7.61}$$

$$s^{*}_{site} \geq 0 \tag{7.62}$$

The application of these same concepts at country depot and CW level are, by analogy, very straightforward.

The solution of this model will return:

- reorder level for each product at each enrolment site and country depot, and for the CW: $s^{h}_{j,u}$, $s^{CD}_{i,u}$, $s_{CW,u}$.
- target level for each product at enrolment site and country depot, and for the CW: $S^{h}_{j,u}$, $S^{CD}_{i,u}$, $S_{CW,u}$
- joint reorder level at each enrolment site and country depot, and the CW: s^{*}_{site}, s^{*}_{CD}, s^{*}_{CW}

Chapter 8

Conclusion

This research was approached with 3 main goals, being the first one the production of a sufficiently general optimization model to describe the material flows in Clinical Trials in a realistic way, providing a systematic method to make the best possible production and distribution decisions. In particular, it was in our interest to allow multiple production runs and the possibility of patient dropouts linked to the quality of the supply chain decisions. Such model was presented in Chapter 4 and its effectiveness was successfully tested in Chapter 6 thru applying it to 3 different clinical trials, covering a range of geographies, arrival rates, time aggregation, shelf lives and number of doses.

The second objective was to contribute to deal with the tractability issues that usually discourage from picking this modeling approach despite how conceptually adequate it is and how efficiently it deals with complex decision-making problems when the number of scenarios is not too big. This objective was satisfied in Chapter 5 by introducing a novel dissimilarity metric, suitable for any Multi-Stage Multi-Echelon Stochastic Programming problem, and picking a convenient clustering method among the existing ones in the unsupervised machine learning literature. The effectiveness of this contribution was also tested in Chapter 6, where we also found considerable regularity in the sample size required to get stable solutions.

The third objective was to set the basis to explore the trade-off between accuracy and simplicity by creating a secondary model that enforces a simpler policy at expenses of some economical efficiency, although with potential benefits to come from the much simpler implementation. This objective was achieved in Chapter 7 with the extension of the main model to enforce a periodic review s-S policy both for a single-product or a multi-product clinical trial, and considering the possibility of joint-replenishment.

Highlighting some of the findings, not only the model outperformed the most popular model existing in the literature but, more importantly, it greatly extended the applicability by covering trials with multiple doses, different arrival processes, limited patience of the patients volunteering to receive treatment, etc. It is also worth to mention that the model was solved in reasonable time, suitable for managerial purposes.

While the case studies analyzed show that a good quality of the inputs leads to better decisions, it is remarkable how, in the event of well wrong input data, the model manages to produce acceptable decisions and accommodate to the situation reasonably well.

With regards to the original method proposed for scenario reduction, it not only reproduced results with extraordinary running time reduction and a minimal accuracy loss, but also the highly linear association between the proposed distance and the change in objective value between scenarios is a good indicator of conceptual coherence.

There is no knowledge up to now about models that approached clinical trial supply chains under stochastic settings which also considered multiple production runs, multiple doses and the possibility of patients dropping out if service level requirements are not met.

Among future research opportunities that we have identified, we believe that the following ones stand out:

- Incorporating the disposal cost to the objective function and the amount and moment to dispose as a decision variable.
- Consider a bayesian update of the arrival rates as a result of the *a priori* distribution and the updated information.
- Testing the extensions presented in Chapter 7.
- Elaborate other extensions enforcing different policies.
- Considering treatments where the number of doses is uncertain.
- Exploring different hierarchies of objectives and elaborating a Pareto-efficiency boundary.

Conclusión

Esta investigación se abordó con 3 objetivos principales, siendo el primero la producción de un modelo de optimización lo suficientemente general como para describir los flujos de material en los ensayos clínicos de una manera realista, proporcionando un método sistemático para tomar las mejores decisiones posibles de producción y distribución. En particular, nos interesaba permitir múltiples ejecuciones de producción y la posibilidad de abandonos de pacientes relacionados con la calidad de las decisiones de la cadena de suministro. Dicho modelo se presentó en el Capítulo 4 y su efectividad se probó con éxito en el Capítulo 6 aplicándolo a 3 ensayos clínicos diferentes, que abarcan una variedad de geografías, tasas de llegada, agregación de tiempo, vida útil y número de dosis.

El segundo objetivo fue contribuir a abordar los problemas de dimensionalidad que suelen desalentar la adopción de este enfoque de modelización a pesar de ser conceptualmente adecuado y de cómo maneja eficientemente problemas complejos de toma de decisiones cuando el número de escenarios no es demasiado grande. Este objetivo se cumplió en el Capítulo 5 al introducir una novedosa métrica de disimilitud, adecuada para cualquier problema de Programación Estocástica Multi-Etapa Multi-Eslabón, y al seleccionar un método de agrupación conveniente entre los existentes en la literatura de aprendizaje automático no supervisado. La efectividad de esta contribución también se probó en el Capítulo 6, donde también encontramos una regularidad considerable en el tamaño de la muestra necesario para obtener soluciones estables.

El tercer objetivo fue sentar las bases para explorar el equilibrio entre precisión y simplicidad al crear un modelo secundario que imponga una política más simple a expensas de cierta eficiencia económica, aunque con posibles beneficios derivados de una implementación mucho más sencilla. Este objetivo se logró en el Capítulo 7 con la extensión del modelo principal para imponer una política s-S de revisión periódica tanto para un ensayo clínico de un solo producto como para un ensayo clínico de múltiples productos, y considerando la posibilidad de reposición conjunta.

Destacando algunos de los hallazgos, no solo el modelo superó al modelo más popular

existente en la literatura, sino que, lo que es más importante, amplió considerablemente la aplicabilidad al abarcar ensayos con múltiples dosis, diferentes procesos de llegada, paciencia limitada de los pacientes voluntarios para recibir tratamiento, etc. También es importante mencionar que el modelo se resolvió en un tiempo razonable, adecuado para la gestión.

Si bien los estudios de caso analizados muestran que una buena calidad de los datos de entrada conduce a mejores decisiones, es notable cómo, en caso de que los datos de entrada sean incorrectos, el modelo logra producir decisiones aceptables y adaptarse razonablemente bien a la situación.

En cuanto al original método propuesto para la reducción de escenarios, no solo reprodujo resultados con gran reducción del tiempo de ejecución y mínima pérdida de precisión, sino que también la asociación altamente lineal entre la distancia propuesta y la variación absoluta de la función objetivo entre escenarios es un buen indicador de coherencia conceptual.

Hasta ahora, no se tiene conocimiento de modelos que hayan abordado las cadenas de suministro de ensayos clínicos bajo condiciones estocásticas que también consideraren múltiples ejecuciones de producción, múltiples dosis y la posibilidad de que los pacientes abandonen el tratamiento en caso de no cumplirse los requisitos de nivel de servicio.

Entre las oportunidades de investigación futura que hemos identificado, creemos que sobresalen las siguientes:

- Incorporar el costo de disposición a la función objetivo y la cantidad y momento de disposición como una variable de decisión.
- Considerar una actualización bayesiana de las tasas de llegada como resultado de la distribución *a priori* y la información actualizada.
- Probar las extensiones presentadas en el Capítulo 7.
- Elaborar otras extensiones que impongan diferentes políticas.
- Considerar tratamientos donde el número de dosis es incierto.
- Explorar diferentes jerarquías de objetivos y elaborar una frontera de eficiencia de Pareto.

Bibliography

Aharon, Ben-Tal, Golany Boaz, and Shtern Shimrit (2009). "Robust multi-echelon multi-period inventory control". In: *European Journal of Operational Research* 199.3, pp. 922–935. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2009.01.058. URL: https://www.sciencedirect.com/science/article/pii/S0377221709002112.

Arthur, David and Sergei Vassilvitskii (Jan. 2007). "K-Means++: The Advantages of Careful Seeding". In: vol. 8, pp. 1027–1035. DOI: 10.1145/1283383.1283494.

Azizi, Vahid and Guiping Hu (2021). "A Multi-Stage Stochastic Programming Model for the Multi-Echelon Multi-Period Reverse Logistics Problem". In: *Sustainability* 13.24. ISSN: 2071-1050. DOI: 10.3390/su132413596. URL: https://www.mdpi.com/2071-1050/13/24/13596.

Bakker, M., J. Riezebos, and R.H. Teunter (Sept. 2012). "Review of inventory systems with deterioration since 2001". English. In: *European Journal of Operational Research* 221.2, pp. 275–284. DOI: 10.1016/j.ejor.2012.03.004.

Beraldi, Patrizia and Maria Elena Bruni (2014). "A clustering approach for scenario tree reduction: an application to a stochastic programming portfolio optimization problem". In: *TOP* 22, pp. 934–949. URL: https://api.semanticscholar.org/CorpusID:121037585.

Bertocchi, Marida, Vittorio Moriggia, and Jitka Dupačová (Feb. 2006). "Horizon and stages in applications of stochastic programming in finance". In: *Annals of Operations Research* 142, pp. 63–78. DOI: 10.1007/s10479-006-6161-3.

Birge, John R. and François Louveaux (1997). *Introduction to Stochastic Programming*. New York, NY, USA: Springer-Verlag.

Bock, Hans-Hermann (2007). "Clustering Methods: A History of k-Means Algorithms". In: *Selected Contributions in Data Analysis and Classification*. Ed. by Paula Brito et al.

Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 161–172. ISBN: 978-3-540-73560-1. DOI: 10.1007/978-3-540-73560-1_15. URL: https://doi.org/10.1007/978-3-540-73560-1_15.

Chan, Carri W. and Linda V. Green (2013). "Handbook of Healthcare Operations Management". In: Springer. Chap. Improving Access to Healthcare: Models of Adaptive Behavior.

Chen, Wei-An (2019). "Drug Supply Chain Optimization for Adaptive Clinical Trials". PhD thesis. Purdue University.

Chen, Ye, Linas Mockus, et al. (2012). "Simulation-optimization approach to clinical trial supply chain management with demand scenario forecast". In: *Computers Chemical Engineering* 40, pp. 82–96. ISSN: 0098-1354.

Chen, Ye, Joseph Pekny, and Gintaras Reklaitis (Aug. 2012). "Integrated Planning and Optimization of Clinical Trial Supply Chain System with Risk Pooling". In: *Industrial Engineering Chemistry Research* 52, pp. 152–165. DOI: 10.1021/ie300823b.

Clariness (2023). *Overcoming clinical trial delays and accelerating clinical trial set-up and enrollment*. Accessed: November 26th 2023. URL: https://www.clinicaltrialsarena.com/downloads/whitepapers/clinical-trial-patient-recruitment/overcoming-clinical-trial-delays/#clariness.

Clemento, Anthony (1999). "New and Integrated Approaches to Successful Accelerated Drug Development". In: *Drug Information Journal* 33, pp. 699–710.

Colvin, Matthew and Christos T. Maravelias (2010). "Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming". In: *European Journal of Operational Research* 203.1, pp. 205–215. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2009.07.022. URL: https://www.sciencedirect.com/science/article/pii/S0377221709005244.

D'Agostino, Ralph and Michael Stephens (1986). *Goodness of fit techniques*. Marcel Dekker.

Deza, M.M. and E. Deza (2009). *Encyclopedia of Distances*. Encyclopedia of Distances. Springer Berlin Heidelberg. ISBN: 9783642002342. URL: https://books.google.es/books?id=LXEezzccwcoC.

Dillon, Mary, Fabricio Oliveira, and Babak Abbasi (2017). "A two-stage stochastic programming model for inventory management in the blood supply chain". In: *International Journal of Production Economics* 187, pp. 27–41. ISSN: 0925-5273.

Dupačová, Jitka, N. Groewe-Kuska, and Werner Roemisch (Jan. 2003). "Scenario Reduction in Stochastic Programming: An Approach Using Probability Metrics". In: *Mathematical Programming* 95, pp. 493–.

Fahimnia, Behnam and Armin Jabbarzadeh (2016). "Marrying supply chain sustainability and resilience: A match made in heaven". In: *Transportation Research Part E: Logistics and Transportation Review* 91, pp. 306–324. ISSN: 1366-5545. DOI: https://doi.org/10.1016/j.tre.2016.02.007. URL: https://www.sciencedirect.com/science/article/pii/S1366554516000296.

Fattahi, Mohammad and Kannan Govindan (2018). "A multi-stage stochastic program for the sustainable design of biofuel supply chain networks under biomass supply uncertainty and disruption risk: A real-life case study". In: *Transportation Research Part E: Logistics and Transportation Review* 118, pp. 534–567. ISSN: 1366-5545. DOI: https://doi.org/10.1016/j.tre.2018.08.008. URL: https://www.sciencedirect.com/science/article/pii/S1366554518301881.

Fattahi, Mohammad, Kannan Govindan, and Esmaeil Keyvanshokooh (2018). "A multi-stage stochastic program for supply chain network redesign problem with price-dependent uncertain demands". In: *Computers Operations Research* 100, pp. 314–332. ISSN: 0305-0548. DOI: https://doi.org/10.1016/j.cor.2017.12.016. URL: https://www.sciencedirect.com/science/article/pii/S0305054817303167.

Fleischhacker, Adam, Anh Ninh, and Yao Zhao (2015). "Positioning Inventory in Clinical Trial Supply Chains". In: *Production and Operations Management* 24.6, pp. 991–1011.

Fleischhacker, Adam and Yao Zhao (2011). "Planning for demand failure: A dynamic lot size model for clinical trial supply chains". In: *European Journal of Operations Research* 211.3, pp. 496–506.

Fleischhacker, Adam J and Yao Zhao (2013). "Contract Development and Manufacturing Costs During Clinical Development of a New Drug". In.

Gangammanavar, Harsha and Suvrajeet Sen (2021). "Stochastic Dynamic Linear Programming: A Sequential Sampling Algorithm for Multistage Stochastic Linear Programming". In: *SIAM Journal on Optimization* 31.3, pp. 2111–2140. DOI: 10.1137/19M1290735. eprint: https://doi.org/10.1137/19M1290735. URL: https://doi.org/10.1137/19M1290735.

Goh, Joel and Evan Porteus (Mar. 2016). "Multi-Echelon Inventory Management Under Short-term Take-or-Pay Contracts". In: *Production and Operations Management* 25, n/a–n/a. DOI: 10.1111/poms.12557.

Haijema, René (2013). "A new class of stock-level dependent ordering policies for perishables with a short maximum shelf life". In: *International Journal of Production Economics* 143.2, pp. 434–439. ISSN: 0925-5273. DOI: https://doi.org/10.1016/j.ijpe.2011.05.021. URL: https://www.sciencedirect.com/science/article/pii/S0925527311002386.

Heitsch, H. and W. Römisch (2009). "Scenario tree reduction for multistage stochastic programs". In: *Computational Management Science* 6, pp. 117–133.

Heitsch, H., W. Römisch, and C. Strugarek (2006). "Stability of Multistage Stochastic Programs". In: *SIAM Journal on Optimization* 17.2, pp. 511–525.

Henrion, René and Werner Roemisch (Sept. 2018). "Problem-based optimal scenario generation and reduction in stochastic programming". In: *Mathematical Programming*. DOI: 10.1007/s10107-018-1337-6.

Hewitt, Mike, Janosch Ortmann, and Walter Rei (Aug. 2022). "Decision-based scenario clustering for decision-making under uncertainty". In: *Annals of Operations Research* 315, pp. 1–25. DOI: 10.1007/s10479-020-03843-x.

Horn, Susan Dadakis (1977). "Goodness-of-Fit Tests for Discrete Data: A Review and an Application to a Health Impairment Scale". In: *Biometrics* 33.1, pp. 237–247. ISSN: 0006341X, 15410420. URL: http://www.jstor.org/stable/2529319 (visited on 11/07/2023).

IM, Schou and C Marschner I (2015). "Methods for exploring treatment effect heterogeneity in subgroup analysis: an application to global clinical trials". In: *Pharm Stat* 14.1, pp. 44–55.

Johansen, S. G. and P. Melchiors (2003). "Can-Order Policy for the Periodic-Review Joint Replenishment Problem". In: *The Journal of the Operational Research Society* 54.3, pp. 283–290.

Kaufman, Leonard and Peter J Rousseeuw (2009). *Finding groups in data: an introduction to cluster analysis*. eng. Wiley Series in Probability and Statistics. WILEY. ISBN: 9780471735786.

— (1987). *Clustering by means of medoids*. Ed. by In: Dodge Y and editor. Amsterdam:

Keutchayan, Julien, Janosch Ortmann, and Walter Rei (2021). *Problem-Driven Scenario Clustering in Stochastic Optimization*. arXiv: 2106.11717 [math.OC].

Kleywegt, Anton J., Alexander Shapiro, and Tito Homem-de-Mello (2002). "The Sample Average Approximation Method for Stochastic Discrete Optimization". In: *SIAM Journal on Optimization* 12.2, pp. 479–502. DOI: 10.1137/S1052623499363220. eprint: https://doi.org/10.1137/S1052623499363220. URL: https://doi.org/10.1137/S1052623499363220.

Lawson, Charles L. and Richard J. Hanson (1995). *Solving Least Squares Problems*. Society for Industrial and Applied Mathematics. DOI: 10.1137/1.9781611971217. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611971217. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611971217.

Lefew, Michael, Anh Ninh, and Vladimir Anisimov (Sept. 2021). "End-to-End Drug Supply Management in Multicenter Trials". In: *Methodology and Computing in Applied Probability* 23.3, pp. 695–709. DOI: 10.1007/s11009-020-09776-. URL: https://ideas.repec.org/a/spr/metcap/v23y2021i3d10.1007_s11009-020-09776-z.html.

Lerma, A. and P. Lichtenstein (2020). *Inventory Optimization in Clinical Trial Supply Chain*.

Lima, Camilo, Susana Relvas, and Ana Barbosa-Póvoa (2018). "Stochastic programming approach for the optimal tactical planning of the downstream oil supply chain". In: *Computers Chemical Engineering* 108, pp. 314–336. ISSN: 0098-1354. DOI: https://doi.org/10.1016/j.compchemeng.2017.09.012. URL: https://www.sciencedirect.com/science/article/pii/S0098135417303174.

Linderoth, Jeff, Alexander Shapiro, and Stephen Wright (Feb. 2006). "The Empirical Behavior of Sampling Methods for Stochastic Programming". In: *Annals of Operations Research* 142, pp. 215–. DOI: 10.1007/s10479-006-6169-8.

Manopiniwes, Wapee and Takashi Irohara (2017). "Stochastic optimisation model for integrated decisions on relief supply chains: preparedness for disaster response". In: *International Journal of Production Research* 55.4, pp. 979–996. DOI: 10.1080/00207543.2016.1211340. eprint: https://doi.org/10.1080/00207543.2016.1211340. URL: https://doi.org/10.1080/00207543.2016.1211340.

Miller, Michael G. (1996). "Optimal Allocation of Resources to Clinical Trials". PhD thesis. Massachusetts Institute of Technology.

Mitchel, Jules T et al. (2022). "The Transformation of Clinical Trials". In: *Applied Clinical Trials* 1/2, pp. 16–20.

Nahmias, Steven (Aug. 1982). "Perishable Inventory Theory: A Review". In: *Operations research* 30, pp. 680–708. DOI: 10.1287/opre.30.4.680.

Newton, William (2023). *Sustainable supply chains: clinical trials in a new era of limited resources*. Accessed: November 26th 2023. URL: https://www.clinicaltrialsarena.com/features/sustainable-clinical-trial-supply-chains/?cf-view.

Ninh, Anh, Michael Lefew, and Vladimir Anisimov (Dec. 2019). "Clinical Trial Simulation: Modeling and Practical Considerations". In: pp. 118–132. DOI: 10.1109/WSC40007.2019.9004916.

Page, Daniel R. (2013). "Generalized Algorithm for Restricted Weak Composition Generation: Generation Algorithm for Second-Order Restricted Weak Compositions". In: *Journal of Mathematical Modelling and Algorithms in Operations Research* 12.4, pp. 345–372. ISSN: 2214-2487.

Parke, Tom (2011). "Simulating Clinical Trials". In: URL: https://api.semanticscholar.org/CorpusID:70939477.

Paulo, Helena et al. (Jan. 2017). "Designing Integrated Biorefineries Supply Chain: Combining Stochastic Programming Models with Scenario Reduction Methods". In: pp. 901–906. ISBN: 9780444639653. DOI: 10.1016/B978-0-444-63965-3.50152-5.

Pauls-Worm, K.G.J. (2016). "Inventory control for a perishable product with non-stationary demand". PhD thesis. Wageningen University.

Pereira, Luis Marcelo (2006). "Critical Considerations about Clinical Trials Simulation". In: *International Journal of Pharmaceutical Medicine* 20, pp. 1–15. URL: https://api.semanticscholar.org/CorpusID:20360118.

Peterson, Magnus et al. (2004). "Optimizing clinical trial supply requirements: simulation of computer-controlled supply chain management". In: *Clinical Trials* 1, pp. 399–412. URL: https://api.semanticscholar.org/CorpusID:22044466.

Römisch, Werner (2009). "Scenario Reduction Techniques in Stochastic Programming". In: *Stochastic Algorithms: Foundations and Applications*. Ed. by Osamu Watanabe and Thomas Zeugmann. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 1–14. ISBN: 978-3-642-04944-6.

Rubio-Herrero, Javier, Anh Ninh, and Michael Lefew (May 2023). "Improving the performance of supply chains in clinical trials with delays: an optimization approach to determining the number of recruitment sites". In: *Annals of Operations Research*, pp. 1–21. DOI: 10.1007/s10479-023-05382-7.

Ruszczyński, Andrzej and A Shapiro (Jan. 2003). *Stochastic Programming*, *Handbook in Operations Research and Management Science*.

Schubert, Erich and Lars Lenssen (2022). "Fast k-medoids Clustering in Rust and Python". In: *Journal of Open Source Software* 7.75, p. 4183. DOI: 10.21105/joss.04183. URL: https://doi.org/10.21105/joss.04183.

Schubert, Erich and Peter J. Rousseeuw (2021). "Fast and eager k-medoids clustering: O(k) runtime improvement of the PAM, CLARA, and CLARANS algorithms". In: *Information Systems* 101, p. 101804. ISSN: 0306-4379. DOI: https://doi.org/10.1016/j.is.2021.101804. URL: https://www.sciencedirect.com/science/article/pii/S0306437921000557.

Shapiro, A (2003). "Inference of statistical bounds for multistage stochastic programming problems". In: *Mathematical Methods of Operations Research* 3, pp. 57–68.

Shu, Jia, Zhengyi Li, and Liya Huang (2013). "Demand selection decisions for a multi-echelon inventory distribution system". eng. In: *The Journal of the Operational Research Society* 64.9, pp. 1307–1313. ISSN: 0160-5682.

Statista (2023). *Statista 2023*. Accessed: November 22th 2023. URL: https://www.statista.com/statistics/732997/number-of-registered-clinical-studies-worldwide/.

Thiers, Fabio Albuquerque (2006). "The Globalization of Clinical Drug Development". PhD thesis. Massachusetts Institute of Technology.

Tohidi, Hossein (2020). "Expert Systems For DecisionMaking inMultistage Healthcare Problems". PhD thesis. North Carolina State University.

Ushakov, Anton V. and Igor Vasilyev (2019). "A parallel heuristic for a k-medoids clustering problem with unfixed number of clusters". In: *2019 42nd International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pp. 1116–1120. DOI: 10.23919/MIPRO.2019.8756919.

Wang, Liangquan and Chaolin Yang (n.d.). "Robust multi-echelon inventory management with multiple suppliers". In: *Naval Research Logistics (NRL)* n/a.n/a (). DOI: https://doi.org/10.1002/nav.22147. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.22147. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.22147.

WHO (2023). *World Health Organization*. Accessed: November 13th 2023. URL: https://www.who.int/health-topics/clinical-trials#tab=tab_1.

Zhang, Yingying, Yi Chai, and Le Ma (2021). "Research on Multi-Echelon Inventory Optimization for Fresh Products in Supply Chains". eng. In: *Sustainability (Basel, Switzerland)* 13.11, pp. 6309–. ISSN: 2071-1050.

Zhao, Hui et al. (2019). "A multi-objective production planning problem with the consideration of time and cost in clinical trials". In: *Expert Systems with Applications* 124, pp. 25–38. ISSN: 0957-4174.

Zheng, Meimei et al. (2021). "A Study on the Optimal Inventory Allocation for Clinical Trial Supply Chains". In: *Applied Mathematical Modelling* 98, pp. 161–184. ISSN:

0307-904X. DOI: https://doi.org/10.1016/j.apm.2021.04.029. URL: https://www.sciencedirect.com/science/article/pii/S0307904X21002298.

Zhou, Wei-Qi, Long Chen, and Hui-Ming Ge (2013). "A multi-product multi-echelon inventory control model with joint replenishment strategy". In: *Applied Mathematical Modelling* 37.4, pp. 2039–2050. ISSN: 0307-904X. DOI: https://doi.org/10.1016/j.apm.2012.04.054. URL: https://www.sciencedirect.com/science/article/pii/S0307904X12002892.