Vrije Universiteit Amsterdam

Master's Thesis

Planning Under Uncertainty: Strategic Decision Making in Supply Chain Network Design

Author: Fergus Hathorn (2722167)

1st supervisor: Alessandro Zocca
Industry supervisor: Paul van Nierop (AIMMS)
2nd reader: Joost Berkhout

A thesis submitted in partial fulfillment of the requirements for the VU Master of Science degree in Business Analytics

July 24, 2023

Planning Under Uncertainty: Strategic Decision Making in Supply Chain Network Design

Author: Fergus Hathorn (2722167)

Vrije Universiteit Amsterdam
Faculty of Science
Business Analytics
De Boelelaan 1081a
1081 HV
Amsterdam
The Netherlands

AIMMS B.V.
Diakenhuisweg 29-35
2033 AP
Haarlem
The Netherlands

Executive Summary

The vulnerability of global supply chain networks has been highlighted in recent times, where disruptions have caused major shortages and exposed the lack of flexibility in these networks. AIMMS provides a supply chain network design product that assists organizations in optimally configuring their supply chains and addressing important business questions that can be answered using mathematical optimization. This study aims to investigate whether supply chain network design decision-making could be enhanced, in terms of robustness and adaptability, by introducing uncertainty to the design phase.

A stochastic programming approach was implemented to address uncertainty in product demand, transportation cost, and throughput in facilities. These sources of uncertainty cover both high-probability-low-impact risks (fluctuations in demand and transport costs) and low-probability-high-impact risks (disruptions to facilities). Scenarios were generated using multivariate probability distributions that capture the behavior of demand and transport cost, and a univariate probability distribution to model disruptive events. The stochastic program's first-stage decisions were focused on facility open/close decisions and allocating products to facilities. Geographical clustering of demand nodes was incorporated to reduce the problem size without significantly altering the supply chain network. To reduce the time to a solution, the Benders decomposition was applied.

When applying the methods to the case of a real distribution network, it was found that the deterministic solution was unable to produce a feasible solution for roughly one-third of the generated scenarios. This was primarily due to a lack of capacity to support throughput when facility disruptions occurred. The stochastic programming solution had higher costs and this was largely due to facility decisions, as transport costs were reduced and handling costs remained consistent between the problems. However, the stochastic solution allowed the network to maintain normal operating performance in all scenarios, even under severe realizations of uncertainty. It was also found that correlations within uncertain parameters had an influence on the objective value, *i.e.,* scenarios that recognized inter-dependencies had a higher variance in the objective. Geographical clustering significantly reduced the solving time while maintaining the problem structure, even for a small number of clusters. The Benders decomposition allowed for larger instances to be solved. However, for this case, investigations revealed that the first-stage decisions stabilized when using at least 30 scenarios, which indicates that solving large instances (>30 scenarios) may not yield results that are *more* insightful.

According to the findings of this study, addressing uncertainty in supply chain network design could provide decision-makers with more robust network configurations that can respond to market changes or catastrophic disruptions. Although incorporating these types of decisions in practice may necessitate a move to a new paradigm of strategic decision-making, the research suggests that tackling uncertainty in supply chain network design should be an important consideration for supply chain managers.

Preface

This thesis has been prepared in partial fulfillment of the Master's program in Business Analytics at the Vrije Universiteit Amsterdam. The objective of conducting a thesis is to demonstrate the student's ability to conduct independent research and to integrate previous academic courses from the degree program during this process.

The research problem for this thesis concerns supply chain network design and the concept of *planning under uncertainty.* Global supply chains have many sources of inherent uncertainty and the methods by which this uncertainty is currently being addressed makes generating robust and resilient supply chain networks a major challenge.

The thesis project has been hosted by *AIMMS B.V.*, a prescriptive analytics software development company headquartered in The Netherlands. AIMMS has a department that focuses on developing a supply chain optimization product, and it is within this department that the thesis project has been situated.

I would like to acknowledge the opportunity provided by AIMMS to conduct a thesis project with such strong support from the company. Namely, Paul van Nierop acted as industry supervisor, adding invaluable insights and guidance in the world of supply chain management, and committed extensive time and resources to ensure that I could conduct my research unhindered. I also wish to extend gratitude to all colleagues at AIMMS who supported the project in various ways.

From an academic perspective, I would like to acknowledge my supervisor, Alessandro Zocca, for the regular discussions, advice, and feedback during the thesis process. Lastly, I thank Joost Berkhout for his contribution as the second reader.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This purpose of this section is to introduce context of the thesis, provide clear description of the problem, aims, objectives, and scope of the investigation, as well as to outline the structure of the thesis. The thesis background is described in Section 1.1 and the problem is defined in Section 1.2. Section 1.3 gives the scope of the thesis before the aim and objectives are defined in Section 1.4. Finally, Section 1.5 provides an outline for the organization of the thesis.

1.1 Company and Product Background

AIMMS develops and deploys prescriptive analytics software that is used by businesses, consultants, and academics. Out-of-the-box products are offered for Supply Chain Network Design (SCND), Center of Gravity (CoG), Inventory Planning, Sales and Operations Planning (S&OP), and Demand Forecasting. Apart from these products, their development platform allows users to create and deploy bespoke applications. The development platform is widely used by academics and operations research modelers.

The SCND product is considered the flagship product and will be the focus of this study. SCND is used to determine, through mathematical optimization, the best location and optimal size of facilities in the supply chain, as well as a detailed breakdown of the optimal flow of products, including throughputs, transportation volume, lanes, and among various other insights. The user is guided through the process first by a data collection process that collects detailed information about the structure and mechanisms of the supply chain, finally prompting the user to populate a spreadsheet with the necessary data points to answer their business questions. The spreadsheet can be uploaded to a cloud platform, whereafter the data are extracted and a Mixed-Integer-Program (MIP) is configured. The user has the option to fine-tune the MIP by relaxing/applying various constraints, inflating/deflating demand, and adjusting any of the input data. The user is then expected to develop a base-case scenario that represents the current state of the supply chain, as well as scenarios that help with answering business questions and/or evaluating the performance of the supply chain under future scenarios. The scenarios can be compared in detail on the cloud platform.

1.2 Project Problem Statement

Recent world events (the COVID-19 pandemic, the Ukraine-Russia war, labor shortages, etc.) have exposed the fragility of supply chains. Disruptions can induce unexpected costs and reduction in service levels, therefore, many business operators are interested in (re)designing supply chains that are robust to disruptions. To address this, AIMMS provides users with the ability to model various scenarios to design a supply chain that is robust to realizations of uncertainty in the future. However, these scenarios are solved independently and are not always sufficient to adequately expose the performance of the supply chain under future uncertainty. This can lead to users being blindsided by external influences on their supply chain which could have been avoided in the design phase if uncertainty within the supply chain network was handled appropriately. Companies that do not build robust supply chains are exposed to a multitude of external factors that could severely affect their ability to meet customer demand. A more sophisticated method of dealing with uncertainty within the supply chain could result in reduced impact from supply chain crises.

1.3 Scope

The scope of this study extends from researching related work and methods to handling uncertainty in SCND, to developing and evaluating an appropriate extension of the existing mathematical model. The deployment or development of a user interface is not within the scope of the study. Apart from minor adjustments for incorporating uncertainty, performing validation on the pre-existing SCND mathematical model does not fall within the scope of this study; hence, it is largely assumed that the underlying program is operating as expected and designed. The proposed solution should highlight the challenges and considerations for implementing a commercial variant of the project, thereby it is not a requirement that the method is fully generalizable. Rather, the proposed solution should be applied to a business case, providing insight into the expected challenges and considerations for developing a commercially viable and generalizable variant of the project. An important part of this project is to deliver practical insights as to how uncertainty can be addressed in SCND, and to provide intuitive analysis regarding the benefit of addressing uncertainty in supply chains. It is not within the scope to present, analyze, or make alterations to the algorithms used to *solve* mathematical programs, rather the modeling approach is explored.

1.4 Aim and Objectives

This study aims to determine whether an out-of-the-box commercial SCND product that is supported by mathematical optimization in a deterministic environment could be enhanced by introducing uncertainty in some of the key quantities. The product is enhanced if the solution forms an SCN that is more robust compared to an SCN that was designed without considering uncertainty. Robustness has been widely understood to be the ability of a supply chain to maintain its normal operations when presented with disruptive scenarios [32].

Several key research objectives are needed to support the study, in particular:

i. Outline the benefits and drawbacks of handling uncertainty in an out-of-the-box product.

ii. Determine which methods of optimization under uncertainty should be considered for SCND through a literature survey, and which of the methods are most suitable for this study.

iii. Determine which methods for evaluating robustness should be considered for SCND through a literature survey; additionally, determine which method is most suitable for this study.

iv. Compare an optimal SCN under uncertainty (from iii) to one that does not consider uncertainty in terms of robustness (from iv).

As the study has a focus on commercial applicability for an out-of-the-box SCND product, several requirements must be considered for the study to contribute to the SCND product in a useful way, in particular:

1. Data requirements for the uncertain elements should be generic data that most users can access.

2. Configuring the model under uncertainty should not add significant workflow to the existing process.

3. The model solve time for commercial solvers should not exceed reasonable limitations.

4. The results of the model under uncertainty should be interpretable and actionable and should be comparable (in terms of network design and resilience) with scenarios generated by optimization without uncertainty.

1.5 Thesis Organization

The thesis is structured as follows. Chapter 2 provides background on preliminary knowledge required for the rest of the thesis, as well as setting a foundation for the study by drawing from related literature. Chapter 3 outlines the modeling approach (based on the conclusions drawn from Chapter 2) and provides a mathematical description of the methods employed. In Chapter 4, appropriate data from an operational supply chain network is used to evaluate the modeling approach, including an introduction to the data, computational experiments, analysis of the resulting network configurations, and an evaluation of the solution *robustness*. Chapter 5 highlights the theoretical and practical implications of the study, as well as the limitations and opportunities for future research. Lastly, Chapter 6 provides a summary of the key methods and findings in the thesis, as well as a reflection on the project.

Chapter 2

Preliminaries and Related Research

In this Chapter, the context of the study is further introduced and a review of relevant literature is conducted. The purpose of this section is to address objectives (ii.) and (iii.) from Section 1.4, *i.e.,* to determine which methods of optimization should be considered for SCND under uncertainty, and to determine how robustness can be evaluated. The process has been categorized into several echelons, which are: an investigation into current trends in SCND (Section 2.1); dealing with uncertainty in SCND, specifically the types of uncertainty that are often considered (Section 2.2); and investigating different mathematical modeling approaches for dealing with uncertainty in SCND (Section 2.3). The literature study provides additional background for the problem of planning under uncertainty and introduces the necessary context as found in academia.

2.1 Trends in Supply Chain Network Design

Supply Chain Management (SCM) is considered to have three planning levels split up by time-horizons [13]: operational planning for daily or weekly activities (*e.g.,* routing decisions), tactical planning for a horizon of around three months to three years (*e.g.,* pricing decisions) and strategic planning for a time horizon of at least three years (*e.g.,* facility location decisions). SCND decisions are typically concerned with strategic decisions, *i.e.,* choosing the number of facilities, locations, and capacities for each facility, as well as transport lanes between nodes within the network [13]. The design of supply chain networks is an expensive decision that should be suitable for several years at a minimum [1]. According to Snyder et al. [26], disruptions in recent years have sparked an increase in academic interest regarding the design of supply chains under uncertainty, in part, associated with the exposure of globalized lean supply chains, which are inherently vulnerable to shortages and sudden changes in demand. This vulnerability is largely caused by a lack of redundancy and has been shown to manifest in a supply chain's (in)ability to "recover from the negative impacts of unknown disruptions and adapt to uncertain future events" [11], which is directly linked to the concept of supply chain resilience. The National Academy of Science defines supply chain resilience as an "ability to prepare and plan for, absorb, recover from, and more successfully adapt to adverse events" [19]. According to Golan et al. [11] the interest in supply chain resilience is growing, citing that the number of relevant review publications has grown notably in recent years. Supply chain resilience incorporates the ability of a supply chain to recover from adverse conditions, which is something that is not normally part of the mathematical programming approach to SCND.

Tomlin [29] introduces two ways in which firms deal with supply chain risks: (1) the passive approach, where firms accept disruption risks; and (2) the active approach where firms apply strategies focused on risk mitigation and contingencies. Firms that accept disruption risks leave themselves exposed to sudden or gradual changes in the supply chain's environment, which can be caused by fluctuation in certain parameters such as demand, or drastic events such as a global pandemic. An example of an active versus passive strategy dates back to 2000 when a fire at a Philips plant in the USA led to the closure of the facility for 6 weeks. The facility was responsible for producing semiconductor components for competing phone brands Nokia and Ericsson. Nokia suffered minimal financial impact with its active risk-mitigation strategy as they were able to scale up production at a secondary facility to meet demand. Ericsson had implemented a single-source, passive risk-mitigation strategy and so suffered enormous financial losses as a result of the disruption.

Designing resilient supply chains can be considered an active risk-mitigation strategy; however, the concept of resilience could be considered broad in the context of planning an SCND under uncertainty using mathematical programming, which is, in general, a static method (no consideration of duration or other dynamic impacts) of modeling the supply chain environment [15]. In this case, it may be more beneficial to consider the concept of developing a robust SCND, which compliments the resilient SCND definition by National-Research-Council [19], producing a supply chain that is built in preparation for and can absorb adverse events. Resilience in the full sense could be better measured by modeling approaches such as simulation, which are known to replicate the behavior of a system [15] and thus it is possible to test the effect of different disruptions and recovery strategies.

Planning under uncertainty in the context of SCND requires knowledge about the types of uncertainty that must be considered to develop a robust and/or resilient supply chain. The type of uncertainty influences the modeling approach, as some types of uncertainty are better suited to dynamic models such as simulation, and others are well suited to static models such as mathematical programs.

2.2 Dealing with Uncertainty in SCND

The reason for incorporating uncertainty in SCND is to generate a network design that can meet performance requirements even during extreme realizations of uncertain parameters. Risks within SCs can be divided into two over-arching types: operational risks and disruption risks [13]. Uncertainty is widely considered to be an inherently unknown quantity, differing from risk which is measurable [7]. In the interest of clarity for this study, the term uncertainty is used concerning parameters with either known or unknown random distributions. In this section sources of operational and disruption uncertainty are discussed in more detail, with specific reference to how each of the sources of uncertainty is related to the categorization suggested by Govindan et al. [13].

2.2.1 Disruption Uncertainty

According to Katsaliaki et al. [15], disruption can be characterized by type, intensity, and duration; it is further stated that disruptions can have a local impact, global impact, or both. Typically, disruption is considered a low-probability, high-impact event such as a natural (or man-made) disaster or social/political unrest that influences the supply chain network [13, 15]. In a review of SCND under uncertainty, it was found that disruption is usually modeled in SCND by assigning failure probabilities to specified parameters, *e.g.,* a facility may have a low probability of being shut down and this can be incorporated in the mathematical formulation of the problem [13]. As disruption events are difficult to forecast, mitigation strategies can be used to reduce the impact of disruptions without the need to predict when and where the disruption will take place. Popular strategies to mitigate disruption include facility fortification [20], where individual facilities are designed to resist anticipated disruptive events, and multi-sourcing [26], where multiple suppliers are included in the network even before a disruption occurs. Handling disruption uncertainty relates predominantly to strategic decision-making, as it concerns high-impact events that cannot be predicted accurately, although recourse actions that support the concept of a resilient supply chain may be implemented on the operational or tactical planning level.

2.2.2 Operational Uncertainty

Operational uncertainty applies to uncertainty around medium-to-high probability, low-to-moderate impact events such as demand fluctuations, supplier shortages, changes to company structures (*e.g.,* buyouts), and fluctuations in various costs (*e.g.,* transportation, facility, supplier costs). It should be noted that [15] showed that low-impact, frequent, and extreme realizations of uncertain parameters impact shareholder wealth more than disruptive events. According to Tosarkani and Amin [30], volatility in transportation cost, holding cost, and forecasting the market's demand are the most challenging issues for decision-makers. A survey completed by [13] showed that demand was the most popular subject of uncertainty in relevant academic studies, followed by activity cost (transportation, holding, production), network facility capacities, supply quantity, and availability of facilities. All except facility availability can be considered operational uncertainty, with facility availability falling under disruption uncertainty as it is typically a low-probability, high-impact event. Operational uncertainty is realized in the operational and tactical level of a supply chain as it concerns short-to-medium-term planning horizons; however, these uncertainties can be useful in strategic decision-making because supply chain networks can be designed to function during fluctuations in uncertain parameters.

2.3 Modeling Approaches

Below is an introduction to well-studied methods for planning supply chain networks under conditions of uncertainty. The approaches have been presented according to the categories devised by Rajagopala et al. [21], where analytical methods are distinguished from simulation-based and other methods. An additional subsection has been included, where some of the popular supply chain risk metrics are discussed. The risk metrics can be used to evaluate to what extent an SCND is robust and (or) resilient to fluctuations in uncertain parameters.

2.3.1 Analytical Approaches

This section provides an introduction to popular approaches to handling uncertainty in the context of SCND. Each of the discussed approaches have drawbacks; however, it is important to recognize the different approaches and under what circumstances they might be advantageous.

Deterministic Mixed-Integer-Programs Dealing with uncertainty when making strategic SCND decisions can be approached by creating scenario representations of probable future events, solving the deterministic Mixed-Integer-Programs (MIPs), and comparing the variables in optimal solutions. As described in [3], the decision-maker may want to know how sensitive the optimal solution is to variations in some parameters. This means that the mathematical program can be solved multiple times, where each solution is associated with an adjustment to the uncertain parameter(s). The optimal solutions for all instances can be compared once a sufficient number of scenarios have been solved. Uncertainty could be analyzed using deterministic MIPs by (1) declaring the uncertain parameter and estimating the parameter's behavior, (2) solving the mathematical program with different deterministic values of the uncertain parameter, and (3) comparing the resulting objective values and variable configurations. A drawback is that this does not necessarily provide a configuration that is optimal under *all* scenarios. If the dimensions of the problem increase, or additional sources of uncertainty are introduced, it becomes prohibitively difficult to handle the inter-dependencies within and between uncertain parameters.

Stochastic Programming In stochastic programming, some parameters may be defined as uncertain or follow a known distribution. Stochastic programs can be formulated with multiple decision stages; however, in many SCND problems, it is sufficient to consider only two stages. In the first stage, a set of decisions are made before revealing the realizations of the uncertain parameter(s). In the second stage (also known as *recourse actions*) decisions are made that optimize the program based on the first-stage decisions and realizations of the uncertain parameter(s) [3]. Using the notation by [3], first-stage decisions can be represented by a vector x, and second-stage by $y(x)$. The realization of the uncertain parameters can be given by vector ξ. Here it should be noted that the decision and event sequence follows $x \to \xi \to y(x)$, where first-stage decisions are made *before* realization of the uncertainty, and second-stage decisions are made *after* this realization.

Stochastic programming is a popular academic approach in the context of SCND [10]. One of the major challenges in stochastic programming is generating and assigning an occurrence probability to a suitable set of scenarios [9], Monte Carlo simulation methods have been used to adequately cover a large number of suitable scenarios [34]. This can be considered a drawback, as a large number of scenarios increases the scale of the problem [6], and this can result in enormous (sometimes prohibitive) computational complexity.

The time-to-solution can be reduced by applying algorithms such as the Benders decomposition, which is a cutting plane algorithm that decentralizes the solving process. To handle solving multiple scenarios, the Sample-Average-Approximation technique is a popular approach to estimating the objective function of a stochastic programming problem. The underlying algorithm can make use of techniques such as Monte-Carlo sampling to approximate the objective function.

According to a review by Govindan et al. [13], stochastic programs that make use of SAA and Benders decomposition are widely used in SCND *literature*, although industry experts state that they are rarely applied in practice. The efficient generation of scenarios is crucial for these types of problems, as the quality of the scenarios has a significant influence on the quality and stability of the optimal solution. Li and Zhang [17] proposed a stochastic programming approach under disruption uncertainty with SAA to solve an SCND problem with facility-dependent failures. In their case, the first stage variables were location decisions, after which, a facility could fail with the pre-defined probability before recourse variables were defined. Fattahi et al. [10] formulated a stochastic programming problem with operational and disruption uncertainties, where demand follows a known normal distribution and where DC location, capacity, allocations, and inventory decisions are all first-stage variables. Disruption events occur through different scenarios, causing specific DCs to become unavailable, at which point allocations and inventory decisions can be re-evaluated as second-stage variables. Tsiakis et al. [31] considered an SCND problem under demand uncertainty, where the facility size, location, and quantities should be determined, along with transportation lanes, production rates, and flows. The authors assumed that demand estimates were generated systematically, where scenarios were composed of optimistic and pessimistic cases. Klibi and Martel [16] proposed an SCND problem with multiple periods, including operational and disruption uncertainty in the form of *multihazards*. They define *multihazards* as a capacity reduction and demand reduction/increase over a random number of periods, where the arrival rate of *multihazards* follows a known distribution, and the severity of the disruption is given by a set of two correlated random variables that control the intensity of the impact and the time to recovery (in periods). The authors found that the network design was indeed more robust, although it was recommended that the Monte Carlo approach to scenario generation was not sufficient as "*all important plausible future facets*" [16] were not adequately represented. A study by Govindan and Fattahi [12] considered an SCND problem for a supply chain under demand uncertainty. The authors defined facility locations and capacities as first-stage variables, second-stage variables were split into multiple periods to allow for time-dependent demand variations.

Stochastic programming is, as evidenced by the literature, a well-studied approach to handling uncertainty in SCND. Scenarios can be generated using sampling strategies like Monte Carlo sampling, or, as suggested by [16], a more sophisticated technique that includes some randomness as well as the typical 'pessimistic' and 'optimistic' scenarios. The scenarios that are used in the stochastic program can contain both operational and disruption uncertainties (or either one), making it a versatile approach to planning under uncertainty.

Robust Optimization The key drivers for robust optimization have been (1) to find an optimal solution under the worst-case scenario, and (2) to reveal an optimal solution that protects against infeasibility, that is, to find an optimal solution that is feasible under any realization of an uncertain parameter (within certain bounds of uncertainty) [13]. Therefore, robust optimization can be considered an *immunization* against uncertainty [2]. It is well known that robust optimization produces conservative solutions [13] (*i.e.,* more costly), although approaches that make use of different uncertainty sets (such as ellipsoids) are known to reduce the extent of this effect. Ben-Tal et al. [2] consider stochastic and robust optimization to be "complementary approaches" for dealing with uncertainty in the parameters of optimization models.

In SCND, knowledge about the probability distributions of disruptions and operational parameters tends to be limited, which makes the accurate estimation of uncertain parameters challenging [33]. As robust optimization requires only knowledge of the extreme realizations of uncertain parameters (rather than the probability distributions), it may be well suited to SCND.

In general, optimizing for the *worst-case* is likely to generate a robust solution that can maintain feasibility during disruptions or fluctuations in uncertain parameters. However, the *worst-case* scenario may have a very low probability of occurring, and may not warrant the capital expenditure required. In such cases, it may be more beneficial to find a solution that can maintain feasibility during extreme realizations of uncertain parameters and also maintain lower network costs when extreme events are not occurring.

2.3.2 Simulation-Based Approaches

Simulation is a popular approach to modeling supply chain risks when it is necessary to capture dynamic behavior within the supply chain [21]. According to Rajagopala et al. [21], popular approaches to simulation in the SCND context include Monte Carlo simulation, system dynamics, agent-based modeling, multi-agent systems, and discrete event simulation. Bueno-Solano and Cedillo-Campos [4] used system dynamics to model the effect of major terrorist attacks on supply chain performance, and how policies such as tightened border restrictions can result in significant inventory build-up, reducing the performance of the supply chain. Schmitt and Singh [24] proposed a method to improve the resilience of supply chain networks by making use of discrete-event simulation to study the influence of disruptions *and* demand uncertainty to analyze different inventory strategies. Colicchia et al. [5] used Monte Carlo simulation to look at uncertainty in "*supply continuity*", where the lead time between suppliers was increased in duration and variation.

Generally, simulation has been used as a tool for designing and evaluating resilient supply chains, mitigating supply chain risks, and reducing the *ripple* effect of these disruptions, rather than configuring a supply chain network that is inherently robust. Furthermore, simulation is typically applied to problems that involve operational or tactical decision-making, as it is a useful tool for modeling dynamic systems where disruptions have both time and severity dimensions.

2.3.3 Evaluating Robustness

An important factor when designing resilient or robust supply chain networks is to measure the level of resilience or robustness. Distinguishing between structural and functional/operational robustness is a popular approach [8, 18]. Structural robustness generally makes use of theoretical graph concepts and considers how important each node and edge is. When several nodes or edges score highly in these measures, it may be that the network is *over-reliant* on these elements, and thus the overall robustness may be lower. Operational (or functional) robustness considers the dynamic processes within the network, given that the network structure is fixed [18]. When measuring operational robustness in a supply chain, many of the popular approaches are inherited from the insurance and financial context [14].

Soni et al. [27] used graph theory to enable the efficient modeling of supply networks that provides information on risks through a single numeric index. Dong [8] developed a robustness index for supply chain networks that considers the structural robustness and functional robustness of a supply chain network. Shukla et al. [25] developed a robustness metric that considered the expected cost of disruption given disruption probabilities and associated costs. A MIP was formulated to minimize cost with the robustness metric in the objective function.

As previously discussed, a *robust* supply chain can maintain normal operating performance in the face of disruption and fluctuations in uncertain parameters. It may also be appropriate to evaluate *robustness* of a supply chain by quantifying to what extent the network can maintain normal operating performance in different scenarios. A generic measure of robustness may not always be applicable, as different supply chain networks have specific objectives and requirements. Therefore, it may be useful to measure robustness from a practical sense, considering the networks' ability to maintain normal operating performance.

2.4 Key Findings from Related Research

This section has discussed which methods of optimization under uncertainty should be considered for SCND, and determining which methods for evaluating robustness should be considered. Furthermore, it was evidenced that strategic planning under uncertainty in SCND may be best approached using an analytical model, as the simulation approach is less adept at network configuration and strategic planning. The reviewed literature indicates that, in general, robust optimization tends to be overly conservative, while stochastic programming is well suited to SCND due to the sequential decision-making process and ability to capture the nature of uncertainty in a more precise manner, compared with robust optimization.

The approach by Fattahi et al. [10] was particularly relevant in capturing both operational and disruption uncertainty in a stochastic programming model and influences the methods proposed in this study. Furthermore, the research by Santoso et al. [23] made use of SAA, Benders decomposition, and multiple sources of uncertainty, which forms an integral part of the foundation of this study.

In terms of measuring the robustness of an SC network, risk and robustness measures are well studied, although there is a lack of consensus on approaches. As such, it is most appropriate to return to the overarching definition of robustness in this study: *the ability for a supply chain to maintain normal operating performance in the face of disruption and fluctuations in uncertain parameters*; this definition motivates measuring robustness by examining to what extent a network configuration can operate normally based on a set of scenarios.

Chapter 3

Modeling Approach

According to the objectives discussed in Section 1.4, and the conclusion of the literature survey, stochastic programming has been selected as the most appropriate method for this study on planning under uncertainty in SCND. In this section a deterministic formulation of the SCND *mixed-integer-problem* (MIP) is constructed (Section 3.1) before addressing uncertainty in a set of parameters using a stochastic programming approach (Section 3.2). Section 3.3 unpacks the method of developing and generating scenarios for the stochastic programming model.

3.1 Deterministic Formulation

The SC network in this study is geographically dispersed, and a predefined set of transport lanes connect the suppliers, resources, and customers. A simplified schematic of the SC network is given in Fig. 3.1.

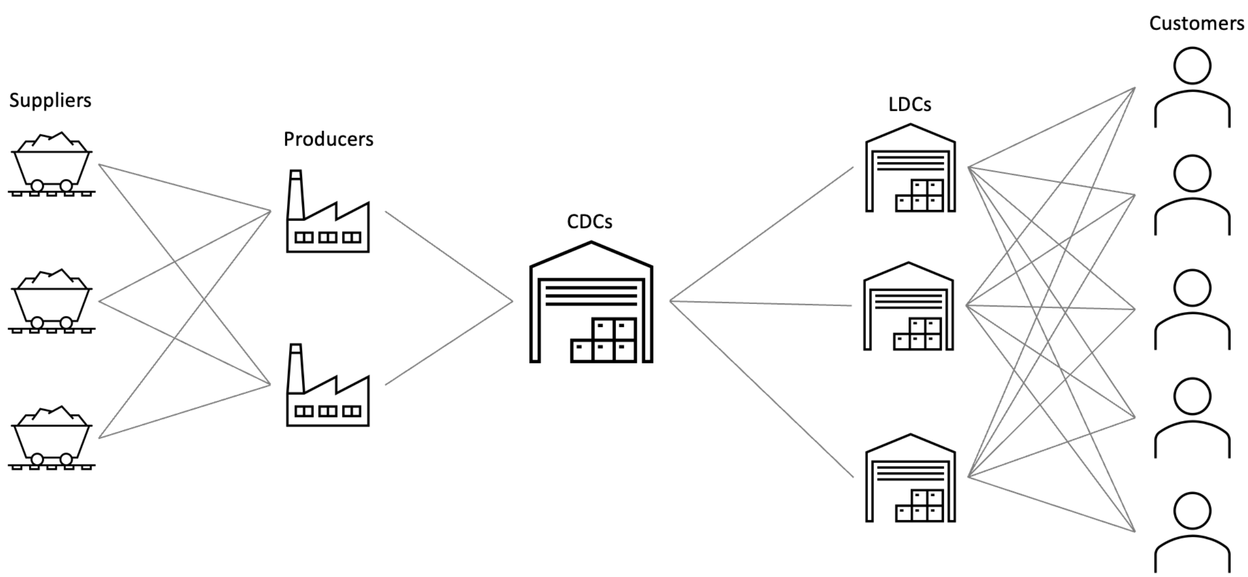

Figure 3.1: Simplified example of the baseline SC network used in this study.

Raw materials travel to production facilities, where they are converted into finished goods. The finished goods are transported to Central Distribution Centers (CDCs), before being dispatched to Local Distribution Centers (LDCs). Finished goods are distributed from LDCs to customers. Optimizing this SC network requires, among others, decisions on facility availability (open or closed), selecting suppliers, selecting transport lanes, and establishing the volume of flow of goods between resources.

The components of the mathematical program associated with a generic SC network are given in Table 3.1, which show sets, parameters, and variables, respectively. Some of the key assumptions are that locations of resources have already been determined, customers can receive finished products from any number of resources, there is no inventory build-up allowed, and all material that enters the network must be delivered to a customer in the form of a finished product, and the capacity of resources cannot be increased.

Sets	Description
C	Customers $c \in C$
T	Transport types $t \in T$
R	Resources $r \in R$
S	Suppliers $s \in S$
P	Products $p \in P$
P_{final}	Final products $p_{final} \in P$
P_{sub}	Sub-products $p_{sub} \in P$
L	Locations $l \in L$
L_{to}	Locations to $l_{to} \in L$
L_{from}	Locations from $l_{from} \in L$
B	Bill of needs $b \in P$
Parameters	**Description**
$D_{p_{\text{final}},c,l}$	Customer c demand for product p_{final} at location l
$C^{\max}_{s,l}$	Overall capacity for supplier s at location l
$C^{\min}_{s,l}$	Minimum amount that supplier s at location l will provide
$C^{\max}_{p,s,l}$	Maximum amount of product p that supplier s at location l can provide
$C^{\min}_{p,s,l}$	Minimum amount of product p that supplier s at location l can provide
$Tr^{\min}_{t,r,l_{\text{from}},l_{\text{to}}}$	Minimum amount of any product that transport t can move from one resource to another
$Tr^{\max}_{t,p,l_{\text{from}},l_{\text{to}}}$	Maximum amount of any product that transport t can move from one resource to another
$Pr^{\min}_{r,l}$	Minimum production capacity for resource r at location l
$Pr^{\max}_{r,l}$	Total production capacity for resource r at location l
$Pr^{\min}_{p,r,l}$	Minimum production capacity for product p at resource r at location l
$Pr^{\max}_{p,r,l}$	Maximum production capacity for product p at resource r at location l
$BON^{\text{created}}_{b,l,r,p}$	Bill of needs b produced at resource l
$BON^{\text{needed}}_{b,l,r,p}$	Bill of needs b needed at resource l
$BON^{\text{cost}}_{b,r,l}$	Production cost of bill of requirements b at resource location.
$R^{\min}_{p,r,l}$	Minimum amount of products p flowing through resource r if resource is in use
$R^{\max}_{p,r,l}$	Maximum amount of products p flowing through resource r if resource is in use
$SC_{p,s,l}$	Cost per product p supplied by supplier s

$NC_{t,p,l_{\text{from}},l_{to}}$	Cost (per distance unit) of using transport t to move a product p from one location to another
$PV_{p,r,l}$	Variable production cost per product p at production facility r
$BP_{b,r,l}$	Cost of production for bill of resources b at facility r
$DCV_{p,r,l}$	Variable DC throughput cost for product p at DC r
$RF_{r,l}$	Fixed cost for keeping resource r available
$RFP_{p,r,l}$	Fixed cost for keeping resource r available for product p
$RC_{r,l}$	Cost of closing resource r
Variables	Description
$DIST_{t,p_{final},r,l_{from},c,l}$	Amount of products p transported by t to customer c from resource r
$DIR_{t,p_{final},s,l_{from},c,l}$	Amount of products p transported by t to customer c from FG supplier s
$S_{t,p,s,l,r,l_{to}}$	Amount of products p transported by t to resource r from supplier s
$SP_{p,s,l}$	Amount of each product p provided by supplier s
$NF_{t,p_{sub},r_{from},r_{to},l_{from},l_{to}}$	Flow of products between suppliers and resources, and between resources
$PROD_{p,r,l}$	Number of product p produced at resource r
$BONP_{b,r,l}$	Bill of needs b for production at resource r, location l
$DCB_{r,l}$	Binary indication of whether resource r is in use
$THRU_{p,r,l}$	Total throughput for product p at DC r
$DCBP_{p,r,l}$	Binary indication of whether resource r is in use for product p

Table 3.1: Sets, Parameters, and Variables in baseline SC network.

The objective of the mathematical program is to minimize the total cost of the SC network, where all cost components from parameters in Table 3.1 are combined. The vector $\boldsymbol{x}$ is shorthand for the vector of flow variables defined in Table 3.2, and are key components in the transportation cost of the network $f_{transport}$, which is measured in the four dimensions of $\boldsymbol{x}$: transport cost, resource cost, handling cost, and production cost. Primary transport costs are given by Eq. (3.1), inter-resource transport costs are given by Eq. (3.2), distribution (or secondary transport) costs are given by Eq. (3.3), and direct transport costs are given by Eq. (3.4).

$$f_{\text{transport}}(\boldsymbol{x}_{\text{transport}}) := \sum_{t,p,s,l_{from},r,l_{to}} NC_{t,p,l_{from},l_{to}} SP_{t,p,s,l_{from},r,l_{to}} + \tag{3.1}$$

$$\sum_{t,p,r,l_{from},l_{to}} NC_{t,p,l_{from},l_{to}} NF_{t,p,r,l_{from},l_{to}} + \tag{3.2}$$

$$\sum_{t,p,r,l_{from},c,l_{to}} NC_{t,p,l_{from},l_{to}} DIST_{t,p,r,l_{from},c,l_{to}} + \tag{3.3}$$

$$\sum_{t,p,r,l_{from},c,l_{to}} NC_{t,p,l_{from},l_{to}} DIR_{t,p,s,l_{from},c,l_{to}}. \tag{3.4}$$

Resource costs consist of the fixed costs associated with operating a facility within the supply chain for a given period, denoted as $f_{resource}(\boldsymbol{y}_{resource})$. These costs are measured for production facilities, CDCs, and LDCs, and are given by (3.5), which is dependent on the binary facility open/close decisions given by vector

Flow Variables $\boldsymbol{x}$	Subset
$S_{p,s,l}$	Supply
$SP_{t,p,s,l_{from},r,l_{to}}$	Flow, Supply, Transport
$NF_{t,p,r,l_{from},l_{to}}$	Flow, Transport
$DIST_{t,p,r,l_{from},c,l_{to}}$	Distributing, Transport
$DIR_{t,p,s,l_{from},c,l_{to}})$	Distributing, Transport
$THRU_{p,r,l}$	Handling
$PROD_{p,r,l}$	Flow, Production
$BONP_{b,r,l}$	Flow, Production
Binary Variables $\boldsymbol{y}$	**Subset**
$PB_{r,l}$	Production
$PB_{p,r,l}$	Production, Product
$DCB_{r,l}$	Resources
$DCBP_{p,r,l}$	Resources, Product

Table 3.2: Flow and binary variables in the SC Network.

$\boldsymbol{y}$ (see binary and flow variables in Table 3.2);

$$\begin{aligned} f_{\text{resource}}(\boldsymbol{y}) := &\sum_{r,l} DCB_{r,l}RF_{r,l} + \sum_{p,r,l} DCBP_{p,r,l}RFP_{p,r,l} + \\ &\sum_{r,l} PB_{r,l}RF_{r,l} + \sum_{p,r,l} PB_{p,r,l}RFP_{p,r,l}. \end{aligned} \tag{3.5}$$

Handling costs are assumed to only occur in the DC facilities and thus are considered a function of throughput. The total handling cost is given by the product of throughput and variable throughput costs in (3.6)

$$f_{\text{handling}}(\boldsymbol{x}_{\text{handling}}) := \sum_{p,r,l} THRU_{p,r,l}DCV_{p,r,l}. \tag{3.6}$$

The total production cost consists of variable production costs and a production cost associated with the bill of materials. This is given by (3.7)

$$f_{\text{production}}(\boldsymbol{x}_{\text{production}}) := \sum_{p,r,l} PROD_{p,r,l}PV_{p,r,l} + \sum_{b,r,l} BONP_{b,r,l}BON^{\text{cost}}_{b,r,l}. \tag{3.7}$$

The total physical cost of the SC network is given by the sum of (3.1)-(3.7), which is the objective that the mathematical optimization program aims to minimize.

Multiple constraints are imposed on the SC Network model to ensure that realistic and required conditions are observed. The equality constraint in (3.8) is imposed to ensure that demand is met precisely. In (3.8), the vector $\boldsymbol{x}_{dist}$ represents the distribution flow variables, with the constraint being given by

$$h_{\text{demand}}(\boldsymbol{x}_{\text{dist}}) := \sum_{t,r,l_{from}} DIST_{t,p,r,l_{from},c,l} + \sum_{t,s,l_{from}} DIR_{t,p,s,l_{from},c,l} - D_{p,c,l} \quad = \quad 0 \quad \forall p,c,l. \tag{3.8}$$

For suppliers, the amount of raw materials provided must be between a minimum and maximum threshold both overall and per product. This requirement is enforced through four constraints that apply a minimum and maximum for the total number of products sourced from a particular supplier. (3.9) demonstrates this in the context of a maximum supply constraint, where $\boldsymbol{x}_{supply}$ is the vector of supply flow variables. (3.10) represents the product level version of (3.9), and is given by

$$g_{\text{supply}_{\max}}(\boldsymbol{x}_{\text{supply}}) := \sum_{p} SP_{p,s,l} + \sum_{t,p,c,l_{to}} DIR_{t,p,s,l,c,l_{to}} - C^{\max}_{s,l} \leq 0 \qquad \forall s,l \tag{3.9}$$

$$g_{\text{supply}^{prod}_{\max}}(\boldsymbol{x}_{\text{supply}}) := SP_{p,s,l} + \sum_{t,c,l_{to}} DIR_{t,p,s,l,c,l_{to}} - C^{\max,\text{prod}}_{p,s,l} \leq 0 \qquad \forall p,s,l. \tag{3.10}$$

It should not be possible for materials or products to build up at a resource unless there is an explicit inventory facility whereby stock can be kept. A flow balance constraint (shown in (3.11)) ensures that the total material in-flow at a resource is equal to the total outflow, that is

$$\begin{aligned} h_{\text{flow}}(\boldsymbol{x}) := \sum_{s} SP_{p,s,l} + \sum_{t,r_{from},l_{from}} NF_{t,p,r_{from},l_{from},r,l} + \\ PROD_{p,r,l} + \sum_{b} BONP_{b,l} BON^{\text{created}}_{b,l,r,p} - (\sum_{b} BONP_{b,l} BON^{\text{needed}}_{b,l,r,p} + \\ \sum_{t,r_{to},l_{to}} NF_{t,p,r_{to},l_{to},r,l} + \sum_{t,c,l_{to}} DIST_{t,p,r,l_{from},c,l}) = 0 \qquad \forall p,r,l. \end{aligned} \tag{3.11}$$

Each DC, similar to suppliers, has a maximum and minimum throughput capacity both overall and for individual products. This restriction is maintained by four constraints that ensure maximum and minimum throughputs are not violated on the net and product levels. The constraint in (3.12) demonstrates this in the context of ensuring that the minimum throughput capacity is met on the *product level*, where $\boldsymbol{y}_{\text{product}} \subset \boldsymbol{y}_{\text{resources}}$ represents the $p \times r \times l$ vector of product resource binary variables. (3.13) is the net level representation of (3.12), given below

$$g_{\text{thru}^{\text{prod}}_{\min}}(\boldsymbol{x}_{\text{handling}}, \boldsymbol{y}_{\text{product}}) := THRU_{p,r,l} - R^{\min}_{p,r,l} DCBP_{p,r,l} \quad \geq \quad 0 \qquad \forall p,r,l \tag{3.12}$$

$$g_{\text{thru}_{\min}}(\boldsymbol{x}_{\text{handling}}, \boldsymbol{y}_{\text{resource}}) := \sum_{p} THRU_{p,r,l} - R^{\min}_{r,l} DCB_{r,l} \quad \geq \quad 0 \qquad \forall r,l. \tag{3.13}$$

As with DCs (above), four constraints are also applied to meet minimum production amounts and to not exceed maximum amounts, both overall and for individual products. (3.14) demonstrates this concept in terms of a minimum total production amount at a given production facility, and (3.15) in terms of minimum total production per product at a facility, that is

$$g_{\text{production}_{\min}}(\boldsymbol{x}_{\text{production}}, \boldsymbol{y}_{\text{production}}) := \sum_{p} PROD_{p,r,l} - Pr^{\min}_{r,l} PB_{r,l} \quad \geq \quad 0 \qquad \forall r,l \tag{3.14}$$

$$g_{\text{production}^{\text{prod}}_{\min}}(\boldsymbol{x}_{\text{production}}, \boldsymbol{y}_{\text{production}}) := PROD_{p,r,l} - Pr^{\min}_{p,r,l} PB_{p,r,l} \quad \geq \quad 0 \qquad \forall p,r,l. \tag{3.15}$$

Transportation is governed by two constraints: a maximum and minimum capacity restriction. (3.16) demonstrates the minimum capacity restriction, which ensures that, for each transport lane and product, at least $Tr^{\min}$ *products* are transported. Again, this constraint must be defined on a net level, as per (3.16), and on a *product level*, (3.17), see below

$$g_{\text{transport}_{\min}}(\boldsymbol{x}_{\text{transport}}) := \sum_{p,r,c} DIST_{t,p,r,l_{from},c,l_{to}} + \sum_{p,r_{from},r_{to}} NF_{t,p,r_{from},l_{from},r_{to},l_{to}} +$$

$$\sum_{p,s,r} SP_{t,p,s,l_{from},r,l_{to}} + \sum_{p,s,c} DIR_{t,p,s,l_{from},c,l_{to}} - Tr^{\min}_{t,r,l_{from},l_{to}} \geq 0 \qquad \forall t, l_{from}, l_{to} \tag{3.16}$$

$$g_{\text{transport}^{\text{prod}}_{\min}}(\boldsymbol{x}_{\text{transport}}) := \sum_{r,c} DIST_{t,p,r,l_{from},c,l_{to}} + \sum_{r_{from},r_{to}} NF_{t,p,r_{from},l_{from},r_{to},l_{to}} +$$

$$\sum_{s,r} SP_{t,p,s,l_{from},r,l_{to}} + \sum_{s,c} DIR_{t,p,s,l_{from},c,l_{to}} - Tr^{\min}_{t,p,r,l_{from},l_{to}} \geq 0 \qquad \forall t, p, l_{from}, l_{to}. \tag{3.17}$$

Finally, two constraints are necessary to ensure that resources operate only under certain conditions, (3.18) ensures that a production facility is only used for a product when the production facility is in use, and (3.19) ensures a product only passes through a DC if the resource is in use; both being given by

$$g_p(\boldsymbol{y}_{\text{production}}) := PB_{p,r,l} - PB_{r,l} \leq 0 \qquad \forall p, r, l \tag{3.18}$$

$$g_d(\boldsymbol{y}_{\text{resource}}) := DCBP_{p,r,l} - DCB_{r,l} \leq 0 \qquad \forall p, r, l. \tag{3.19}$$

The MIP in this section can be solved to find the optimal allocation of resources within the SC Network, as well as which suppliers and transport lanes best support the resources. Different scenarios can be tested to evaluate whether the SC network can operate under different circumstances, *e.g.,* increased demand or variance in transport costs. It may be the case that many scenarios are necessary for an informed decision, which can lead to a challenging and inefficient analysis.

3.2 Two-Stage Stochastic Programming Approach

A two-stage stochastic programming approach to SCND allows some decisions (first-stage) to be made without knowledge of uncertain parameter realizations, while other decisions (second-stage) can be made once the uncertain parameters have been realized. It is assumed that second-stage variables can be adjusted on the operational or tactical planning level, meaning that they can be used as recourse actions to account for variations in uncertain parameters.

The classification of first-stage and second-stage variables could depend on the supply chain, *e.g.,* a responsive supply chain where production is driven by customer demand may have production and distribution variables as second-stage, while facility open/close decisions would remain fixed. Conversely, a supply chain supporting a product that is pushed into the market may not have as many recourse actions, as, for example, supply and/or production levels are decided before customer demand can be accurately forecast. In the latter case, first-stage decisions could extend to variables such as production amount, inter-resource transport, and throughput at CDCs.

The stochastic programming approach in this study assumes that supplier contracts and open/close decisions for facilities should be made as first-stage decisions, as these decisions best complement the strategic planning horizon in this study. Production decisions, throughput, transport type, and transport route decisions are considered second-stage decisions, as they fall under the tactical and operational planning horizon. It is assumed that demand $\boldsymbol{d}$, per-mile-transport costs $\boldsymbol{t}$, and the maximum throughput capacity of each DC $\boldsymbol{c}$ are independent stochastic parameters. It is further assumed that parameters $\boldsymbol{d}$ and $\boldsymbol{t}$ have known multivariate distributions. The parameter $\boldsymbol{c}$ represents a probability distribution that models the behavior of throughput capacity at a facility, where there is a high probability of no disruption $c \approx 1$ and a lower probability of $c << 1$. The stochastic data vector is given by $\boldsymbol{\xi} = (\boldsymbol{d}, \boldsymbol{t}, \boldsymbol{c})$. A realization of this vector is given by $\xi = (d, t, c)$. The two-stage stochastic program can be formulated as follows:

$$
\begin{aligned}
\min \quad & f_{\text{resources}}(\boldsymbol{y}) + \mathbb{E}Q(\boldsymbol{y}, \boldsymbol{\xi}) \\
\text{s.t.} \quad & g_p(\boldsymbol{y}_{\text{production}}) \leq 0 \\
& g_d(\boldsymbol{y}_{\text{resource}}) \leq 0 \\
& \boldsymbol{y} \in \{0, 1\}
\end{aligned}
\tag{3.20}
$$

Where $Q(y, \xi)$ is the optimal value of the second stage problem given by (3.21).

$$
\begin{aligned}
\min \quad & f_{\text{transport}}(\boldsymbol{x}_{\text{transport}}) + f_{\text{handling}}(\boldsymbol{x}_{\text{handling}}) + f_{\text{production}}(\boldsymbol{x}_{\text{production}}) \\
\text{s.t.} \quad & h_{\text{demand}}(\boldsymbol{x}_{\text{dist}}) = 0 \\
& h_{\text{flow}}(\boldsymbol{x}) = 0 \\
& g_{\text{supply}_{\text{max}}}(\boldsymbol{x}_{\text{supply}}) \leq 0 \\
& g_{\text{supply}^{\text{prod}}_{\text{max}}}(\boldsymbol{x}_{\text{supply}}) \leq 0 \\
& g_{\text{supply}_{\text{min}}}(\boldsymbol{x}_{\text{supply}}) \geq 0 \\
& g_{\text{supply}^{\text{prod}}_{\text{min}}}(\boldsymbol{x}_{\text{supply}}) \geq 0 \\
& g_{\text{thru}_{\text{max}}}(\boldsymbol{x}_{\text{handling}}, \boldsymbol{y}_{\text{RP}}) \geq 0 \\
& g_{\text{thru}_{\text{min}}}(\boldsymbol{x}_{\text{handling}}, \boldsymbol{y}_{\text{product}}) \leq 0 \\
& g_{\text{thru}^{\text{prod}}_{\text{max}}}(\boldsymbol{x}_{\text{handling}}, \boldsymbol{y}_{\text{product}}) \geq 0 \\
& g_{\text{thru}^{\text{prod}}_{\text{min}}}(\boldsymbol{x}_{\text{handling}}, \boldsymbol{y}_{\text{product}}) \leq 0 \\
& g_{\text{production}_{\text{min}}}(\boldsymbol{x}_{\text{production}}, \boldsymbol{y}_{\text{production}}) \geq 0 \\
& g_{\text{production}_{\text{max}}}(\boldsymbol{x}_{\text{production}}, \boldsymbol{y}_{\text{production}}) \leq 0 \\
& g_{\text{production}^{\text{prod}}_{\text{min}}}(\boldsymbol{x}_{\text{production}}, \boldsymbol{y}_{\text{production}}) \geq 0 \\
& g_{\text{production}^{\text{prod}}_{\text{max}}}(\boldsymbol{x}_{\text{production}}, \boldsymbol{y}_{\text{production}}) \leq 0 \\
& g_{\text{transport}_{\text{min}}}(\boldsymbol{x}_{\text{transport}}) \geq 0 \\
& g_{\text{transport}_{\text{max}}}(\boldsymbol{x}_{\text{transport}}) \leq 0 \\
& g_{\text{transport}^{\text{prod}}_{\text{min}}}(\boldsymbol{x}_{\text{transport}}) \geq 0
\end{aligned}
\tag{3.21}
$$

$$g_{\text{transport}^{\text{prod}}_{\text{max}}}(\boldsymbol{x}_{\text{transport}}) \leq 0$$
$$\boldsymbol{x}_{\text{transport}} \in \mathbb{R}_{+}^{tpl_{\text{from}}l_{\text{to}}}$$
$$\boldsymbol{x}_{\text{dist}} \in \mathbb{R}_{+}^{pcl}$$
$$\boldsymbol{x}_{\text{supply}} \in \mathbb{R}_{+}^{sl}$$
$$\boldsymbol{x}_{\text{handling}}, \boldsymbol{x}_{\text{production}} \in \mathbb{R}_{+}^{prl}$$

3.2.1 SAA Strategy

To evaluate the expectation of the second-stage problem in the objective (3.21), the SAA approach is used. That is, the stochastic vector $\boldsymbol{\xi}$ contains random samples $\xi_1, \xi_2, \ldots, \xi_N$, where N is the number of scenarios that are evaluated in the expectation of (3.21). If the parameters in $\boldsymbol{\xi}$ follow a continuous distribution, and N is sufficiently large, it can be stated that

$$\mathbb{E}[Q(y, \xi_n)] \approx \frac{1}{N} \sum_{n=1}^{N} Q(y, \xi_n). \tag{3.22}$$

Thus, for N large enough, the objective in (3.20) can be stated as:

$$\min_{y \in \boldsymbol{y}} \{ f_{\text{resources}}(y) + \frac{1}{N} \sum_{n=1}^{N} Q(y, \xi_n) \}. \tag{3.23}$$

The rate of convergence for (3.22) is exponential in N, such that the number of scenarios required for a good approximation of (3.21) can, in many cases, be modest [23].

3.2.2 Benders Decomposition

Benders decomposition is a cutting plane algorithm that is suited to SAA problems, as an SAA problem is concerned with minimizing the summation of N convex non-closed form functions. The Benders decomposition approach requires the declaration of a master problem and a set of sub-problems. The master problem represents a high-level version of the stochastic programming problem and captures the main decision variables. The master problem also provides the initial solution that serves as a starting point for the algorithm. As the algorithm iterates, solutions from solving the sub-problems are incorporated to improve the overall solution quality. The sub-problems address subsets of the variables and constraints and represent lower-level decisions. The sub-problems are more manageable to solve and can be solved separately from the master problem. This problem separation means that Benders decomposition can efficiently exploit decomposed SAA problems into smaller, tractable components. Below, the Benders decomposition for (3.23) is stated in brief, adapted from the study by Santoso et al. [23].

Step 1: Initialize the decomposition with lower bound $lb = -\infty$ and upper bound $ub = \infty$, as well as a counter to keep track of iterations, denoted as i. The objective value can be denoted as $\hat{f}(\tilde{y})$, where $\tilde{y}$ is the current best feasible solution.

Step 2: Define and solve the master problem given by:

$$
\begin{aligned}
\text{lb} := \min_{\boldsymbol{y},\theta} \quad & f_{\text{resources}}(\boldsymbol{y}) + \theta \\
\text{s.t.} \quad & g_p(\boldsymbol{y}_{\text{production}}) \leq 0 \\
& g_d(\boldsymbol{y}_{\text{resource}}) \leq 0 \\
& \boldsymbol{y} \in \{0,1\} \\
& \theta \geq a_j^\top y + b_j, \quad j = 1, \ldots, i,
\end{aligned}
\tag{3.24}
$$

where θ is lb estimate of (3.22), enforced by cuts of the form $\theta \geq a_k^\top + b_k$. It follows also that the feasible solution of (3.24) for iteration i is y_i.

Step 3: Solve the sub-problem given by (3.21) for each scenario $n = 1, \ldots, N$ given the solution y_i is feasible and an uncertainty realization ξ_n, the resulting objective is $\hat{f}(y_i)$. Recall ub is the upperbound solution, if $\hat{f}(y_i) < ub$, then the current solution $\tilde{\boldsymbol{y}} := y_i$, and $ub := \hat{f}(\tilde{\boldsymbol{y}})$.

Step 4: Check the difference between ub and lb, that is, if $ub - lb < \delta_{\text{critical}}$, where δ_{critical} is the tolerance. When $ub - lb < \delta_{\text{critical}}$ is true, $\tilde{y}$ can be considered optimal and the algorithm is terminated, otherwise, the next step is conducted.

Step 5: The cut coefficients $a_{i+1}^\top$ and $b_{i+1}^\top$ must be computed and updated using the optimal dual solutions from (3.21) for each $n = 1, \ldots, N$ and for each sub-problem about y_i, ξ_n. This optimality cut aims to improve the $a_{i+1}^\top + b_{i+1}^\top = \frac{1}{N} \sum_{n=1}^{N} Q(y_i, \xi_n)$ estimate from (3.24), and it is exact for a given y_i. After computing the cut coefficients, return to Step 1.

It should be noted that the Benders decomposition is a finite scheme, although in some cases the convergence behavior can require a prohibitively large number of iterations before converging to the optimal solution, which may even result in longer computation times compared with a solving algorithm without a Benders decomposition. For this reason, the tolerance δ_{critical} (between a MIP solution and its linear program relaxation) can be set such that convergence to a near-optimal solution can occur within a reasonable process time.

3.3 Generating Scenarios

The random parameters (d, t, c) are assumed to follow independent, continuous, probability distributions. These parameters must be modeled in a way that provides a representative and reduced set of scenarios. To address this, a multivariate sampling strategy is implemented, where correlation is independently introduced for *within* demand and transportation cost random variables. The input data for the model proposed in Section 3.2 is introduced in this section, along with the statistical properties of each random parameter. Finally, the method of combining scenarios containing realizations of random parameters is discussed.

In the deterministic model from Section 3.1, the uncertain parameters are not distinguished as random variables, rather they assume a deterministic value. The value for each demand ($D_{p,c,l}$), transport per distance-unit cost ($NC_{t,p,l_{\text{from}},l_{\text{to}}}$), and facility capacity ($R^{\max}_{p,r,l}$) in the deterministic model is assumed to represent the expected value for each associated random parameter.

3.3.1 Demand Uncertainty

Let $\tilde{D}_{p,c,l}$ represent the random parameter for demand and $D_{p,c,l}$ the demand in the deterministic case. It is first assumed that for each $p \in P$, $c \in C$, demand follows an independent normal distribution such that $\tilde{D}_{p,c,l} \sim \mathcal{N}(\mu_{p,c}, \sigma^2_{p,c})$, where $\mu_{p,c} = D_{p,c,l}$, and $\sigma_{p,c} = \lambda_D D_{p,c,l}$. The coefficient λ_D controls the variance of $\tilde{D}_{p,c,l}$ such that the standard deviation of demand is proportional to the expected demand $\mu_{p,c}$.

Using univariate normal distribution to simulate demand for each p, c could mean that an excessive number of scenarios are required to achieve the aim of providing a representative and reduced set of scenarios, especially when combining these samples with other independent sources of uncertainty. To address this challenge, one approach that can be used is to employ a multivariate normal distribution for sampling demand. Furthermore, it may be the case that a multivariate distribution is more realistic than multiple univariate distributions. It is therefore necessary to introduce a covariance matrix and subsequently correlation for demand between customers and products. Sampling from a multivariate distribution is more likely to produce a representative set of scenarios because the correlation between customer and product demand can be controlled, which could result in more realistic scenarios and less redundant scenarios (if the correlation between customer demand is present in the underlying customer behavior). The level of correlation should depend on the type of supply chain network that is being modeled, as some products are known to have higher demand correlations than others (*e.g.,* ice-cream sales are known to increase with warmer weather, which often affects entire regions). Introducing some correlation within demand could realize a more representative set of scenarios by reducing the number of unrealistic, duplicate, or redundant scenarios that are generated.

The proposed multivariate normal distribution can be denoted as $\mathcal{N}(\mu, \Sigma)$, where Σ is the positive semi-definite (*p.s.d.*) covariance matrix that measures how random variables change together. Say X, Y are random variables, the covariance between them can be given by $C(X, Y)$. In the case where the variables are independent, $C(X, Y) = 0$. The variance of X is equal to $C(X, X)$. In this case, the variables in the covariance matrix are each $\tilde{D}_{p,c,l}$. The covariance matrix $\Sigma \in \mathbb{R}^{pc}$ was constructed by first taking a diagonal $I_{pc}\boldsymbol{\sigma}$, where $\boldsymbol{\sigma}$ is the $|P||C| \times 1$ vector of $\sigma_{p,c}$. Let $\tilde{n} = |P||C|$ and henceforth the two-dimensional arrays $\mu_{p,c}$ and $\sigma_{p,c}$ can be considered one-dimensional arrays with combinations *pc*, such that a value of μ_{pc} or σ_{pc} can be given by μ_i or σ_i, respectively, and $i = 1, \ldots, \tilde{n}$. Let entries into the Σ matrix be given by Σ_{ij}, $i, j = 1, \ldots, \tilde{n}$, such that $\Sigma_{ij} = C(\sigma^2_i, \sigma^2_j)$. As values for $C(\sigma^2_i, \sigma^2_j)$ are not known empirically, it is necessary to construct reasonable values that ensure the *p.s.d.* property of Σ holds. As the dimensions of Σ are fairly large ($\tilde{n}$ can have a magnitude of several thousand, or even higher), construction of a *p.s.d.* matrix is non-trivial. A series of trials showed that (3.25) consistently resulted in Σ with *p.s.d.* properties, see (3.25) below:

$$\Sigma_{ij} := C(\sigma_i^2, \sigma_j^2) = \rho \min\{\sigma_i^2, \sigma_j^2\}, \qquad i, j = 1, \ldots, \tilde{n};\ i \neq j. \tag{3.25}$$

An advantage in selecting the $\min\{\sigma_i^2, \sigma_j^2\}$ function in (3.25) is that it prevents samples from smaller μ_i incurring abnormally an abnormally large variance. This decision means that when $\mu_i \approx \mu_j$, the correlation between samples will approximate ρ, while when $\mu_i >> \mu_j$ or $\mu_i << \mu_j$ $(i \neq j)$, correlation will be strictly lower than ρ.

Sampling from a multivariate normal distribution can be done using affine transformations of a normal distribution. Say X is normally distributed, then Y is also normally distributed if $Y = LX + w$, where L is a linear transformation and w represents a vector; specifically, $Y \sim (w + L\mu_X, L\Sigma_X L^\top)$. This can be proven as per the notes by Roelants [22] shown in (3.26)-(3.32).

Proof. Demonstrating that an affine transformation of vector $Y = LX + w$ using $X \sim \mathcal{N}(\mu_X, \Sigma_X)$ results in Y following a normal distribution with μ_Y and Σ_Y in terms of μ_X, Σ_X, and the linear transformation L, respectively [22].

$$\mu_Y = \mathbb{E}Y = \mathbb{E}[LX + w] = L\mu_X + w \tag{3.26}$$
$$\Sigma_Y = \mathbb{E}[(Y - \mu_Y)(Y - \mu_Y)^\top] \tag{3.27}$$
$$= \mathbb{E}[(LX + w - L\mu_X - w)(LC + w - L\mu_X - w)^\top] \tag{3.28}$$
$$= \mathbb{E}[(L(X - \mu_X))(L(X - \mu_X))^\top] \tag{3.29}$$
$$= \mathbb{E}[L(X - \mu_X)(X - \mu_X)^\top L^\top] \tag{3.30}$$
$$= L\mathbb{E}[(X - \mu_X)(X - \mu_X)^\top]L^\top \tag{3.31}$$
$$= L\Sigma_X L^\top \tag{3.32}$$

□

The proof in (3.26)-(3.32) provides a framework for sampling from a multivariate normal distribution. The sampling process is initiated by sampling $X \sim \mathcal{N}(0, I)$, where I is the covariance matrix Σ_X, which is an identity matrix. It is known that Y can be sampled by sampling X with an affine transformation applied. Thus, $\Sigma_Y = LIL^\top = LL^\top$, and $\mu_Y = w$ as μ_X follows the properties of a standard normal distribution and is set to zero. The Cholesky decomposition can be used to find L. Algorithm 1 was derived from the notes by Roelants [22], and can be implemented to generate samples from a multivariate normal distribution. Testing in Python revealed that an implementation for sampling based on the above approach rather than using the `numpy.random.multivariate_normal` library resulted in significant improvements in computation time. Results from computational experiments can be found in Appendix A.1.

Lastly, the demand parameter must take an integer form for most Network Design problems, as modeling partial products is not always practical. This posed a challenge, as the generated multivariate normal distribution provides the mechanism to sample continuous variables. For this study, it was found that rounding the sampled values to the nearest integer did not significantly alter the properties of the distribution; *i.e.,* the rounded sample had approximately the same mean and variance as the continuous sample.

Algorithm 1 Sampling from Multivariate Normal Distribution [22]

Require:

`d`: Number of dimensions

`mu`: `d`-dimensional mean vector

`covariance`: $d \times d$ *p.s.d.* covariance matrix from (3.25)

`n_samples`: Number of samples to draw

1: Create L: Lower triangular matrix obtained from Cholesky decomposition of `covariance`: $L =$ Cholesky(`covariance`)

2: Sample an array of size (`d`, `n_samples`) represented as K from X, $X \sim \mathcal{N}(0, I_{dd})$, where I_{dd} is an identity matrix of size `d` $\times$ `d`

3: Apply affine transformation to obtain Y: $Y = L \cdot K +$ `mu`

Ensure: Y: Samples from the multivariate normal distribution are strictly non-negative with mean `mu` and covariance matrix `covariance`

3.3.2 Transport Cost Uncertainty

In a similar approach to Section 3.3.1, the per distance-unit transport cost $NC_{t,p,l_{\text{from}},l_{\text{to}}}$ is represented by the random parameter $\tilde{NC}_{t,p,l_{\text{from}},l_{\text{to}}} \sim \mathcal{N}(\mu_{t,p,l_{\text{from}},l_{\text{to}}}, \sigma_{t,p,l_{\text{from}},l_{\text{to}}})$, with $\mu_{t,p,l_{\text{from}},l_{\text{to}}} = NC_{t,p,l_{\text{from}},l_{\text{to}}}$ and $\sigma_{t,p,l_{\text{from}},l_{\text{to}}} = \lambda_T NC_{t,p,l_{\text{from}},l_{\text{to}}}$, where λ_T is the coefficient that controls the magnitude of variance. To a generate representative and reduced set of scenarios, the method for constructing and covariance matrix and sampling from a multivariate Gaussian distribution from Algorithm 1 was used, with $NC_{t,p,l_{\text{from}},l_{\text{to}}}$ as the `d`-dimensional mean vector.

The motivation for the correlation between per-distance-unit transportation costs can be considered more than a convenient approach to attain a reduced set of scenarios. The correlation factor ρ is introduced in (3.25), and it controls the level of correlation that exists between $NC_{t,p,l_{\text{from}},l_{\text{to}}}$ of similar magnitude. It is realistic to assume that when transport costs change along one route, they are likely to change proportionally along other routes, as these costs are strongly driven by parameters such as fuel price. These correlations may be stronger when modeling a supply chain contained in a country or region, while global supply chains may have a weaker correlation between transportation costs due to the variation of fuel prices globally. Additionally, the transportation types may also influence the correlation factor ρ because the change in fuel price may not affect some modes of transport as significantly as others, *e.g.,* rail network costs vary compared to trucking costs. As such, the parameter ρ should be chosen depending on the scale and transportation scheme of the supply chain.

3.3.3 Facility Capacity Uncertainty

The purpose of introducing uncertainty in facility capacities is to model events where disruptions of varying severity prevent facilities from fulfilling their maximum theoretical throughput. Disruptions of significant magnitude could occur with different probability among facilities; furthermore, the expected severity of disruption may also depend on the location of the facility (for socio-economic, political, and geographical reasons). As such, the probability of realizing a disruption and the severity of the realized disruption *could* be modeled as a joint-probability distribution. However, generating a representative and a reduced set of scenarios from a joint-probability distribution could be challenging. The proposed solution to this is to model throughput capacity as a random variable belonging to a known distribution. A requirement for this distribution is that, when sampling from the distribution, there should be a high probability that the realized capacity is not significantly lower than the theoretical maximum capacity. Disruption should be a rare event, although the *rarity* must be adjustable, *e.g.,* in some cases it may be necessary to model more frequent disruptions.

An appropriate distribution, and the one chosen for this study, is the Continuous Bernoulli distribution, given by the notation $\mathcal{CB}(\lambda)$, where $\lambda \in (0,1)$ is the shape parameter. This distribution is supported in [0,1], and Fig. 3.2 shows the continuous distribution function, $F_{\mathcal{CB}(\lambda)}(b)$, for different λ. In this case, the value of $F_{\mathcal{CB}(\lambda)}(b)$ represents the facility capacity. The sampling strategy is based on the notion of inverse transform sampling, which essentially maps uniform random numbers between 0 and 1 to a desired probability distribution by evaluating the cumulative distribution function. It was implemented by first drawing samples for $b \sim \text{Uniform}(b_l, b_u)$, where b_l is the lower-bound value for b and $b_u = 1$ is the upper-bound. Samples from the cumulative $F_{\mathcal{CB}(\lambda)}(b)$ distribution are then given by:

$$F_{\mathcal{CB}(\lambda)}(b) = \begin{cases} b & \text{if } \lambda = \frac{1}{2} \\ \frac{\lambda^b(1-\lambda)^{1-b}+\lambda-1}{2\lambda-1} & \text{otherwise.} \end{cases} \tag{3.33}$$

Another adjustable parameter is the probability of a shutdown lasting for an entire strategic decision-making period. That is, it should be possible to control the maximum expected disruption for a facility, *e.g.,* a drawdown in the capacity of 50%, as it is not certain when (within the modeled period) the failure occurs. It may be that a facility operates at full capacity for a portion of the period before a catastrophic failure occurs, preventing any more products from flowing through the facility. Even in this case, despite a catastrophic failure, the facility still provides some throughput during the period. Therefore, a theoretical minimum throughput (as a proportion of maximum throughput capacity) was instated. This theoretical minimum capacity (because of a failure) corresponds with $b_l = 0.1$, where $F_{\mathcal{CB}(0.001)}(b_l) \approx 0.5$. A realization of capacity as 0.5 can be interpreted as the following: for a given period and facility, the maximum allowed throughput is limited to half of the facility's actual maximum throughput, this could be due to a sustained reduction in capacity during the period, or due to a single catastrophic event that prevented any throughput after some time had elapsed, among other potential scenarios. The shape parameter $\lambda = 0.001$ was selected for this study, however, it is important to note that the choice of λ could vary for different supply chain networks and facilities. For example, if facilities have extremely low probabilities of experiencing any form of failure, one might choose $\lambda < 0.001$. If the probability of experiencing failure during a given period is high, then $\lambda > 0.001$ may be appropriate. The choice of λ has a direct influence on the likelihood of capacity realizations substantially lower than the theoretical maximum capacity. The histogram in Fig. 3.3 shows the distribution of a sample size of 1000 in terms of the resulting failure severity (where the failure severity is the proportion of throughput that is allowed for a given scenario).

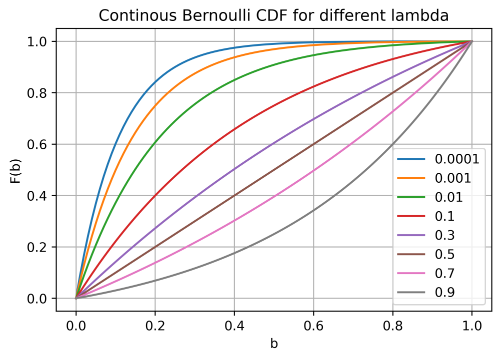

Figure 3.2: CDF for Continuous Bernoulli distribution for different λ.

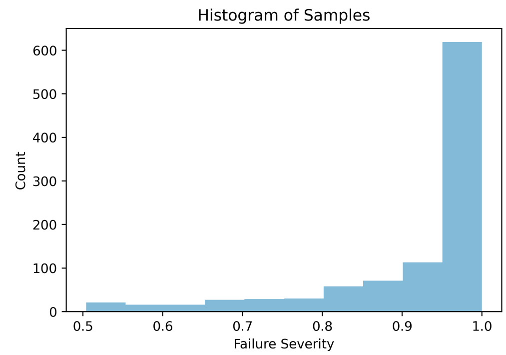

Figure 3.3: Histogram showing 1000 samples from $\mathcal{CB}(0.001)$ with a lower limit of ≈ 0.5

3.3.4 Generating Scenarios from Independent Samples

Random samples of length N were generated for each uncertain parameter (with a total of M parameters) and were stacked into a $M \times N$ matrix denoted as S. Each column of the matrix S represents a scenario ξ_n, such that there are a total of N scenarios consisting of M independent realizations of uncertain parameters. Based on the principals of SAA described in Section 3.2.1, the quality of the solution should improve (get closer to the true objective value) exponentially with larger N, thus, it is possible that a reasonable approximation of the true objective could be achieved with a computationally manageable choice of N. Fig. 3.4 provides a schematic demonstration of how scenarios are formed, where each s_{mn} represents a random sample $n = 1, \ldots, N$ from uncertain parameter $m = 1, \ldots, M$. From the resulting matrix, a scenario ξ_n is formed by taking column-wise combinations of independent random samples, creating a set of scenarios $\boldsymbol{\xi} = \xi_1, \ldots, \xi_N$

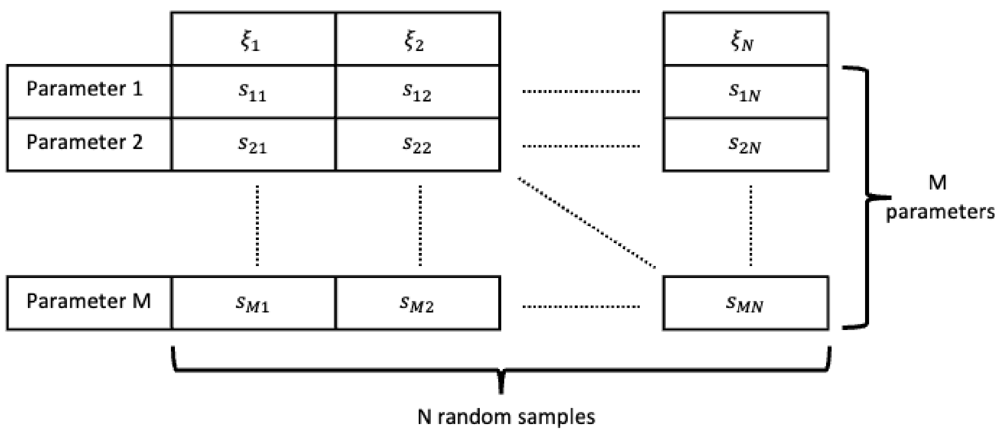

Figure 3.4: Method of combining independent samples into scenarios.

Chapter 4

Experiments

At the core of any SCND mathematical program should be a set of business questions that need to be answered. These business questions (such as which facilities should be used, which facilities should serve certain customers, which transport lane should be used between two facilities, etc.) can prompt the type of data that is required to answer the appropriate questions. In this section, the methods described in Section 3 were applied to a real set of SCND business questions with corresponding data. In Section 4.1 the business case and questions are unpacked, after which Section 4.2 describes the data input for the stochastic programming model. A set of experiments are conducted in Section 4.3, and the resulting network configurations are presented in Section 4.4. Section 4.5 is focused on evaluating the robustness of the network configurations, and Section 4.6 investigates the effect of correlation on the solution.

4.1 Case Study: *Mystery Merchants*

In this case, the supply chain network belongs to *Mystery Merchants*, who operate their SC network solely in The Netherlands. The company name is an alias used for confidentiality reasons. *Mystery Merchants* distributes products to a large number of stores in The Netherlands, and does this via a two-tiered DC structure. The first tier is responsible for processing products from suppliers, and the second tier is responsible for distributing products to customers. *Mystery Merchants* would like to determine: (1) which DCs should be open; (2) which products should be processed in each DC; and (3) whether it is possible to reduce the aggregated driving distance. The Company has noticed that errors in demand forecasting have resulted in unforeseen expenses in transportation and processing, and they would like to design their network to handle fluctuations in demand without making changes to the network structure which lessens the effect on their operating expenses. The Company has also determined that estimates for per-kilometer transportation costs have been inaccurate, especially due to external events that have caused volatility in oil prices. As such, it would be preferable to design an SC Network that is fairly robust to fluctuations in per-kilometer transport costs. Lastly, The Company suffered from various disruptions at their facilities - primarily due to labor shortages and subsequent strikes. They would like to optimize their network such

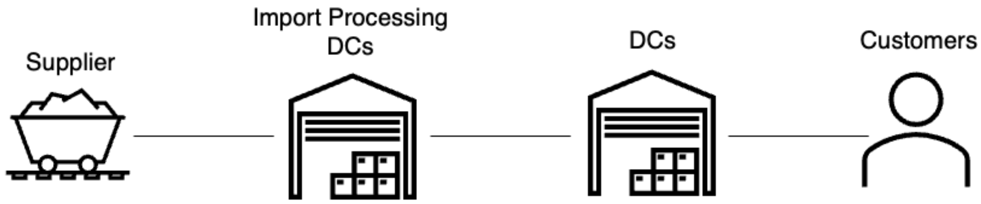

Figure 4.1: Basic SC Network schematic for *Mystery Merchants*.

that facility disruptions do not have such a severe impact on the entire SC Network and its ability to meet service-level agreements. It would be desirable for *Mystery Merchants* to have a more robust SC Network that provides them with the ability to meet customer demand even when estimates of demand, transport cost, and facilities are not consistent with forecasts.

Figure 4.1 shows a schematic of *Mystery Merchants*' SC network, where a supplier provides products that enter the SC Network and flow into Import Processing DCs. From this point, products are sent to DCs which are responsible for distributing the products to customers. Note, there is no production and only a single supplier. In reality, *Mystery Merchants* uses many suppliers, but for modeling purposes, it is sufficient to assume a single supply source because most suppliers are outside of The Netherlands. Note also that one transport lane exists between each node in the network, as transport lane decisions are not a priority for *Mystery Merchants*.

Mystery Merchants does not have data to estimate the probability distribution of their customer demand patterns, per-kilometer transport costs, nor to determine the probability of failure events among facilities. As such, demand expectations and per-kilometer transport costs fitted to a Normal distribution, and the proposed Continuous Bernoulli for facility capacity can be considered sensible choices for *Mystery Merchants* to address their business questions.

4.2 Data Description

As *Mystery Merchants* does not have production facilities in their SC network, data was provided for all supply, warehousing, and distribution parameters from Table 3.1. This data can be grouped into four core components: (1) facility data, (2) product data, (3) customer data, and (4) transportation data. The data associated with each of these components is discussed in this section.

4.2.1 Facility Data

Mystery Merchants has nine distribution facilities, of which three are responsible for processing products that arrive from suppliers (CDCs), while the remaining six facilities are used for distributing products to customers (LDCs). Table 4.1 indicates which cities the DCs are located in, as well as their capacity and costs. Figure 4.2 shows the geographical distribution of facilities within The Netherlands. The two

Figure 4.2: Geographical distribution of DCs within The Netherlands. The three purple locations are CDCs, the rest are LDCs.

LDCs with the highest capacity are located in Amsterdam and Utrecht, which are fairly central in terms of population density for The Netherlands. Outlying LDCs such as Groningen, Tilburg, and Eindhoven have the lowest throughput capacity. The CDCs are strategically placed in logistics and import hubs. Note that the throughput variable cost is averaged overall products for a given facility. The fixed cost refers to the cost of opening and operating a facility for a period, regardless of the volume and type of products that flow through the facility.

DC Type	City	Capacity (products)	Fixed Cost (€)	Variable Cost (€/product)
LDC	Amsterdam	10,000,000	3,825,000	0.51
CDC	Rotterdam	6,000,000	1,080,000	0.6
CDC	Schiphol	6,000,000	1,080,000	0.6
CDC	Breda	8,000,000	656,250	0.6
LDC	Utrecht	10,000,000	1,579,500	0.48
LDC	The Hague	5,000,000	1,080,000	0.516
LDC	Groningen	3,000,000	1,080,000	0.516
LDC	Eindhoven	3,000,000	1,080,000	0.516
LDC	Tilburg	3,000,000	1,080,000	0.516

Table 4.1: List of CDCs and LDCs in different cities with capacity, fixed cost, and variable cost.

4.2.2 Product Data

Mystery Merchants distributes more than 1000 products. Each product belongs to a department, and there are a total of 20 departments. For strategic decision-making, it is considered sufficient to model the network based on departmental product flow, rather than modeling the network on the individual product level. Table 4.2 shows the demand for each product, where each row represents the product demand for a different department. Note that the demand is expressed in terms of product-department units, where each unit corresponds to a product within its respective department, individual products are not distinguished.

Product	Expected Demand (department level units)
D1	818,759
D2	431,536
D3	760,965
D4	973,372
D5	757,060
D6	443,856
D7	589,575
D8	477,959
D9	549,151
D10	480,930
D11	387,674
D12	404,024
D13	386,340
D14	514,936
D15	686,330
D16	419,676
D17	504,213
D18	677,175
D19	479,214
D20	660,951

Table 4.2: List of department products and their corresponding total demand.

4.2.3 Customer Data

Customers are aggregated such that each *customer* is represented on the town or city level. Consequently, cities and towns with more stores would generally have more demand than those with few stores. There are 400 customer locations on the aggregated level. Figure 4.3 shows the location and relative demand of each customer in The Netherlands (portrayed by the *size* of the location marker). Demand is predominantly in the Randstad region (Amsterdam, Rotterdam, Utrecht, and The Hague).

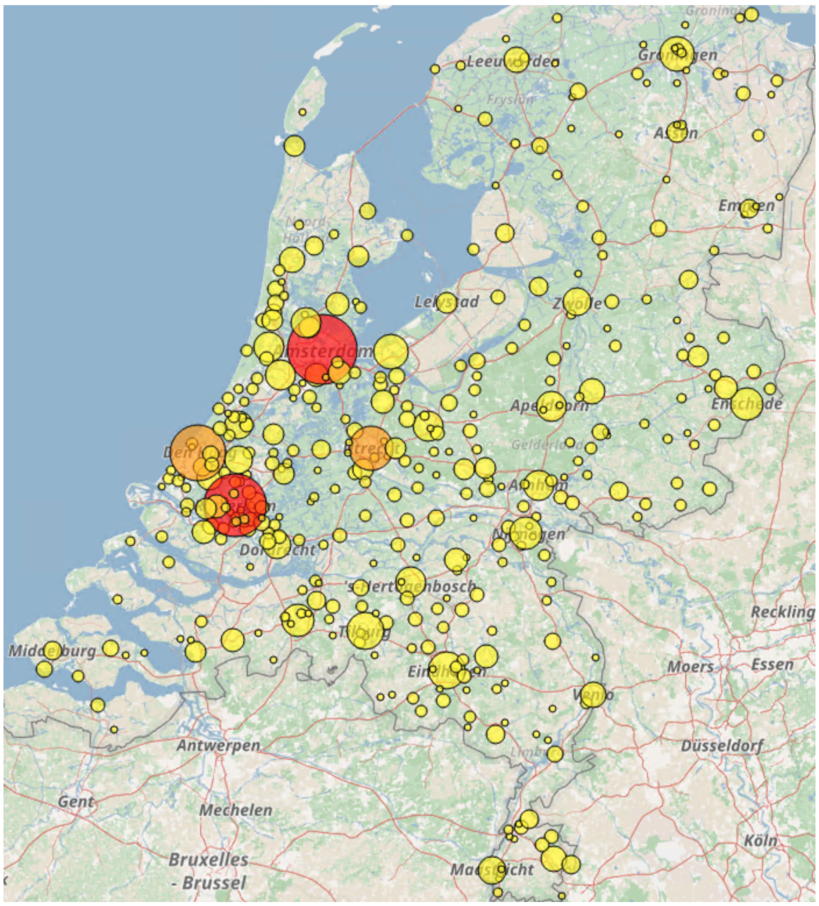

Figure 4.3: Geographical distribution of customers in The Netherlands. Markers are sized and colored proportionally to the demand level (darker colors and larger markers indicate more demand).

4.2.4 Transportation Data

As previously discussed, decisions regarding transportation type are not part of the core business questions that *Mystery Merchants* is investigating. The transportation data involves the product-department per-kilometer cost of transport and links between the *supplier* and each CDC, a fully connected network between CDCs and LDCs, and again a fully connected network between LDCs and customers. The per-kilometer transport cost for each product department is the same for each route, *i.e.,* per-kilometer distribution costs are equivalent to inter-resource and sourcing per-kilometer costs. In all cases, there exists only one form of transport between nodes in the network.

4.3 Computational Results

In this section, the methods proposed in Section 3 are implemented on *Mystery Merchants* data, and the computational results are presented, specific reference is made to solving time. In Section 4.3.1, an approach to simplify *Mystery Merchants*' model is proposed, as this could lead to lower computational cost at the expense of decision granularity. In Section 4.3.2 Bender's decomposition is also compared with the *monolithic* (standard) stochastic program solving approach in terms of computation time.

4.3.1 Clustering Experiment

As discussed in Section 2, a drawback of stochastic programming is the computational complexity and the long solving times. The result is a trade-off between the solution quality (which improves with scenarios) and solving time. This study proposes an artificial reduction in problem size as an approach to reduce computation time and increase the number of scenarios that can be considered. Much of the complexity in a stochastic programming problem can be attributed to the recourse decisions (the second stage in this case). In this context, the recourse decisions are focused on product flow through the network. To simplify the product flow, customers can be grouped into clusters. The clusters were formed on geographical position using the `KMeans` algorithm by `sklearn`. The number of clusters K is the parameter that controls the number of customer clusters that are generated. As K increases, the objective value should converge to the solution found without clustering. Figure 4.4 shows a comparison between K and the corresponding *delta* hereafter δ_c. The value for δ_c was calculated as the absolute difference between transport costs in the clustered solution compared with the full solution (no clustering) and expressed as a percentage. It is clear that δ_c decreases as K becomes larger and the clustered problem begins to converge to the full problem. The rapid decrease of δ_c could indicate that a high-quality surrogate of the full problem could be approximated using relatively few clusters. Further motivation for using a clustered approach for high-level decision-making is given by Figure 4.4, which also shows how solving time increases with K at a polynomial rate. The cost breakdown, solving time, and problem size (in terms of constraints and variables) for different K can be

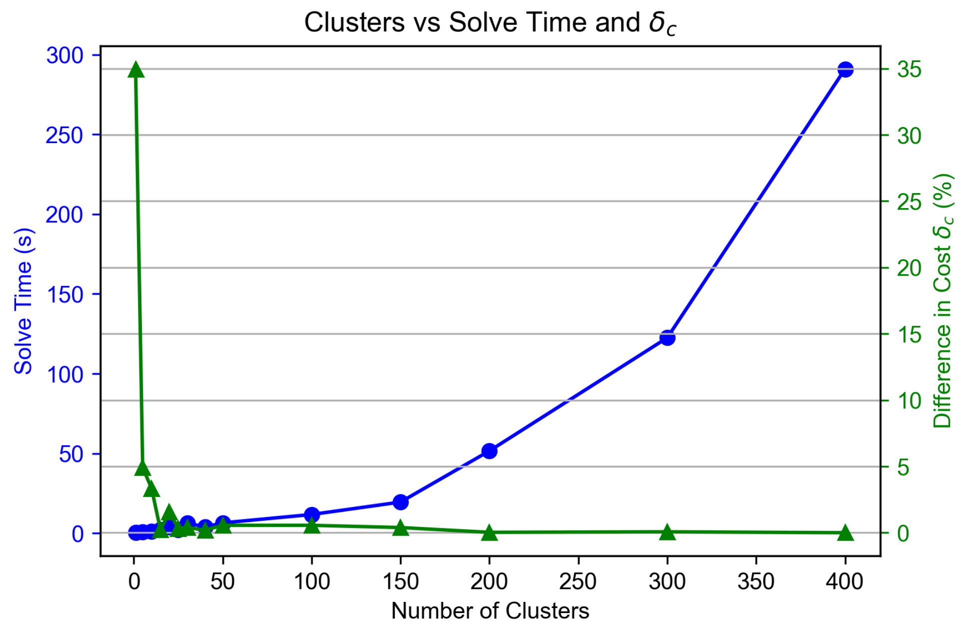

Figure 4.4: Comparison of δ_c, solve time, and number of clusters.

seen in Appendix A.2. Transport costs are the main source of differing objective values between clustered solutions and unclustered solutions.

4.3.2 Algorithmic Approach

Section 3.2.2 outlined Bender's decomposition which is a well-known approach to increasing the rate of convergence for stochastic programming problems. The Benders decomposition was applied to the two-stage stochastic program from Section 3.2 using *CPLEX*'s Benders algorithm [28]. An experiment was set up to observe the difference in solving time for identical problems with and without Benders decomposition. In this context, the approach without using Benders is referred to as the *monolithic* algorithm.

The results from Section 4.3.1 indicate that it may be worthwhile reducing the size of the problem using customer clustering. It follows that the computational performance of Benders should be compared with the monolithic variant based on the number of scenarios N *and* on the number of clusters K. It was observed that the Benders approach converged relatively fast to an integer solution within a reasonable relative optimality gap δ_{gap}; however, the rate of converge slowed dramatically for small δ_{gap}. Note that δ_{gap} is the difference between the best-known integer solution and the linear relaxation given as a proportion of the best-known integer solution. Santoso et al. [23] reported this "tailing off behavior" and found that it was due to almost all incumbent solutions having the same major configuration, which in this case consists of facility-related binary decisions. The minor configuration of the solutions has a smaller influence on the objective, and thus the upper bound does not change significantly between iterations. Furthermore, it may also be the case that the optimal solution has been found even when δ_{gap} is larger, but the best LP bound still needs to close in on the MIP solution. Thus, the solving process could be terminated when δ_{gap} is sufficiently small. In general, industry experts consider $\delta_{\text{gap}} = 0.00001$ to be sufficiently small for the optimal solution in deterministic problems, even when minor configuration decisions form part of important business questions. When the SC Network Design models are particularly large, difficult to solve, or time to a solution is important, a gap of $\delta_{\text{gap}} = 0.001$ is considered adequate. In the context of planning under uncertainty and making strategic SC Network Design decisions, a larger δ_{gap} may be sufficient, particularly when the minor configuration decisions are not primary business questions. For solving the mathematical program in Section 3.2, $\delta_{\text{gap}} = 0.0005$ was used. This ensured that the major configuration decisions would be consistent with a solution to optimality at the expense of marginally sub-optimal minor configuration decisions. Note that for all experiments, a soft cut-off limit was imposed at 30,000 seconds (8.3 hours), and a hard limit at 50,000 seconds. The soft limit was enforced in experiments that included clustering in the following way: for each cluster size, the stochastic program was solved for increasing numbers of scenarios until the solving time exceeded 30,000 seconds, at which point no further experiments were initiated for that cluster, and the experiments with the next cluster size were initiated. Therefore, it may be that solve times exceed the 30,000s limit, however, no subsequent solves would be attempted for that K. The hard limit was the maximum time that the solving process was conducted for, *i.e.,* the solver was interrupted after 50,000 seconds.

Figure 4.5 shows the running time for different combinations of clusters K and scenarios N. This figure indicates how running time scales up with the number of clusters and scenarios. The color gradient in Figure 4.5 is on a logarithmic scale to better visualize the rapidly increasing running times. Note that sometimes it may appear that the soft limit was not adhered to, *e.g.,* K=200, in these cases no solution was found after the hard limit was reached, and thus no additional experiments were attempted for that K

Based on the running time results presented in Figure 4.5, it is clear that (for this problem) the monolithic program is heavily influenced by K and N. The proposed Benders decomposition was implemented to reduce the time to a solution. Figure 4.6 shows a comparison between running time for the monolithic program and the Benders approach for specific K. The Figure shows what appears to be erratic behavior in running time, as one can observe dramatic increases in solve time for $N = 50$, $K = 100$, followed by a reduction in running time when $N = 100$. For $K = 10$, the same behavior can be observed, although the increase when $N = 50$ is less extensive.

To further evaluate the scaling of running time with and without Benders decomposition, the data from Figure 4.6 was transformed to a logarithmic scale and is shown in Figure 4.7. Interestingly, the *scaling-up* behavior appears to be similar when examined on a logarithmic scale, providing additional insight into how running time scales for different K and N when using Benders decomposition.

4.3.3 Sample Average Approximation Quality

Clustering customer locations and using a Benders decomposition appears to permit the evaluation of a large set of scenarios. Despite the ability to solve a large number of scenarios, the number of scenarios that

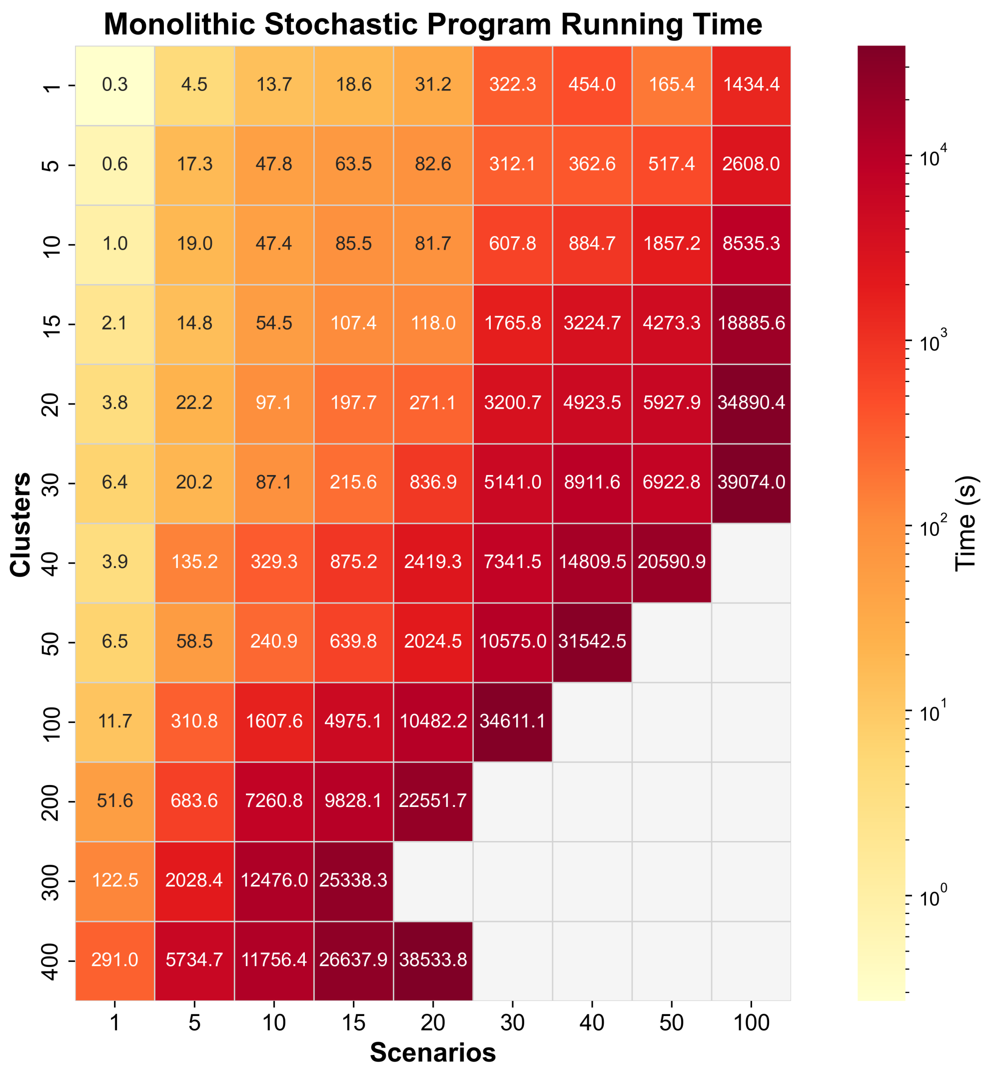

Figure 4.5: Heatmap showing the solving time for different combinations of clusters and scenarios.

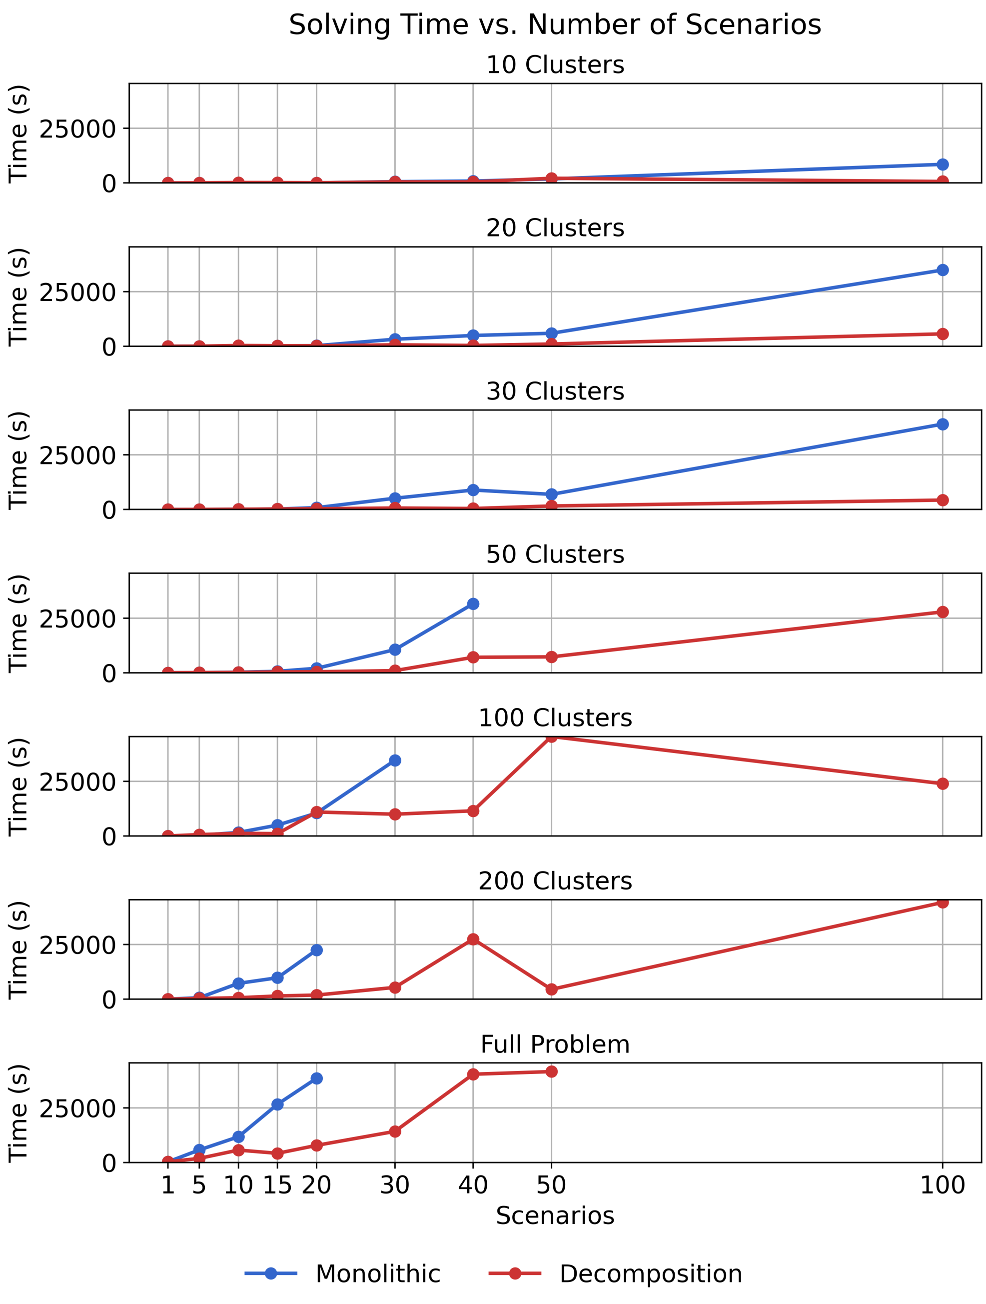

Figure 4.6: Comparison between solving time for different K.

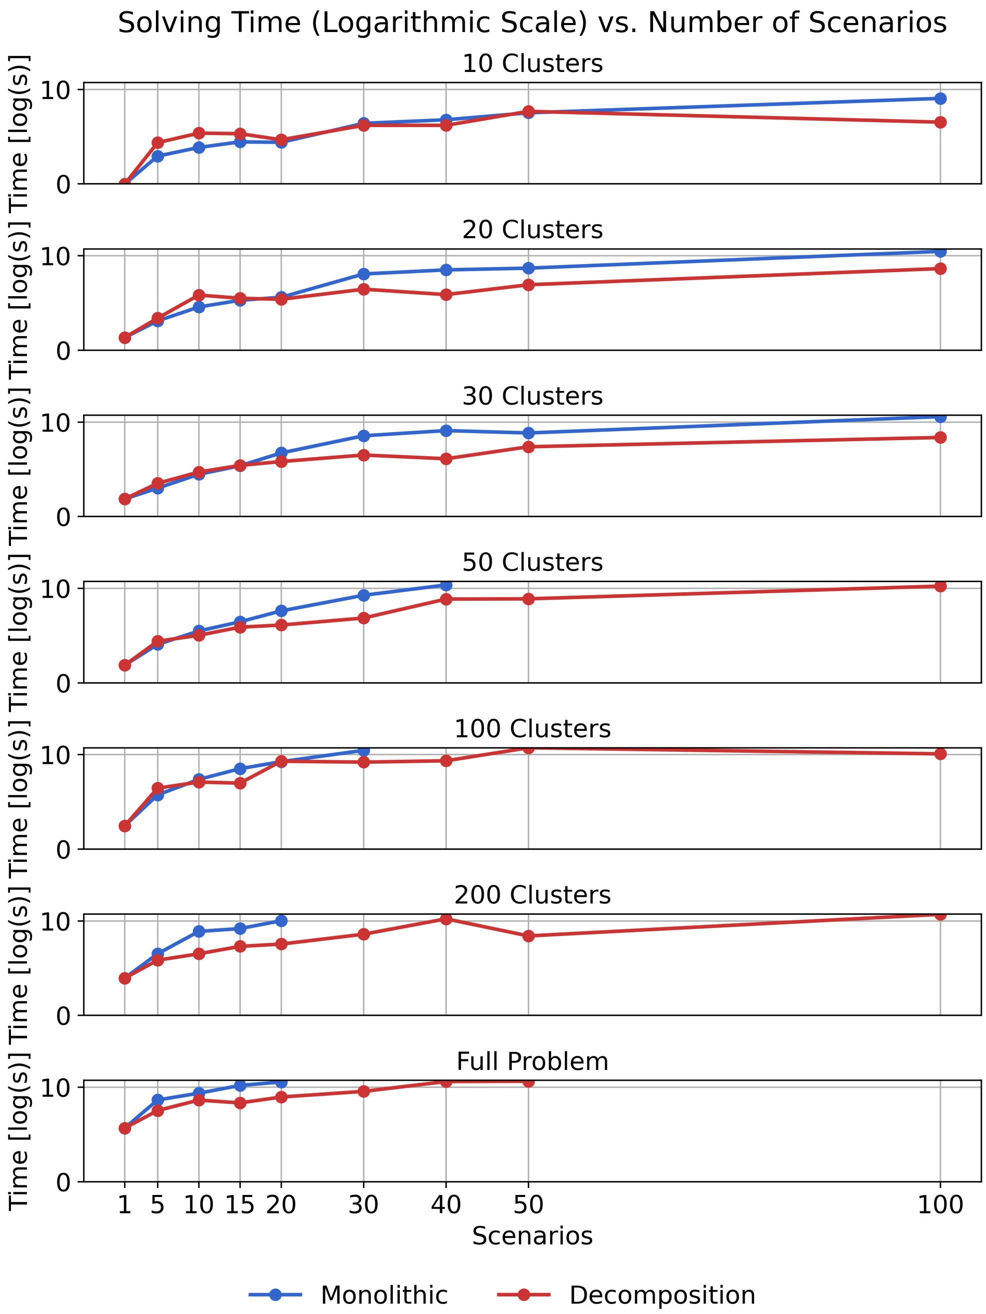

Figure 4.7: Comparison between solving time for different K on logarithmic scale.

are necessary for convergence to a consistent, optimal first-stage decision remains unclear, *i.e.,* how many scenarios should be evaluated for a good approximation of the objective function. It was assumed that the first-stage decisions related to the experiment with the highest solvable (based on the time limitations discussed earlier) K and N values represent the *optimal* first-stage decisions. Thus, the *optimal* first-stage decisions were those associated with $K = 400$ (the full problem) and $N = 50$. The assumption was formulated based on a mechanism behind SAA (Section 3.2.1), whereby the approximation of the objective function improves exponentially as N increases. It is also expected that an increase in K represents a more realistic solution, although the experiments in Section 4.3.1 show that smaller K values can provide a very good approximation of the full problem. The notation of *Similarity* between an incumbent solution and the optimal solution for different K and N can be given by the Jaccard similarity coefficient J:

$$J_i = \frac{|A_i \cap B_i|}{|A_i \cup B_i|}. \tag{4.1}$$

The Jaccard coefficient J_i represents the intersection between the set of optimal first-stage decisions and the incumbent set of first-stage decisions measured and given as a proportion of the total number of elements in the optimal solution set and the incumbent solution set. In (4.1), A_i is the optimal solution set for variable i, B_i is the incumbent solution set for variable i, and $i \in \{DCB_{r,l}, DCBP_{p,r,l}\}$. According to the principles behind the SAA algorithm, the stochastic programming solution using SAA improves exponentially with the number of scenarios. Therefore, it is worthwhile to investigate the point at which first-stage decisions remain consistent, as this can indicate how many scenarios are necessary to reach (what is considered to be) the *optimal* solution. Figure 4.8 shows this convergence for different K. By examining Figure 4.8 it is clear that (for most K) the optimal set of facility open/close decisions is reached after solving approximately 30 scenarios. Interestingly, for $K = 40$ the optimal facility open/close decisions are not reached at $N = 30$, or even when $N = 50$. When examining Figure 4.5 there appears to be a delineation between the solving time at $K = 40$ for lower values of N. It may be that the $K = 40$ configuration of customer clusters is particularly difficult to solve.

The same analysis as shown in Figure 4.8 was repeated for the $DCBP_{p,r,l}$ variable. This variable has an open/close decision for each product department and each facility. Figure 4.9 shows how the solutions converge to what is considered the *optimal* configuration as N increases. Note that the convergence is less consistent compared to $DCB_{r,l}$, although it can be observed that $N = 30$ appears to mark the point at which most values of K reach a solution that is close to *optimal*. Interestingly, some of the plots show a slight reduction in similarity after the $N = 30$ point. Despite how scenarios are constructed (when more scenarios are added they are merged with existing scenarios), it should be noted that the results in Figures 4.8-4.9 are not necessarily trivial. Adding more scenarios could still change the optimal network configuration, thus the concept of *converging* to an *optimal* configuration while adding scenarios can provide insights regarding the benefits or drawbacks of solving problems with an (unnecessarily) large number of scenarios.

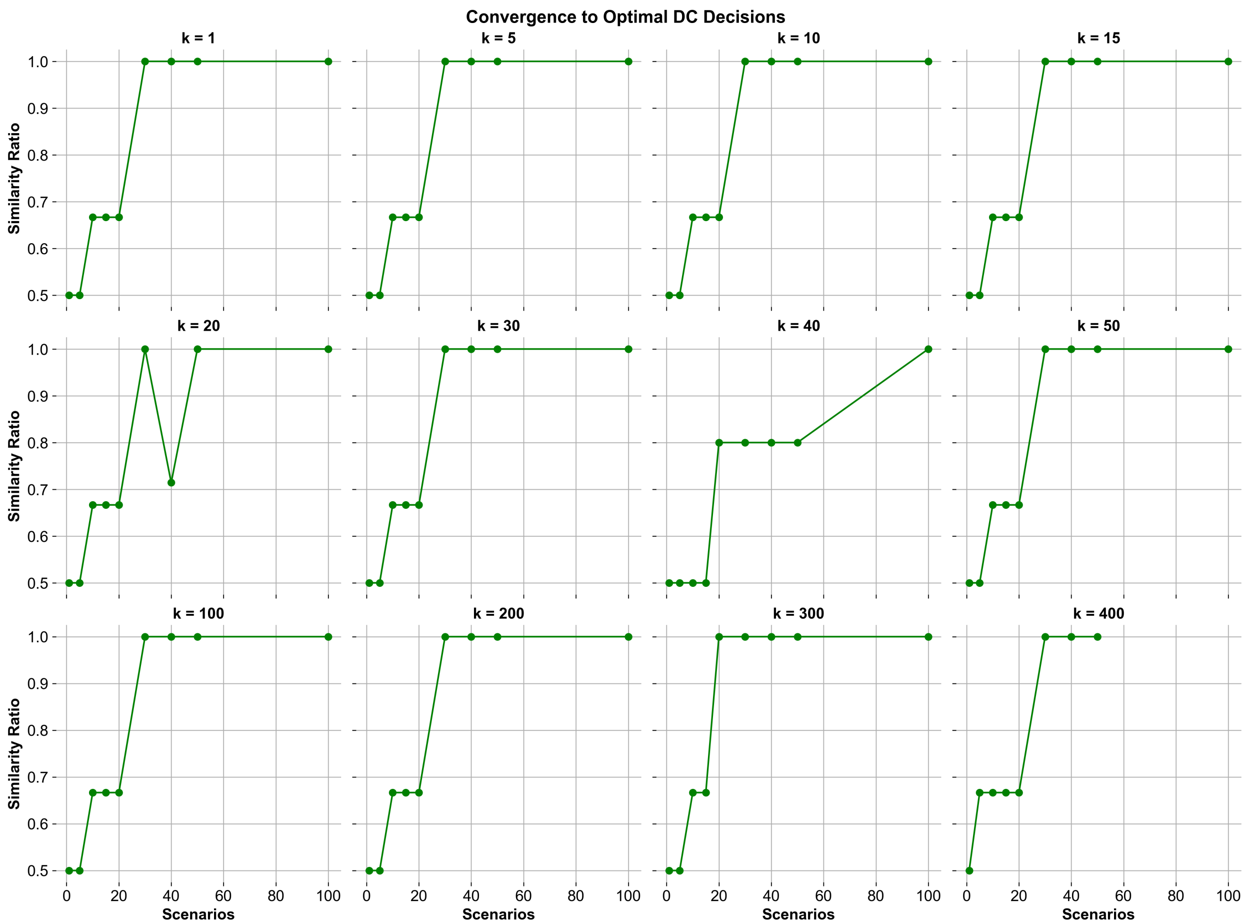

Figure 4.8: Plots showing convergence to the optimal $DCB_{r,l}$ for each K.

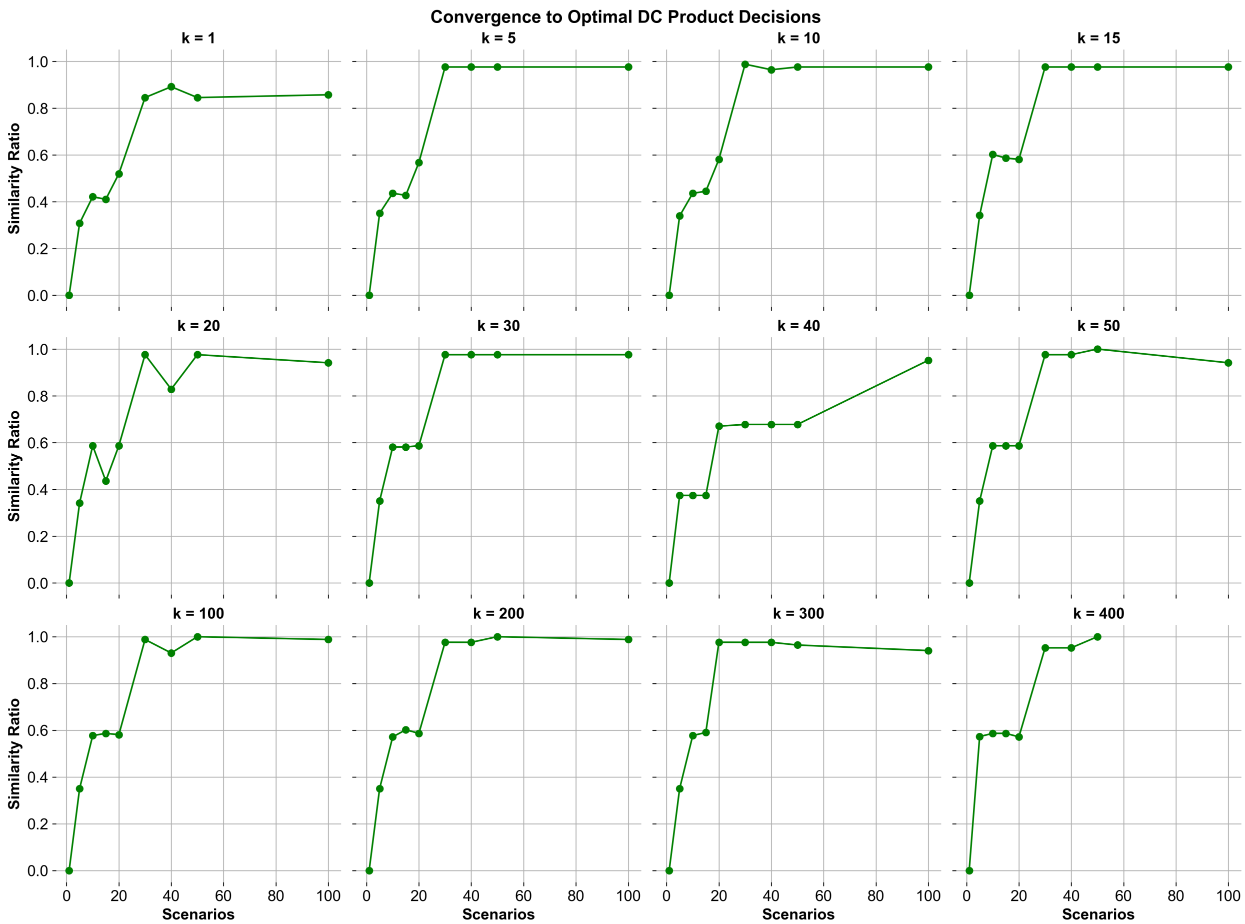

Figure 4.9: Plots showing convergence to the optimal $DCBP_{p,r,l}$ for each K.

Figure 4.10: Deterministic DC configuration (purple CDCs, red LDCs).

Figure 4.11: Stochastic DC configuration (purple CDCs, red LDCs).

4.4 Network Configuration Comparison

The results discussed in Section 4.3 highlighted the impact of clustering, applying a Benders decomposition, and the utility of examining more scenarios in the SAA approach. In this Section, the network configurations that were generated using stochastic programming are compared with configurations generated using a deterministic approach. In particular, the decisions relating to *Mystery Merchants*' core business question are addressed.

4.4.1 Facility Open/Close Decisions

The deterministic problem makes facility open/close decisions based on the expectation that no disruptions occur during the strategic decision horizon. Furthermore, it is expected that demand and transportation costs do not differ from their expected values. The resulting configuration is shown in Figure 4.10. The stochastic programming configuration (using the same *optimal* stochastic solution from Section 4.3.3), is shown in Figure 4.11. Interestingly, the solution in Figure 4.11 makes use of the two largest LDCs (in terms of maximum capacity) and all three CDCs. The deterministic result in Figure 4.10 makes use of only two of the CDCs, as well as one *large* and one *small* LDC.

4.4.2 Department-DC Assignment

Mystery Merchants' second core business question was which DCs should department products be assigned to. It may be important to note that this problem was constructed based on department-level product aggregation, so when multiple facilities are assigned products from the same department, it does not necessarily mean that all products from that department have been assigned to the facility. It follows that decisions on a tactical or operational level should still be made to determine which products should be distributed through each facility. Nonetheless, departmental-product assignment decisions can be important on a strategic decision level. The department assignment decisions are detailed in Table 4.3. It can be observed in the deterministic solution that there are several departments whose products are assigned to only one CDC and one LDC. This means that the department may not meet demand if either of the facilities is disrupted. The stochastic optimal solution reveals that each department is assigned to at least two CDCs and two LDCs (except department 1 which is only assigned to one CDC). This builds redundancy in the network and means that service level failure cannot occur from a single point of failure. The stochastic solution is inherently more robust than the deterministic one, as a single point of failure cannot severely disrupt the ability of the network to meet customer demand. Furthermore, one of the facilities (Breda CDC) is opened by has very low utilization by product departments. This means that there is another point of redundancy in the network that could (under extreme circumstances) carry additional products, should the other CDCs be severely disrupted.

Dept.	**Deterministic**				**Stochastic**				
	CDC		LDC		CDC			LDC	
	Breda	Schiphol	Tilburg	Utrecht	Breda	R'dam	Schiphol	Utrecht	A'dam
D1	1	1	1	1	0	0	1	1	1
D2	1	1	1	1	0	1	1	1	1
D3	1	1	1	1	0	1	1	1	1
D4	1	1	1	1	1	1	0	1	1
D5	1	0	1	1	0	1	1	1	1
D6	1	1	1	1	0	1	1	1	1
D7	1	0	1	1	0	1	1	1	1
D8	1	1	1	1	0	1	1	1	1
D9	1	1	1	1	0	1	1	1	1
D10	1	1	1	1	0	1	1	1	1
D11	1	1	1	1	1	1	1	1	1
D12	1	0	1	1	0	1	1	1	1
D13	1	1	1	1	0	1	1	1	1
D14	1	1	1	1	0	1	1	1	1
D15	1	1	1	1	0	1	1	1	1
D16	1	1	1	1	0	1	1	1	1
D17	1	1	1	1	0	1	1	1	1
D18	1	1	1	1	0	1	1	1	1
D19	1	0	1	1	1	1	1	1	1
D20	1	0	0	1	0	1	0	1	1

Table 4.3: Product department assignments to DCs for the deterministic and stochastic problem, a 1 indicates that the product department is assigned to that DC, otherwise a 0 is given.

4.5 Evaluating the Robustness of the Deterministic Solution

The purpose of this section is to evaluate the robustness of the deterministic solution by solving it under different scenarios. The optimal solution under a deterministic environment may be sensitive to even small changes in some parameters. For example, a deterministic solution may propose running each (or a subset) of the facilities near-maximum capacity, as this is generally the most cost-effective utilization of resources. It follows that even a small change in demand, or a disruption at a facility, could result in the inability to meet demand. It should be noted that the influence of facility capacity constraint presents itself twofold: (1) uncertainty regarding capacity can realize reduced capacities as disruptions even with lowered demand, (2) increased demand could exceed even the theoretical maximum capacity (without disruption). It is therefore expected that the influence of lower capacity within scenarios would have a significant effect on the feasibility of solutions.

4.5.1 Method for Measuring Robustness

A methodology that involves evaluating the deterministic solution under different scenarios was developed to assess the robustness of these solutions. The variables corresponding with the first stage of the stochastic program are particularly important, as these cannot be adjusted based on uncertain realizations. The proposed approach allows for the performance of the deterministic solution to be gauged under conditions of uncertainty. Furthermore, it can provide insights into the adaptability and reliability of the deterministic solution when compared with the stochastic solution. The steps followed in this analysis are: (1) solving the SCND problem deterministically and retaining the first-stage solutions; (2) solving the SCND problem with fixed first-stage variables for a set of scenarios; (3) collecting the resulting costs and whether or not the problem was feasible; (4) analyze data extract insights regarding the influence of uncertain parameters on the costs and feasibility of solutions; and lastly (5) a cost breakdown between the stochastic and feasible deterministic results. Additionally, steps (1)-(4) were repeated except the first-stage decisions from step (1) were not enforced in subsequent scenario-dependent deterministic solves, which allows the gathering of scenario-based insights under perfect information.

The proposed methods allow for a systematic evaluation of the robustness of the deterministic solution and provide insight into the need for addressing uncertainty in the design phase.

4.5.2 Deterministic Solution Robustness

Solving the deterministic SCND problem revealed that is optimal to open four facilities, two CDCs, and two LDCs. In the case of the LDCs, one large facility and one smaller facility were opened (`Utrecht_DC` and `Tilburg_DC`, respectively). The two CDCs (`Breda_DC` and `Schiphol_DC`) . The first-stage variables (DC open/close and product-DC allocations) were then fixed, and the problem was solved independently for each

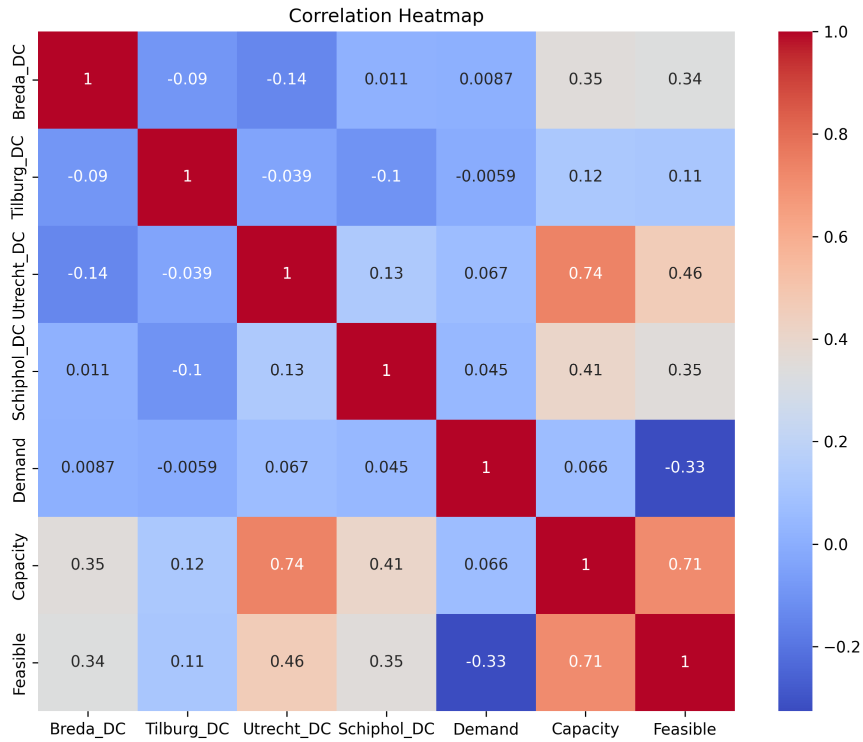

Figure 4.12: Correlation heatmap for analyzing the drivers of feasibility.

scenario generated as per Section 3.3. For each independent solution, the costs (from the objective function) were recorded and infeasible solutions were flagged.

A high-level analysis on cases $K = 20, 100, 200$ (each with 100 scenarios) found that approximately 30% of scenarios induced infeasibility when using the configuration from the deterministic solution. To further investigate this, the relationship between scenario properties and the feasibility of the problem is examined. Each scenario was associated with the feasibility status, where the scenario was decomposed into facility-level capacity coefficients, the overall capacity of the facilities, and total demand. The overall capacity was calculated by taking the minimum combined capacity for CDCs and LDCs, respectively. Being able to identify the scenarios (or properties of scenarios) that are associated with infeasibility allows a better understanding of the limitations of the deterministic solution. Figure 4.12 shows the correlation between the decomposed scenarios and the feasibility of the problem. That is the capacity for each DC, the total demand, and the maximum capacity of the network.

The capacity for each DC (`Breda_DC`, `Tilburg_DC`, `Utrecht_DC`, `Schiphol_DC`) is a value in [0,1] as per the sampling description in Section 3.3.3. The `Demand` entry is the total demand for each scenario (summation for all customers, all products). Finally, the `Capacity` entry is the maximum capacity of the network under

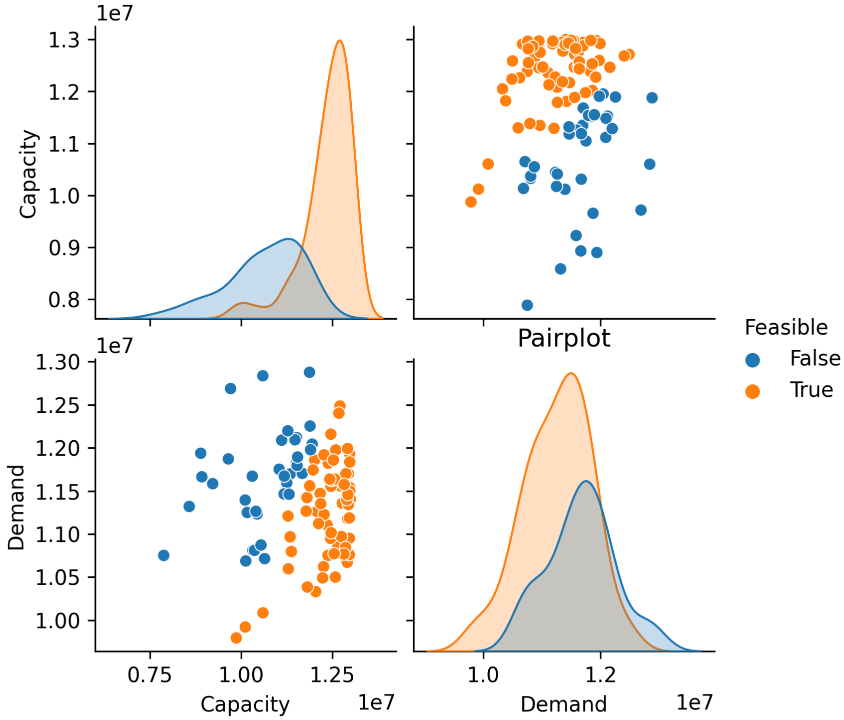

Figure 4.13: Pairplot comparing problem feasibility in terms of Demand and Capacity.

a given scenario, which is calculated by taking the lesser of combined LDC capacity or CDC capacity. In Figure 4.12, the data was generated by solving the deterministic problem with $K = 200$ and $N = 100$; the experiment was repeated for various cases of K but the results given in Figure 4.12 are representative of the trends. It is clear that `Capacity` has the highest correlation with the `Feasible` variable. `Demand` shows a negative correlation with feasibility, which is expected, as when demand increases it is more likely that customers will not be serviced. Interestingly, the `Utrecht_DC` variable is highly correlated (0.74) with `Capacity`, which indicates that this DC heavily influences the overall capacity of the network. Similarly, the variable `Utrecht_DC` also shares a significant correlation with `Feasible`, indicating that the capacity of `Utrecht_DC` could have a large influence on the feasibility of the problem.

To further investigate the relationship between `Capacity`, `Demand`, and `Feasible`, Figure 4.13 was generated. Figure 4.13 shows the relationship between `Capacity` and `Demand` in terms of feasibility. The density peaks of the `Capacity` feature exhibit distinct separation, with lower capacities being predominantly associated with infeasible solutions. In contrast, the distributions of `Demand` show a less pronounced difference between infeasible and feasible scenarios, although higher demand values are more commonly observed in the infeasible cases.

The distribution of realized capacity coefficients between infeasible and feasible solutions can be seen in Figure 4.14. This figure shows additional evidence that the feasibility of a deterministic solution is strongly associated with instances with higher capacities.

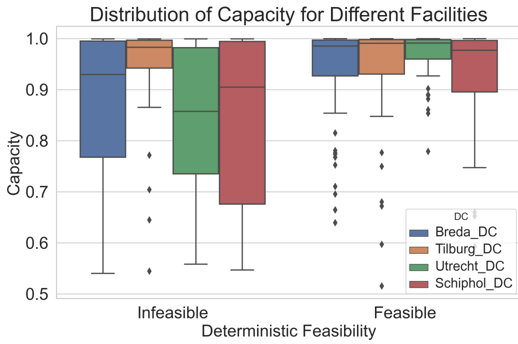

Figure 4.14: Capacity distribution and feasibility between different facilities.

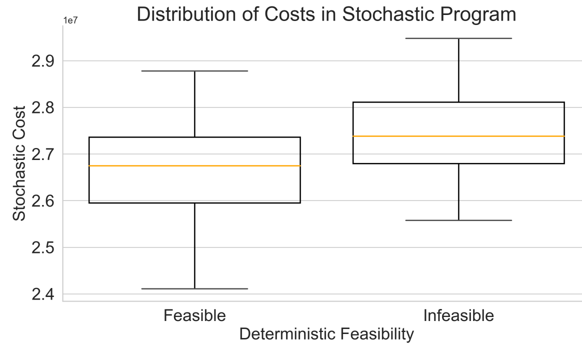

Figure 4.15: Distribution of the stochastic objective value between feasible and infeasible deterministic solutions.

There appears to be a relationship between feasibility and both the `Demand` and `Capacity` variables. The relationship between the objective function and feasibility remains unexplored, although this is challenging because infeasible solutions in the deterministic problem lack an associated cost. Nonetheless, comparing the objective of the stochastic problem with the feasibility of the deterministic problem could provide valuable insights. Specifically, it can indicate if the infeasible deterministic solutions are associated with a particularly costly subgroup of solutions. Figure 4.15 shows the distribution of the stochastic objective value for feasible and infeasible solutions. Note that the costs appear to be higher in the *Infeasible* cases.

Thus far, relationships within the decomposed scenarios have been analyzed with respect to their influence on feasibility within the deterministic solution. Although this is an important aspect to consider, it may also be useful to understand the impact that uncertainty has on the objective value. When considering the stochastic program results, a Pearson Correlation Coefficient of 0.992 was calculated between `Demand` and the objective value. Similarly, the same correlation for the deterministic problem was found to be 0.996. The correlation coefficient for the stochastic program was recalculated using only datapoints that were associated with deterministically feasible solutions and the result was approximately 0.996, in line with the deterministic problem. Facility capacity had a lower correlation with the objective value, at 0.43 and 0.41 for the stochastic and deterministic problems, respectively (when adjusted to exclude infeasible data points). Figures 4.16-4.17 shows the relationship between the objective value and both `Demand` and `Capacity`. It is evident that `Demand` has a strong positive correlation with the objective; however, the correlation observed between `Capacity` and the objective does not appear well distinguished.

Based on Figures 4.16-4.17, it is clear that the cost of the stochastic solution is systematically higher than the deterministic variant. This observation could be considered trivial, as the stochastic solution has higher fixed costs due to having more facilities. Figures 4.18-4.19 show the distribution of handling and transport costs, respectively, for both the deterministic and stochastic problems.

There is not a large difference between these costs between deterministic and stochastic problems the

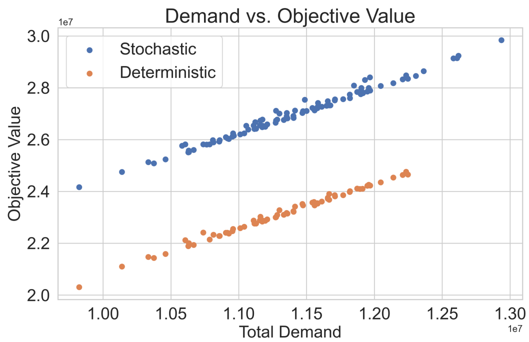

Figure 4.16: Relation between demand and the objective value.

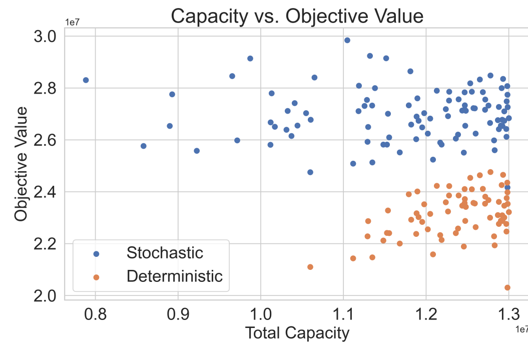

Figure 4.17: Relation between capacity and the objective value.

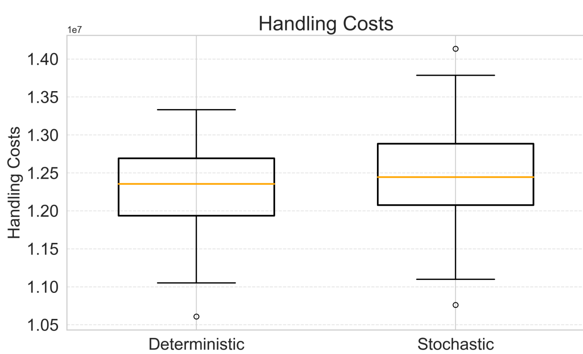

Figure 4.18: Distribution of handling costs for stochastic and deterministic problems.

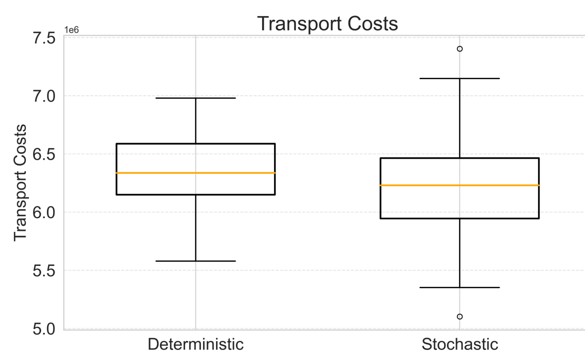

Figure 4.19: Distribution of transport costs for stochastic and deterministic problems.

transportation costs in the stochastic model (Figure 4.19) appear to be lower than the deterministic model. This is likely due to the stochastic solution using more DCs, thus it is likely that the travel distance between customers may reduce in some cases. Handling costs in the stochastic problem appear to be very similar to the deterministic problem (Figure 4.18), although the stochastic problem handling costs are slightly higher.

A final cost comparison was made between the stochastic approach and a deterministic approach, except in this case the deterministic problem was solved without constraining the first-stage decisions. That is, the deterministic problem was solved for each scenario under the assumption of having perfect information. Figure 4.20 compares the objective value, transport cost, and handling cost for the deterministic problem under perfect information with the respective stochastic problem costs. The resulting costs are consistent with the previous findings, where the objective function value is higher for the stochastic problem, handling costs are approximately equal, and transport costs are marginally lower for the stochastic solution.

To formalize the statements made in this section, it is necessary to perform statistical tests. The tests are focused on evaluating whether there is a difference between data associated with *infeasible* or *feasible* deterministic solutions. The two sample t-test was used to test the hypothesis that the means of both samples are equal, under the assumption that the data are normally distributed. In the case that the data are not normally distributed, a Mann-Whitney U non-parametric test was used. The Mann-Whitney U test has a null hypothesis that for samples A and B, a randomly selected datapoint from sample A has an equal probability of being greater than or less than a random datapoint from sample B. The `scipy.stats` package in Python was used to conduct the statistical tests. In all cases, a significance level of 0.05 was used. A Shapiro-Wilks test was used to evaluate whether data can be assumed to follow a normal distribution. It was found that the assumption of normality of the objective function data could not be rejected, thus a t-test can be applied. Likewise, with Demand data, the assumption of normality could not be rejected. The assumption of normality within the Capacity data was rejected, thus the Mann-Whitney U test was applied to this case. Table 4.4 displays the p-values obtained for each test, revealing significant differences in scenario data between feasible and infeasible deterministic problems for all tested parameters, except for `Tilburg_DC`. The results suggest that `Tilburg_DC` does not have a significant influence on feasibility.

Table 4.4: Results from statistical testing between deterministically infeasible and feasible subgroups. Significant (at a level of 0.05) results are indicated by *.

Feature	Test	p-value
Objective Value	t-test	$0.004*$
Capacity	Mann-Whitney U	$3.3 \times 10^{-14}*$
Demand	t-test	$0.019*$
Breda_DC	Mann-Whitney U	$0.048*$
Schiphol_DC	Mann-Whitney U	$0.01*$
Utrecht_DC	Mann-Whitney U	$2.5 \times 10^{-6}*$
Tilburg_DC	Mann-Whitney U	0.5

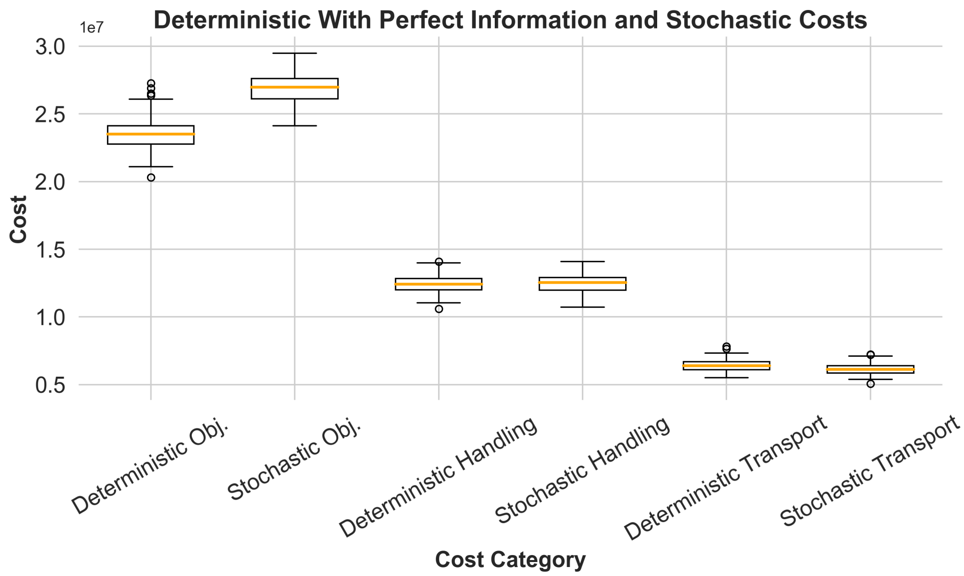

Figure 4.20: Comparison between costs of the deterministic problem under perfect information and the stochastic problem.

Statistical tests were conducted to determine whether the handling and transport costs differ significantly for stochastic and deterministic problems. The hypothesis that the data followed a normal distribution could not be rejected (with a Shapiro-Wilks test) at a significance level of 0.05), and thus the t-test was applied. It was found that transportation costs differed significantly (p-value of 0.015) between deterministic and stochastic results, even when comparing only feasible solutions. In contrast, the handling costs did not differ significantly. Despite failing to reject the null hypothesis that the data (transport costs and handling costs) follows a normal distribution, a Mann-Whitney U test was also conducted. The results of the Mann-Whitney U test reinforced the conclusions from the t-test. The statistical tests were performed again on the case where the deterministic problem was solved without imposing any constraints on first-stage decisions. In this case, it was observed that there was a significant difference in transport costs between the two scenarios, whereas handling costs were found to be approximately equal.

4.5.3 Key Takeaways from Deterministic Solution Robustness

In Section 2, the concept of a *robust* supply chain was defined as a supply chain that can maintain normal operating performance under different scenarios. According to the findings from Section 4.5.2, the feasibility of the proposed *optimal* deterministic solution is unable to produce a feasible solution in approximately 30% of the scenarios that were tested. The scenarios were decomposed to determine which aspects most influenced the feasibility of the problem (in the deterministic model) and the costs (captured by the objective value). It was found that there was a strong positive correlation between feasibility and the total capacity of the resources. A weaker (although still significant) correlation existed between feasibility and total demand, although in this case demand was negatively correlated with feasibility indicating that an increase in demand was likely to induce infeasibility of the problem. Thus, capacity was the key driver of the deterministic problem feasibility. When considering the objective value cost for the deterministic and stochastic approaches, it was clear that the stochastic programming approach resulted in higher total costs. Interestingly, when these costs were corrected for the change in resource cost (which is consistent among all scenarios), the difference became less distinguished. From the statistical perspective, handling costs can be considered equal for both the deterministic and stochastic approach, which could be considered a trivial finding as the same number of products pass through the network and a particular item can only be processed through one CDC and LDC. Given that handling costs do not vary drastically between facilities, it is expected that this result is consistent for both approaches. In contrast, it was observed that transportation costs differed between approaches with statistical significance, the lower costs being associated with the stochastic solution. This could be due to there being more resources, which could mean reduced distances between resources and customers. Additionally, it could be that the resource locations selected in the deterministic solution are not well positioned when considering fluctuations in demand levels. This marks the first occasion that *Mystery Merchants*' third core business question has been addressed directly: it does appear that the stochastic solution reduces the aggregated driving distance. Lastly, there was a strong positive correlation observed between demand and the objective value in both approaches, indicating that while capacity is a key driver for feasibility, demand is a key driver in the total cost of the network.

These findings indicate that the increased costs observed in the stochastic objective value, specifically, can be attributed to resource fixed costs. This bolsters the case for using stochastic programming to generate *robust* network configurations.

4.6 Effect of Correlations in Uncertain Parameters

The results that have been outlined until this point have assumed that there exists a correlation between customer demand levels, as well as a correlation between transportation costs within the network. For all previous experimentation, the correlation between samples has approximated $\rho = 0.75$. An additional experiment was designed to evaluate how the choice of ρ may influence the findings in Section 4. It was previously stated that `Demand` had a significant correlation with the objective value of the deterministic problem, while capacity had a stronger correlation with the feasibility. As such, it is expected that reducing ρ would consequentially reduce the variance of objective value realizations. This could result in a reduced correlation between the objective value and the demand level. Furthermore, it is expected that reducing ρ would reduce the correlation between `Demand` and feasibility, as it is less probable that demand will collectively rise higher than expected capacity limitations.

Indeed, when the experiment from Section 4.5.2 was repeated (solving the deterministic model for each independent scenario), the correlation between `Demand` and `Feasible` became insignificant (-0.08), while the correlation between `Capacity` and `Feasible` remained strong (0.84) (see Appendix A.3 for the correlation

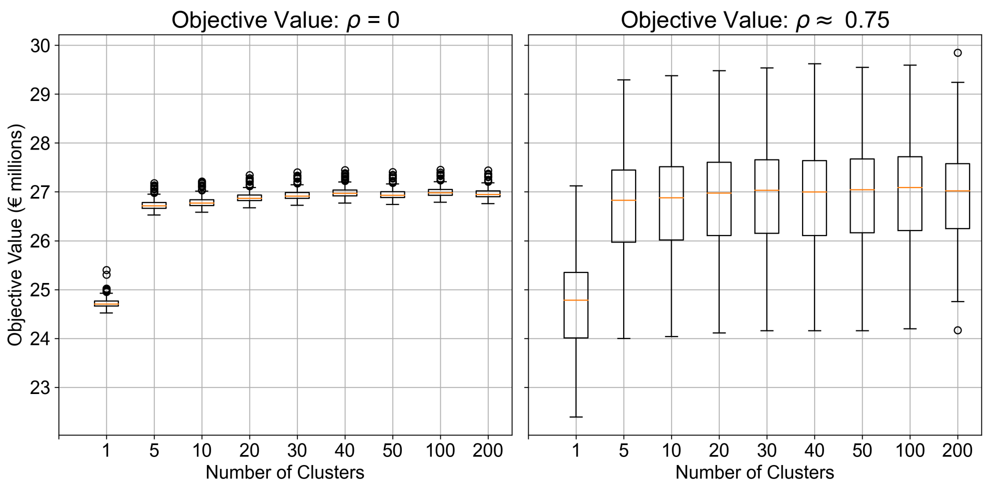

Figure 4.21: Distribution of objective value for scenarios with and without correlation within uncertain parameters.

plot). Treating `Demand` as a multivariate random variable with $\rho > 0$ can introduce additional load on the network during periods of higher demand, inducing infeasibility in some cases. Interestingly, the stochastic programming approach did not yield any differences in first-stage decisions with and without $\rho > 0$. It is possible that the stochastic approach can produce a network configuration that can operate normally, even under disruptions, and may also be robust to some inter-dependencies within parameters that are not explicitly captured by the model.

The stochastic problem was solved for scenarios under $\rho = 0$ for each $K = \{1, 5, 10, 20, 30, 40, 50, 100, 200\}$ and $N = 100$. The objective value associated with each scenario and each K was stored. Figure 4.21 shows the distribution of objective values between uncorrelated and correlated samples. It can be observed that the scenarios that do not correlate demand and transport costs are distributed with low variance. Conversely, when $\rho \approx 0.75$ the resulting objective values are distributed with higher variance. Although the multivariate distribution has been artificially generated, the contrast between costs when $\rho = 0$ and $\rho > 0$ demonstrates that it is important to consider any co-dependencies within uncertain parameters.

Chapter 5

Discussion

In this chapter, the implications of the proposed methods for addressing uncertainty in the context of SCND will be discussed. The theoretical implications are outlined in Section 5.1, which is focused on examining how the findings contribute to existing understandings of SCND under uncertainty and evaluating the implications of the findings for further theoretical development in SCND. Section 5.2 underscores the practical implications of the findings, that is:

- Discussing how the findings can be applied in supply chain management;
- Exploring the benefits and challenges regarding implementation and maintenance of the proposed methods;
- Identify guidelines that could assist practitioners to improve SCND under uncertainty.

Section 5.4 is focused on highlighting the limitations of the study, including challenges and constraints in the research process, limitations in data collection, and potential improvement areas for future studies. Lastly, Section 5.3 discusses the opportunity for future research, specifically:

- Research questions that arose during the study;
- Focus areas for further investigation;
- A discussion regarding how addressing the research gaps could contribute to furthering knowledge in supply chain management.

5.1 Theoretical Implications

The focus of this section is to outline how this research complements the existing research landscape of SCND under uncertainty and to articulate the implications of this research on further theoretical development within SCND. Section 5.1.1 highlights the theoretical contributions, while Section 5.1.2 discusses the implications.

5.1.1 Theoretical Contribution

The primary theoretical contributions of this study stem from five core topics: geographical clustering of customers, application of Benders decomposition, introducing a multivariate distribution for uncertain parameters, simultaneous consideration of multiple sources of uncertainty, and analysis of the convergence process to stable, optimal, first-stage decisions. Each of these topics will be discussed below.

Geographical clustering It was found that applying geographical clustering to customers in the SC network and solving the problem resulted in significantly reduced solving time. Importantly, the difference in objective value for clustered problems versus the full problem became inconsequential (to the first stage problem) with less than 20 clusters. This indicates that the problem structure was well preserved, even when few customer clusters were used. From a theoretical standpoint, it may be interesting to consider clustering techniques to simplify large problem instances that do not necessarily benefit from the complexity added by modeling many customer nodes.

Benders decomposition Although the Benders decomposition is a well-studied technique that has been applied to SCND extensively, it is not a technique that is well-used in practical applications. This study demonstrated the conditions under which the Benders decomposition reduced the solving time. In conjunction with the geographical clustering approach, the application of the Benders decomposition made it possible to solve otherwise prohibitively large problem instances (up to 100 scenarios). While the solving time without Benders decomposition exhibited a polynomial increase with the number of clusters, the same trend was not found when applying Benders decomposition. Increasing solving time was observed with the number of clusters, although this increase was less consistent and less pronounced compared to solving the stochastic programs without Benders decomposition.

Multivariate distribution It was found that introducing correlation within uncertain parameters had a significant influence on the objective value, while the first-stage decisions remained unchanged. This is a valuable insight, as accurately modeling the uncertain parameters could significantly improve the ability of organizations to prepare for fluctuations in costs. It may also be the case that these inter-dependencies are more realistic than modeling independent probability distributions; if this is the case then the results found when analyzing the influence on the objective function when introducing correlation should motivate the necessity to introduce realistic dependencies within uncertain parameters.

Multiple sources of uncertainty Supply chain networks have inherent uncertainty from a multitude of sources. As such, planning under uncertainty may benefit from addressing multiple sources of uncertainty simultaneously. The results from this study showed that disruptive uncertainty was the leading cause of infeasibility in the deterministic problem, while operational uncertainty influenced costs. When planning under uncertainty, it could be considered important to address both categories of uncertainty (as defined in this study) to fully understand the impact on the objective function and the optimal network configuration.

Convergence of SAA It was found that optimal first-stage decisions became consistent when solving 30 or more scenarios. It has been suggested and shown that the objective value using SAA approximates the true objective value as the number of scenarios N increases. This approximation becomes more accurate at an exponential rate with N. In the experiments from this study, it was shown that the objective value increased until $N = 30$, whereafter there appeared to be a plateau. Similar behavior was observed with the number of clusters K, where first-stage decisions were optimal and the objective value plateaued from $K = 5$. Thus, the smallest problem that provides an adequate representation of the problem can be defined as $K = 5$, $N = 30$. Although this will likely change between different networks, it is valuable to understand how many scenarios are necessary to build a robust network.

5.1.2 Implications of Theoretical Contributions in SCND

There are associated theoretical implications for each contribution. These implications can manifest as a trade-off, where the user or modeler must decide how best to make use of the contribution based on their organizational aims and objectives, or the requirements for a specific project. The implications associated with each contribution are discussed below.

Geographical clustering Applying clustering techniques to simplify large problem instances could raise opportunities from a theoretical perspective. Grouping customers or other network nodes based on criteria such as geographical position (or even demand characteristics) can reduce the complexity of the problem without compromising the exactness of first-stage decisions. This could provide avenues for developing alternative problem formulations that leverage unsupervised learning techniques such as clustering to make solving large-scale network design problems more efficient.

Benders decomposition This study has shown the conditions where Benders decomposition reduces solving time, and brings to attention the possible under-utilization of Benders decomposition in real-world applications of SCND under uncertainty, despite being extensively studied in academia. The application of Benders decomposition was particularly useful when combined with geographical clustering of clusters, which highlights the potential to solve what can be computationally prohibitive problems. An interesting finding is that the trends in solving time differ between Benders decomposition and monolithic approaches. Solving with Benders decomposition reveals relatively weak trends while solving the monolithic problem revealed clear polynomial trends when increasing clusters or scenarios. These findings can be considered valuable insights for both practitioners and researchers, as they highlight the case for practical applicability of Benders decomposition, as well as the associated computational efficiency.

Multivariate distribution This study highlighted the impact of introducing correlation within uncertain parameters on both the objective value and the network configuration. Interestingly, the first-stage decisions were consistent between correlated and uncorrelated scenarios, despite a significant impact on the objective value. The importance of accurately capturing the dynamics of the uncertain parameters is highlighted in these results, as fluctuations in cost can become increasingly severe as correlations increase. Inter-dependencies in real-world supply chains should therefore be captured for more realistic scenario modeling.

Multiple sources of uncertainty The findings in this study revealed that the primary cause for infeasibility when passing scenarios through a deterministic problem was disruptive uncertainty. Operational uncertainty had a significant impact on the network costs. It is important to recognize *and* incorporate disruptive and operational uncertainties to provide supply chain decision-makers with a comprehensive understanding of how these uncertainties influence the network costs and robustness of network configurations. Much emphasis has been placed on addressing multiple sources of uncertainty, underscoring the need for a comprehensive approach to strategic decision-making (in the context of SCND) that improves the robustness, and consequently resilience, of supply chain networks.

Convergence of SAA This study found important implications for the convergence of SAA in SCND. *Optimal* first-stage decisions were stable when solving for 30 scenarios or more. Previous studies have indicated that the true objective value is approached as the number of scenarios increases. Furthermore, it was found that the objective value increased until reaching 30 scenarios, whereafter it reached what appears to be a plateau (although small increases were still observed). This showed that understanding the appropriate number of scenarios for adequately representing the behavior of uncertain parameters is valuable for building robust supply chain networks without prohibitive time-to-solution.

5.2 Practical Implications

This section discusses the practical implications of the findings from this research. Section 5.2.1 considers how the finding may be applicable in supply chain management, Section 5.2.2 is focused on the challenges regarding implementing and operating the proposed method for handling uncertainty in SCND, and Section 5.2.3 offers guidelines for improving SCND while considering uncertainty.

5.2.1 Application of Findings in Supply Chain Management

It is clear that the deterministic approach to solving SCND problems does not adequately account for uncertainty in influential parameters; furthermore, generating scenarios and solving them independently has limited decision-making insights as solutions should be considered independent. Applying stochastic programming in SCND can allow decision-makers to develop network configurations that allow normal operations to be executed in fluctuating conditions. Essentially, stochastic programming could assist in creating supply chain networks that can adapt to changes in market conditions and disruptive events.

Adaptation to dynamic environments is critical in modern supply chains. The findings in this study show that addressing uncertainty in the design phase can generate robust network configurations that can adapt to significant disruptions and fluctuations. Such an approach could guide organizations to develop supply chain networks that are capable of fulfilling operational requirements in dynamic environments, rather than optimizing for independent scenarios or market expectations.

It is possible that comprehensive risk mitigation strategies could be developed and implemented, provided that key areas of uncertainty can be recognized and modeled appropriately. Proactive risk mitigation could decrease the impact of disruptions and fluctuations, which makes the network robust and inherently provides a platform for resilience.

Recognizing key uncertainties may be challenging, as it requires line-of-sight and information sharing between all echelons in supply chain networks. As such, the application of comprehensive risk mitigation strategies could motivate collaborative behavior between organizations.

5.2.2 Implementation and Operational Challenges

Although incorporating risk mitigation strategies, such as stochastic programming, has been shown to have significant benefits, implementing and operating such strategies have substantial challenges.

Stochastic programming solutions are sensitive to the input data, particularly when making use of SAA. The probabilistic approach proposed in this study requires that uncertain data can be modeled using statistical methods, such as estimating the underlying probability distribution of uncertain parameters. Making such estimations should not be attempted without adequate statistical analysis of the uncertain parameter. This is because it remains unclear whether sampling from a probability distribution that does not represent the underlying uncertain parameter would yield accurate or useful results. As such, it is a requirement that there should be sufficient high-quality data availability to produce useful risk mitigation strategies. Not all organizations have suitable data management systems to properly store, analyze, and use the appropriate data. Thus, it is likely that substantial workflow would be required to develop organization-wide data management systems that complement sophisticated strategic planning methods.

Until recent times, supply chains have predominantly been designed to minimize operational costs while reaching required operational performance. Major disruptions and events (*e.g.,* COVID-19, Ukraine-Russia war, shipping canal blockages, etc.) have highlighted the importance of a paradigm shift to designing supply chains that not only keep operational costs low but also provide resilience through robust network configuration and adaptability. This shift implies that it is likely that resilient supply chain networks will require redundancies that may have previously been considered wasteful expenditure, *e.g.,* operating more facilities with lower utilization levels, maintaining multiple supply routes, or avoiding opening facilities in unstable regions. Strategic decisions that balance these redundancies with costs could require careful change management and organizational alignment to support the fundamental shift in planning.

Along with change management and organizational alignment, it may also be useful to ensure that the concept of planning under uncertainty is understood (at least on a high level) by decision-makers. The concept of addressing multiple scenarios simultaneously may be challenging to distinguish from solving each scenario independently and comparing the results. The solution provided by the stochastic programming approach is likely more expensive, and therefore decision-makers may be inclined to infer 'optimal' decisions from a set of independent scenario configurations. Therefore, the methods used to solve SCND problems with uncertainty must be packaged in a way that highlights and motivates the core functionality and benefit of using such an approach (this relates to the *change management* aspect addressed earlier).

Building a representative set of scenarios and a method for simplifying the network (*e.g.,* geographical clustering of customers) may be challenging. It was shown that these operations resulted in the ability to generate robust network solutions without exorbitant time-to-solution. Additional challenges may arise when considering which sources of uncertainty are most useful to study. There is a direct trade-off between the time-to-solution and the number of uncertain parameters in the model. More uncertain parameters are likely to require more scenarios for comprehensive representation, thus resulting in larger problem instances (possibly prohibitively large).

5.2.3 Guidelines for Improving SCND Under Uncertainty

Based on the discussion around the practical applications of the findings in this study (Section 5.2.1) and challenges (Section 5.2.2), a set of guidelines was developed to enhance the practice of SCND. The guidelines are given below:

- Ensure that sources of uncertainty are *recognized* and *prioritized* in terms of their expected influence on strategic planning decisions.

- Evaluate the resilience and flexibility of existing networks by performing *robustness* analyses on incumbent decisions.

- Consider *inter-dependencies* when generating scenarios, these dependencies can have a significant impact on network costs and potentially the robustness of the network.

- Ensure that *complexity* is balanced with *robustness*. Striving for highly detailed and complex network models that minimize cost may not always provide robust solutions (a phenomenon known as *over-optimization*). Rather find a balance between high-detail, cost-optimized models, and robust configurations that mitigate risk on all planning levels.

- Foster *collaboration* between entities within the supply chain network. Establishing effective communication channels and mechanisms to promote information sharing could provide an opportunity to respond to changes through (near) real-time visibility.

- Invest in infrastructure to *collect, store, and analyze* data associated with supply chain network parameters and behavior.

- Reconsider long-term strategic decisions to align with key business questions that recognize influential uncertainties, this can mean considering market dynamics, forecasting trends, and anticipating potential disruptions, as well as ensuring that decision-makers are not penalized for making robust decisions that may come at a higher cost.

5.3 Limitations

This section addresses the limitations of this research. For this purpose, the limitations have been split into three overarching categories: Section 5.3.1 considers the challenges in the research process; Section 5.3.2 looks at the limitations in data availability and collection; while Section 5.3.3 is focused on the potential improvement areas for this research.

5.3.1 Challenges in the Research Process

Selecting an approach to handling uncertainty in SCND introduces challenges that depend mainly on the proposed method, rather than the core problem. As other methods were not implemented, it is undetermined how stochastic programming compares with alternatives (such as robust optimization, or a simulation approach). The chosen method was kept as general as possible to maintain applicability to different supply chain networks. As such, some assumptions and simplifications were introduced. Two of the most influential assumptions were (1) that the problem can be simplified into a *first* and *second* stage, and (2) the assignment of variables to each of these stages. For (1), a major limitation in computational tractability is that additional stages add computational complexity, which can quickly cause the size of the problem to explode. However, it may be an oversimplification to only consider two stages. Supply chain experts and relevant literature (discussed in Section 2) do agree that two stages can adequately represent a large portion of SCND problems. For (2), there may be alternative approaches to the one suggested in this study. When supply chains are demand-driven (meaning products flow through the system based on customer orders) then it may be suitable to model first-stage decisions solely as facility locations. However, when a supply chain is proactively distributing products into the market, it may be useful to consider facility locations and stock levels as first-stage decisions, for example. Another limitation of this study is it is lacking *comprehensive* external validity, *i.e.,* only one case study was performed. To better understand the findings that were presented, it could be beneficial to compare the results between different supply chain networks, especially those with multiple suppliers and production facilities.

5.3.2 Limitations in Data Collection

Although the static data for facilities, transport lanes, and customers for *Mystery Merchants* was available, there was no data to provide insights on the uncertain parameters. As such, the scenarios were generated using artificial data. Assumptions were made to produce the behavior of uncertain parameters, which may not be applicable in reality. In particular, a probability distribution was used as a proxy for facility capacity. In reality, fitting a probability distribution to facility capacity may be difficult or even impossible, because disruption events are notoriously challenging to predict and model. Furthermore, all facilities were modeled using the same probability distribution - likely an oversimplification. A more comprehensive result may be achieved by collaborating with facility managers and decision-makers to determine facility-specific disruption probabilities.

5.3.3 Potential Improvement Areas

Three key areas of improvement were identified from Sections 5.3.1-5.3.2. It is expected that addressing these improvement areas would further bolster the results and applicability of this study. The key improvement areas are:

1. Utilize *historical data* from uncertain parameters to estimate probability distributions. This would assist with developing realistic and representative scenarios, which may significantly alter the resulting network configuration and cost.

2. *Collaborate* with the case subject (*Mystery Merchants*) to critically evaluate the assumptions and subsequent modeling approach. It may be valuable to get feedback on the proposed methodology, it could be that important elements of the network design have not been adequately addressed, while others may have been over-emphasized.

3. Evaluate the approach on *more* cases. Further evaluation would indicate whether the results from this study generalize across different problem contexts.

5.4 Future Research

This section is aimed at addressing the opportunities and avenues for future research studies that consider SCND under uncertainty. This has been divided into three main areas: Section 5.4.1 covers the research questions that emerged during the study; Section 5.4.2 considers the specific focus areas for further investigation, which can include elements of this research that did not fit within the scope of the study; and Section 5.4.3 discusses the opportunities for furthering the development of knowledge when considering uncertainty in SCND.

5.4.1 Emerging Research Questions

Several research questions arose during the study, with some stemming from the limitations in Section 5.3. The most urgent research questions are listed below:

- How does the proposed approach methodology compare to other approaches, such as robust optimization, in terms of resulting network configurations and computational efficiency?
- What are the implications and trade-offs of using discrete scenarios instead of a probabilistic approach in SCND under uncertainty?
- How do the proposed methods perform when applied to different supply chain networks, and which of the findings generalize across network types?
- How do different probability distributions, when used to generate scenarios, influence the robustness of the obtained solution? What are the impacts of different distributions on the resulting network configurations and subsequent robustness?
- How can the proposed methods be applied in a user-friendly manner that facilitates understanding of the problem and interpretation of the results? What decision support tools can be developed to enhance decision-making in SCND under uncertainty?

5.4.2 Focus Areas for Further Investigation

Some aspects of this study were not investigated comprehensively, therefore further investigation could focus on addressing these aspects. First, an exhaustive sensitivity analysis could be conducted by evaluating how each of the uncertain parameters influences the objective value and network configuration. Secondly, additional sources of uncertainty could be considered, *e.g.,* transport lane disruptions. Thirdly, introducing dependencies between uncertainties (rather than *within*) should be considered, *e.g.,* high transport costs has *some* effect on demand. Additionally, the long-term implications of the proposed methods (or similar methods) could be analyzed in terms of resilience, market competitiveness, and overall supply chain performance.

5.4.3 Furthering Knowledge in Supply Chain Management

In addition to the *out-of-scope* (Section 5.4.1) and *in-scope* (Section 5.4.2) suggestions for future research, it is also recommended that these studies are complemented by additional elements that could further progress the field of supply chain management. These additional elements are:

- Integrating predictive analytics methods with real-time data, enabling proactive risk mitigation and recovery from disruptions while they evolve;
- Application simulation-based techniques, such as discrete event simulation, to better distinguish the effect of duration and intensity from dynamic and complex disruptions;

- Exploration of the impact of global trends, such as climate change, political relationships, or disruptive technologies, on SCND; and, subsequently, the development of strategies that further enhance resilience within supply chains.

Chapter 6

Conclusion

This study aims to determine whether introducing uncertainty in key quantities of an out-of-the-box commercial SCND product could enhance decision-making. The method used to investigate this aim was defined and supported by the literature review and preliminary background information. It was found that the most applicable approach to the problem was to use stochastic programming with uncertainty in demand, transport costs, and facility capacity. This approach aims to generate a robust network configuration, *i.e.,* fixing an optimal set of binary decisions that allow the network to function under fluctuations of uncertain parameters. Data from a distribution network were used to evaluate this method. The scenario data were, however, generated using artificial representations of uncertainty. The proposed stochastic programming method was applied to the company data and the results were analyzed in terms of computational properties, associated network configurations, and also practical implications. The study was evaluated in terms of theoretical and practical contributions, limitations, and opportunities for future research. This section also provided an opportunity to synthesize the results and extract practical insights into how this study contributes to the further development of products offered by AIMMS.

The main findings showed that using stochastic programming could result in significant improvements to the robustness of supply chain networks. The network from the pre-existing deterministic model was unable to operate in a large portion of scenarios that were tested from the stochastic programming model. It was also found that transportation costs were reduced in the stochastic approach, indicating less distance covered by transportation mediums. Additionally, inter-dependencies should not be overlooked in the modeling process, as they had a significant influence on the variance of network costs. In terms of computation time, the stochastic programming model could have prohibitively long solving times, particularly when a large number of scenarios were used. Efforts to reduce this solving time were successfully employed, namely Benders decomposition and geographical clustering of customers. It was observed that the clustering approach did not compromise the structure of the problem. Even with a small number of customer clusters, all first-stage decisions were consistent with the full problem.

The results could be largely dependent on the company data and the scenario generation approach. Thus, it was suggested that further analysis using different supply chain networks and historic data could enhance the holistic conclusions and generalizability of results. However; the results and conclusions in this study indicate that incorporating uncertainty in supply chain network design provides a notable opportunity to enhance decision-making in supply chain network design.

Bibliography

[1] R. Aldreghetti, D. Battini, D. Ivanov, and I. Zennaro. Costs of resilience and disruptions in supply chain network design models: A review and future research directions. *International Journal of Production Economics*, 235, 2021.

[2] A Ben-Tal, L.E. Ghaoui, and A Nemirovski. *Robust Optimization*, volume 28 of *Princeton Series in Applied Mathematics*. Princeton University Press, 2009. ISBN 978-1-4008-3105-0.

[3] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, 2011. ISBN 978-1-4614-0236-7.

[4] A. Bueno-Solano and M.G. Cedillo-Campos. Dynamic impact on global supply chains performance of disruptions propagation produced by terrorist acts. *Transportation Research Part E: Logistics and Transportation Review*, 61(C):1–12, 2014.

[5] C. Colicchia, F. Dallari, and M. Melacini. Increasing supply chain resilience in a global sourcing context. *Production Planning & Control*, 21(7):680–694, 2010.

[6] T.G. Crainic, M. Hewitt, F. Maggioni, and W. Rei. Partial Benders Decomposition: General Methodology and Application to Stochastic Network Design. *Transportation Science*, 55:414–435, 2020.

[7] S. Dani. *Predicting and Managing Supply Chain Risks*, pages 53–66. Springer US, 2009. ISBN 978-0-387-79934-6.

[8] M. Dong. Development of supply chain network robustness index. *International Journal of Services Operations and Informatics*, 1:54, 2006.

[9] M. Fattahi and K. Govindan. 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. *Transportation Research Part E: Logistics and Transportation Review*, 118:534–567, 2018.

[10] M. Fattahi, K. Govindan, and R. Maihami. Stochastic optimization of disruption-driven supply chain network design with a new resilience metric. *International Journal of Production Economics*, 230, 2020.

[11] M.S. Golan, L.H. Jernegan, and I. Linkov. Trends and applications of resilience analytics in supply chain modeling: systematic literature review in the context of the COVID-19 pandemic. *Environment Systems and Decisions*, 40:222–243, 2020.

[12] K. Govindan and M. Fattahi. Investigating risk and robustness measures for supply chain network design under demand uncertainty: A case study of glass supply chain. *International Journal of Production Economics*, 183:680–699, 2017.

[13] K. Govindan, M. Fattahi, and E. Keyvanshokooh. Supply chain network design under uncertainty: A comprehensive review and future research directions. *European Journal of Operational Research*, 263: 108–141, 2017.

[14] I. Heckmann, T. Comes, and S. Nickel. A critical review on supply chain risk – Definition, measure and modeling. *Omega-international Journal of Management Science*, 52:119–132, 2015.

[15] K. Katsaliaki, P. Galetsi, and S. Kumar. Supply Chain Disruptions and Resilience: A Major Review and Future Research Agenda. *Annals of Operations Research*, 319, 2021.

[16] Walid Klibi and Alain Martel. Modeling approaches for the design of resilient supply networks under disruptions. *International Journal of Production Economics*, 135:882–898, 2012.

[17] X. Li and K. Zhang. A sample average approximation approach for supply chain network design with facility disruptions. *Computers & Industrial Engineering*, 126:243–251, 2018.

[18] Judit Mesterné Monostori. Supply chains robustness: Challenges and opportunities. *Procedia CIRP*, 67:110–115, 2018.

[19] National-Research-Council. *Disaster Resilience: A National Imperative*. The National Academies Press, 2012. ISBN 978-0-309-26150-0.

[20] Xuwei Qin, X. Liu, and Lixin Tang. A two-stage stochastic mixed-integer program for the capacitated logistics fortification planning under accidental disruptions. *Computers and Industrial Engineering*, 65: 614–623, 2013.

[21] V. Rajagopala, S.P. Prasanna, and M. Gohb. Decision-making models for supply chain risk mitigation: A review. *Computers and Industrial Engineering*, 113:646–682, 2017.

[22] Peter Roelants. Multivariate normal distribution. Website, September 28 2018. URL `https://peterroelants.github.io/posts/multivariate-normal-primer/`. Accessed on 6 June 2023.

[23] T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro. A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, 167:96–115, 2005.

[24] A.J. Schmitt and M. Singh. A quantitative analysis of disruption risk in a multi-echelon supply chain. *International Journal of Production Economics*, 139(1):22–32, 2012.

[25] A. Shukla, V.A. Lalit, and V. Venkatasubramanian. Optimizing efficiency-robustness trade-offs in supply chain design under uncertainty due to disruptions. *International Journal of Physical Distribution & Logistics Management*, 41:623–647, 2011.

[26] L.V. Snyder, Z. Atan, P. Peng, Y. Rong, A.J. Schmitt, and B. Sinsoysal. OR/MS models for supply chain disruptions: a review. *IIE Transactions*, 48:89–109, 2016.

[27] U. Soni, V. Jain, and S. Kumar. Measuring supply chain resilience using a deterministic modeling approach. *Computers and Industrial Engineering*, 74:11–25, 2014.

[28] IBM Optimization Studio. Ibm ilog cplex optimization studio documentation - benders' algorithm. `https://www.ibm.com/docs/en/icos/20.1.0?topic=optimization-benders-algorithm`, 2021. Accessed: June 13, 2023.

[29] B. Tomlin. On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks. *Management Science*, 52:639–657, 2006.

[30] B.M. Tosarkani and S.H. Amin. A possibilistic solution to configure a battery closed-loop supply chain: Multi-objective approach. *Expert Systems with Applications*, 92:12–26, 2018.

[31] Panagiotis Tsiakis, Nilay Shah, and Constantinos C. Pantelides. Design of Multi-echelon Supply Chain Networks under Demand Uncertainty. *Industrial & Engineering Chemistry Research*, 40:3585–3604, 2001.

[32] K. Zhao, K.P. Scheibe, J. Blackhurst, and A. Kumar. Supply chain network robustness against disruptions: Topological analysis, measurement, and optimization. *IEEE Transactions on Engineering Management*, 66:127–139, 2019.

[33] S. Zokaee, A. Jabbarzadeh, B. Fahimnia, and S.J. Sadjadi. Robust supply chain network design: an optimization model with real world application. *Annals of Operations Research*, 257:15–44, 2017.

[34] H. Üster and G. Memişoğlu. Biomass Logistics Network Design Under Price-Based Supply and Yield Uncertainty. *Transportation Science*, 52:474–492, 2017.

Appendix

A.1 Sampling Experiment

This experiment compares sampling using Numpy's `random.multivariate_normal` library (`library` implementation) and the method that was implemented using an affine transformation and Cholesky decomposition from Section 3.3.1 (`custom` implementation). The covariance matrix is of size $n \times n$, where n is the size of the vector of expected values in the uncertain parameter. The experiment was conducted for a subset of $n = 1, \ldots, 5000$, where for each n a vector of size n was sampled using the `numpy.random.uniform` distribution between [1000, 10000]. A covariance matrix of size $n \times n$ was then generated, and used as input for sampling an n matrix from both the `library` implementation and the `custom` implementation The time to sample for each case was recorded and the associated samples that were generated were stored. Figure 1a shows the comparison between the `custom` and `library` implementations, where it is clear that the `custom` implementation performs significantly better than the `library` implementation, particularly when n is large. Figure 1b shows the same comparison as Figure 1a, except with a logarithmic scale in time. This figure enables further observation of the rate of increase. Lastly, Figure 1c shows the comparison of actual samples (for $n = 5000$ in this case), indicating no significant or observable difference in the resulting sample distributions. The same observation was made for each n. Each boxplot in Figure 1c has been formed by taking the mean over 100 scenarios (samples) of $n = 5000$ in order to ensure an accurate comparison of the realized distributions.

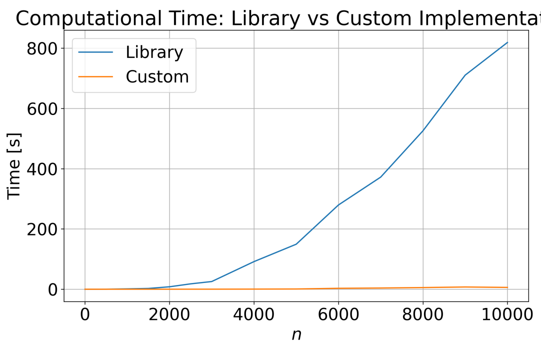

(a) Computational Time: `Library` vs `Custom` Implementation

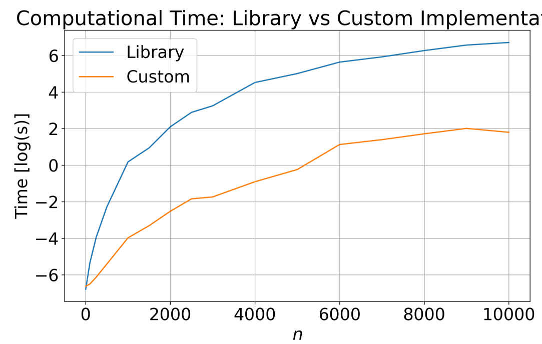

(b) Computational Time: `Library` vs `Custom` Implementation (Log Scale)

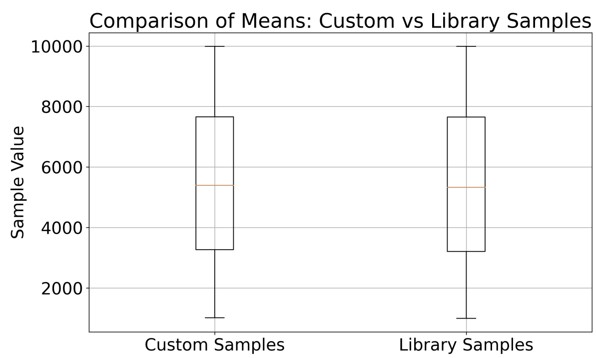

(c) Comparison of Means: Custom vs Library Samples

Figure 1: Comparison of Computational Time and Sample Means

A.2 Clustering Experiment

Table 1 shows the solving time, total cost, and percentage difference in *transport* cost between solving for the clustered problem versus the full problem. Note that transport cost is the only cost that differs noticeably between different clusters.

Clusters	Variables (Integer)	Constraints	Time [s]	Total [€]	% Delta
1	2044 (564)	2204	0.27	21,009,456.51	34.96%
5	2524 (564)	2764	0.63	22,911,138.91	4.91%
10	3124 (564)	3464	0.98	23,010,233.45	3.34%
15	3724 (564)	4164	2.13	23,208,184.83	0.21%
20	4324 (564)	4864	3.81	23,121,749.03	1.58%
30	5524 (564)	6264	6.4	23,196,191.93	0.40%
40	6724 (564)	7664	3.92	23,209,664.27	0.19%
50	7924 (564)	9064	6.45	23,186,119.40	0.56%
100	13924 (564)	16064	11.7	23,257,338.82	0.56%
200	25925 (564)	30064	51.58	23,219,508.25	0.03%
300	37924 (564)	44064	122.5	23,226,138.91	0.07%
400	49924 (564)	58064	291	23,221,680.13	0.00%

Table 1: Results for clustering customer locations.

A.3 Correlation Analysis of Results for Univariate Uncertainty Distribution

Fig. 2 shows the correlation heatmap that considers the correlation between uncertain parameters and the feasibility of the problem. In this case, the uncertain parameters are sampled from independent univariate distributions, and it is clear that demand does not have a significant relationship with the feasibility of the problem, while capacity remains strongly correlated. This differs from the case when multivariate distributions are used to sample demand and capacity, as those results indicate that demand and capacity are strongly correlated with problem feasibility.

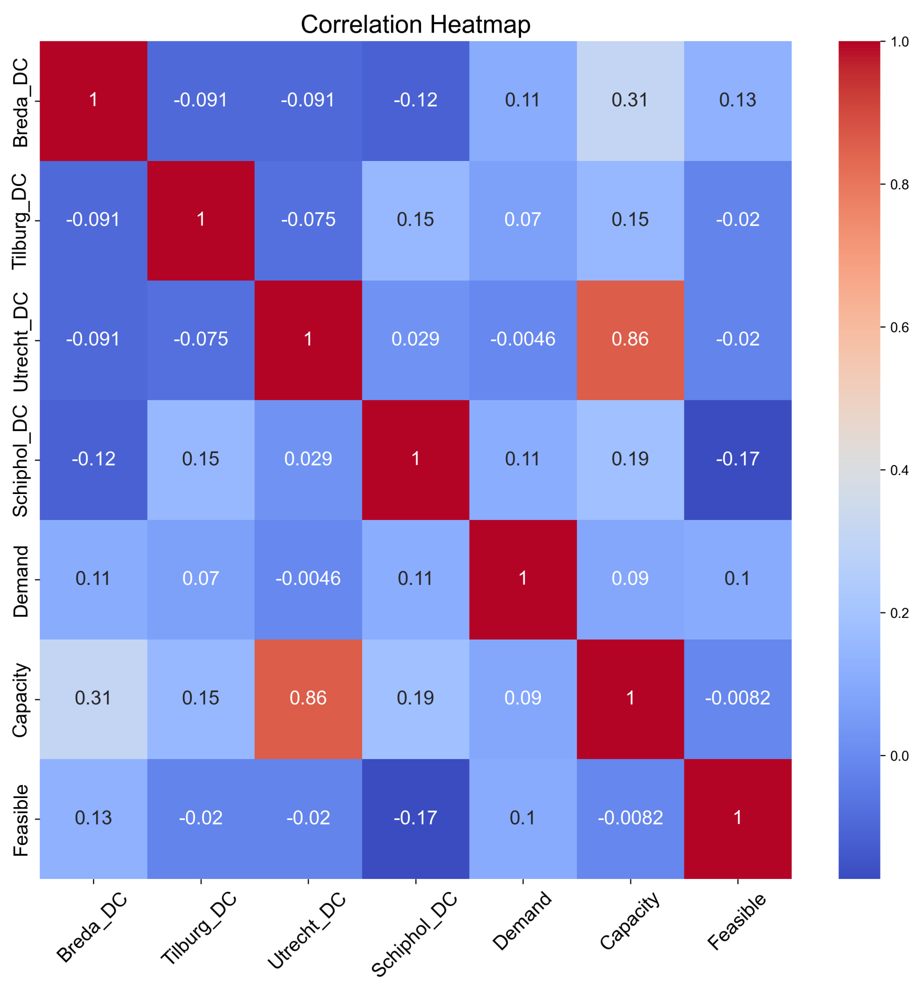

Figure 2: Correlation plot for uncorrelated uncertainty.