


ROBUST RESOURCE ALLOCATION UNDER UNCERTAINTY



A THESIS SUBMITTED TO
THE GRADUATE SCHOOL OF ENGINEERING AND SCIENCE
OF BILKENT UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF
MASTER OF SCIENCE
IN
INDUSTRIAL ENGINEERING


By
Ali Eren Demir
July 2024

ROBUST RESOURCE ALLOCATION UNDER UNCERTAINTY

By Ali Eren Demir

July 2024

We certify that we have read this thesis and that in our opinion it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Science.



Mustafa Çelebi Pınar(Advisor)

İtir Karaesmen Aydın

Özlem Çavuş İyigün

Approved for the Graduate School of Engineering and Science:

Orhan Arıkan
Director of the Graduate School

ABSTRACT

ROBUST RESOURCE ALLOCATION UNDER UNCERTAINTY

Ali Eren Demir
M.S. in Industrial Engineering
Advisor: Mustafa Çelebi Pınar
July 2024

Current methods for determining optimal capacity controls under uncertainty, such as stochastic optimization, often assume a known distribution for unknown parameters. This paper presents a novel approach using robust optimization to address the stochastic resource allocation problem in airline seat-inventory control. Our static formulations account for demand dependencies, offering a streamlined alternative to existing customer-choice models in revenue management literature. We analyze the structure of our proposed formulations, and provide insights on several robust counterparts of the seat-inventory control problem, considering various measures of robustness.

We introduce algorithms based on these robust formulations to derive actionable results. Through extensive simulations focused on seat allocation problems within the revenue management domain, our proposed formulations demonstrate significantly improved worst-case performances. Notably, even under favorable scenarios, the performance of our solutions are comparable to those of the existing methods in the revenue management literature. By providing protection against forecasting errors in demand distribution parameters and offering improved booking limit controls when demand falls below expected value, our formulations demonstrate superior revenue retention compared to existing methods in our comparative analyses.

Keywords: Robust optimization, resource allocation, capacity controls, revenue management, airline seat-inventory control, booking limits, worst-case analysis.

ÖZET

BELİRSİZLİK ALTINDA SAĞLAMCI KAYNAK PAYLAŞTIRMA

Ali Eren Demir
Endüstri Mühendisliği, Yüksek Lisans
Tez Danışmanı: Mustafa Çelebi Pınar
Temmuz 2024

Belirsizlikler altında optimal kapasite denetimini belirleme için mevcut yöntemler genellikle bilinmeyen parametreler için bilinen bir dağılımı varsayar, bu bir örnek rassal doğrusal programlamadır. Bu tez, hava yolu gelir yönetimi ve kapasite denetimindeki rassal kaynak paylaşma problemini ele almak için sağlamcı optimizasyon biçimlendirmelerini kullanarak yeni bir yaklaşım sunar. Dürük biçimlendirmelerimiz, bağımlı talepleri hesaba katarak, gelir yönetimi literatüründeki mevcut tüketici-tercihi modellerine alternatif bir çözüm sunar. Önerdiğimiz biçimlendirmelerin yapısını inceleyerek, kapasite denetimi problemine çeşitli sağlamcılık ölçütleri göz önüne alınarak alternatif optimizasyon yöntemleri sunuyoruz.

Bu sağlamcı biçimlendirmelere dayalı algoritmaları tanıtarak uygulanabilir sonuçlar elde ediyoruz. Gelir yönetimi alanındaki koltuk paylaşma problemlerine odaklanan geniş kapsamlı benzetimler yoluyla, önerdiğimiz biçimlendirmeler önemli ölçüde geliştirilmiş en kötü durum performansları sergiler. Özellikle, olumlu senaryolarda bile, çözümlerimizin performansı gelir yönetimi literatüründeki mevcut yöntemlerin performansı ile kıyaslanabilir seviyededir. Talep dağılım parametrelerini oluşturan tahmin hatalarının olası etkilerine karşı koruma sağlayarak, talebin beklenen değerin altında gerçekleştiği senaryolarda iyi çalışan metodlar sunarak, biçimlendirmelerimiz, talep gerçekleşmesi beklenenden düşük olduğunda mevcut çözümlerden daha iyi gelirler elde eder.

Anahtar sözcükler: Sağlamcı optimizasyon, kaynak paylaşma, kapasite denetimi, gelir yönetimi, hava yolu kapasite yönetimi, rezervasyon sınırlama kontrolleri, en kötü durum irdelemesi.

Acknowledgement

I am deeply indebted to my advisor, Prof. Mustafa Çelebi Pınar, for his invaluable guidance, encouragement, and unwavering support throughout the course of this research. His expertise, patience, and insightful feedback have been instrumental in shaping this thesis. I am profoundly grateful for his mentorship, and I hope to carry this influence into my future endeavors.

A special thanks to Prof. İtir Karaesmen Aydın. Her profound knowledge in this field has been crucial to my research. I am grateful for her support in theorizing, writing, and editing, as well as for being a part of my jury. Her contributions have significantly enhanced the quality of this work.

I would also like to extend my sincere gratitude to Özlem Çavuş İyigün for graciously agreeing to be one of my jurors and for her valuable insights.

I am deeply thankful for the unwavering support of my friends, who have stood by me through the highs and lows of this journey. In particular, I extend my heartfelt thanks to Kemal Dağ, Berke Aras, Göksü Ece Okur, and Zeynep Göze Gürkan for their exceptional support and encouragement. Your friendship has been a cornerstone of my strength. I am grateful to Mustafa Çolak, Emre Çakır, Ali İlhan Haliloğlu, Elif Rana Yöner, Deniz Şahin, Elif Sena Işık, and Gizem İkizler for their invaluable camaraderie and assistance.

To my family, whose love, encouragement, and understanding have sustained me through every challenge. In particular, I am profoundly thankful to my brother Emre Demir for his unwavering belief in my abilities, which has served as a constant source of motivation.

I would like to acknowledge the financial support of Tübitak, whose assistance has played a significant role in facilitating this research endeavor.

Lastly, my heartfelt gratitude goes out to all those who have lent me their ears and time, patiently listening to me discuss my work.

Contents

1	Introduction	1
1.1	Structure of Thesis	3
2	Problem Definition and Literature Review	4
2.1	Problem Definition	4
2.2	Literature Review	9
2.2.1	Single-Resource Revenue Management	9
2.2.2	Robust Optimization and Worst-Case Analysis	11
2.2.3	Robust Optimization and Distribution-Free Methods in Revenue Management	12
3	A New Robust Formulation	14
3.1	Introduction of Robustness	15
3.2	Uncertainty Set: ℓ_∞ -norm	17
3.3	Uncertainty Set: ℓ_1 -norm	20

3.3.1	Analysis of a Special Case: Independent Demand	21
3.3.2	Solution Structures Under ℓ_1 -norm	30
3.4	Budgeted Uncertainty Set	34
3.5	Interval Uncertainty	38
3.6	Observations on Effect of Uncertainty	40
3.7	Finding Booking Limits	41
3.7.1	Illustration	44
4	Computational Experiments	46
4.1	The Policies Tested in the Experiments	47
4.2	Simulation Setup	48
4.3	Performance Metrics	49
4.4	The Effect of Uncertainty Set Size	50
4.4.1	Changing Uncertainty Sizes Under Independent Demand Data	51
4.4.2	Changing Uncertainty Sizes Under Dependent Data	57
4.4.3	Choosing The Size of the Set	65
4.5	Experiments to Benchmark the Robust Methods to Others	67
4.5.1	Two-Product Performance Comparison	68
4.5.2	Comparison Under Multi-Product Scenario	94

4.5.3	Forecasting Errors Leading to Booking Limit Calculations	113
4.6	Summary of Observations	118
5	Conclusion and Future Research	120
A	Equivalent Linear Programming Examples of The Robust Models with $n = 2$	127
B	Effect of Standard Deviation - I	130
C	Effect of Standard Deviation - II	134
D	Analysis of 2-Class Example	137
E	Uniformly Distributed Demand	138
F	On The Robustness of The Realized Revenues	141
G	Partitioned Booking Limit Examples	143

List of Figures

3.1	Graphical Interpretation of Changing Isoprofit Line	31
3.2	Graphical Interpretation of Changing Constraint Line	32
3.3	Graphical Interpretation of Possible Feasible Regions	33
3.4	Objective Value Comparison for Four-Class Example Under In- creasing Uncertainty Set Sizes.	41
4.1	Booking Limit and Revenue Comparison of <i>Algorithm-NL_∞</i> and EMSRb Policies Under Uniform Arrival Pattern.	51
4.2	Booking Limit and Revenue Comparison of <i>Algorithm-NL₁</i> and EMSRb Policies Under Uniform Arrival Pattern.	53
4.3	Booking Limit and Revenue Comparison of <i>Algorithm-NL_{B₂}</i> and EMSRb Policies Under Uniform Arrival Pattern.	54
4.4	Booking Limit and Revenue Comparison of <i>Algorithm-NL_{B₁}</i> and EMSRb Policies Under Uniform Arrival Pattern.	54
4.5	Revenue Comparison of <i>Algorithm-NL₁</i> and EMSRb Policies Under LBH Arrival Pattern.	56

4.6	Booking Limit and Revenue Comparison of <i>Algorithm-NL_∞</i> and EMSRb Policies Under Uniform Arrival and Cholesky Decomposition.	59
4.7	Booking Limit and Revenue Comparison of <i>Algorithm-NL₁</i> and EMSRb Policies Under Uniform Arrival and Cholesky Decomposition	60
4.8	Booking Limit and Revenue Comparison of <i>Algorithm-NL_{B₂}</i> and EMSRb Policies Under Uniform Arrival and Cholesky Decomposition	61
4.9	Booking Limit and Revenue Comparison of <i>Algorithm-NL_∞</i> and EMSRb Policies Under Uniform Arrival and Square Root Method.	63
4.10	Booking Limit and Revenue Comparison of <i>Algorithm-NL₁</i> and EMSRb Policies Under Uniform Arrival and Square Root Method.	64
4.11	Booking Limit and Revenue Comparison of <i>Algorithm-NL_{B₂}</i> and EMSRb Policies Under Uniform Arrival and Square Root Method.	65
4.12	Relative Performance of Different Policies Against The Littlewood's Rule Under Changing Fare Ratio Levels, LBH Arrival Pattern.	70
4.13	Relative Performance of Different Policies Against The Littlewood's Rule Under Changing Demand Factor Levels, LBH Arrival Pattern.	72
4.14	Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under LBH Arrival Pattern.	76
4.15	Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under Uniform Arrival Pattern.	77
4.16	Realized Revenue Histograms at $(\sigma_1, \sigma_2) = (20, 20)$ When The Fare Ration Is 0.7, Under LBH Arrival Pattern.	79

4.17 Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case.	81
4.18 Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case Variant With DF of 1.6.	82
4.19 Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case Variant With DF of 1.6, Low-fare Deviation of 20.	83
4.20 Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case Variant With DF of 1, Low-fare Deviation of 5.	84
4.21 Percentile vs. Revenue. LBH Arrival Pattern Under The First Case.	85
4.22 Percentile vs. Revenue. LBH Arrival Pattern Under The First Case Variant with Low-fare Deviation of 20.	86
4.23 Percentile vs. Revenue. LBH Arrival Pattern Under 45-15 Demand.	87
4.24 Percentile vs. Revenue. Five-capacity, Full Negative Correlation Example Using The Square Root of The Covariance Matrix.	89
4.25 Percentile vs. Revenue. Eighty-capacity, Full Negative Correlation Example Using The Square Root of The Covariance Matrix.	90
4.26 Percentile vs. Revenue. Eighty-capacity, 0.7 Fare Ratio Example Using The Square Root of The Covariance Matrix.	91
4.27 Percentile vs. Revenue. Five-capacity, Full Negative Correlation Example Using Cholesky Decomposition of The Covariance Matrix.	92
4.28 Percentile vs. Revenue. Eighty-capacity, Full Negative Correlation Example Using The Cholesky Decomposition of The Covariance Matrix.	93

4.29 Percentile vs. Revenue. Eighty-capacity, 0.7 Fare Ratio Example Using The Cholesky Decomposition of The Covariance Matrix. . .	94
4.30 Realized Revenue Histograms for Table 3.1, Under LBH Arrival Pattern.	98
4.31 Realized Revenue Histograms for Table 4.8, Under LBH Arrival Pattern.	99
4.32 Realized Revenue Histograms for Table 4.8, Under Uniform Arrival Pattern.	100
4.33 Percentile vs. Revenue. LBH Arrival Pattern for Table 3.1.	102
4.34 Percentile vs. Revenue. LBH Arrival Pattern for Table 4.8.	103
4.35 Percentile vs. Revenue. Uniform Arrival Pattern for Table 4.8. . .	104
4.36 Percentile vs. Revenue. Uniform Arrival Pattern for Table 4.8, With 120 Unit Capacity.	105
4.37 Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 3.1 Using The Square Root of The Covariance Matrix.	107
4.38 Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 4.8 Using The Square Root of The Covariance Matrix.	108
4.39 Percentile vs. Revenue. LBH Arrival Pattern for Two-negative Block Correlated Table 3.1 Using The Square Root of The Covari- ance Matrix.	109
4.40 Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 3.1 Using The Cholesky Decomposition of The Covariance Matrix.	110

4.41	Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 4.8 Using The Cholesky Decomposition of The Covariance Matrix.	111
4.42	Percentile vs. Revenue. LBH Arrival Pattern for Two-negative Blocks Correlated Table 3.1 Using The Cholesky Decomposition of The Covariance Matrix.	112
4.43	Percentile vs. Revenue. LBH Arrival Pattern for Table 3.1.	114
4.44	Percentile vs. Revenue of Under Estimating The Mean Demand Value using Table 3.1 Under LBH Arrival Pattern.	115
4.45	Percentile vs. Revenue of Over Estimating The Mean Demand Value using Table 3.1 Under LBH Arrival Pattern.	116
4.46	Percentile vs. Revenue of Over and Under Estimating The Variance using Table 3.1 Under LBH Arrival Pattern.	117
B.1	Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.5 Under LBH Arrival Pattern.	132
B.2	Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.5 Under Uniform Arrival Pattern.	133
C.1	25th Percentile of The Realized Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under Uniform Arrival Pattern.	135
C.2	25th Percentile of The Realized Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under LBH Arrival Pattern.	136
E.1	Percentile vs. Revenue. LBH Arrival Pattern Under 40-80 Interval.	139

E.2	Percentile vs. Revenue. LBH Arrival Pattern Under 40-80 Interval Variant with Standard Deviations of 60 and 30.	140
G.1	Percentile vs. Revenue for Table 3.1 Under LBH Arrival Pattern, $C = 100$	143
G.2	Percentile vs. Revenue for Table 4.8 Under LBH Arrival Pattern, $C = 100$	144

List of Tables

2.1	Notation For Frequently Used Parameters In This Thesis.	6
3.1	Single-resource, four-class example [1].	44
4.1	Changing Uncertainty in <i>Algorithm-NL₁</i> , LBH Arrival Pattern for Table 3.1.	66
4.2	Booking Limits and Fare Ratios, Products with a Mean of 60, Standard Deviation of 30.	69
4.3	Booking Limits and Demand Factor, Products with a Mean of 60, Standard Deviation of 30.	71
4.4	Booking Limit Controls Under Changing Standard Variations When the Fare Ratio Is 0.7.	74
4.5	Booking Limits Under Different Demand Factor Levels for Table 3.1.	95
4.6	Mean Revenues with 95% CI Under LBH Arrival Pattern For Table 3.1 with SNR Column.	96
4.7	Mean Revenues with 95% CI Under Uniform Arrival Pattern For Table 3.1 with SNR Column.	97

4.8	Single-resource, four-product example, evenly spaced fares [1]. . .	103
B.1	Booking Limit Controls Under Changing Standard Variations When the Fare Ratio Is 0.5.	130
D.1	Two-product and 5 Units of Capacity Example, All Possible Book- ing Limit Controls	137
F.1	Metrics Related to The Distribution of Realized Revenues.	142

Chapter 1

Introduction

In contemporary business operations, revenue management has emerged as a fundamental strategic tool for organizations seeking to optimize profitability. Centered on the principles of pricing and capacity allocation, revenue management plays a pivotal role across various industries, including hospitality, airlines, retail, and manufacturing. At its core, revenue management involves apportioning finite resources, such as inventory or production capacity, across multiple products or classes amid uncertain and dynamic demands. This thesis delves into this challenge, focusing on the capacity allocation problem and offering a novel perspective infused with the principles of robust optimization, aiming to make decisions that are effective under a range of possible future scenarios.

The capacity allocation problem, also known as the seat-inventory control or the booking control problem in airline revenue management, addresses how to distribute limited resources effectively across products or classes in order to maximize revenues. One challenge lies in demand, it is unknown and fluctuates over time. Conventional methods often rely on deterministic approximations, which assume an idealized and clear view of future demand. However, real-world demand is inherently variable, influenced by factors such as seasonality, market volatilities, and unforeseen events. Such unpredictability can lead to suboptimal resource allocation and lost revenue opportunities.

To mitigate the risks associated with demand variability, companies often create artificial products (classes) that target different customer segments. By diversifying their product offerings and assigning different prices to products, businesses can better manage demand variability and maximize resource utilization. For instance, instead of relying on a single high-price product with low demand, companies develop a range of products, including lower-price options, to appeal to a broader customer base. This strategy helps balance demand across various segments and ensures more consistent utilization of available resources. Consequently, this necessitates a system to determine how much resource should be allocated to each product. Without carefully chosen limits on capacity allocated to lower-price segments/products, a company has the risk of selling out all of their capacity at low price.

In this context, the decision-maker defines a set of products aimed at different customer types or segments. Each product, while using the same resources, has a different demand probability structure. In this thesis, we use the terms class, product, fare-class interchangeably. The capacity allocation problem involves assigning a portion of the resources to a set of predefined products, each offering different values to customers and generating different levels of revenue for the organization. An example from the airline industry clarifies this: When we used the online booking tool of Turkish Airlines (THY) on 23 June, 2024 to book a one-way, economy ticket from Ankara to Istanbul on August 2, the same flight has seats available in EcoFly class for \$76.12, ExtraFly for \$87.69, and PrimeFly class for \$89.82. Clearly these three classes use the same resources (economy seats) on this direct flight but they have different “restrictions” such as baggage allowance (15 kg for EcoFly vs. 25 kg for PrimeFly) and refundability (full refund for PrimeFly, \$6.85 penalty for EcoFly). THY’s problem is to determine how many seats to sell in these three classes given the seat capacity of this flight.

The set of resources varies across industries. In airline revenue management, the resource is the available number of seats on flight legs. In other contexts, resources could include available hotel rooms, rental cars, or container space on a ship. Typically, the cost of allocating a unit of capacity to each product is the same and is negligible. In addition, whether a flight is full or empty, the total cost

of operating the flight is considered a sunk cost. Therefore, the focus in revenue management is on maximizing the revenue generated rather than maximizing profits considering the cost structure.

The primary objective of this thesis is to explore the capacity allocation problem through the lens of robust optimization, aiming to develop strategies that enhance decision-making under uncertainty. This research addresses the gap in existing literature by providing a comprehensive analysis of robust optimization techniques in the context of capacity allocation, offering new insights and practical applications.

1.1 Structure of Thesis

In Chapter 2 we discuss the problem in detail and give the relevant literature regarding airline revenue management, robust optimization, and joint works in these fields. First, single-leg airline revenue management methodologies are reviewed, which is followed by a brief introduction to the robust optimization literature. Finally, research on revenue management that uses the robust optimization framework is discussed. In Chapter 3, we introduce our new formulations influenced by robust optimization and derive their linear programming counterparts. The properties of these formulations are also discussed in that chapter. In Chapter 4, results of computational experiments are presented. The performance of our new formulations are compared to the existing methods in the literature under scenarios with correlated and uncorrelated demand. The thesis is then concluded and future research directions are discussed in Chapter 5.

Chapter 2

Problem Definition and Literature Review

In this chapter, we discuss the history of stochastic resource allocation problem and its applications. Specifically, applications related to airline revenue management are discussed in more detail as well as robust optimization developments in the area.

2.1 Problem Definition

Demand uncertainty is one of the most significant factors in resource allocation decisions. Decisions must be made before realizing any demand, and finding the optimal capacity control variables requires stochastic programming models. Assuming the decision maker is risk-neutral, a large proportion of the literature focuses on maximizing the expected revenue.

Due to the scope of this thesis, we discuss the problem with only a single resource (e.g. economy seats on one flight). In the context of airline revenue management, also known as single-leg capacity controls, this problem with only

one resource was first examined by [2], assuming only two products are offered. The first product's fare r_1 is assumed to be greater than or equal to that of the fare of the second product r_2 , i.e. $r_1 \geq r_2$. In this setting, the demand for the lower-fare class is high enough to fill the capacity, C , of the resource with a probability of 1, $P(D_2 \geq C) = 1$. The demand for low-fare class is assumed to materialize before the demand of the high-fare class.

In this simplified setting, optimal controls that maximize the expected revenues can be easily calculated by finding the protection level of the high-fare class. The protection level, PL_1 , is the amount of capacity that will be preserved for the higher-fare class. This is identical to limiting the number of accepted lower-fare classes, which is also called the booking limit, denoted as b_2 . When a protection level is determined, the booking limit for the lower-fare class equals the total capacity C minus the protection level, i.e., $b_2 = C - PL_1$.

To calculate the optimal protection level policy, [2] introduced the notion of *expected marginal revenue*. If the decision to protect the x^{th} unit for the higher-fare class is made, its expected marginal revenue is equal to $r_1 P(D_1 \geq x)$, where D_1 is the random demand of this class. Note that, this term is decreasing in x ; as the amount of protection for the higher-fare class increases, its expected marginal revenue will decrease. The optimal control for the resource allocation is simply to allocate a seat to the lower-fare class (i.e. not protect the seat for high-fare class) if and only if the expected marginal revenue of protecting the seat for higher-fare class is not higher than the revenue of the lower-fare product:

$$r_2 \geq r_1 P(D_1 \geq x). \quad (2.1)$$

In the optimal booking limit policy for the two-product single-resource setting, we find the optimal protection level of the first product, PL_1^* , using (2.1). This level satisfies the following conditions in equation (2.2) below if the demand distribution is discrete. Moreover, when the distribution of the product demand is known to be continuous, optimal controls can be found solving equation (2.3), where F_1 is the cumulative probability function (cdf) of high-fare demand D_1 and

F_1^{-1} is the inverse of this cdf:

$$r_2 < r_1 P(D_1 \geq PL_1^*), \quad r_2 \geq r_1 P(D_1 \geq PL_1^* + 1), \quad (2.2)$$

$$r_2 = r_1 P(D_1 > PL^*) \iff PL^* = F_1^{-1} \left(\frac{r_1 - r_2}{r_1} \right). \quad (2.3)$$

However, optimal controls of the two-product setting do not easily extend to the multi-product setting. To solve such instances, stochastic dynamic programming models are introduced. First, some notation must be introduced. The set N denotes the set of products. The number of products is thus equal to the cardinality of this set and denoted with $|N| = n$. For convenience, the set of products is ordered by fare, that is, $r_1 \geq r_2 \geq \dots \geq r_n$. The notations used for various concepts are listed in Table 2.1 for future reference.

Concept	Notation	Description (if needed)
Vector	Bold	j^{th} element of vector $\mathbf{D}$ is D_j
Index of Products	j	
Set of Products	N	
Number of Products	n	$ N = n$
Random Variables (RV)	Upper-Case	RV for demand of product j is D_j
RV Realization	Lower-Case	a realization of D_j is d_j
Mean Value of an RV	$\bar{d}_j, \bar{\mathbf{d}}$	mean value of RV D_j is $\bar{d}_j$
Resource Capacity	C	single-resource capacity
Resource Usage	a_{1j}	row-vector, capacity used by product j
Fares or Revenues	$r_j, \mathbf{r}$	where $r_1 \geq r_2 \geq \dots \geq r_n$
Booking Limit	b_j or $\mathbf{b}$	
Protection Level	$PL_j, \mathbf{PL}$	
Decision Variable	$y_j, \mathbf{y}$	

Table 2.1: Notation For Frequently Used Parameters In This Thesis.

The capacity of the resource is denoted by capital C . Lastly, D_j is the random variable for the demand for product j , $\forall j \in N$, with $\bar{d}_j$ denoting the expected value of this distribution. All vectors are denoted in bold. For instance, $\mathbf{D} = (D_1, \dots, D_n)^T$.

Assuming the demand arrives in the order of fare classes, starting with the lowest fare class, a stochastic dynamic programming model can be formulated as

follows: Let $V_j(c)$ denote the value function at the start of time period or stage j , with c representing the available resources. As the demand D_j of time period j is realized, the optimal controls can be decided. The Bellman equation with the boundary conditions $V_0(c) = 0$ for $c = 1, \dots, C$ is:

$$V_j(c) = E \left[\max_{0 \leq \tilde{s}_j \leq \min\{D_j, c\}} \{r_j \tilde{s}_j + V_{j-1}(c - \tilde{s}_j)\} \right]. \quad (2.4)$$

where $\tilde{s}_j$ is the amount of resources to be sold in class j . This formulation is taken from the survey of models and methods in [3]. In each stage j , $j = n, n-1, \dots, 1$, demand is D_j . The control $\tilde{s}_j$ is decided, which corresponds to how much of the demand realization that will be accepted given the amount of capacity left, c , after a value is realized for D_j . Based on the control $\tilde{s}_j$, remaining capacity is updated to $c - \tilde{s}_j$, the amount of realized revenue is observed. The expected revenue maximizing controls are found solving the recursive model in Equation (2.4). However, this method yields a computationally complex and costly solution. Fortunately, the stochastic resource allocation problem can easily be approximated by deterministic and/or static models. More complex approximations, including simpler dynamic programming formulations and heuristics founded on stochastic programming, have also been proposed and analyzed in the literature. The recursive model above assumes the demand in each fare class materializes in only one period and demand for low-fare products arrive first.

While unsophisticated, the “Deterministic Linear Program” *DLP* is a pivotal model for the dynamic booking control problem. It serves as an upper bound to a stochastic programming model and is used in various heuristics, mainly bid-price controls [1]. This deterministic and static approach assumes known demand, typically taken as the expected value of the demand distribution. When the demand distribution is assumed to be known, the DLP provides a definite upper bound to the stochastic dynamic programming problem, as it represents the best-case scenario where demand matches the expected value precisely [4]. However, when the true demand distribution is unknown and only the expected value is considered, the DLP cannot be utilized effectively. Despite not providing a tight upper bound to stochastic dynamic programming models, the DLP is widely adopted in both scholarly literature and practical applications [4].

The methodology proposed in this paper is also based on the DLP. We present the formulation of the DLP for the single-resource problem below, where maximum possible allocation is taken as the expected value of the product demand

$$\begin{aligned}
& \text{maximize} && \sum_{j \in N} r_j x_j \\
& \text{subject to} && \sum_{j \in N} x_j \leq C, \\
& && 0 \leq x_j \leq \bar{d}_j, \quad \forall j \in N.
\end{aligned} \tag{2.5}$$

and the decision variable x_j is the amount of resource sold to class j , $j \in N$. The DLP which was introduced for the multi-resource problem is beyond the scope of this thesis. When it is formulated for the single-resource problem, we see that it takes the form of a bounded knapsack problem where the optimal solution is: $x_1^* = \min(\bar{d}_1, C)$, $x_j^* = \min(\bar{d}_j, \sum_{k=1}^{j-1} x_k^*)$.

The assumption of DLP to disregard the probabilistic nature of the demand is naive. To address the challenges posed by demand uncertainty, we propose a novel formulation based on robust optimization. Our primary objective is to expand upon the DLP and pursue robust solutions within its framework. Robust optimization is powerful framework for decision-making under uncertainty by considering an ensemble of plausible scenarios. This approach ensures that the chosen strategy performs well across all conceivable scenarios, allowing decision-makers to handle the uncertainty of demand and circumvent the perils of underestimation or overestimation due to inaccurate estimation of demand probabilities.

We also would like to point out how the optimal solution of DLP is used in booking control. The optimal solution, $\mathbf{x}^*$ can be used or the model could be used to develop various heuristics, notably the bid-price heuristics. x_j^* is interpreted as the booking limit for product j , where the decision is to allocate (or save) x_j^* seats for this particular class. On the other hand, bid-price controls utilize the dual of DLP, leveraging shadow price information. The shadow price of the capacity control is used as a threshold value to determine which requests to accept or reject.. DLP is solved multiple times over a planning horizon as the capacity changes, and shadow prices determine which classes are allowed to make reservations.

2.2 Literature Review

2.2.1 Single-Resource Revenue Management

The stochastic capacity allocation problem in revenue management was first discussed in a two-product setting with known demand distributions by [2] as we mentioned before. The optimality condition equations (2.2) and (2.3) came to be known as the “Littlewood’s Rule” which states that the x -th request from the lower-fare class should be accepted if and only if the stated probabilistic condition in equation (2.1) holds, providing the optimal booking limit for the lower-fare class and the protection level for the high-fare class.

To address more realistic cases, solution strategies for n -product settings were developed. Static models can approximate the optimal solution, but unlike the two-product case, static models cannot guarantee an optimal solution for n -product settings. Dynamic programming can solve for the optimal solution, but due to the curse of dimensionality, such problems are impractical for realistic applications. Instead, various heuristics are used, notably the marginal revenue-based policies developed by Belobaba, known as the “Expected Marginal Seat Revenue” (EMSR) heuristics “a” and “b.” Belobaba generalized Littlewood’s Rule for multi-fare problems, calculating protection for higher-fare products against relatively lower-fare ones, and summing the protection levels in EMSRa [5]. Later, for the nested product structure, EMSR’s second version, EMSRb, was introduced, calculating protection levels for a product by treating all other lower-fare products as a single product [6].

To solve multi-fare problems, bid-price heuristics were also used. Versions of bid-price heuristics, which in the simplest case uses shadow prices obtained from solving linear programming models, have been influential for major airlines [7]. The method was first introduced in [8] and was later discussed in [9]. An overview of bid-price heuristics in the field, and the use of static linear programs in seat allocation problems are discussed in [10].

While heuristics gained attention for solving the dynamic multi-fare problems, optimal solutions were also investigated [11, 12]. However, due to long booking horizons and large number of products, these approaches suffer from Bellman’s curse of dimensionality. To make the problem tractable, the value function used in dynamic programming is approximated using different methodologies, which proved better than some heuristics such as the bid-price methods [13]. In the dynamic setting, there have been works incorporating approximations of the value function with bid-price heuristics as well [7, 14].

The optimal solution to the original problem should not consider any arrival pattern and be free of policies regarding enforcing booking limits. Research in the field relaxes such constraints. Earlier works assume an arrival pattern where the lowest-fare demand is realized first, called the low-before-high (LBH) pattern [15–17]. Subsequent research relaxed this arrival pattern, though demand realizations for different products still assumed to be uninterrupted by other products [18].

In terms of booking limit policies in the single-resource setting, most discussed ones are partitioned booking limits and nested booking limits. More advanced policies are also introduced, but such research is mainly related to network settings, which are beyond the scope of this thesis. We refer the readers to various nesting and class grouping policies, but not discuss them by incorporating them to our formulations [15, 19, 20]. Unlike the complexities of the network problems, in the single-resource case, when the arrival pattern is assumed to be sequential, i.e. arrivals grouped by fare class but not necessarily ordered LBH, a nested booking limit policy where nesting occurs by fare class is optimal [1]. We also focus on nested-by-fare booking limit policies in our experiments.

To account for the uncertainty in demand, most of the earlier research assumed a known demand distribution. Maximizing expected revenue is one of the objectives that has gotten the most attention in the literature. Using the full demand distribution information led researchers to employ mathematical models even for the complex multi-resource (or network) settings [1, 21]. When the demand distribution is not known, iterative methods that use information from previously admitted products in a given instance rather than historical demand information

are shown to be effective [22].

Proposed algorithms utilizing demand distributions and sequences of arrivals depend heavily on accuracy of demand forecasts. Possible forecasting errors are proven to be detrimental as shown in [23]. To eliminate possible forecasting errors, different techniques such as data-driven optimization were proposed [20, 22, 24].

2.2.2 Robust Optimization and Worst-Case Analysis

Another approach to dealing with forecasting errors or deviation from historical demand realization norms is robust optimization. In robust models, uncertain parameters are not assumed to be supported by known probability distribution functions. Instead, a set of possible values which the uncertain parameter can take is built. Over this set of possible realization of the uncertain parameters, a solution that is feasible for all realizations is found, and the best such solution is the robust optimal solution.

Robust optimization literature started to grow towards the end of the 1990s. Leading researchers began applying the concept to semidefinite programming [25, 26], least-squares problems [27], linear programming [28, 29], and convex optimization [30]. Many researchers turned to robust optimization following rapid developments in the field, especially in dealing with uncertainty in linear programs. After a seminal paper [31] solved issues regarding the conservativeness of robust solutions, the field gained popularity and other research fields begun applying its findings.

Analysis of distribution-free methods and their applications existed before early twenty-first-century developments under different objectives, such as maximizing the minimum profit [32–34]. Early developments in the field are summarized in the book [35]. For more recent developments and improvements in the subfield of Adaptive Robust Optimization, we draw from recent collective work in the field [36].

2.2.3 Robust Optimization and Distribution-Free Methods in Revenue Management

Using developments in robust optimization literature, the assumption that the true demand distribution is known is relaxed in the revenue management problems. Several works investigated robust optimization techniques in the setting of airline revenue management, assuming as little as possible regarding the true demand distributions. Seminal works in the field worked with the assumption that demand for each class is realized within a prespecified interval or utilized other robust approaches when demand is assumed to be discrete [37–39]. Also in [40], several robust optimization techniques were investigated for inventory problems under supply chain and revenue management domains.

Perakis and Roels propose two different objective functions to solve the stochastic capacity allocation problem in the context of airline seat allocation [37]. They find solutions that are robust by (i) maximizing the minimum possible revenue given all possible demand scenarios, and (ii) minimizing the maximum possible regret. Lan et al. choose absolute regret and competitive ratio as the objective function in their seat allocation models [38]. While we do not consider such objectives, our proposed methodologies are comparable as all assume that true demand distribution may not be known in advance and seek solutions with limited demand information. Additionally, Birbil et al. introduces robustness in single-resource setting where the demand distribution is assumed to be discrete [39]. We assume continuous demand distributions similar to two previously mentioned papers [37, 38].

A 2017 survey summarizes findings in the field of robust revenue management, focusing on risk-averse decision-making [41]. This work categorizes the models by their objective functions and types of risk-aversion strategies and investigates the field under capacity controls or dynamic pricing methods.

More recently, linear programs are developed to find robust booking limits, with the objective of minimizing maximum regret for network problems [42].

Other applications of robust optimization in the airline industry include crew scheduling [43]. Recently, forecasting errors in airline revenue management are also re-investigated [44], prior to discussions of robust forecasting of demand distribution parameters.

Our work introduces novel robust solutions to the single-leg capacity control problem in revenue management. While existing research that only assume demand lies in known intervals or works with discrete demand, we take a variation of the DLP formulation as the basis for our robust formulations. Demand intervals as in [37, 38] can easily be incorporated to our framework. We are also able to account for correlation in demand across different classes whereas majority of the research on the single-resource and network problems assume demand in each class is statistically independent of demand in other classes.

Chapter 3

A New Robust Formulation

The DLP, despite its simplicity, has been shown to be effective when benchmarked against optimal or other heuristic solutions to the seat allocation problem. Booking control policies derived from the DLP achieve high performance (e.g. total revenues) in various simulations. As previously mentioned, our proposed formulation is based on DLP, introduced in (2.5). In DLP, for each product, we have a decision variable x_j , which denotes the amount of resources sold to product j , hence the capacity allocated to product j . First, we change the decision variable in the DLP and introduce a new decision variable y_j . The variable y_j , together with the expected value of the demand for product j , correspond to the amount of sales of product j , hence the capacity to be allocated for product j . y_j is defined in the range $[0, 1]$, and can be interpreted as the fraction of the mean demand of class j to be met.

So, the relation between the variable x_j in (2.5) and our new variable y_j is $y_j \bar{d}_j = x_j$, for all $j \in N$. Using this transformation, a new formulation, referred to as the “*Modified Deterministic Linear Program*”, MLP is obtained:

$$\begin{aligned} & \underset{\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}}{\text{maximize}} && \sum_{j \in N} r_j y_j \bar{d}_j \\ & \text{subject to} && \sum_{j \in N} a_{1j} y_j \bar{d}_j \leq C. \end{aligned} \tag{3.1}$$

Notice that there is a one-to-one correspondence between the solutions of DLP

and MLP and the formulations are equivalent.

3.1 Introduction of Robustness

In the MLP, parameters for demand of products appear in both the objective function and the constraints. Previous research that has dealt with robust versions of the same problem, incorporated robustness only in the constraints [38] [37], or only in the objective function [39]. By using the variable transformation, we are able to change the way uncertain parameters affect the optimal solution: we incorporate uncertainty to both the objective function and the constraints.

Ideally, any solution to the seat allocation problem should achieve high revenues and high resource utilization. It is not reasonable to leave capacity unused if there is demand. The robust or distribution-free approaches try to achieve the highest revenues (or lowest regret) under the worst-case scenarios, i.e. the goal is to maximize the minimum possible revenue (or minimize the maximum regret). We employ the same approach in our formulation: the objective function is to maximize the minimum possible revenue over all possible demand scenarios. On the other hand, the capacity constraint should hold for all possible demand scenarios. Given the nature of the constraint, it is enough for the maximum possible resource use to not exceed the capacity. Therefore, we work with the following formulation:

$$\begin{aligned} z^* = \quad & \underset{\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}}{\text{maximize}} \quad \left\{ \underset{\mathbf{D} \in \mathcal{U}}{\text{minimize}} \sum_{j \in N} r_j y_j D_j \right\} \\ & \text{subject to} \quad \left\{ \max_{\mathbf{D} \in \mathcal{U}} \sum_{j \in N} a_{1j} y_j D_j \right\} \leq C. \end{aligned} \tag{3.2}$$

In the robust formulation (3.2), the set $\mathcal{U}$ is referred to as the “uncertainty set” and includes all possible demand scenarios. Also note that, the constraint set $\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}$ corresponds to $2n$ constraints in the form $0 \leq y_j$ and $y_j \leq 1$, for all $j \in N$.

Proposition 3.1.1. *z^* is non-decreasing in capacity of the resource, i.e., C .*

Proof. We observe that when the capacity of the resource increases, the feasible region of this maximization problem expands. An optimal solution to the problem with capacity C is also feasible for the problem with capacity $C + \xi$ for any $\xi > 0$. Therefore, the optimal objective function value of the problem with capacity C cannot be lower than the optimal objective function value of the problem with capacity $C + \xi$. So, z^* does not decrease as C increases. $\square$

In the remainder of the thesis, we investigate uncertainty sets that are defined by different norms. For an arbitrary ℓ_s -norm, where s defines the norm, generalized version of the uncertainty set of interest is defined by

$$\mathcal{U}_s = \{\bar{\mathbf{d}} + P\mathbf{u} : \|\mathbf{u}\|_s \leq \epsilon\}. \quad (3.3)$$

In (3.3), we have $\bar{\mathbf{d}} \in \mathbb{R}^n$, $\mathbf{u} \in \mathbb{R}^n$ and $P \in \mathbb{R}^{n \times n}$. Firstly, $\bar{\mathbf{d}}$ is the vector of mean demand information which is estimated before the capacity allocation decisions are made. The parameter matrix P , which we will refer to as the “*Correlation Parameter Matrix*”, represents the matrix square root of the covariance matrix for the offered products, i.e., assuming the covariance matrix is defined by Σ , we have $\Sigma = PP$. However, in the subsequent chapters, the correlation parameter is also taken as the Cholesky Decomposition of the covariance matrix, where P is a lower-triangular matrix such that $\Sigma = PP^T$.

The parameter ϵ defines the size of the uncertainty set in (3.3). Basically, the higher the ϵ value, the larger the region for the demand scenarios and demand values are allowed to be farther away from the estimated mean demand. When the size of the uncertainty set is defined by only one norm, ϵ is used as the only parameter that defines the size of the uncertainty set. However, in the Budgeted Uncertainty formulation in Section 3.4, we use two norms in the same model and in that case we use two parameters Γ and ρ for the uncertainty set sizes.

Remark 3.1.2. *If the uncertainty set is assumed to consist only of the estimated mean demand, i.e., when $\epsilon = 0$, then the problem in (3.2) is equivalent to the DLP.*

3.2 Uncertainty Set: ℓ_∞ -norm

In this section, we incorporate the ℓ_∞ -norm to the uncertainty set formulation in equation (3.2) as follows:

$$\mathcal{U}_\infty = \{\bar{\mathbf{d}} + P\mathbf{u} : \|\mathbf{u}\|_\infty \leq \epsilon\}. \quad (3.4)$$

By definition, the ℓ_∞ -norm is the greatest value of the absolute values of the components of a vector, i.e. $\|\mathbf{u}\|_\infty = \max_j |u_j|$. We use the properties of the ℓ_∞ -norm to reduce our robust formulation to a linear program.

Proposition 3.2.1. *If the uncertainty set $\mathcal{U}_\infty$ is used, then the robust formulation in equation (3.2) is equivalent to the following linear program:*

$$\begin{aligned} z_\infty^*(\epsilon) = & \underset{\mathbf{y}, \mathbf{t}, \mathbf{w}}{\text{maximize}} && \mathbf{r}^T \text{diag}(\mathbf{y}) \bar{\mathbf{d}} - \epsilon \sum_{j \in N} t_j \\ & \text{subject to} && a_1 \cdot \text{diag}(\mathbf{y}) \bar{\mathbf{d}} + \epsilon \sum_{j \in N} w_j \leq C, \\ & && -t_j \leq (\mathbf{r}^T \text{diag}(\mathbf{y}) P)_j \leq t_j, \quad \forall j \in N, \\ & && -w_j \leq (a_1 \cdot \text{diag}(\mathbf{y}) P)_j \leq w_j, \quad \forall j \in N, \\ & && 0 \leq y_j \leq 1, \quad \forall j \in N, \\ & && t_j, w_j \geq 0, \quad \forall j \in N. \end{aligned} \quad (3.5)$$

Proof. First, we write the constraint that is provided in the robust formulation in equation (3.2) as

$$\left\{ \max_{\mathbf{u}: \|\mathbf{u}\|_\infty \leq \epsilon} \sum_{j \in N} a_{1j} y_j (\bar{\mathbf{d}} + P\mathbf{u})_j \right\} \leq C. \quad (3.6)$$

Taking the parts unrelated to the maximization problem out, the left-hand side of the constraint is simplified to

$$\sum_{j \in N} a_{1j} y_j \bar{d}_j + \left\{ \max_{\mathbf{u}: \|\mathbf{u}\|_\infty \leq \epsilon} \sum_{j \in N} a_{1j} y_j (P\mathbf{u})_j \right\} \leq C. \quad (3.7)$$

Moreover, the term $\max_{\mathbf{u}: \|\mathbf{u}\|_\infty \leq \epsilon} \sum_{j \in N} a_{1j} y_j (P\mathbf{u})_j$ has a closed-form solution that is obtained as:

$$\max_{\mathbf{u}: \|\mathbf{u}\|_\infty \leq \epsilon} \sum_{j \in N} a_{1j} y_j (P\mathbf{u})_j = \max_{\mathbf{u}: \|\mathbf{u}\|_\infty \leq \epsilon} a_1 \cdot \text{diag}(\mathbf{y})(P\mathbf{u}), \quad (3.8)$$

$$= \max_{\mathbf{u}: \|\mathbf{u}\|_\infty \leq \epsilon} [a_1 \cdot \text{diag}(\mathbf{y})P] \mathbf{u}, \quad (3.9)$$

$$= \max_{\mathbf{u}: |u_j| \leq \epsilon, \forall j \in N} \sum_{j \in N} [a_1 \cdot \text{diag}(\mathbf{y})P]_j u_j, \quad (3.10)$$

$$= \sum_{j \in N} \max_{\mathbf{u}: |u_j| \leq \epsilon} [a_1 \cdot \text{diag}(\mathbf{y})P]_j u_j, \quad (3.11)$$

$$= \epsilon \sum_{j \in N} |a_1 \cdot \text{diag}(\mathbf{y})P|_j, \quad (3.12)$$

$$= \epsilon \|a_1 \cdot \text{diag}(\mathbf{y})P\|_1. \quad (3.13)$$

We explain the steps taken in equations (3.8) to (3.13) next.

In equation (3.8), the summation on the left-hand side is restated using the matrix multiplication notation on the right. In the same equation, we introduce a new notation: $\text{diag}(\cdot)$ operator maps a vector, to a square matrix with both dimensions equal to the input vector length. Moreover, the diagonal elements of the output matrix respective takes the values of the elements of the input vector, and all other entries of the output matrix are zero. To get equation (3.17), associative property of matrix multiplication is used. In equation (3.10), we have written the matrix multiplication as a summation.

In equation (3.11), we interchange maximization with summation because the function separates by u_j . Next, intuitively, to maximize each term, we set $u_j^* = -\epsilon$, if $[a_1 \cdot \text{diag}(\mathbf{y})P]_j \leq 0$, $u_j^* = \epsilon$ otherwise to get (3.12). Lastly, the summation in (3.12) is simply the ℓ_1 -norm, hence in equation (3.13) we use the operator $\|\cdot\|_1$.

Finally, using this derivation, the capacity constraint in equation (3.7) rewritten in a simplified manner.

$$a_1 \cdot \text{diag}(\mathbf{y})\bar{\mathbf{d}} + \epsilon \|a_1 \cdot \text{diag}(\mathbf{y})P\|_1 \leq C, \quad (3.14)$$

or equivalently as:

$$\sum_{j \in N} a_{1j} y_j \bar{d}_j a_{1.} + \epsilon \sum_{j \in N} [a_{1.} \text{diag}(\mathbf{y}) P]_j \leq C. \quad (3.15)$$

Applying the same steps above to the inner maximization problem in the objective function in equation (3.2), we get

$$\min_{\mathbf{D} \in \mathcal{U}_\infty} \sum_{j \in N} r_j y_j D_j = \min_{|u_j| \leq \epsilon, \forall j \in N} \sum_{j \in N} r_j y_j (\bar{\mathbf{d}} P \mathbf{u})_j, \quad (3.16)$$

$$= \sum_{j \in N} r_j y_j \bar{d}_j + \min_{|u_j| \leq \epsilon, \forall j \in N} \sum_{j \in N} r_j [\text{diag}(\mathbf{y}) P]_j u_j, \quad (3.17)$$

$$= \sum_{j \in N} r_j y_j \bar{d}_j - \epsilon \sum_{j \in N} r_j |\text{diag}(\mathbf{y}) P|_j. \quad (3.18)$$

Therefore the robust problem with ℓ_∞ -norm uncertainty set is equivalent to

$$\begin{aligned} z_\infty^*(\epsilon) = & \underset{\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}}{\text{maximize}} \quad \mathbf{r}^T \text{diag}(\mathbf{y}) \bar{\mathbf{d}} - \epsilon \|\mathbf{r}^T \text{diag}(\mathbf{y}) P\|_1 \\ & \text{subject to} \quad a_{1.} \text{diag}(\mathbf{y}) \bar{\mathbf{d}} + \epsilon \|a_{1.} \text{diag}(\mathbf{y}) P\|_1 \leq C. \end{aligned} \quad (3.19)$$

The resulting model can be further simplified and reformulated as a linear programming model. Two new vectors of decisions variables, $\mathbf{t}$ and $\mathbf{w}$, are introduced to account for two separate ℓ_1 -norm operators, one in the objective function and one in the constraint. Each element of these new decision variable vectors correspond to one element of the vector inside the norm operators.

$$\begin{aligned} z_\infty^*(\epsilon) = & \underset{\mathbf{y}, \mathbf{t}, \mathbf{w}}{\text{maximize}} \quad \mathbf{r}^T \text{diag}(\mathbf{y}) \bar{\mathbf{d}} - \epsilon \sum_{j \in N} t_j \\ & \text{subject to} \quad a_{1.} \text{diag}(\mathbf{y}) \bar{\mathbf{d}} + \epsilon \sum_{j \in N} w_j \leq C, \\ & \quad -t_j \leq (\mathbf{r}^T \text{diag}(\mathbf{y}) P)_j \leq t_j, \quad \forall j \in N, \\ & \quad -w_j \leq (a_{1.} \text{diag}(\mathbf{y}) P)_j \leq w_j, \quad \forall j \in N, \\ & \quad 0 \leq y_j \leq 1, \quad \forall j \in N, \\ & \quad t_j, w_j \geq 0, \quad \forall j \in N, \end{aligned} \quad (3.20)$$

□

A parametric version when there are only two products under independent demands is available in Appendix A, Example B.1.

3.3 Uncertainty Set: ℓ_1 -norm

Now, we model the uncertainty set using the ℓ_1 -norm. Although the only difference with the previous section is in the norm, resulting linear programming (LP) formulation is not the same as in the previous section. Consequently, the optimal solutions are not expected to be the same. Also, the uncertainty set $\mathcal{U}_\infty$ allows all product demands to deviate from their mean upto a certain threshold, whereas $\mathcal{U}_1$ considers the total deviations (added up for all products) upto a threshold. Below, we present $\mathcal{U}_1$:

$$\mathcal{U}_1 = \{\bar{\mathbf{d}} + P\mathbf{u} : \|\mathbf{u}\|_1 \leq \epsilon\} \quad (3.21)$$

Proposition 3.3.1. *If the uncertainty set $\mathcal{U}_1$ is used, then robust formulation in equation (3.2) is equivalent to the following linear program:*

$$\begin{aligned} z_1^*(\epsilon) = & \underset{\mathbf{y}, \rho_{obj}, \rho_{constr}}{\text{maximize}} && \sum_{j \in N} r_j y_j \bar{d}_j - \epsilon \rho_{obj} \\ & \text{subject to} && a_1 \cdot \text{diag}(\mathbf{y}) \bar{\mathbf{d}} + \epsilon \rho_{constr} \leq C, \\ & && (\mathbf{r}^T \text{diag}(\mathbf{y}) P)_j \leq \rho_{obj}, \quad \forall j \in N, \\ & && (a_1 \cdot \text{diag}(\mathbf{y}) P)_j \leq \rho_{constr}, \quad \forall j \in N, \\ & && 0 \leq y_j \leq 1, \quad \forall j \in N, \\ & && \rho_{obj}, \rho_{constr} : \text{free}. \end{aligned} \quad (3.22)$$

Proof. The maximization problem on the left-hand-side of the capacity constraint $\max_{\mathbf{D} \in \mathcal{U}_1} \sum_{j \in N} a_{1j} y_j D_j \leq C$ is simplified as follows. Let $j^* = \arg \max_{j \in N} |a_{1j} \mathbf{y} P|_j$. Then, assume that $|u_{j^*}| = \epsilon$. The optimality of this solution to the problem $\max_{\|\mathbf{u}\|_1 \leq \epsilon} \sum_{j \in N} a_{1j} y_j (P\mathbf{u})_j$ can be seen by considering the contrapositive.

Assume to the contrary that $|u_{j^*}| = \epsilon$ is not the optimal solution. Then, there exists a positive real number α such that for some index j' , we have $|u_{j'}| = \alpha$. However, it is immediately evident that the sum $\sum_{j \in N} a_{1j} y_j (P\mathbf{u})_j$ would be greater or equal if $|u_{j^*}|$ were to be increased by α while $|u_{j'}|$ were to be decreased by α , as $|[a_{1j} \mathbf{y} P]_{j^*}| \geq |[a_{1j} \mathbf{y} P]_{j'}|$. Therefore,

$$\max_{\mathbf{D} \in \mathcal{U}_1} \sum_{j \in N} a_{1j} y_j D_j = a_1 \cdot \text{diag}(\mathbf{y}) \bar{\mathbf{d}} + \epsilon \|a_1 \cdot \text{diag}(\mathbf{y}) P\|_\infty \leq C. \quad (3.23)$$

Using the same reasoning, the inner minimization term in the objective function can be reduced to $\min_{\mathbf{D} \in \mathcal{U}_1} \sum_{j \in N} r_j y_j D_j = \sum_{j \in N} r_j y_j D_j - \epsilon \|\mathbf{r}^T \text{diag}(\mathbf{r})P\|_\infty$. The robust model under the uncertainty set $\mathcal{U}_1$ is then reformulated as the following:

$$\begin{aligned} z_1^*(\epsilon) = & \underset{\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}}{\text{maximize}} && \mathbf{r}^T \text{diag}(\mathbf{y})\bar{\mathbf{d}} - \epsilon \|\mathbf{r}^T \text{diag}(\mathbf{y})P\|_\infty \\ & \text{subject to} && a_1 \cdot \text{diag}(\mathbf{y})\bar{\mathbf{d}} + \epsilon \|a_1 \cdot \text{diag}(\mathbf{y})P\|_\infty \leq C. \end{aligned} \quad (3.24)$$

While this model is not linear, an equivalent linear programming formulation can be found by introducing two additional decision variables ρ_{obj} and ρ_{constr} that will correspond to the maximal element of the vectors $\mathbf{r}^T \text{diag}(\mathbf{y})P$ and $a_1 \cdot \text{diag}(\mathbf{y})P$, respectively. These variables are used to linearize the ℓ_∞ -norms. Hence, we get

$$\begin{aligned} z_1^*(\epsilon) = & \underset{\mathbf{y}, \rho_{obj}, \rho_{constr}}{\text{maximize}} && \sum_{j \in N} r_j y_j \bar{d}_j - \epsilon \rho_{obj} \\ & \text{subject to} && a_1 \cdot \text{diag}(\mathbf{y})\bar{\mathbf{d}} + \epsilon \rho_{constr} \leq C, \\ & && (\mathbf{r}^T \text{diag}(\mathbf{y})P)_j \leq \rho_{obj}, \quad \forall j \in N, \\ & && (a_1 \cdot \text{diag}(\mathbf{y})P)_j \leq \rho_{constr}, \quad \forall j \in N, \\ & && 0 \leq y_j \leq 1, \quad \forall j \in N, \\ & && \rho_{obj}, \rho_{constr} : \text{free}. \end{aligned} \quad (3.25)$$

□

A parametric version when there are only two products under independent demands can be found in Appendix A, Example B.2.

3.3.1 Analysis of a Special Case: Independent Demand

When the demand of product j is assumed to be independent of the demand of other products, then P is a diagonal matrix with $P_{jj} = \sigma_j$ where σ_j is the standard deviation of demand for product j .

Then, the formulation in equation 3.25 is equivalent to the following LP model:

$$\begin{aligned}
z_1^*(\epsilon) = & \underset{\mathbf{y}, \rho_{obj}, \rho_{constr}}{\text{maximize}} && \sum_{j \in N} r_j y_j \bar{d}_j - \epsilon \rho_{obj} \\
& \text{subject to} && \sum_{j \in N} a_{1j} y_j \bar{d}_j + \epsilon \rho_{constr} \leq C, \\
& && r_j y_j \sigma_j \leq \rho_{obj}, && \forall j \in N, \\
& && a_{1j} y_j \sigma_j \leq \rho_{constr}, && \forall j \in N, \\
& && 0 \leq y_j \leq 1, && \forall j \in N, \\
& && \rho_{obj}, \rho_{constr} : \text{free}.
\end{aligned} \tag{3.26}$$

because $a_1 \cdot \text{diag}(\mathbf{y}) \bar{\mathbf{d}} = \sum_{j \in N} a_{1j} y_j \bar{d}_j$, $r^T \text{diag}(\mathbf{y}) \bar{\mathbf{d}} = \sum_{j \in N} r_j y_j \bar{d}_j$, and $(a_1 \cdot \text{diag}(\mathbf{y}) P)_j = a_{1j} y_j \sigma_j$.

Importantly, this linearized version of the model has a significant advantage as it can be solved using commercially available LP solvers.

Now, we explore a set of essential properties that apply to the specific scenario of interest. These proposed properties provide insights on key aspects of our discussion and play a pivotal role in our analysis. Not only these properties are important to understand the underlying complexities of the formulation, they provide reasoning for investigating the experimental performance of the formulations rather than their stochastic properties.

The range of values that ϵ can take for formulation (3.26) to yield non-trivial solutions depends on the correlation parameter and the expected demand of the products. It is worth noting that while ϵ can be assumed not to be in these ranges, the solutions outside of this interval may not be of particular interest for the given problem. Specifically, the range of values of ϵ , where the optimal solution is trivial and all zero is not of any interest.

Proposition 3.3.2. (i) If $\epsilon \geq \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$, the optimal solution is $y_j^* = 0$ for all $j \in N$. (ii) $z_1^*(\epsilon) \geq 0$ for any $\epsilon \geq 0$. (iii) If $\epsilon < \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$, then $z_1^*(\epsilon) \geq 0$.

Proof. We prove each part in turn. Part (i): We argue that if ϵ is selected within the given interval, objective function turns negative for non-zero values. Assume

that $\epsilon = \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$, and without loss of generality (WLOG) assume that there exists a solution $\mathbf{y}^*$ such that the capacity is sufficient for the given value of ϵ . Now, let i be a product index such that $r_i P_{ii} y_i^* \geq r_l P_{ll} y_l^* \forall l \in N$. Finally, under these assumptions, the objective function for this solution can be expressed as the following.

$$\begin{aligned} \sum_{j \in N} r_j \bar{d}_j y_j^* - \epsilon r_i P_{ii} y_i^* &= \sum_{j \in N} r_j \bar{d}_j y_j^* - \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} r_i P_{ii} y_i^*, \\ &= \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} (r_j y_j^* P_{jj} - r_i y_i^* P_{ii}), \\ &\leq 0, \end{aligned}$$

where the last inequality follows from the initial assumption that $r_i P_{ii} y_i^* \geq r_l P_{ll} y_l^* \forall l \in N$. Moreover, under the same assumptions, if $\epsilon = \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} + \gamma$, where $\gamma > 0$, the objective function can be written as:

$$\begin{aligned} \sum_{j \in N} r_j \bar{d}_j y_j^* - \epsilon r_i P_{ii} y_i^* &= \sum_{j \in N} r_j \bar{d}_j y_j^* - \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} r_i P_{ii} y_i^* - \gamma r_i P_{ii} y_i^*, \\ &= \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} (r_j y_j^* P_{jj} - r_i y_i^* P_{ii}) - \gamma r_i P_{ii} y_i^*, \\ &\leq 0. \end{aligned}$$

The last inequality follows from the previous assumptions and the non-negativity of the last term. Therefore, for any value of ϵ greater than $\sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$, non-zero solutions will yield non-positive objective values compared to the zero objective value of the solution with all decision variables equal to zero.

Part (ii): $y_j' = 0 \forall j \in N$, $\rho'_{constr} = \rho'_{obj} = 0$ is a feasible solution with $z_1'(\epsilon) = 0$, $\forall \epsilon \geq 0$.

Part (iii): If ϵ is selected within the given range, we argue that there exists a solution $\mathbf{y}^*$ such that the objective value is greater than 0. Assume that $\epsilon =$

$\sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} - \gamma$, where $0 < \gamma < \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$. Also, WLOG, assume that for such solution, the capacity is sufficient. This is due to the fact that for a solution vector, if the capacity is not sufficient, all of the elements of the decision variable vector could be decreased until the capacity criterion is met. Now, for the given solution vector, assume that there exists a product vector i such that $r_i P_{ii} y_i^* \geq r_l P_{ll} y_l^* \forall l \in N$. The objective function is stated as the following.

$$\begin{aligned} \sum_{j \in N} r_j \bar{d}_j y_j^* - \epsilon r_i P_{ii} y_i^* &= \sum_{j \in N} r_j \bar{d}_j y_j^* - \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} r_i P_{ii} y_i^* + \gamma r_i P_{ii} y_i^*, \\ &= \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} (r_j y_j^* P_{jj} - r_i y_i^* P_{ii}) + \gamma r_i P_{ii} y_i^*, \\ &\geq 0. \end{aligned}$$

The last inequality follows from the existence of a solution y' in which $\forall j \in N$ and $\forall i \in N$ we have $r_j y'_j P_{jj} = r_i y'_i P_{ii}$. If the existence of such a solution is proved, the optimal value of the problem when $\epsilon = \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} - \gamma$ is trivially observed to be non-negative.

For the last inequality, assuming such a solution will equate the first term of the left-hand side to zero. In addition, the function on the left-hand side has a strictly positive term if the price and the standard deviation of product with index i are not zero. Moreover, if both of the mentioned parameters are non-zero, the last inequality can be expressed by a strictly greater than sign instead of greater than equal to sign. To conclude this proof, we must show the existence of such solution.

Assume a solution where all terms satisfy $r_i y'_i P_{ii} > 0$ and are equal to each other. The left-hand side of the capacity constraint, for an arbitrary index i , can be written as $\sum_{j \in N} a_j \bar{d}_j y'_j + \epsilon a_i P_{ii} y'_i$. As we assume positivity of the terms, $y_j > 0 \forall j \in N$, if given left-hand side is larger than the available capacity, all of the elements of y' can be decreased by multiplying with a constant, i.e., new solution that will satisfy the capacity constraint and has the term equality condition will

be $y'' = y'\alpha$, where $\alpha = \frac{c}{\sum_{j \in N} a_j d_j y'_j + \epsilon a_i P_{ii} y'_i}$. As there exists a feasible solution where $\forall i \in N$ we have $r_j y'_j P_{jj} = r_i y'_i P_{ii}$, when $\epsilon < \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$, the problem has a non-negative optimal solution. $\square$

Proposition 3.3.3. *When $\epsilon < \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$, the objective function is decreasing in ϵ .*


Proof. In case that ϵ is increased, capacity constraint will have an extra positive term, let this increase in ϵ be γ , where $\gamma \geq 0$, and call $\epsilon' = \epsilon + \gamma$. For the optimal solution y^* when $\epsilon = \epsilon'$, WLOG, assume that there is a product index i such that $a_i P_{ii} y_i \geq a_l P_{ll} y_l \forall l \in N$.

$$\begin{aligned} \sum_{j \in N} a_j \bar{d}_j y_j^* + \epsilon' a_i P_{ii} y_i^* &= \sum_{j \in N} a_j \bar{d}_j y_j^* + \epsilon a_i P_{ii} y_i^* + \gamma a_i P_{ii} y_i^*, \\ &> \sum_{j \in N} a_j \bar{d}_j y_j^* + \epsilon a_i P_{ii} y_i^*. \end{aligned}$$

It can be deducted that, as the ϵ increases, any non-zero vector of y will yield a larger left-hand side value in the capacity constraint. In other words, if $\gamma > 0$ and at least one element of y is non-zero, we have, as the updated capacity, $c' = c - \gamma a_i P_{ii} y_i^*$. Due to non-negativity of $a_i P_{ii} y_i^*$, we also have $c > c'$. Based on the previous proposition, a decrease in the available capacity will decrease the objective value. $\square$

Proposition 3.3.4. *In a two-product problem, all of the capacity is utilized in the optimal solution if the equality $2\bar{d}_2 r_1 P_{11} + (\bar{d}_1 - C)r_2 P_{22} \geq 0$ holds, when $\epsilon < \sum_{j \in N} \frac{\bar{d}_j}{P_{jj}}$.*

Proof. We provide a graphical interpretation of the setting when the mentioned inequality does not hold. The interval $2\bar{d}_2 r_1 P_{11} + (\bar{d}_1 - c)r_2 P_{22} \geq 0$ is favorable as in case that the inequality does not hold, and if ϵ takes values between $\frac{\bar{d}_2}{P_{22}} \leq \epsilon \leq \frac{r_2 P_{22} c - \bar{d}_1 r_2 P_{22} - \bar{d}_2 r_1 P_{11}}{P_{22} P_{11} r_1}$, there will be solutions that do not utilize all of the capacity. Below graph is an example of such case.



In this graph, gray lines, which span out from the point at $y_1 = 1$ are the isoprofit lines where the other two lines which are black, depict the capacity constraint. If the given equality does not hold, for some ϵ values from the given interval, optimal solution will be the black point. Intuitively, increasing the value of y_2 decreases the objective value in the given interpretation, this is due to having ϵP_{22} multiplier being larger than $\bar{d}_j$ so that the coefficient of y_2 in the objective function is negative when $r_1 P_{11} y_1 \leq r_1 P_{11} y_1$. $\square$

Remark 3.3.5. *The generalization of Proposition (3.3.4) to problems with more than two classes is not given. The optimal solution may not utilize all of the capacity in multi-class problems as well.*

Proposition 3.3.6. *In the two-product problem, if $\bar{d}_1 + \bar{d}_2 \geq C$ and $2\bar{d}_2 r_1 P_{11} + (\bar{d}_1 - C)r_2 P_{22} \geq 0$ hold, the optimal solution is obtained in closed-form.*


- For $\epsilon < \frac{r_1 - r_2}{r_1 + r_2} \frac{\bar{d}_1}{P_{11}}$: $y_1^* = 1$, $y_2^* = \frac{C - \epsilon P_{11} - \bar{d}_1}{d_2}$.
- For $\frac{r_1 - r_2}{r_1 + r_2} \frac{\bar{d}_1}{P_{11}} \leq \epsilon < \epsilon^+$: $y_1^* = y_2^* \frac{P_{22}}{P_{11}}$, $y_2^* = \frac{C}{\bar{d}_1 \frac{P_{22}}{P_{11}} + \bar{d}_2 + \epsilon P_{22}}$, where ϵ^+ is the positive root of the quadratic equation in ϵ , $\epsilon^2 r_1 P_{11} P_{22} + \epsilon(r_1 \bar{d}_2 P_{11} - r_1 \bar{d}_1 P_{22}) + \bar{d}_1 \bar{d}_2 (r_2 - r_1) = 0$, which always has one non-negative root if $r_1 \geq r_2$.
- For $\epsilon^+ \leq \epsilon < \frac{\bar{d}_1}{P_{11}} + \frac{\bar{d}_2}{P_{22}}$: $y_1^* = y_2^* \frac{r_2}{r_1}$.
- For $\frac{\bar{d}_1}{P_{11}} + \frac{\bar{d}_2}{P_{22}} \leq \epsilon$: $y_1^* = y_2^* = 0$.

Proof. Critical points which the solution structure changes and the structure of the optimal solution are derived via graphs. Firstly, there are 2 inner maximization problems, $\max_{j \in N} r_j y_j P_{jj}$ and $\max_{j \in N} a_j y_j P_{jj}$. So, the outer problem can have 3 different structures. This is due to, if $y_1 \frac{P_{11}}{P_{22}} \geq y_2$, solutions to the inner maximization problems are known. Also, if $y_2 > y_1 \geq y_2 \frac{r_2}{r_1}$ inner problems have different solutions and if $y_2 \frac{r_2}{r_1} > y_1$ these problems result differently, leading 3 different solutions for the outer problem.

We now prove for the first case, $y_1 \frac{P_{11}}{P_{22}} \geq y_2$. Assuming that the optimal solution is known to have such relation, solving the problem for arbitrarily chosen price and demand parameters reduces to a simple Linear Programming problem with 2 variables and 1 constraint:

$$\begin{aligned} & \underset{0 \leq y \leq 1}{\text{maximize}} && r_1 y_1 \bar{d}_1 + r_2 y_2 \bar{d}_2 - \epsilon r_1 P_{11} y_1 \\ & \text{subject to} && \bar{d}_1 y_1 + \bar{d}_2 y_2 + \epsilon P_{11} y_1 \leq c. \end{aligned}$$

Below is the graphical interpretation for the reduced mathematical programming problem that was introduced for the $y_1 \frac{P_{11}}{P_{22}} \geq y_2$.




In this graph, the line passing through the origin, $y_1 \frac{P_{11}}{P_{22}} = y_2$, limits the relation which y_1 and y_2 will have in this case. The other dashed line passing through both depicted points is due capacity constraint. Shaded area is the feasible region

in the given y relation. Based on changing ϵ values, both the capacity constraint line and the isoprofit lines are expected to have changing slopes. Based on the relation between the slopes of these two lines, optimal solution will overlap with the circle or the square point on the graph. The latter line, which was created by the constraint has a slope of $\frac{\bar{d}_1 + \epsilon P_{11}}{d_2}$, while the isoprofit lines have the slope of $\frac{r_1 \bar{d}_1 - \epsilon r_1 P_{11}}{r_2 d_2}$. We now equate these slopes:

$$\begin{aligned}\frac{\bar{d}_1 + \epsilon P_{11}}{d_2} &= \frac{r_1 \bar{d}_1 - \epsilon r_1 P_{11}}{r_2 d_2} \\ r_2 \bar{d}_2 \bar{d}_1 + \epsilon r_2 P_{11} \bar{d}_2 &= r_1 \bar{d}_2 \bar{d}_1 - \epsilon r_1 P_{11} \bar{d}_2 \\ (r_1 + r_2) \bar{d}_2 \epsilon P_{11} &= \bar{d}_2 \bar{d}_1 (r_1 - r_2) \\ \epsilon &= \frac{r_1 - r_2}{r_1 + r_2} \frac{\bar{d}_1}{P_{11}}\end{aligned}$$

Based on this finding, if the ϵ is equal to the value above, line between square and the circle points on the graph are alternative optima. Moreover, note that, if ϵ is larger than this value, isoprofit lines will have steeper slopes while the constraint line is less steep. So, larger ϵ values will make the square point optimal, whereas, if ϵ is smaller than this threshold, the circle point will be the optimal point.

On the second interval, while the objective function remains the same, constraint changes compared to the first interval for possible ϵ values. In this interval, the constraint is $\bar{d}_1 y_1 + \bar{d}_2 y_2 + \epsilon y_2 P_{22} \leq c$. Now there is a similar problem to solve, as plotted below.



On the above graph, shaded region denotes the feasible region for the given interval. If the slope of constraint line and the isoprofits lines are equal, the line segment between the circle and the square points yield the same objective value. So, we will equate the slopes of these lines:

$$\frac{r_1 \bar{d}_1 - \epsilon r_1 P_{11}}{r_2 \bar{d}_2} = \frac{\bar{d}_1}{d_2 + \epsilon P_{22}}$$

$$r_1 \bar{d}_1 \bar{d}_2 - \epsilon r_1 \bar{d}_2 P_{11} + \epsilon r_1 \bar{d}_1 P_{22} - \epsilon^2 r_1 P_{11} P_{22} = r_2 \bar{d}_1 \bar{d}_2$$

$$\epsilon^2 r_1 P_{11} P_{22} + \epsilon(r_1 \bar{d}_2 P_{11} - r_1 \bar{d}_1 P_{22}) + \bar{d}_1 \bar{d}_2 (r_2 - r_1) = 0$$

We now investigate the determinant of this equation, show that it has 2 real roots. Moreover, assuming the function has 2 real roots, we have $\bar{d}_1 \bar{d}_2 (r_2 - r_1) < 0$ as $r_1 \geq r_2$ and $r_1 P_{11} P_{22} > 0$, one of the roots is non-positive and the other is non-negative valued. This is due the fact that multiplication of the roots is equal to dividing the intercept variable by the coefficient of the quadratic term, which are positive and negative values, respectively. Now we show that the determinant is non-negative, i.e., there in fact exist 2 real roots.

$$\begin{aligned}
& r_1^2 \bar{d}_2^2 P_{11}^2 - 2r_1^2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} + r_1^2 \bar{d}_1^2 P_{22}^2 - 4r_1 r_2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} + 4r_1^2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} \\
&= r_1^2 \bar{d}_2^2 P_{11}^2 + 2r_1^2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} + r_1^2 \bar{d}_1^2 P_{22}^2 - 4r_1 r_2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} \\
&\geq r_1^2 \bar{d}_2^2 P_{11}^2 + r_1^2 \bar{d}_1^2 P_{22}^2 - 2r_1 r_2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} \\
&\geq r_2^2 \bar{d}_2^2 P_{11}^2 + r_1^2 \bar{d}_1^2 P_{22}^2 - 2r_1 r_2 \bar{d}_1 \bar{d}_2 P_{11} P_{22} \\
&= (r_2 \bar{d}_2 P_{11} - r_1 \bar{d}_1 P_{22})^2 \\
&\geq 0
\end{aligned}$$

The positive real root of the equation under investigation corresponds to ϵ^+ . So, utilizing the discussion as above, if ϵ is less than this value, based on the slopes of the lines mentioned, the square point is the optimal solution. If ϵ is larger, model tends to choose the circle point, structure of which is given in the next interval for ϵ . Also note that, in Proposition (3.3.2), we have proved that if $\frac{\bar{d}_1}{P_{11}} + \frac{\bar{d}_2}{P_{22}} \leq \epsilon$, optimal values are $y_1^* = y_2^* = 0$. $\square$

3.3.2 Solution Structures Under ℓ_1 -norm

Under the ℓ_1 -norm, the formulation presents unique challenges due to the presence of the maximum operator in both the objective function and the constraints, due to these operators, changing the size of the uncertainty set ϵ , greatly affects the optimal solution and creates different structures of closed-form optimal solutions. To provide context on these intricate cases, we offer insights by delving into a specific scenario, the two-product single-resource case where the matrix P is diagonal, i.e. the demand for each product is independent of the other. Subsequently, in this thesis, we refer to graph-based solutions to establish the proofs underpinning the properties of the proposed formulation, which is similar to the investigation of this subsection. These graphical representations provide a tangible framework for elucidating the nuances of the problem and aid in comprehending its underlying complexities.

Firstly, we discuss the improvement directions and the concept of isoprofit lines

pertinent to the problem at hand. It is essential to recognize that the configuration of isoprofit lines is contingent upon the ℓ_∞ -norm on the objective function. In the mentioned special case, we have $\|\mathbf{r}^T \text{diag}(\mathbf{y})P\|_\infty = \max \{r_1 P_{11}y_1, r_2 P_{22}y_2\}$. Visualizing this in a two-dimensional graph, the isoprofit lines and their orientations dynamically evolve based on the interplay between the variables y_1 and y_2 . This relationship is defined by the line equation $y_1 \frac{r_1 P_{11}}{r_2 P_{22}} = y_2$, which is denoted by the dashed line in the figure below.




Figure 3.1: Graphical Interpretation of Changing Isoprofit Line

In the graph above lines indicated by the legend entries are the isoprofit lines that are spanned by two cases of the maximization term in the objective. When examining a given solution vector, $\mathbf{y}^*$, if it is known that $y_1^* \frac{r_1 P_{11}}{r_2 P_{22}} \geq y_2^*$, it becomes possible to express the solution of the maximum operator within the objective function. Then, isoprofit lines for the respective feasible region becomes well-defined, as depicted by the densely dotted line on Figure 3.1.

However, when $y_1^* \frac{r_1 P_{11}}{r_2 P_{22}} < y_2^*$, the maximum operator undergoes a transformation, leading to a different linear equation within the objective function, resulting in isoprofit lines that represented by the line with dashes and dots. Consequently, solutions with the former relationship, situated above the dashed line, exhibit improvements directed along the arrow on the isoprofit line that was mentioned first, whereas solutions characterized by the latter relationship progress according to

the arrow on the second line in Figure 3.1.

Transitioning to the constraint-specific relations, the formulation incorporates a piecewise linear constraint predicated on the following relationship among the decision variables $y_1 \frac{P_{11}}{P_{22}} = y_2$, using the assumption that $a_1 = a_2 = 1$. Notably, under the assumption of $r_1 \geq r_2$, the slope of the boundary where the isoprofit lines intersect is steeper than the slope of the boundary corresponding to the constraint line.




Figure 3.2: Graphical Interpretation of Changing Constraint Line

In Figure 3.2, dashed line is a representation of $y_1 \frac{P_{11}}{P_{22}} = y_2$, while the densely dashed and dash-dot lines are capacity constraint lines. Similar to the previous discussion, if a solution vector has a known position with respect to the real line divided by $y_1 \frac{P_{11}}{P_{22}}$ and y_2 , the feasible region can be expressed as the solution to the maximization operator is known. Based on the observations from these two graphs, we conclude that for different pairs of y_1 and y_2 , the feasible region of the problem change. This is due to the fact that these decision variables change the inner optimization problem results, hence change the feasible region as well as the improvement direction. Mainly, for the 2-product single-resource case, there are 3 different feasible regions, which are depicted below.

The numbered regions on Figure 3.3 depicts the feasible regions when a given




Figure 3.3: Graphical Interpretation of Possible Feasible Regions

solution vector satisfies additional constraints given for each region: $y_1 \frac{P_{11}}{P_{22}} < y_2$ for the first region; $y_1 \frac{r_1 P_{11}}{r_2 P_{22}} < y_2$, $y_1 \frac{P_{11}}{P_{22}} \geq y_2$ for the second region and $y_1 \frac{r_1 P_{11}}{r_2 P_{22}} \geq y_2$ for the last region. Basically, the optimal solution to *Region 2* can be found using a reformulated version of (3.25), which is the following LP.

$$\begin{aligned}
z_1^*(\epsilon) = \quad & \text{maximize} \quad r_1 \bar{d}_1 y_1 + r_2 \bar{d}_2 y_2 + -\epsilon r_1 P_{11} y_1 \\
& \text{subject to} \quad a_{11} \bar{d}_1 y_1 + a_{12} \bar{d}_2 y_2 + \epsilon a_{12} P_{22} y_2 \leq C, \\
& \quad y_1 \frac{r_1 P_{11}}{r_2 P_{22}} < y_2, \\
& \quad y_1 \frac{P_{11}}{P_{22}} \geq y_2, \\
& \quad 0 \leq y_1 \leq 1, \\
& \quad 0 \leq y_2 \leq 1.
\end{aligned} \tag{3.27}$$

Similarly, for the other regions, a model could be written in a linear form while adding constraints that define the relation between y_1 and y_2 , changing the second and the third constraint of (3.27). For this simple example, solutions of the three proposed model could be compared, and the solution vector which yields the highest objective value as a result to these problems will be the optimal solution to the original problem. However, complexity of the problem increases with increase in the number of products. As it is challenging to define all the possible solution structures or their existence for the problem, we refer to the

simulation studies to show the existence of different booking limit controls and optimal y_i values that could be obtained using different ϵ values.

While the discussion here could be a starting point for solving large instances, based on numerous linear programs on different possible relations between the y_i values, and comparing the objective values for each, we propose a different method for handling large instances. Mainly, as the non-linear problem is trivially linearized by adding new variables and constraints, solving the linear programming equivalent of the robust formulations is not challenging at all to call for such approach.

3.4 Budgeted Uncertainty Set

Having discussed uncertainty sets that are generated by two norms exclusively, we now consider their intersection. The budgeted uncertainty set is more flexible in adjusting the conservativeness of the results [31]. This set is also called the budget uncertainty set when first investigated. In the previous cases, we had only one variable to adjust the size of the uncertainty set and as discussed many times in the literature, one disadvantage is the limited control on the shape and the size of the uncertainty set. Budgeted uncertainty set employs two parameters and has more control.

Another advantage of the budgeted uncertainty set is that it combines the interpretations of two previous uncertainty sets to guide a decision maker in deciding the uncertainty parameters. The cardinality constrained uncertainty set introduced in [31], allows at most a prespecified number of constraints to deviate from its nominal value, enabling decision maker to have strong control on the conservativeness of the formulation. While the cardinality constrained uncertainty set uses ℓ_0 -norm, we utilize ℓ_1 -norm to account for determining the amount of possible deviations. $\mathcal{U}_B$ denotes the budgeted uncertainty set employed in this work.

$$\mathcal{U}_B = \{\bar{\mathbf{d}} + P\mathbf{u} : \|\mathbf{u}\|_\infty \leq \rho, \|\mathbf{u}\|_1 \leq \Gamma\}. \quad (3.28)$$

The uncertainty set (3.28) have the same robust counterpart as the cardinality constrained uncertainty set built using ℓ_0 and ℓ_∞ -norms [36]. In (3.28), Γ corresponds to the maximum number of coefficients that can deviate from its nominal value. We use ρ to denote the set size control with respect to ℓ_∞ -norm.

Given the uncertainty set, the corresponding constraint is the following. Moreover, for the discussions under this subsection, we only provide robust counterparts for the single-leg setting, and assume that the incidence matrix is only a vector of ones and can be ignored.

Proposition 3.4.1. *Consider the constraint and the uncertainty budgets,*

$$(\bar{\mathbf{d}} + P\mathbf{u})\mathbf{y} \leq C, \quad \forall \mathbf{u} : \|\mathbf{u}\|_\infty \leq \rho, \|\mathbf{u}\|_1 \leq \Gamma. \quad (3.29)$$

If there exists a vector $\mathbf{x}$ that satisfies:

$$\bar{\mathbf{d}}^T \mathbf{y} + \rho \|\mathbf{x}\|_1 + \Gamma \|P^T \mathbf{y} - \mathbf{x}\|_\infty \leq C, \quad (3.30)$$

then the constraint in (3.29) is equivalent to (3.30).

Proof. We replicate the approach described in [31] in this proof. We know that if (3.29) holds, the following constraint holds:

$$\max_{\mathbf{u} : \|\mathbf{u}\|_\infty \leq \rho, \|\mathbf{u}\|_1 \leq \Gamma} (\bar{\mathbf{d}} + P\mathbf{u})\mathbf{y} \leq C, \quad (3.31)$$

Consider the maximization problem on the left-hand-side of the constraint. The dual of that problem can be formulated as follows.

$$\min_{\mu \geq 0, \lambda \geq 0} \left\{ \max_{u \in \mathbb{R} : u \geq 0} \{ (\bar{\mathbf{d}} + Pu)^T \mathbf{y} + \mu(\rho - \|u\|_\infty) + \lambda(\Gamma - \|u\|_1) \} \right\} \quad (3.32)$$

Moreover, for (3.32), we replace the norm operators with corresponding minimization problems for later use.

$$\min_{\mu \geq 0, \lambda \geq 0} \left\{ \bar{\mathbf{d}}^T \mathbf{y} + \mu\rho + w\Gamma + \max_{u \in \mathbb{R} : u \geq 0} \left\{ (Pu)^T \mathbf{y} - \mu \min_{\|l_1\|_1 \leq 1} l_1^T u - \lambda \min_{\|l_2\|_\infty \leq 1} l_2^T u \right\} \right\} \quad (3.33)$$

For this minimization problem, minimization terms generated from ℓ_1 and ℓ_∞ -norms and the inner maximization term can be interchanged. This is allowed based on a corollary discussed in [45]. As the functions of inner-most minimization terms are continuous concave-convex functions on their effective domains, and the feasible regions of both the maximization and the minimization problems are bounded sets, in short, we have “min max” = “max min”.

The inner minimization terms can also be interchanged with maximization and expressed under the same outer minimization problem, leading to the following problem which is equivalent to the dual problem in equation (3.32):

$$\min_{\mu \geq 0, \lambda \geq 0, \|l_1\|_1 \leq 1, \|l_2\|_\infty \leq 1} \left\{ \bar{d}^T y + \mu \rho + \lambda \Gamma + \max_{u \in \mathbb{R}: u \geq 0} \{ (Pu)^T y - \mu l_1^T u - \lambda l_2^T u \} \right\} \quad (3.34)$$

Now, the inner maximization term is fixed to be zero. As otherwise, maximization over u will tend to infinity as well as the outer minimization problem. We started with a dual of a maximization problem in the constraints, if the outer minimization term in the above formulation tends to infinity, so will the left-hand side in (3.31), meaning that the constraint is infeasible. Adding the discussed constraints, the problem can be reformulated as the following optimization problem.

$$\begin{aligned} \min \quad & \bar{d}^T y + \mu \rho + \lambda \Gamma \\ \text{subject to} \quad & P^T y - \mu l_1 - \lambda l_2 = 0, \\ & \|l_1\|_1 \leq 1, \\ & \|l_2\|_\infty \leq 1, \\ & \mu, \lambda \geq 0. \end{aligned} \quad (3.35)$$

Let $\mu l_1 = \beta_1$ and $\lambda l_2 = \beta_2$ as a change of variables in this formulation. Then, the second constraint is now written as $\|\beta_1\|_1 \leq \mu$, and the third one is $\|\beta_2\|_\infty \leq \lambda$; while the first one is simply $P^T y - \beta_1 - \beta_2 = 0$. The domain of these new variables is unrestricted. This due to the domains of μ and l_1 for β_1 ; λ and l_2 for β_2 . As μ and λ are non-negative numbers; l_1 and l_2 can take positive and negative values, their respective multiplications will be unrestricted.

Note that the minimization problem has the objective function $\bar{d}^T y + \mu\rho + \lambda\Gamma$, where ρ and Γ are positive coefficients. So, μ and λ will take their lowest possible values, which in this case equal to $\|\beta_1\|_1$ and $\|\beta_2\|_\infty$, respectively.

$$\begin{aligned} \min \quad & \bar{d}^T y + \rho\|\beta_1\|_1 + \Gamma\|\beta_2\|_\infty \\ \text{subject to} \quad & P^T y - \beta_1 - \beta_2 = 0, \\ & \beta_1, \beta_2 : \text{free}. \end{aligned} \tag{3.36}$$

We substitute $\beta_2 = P^T y - \beta_1$, the problem can be reformulated under one decision variable which is unrestricted. So, the derived problem is equivalent to (3.32), and given below.

$$\min_{\beta: \text{free}} \bar{d}^T y + \rho\|\beta\|_1 + \Gamma\|P^T y - \beta\|_\infty \tag{3.37}$$

□

Note that, to ensure a clear and concise representation of the formulation, we now introduce a notation for new parameters. Specifically, we define vector $\bar{d}' = \bar{d} \text{diag}(\mathbf{r})$ and matrix P' where the i -th row, j -th column correspond to $(P_{ij}r_i)$.

Proposition 3.4.2. *If the uncertainty set $\mathcal{U}_B$ is used, then the robust formulation in equation (3.2), is equivalent to the following linear program:*

$$\begin{aligned} z_B^*(\Gamma, \rho) = \quad & \underset{\mathbf{y}, \mathbf{t}, \mathbf{w}, \mathbf{x}, \mathbf{x}', \rho_{obj}, \rho_{constr}}{\text{maximize}} \quad \sum_{j \in N} r_j y_j \bar{d}_j - \rho \sum_{j \in N} t_j - \Gamma \rho_{obj} \\ & \text{subject to} \quad \sum_{j \in N} a_{1j} y_j \bar{d}_j + \rho \sum_{j \in N} w_j + \Gamma \rho_{constr} \leq C, \\ & (P'^T y + \beta')_j \leq \rho_{obj}, \quad \forall j \in N, \\ & (P^T y - \beta)_j \leq \rho_{constr}, \quad \forall j \in N, \\ & -t_j \leq \beta'_j \leq t_j, \quad \forall j \in N, \\ & -w_j \leq \beta_j \leq w_j, \quad \forall j \in N, \\ & 0 \leq y_j \leq 1, \quad \forall j \in N, \\ & t_j, w_j \geq 0, \quad \forall j \in N, \\ & \rho_{obj}, \rho_{constr} : \text{free}, \quad \forall j \in N, \\ & \beta_j, \beta'_j : \text{free}, \quad \forall j \in N. \end{aligned} \tag{3.38}$$

Proof. Based on Proposition (3.4.1), we now find the linear equivalent of the inner minimization term related to the objective function of the formulation provided in equation (3.2), under the budgeted uncertainty set given in equation (3.28). This problem can be reformulated by finding its dual using a similar approach as before. The dual of the problem will yield a maximization problem, which is directly incorporated under the outer maximization problem. After finding linear counterparts of the terms introduced by the robust optimization approach, we report the following non-linear model as a robust counterpart to the proposed formulation under (3.28).

$$\begin{aligned} z_B^*(\Gamma, \rho) = & \underset{0 \leq \mathbf{y} \leq \mathbf{1}; \beta, \beta': \text{free}}{\text{maximize}} && \bar{d}^T \mathbf{y} - \rho \|\beta'\|_1 - \Gamma \|P'^T \mathbf{y} + \beta'\|_\infty \\ & \text{subject to} && \bar{d}^T \mathbf{y} + \rho \|\beta\|_1 + \Gamma \|P^T \mathbf{y} - \beta\|_\infty \leq C. \end{aligned} \quad (3.39)$$

Without the inner optimization problems, the formulation is, still, non-linear. Similar to previous sections, its equivalent linear programming formulation can be found by introducing additional decision variables that were used to account for the non-linear terms. $\square$

A parametric version when there are only two products under independent demands can be found in Appendix A, Example B.3.

3.5 Interval Uncertainty

In this section, we investigate the interval uncertainty for the demand of each product. Note that, if the uncertainty set is described as an interval of demand for each product or class, dependency among the demand of different classes will be ignored and demand will be assumed to be statistically independent for each class. We build the uncertainty set using the given upper bounds and lower bounds for the demand for each product:

$$\mathcal{U}_{int} = \{\mathbf{l} \leq \mathbf{d} \leq \mathbf{v}\} \quad (3.40)$$

where, $\mathbf{l}$ and $\mathbf{v}$ stand for vectors for the lower and upper bounds on the demand for each product, respectively. Similar to the previous chapters, $\mathbf{d}$ can be expressed as $\bar{\mathbf{d}} + P\mathbf{u}$ where the vector $\mathbf{u}$ explains uncertainty. We now write the mean demand for each product as $\bar{d}_j = \frac{l_j + v_j}{2}$. Also, the correlation parameter matrix is expressed as $P = \text{diag}(\frac{\mathbf{v}-\mathbf{l}}{2})$ [36]. Finally, the resulting uncertainty set, given the described parameters is the following.

$$\mathcal{U}_{int} = \{\bar{\mathbf{d}} + P\mathbf{u} : -\mathbf{1} \leq \mathbf{u} \leq \mathbf{1}\} \quad (3.41)$$

Proposition 3.5.1. *If the uncertainty set $\mathcal{U}_{int}$ is used, then the robust formulation in equation (3.2) is equivalent to the following linear program:*

$$\begin{aligned} z_{int}^* = & \underset{\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}}{\text{maximize}} && \sum_{j \in N} r_j y_j l_j \\ & \text{subject to} && \sum_{j \in N} a_{1j} y_j v_j \leq C. \end{aligned} \quad (3.42)$$

Proof. Rewriting the original mathematical model based on this uncertainty set, we simply have the following structure after the substitution $\mathbf{d} = \bar{\mathbf{d}} + P\mathbf{u}$ is implemented:

$$\begin{aligned} & \underset{\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}}{\text{maximize}} && \left\{ \underset{-\mathbf{1} \leq \mathbf{u} \leq \mathbf{1}}{\text{minimize}} \sum_{j \in N} r_j y_j (\bar{\mathbf{d}} + P\mathbf{u})_j \right\} \\ & \text{subject to} && \left\{ \max_{-\mathbf{1} \leq \mathbf{u} \leq \mathbf{1}} \sum_{j \in N} a_{1j} y_j (\bar{\mathbf{d}} + P\mathbf{u})_j \right\} \leq C. \end{aligned} \quad (3.43)$$

Intuitively, minimization problem in the objective function and the maximization problem in the constraint are trivial: As P , $\mathbf{r}$ and $\mathbf{y}$ are non-negative vectors and matrices, the solution vectors to both problem will be $\mathbf{u}^* = -\mathbf{1}$ in the objective and $\mathbf{u}^* = \mathbf{1}$ in the constraints.

Using the definitions of the parameters, we first provide a linear formulation for the objective function. As all of the uncertainty variables in this inner problem are said to take the value of -1 , the linearized formulation of the objective function

of (3.43) is derived below.

$$\begin{aligned}
\text{minimize}_{-1 \leq \mathbf{u} \leq 1} \sum_{j \in N} r_j y_j (\bar{\mathbf{d}} + P\mathbf{u})_j &= \sum_{j \in N} r_j y_j (\bar{d}_j - P_{jj}) \\
&= \sum_{j \in N} r_j y_j \left(\frac{l_j + v_j}{2} - \frac{v_j - l_j}{2} \right) \\
&= \sum_{j \in N} r_j y_j l_j
\end{aligned} \tag{3.44}$$

Using a similar argument, one can find that the LHS of the constraints are now equal to $\sum_{j \in N} a_{1j} y_j u_j$, $\forall i \in K$. Rewriting the closed-form solutions to the inner optimization problems, we find the LP equivalent of (3.43), presented below.

$$\begin{aligned}
z_{Int}^* &= \text{maximize}_{\mathbf{0} \leq \mathbf{y} \leq 1} \sum_{j \in N} r_j y_j l_j \\
&\text{subject to } \sum_{j \in N} a_{1j} y_j v_j \leq C.
\end{aligned} \tag{3.45}$$

□

A parametric version when there are only two products under independent demands can be found in Appendix A, Example B.4.

Remark 3.5.2. *LP given in Formulation (3.45) is a continuous (or fractional) knapsack problem, that is easily solved [46]. The classes are ordered by $\frac{r_j l_j}{a_{1j} v_j}$. In the descending order, products are included in the knapsack as long as there is enough capacity.*

3.6 Observations on Effect of Uncertainty

The objective function value of the robust formulation decrease in C as we proved in proposition (3.1.1). When no uncertainty is assumed, objective values of the robust models are equal to the objective value of DLP. As the uncertainty increases, objective values of the robust formulations decrease. We illustrate this with an example below.

This example has four products with $\mathbf{r}^T = (1050, 567, 534, 520)$, $\bar{\mathbf{d}}^T = (17.3, 45.1, 39.6, 34)$, $C = 100$, $\sigma^T = (5.8, 15, 13.2, 11.3)$, and non-diagonal elements of P equal to 0 [1]. The uncertainty set parameters are changing from 0 to 15. The objective function values of DLP and our robust models are plotted in Figure 3.4.




Figure 3.4: Objective Value Comparison for Four-Class Example Under Increasing Uncertainty Set Sizes.

3.7 Finding Booking Limits

While we have discussed the properties of the model and the structure of the solution vector, it is important to note that the decision variables represented by y may not inherently align with direct interpretations in the context of booking controls. To make practical sense of the optimal solutions, we must convert these decision variables into booking controls, specifically, booking limits or protection levels.

Our proposed methodology involves deriving booking limits for products based on the optimal values $\mathbf{y}^*$, as we interpret the solutions as representing the sales volume by product, computed as a percentage of the estimated (or mean) demand. Below, we give two algorithms that translate the results of our formulations into actionable booking controls.

We have discussed the use of DLP in the previous chapters, using the solutions of the MLP model given in formulation (3.1) is similar. When the MLP is solved, expected demand values are multiplied with optimal $\mathbf{y}^*$ values. Each multiplication, $\bar{d}_j y_j^*$, corresponds to the booking limit of the products. As there is an equivalence between DLP and MLP, after variable manipulation, both methods find the same booking limits.

The first booking controls we investigate are partitioned booking limits. In partitioned booking limits, products are treated individually. Each available product j is given a specific booking limit b_j^P . Moreover, while these limits are used for a single product, changes in number of accepted bookings of product j does not affect the booking limit of other products. This way, capacity is “partitioned” and each product gets resources dedicated to it. Capacity allocated to a product, if unused, cannot be used to meet the demand of other products when capacity is partitioned.

Algorithm PL Finding Partitioned Booking Limits

- Step 1: Solve the LP equivalent of the robust seat-allocation problem to find the optimal values $\mathbf{y}^*$.
 - Step 2: For each product $j \in N$, calculate $\bar{d}_j y_j^*$.
 - Step 3: Compute the multiplier ψ : $\psi = \frac{C}{\sum_{j \in N} \bar{d}_j y_j^*}$.
 - Step 4: The partitioned booking limit b_j^P of class j is computed as: $b_j^P = \psi \bar{d}_j y_j^*, \forall j \in N$.
-

Above algorithm, works similar to converting MLP solutions to booking limits.

However, in the robust counterparts, we introduce a penalty for positive y_j values. After such penalty, booking limits of the products may not necessarily use up all the capacity. We normalize the protection levels for each product by creating an artificial constant ψ that allocates unused capacity space proportional to the products. While any unused capacity can be allocated based on various rules or methods, using the normalization with ψ , we are able to protect the relative ratios of booking limits obtained from the robust model while using all of the capacity.

The second type of booking controls are nested booking limits. As mentioned earlier, we assume that products are ordered by fare, thus nesting refers to nesting-by-fare. Through nesting, decision makers are able to sell the unused capacity that is allocated to low-fare classes to meet the demand of high-fare classes. The nested booking limit for product j is denoted as b_j^N and is obtained using the algorithm below.

Algorithm NL Finding Nested Booking Limits

- Step 1: Solve the LP equivalent of the robust seat-allocation problem to find the optimal values $\mathbf{y}^*$.
 - Step 2: For each product $j \in N$, calculate $\bar{d}_j y_j^*$.
 - Step 3: Compute the multiplier ψ : $\psi = \frac{C}{\sum_{j \in N} \bar{d}_j y_j^*}$.
 - Step 4: Compute the partitioned booking limits as $b_j^P = \psi \bar{d}_j y_j^*$, $\forall j \in N$.
 - Step 5: The nested booking limit on products $\{j, \dots, n\}$ is calculated using $b_j^N = \sum_{i=j}^{|N|} b_i^P$.
-

The first step in both algorithms requires an LP to be solved. However, in some instances, solved LPs might yield alternative optima, which can also occur in solving DLP. In the context of airlines, choosing one of the alternative optima is made based on the strategic business decisions, for example, focusing or allocating more to a specified customer type. However, we simply use the solution vectors

that the LP solver yields.

3.7.1 Illustration

To illustrate how the booking limits of the proposed algorithms are calculated, we rely on data from a single-resource four-product example from [1]; this example is commonly used in the literature. The demand information is summarized in the following table. Capacity of the resource is $C = 100$.

j	r_j	$\bar{d}_j$	σ_j
1	1050	17.3	5.8
2	567	45.1	15
3	534	39.6	13.2
4	520	34	11.3

Table 3.1: Single-resource, four-class example [1].

The optimal solutions for all of the discussed methods depend on the selected level of uncertainty set size. To better illustrate the algorithm, we present a booking limit policy for a fixed parameter ϵ and one of the LP formulations for step 1 of the algorithm. Discussed algorithms are investigated under changing set sizes in the next chapter. We use *Algorithm-NL₁* with $\epsilon = 11$, and find the booking limits that correspond to this particular robust solution.

(Step 1) Given the parameters, we formulate the problem in equation (3.25). The optimal solution vector is $(y_1^*, y_2^*, y_3^*, y_4^*) = (0.403, 0.289, 0.349, 0.418)$.

(Step 2) Calculate $\bar{\mathbf{d}}\mathbf{y}^* = (6.979, 13.028, 13.803, 14.216)$.

(Step 3) Compute the total: $T = 6.979 + 13.028 + 13.803 + 14.216 = 48.026$, and calculate the ratio: $\psi = \frac{100}{48.026} = 2.082$.

(Step 4) Calculate $\mathbf{b}^P = \psi \bar{\mathbf{d}}\mathbf{y}^* = (14.5, 27.1, 28.7, 29.6)$.

(Step 5) Calculate $b_j^N = \sum_{i=j}^4 b_i^P$. $\mathbf{b}^N = (100, 85.4, 58.3, 29.6)$.

In most cases, the LP formulations differ in exact $\mathbf{y}^*$, however, the relative ratio of the values of the optimal solutions remain the same. Hence, while based on possible uncertainty set sizes, there are many optimal solutions, most correspond to the same booking limits. We conclude this section with a result on two different solutions to robust formulations that yield the same booking limits.

Proposition 3.7.1. *Suppose $\mathbf{y}^*$ and $\tilde{\mathbf{y}}^*$ are the optimal solutions to two robust formulations that differ in their uncertainty sets, and $\mathbf{y}^* \neq \tilde{\mathbf{y}}^*$. The partitioned (nested) booking limits obtained by Algorithm-PL (NL) using $\mathbf{y}^*$ are equal to those obtained by using $\tilde{\mathbf{y}}^*$ if and only if $\frac{y_i^*}{y_j^*} = \frac{\tilde{y}_i^*}{\tilde{y}_j^*}$ for all $i \in N, j \in N$.*

Proof. If we have $\frac{y_i^*}{y_j^*} = \frac{\tilde{y}_i^*}{\tilde{y}_j^*}$ for $i \neq j$ then we also have $\frac{y_i^*}{\tilde{y}_i^*} = \frac{y_j^*}{\tilde{y}_j^*}$, meaning that there exists a non-negative coefficient, call α , such that $\alpha = \frac{y_j^*}{\tilde{y}_j^*}$ for all $j \in N$.

Consequently, we have, $\sum_{j \in N} \bar{d}_j y_j^* = \alpha \sum_{j \in N} \bar{d}_j \tilde{y}_j^*$. Then, immediately, the multiplier ψ^* computed in of step 3 of Algorithm-PL (NL) corresponding to $\mathbf{y}^*$ and $\tilde{\psi}$ corresponding to $\tilde{\mathbf{y}}$ has to satisfy $\psi^* = \alpha \tilde{\psi}$.

As a result, $b_j^P = \psi^* \bar{d}_j y_j^* = \alpha \tilde{\psi} \bar{d}_j y_j^* = \tilde{\psi} \bar{d}_j (\alpha y_j^*) = \tilde{\psi} \bar{d}_j (\tilde{y}_j^*) = \tilde{b}^P$. Moreover, since the partitioned booking limits $b_j^P = \tilde{b}_j^P$, the nested booking limits obtained in step 5 of Algorithm-NL also satisfy $b_j^N = \tilde{b}_j^N$ for all $j \in N$. $\square$

Chapter 4

Computational Experiments

In this chapter, we compute the performance of the booking limits obtained from our robust formulations using simulations. We also provide insights on the choice of uncertainty set parameters and benchmark our methods against other well-known methods. In this chapter, we strive to find the effectiveness of the robust solutions. While a stochastic analysis on the expected revenue to be achieved could be theoretically argued, we are interested in the robustness of the solutions as well as the mean revenues achieved. The robust formulations yield booking limit controls through an algorithm, and the performance of resulting controls are not easily understandable. So, we first investigate the booking limit controls under changing uncertainty set sizes to find a better performing parameter for these sets. Then, using methods which are extensively studied, we benchmark our methods to provide insights regarding their possible practical applications.

We first introduce the policies that will be investigated. Then, simulation setups are discussed, including the number of simulation runs in each experiment and the methods of mimicking flight booking sequences. In the next section, used metrics for the performances of the methods are listed as well as their importance and interpretations. Before starting to benchmark the robust formulations, we suggest selecting parameters for the uncertainty set sizes for each of them. In

the first section of experiments, we examine the two-product setting under independent and dependent demand assumptions. Providing insights into the effect of fare ratio and demand factor on the robust models. The second set of experiments includes similar experiments for multi-product problems, i.e. $n > 2$. Then we investigate how the methods explained in the first section perform when the mean of the demand distributions are over or underestimated. Finally, we conclude the chapter with a summary of our observations.

4.1 The Policies Tested in the Experiments

In this section, we introduce the methods that we use in the computational experiments. New robust formulations are compared to EMSRb and DLP methods. While both are heuristics, they are commonly used in the literature as benchmarks. Booking limit control derivation using DLP was discussed in the previous chapters; for a complete implementation of EMSRb, see [1]. Also, in two-product examples, Littlewood's Rule is used to find the optimal controls, which is the equivalent of EMSRb when $n = 2$. In the list below, we present the notation for the booking limit controls derived from the robust formulations.

- *Algorithm-NL_∞*: This algorithm computes nested booking limits using *Algorithm-NL*, where the LP equivalent of the formulation under ℓ_∞ -norm, given in (3.20) is solved in step 1.
- *Algorithm-NL₁*: Computes nested booking limits; this algorithm solves the LP in (3.25).
- *Algorithm-NL_B*: In this case, the LP problem addressed in step 1 of the second algorithm is (3.38). Also, note that, the subscript of B denotes the selected ρ level.
- *Algorithm-NL_I*: Although not explored as extensively as other methodologies, this algorithm corresponds to solving the LP problem (3.45) during the first step.

4.2 Simulation Setup

We simulate the booking process for commercial flights in each of our experiments. At each instance, we randomly generate the total number of booking requests for each class based on the specified demand distributions and parameters. Then, we determine the sequence in which the booking requests arrive. Moreover, the policies that are benchmarked against each other observe the same number and sequence of booking requests. Each arriving request is accepted or rejected based on the booking limit policies.

We refer to two simulation setups. In the first setup, which we report the percentile graphs (explained in the following section), in each metric calculation, we take 600 simulation runs and calculate various performance indicators such as percentile levels of the realized revenues. After conducting 30 metric calculations, the mean of the metrics found in these 30 calculations are reported. In the second setup, which is used to report metrics such as the mean of the realized revenues or the sample standard deviation of the realized revenues, 2000 simulation runs are taken to report the mentioned metrics.

Simulation strategies also involve the exploration of different customer arrival patterns. Firstly, as a benchmark, we utilize the “low-before-high” (*LBH*) arrival pattern, where demand for the product with the lowest fare is realized first, followed by the second lowest fare to the highest fare, sequentially. Secondly, we employ a version of random arrival patterns. Initially, we treat customers in the same fare class as indistinguishable objects. Then, we assume that these objects are placed in a box, in numbers that represent the demand realization generated from the random variable based on the demand distribution. When a customer arrives, an object is randomly selected from the box, determining the fare class of the arrived customer. We refer to this approach as the “Uniformly Shuffled Arrival Pattern,” abbreviated as *Uniform Arrival Pattern*. Next, we discuss metrics that experiments yield, together with their importance in understanding the performance of different booking limit controls.

4.3 Performance Metrics

Results of the simulations are explained by various metrics. Even though the robust formulations are proposed to provide better worst-case performance, we use the expected revenue to contrast methods maximizing the expected revenue and worst-case performance.

- Mean revenue and 95% confidence interval (CI), where the true mean is located, calculated using the sample standard deviation of the realized revenues of the simulation runs.
- Distribution of revenues: The distributions offer valuable information about the range of achievable policy performances, from minimum (worst-case) to maximum (best-case). Our approach aligns with previous studies, such as [38], which emphasize the importance of analyzing the distributions of the realized revenues for informed decision-making in revenue management. We rely on histograms and percentile calculations of the realized revenues.
- A version of signal-to-noise ratio (SNR), which is calculated as $\frac{\mu+1}{s+1}$, where μ is average revenue and s is the standard deviation of revenue. While a good indicator of finding high-mean and low-variance methods, good and bad deviations are penalized equally and the trade-off between mean and standard deviation is not apparent. These are among the limitations of SNR [47]. SNR is also related to the inverse of coefficient of variation, a measure of dispersion of a probability distribution.

We also include two methods of finding the skewness and kurtosis of the realized revenues, which can be helpful in determining the dispersion of realized revenues.

- Skewness: A measure of the asymmetry of the probability distribution of a real-valued random variable. We employ two calculations for determining

the skewness. First one is referred as the Fisher-Pearson skewness:

$$\text{Skewness} = \frac{n}{(n-1)(n-2)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s} \right)^3,$$

where n is the sample size, x_i are the sample values, $\bar{x}$ is the sample mean, and s is the sample standard deviation. Second one is a quantile-based definition of skewness and calculated as $\frac{(q_{90}+q_{10})/2-q_{50}}{(q_{90}-q_{10})/2}$, where q_i is the i -th percentile [48].

- Kurtosis: A measure of the “tailedness” of the probability distribution of a real-valued random variable. High kurtosis refers to having the observed values around the mean, and peaked, compared to the standard normal distribution. Again, we use two methods of determining the kurtosis, first one is the following:

$$\text{Kurtosis} = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s} \right)^4 - \frac{3(n-1)^2}{(n-2)(n-3)},$$

where n is the sample size, x_i are the sample values, $\bar{x}$ is the sample mean, and s is the sample standard deviation. Similarly, a quantile-based definition of kurtosis can also be introduced, which is equal to $\frac{q_{90}-q_{10}}{q_{75}-q_{25}}$ [48].

4.4 The Effect of Uncertainty Set Size


In this section, we start by exploring different booking limit policies obtained after manipulating the size parameters for the uncertainty sets, namely, ϵ for (3.20) and (3.25); ρ and Γ for (3.38). Initially, we present the performance of potential booking limit policies under the assumption that the demand for the products are distributed independently from each other. In the subsequent subsection, we introduce correlation between the products and investigate two methods of incorporating such information into our proposed methodologies. Note that, in this section, box and whisker plots refer to mean revenue and 95% CI and line plots refer to the SNR.

4.4.1 Changing Uncertainty Sizes Under Independent Demand Data

The following experiments adhere to the demand structure outlined in Table 3.1 and the arrival sequence is uniform. For each level of uncertainty set size, we conducted 2000 simulation runs after determining the booking limits. Subsequently, we present the mean revenue obtained by each policy at different uncertainty set sizes, along with their respective 95% confidence intervals. Additionally, to comment on the robustness of the solutions, we provide the SNR as well.

Control	b_1	b_2	b_3	b_4
$\epsilon = 0$	95	78	33	0
$\epsilon = 0.25$	95	76	27	0
$\epsilon = 0.5$	95	75	22	0
$\epsilon = 0.75$	95	73	17	0
$\epsilon = 1$	95	72	12	0
$\epsilon = 1.25$	95	70	7	0
$\epsilon = 1.5$	95	69	1	0
$\epsilon = 1.75$	95	68	0	0
$\epsilon = 2$	95	66	0	0
$\epsilon = 2.25$	95	65	0	0
$\epsilon = 2.5$	95	63	0	0
$\epsilon = 2.75$	95	62	0	0
$\epsilon = 3$	95	95	68	68
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels



(b) Mean Revenue, 95% CI of Revenue and SNR at Different ϵ Levels

Figure 4.1: Booking Limit and Revenue Comparison of $Algorithm-NL_\infty$ and EMSRb Policies Under Uniform Arrival Pattern.

In Figure 4.1, we observe that $Algorithm-NL_\infty$ performs less effectively than the widely recognized EMSRb policy in terms of both the variance and mean of realized revenues. This discrepancy can be attributed to the conservative nature of the robust solutions. While $Algorithm-NL_\infty$ prohibits any customers from the lowest fare class from booking at any given ϵ level, this is not the


case with the EMSRb policy. Furthermore, seat allocation for the second lowest fare class decreases as the uncertainty size increases in the proposed algorithm. It's noteworthy that as the number of reserved seats for the lowest fare classes decreases, mean revenues also decrease, indicating that a high number of rejected customers from the low fare classes hinders the performance of *Algorithm-NL_∞*.

Interestingly, at $\epsilon = 3$, *Algorithm-NL_∞* demonstrates strong performance with a 95% confidence interval of $58,695 \pm 276.48$, compared to $58,236 \pm 548.11$ achieved by the EMSRb policy. However, in this instance, the proposed algorithm suggests a booking limit policy that resembles a policy with no booking limit control. As the three products with the lowest fare yield similar revenue per sale, rejecting one but accepting others would not create significant discrepancies, underscoring the importance of determining the booking limit for the product with the second highest fare, protecting the first product. Since at $\epsilon = 3$, EMSRb and the proposed formulation have similar b_2 values, revenue realizations are comparable.

Other than the uniqueness of the control at $\epsilon = 3$, *Algorithm-NL_∞* fails to provide more robust controls than the EMSRb policy. Since the mean is lower and the standard deviation is higher for the realized revenues of the *Algorithm-NL_∞*, this is expected.

Control	b_1	b_2	b_3	b_4
$\epsilon = 0$	95	78	33	0
$\epsilon = 1$	95	76	51	25
$\epsilon = 2$	95	81	54	27
$\epsilon = 3$	95	81	54	27
$\epsilon = 4$	95	81	55	28
$\epsilon = 5$	95	81	55	28
$\epsilon = 6$	95	81	55	28
$\epsilon = 7$	95	81	55	28
$\epsilon = 8$	95	81	55	28
$\epsilon = 9$	95	81	55	28
$\epsilon = 10$	95	81	55	28
$\epsilon = 11$	95	81	55	28
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels




(b) Mean Revenue, 95% CI of Revenue and SNR at Different ϵ Levels

Figure 4.2: Booking Limit and Revenue Comparison of *Algorithm-NL₁* and EMSRb Policies Under Uniform Arrival Pattern.

In the second figure, Figure 4.2, we observe better results achieved using *Algorithm-NL₁* compared to the EMSRb policy. Notably, decision makers are less sensitive to adjustments in the measure of uncertainty when using this algorithm, as it yields the same booking limit controls for $\epsilon \in [4, 11]$. Furthermore, the consistent controls achieved by *Algorithm-NL₁* surpass that of the EMSRb policy in terms of mean revenues and the size of the confidence interval, which is more favorable when small. In return, achieved SNR is better for this algorithm. For instance, in the experiment conducted at $\epsilon = 7$, *Algorithm-NL₁* reports a 95% confidence interval of $58,002 \pm 399.22$, while EMSRb reports $57,267 \pm 605.47$. This represents an improvement of over 1.2% in mean revenue and a remarkable 34.06% decrease in the size of the confidence interval.

Control	b_1	b_2	b_3	b_4
$\Gamma = 0$	95	78	33	0
$\Gamma = 1$	95	76	51	25
$\Gamma = 2$	95	81	54	27
$\Gamma = 3$	95	81	54	27
$\Gamma = 4$	95	72	49	25
$\Gamma = 5$	95	71	49	25
$\Gamma = 6$	95	69	47	24
$\Gamma = 7$	95	67	46	23
$\Gamma = 8$	95	66	0	0
$\Gamma = 9$	95	66	0	0
EMSRb	95	78	44	12

(a) Booking Limits at Different Γ Levels




(b) Mean Revenue, 95% CI of Revenue and SNR at Different Γ Levels

Figure 4.3: Booking Limit and Revenue Comparison of $Algorithm-NL_{B_2}$ and EMSRb Policies Under Uniform Arrival Pattern.

Control	b_1	b_2	b_3	b_4
$\Gamma = 0$	95	78	33	0
$\Gamma = 1$	95	76	51	25
$\Gamma = 2$	95	75	50	24
$\Gamma = 3$	95	73	50	25
$\Gamma = 4$	95	72	12	0
$\Gamma = 5$	95	72	12	0
$\Gamma = 6$	95	72	12	0
$\Gamma = 7$	95	72	12	0
$\Gamma = 8$	95	72	12	0
$\Gamma = 9$	95	72	12	0
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels



(b) Mean Revenue, 95% CI of Revenue and SNR at Different Γ Levels

Figure 4.4: Booking Limit and Revenue Comparison of $Algorithm-NL_{B_1}$ and EMSRb Policies Under Uniform Arrival Pattern.

Lastly, we can address Figures 4.3 and 4.4 together due to their similarity. In both of these experiments, as Γ increases, representing the size of the uncertainty set with respect to the ℓ_1 -norm, the uncertainty set becomes mainly defined by the ℓ_∞ -norm. For smaller Γ levels, the results are comparable to those obtained with *Algorithm-NL₁*. This is because if Γ is respectively smaller than ρ , *Algorithm-NL_{B _{ρ}}* effectively becomes *Algorithm-NL₁* by definition. Furthermore, as Γ increases, this algorithm transitions to *Algorithm-NL _{∞}* . This transition is evident when comparing Figures 4.3 and 4.4. In the former, the uncertainty size is more restricted due to a larger ρ value, and the algorithm yields promising results up to $\Gamma = 7$, similar to the outcomes in Figure 4.2. However, when ρ is decreased, the uncertainty set resembles one spanned only by the ℓ_∞ -norm, and the experiment results align similarly to those in Figure 4.1 for a larger interval of Γ levels. So, following the pattern and the SNR, small integers chosen between 1 and 3 for both of the parameters of this algorithm, seem to yield better results.

As mentioned, the four experiments above assume a uniform arrival pattern. In most cases, changing the arrival pattern to LBH does not significantly affect the relative performance of the proposed algorithms versus the EMSRb policy. However, it is noteworthy to highlight the performance of *Algorithm-NL₁* relative to the EMSRb policy when the arrival pattern is assumed to be LBH.




Figure 4.5: Revenue Comparison of *Algorithm-NL₁* and EMSRb Policies Under LBH Arrival Pattern.

In Figure 4.5, we observe that, compared to Figure 4.2, the solutions obtained from *Algorithm-NL₁* are slightly less favorable in contrast to the EMSRb policy. However, even in this case, *Algorithm-NL₁* reports a 95% confidence interval which is smaller than the one reported by EMSRb. While the size of the confidence interval obtained from EMSRb varies from 1046.72 to 1277.06, for *Algorithm-NL₁* it varies from 656.6 to 828.6. Moreover, during the experiment at $\epsilon = 9$, the experiment with the largest size of confidence interval reported by both methodologies, a 35.12% decrease in the size of the confidence interval is achieved by *Algorithm-NL₁* over the EMSRb policy. As there is a direct correspondence between the size of the confidence interval and sample standard deviation, similar to the experiment under Uniform Arrival Pattern, we observe better SNR for *Algorithm-NL₁*. As it is the only algorithm that consistently yields better SNR, hence more robust than EMSRb, it is more promising than the rest of the robust algorithms.

4.4.2 Changing Uncertainty Sizes Under Dependent Data

Following the previous experiments with independent demand structures, we now investigate the performance of our formulations under dependent demand. As introduced, we refer to two distinct variations of incorporating the correlation between the demands for different products, which are investigated separately in the following subsections.

For this section, we assume a perfectly correlated demand structure, which typically corresponds to a correlation matrix where all the elements are equal to one. However, due to our proposed methods of deriving the correlation parameter matrix P , such a correlation matrix is infeasible as its determinant is zero, and it does not have a square root or a Cholesky Decomposition. Therefore, in the following experiments, we refer to the correlation matrix of perfectly correlated products as a matrix with ones on the diagonal and all other elements equal to 0.99. By using this perturbation, we can approximate the correlation parameter using both proposed methods.

It's important to note that both methods of finding the correlation parameter use the covariance matrix of the products as input, and we assume the correlation matrix to be given. Fortunately, deriving the covariance matrix from the correlation matrix is a trivial task, given the standard deviations of the demand for each product. Assuming a set denoted by $\mathbb{X}$ includes the set of random variables defining the product demands, the covariance matrix is denoted as $\Sigma = Cov(\mathbb{X})$, and the correlation matrix is denoted as $Corr(\mathbb{X})$. The relationship between these two matrices is given below, where σ is the vector of standard deviations of the demand distributions for products.

$$\Sigma = diag(\sigma) \cdot Corr(\mathbb{X}) \cdot diag(\sigma) \quad (4.1)$$

Finally, utilizing the data from Table 3.1 and equation (4.1), we compute and

report the covariance matrix for the experiments as follows:

$$\Sigma = \begin{bmatrix} 33.64 & 86.13 & 75.79 & 64.88 \\ 86.13 & 225.0 & 196.02 & 167.81 \\ 75.79 & 196.02 & 174.24 & 147.67 \\ 64.88 & 167.81 & 147.67 & 127.69 \end{bmatrix} \quad (4.2)$$


Correlation Parameter Derived Using Cholesky Decomposition

To reiterate the Cholesky Decomposition, the goal is to find a lower triangular matrix, P , such that $\Sigma = PP^T$. The Cholesky Decomposition of the covariance matrix in equation (4.2) is computed below. Each component of the matrix is rounded to three decimal places. All the robust algorithms in this subsection use this matrix in their uncertainty sets.

$$P_{Cholesky} = \begin{bmatrix} 5.8 & 0 & 0 & 0 \\ 14.85 & 2.116 & 0 & 0 \\ 13.068 & 0.926 & 1.615 & 0 \\ 11.187 & 0.793 & 0.459 & 1.304 \end{bmatrix} \quad (4.3)$$

Control	b_1	b_2	b_3	b_4
$\epsilon = 0$	95	78	33	0
$\epsilon = 0.25$	95	76	27	0
$\epsilon = 0.5$	95	74	21	0
$\epsilon = 0.75$	95	73	15	0
$\epsilon = 1$	95	71	10	0
$\epsilon = 1.25$	95	70	4	0
$\epsilon = 1.5$	95	68	0	0
$\epsilon = 1.75$	95	67	0	0
$\epsilon = 2$	95	65	0	0
$\epsilon = 2.25$	95	64	0	0
$\epsilon = 2.5$	95	62	0	0
$\epsilon = 2.75$	95	0	0	0
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels




(b) Mean Revenue, 95% CI of Revenue and SNR at Different ϵ Levels

Figure 4.6: Booking Limit and Revenue Comparison of $Algorithm-NL_\infty$ and EMSRb Policies Under Uniform Arrival and Cholesky Decomposition.

In Figure 4.6, similar to the independent case, we conclude that while $Algorithm-NL_\infty$ provides extremely conservative solutions by reserving a significant amount of capacity for the higher fares, these solutions are not consistent with changing uncertainty set sizes. For smaller ϵ values, the algorithm is comparable to the EMSRb policy in terms of mean revenues but not in terms of SNR. However, small changes in the size parameter deteriorate the simulation performance and have a significant effect on the booking limit policy. Which is not favorable, as it increases the need for sensitivity analysis before using the policy, increasing the computational costs.

Control	b_1	b_2	b_3	b_4
$\epsilon = 0$	95	78	33	0
$\epsilon = 0.25$	95	76	27	0
$\epsilon = 0.5$	95	75	22	0
$\epsilon = 0.75$	95	73	17	0
$\epsilon = 1$	95	72	12	0
$\epsilon = 1.25$	95	71	7	0
$\epsilon = 1.5$	95	69	2	0
$\epsilon = 1.75$	95	68	0	0
$\epsilon = 2$	95	66	0	0
$\epsilon = 2.25$	95	65	0	0
$\epsilon = 2.5$	95	63	0	0
$\epsilon = 2.75$	95	62	0	0
$\epsilon = 3$	95	95	5	5
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels




(b) Mean Revenue, 95% CI of Revenue and SNR at Different ϵ Levels

Figure 4.7: Booking Limit and Revenue Comparison of *Algorithm-NL₁* and EMSRb Policies Under Uniform Arrival and Cholesky Decomposition

Figure 4.7 reports similar results to the previous table-plot figure. In the independent case, *Algorithm-NL₁* was quite favorable due to its insensitivity to changing ϵ levels and comparable, and moreover, better policy performance compared to EMSRb. However, when correlation is incorporated, these encouraging properties diminish. The conservativeness in saving allocation for the higher fare classes still poses problems similar to the preceding experiment.

Control	b_1	b_2	b_3	b_4
$\Gamma = 0$	95	78	33	0
$\Gamma = 0.5$	95	75	22	0
$\Gamma = 1$	95	72	12	0
$\Gamma = 1.5$	95	69	2	0
$\Gamma = 2$	95	66	0	0
$\Gamma = 2.5$	95	66	0	0
$\Gamma = 3$	95	66	48	0
$\Gamma = 3.5$	95	66	65	41
$\Gamma = 4$	95	66	65	40
$\Gamma = 4.5$	95	66	65	41
$\Gamma = 5$	95	65	64	40
$\Gamma = 5.5$	95	65	0	0
$\Gamma = 6$	95	65	0	0
$\Gamma = 6.5$	95	65	0	0
EMSRb	95	78	44	12

(a) Booking Limits at Different Γ Levels



(b) Mean Revenue, 95% CI of Revenue and SNR at Different Γ Levels

Figure 4.8: Booking Limit and Revenue Comparison of *Algorithm-NL_{B₂}* and EMSRb Policies Under Uniform Arrival and Cholesky Decomposition

Interestingly, while the previous algorithms failed to match their performances in the independent cases, *Algorithm-NL_{B₂}* remains intriguing in the experiment with dependent demand structure. Previously, *Algorithm-NL_{B_r}* was mentioned to mimic either *Algorithm-NL_∞* or *Algorithm-NL₁*, however, in this experiment, one can observe that at $\Gamma \in [3.5, 5]$, *Algorithm-NL_{B₂}* reports unique booking limit controls. Moreover, these unique controls are comparable to the EMSRb policy in contrast to controls from the previous algorithms, and in terms of mean revenues, they are much closer to the benchmark. Additionally, they report SNRs that are slightly better than the EMSRb policy. However, finding this subset of policies that are favorable could be different for different product demand information and may not even be present, imposing a need for sensitivity analysis in practice.

Correlation Parameter Derived Using The Square Root of The Covariance Matrix

In this subsection, as previously mentioned, we use the square root of the covariance matrix, Σ , as the correlation parameter. Which is the matrix P such that $\Sigma = PP$. Such square root can be calculated by diagonalization using the eigenvalues and the corresponding eigenvectors of the covariance matrix. One can choose matrix V as a matrix with the eigenvectors of Σ as columns and assume a diagonal matrix D with diagonal elements equal to the eigenvalues. Then, by finding the square root of each element of D , we find a matrix S having the relationship $D = SS$. Moreover, the square root of Σ is equal to the following matrix multiplication, which can be used to find P .


$$\Sigma = PP = (VSV^{-1})(VSV^{-1}) \quad (4.4)$$

Utilizing the aforementioned diagonalization method to compute the square root of the covariance matrix, we report the correlation matrix where entries rounded to three decimal places that will be utilized in this subsection as the following.

$$P_{SquareRoot} \begin{bmatrix} 2.021 & 3.524 & 3.131 & 2.709 \\ 3.524 & 10.312 & 7.807 & 6.730 \\ 3.131 & 7.807 & 8.239 & 5.968 \\ 2.709 & 6.730 & 5.968 & 6.280 \end{bmatrix} \quad (4.5)$$

Control	b_1	b_2	b_3	b_4
$\epsilon = 0$	95	78	33	0
$\epsilon = 0.25$	95	75	23	0
$\epsilon = 0.5$	95	72	13	0
$\epsilon = 0.75$	95	69	3	0
$\epsilon = 1$	95	67	0	0
$\epsilon = 1.25$	95	64	0	0
$\epsilon = 1.5$	95	95	7	0
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels




(b) Mean Revenue, 95% CI of Revenue and SNR at Different ϵ Levels

Figure 4.9: Booking Limit and Revenue Comparison of $Algorithm-NL_\infty$ and EMSRb Policies Under Uniform Arrival and Square Root Method.

Our conclusion on this experiment is again similar to the previous discussion regarding $Algorithm-NL_\infty$. The performance of this algorithm is still not encouraging for further investigation, and moreover, consistency concerning the size of the uncertainty set and booking limit policies requires the need for sensitivity analysis and the fact that it fails to provide better SNR yields, do not possess a strong argument for the use of this algorithm under the dependent demand formulation.

Control	b_1	b_2	b_3	b_4
$\epsilon = 0$	95	78	33	0
$\epsilon = 0.5$	95	76	26	0
$\epsilon = 1$	95	74	48	0
$\epsilon = 1.5$	95	73	73	21
$\epsilon = 2$	95	71	71	15
$\epsilon = 2.5$	95	69	69	10
$\epsilon = 3$	95	67	67	4
$\epsilon = 3.5$	95	66	66	1
$\epsilon = 4$	95	64	64	0
$\epsilon = 4.5$	95	62	62	0
$\epsilon = 5$	95	95	95	32
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels




(b) Mean Revenue, 95% CI of Revenue and SNR at Different ϵ Levels

Figure 4.10: Booking Limit and Revenue Comparison of *Algorithm-NL₁* and EMSRb Policies Under Uniform Arrival and Square Root Method.

Using *Algorithm-NL₁* by deriving the correlation parameter from the Cholesky Decomposition was not particularly interesting; however, Figure 4.10 is promising. While the consistency of the policies under different ϵ values is an issue, there are well-performing booking limit controls found using this algorithm with the square root of the covariance matrix. It's noteworthy that the booking limit control from the EMSRb policy is (95, 78, 44, 12), and comparing this to *Algorithm-NL₁* at $\epsilon = 1.5$, (95, 73, 73, 21), we almost achieve the same performance with two distinct controls. Moreover, again unlike the experiment with Cholesky Decomposition, we were able to find booking controls that yield higher or the same SNR, which is favorable for robust formulations.

Control	b_1	b_2	b_3	b_4
$\Gamma = 0$	95	78	33	0
$\Gamma = 0.5$	95	76	26	0
$\Gamma = 1$	95	74	48	0
$\Gamma = 1.5$	95	73	73	21
$\Gamma = 2$	95	71	71	15
$\Gamma = 2.5$	95	69	69	51
$\Gamma = 3$	95	68	68	54
$\Gamma = 3.5$	95	66	57	57
$\Gamma = 4$	95	65	60	60
$\Gamma = 4.5$	95	63	63	63
$\Gamma = 5$	95	62	62	62
EMSRb	95	78	44	12

(a) Booking Limits at Different ϵ Levels



(b) Mean Revenue, 95% CI of Revenue and SNR at Different Γ Levels

Figure 4.11: Booking Limit and Revenue Comparison of *Algorithm-NL_{B₂}* and EMSRb Policies Under Uniform Arrival and Square Root Method.

Figure 4.11, presents a similar findings as the previous figure. Firstly, while the consistency is still an issue, at mid-levels for possible Γ values, well-performing policies can be found in terms of SNR and mean revenues. Note that even though the mean revenue performance deteriorates as Γ increases, SNR remains quite similar if not higher. However, note that, most booking limits in this figure resembles a policy which do not enforce any booking limits. So, in certain demand factor scenarios, the algorithms might be less and less preferable, i.e., when the demand is much higher than the available capacity and lower fare classes book first. Also, the booking limit policy at $\Gamma = 1.5$ is the same as *Algorithm-NL₁* control at $\epsilon = 1.5$, which was noted in the analysis of Figure 4.10.

4.4.3 Choosing The Size of the Set

As shown, there are multiple booking policies depending on the selected parameters. In this sub-discussion we provide more details on selecting better performing

uncertainty set size parameter for *Algorithm-NL₁* under independent demand using more metrics that might provide insights into the robustness of different policies. In the next table, we compare possible controls under changing uncertainty size for the mentioned algorithm under three new metrics: minimum revenue achieved, maximum regret realized from the perfect hindsight model (PH), and average percent deviation from the PH, where PH is a model where the random demand is assumed to be known in advance.

Control	b_1	b_2	b_3	b_4	Minimum Revenue	Maximum Regret	% Deviation From PH
$\epsilon = 0.5$	119	101	67	34	55394	18720	3.54
$\epsilon = 1$	119	100	67	33	54860	19240	3.99
$\epsilon = 1.5$	119	99	66	33	54293	19240	4.48
$\epsilon \in [2, 3]$	119	101	67	34	55394	18720	3.54
$\epsilon \in [3.5, 6.5]$	119	102	69	35	55961	18200	3.02
$\epsilon \in [7, 11.5]$	119	102	69	35	55961	18200	3.02
EMSRa	119	102	80	63	55961	10753	2.33
EMSRb	119	102	68	36	55961	17680	3.00

Table 4.1: Changing Uncertainty in *Algorithm-NL₁*, LBH Arrival Pattern for Table 3.1.

It's worth noting that, while the booking limits of *Algorithm-NL₁* may appear identical for the controls derived from our formulation at $\epsilon \in [3.5, 6.5]$ and $\epsilon \in [7, 11.5]$ intervals, we have opted to distinguish them due to slight disparities in solution structures when rounding is not applied. To provide clarity, within the former interval, the booking limits $(b_1^N, b_2^N, b_3^N, b_4^N)$ are approximated as $(119, 101.57, 69.04, 34.57)$.

The largest possible ϵ value that yield a non-trivial solution in independent demand scenario is better, compared to the booking limit policies achieved using smaller ϵ values for the same algorithm. So, in the context of *Algorithm-NL₁*, we recommend selecting ϵ as the largest permissible value within its range. This is substantiated through the empirical study in Table 4.1, where we have observed that among the various controls generated using different ϵ values, as ϵ increases,

performance of the algorithm consistently yields more superior results. Moreover, in Table 4.1 in out of all three metrics given, highest possible ϵ level strictly dominate other levels.

Consequently, for the subsequent experiments under independent demand distributions, we have made the decision to fix ϵ at the value of $\sum_{j \in N} \frac{\bar{d}_j}{P_{jj}} - 0.001$ for *Algorithm-NL₁*. The inclusion of the negative term is intended to ensure that the objective value remains distinct from a vector of all zeros. Also, for other algorithm variants, set sizes are mostly taken as 1 or 2, as they yield better controls other possible set size selections, which can be justified in the previous examples of this section without the detailed discussion as the case with *Algorithm-NL₁*.

In terms of applying the robust algorithms under dependent demand, however, there is not a clear better performing controls based on the uncertainty set sizes as we have with the independent demand distributions. Under this setting, for most versions of the proposed algorithms, choosing the uncertainty set size small seem to perform better in terms of mean revenue, standard deviation and the SNR metrics, but still, the numerical studies do not suggest a dominant set size.

4.5 Experiments to Benchmark the Robust Methods to Others

In this section, we investigate the booking limit controls of the policies using DLP and EMSRb, and benchmark the robust formulations in more detail. The first three sections of this chapter outline the experimental setup, performance metrics, and the policies used. We first investigate two-product scenarios (where EMSRb equals to Littlewood's Rule), then multi-product scenarios. Both of subsections start by investigating the effect of various parameters such as demand factor, standard deviation, and fare ratio on the booking limit controls as well as on the resulting performance. We then provide a few general realistic scenarios.

4.5.1 Two-Product Performance Comparison

We aim to elucidate the performance differences offered by the existing methods in the literature and our proposed formulations under a setting where only two products are offered. Throughout the results presented in this section, the average of 30 experiments is reported. In each of these experiments, we follow the first simulation setup discussed in Section 4.2, where 30 individual experiments with 600 simulation runs are conducted for each figure.

The investigation of the two-product case separately is due to the easier comparison of the proposed algorithms against the optimal solution. The initial experiments assume Uniformly Shuffled Arrival Pattern. Later, some of the experiments are repeated under the LBH pattern. It's important to note that in the two-product setting with the LBH pattern, EMSRb control is simply the optimal solution to the expectation maximization problem, following Littlewood's rule. This allows us to benchmark the robust formulation against the optimal control.

For the rest of this subsection, the uncertainty set size for *Algorithm-NL₁* is the largest possible value that ϵ can reach without forcing all variables to be zero, namely, $\epsilon = \frac{\bar{d}_1}{P_{11}} + \frac{\bar{d}_2}{P_{22}} - 0.01$. For algorithms *Algorithm-NL_B* and *Algorithm-NL_∞*, the Γ and ϵ parameters are selected relatively small, usually chosen between 1 or 2. This selection process is drawn from the experiments conducted based on changing uncertainty set sizes at the beginning of this chapter.

Fare Ratio and Demand Factor Effect Analysis

In this part of the subsection, we investigate the effect of the amount of difference between the fares of two products on the performances of the proposed algorithms compared to EMSRb with $n = 2$. First, we fix the high-fare at 1000 and change the low-fare from 100 to 1000 while keeping the demand structure for the two products at a mean of 60 and a standard deviation of 30, where the capacity of the resource is 100.

Before delving into the performance comparison, we first report the changing booking limit policies under this experimental setting. In the following table, we present the booking limit for the low-fare product for a given fare ratio.

Fare Ratio	<i>Algorithms</i>			Littlewood's Rule	DLP
	$NL_1^{3.99}$	NL_∞^1	$NL_{B_1^2}$		
0.1	91	10	10	2	40
0.2	83	10	10	15	40
0.3	77	10	10	24	40
0.4	71	10	10	32	40
0.5	67	10	10	40	40
0.6	62	10	10	48	40
0.7	59	10	10	56	40
0.8	56	10	10	65	40
0.9	53	10	10	78	40
1	50	10	10	133	40

Table 4.2: Booking Limits and Fare Ratios, Products with a Mean of 60, Standard Deviation of 30.

Table 4.2 highlights one of the shortcomings of the proposed algorithms. While in some scenarios we have observed that the proposed algorithms perform well, there have been cases where they lacked performance compared to standard methods. This table may shed light on the reasoning behind these changing performance relations.

In a setting with two products, when the fares of the products are closer, these products should effectively become the same entity. This is evident in Littlewood's Rule, as this optimal policy starts by limiting the low-fare products when the fare ratio is small, and gradually ceases to limit the low-fare product as the fares of the products approach parity. This concept may seem intuitively trivial, however, the proposed algorithms do not consistently adhere to this logic.

While two of the algorithms does not change its behavior, *Algorithm-NL₁* acts in the opposite direction. The performance of these booking limits under changing

fare ratios are visualized in a mean revenue experiment, provided next.




Figure 4.12: Relative Performance of Different Policies Against The Littlewood's Rule Under Changing Fare Ratio Levels, LBH Arrival Pattern.

In the experiment shown in Figure 4.12, 2000 simulation runs are conducted for each fare ratio level, and the mean revenue is calculated. Then, the relative performance of different methods against Littlewood's Rule is reported. As expected, the counter-intuitive booking limits of *Algorithm-NL₁* only begin to approach the optimal policy when the fare ratio is high, where booking limit policies start to have less and less importance, as shown by the optimal policy at fare ratio 1 which is more than the capacity. Note that, theoretically, at this level, protection level from the Littlewood's rule approach minus infinity.

Also, algorithms NL_∞ and NL_B deteriorate in performance as the fare ratio approaches to 1, as they start failing to utilize the capacity. For instance, at fare ratio 1, Littlewood's accept any request but these algorithms only accept 10 request from one of the classes, thus capacity of 100 is not fulfilled with the given demand information.

Next, we explore how the demand factor impacts the performance of these algorithms. In this case, the fares of the products are fixed at 1000 and 500, with mean and standard deviation values of 60 and 30 for each product, respectively. The capacity of the resource varies from 150 to 60, allowing us to investigate booking limits for demand factors ranging from 0.8 to 2.

Demand Factor	<i>Algorithms</i>			Littlewood's Rule	DLP
	NL_{∞}^1	NL_1	NL_{B_1}		
0.80	100	60	60	90	75
0.86	93	50	50	80	70
0.92	87	40	40	70	65
1.00	80	30	30	60	60
1.09	73	20	20	50	50
1.20	67	10	10	40	40
1.33	60	0	0	30	30
1.50	53	0	0	20	20
1.71	47	0	0	10	10
2.00	40	0	0	0	0

Table 4.3: Booking Limits and Demand Factor, Products with a Mean of 60, Standard Deviation of 30.

In Table 4.3, the algorithms follow an intuitive reasoning similar to Littlewood's Rule, which was not the case with changing fare ratios. The optimal policy determines a protection level for the high-fare product, and the booking limit is calculated by subtracting this value from the capacity. So, the protection level is the same for every demand factor level in the optimal policy, and as the capacity decreases, so does the amount of capacity left for allocating the low-fare product. All variations of the proposed algorithms tend to decrease the booking limit of the low-fare product as there is less total capacity available, which is intuitive.




Figure 4.13: Relative Performance of Different Policies Against The Littlewood's Rule Under Changing Demand Factor Levels, LBH Arrival Pattern.

Surprisingly, Figure 4.13 shows deteriorating performance of *Algorithm-NL₁* as the demand factor increases. In the previous paragraph, we explained that the changes in the booking limit of the low-fare product are expected and intuitively explainable. However, while the booking limits under changing fare ratios were counter-intuitive, as shown in Figure 4.12, the proposed algorithms performed better compared to Figure 4.13.

While the limits change following a similar reasoning to the optimal policy, the proposed algorithms either fail to react to these changes or react too quickly. Specifically, *Algorithm-NL₁* gives a booking limit of 100 at demand factor level 0.8, while the optimal policy suggests 90, corresponding to a 10 unit difference in a 150 unit capacity. However, at demand factor level 2, the former method limits the low-fare at 40 while the latter does not even allow admitting any low-fare products, corresponding to a 40 unit difference in a 60 unit capacity. Resulting in more deviation from the booking limit of the optimal policy as the demand factor increases.

Standard Variation Effect Analysis

In this subsection, we investigate the effects of standard deviation on the booking limit policies. In the experimental setting, we assume that the mean demand of each product is 60 and is a normal distribution while the capacity is fixed at 100. Different from the previous experiments, for each product, we vary the standard deviation between 5, 10, 15, and 20.

We also consider two sets of fares, 1000-500 and 1000-700. In the former setting, DLP and the Littlewood's Rule always protect the same number of seats for the higher-fare class. Moreover, under this fare ratio of 0.5, for any combination of (σ_1, σ_2) , the protection from the rule is always equal to the mean of the higher-fare class product. Hence, we provide our analysis on this fare ratio in Appendix B and discuss the more dynamic setting of fare ratio of 0.7 in this part.

Below, we provide the booking limits under different combinations of (σ_1, σ_2) . Once again, the parameters for the sizes of the uncertainty sets are kept the same. To reiterate, for *Algorithm-NL_∞*, $\epsilon = 1$; for *Algorithm-NL₁*, ϵ is the highest possible value; and for *Algorithm-NL_B*, $\rho = 1$ and Γ is equal to the ϵ of *Algorithm-NL₁*.

(σ_1, σ_2)	<i>Algorithms</i>			Littlewood's Rule	DLP
	NL_∞^1	NL_1	NL_{B_1}		
(5, 5)	35	59	35	43	40
(10, 5)	32	74	32	45	40
(15, 5)	28	81	28	48	40
(20, 5)	70	85	70	50	40
(5, 10)	33	42	33	43	40
(10, 10)	30	59	30	45	40
(15, 10)	26	68	26	48	40
(20, 10)	73	74	73	50	40
(5, 15)	32	32	32	43	40
(10, 15)	29	49	29	45	40
(15, 15)	25	59	25	48	40
(20, 15)	21	66	21	50	40
(5, 20)	30	26	30	43	40
(10, 20)	27	42	27	45	40
(15, 20)	24	52	24	48	40
(20, 20)	20	59	20	50	40

Table 4.4: Booking Limit Controls Under Changing Standard Variations When the Fare Ratio Is 0.7.

It's noteworthy that both *Algorithm-NL_∞* and *Algorithm-NL_B* yield the same booking limits. However, while these two algorithms decrease the booking limit of the lower-fare product as the variation of the first product increases in most cases, *Algorithm-NL₁* behaves conversely by consistently increasing the booking limit.

The robust models are expected to protect against possible adverse scenarios that existing methods tend to overlook due to their low probability of occurrence. In the case of increasing variation of the higher-fare product, the probability of realizing lower demand for this class increases. Hence, it is expected from the robust solutions to allow more bookings from the lower fare classes, thus protecting revenue from such worst-case scenarios.

Conversely, all three algorithms tend to decrease the booking limit when the

variation of the second product increases. In such a scenario, maintaining the booking limit policy unchanged would be more intuitive. In terms of robustness, a decrease in the booking limit due to increased variation of the lower-fare product implies that the model may be unable to protect the decision maker in the event of an extremely low demand realization for this class. However, as the probability of higher demand realizations for the lower-fare class increases due to heavier tail probabilities, the booking limit could be more readily exhausted by this fare, resulting in fewer capacities being allocated to the higher-fare class. Thus, decreasing the booking limit of the lower-fare product when the standard deviation of this product increases still adheres to the rationale behind robust models.

Now, we can observe that this is the case with the Littlewood's Rule as well. However, while the robust models dramatically change this decision, Littlewood's is more subtle. For example, from $(\sigma_1, \sigma_2) = (5, 5)$ to $(\sigma_1, \sigma_2) = (20, 5)$, the booking limits from *Algorithm-NL_∞* change from 35 to 70, while the booking limits found using Littlewood's only change from 43 to 50.

Another observation is that if σ_1 is not changing, an increase in σ_2 is disregarded by the Littlewood's Rule. While we explained that a changing booking limit against such a change might be unimportant, under this fare ratio, we can experimentally observe both strategies.




Figure 4.14: Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under LBH Arrival Pattern.

In Figure 4.14, it becomes more evident that, excluding a few comparisons, the increase in the standard deviation of both products enhances the performance of robust models relative to the Littlewood's Rule. Particularly, *Algorithm- NL_1* starts to outperform the mean revenue achieved by this policy in some high-variation scenarios.

When the variation of the demand for the lower-fare product is high and the variation of the higher-fare product is mid-to-high, *Algorithm- NL_1* tends to be

more preferable. However, low variation in the lower-fare product but high variation in the higher-fare product seems to deteriorate its performance. This can be explained by the very low booking limit of 26, which is 17 units lower than the optimal policy. As the mean of the higher-fare product is 60 and with a standard deviation of 5, leaving 74 capacity would result in a lot of unused capacity, highlighting one downside of *Algorithm-NL₁*.




Figure 4.15: Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under Uniform Arrival Pattern.


Once more, in the uniform arrival pattern setting, robust models perform better compared to the existing methods. Except for two scenarios, which were

the most problematic cases for them in the previous experiment, *Algorithm-NL₁* either performed very close to Littlewood’s Rule or outperformed it. While increasing σ_2 does not seem to have an effect on the performances except at $\sigma_1 = 5$ for *Algorithm-NL₁*, as σ_1 increases, *Algorithm-NL₁* becomes much better compared to Littlewood’s, demonstrating efficiency in many scenarios, especially in scenarios where the demand variation is high.


Similar to the percentile experiments conducted before to understand the performances of the methods under good or bad scenarios, we have prepared bar plot figures for the experimental setting with the fare ratio of 0.7 that report the 25th percentile of the realized revenues to highlight the worst-case performance under the effect of standard deviation, which can be found in Appendix C. One observation from these tables was that in some cases, the 25th percentile of the realized revenues from the *Algorithm-NL₁* controls was higher than its mean, while this was observed in much fewer cases with Littlewood’s Rule. Based on the large negative difference between the 25th percentile and the mean of the revenues from *Algorithm-NL₁*, as well as the large positive difference from Littlewood’s Rule, we decided to further investigate this phenomenon, in the next part.

Analysis on The Distribution of The Realized Revenues

In this part, we provide a histogram for realized revenues derived from one of the experiments of the previous part. The expected demand for each product is again 60, standard deviations are fixed at $(\sigma_1, \sigma_2) = (20, 20)$.



(a) Histogram for *Algorithm-NL₁*



(b) Histogram for EMSRb

Figure 4.16: Realized Revenue Histograms at $(\sigma_1, \sigma_2) = (20, 20)$ When The Fare Ration Is 0.7, Under LBH Arrival Pattern.

In Figure 4.16, revenue realizations of the 2000 simulation runs at $(\sigma_1, \sigma_2) = (20, 20)$ are shown based on their frequencies. We have commented that high

standard deviations for both products are scenarios that are more preferable for the robust models. In this figure, the first thing to note is that the tail probabilities of revenue realizations under 80,000 are much heavier for the EMSRb policy, with the first histogram showing a noticeably larger negative skewness. The most frequent realization bin is higher for EMSRb; however, in good-case scenarios, *Algorithm-NL₁* seems to distribute the realized revenue to a few bins above the 80,000 revenue level. By having a lighter left-hand tail, we conclude that instead of bad revenue realizations, *Algorithm-NL₁* distributes these occurrences to average-to-good case scenarios, ultimately protecting the decision maker.

Illustrative Examples

After individually investigating the effect of various parameters on the booking limit controls, we now begin with few generic examples. First, products generate 1000 and 500 units of revenue. The demands of the products have mean values of 30 and 60, with standard deviations of 10 and 5, respectively. Additionally, the capacity of the resource is assumed to be 80, corresponding to a demand factor of 1.125. This experiment setting will be referred to as the *First Case*.




Figure 4.17: Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case.

In Figure 4.17, the policies obtained using $Algorithm-NL_1$ and $Algorithm-NL_\infty$ are not the same but their percentile performances are the same. Essentially, among the realized revenues, when the demand realization is unfavorable, i.e., achieved revenues are lower than the overall experiments, robust solutions protect the decision maker. One can interpret the 50th percentile as the median of realized revenues, where our solutions fell 1.56% behind the EMSRb policy; however, for more unfavorable realizations, for instance, at the 10th percentile level, they are 10.73% better than this policy. This highlights their importance as they do not deteriorate the “average performance” but protect against “worst-case scenarios.”

Based on the previous sections, proposed solutions show greater promise at a demand factor of 1.6. Thus, to experiment under different demand factors, the mean demand of the low-fare product is changed. To achieve the desired demand factor level, the mean demand of the low-fare product is fixed at 98, while all other parameters are kept the same.




Figure 4.18: Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case Variant With DF of 1.6.

Figure 4.18 is not as promising as the previous one. While the proposed methodologies protect from worst-case scenarios, their performance relative to the average scenario and slightly better worst-case scenarios is less favorable. An improvement of 9.73% is realized for the 10th percentile, but a deterioration of 8.84% is observed for the average scenario compared to the EMSRb policy. Still, range of realized revenues is still much lower for $Algorithm-NL_1^{22.59}$. 80% of the realized revenues fall in $[46, 331, 53, 750]$, while the same percentage revenues realized using the EMSRb policy fall in $[42, 223, 55, 000]$.

Now, further updates to the demand are investigated, mainly the standard deviation for the low-fare product is now taken as 20, while other parameters are similar to the second experiment, where the demand factor is 1.6.




Figure 4.19: Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case Variant With DF of 1.6, Low-fare Deviation of 20.

While Figure 4.19 is quite similar to the previous ones of this chapter, realized revenues for the first two proposed algorithms are starting to differ from each other. For this experiment, *Algorithm-NL₁* is stochastically dominates *Algorithm-NL_∞*, by observation. Moreover, while the performance deterioration for the average scenario remains at similar levels, specifically 8.45% for Figure 4.19, the worst-case scenario performance of *Algorithm-NL₁* is 12.73% better than the EMSRb policy, at the 10th percentile level.

Now, we changed the mean demand for the low-fare product to be 50 to achieve a demand factor of 1. Also, the standard deviation for this product is decided to be 5, while other parameters are similar to the previous experiment.




Figure 4.20: Percentile vs. Revenue. Uniformly Shuffled Arrival Pattern Under The First Case Variant With DF of 1, Low-fare Deviation of 5.

The differences in performances of policies are smaller in Figure 4.20 compared to the previous experiments; however, it is quite important for the thesis. In this experiment, $Algorithm-NL_1$ and $Algorithm-NL_\infty$ stochastically dominate the EMSRb policy, achieving better or equal performance at all percentile levels. While this does not hold in general, the existence of such experimental settings is intriguing for the sake of these new methodologies.

In the following experiments, we investigate the performances of the mentioned controls under LBH arrival pattern, where the EMSRb is optimal. In the first of such experiments, we set the parameters similar to the first case of random distribution pattern. The capacity of the resource is again fixed at 80. As our and the benchmark methods do not yield different booking limits based on customer arrival patterns, resulting booking limit controls are similar to the ones found earlier.




Figure 4.21: Percentile vs. Revenue. LBH Arrival Pattern Under The First Case.

In Figure 4.21, the performance of our formulation in LBH arrival seems to be slightly deteriorated from the uniform arrival experiment of Figure 4.17. This might be due to the fact that the booking limit of algorithms $Algorithm-NL_\infty$ and $Algorithm-NL_1$ for the lower fare class is higher than the EMSRb heuristic. When the low-fare demand is realized first, much of the capacity is exhausted by this fare, and robust solutions do not protect as much capacity for the high-fare as the EMSRb policy. Still, in the worst-case scenario realizations, robust solutions are more favorable compared to the optimal policy; showing smaller range of realizable revenues.

Previously, the last experiment on the variants of the *First Case* assumed a high variation for the low-fare product. All of the parameters are the same as the original setting, except the standard deviation of the second product is now 20.




Figure 4.22: Percentile vs. Revenue. LBH Arrival Pattern Under The First Case Variant with Low-fare Deviation of 20.

Increasing the standard deviation of the low-fare class in the LBH arrival case, Figure 4.22, reported a similar result to the uniform arrival case of Figure 4.20. Surprisingly, the former assumed a relatively high standard deviation while the latter assumed low. This may be attributed to the notion of robust solutions. As the demand structure becomes somewhat extreme, and maybe unrealistic, robust solutions protect the high-fare better than the EMSRb policy. While the EMSRb policy calculates the protection level for the high-fare class, the demand distribution of the low-fare class is not as important as in the proposed solutions, hence changing the distribution characteristics of the low-fare class is better adapted in robust solutions.

For completeness, while the experiments are quite similar, different demand structures are investigated. Now, we assume the same product prices, however, mean demand value is fixed at 45 and the standard deviation is fixed at 15 for the both products, keeping the capacity of the resource at 80. This setting will be referred to as *45-15 Demand*.




Figure 4.23: Percentile vs. Revenue. LBH Arrival Pattern Under 45-15 Demand.

Similar discussions as the above figures could be made for Figure 4.23. While the proposed solutions are not as favorable as the EMSRb policy for the average cases, they protect the realizable revenue for the worst-case scenarios.

To follow the previous works on robust seat-inventory problems, we also considered demand distributions that are known to be uniform. Experiments under this assumption are discussed in Appendix E.

4.5.1.1 Dependent Demand for Two-product Settings

In this subsection, we will discuss the performance of our methodologies under correlated demand structures using generic examples. Initially, we created a virtual scenario where the optimal policy could be determined, and comment on performances with respect to the optimal controls. Using two different correlation parameter calculation methods, we provide performance comparisons under different scenarios.

In the first experiments of the series below, we start with a perfectly uncorrelated demand structure, where the revenues of products are 100 and 50. The demand for each product has a mean of 5 and a standard deviation of $\sqrt{2}$. Below, we calculate the covariance matrix using the standard deviation and correlation information, which will be used to determine the correlation parameter for the robust formulations.

$$\Sigma = \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{bmatrix} \cdot \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \cdot \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{bmatrix} = \begin{bmatrix} 2 & -2 \\ -2 & 2 \end{bmatrix} \quad (4.6)$$

Due to its simplicity, every scenario can be mapped for this example. As there is only 5 units of capacity and two products, the booking limit decision for the lower-fare product will be an integer between 0 and 5. All 6 of the possible controls are investigated in Appendix D. For all possible demand realizations, under the LBH assumption, realized revenue is calculated, and the expected revenue as well as the range of revenues are reported in this appendix content.

It's noteworthy that for this example, the two benchmarks, EMSRb and DLP, find the same booking limit. While these methodologies do not incorporate correlation information, both find the booking limit of the lower-fare product as 2. However, the booking limits of 3 and 4 correspond to the optimal expected revenue maximization solution, with booking limit 4 having the best worst-case performance.

Correlation Parameter Derived Using The Square Root of The Covariance Matrix

Using the square root of the covariance matrix to compute the correlation parameter P , as previously discussed, we perturb the assumed correlation matrix. Products are assumed to be perfectly negatively correlated, meaning the non-diagonal elements of the correlation matrix are set to -1. However, such a matrix is not positive definite, and finding the square root is not feasible. Therefore, in our calculation, the non-diagonal elements of the correlation matrix are assumed

to be -0.99. In the first experiment, we investigate the example with a capacity of five, comparing possible booking limit controls. Below, we report the square root of the matrix (4.6).

$$P_{SquareRoot} = \begin{bmatrix} 1.06820755 & -0.92678619 \\ -0.92678619 & 1.06820755 \end{bmatrix} \quad (4.7)$$




Figure 4.24: Percentile vs. Revenue. Five-capacity, Full Negative Correlation Example Using The Square Root of The Covariance Matrix.

In Figure 4.24, we observe three of the six possible controls. While they seem to perform similarly in good cases, the EMSRb policy is inferior to the optimal solutions during worst cases. Out of 600 simulation runs, EMSRb realized a revenue of 200 in 122 runs, whereas in these worst-case scenarios, the algorithms realized 250 in all 122 experiments, and the alternative optimal control $b_2 = 4$ realized 300. Although EMSRb has a better median value, it performed, based on the mean, 3.42% less than the optimal controls.

Given the total capacity is limited to 5, rounding creates large changes in the

resulting controls. Therefore, we now assume a capacity of 80, with both products distributed with a standard distribution with a mean of 60 and a standard deviation of 20. The higher-fare product is priced at 1000 units, while the lower-fare is at 500 units. Additionally, we assume that the products are perfectly negatively correlated.




Figure 4.25: Percentile vs. Revenue. Eighty-capacity, Full Negative Correlation Example Using The Square Root of The Covariance Matrix.

Firstly, Figure 4.25 reveals a notable difference between the EMSRb policy and the algorithms. While EMSRb spans a wider range of realized revenues, from 44,480 to 70,000, our algorithms' realizations fall within a narrower range, from 52,980 to 62,991. At the lowest percentile level, the robust models outperform the EMSRb policy by 19.1% in terms of revenues.

In an attempt to distinguish between DLP and EMSRb, we set the fares of the products to 1000 and 700 units, while keeping all other parameters constant from the previous setting. However, we now assume that perfect negative correlation does not exist; instead, non-diagonal entries of the correlation matrix are assumed to be -0.5.




Figure 4.26: Percentile vs. Revenue. Eighty-capacity, 0.7 Fare Ratio Example Using The Square Root of The Covariance Matrix.

At this fare ratio, as expected, DLP and EMSRb differentiate. While both promise better revenues for good-case scenarios, they still fail to adequately protect the decision maker from worst-case scenarios. Moreover, the wide range of realized revenues remains a major issue for these existing methods.

Correlation Parameter Derived Using The Cholesky Decomposition of The Covariance Matrix

We now utilize the Cholesky Decomposition to compute the correlation parameter P , as previously discussed. Experimental settings mirror those of the previous three experiments, examining the performance of the robust formulations under

a different method of calculating the correlation parameter. The Cholesky Decomposition of the covariance matrix (4.6) is computed as the following matrix:

$$P_{Cholesky} = \begin{bmatrix} 1.41421356 & 0.00000000 \\ -1.40007143 & 0.19949937 \end{bmatrix} \quad (4.8)$$




Figure 4.27: Percentile vs. Revenue. Five-capacity, Full Negative Correlation Example Using Cholesky Decomposition of The Covariance Matrix.

The results presented in Figure 4.27 reveal similar trends to the first experiment with dependent demand structures under the square root of the covariance matrix. As all methods depicted in Figure 4.24 and 4.27 found the same policies, observations do not change. Particularly, the EMSRb policy tends to exhibit higher or equal percentile values after the 40th percentile compared to the optimal controls, its mean realized revenue falls short by 3.22% when compared to the optimal solutions. Notably, the range of observed revenues is significant; while the optimal solutions consistently achieve revenues of at least 300 or 250, the EMSRb policy may only reach 200. Out of 600 simulation runs, EMSRb realized a revenue of 200 in 129 instances, whereas the algorithms achieved 250 in all of these runs.

Once again, we set the capacity to 80 to mitigate the impact of control rounding. The products are normally distributed with a mean of 60 and a standard deviation of 20, while the correlation matrix is characterized by a value of -0.99 for each non-diagonal entry, with a product fare ratio of 0.5.




Figure 4.28: Percentile vs. Revenue. Eighty-capacity, Full Negative Correlation Example Using The Cholesky Decomposition of The Covariance Matrix.

In Figure 4.28, we observe a result similar to Figure 4.25 due to the similar booking limits. While the good-case scenarios of the robust models may lack performance compared to the EMSRb policy, the revenue realizations are much more stable, with a significantly smaller range. Continuing with the same experimental setting, we adjust the fare ratio to 0.7, and the non-diagonal elements of the correlation matrix are set to -0.5 instead of -1.




Figure 4.29: Percentile vs. Revenue. Eighty-capacity, 0.7 Fare Ratio Example Using The Cholesky Decomposition of The Covariance Matrix.

Once more, due to the similarity in booking limit controls as seen in Figure 4.26, we observe a very similar situation in Figure 4.29. It's worth emphasizing that using both methods of deriving the correlation parameter, the robust formulations offer much more stable revenue realizations and effectively shield the decision maker from worst-case scenarios.

4.5.2 Comparison Under Multi-Product Scenario

In this section, we discuss the performance of the proposed algorithms in multi-product settings. In the first section of this chapter, different booking limit policies under changing uncertainty set sizes were also discussed in a similar setting, however, now we systematically select uncertainty parameters and comment on the possible revenues achieved.

Demand Factor Analysis

In the subsequent tables, we conduct a performance comparison between the praised *Algorithm-NL₁* and the EMSRb policy, as we observe this algorithm to be better performing than the rest of the robust models. For each of the capacity levels, we consider 2000 sample demand realizations. Based on the booking limits generated by both methods, we calculate and report the average revenue, along with the corresponding 95% confidence interval. Initially we provide a table showing the booking limit controls of the two methodologies under different demand factor levels. Experiment follows the demand information of Table 3.1.

C	DF	EMSRb Controls	<i>Algorithm-NL₁</i> Controls
68	2.00	(68, 51, 17, 0)	(68, 58, 40, 20)
75	1.81	(75, 58, 24, 0)	(75, 64, 44, 22)
85	1.60	(85, 68, 34, 2)	(85, 73, 50, 25)
95	1.43	(95, 78, 44, 12)	(95, 81, 55, 28)
105	1.29	(105, 88, 54, 22)	(105, 90, 61, 31)
115	1.18	(115, 98, 64, 32)	(115, 98, 67, 34)
125	1.08	(125, 108, 74, 42)	(125, 107, 73, 37)
135	1.00	(135, 118, 84, 52)	(135, 115, 79, 40)
145	0.93	(145, 128, 94, 62)	(145, 124, 85, 43)
155	0.87	(155, 138, 104, 72)	(155, 132, 90, 46)
165	0.82	(165, 148, 114, 82)	(165, 141, 96, 49)

Table 4.5: Booking Limits Under Different Demand Factor Levels for Table 3.1.

In the following two tables, we conduct the aforementioned experiments. In the first table, the customer arrival pattern is assumed to be LBH and in the second one, customers arrive in uniform arrival pattern.

C	DF	EMSRb	NL_1	SNR-EMSRb	SNR- NL_1
68	2.00	43,258.62 $\pm$ 191.49	41,570.12 $\pm$ 55.69	9.901	32.71
75	1.81	46,965.95 $\pm$ 207.97	45,671.80 $\pm$ 82.30	9.897	24.32
85	1.60	52,183.64 $\pm$ 213.58	51,433.06 $\pm$ 90.39	10.71	24.94
95	1.43	57,203.77 $\pm$ 219.45	57,238.64 $\pm$ 135.17	11.42	18.56
105	1.29	62,504.32 $\pm$ 216.15	62,565.65 $\pm$ 167.05	12.67	16.41
115	1.18	67,436.68 $\pm$ 227.47	67,531.60 $\pm$ 219.14	12.99	13.51
125	1.08	71,717.64 $\pm$ 288.15	71,581.19 $\pm$ 297.21	10.91	10.56
135	1.00	75,389.62 $\pm$ 343.53	74,918.36 $\pm$ 357.19	9.618	9.192
145	0.93	78,459.18 $\pm$ 417.03	77,927.83 $\pm$ 422.49	8.246	8.081
155	0.87	80,701.51 $\pm$ 482.65	80,068.63 $\pm$ 477.40	7.328	7.351
165	0.82	81,727.43 $\pm$ 539.00	81,235.59 $\pm$ 529.74	6.645	6.721

Table 4.6: Mean Revenues with 95% CI Under LBH Arrival Pattern For Table 3.1 with SNR Column.

From this experiment, considering the LBH arrival pattern, it's evident that, with the exception of the capacity levels of 125, 135, and 145, our formulation consistently outperforms the heuristic in terms of confidence interval width. Additionally, in some cases, our formulation yields higher average revenues. It's worth noting that while the relative performance of our method declines with increasing demand factors, the robust formulation stands out for its potential to reduce the observed revenue variability, offering an impressive decrease of up to 70.9% at a capacity level of 68, when using confidence interval width as the variability metric. However, it's important to highlight that EMSRb exhibits a minor improvement of only 3.9% in terms of revenue variability at the capacity level of 135.

C	DF	EMSRb	NL_1	SNR-EMSRb	SNR- NL_1
68	2.00	43,438.91 $\pm$ 199.81	42,275.90 $\pm$ 72.45	9.528	25.57
75	1.81	47,143.58 $\pm$ 212.66	46,491.39 $\pm$ 91.87	9.716	22.18
85	1.60	52,316.51 $\pm$ 230.77	52,297.09 $\pm$ 120.99	9.936	18.94
95	1.43	57,582.50 $\pm$ 224.81	57,900.91 $\pm$ 149.03	11.23	17.03
105	1.29	62,848.03 $\pm$ 230.28	63,259.07 $\pm$ 186.33	11.96	14.88
115	1.18	67,560.39 $\pm$ 250.10	67,724.52 $\pm$ 242.49	11.84	12.24
125	1.08	72,347.70 $\pm$ 301.68	72,101.57 $\pm$ 308.14	10.51	10.26
135	1.00	76,155.24 $\pm$ 347.16	75,458.33 $\pm$ 354.23	9.614	9.336
145	0.93	79,097.44 $\pm$ 434.20	78,331.96 $\pm$ 430.60	7.984	7.973
155	0.87	80,962.61 $\pm$ 507.63	80,170.83 $\pm$ 493.72	6.990	7.117
165	0.82	81,433.45 $\pm$ 554.80	80,865.92 $\pm$ 537.54	6.433	6.593


Table 4.7: Mean Revenues with 95% CI Under Uniform Arrival Pattern For Table 3.1 with SNR Column.

In the second experiment, we extended our analysis to a more realistic scenario where we made no assumptions about customer arrival patterns. As discussed earlier, we employed uniformly distributed customer arrivals, meaning that any type of customer could arrive at any point during the experiment, subject to the leftover booking limit controls.


In this experiment, akin to the LBH pattern, with the exceptions of capacity levels 125 and 135, our formulation excels in terms of achieving narrower 95% confidence intervals. The width of these intervals can be improved by up to a remarkable 63.7% when compared to the results obtained by the EMSRb method. However, it's worth noting that in the best-case scenario for EMSRb, it outperforms our proposed methodology by a mere 2% in terms of confidence interval width.

Additionally in both of these experiments, the SNR is better for *Algorithm- NL_1* except for few demand factor levels, where the difference is quite small. As the capacity of the resource increases, the robustness of the robust algorithm gets closer to the EMSRb policy, while still being more robust, highlighting lower revenue variability when using robust algorithms.

Analysis on The Distribution of The Realized Revenues




(a) Histogram for $Algorithm-N_1$




(b) Histogram for EMSRb

Figure 4.30: Realized Revenue Histograms for Table 3.1, Under LBH Arrival Pattern.




(a) Histogram for *Algorithm-N₁*




(b) Histogram for EMSRb

Figure 4.31: Realized Revenue Histograms for Table 4.8, Under LBH Arrival Pattern.



(a) Histogram for $Algorithm-N_1$



(b) Histogram for EMSRb

Figure 4.32: Realized Revenue Histograms for Table 4.8, Under Uniform Arrival Pattern.

In this part, we provide 3 histograms of realized revenues under various scenarios of the multi-product setting. In all three cases *Algorithm- N_1* has a lighter left-hand tail, as we observed this phenomenon for the two-product setting. Especially in Figure 4.31, revenue realizations under 60,000 are rare, while the EMSRb policy has multiple realizations, even below 50,000. Similarly, in Figure 4.32, *Algorithm- N_1* has most of its realizations concentrated above the 70,000 revenue mark, whereas EMSRb has considerable amount of realizations below 65,000. These experiments highlight the robustness of the robust algorithm *Algorithm- N_1* as across multiple experimental settings, we were able experimentally show that the algorithm is less likely to provide revenue realizations that are worse than its theoretical mean.

Illustrative Examples

To explore realistic scenarios in more detail, we again focus on the four-product single-resource example of Table 3.1, and conduct an experiment where we report the percentile of the achieved revenues during all simulation runs, using the same methodology to that of the two-product case.




Figure 4.33: Percentile vs. Revenue. LBH Arrival Pattern for Table 3.1.

Based on previous experiments, we know that the proposed algorithms get closer to the EMSRb and DLP methods, but slightly perform worse in terms of mean revenues. Although this is the case, in Figure 4.33, performance of these algorithms are comparable to those of EMSRb in the LBH arrival pattern. Especially, controls of $Algorithm-NL_{B_2}^2$ and $Algorithm-NL_1^{11.98}$, they protect from the worst-case scenarios compared to the EMSRb policy. The main difference between them and EMSRb is the protection level concerning two lowest-fare products. The proposed algorithms allow more lower fare classes to book, and in an LBH scenario, when the total demand is low and away from the the mean of the demand distributions, they are able to fill the capacity more. I such scenario, as the DLP and EMSRb largely limits these lowest two fares, they leave too many empty space when the bookings requests end.

Next we try a different demand scenario. Similar to Table 3.1, there is another example with evenly spaced fares [1]. This data structure is provided on the below table.

j	r_j	$\bar{d}_j$	σ_j
1	1050	17.3	5.8
2	950	45.1	15
3	699	39.6	13.2
4	520	34	11.3

Table 4.8: Single-resource, four-product example, evenly spaced fares [1].

Data of Table 4.8 has only 2 different entries compared to 3.1. While the demand distributions do not change, fares of the second and the third products are different, 567 change to 950 and 534 changed to 699, respectively. Previously, as the fares are quite close for three products, in terms of performance and in a nested by fare setting, booking limit of the second product was the most important decision. When the policies gave controls that have similar booking limits for this product, performances were close. However, in the evenly spaced fare example, booking limit decision of all products are more equally important.




Figure 4.34: Percentile vs. Revenue. LBH Arrival Pattern for Table 4.8.

The proposed algorithms and EMSRb reacted differently to the fare change. While the algorithms increased the booking limits of two most lower face classes, on the contrary, EMSRb decreased their booking limit. Especially, the lowest fare product has a booking limit of 3 when the EMSRb policy is used, compared to 36 and 28 from the algorithms. Also, algorithms decreased the booking limit of the second product while EMSRb increased its booking limit.

Figure 4.34, is not a promoter of the proposed algorithms in terms of realized revenues. However, one must note that while the robust formulations gave controls that outperform the existing methods in the worst-case scenarios, range of revenues realized is much smaller for them. For $Algorithm-NL_1^{11.98}$, 80% of the revenues fell in an interval with width 11,523, while this width is 24,446 for the EMSRb policy.

Based on the poor median revenue performance of our methods in the evenly spaced fare example, we decided to focus on their performance in random arrival pattern. The same scenario was used for the below figure and the booking limits are kept to be the same, as none of the methods consider arrival patterns.




Figure 4.35: Percentile vs. Revenue. Uniform Arrival Pattern for Table 4.8.

As expected, in Figure 4.35, performances of the proposed algorithms gotten closer to that of EMSRb and DLP, and upto the 30th percentile, managed to perform better than the DLP. However, still the protection form the worst-cases is not as eminent as the experiments with data structure Table 3.1. While the range of realized revenues gotten closer for all methods, still, robust formulation are able to provide much stable revenue realizations under changing scenarios.




Figure 4.36: Percentile vs. Revenue. Uniform Arrival Pattern for Table 4.8, With 120 Unit Capacity.

In Figure 4.36, a similar experiment is conducted, where only the capacity is changed to 120 compared to the previous experiment. We observe a similar result, where the robust models are better in worse scenarios and able to realize revenue in a much smaller range of values. However, when the capacity is increased, all the methods gotten loser, while the median revenue performance of the EMSRb is not as better as the previous ase with 100 capacity, similarly, range of realized revenues is not as large compared to the robust solutions. Highlighting that, when the demand factor is lower, good-case performances of the robust formulations get closer to the existing methods, however, they still provide protection when

the demand realization are adverse.

4.5.2.1 Dependent Demand for Multi-product Settings

Similar to the discussion with two-product scenarios, now we examine the performance of the various methods under correlated demand structure, however, in this subsection scenarios are more realistic, i.e., multiple number of products exist. Again, similar to the discussions on dependent demand structure for two-product setting, we refer to two versions of finding the correlation parameter P , both are discussed within their own subsections in the subsequent parts.

Using The Square Root of The Covariance Matrix

In the first experiment setting, we again refer to the data structure of Table 3.1, however, in this case we assume that all of the products are fully correlated. An example scenario would be, from one location to another location, demand for each product could be positively correlated with other products. Similar to two-case, the square root of the covariance matrix does not exist if all entries of the correlation matrix are 1. So, we assume that non-diagonal elements of the correlation matrix is 0.99.




Figure 4.37: Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 3.1 Using The Square Root of The Covariance Matrix.

In the independent experiments of the same data, the proposed algorithms were much more comparable to the EMSRb method. However, when the correlation parameter is introduced, controls from the proposed algorithms started to lag behind these other methods. While EMSRb and DLP used in Figure 4.37 do not use the correlation information, algorithms that use this information reported worse controls.




Figure 4.38: Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 4.8 Using The Square Root of The Covariance Matrix.

In Figure 4.38, $Algorithm-NL_1$ seem to better relative to EMSRb compared to Figure 4.37. Which is the opposite of the experiment conducted with the independent data, where $Algorithm-NL_1$ was better in terms of relative performance against EMSRb.

In the next experiment, first two and the last two products are assumed to be substitutes of each other, i.e., these groups of products are negatively correlated. Similar to previous discussions, to find the square root of the covariance matrix, the correlation parameter is calculated by assuming that the correlation matrix is the following.

$$Corr(\mathbb{X}) = \begin{bmatrix} 1.00 & -0.99 & 0.00 & 0.00 \\ -0.99 & 1.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 1.00 & -0.99 \\ 0.00 & 0.00 & -0.99 & 1.00 \end{bmatrix} \quad (4.9)$$




Figure 4.39: Percentile vs. Revenue. LBH Arrival Pattern for Two-negative Block Correlated Table 3.1 Using The Square Root of The Covariance Matrix.

While DLP managed to provide better controls in the previous experiments of this subsection, in Figure 4.39 the proposed algorithms managed to provide better performance. Moreover, generally $Algorithm-NL_1$ was better compared to others, however, in this case, $Algorithm-NL_\infty$ provided a very close performance in worst case scenarios and provided the best solution for the best case scenarios. In Figure 4.39, EMSRb still proves to be the best solution method, the proposed algorithms do not lag behind as much as the fully correlated experiment settings, relative to the EMSRb policy.

Using The Cholesky Decomposition of The Covariance Matrix

In this subsection, similar to the previous one, we start with fully correlated demand assumption using the demand structure of Table 3.1. The correlation matrix is constructed following the perturbation discussions before.




Figure 4.40: Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 3.1 Using The Cholesky Decomposition of The Covariance Matrix.

In Figure 4.40, we observe a similar downside of the proposed algorithms that we have also observed in the previous subsection. When we used the square root of the covariance matrix as the correlation parameter, resulting controls were also extreme. Extremity in this case refers to the booking limit controls that resemble no booking limit controls. Using Cholesky Decomposition for this case reports a worse performance mainly due to algorithms limit too much of the lower fare class allocation while EMSRb policy allow them to book. When the demand realization is one of the worst cases, i.e., total demand is low, *Algorithm-NL₁³* will not be able to fill the capacity, it only allows 10 bookings from the lower 2 classes, leaving 90 space to higher 2 classes. Assuming no correlation realizing more than 80 arrivals from the first two classes is less than 0.14, leaving too many empty capacity.




Figure 4.41: Percentile vs. Revenue. LBH Arrival Pattern for Fully Correlated Table 4.8 Using The Cholesky Decomposition of The Covariance Matrix.

Similar to Figure 4.40, in Figure 4.41, we observe extreme booking limits by the proposed algorithms, leaving too any empty seats in worst case scenarios. Although, while in evenly distributed fare experiments of the independent demand assumption, the algorithms detrimented in performance compared to the standard example, in this case they gotten better relative to the EMSRb performance. We observe that while introducing correlation the algorithms turn booking limits extreme, using this information might increase their potential in comparison to existent methods.

Now, we refer to the below correlation matrix in the next experiment.

$$Corr(\mathbb{X}) = \begin{bmatrix} 1.00 & -0.99 & 0.00 & 0.00 \\ -0.99 & 1.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 1.00 & -0.99 \\ 0.00 & 0.00 & -0.99 & 1.00 \end{bmatrix} \quad (4.10)$$




Figure 4.42: Percentile vs. Revenue. LBH Arrival Pattern for Two-negative Blocks Correlated Table 3.1 Using The Cholesky Decomposition of The Covariance Matrix.

First thing to notice in Figure 4.42 is that the algorithms provided controls that are of interest to the purpose of this thesis. $Algorithm-NL_1$ is showing better worst case scenario performance, even though relative to EMSRb policy it is small. Even up to 40th percentile levels, $Algorithm-NL_1$ manages to outperform EMSRb, protecting the decision maker in regards to unwanted demand realizations.

In Appendix F, few of the examples discussed in the previous experiments are presented under various robustness metrics including standard deviation, skewness, kurtosis, and SNR. Also in Appendix G, experiments that test partitioned booking limit policies are given. The booking limits for robust models are found using the variants of $Algorithm-PL$.

4.5.3 Forecasting Errors Leading to Booking Limit Calculations

Now we investigate the effects of estimation errors of the demand when computing booking limits. We also report and comment on the performances achieved by different methodologies when such estimation errors occur. For this section the experiments only done under independent data structure and uses the data from Table 3.1.

First we start by forecasting errors for the mean demand value. The booking limits are calculated using the aforementioned data structure, however, when taking instances of the random variables in the simulations, mean demand of the products are multiplied with a constant α , keeping the standard deviations same.

In the first figure of this section, we provide a replication of the experiment of Figure 4.33. In the following experiments, when we assume that the demand is either over or under estimated, controls from the policies will be the same as this experiment. However, as mentioned, demand of each product, when simulating realizations, are assumed to be $\alpha \bar{d}_j \forall j \in N$.




Figure 4.43: Percentile vs. Revenue. LBH Arrival Pattern for Table 3.1.

Now given the experiment with complete and true information regarding the demand structures is provided, we provide an example setting where the mean of the demand is under estimated. In the next figure α is taken as 1.15 and 1.25.




Figure 4.44: Percentile vs. Revenue of Under Estimating The Mean Demand Value using Table 3.1 Under LBH Arrival Pattern.

Before commenting on the results, we investigate the effects of over estimating the demand by taking α as 0.75 and 0.85 in two experiments. It is important to note that, the percentile levels that the robust models start to show better revenues, is being shifted to lower levels as the error constant increased. Lowest percentile that the EMSRb policy beats $Algorithm-NL_1$ at 50th percentile when the $\alpha = 1$, at 30th and 20th when $\alpha = 1.15$ and $\alpha = 1.25$, respectively.




Figure 4.45: Percentile vs. Revenue of Over Estimating The Mean Demand Value using Table 3.1 Under LBH Arrival Pattern.

It is important to note that, in contrast to the previous figure, the percentile levels that the robust models start to show better revenues, is being shifted to higher levels as the error constant decreased. Lowest percentile that the EMSRb policy beats $Algorithm-NL_1$ at 50th percentile when the $\alpha = 1$, at 80th and 90th when $\alpha = 0.85$ and $\alpha = 0.75$, respectively.

In these experiments, value of the robust models are shown. In all of the percentile experiments, actually, the left side corresponds to the worst cases, which we also referred to in the previous explanations. On the other hand the right side of the percentile figures correspond to the good cases. When the mean of all demand distributions are multiplied with a constant, focus of these percentile figures shift. Basically, when we assume a demand structure and find booking limits, but then in the experimental studies, assume that the actual revenue was 0.75 times the estimated value, we shift to the worst cases, where the capacity may not be satisfied as intended. So in the above experiments, we can actually observe that when the error constant is less than 1, we shift to the left part of the original percentile vs. revenue experiments. A similar argument is applicable for the error

constants of larger than 1. As a result, the robust models are promising when the mean demand is overestimated, their efficiency in worst-case performance shift to expected or good scenarios. However, when the mean demand is realized more than the forecast, their performance diminish compared to the EMSRb policy, which is expected as there is a direct correspondence between good scenarios and underestimated mean demand.

Now, we investigate the effect of forecasting errors related to the standard deviation of the demand distributions. While as not detrimental as the errors related to the forecasting the mean, extreme mistakes might be visible through the percentile experiments.




Figure 4.46: Percentile vs. Revenue of Over and Under Estimating The Variance using Table 3.1 Under LBH Arrival Pattern.

Similar to our discussions with the over and under estimation of the mean of the demand distribution, when the standard deviation is underestimated, robust models tend to work better compared to their counterparts. On the contrary, when the demand is much more stable than it was as estimated, robust models tend to lose their performance. Again, this is an expected result, as the robust

models promise to work to protect the decision maker when the demand realizations are extreme.

4.6 Summary of Observations

- The uncertainty set size is hard to define, but in the independent demand case, *Algorithm-NL₁* provide consistent booking limits. However, in the correlated demand case, defining the uncertainty set size is problematic, requiring a need for sensitivity analysis.
- For the independent demand scenario, the proposed algorithm performs well, with *Algorithm-NL₁* almost surpassing the expected revenue of the EMSRb policy in some realistic scenarios even though the method is developed for worst-case performance.
- In the dependent demand scenario, the controls of the robust models are overly conservative, leading to too many rejections and suboptimal performance.
- In terms of robustness, based on the SNR, our models perform well due to the low sample standard deviation of the realized revenues. Also, because the standard deviations are lower, the 95% CIs for the mean revenue are narrower compared to those of EMSRb. This indicates a more consistent range of possible revenue realizations.
- In histograms of realized revenues, the robust models have a lighter left-hand tail, indicating a smaller probability of realizing revenues that are far below the mean of the realized revenue.
- In terms of worst-case performance based on lower-percentile levels, our models consistently outperform EMSRb, with *Algorithm-NL₁* always achieving a higher 10th percentile level.
- Under LBH arrival patterns, existing models, especially EMSRb, perform well, and our models are comparable. However, when the arrival pattern

is uniform, our models retain more of their performance compared to the benchmarks.

- When demand decreases due to disturbances and the true mean demand is overestimated, robust models outperform EMSRb in terms of worst-case, average-case performance, as well as mean revenue.



Chapter 5

Conclusion and Future Research

We proposed new robust methods for single-leg revenue management. All our robust models reduce to a linear program which can be solved in practical applications, where the dynamic models get too computationally complicated. We show that the effectiveness using various simulations highlighting the mean revenue performance and worst-case analysis. The robust models, consistently provided better worst-case performance compared to the benchmark methods.

While DLP only uses the first moment information of the demand distribution, our robust formulations use first, second, and the product moments. Enabling streamlined solution methodologies under dependent demand for different products.

The work on this thesis follows the previous work on the field, and provides a solution to the seat-inventory problem in compact and fast manner, enabling large instances to be quickly investigated.

In the case of overestimating the expected demand, the robust models prove more effective than the classical models EMSRb and DLP. They provide protection for the decision makers in the event of an unexpected scenario where the demand for all products significantly decrease.

One downside to the robust formulations is the effect of the uncertainty set parameters. While we have recommended a selection scheme for these parameters, for some uncertainty sets, these parameters greatly change the booking limit controls. As it is hard to understand the effect of these parameters, future research on selecting these parameters or a new line of research where such parameter is not to be determined would be an interesting work.

Moreover, the performance of the robust models in some experiments are not as good as their performance in other experiments compared to the benchmark methods. More experiments are needed to further test the performance of the robust models under specific demand information scenarios, especially those with correlated demand and inaccurate forecasts. In practice, major events that effect airline demand such as the COVID-19 pandemic and the aftermath of September 11th terrorist attacks, led to unprecedented changes in demand, correlated across all fare classes.

Bibliography

- [1] K. Talluri and G. van Ryzin, *The Theory and Practice of Revenue Management*. International Series in Operations Research & Management Science, Springer US, 2005.
- [2] K. Littlewood, “Special issue papers: Forecasting and control of passenger bookings,” *Journal of Revenue and Pricing Management*, vol. 4, no. 2, pp. 111–123, 2005.
- [3] K. T. Talluri, G. J. van Ryzin, I. Z. Karaesmen, and G. J. Vulcano, “Revenue management: Models and methods,” in *2008 Winter Simulation Conference*, pp. 145–156, 2008.
- [4] G. Gallego and H. Topaloglu, *Revenue Management and Pricing Analytics*. International Series in Operations Research & Management Science, Springer New York, 2019.
- [5] P. P. Belobaba, “Application of a probabilistic decision model to airline seat inventory control,” *Operations Research*, vol. 37, no. 2, pp. 183–197, 1989.
- [6] P. Belobaba and A. Y. M. S. Group, *Optimal Vs. Heuristic Methods for Nested Seat Allocation: Presentation to the AGIFORS Yield Management Study Group, Brussels, Belgium, May 4, 1992*. AGIFORS Yield Management Study Group, 1992.
- [7] D. Adelman, “Dynamic bid prices in revenue management,” *Operations Research*, vol. 55, no. 4, pp. 647–661, 2007.

- [8] R. W. Simpson, “Using network flow techniques to find shadow prices for market and seat inventory control,” Memorandum M89-1, MIT Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA, 1989.
- [9] E. L. Williamson, *Airline Network Seat Control*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1992.
- [10] K. T. Talluri and G. J. van Ryzin, “An analysis of bid-price controls for network revenue management,” *Management Science*, vol. 44, no. 12, pp. 1577–1593, 1998.
- [11] T. C. Lee and M. Hersh, “A model for dynamic airline seat inventory control with multiple seat bookings,” *Transportation Science*, vol. 27, no. 3, pp. 252–265, 1993.
- [12] C. J. Lautenbacher and S. Stidham, “The underlying markov decision process in the single-leg airline yield-management problem,” *Transportation Science*, vol. 33, no. 2, pp. 136–146, 1999.
- [13] D. Bertsimas and I. Popescu, “Revenue management in a dynamic network environment,” *Transportation Science*, vol. 37, no. 3, pp. 257–277, 2003.
- [14] H. Topaloglu, “Using lagrangian relaxation to compute capacity-dependent bid prices in network revenue management,” *Operations Research*, vol. 57, no. 3, pp. 637–649, 2009.
- [15] R. E. Curry, “Optimal airline seat allocation with fare classes nested by origins and destinations,” *Transportation Science*, vol. 24, no. 3, pp. 193–204, 1990.
- [16] R. D. Wollmer, “An airline seat management model for a single leg route when lower fare classes book first,” *Operations Research*, vol. 40, no. 1, pp. 26–37, 1992.
- [17] S. L. Brumelle and J. I. McGill, “Airline seat allocation with multiple nested fare classes,” *Operations Research*, vol. 41, no. 1, pp. 127–137, 1993.

- [18] L. W. Robinson, “Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes,” *Operations Research*, vol. 43, no. 2, pp. 252–263, 1995.
- [19] Z. Chi, *Airline yield management in a dynamic network environment*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1995.
- [20] D. Bertsimas and S. de Boer, “Simulation-based booking limits for airline revenue management,” *Operations Research*, vol. 53, no. 1, pp. 90–106, 2005.
- [21] S. V. de Boer, R. Freling, and N. Piersma, “Mathematical programming for network revenue management revisited,” *European Journal of Operational Research*, vol. 137, no. 1, pp. 72–92, 2002.
- [22] G. van Ryzin and J. McGill, “Revenue management without forecasting of optimization: An adaptive algorithm for determining airline seat protection levels,” *Management Science*, vol. 46, no. 6, pp. 760–775, 2000.
- [23] L. R. Weatherford and P. P. Belobaba, “Revenue impacts of fare input and demand forecast accuracy in airline yield management,” *The Journal of the Operational Research Society*, vol. 53, no. 8, pp. 811–821, 2002.
- [24] G. van Ryzin and G. Vulcano, “Simulation-based optimization of virtual nesting controls for network revenue management,” *Operations Research*, vol. 56, no. 4, pp. 865–880, 2008.
- [25] A. Ben-Tal and A. Nemirovski, “Robust truss topology design via semidefinite programming,” *SIAM Journal on Optimization*, vol. 7, no. 4, pp. 991–1016, 1997.
- [26] L. El Ghaoui, F. Oustry, and H. Lebret, “Robust solutions to uncertain semidefinite programs,” *SIAM Journal on Optimization*, vol. 9, no. 1, pp. 33–52, 1998.
- [27] L. El Ghaoui and H. Lebret, “Robust solutions to least-squares problems with uncertain data,” *SIAM Journal on Matrix Analysis and Applications*, vol. 18, no. 4, pp. 1035–1064, 1997.

- [28] A. Ben-Tal and A. Nemirovski, “Robust solutions of uncertain linear programs,” *Operations Research Letters*, vol. 25, no. 1, pp. 1–13, 1999.
- [29] A. Ben-Tal and A. Nemirovski, “Robust solutions of linear programming problems contaminated with uncertain data,” *Mathematical Programming*, vol. 88, pp. 411–424, 01 2000.
- [30] A. Ben-Tal and A. Nemirovski, “Robust convex optimization,” *Mathematics of Operations Research*, vol. 23, no. 4, pp. 769–805, 1998.
- [31] D. Bertsimas and M. Sim, “The price of robustness,” *Operations Research*, vol. 52, no. 1, pp. 35–53, 2004.
- [32] H. E. Scarf, “A min-max solution to an inventory problem,” in *Studies in Mathematical Theory of Inventory and Production* (K. J. Arrow, S. Karlin, and H. E. Scarf, eds.), pp. 201–209, Stanford, CA: Stanford University Press, 1958.
- [33] G. Gallego and I. Moon, “The distribution free newsboy problem: Review and extensions,” *The Journal of the Operational Research Society*, vol. 44, no. 8, pp. 825–834, 1993.
- [34] J. K. R. Guillermo Gallego and D. Simchi-Levi, “Minimax analysis for finite-horizon inventory models,” *IIE Transactions*, vol. 33, no. 10, pp. 861–874, 2001.
- [35] A. Ben-Tal, L. Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton Series in Applied Mathematics, Princeton University Press, 2009.
- [36] D. Bertsimas and D. den Hertog, *Robust and Adaptive Optimization*. Dynamic Ideas LLC, 2022.
- [37] G. Perakis and G. Roels, “Robust controls for network revenue management,” *Manufacturing & Service Operations Management*, vol. 12, no. 1, pp. 56–76, 2010.
- [38] Y. Lan, H. Gao, M. O. Ball, and I. Karaesmen, “Revenue management with limited demand information,” *Management Science*, vol. 54, no. 9, pp. 1594–1609, 2008.

- [39] S. I. Birbil, J. B. G. Frenk, J. A. S. Gromicho, and S. Zhang, “The role of robust optimization in single-leg airline revenue management,” *Management Science*, vol. 55, no. 1, pp. 148–163, 2009.
- [40] A. Thiele, *A robust optimization approach to supply chains and revenue management*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [41] J. Gönsch, “A survey on risk-averse and robust revenue management,” *European Journal of Operational Research*, vol. 263, no. 2, pp. 337–348, 2017.
- [42] J. An, A. Mikhaylov, and S.-U. Jung, “A linear programming approach for robust network revenue management in the airline industry,” *Journal of Air Transport Management*, vol. 91, p. 101979, 2021.
- [43] A. Teymouri, H. Sahebi, and M. Pishvaei, “Airline operational crew-aircraft planning considering revenue management: A robust optimization model under disruption,” *International Journal of Industrial Engineering Computations*, vol. 14, no. 2, pp. 381–402, 2023.
- [44] T. Goncalves and B. Almada-Lobo, “Enhancing robustness to forecast errors in availability control for airline revenue management,” *Journal of Revenue and Pricing Management*, 2024.
- [45] R. Rockafellar, *Convex Analysis*. Princeton mathematical series, Princeton University Press, 1970.
- [46] G. B. Dantzig, “Discrete-variable extremum problems,” *Operations Research*, vol. 5, no. 2, pp. 266–277, 1957.
- [47] M. Doumpos, C. Zopounidis, and E. Grigoroudis, eds., *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, vol. 237 of *International Series in Operations Research & Management Science*. Springer Cham, 1 ed., 2016. Springer International Publishing Switzerland 2016.
- [48] V. Voudouris, K. Matsumoto, J. Sedgwick, R. Rigby, D. Stasinopoulos, and M. Jefferson, “Exploring the production of natural gas through the lenses of the ACEGES model,” *Energy Policy*, vol. 64, no. C, pp. 124–133, 2014.

Appendix A

Equivalent Linear Programming Examples of The Robust Models with $n = 2$

In a two-product problem, we need the vectors $(r_1, r_2), (\bar{d}_1, \bar{d}_2)$ and the matrix $P = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix}$ to formulate the robust models. When the demand of classes 1 and 2 are not correlated (i.e. independent demand distributions), then

$$P = \begin{bmatrix} P_{11} & 0 \\ 0 & P_{22} \end{bmatrix}$$

Example B.1. The demand is independent in this example. The ℓ_∞ -norm is used in the uncertainty set of the robust formulation. The equivalent linear programming formulation is:

$$\begin{aligned}
& \underset{0 \leq y_1, y_2 \leq 1}{\text{maximize}} && r_1 \bar{d}_1 y_1 + r_2 \bar{d}_2 y_2 - \epsilon(t_1 + t_2) \\
& \text{subject to} && a_{11} \bar{d}_1 y_1 + a_{12} \bar{d}_2 y_2 + \epsilon(w_1 + w_2) \leq C, \\
& && -t_1 \leq r_1 P_{11} y_1 \leq t_1, \\
& && -t_2 \leq r_2 P_{22} y_2 \leq t_2, \\
& && -w_1 \leq a_{11} P_{11} y_1 \leq w_1, \\
& && -w_2 \leq a_{12} P_{22} y_2 \leq w_2, \\
& && t_1, t_2, w_1, w_2 \geq 0.
\end{aligned} \tag{A.1}$$

Example B.2. The demand is independent in this example. The ℓ_1 -norm is used in the uncertainty set of the robust formulation. The equivalent linear programming formulation is:

$$\begin{aligned}
& \underset{0 \leq y_1, y_2 \leq 1}{\text{maximize}} && r_1 \bar{d}_1 y_1 + r_2 \bar{d}_2 y_2 - \epsilon \rho_{obj} \\
& \text{subject to} && a_{11} \bar{d}_1 y_1 + a_{12} \bar{d}_2 y_2 + \epsilon \rho_{constr} \leq C, \\
& && r_1 P_{11} y_1 \leq \rho_{obj}, \\
& && r_2 P_{22} y_2 \leq \rho_{obj}, \\
& && a_{11} P_{11} y_1 \leq \rho_{constr}, \\
& && a_{12} P_{22} y_2 \leq \rho_{constr}, \\
& && \rho_{obj}, \rho_{constr} : \text{free}.
\end{aligned} \tag{A.2}$$

Example B.3. The demand is independent in this example. The ℓ_∞ and ℓ_1 -norms are used in the uncertainty set of the robust formulation to build a budgeted uncertainty set. The equivalent linear programming formulation is:

$$\begin{aligned}
& \underset{0 \leq y_1, y_2 \leq 1}{\text{maximize}} && r_1 \bar{d}_1 y_1 + r_2 \bar{d}_2 y_2 - \rho(t_1 + t_2) - \Gamma \rho_{obj} \\
& \text{subject to} && a_{11} \bar{d}_1 y_1 + a_{12} \bar{d}_2 y_2 + \rho(w_1 + w_2) + \Gamma \rho_{constr} \leq C, \\
& && P'_{11} y_1 + x'_1 \leq \rho_{obj}, \\
& && P'_{22} y_2 + x'_2 \leq \rho_{obj}, \\
& && P_{11} y_1 + x_1 \leq \rho_{constr}, \\
& && P_{22} y_2 + x_2 \leq \rho_{constr}, \\
& && -t_1 \leq x'_1 \leq t_1, \\
& && -t_2 \leq x'_2 \leq t_2, \\
& && -w_1 \leq x_1 \leq w_1, \\
& && -w_2 \leq x_2 \leq w_2, \\
& && t_1, t_2, w_1, w_2 \geq 0, \\
& && \rho_{obj}, \rho_{constr} : \text{free}, \\
& && x'_1, x'_2, x_1, x_2 : \text{free}.
\end{aligned} \tag{A.3}$$

Example B.4. The demand is independent in this example. Intervals are used to define the uncertainty set of the robust formulation. The equivalent linear programming formulation is:

$$\begin{aligned}
& \underset{0 \leq y_1, y_2 \leq 1}{\text{maximize}} && r_1 l_1 y_1 + r_2 l_2 y_2 \\
& \text{subject to} && a_{11} u_1 y_1 + a_{12} u_2 y_2 \leq C.
\end{aligned} \tag{A.4}$$

Appendix B

Effect of Standard Deviation - I

(σ_1, σ_2)	<i>Algorithms</i>			Littlewood's Rule	DLP
	NL_∞^1	NL_1	NL_{B_1}		
(5 , 5)	35	67	35	40	40
(10, 5)	32	80	32	40	40
(15, 5)	28	86	28	40	40
(20, 5)	24	89	24	40	40
(5 , 10)	33	50	33	40	40
(10, 10)	30	67	30	40	40
(15, 10)	26	75	26	40	40
(20, 10)	22	80	22	40	40
(5 , 15)	32	40	32	40	40
(10, 15)	29	57	29	40	40
(15, 15)	25	67	25	40	40
(20, 15)	21	73	21	40	40
(5 , 20)	30	33	30	40	40
(10, 20)	27	50	27	40	40
(15, 20)	24	60	24	40	40
(20, 20)	20	67	20	40	40

Table B.1: Booking Limit Controls Under Changing Standard Variations When the Fare Ratio Is 0.5.

Starting with Table B.1, Littlewood’s Rule suggests that a booking limit of 40 is optimal for any (σ_1, σ_2) combination. However, while these two algorithms decrease the booking limit of the lower-fare product as the variation of the first product increases, *Algorithm-NL₁* behaves conversely by increasing the limits.

While not providing mathematical reasoning, through observation and intuition, robust models are expected to protect against possible adverse scenarios that existing methods tend to overlook due to their low probability of occurrence. In the case of increasing variation of the higher-fare product, the probability of realizing lower demand for this class increases. Hence, it is expected from the robust solutions to allow more bookings from the lower fare classes, thus protecting revenue from such worst-case scenarios.

All three algorithms tend to decrease the booking limit when the variation of the second product increases. In such a scenario, maintaining the booking limit policy unchanged would be more intuitive. In terms of robustness, a decrease in the booking limit due to increased variation of the lower-fare product implies that the model may be unable to protect the decision maker in the event of an extremely low demand realization for this class. However, as the probability of higher demand realizations for the lower-fare class increases due to heavier tail probabilities, the booking limit could be more readily exhausted by this fare, resulting in fewer capacities being allocated to the higher-fare class. Thus, decreasing the booking limit of the lower-fare product when the standard deviation of this product increases still adheres to the rationale behind robust models.




Figure B.1: Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.5 Under LBH Arrival Pattern.

From Table B.1, it's apparent that *Algorithm-NL₁* consistently allowed more bookings for the lower-fare class compared to the optimal policy, resulting in lower achieved revenues. However, as observed in previous experiments, this tendency enables robust models to provide relatively stable revenues with smaller variations. While a negative correlation is noticeable between the value of σ_1 and mean revenues, this trend doesn't necessarily extend to the value of σ_2 . Only the performance of *Algorithm-NL₁* relative to other methods seems to be positively correlated with this value, a promising feature where it's the only method affected by the variation in the lower-fare product. The next figure presents the same

experiment under a uniform arrival pattern, offering insights into the performance of robust models under more realistic scenarios.




Figure B.2: Mean Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.5 Under Uniform Arrival Pattern.

The sensitivity to changing variation under realistic scenarios is more apparent in Figure B.2. At every fixed σ_2 , while other methods show decreasing mean revenues as σ_1 increases, *Algorithm-NL₁* is more resilient. It even manages to report better mean revenues under high σ_1 , which is not expected from the proposed robust models. However, the relative performance against Littlewood's Rule of other robust models seems to be negatively correlated with both σ_1 and σ_2 .



Appendix C

Effect of Standard Deviation - II




Figure C.1: 25th Percentile of The Realized Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under Uniform Arrival Pattern.




Figure C.2: 25th Percentile of The Realized Revenues Under Changing Standard Deviations When the Fare Ratio Is 0.7 Under LBH Arrival Pattern.

Appendix D

Analysis of 2-Class Example

Controls	Demand Scenarios:	(1,5)	(2,4)	(3,3)	(4,2)	(5,1)	Expected Revenue	Revenue Range
$b_2 = 5$	No. Admitted Revenue	(0,5) \$250	(1,4) \$300	(2,3) \$350	(3,2) \$400	(4,1) \$450	\$350	[\$250, \$450]
$b_2 = 4$	No. Admitted Revenue	(1,4) \$300	(1,4) \$300	(2,3) \$350	(3,2) \$400	(4,1) \$450	\$360	[\$300, \$450]
$b_2 = 3$	No. Admitted Revenue	(1,3) \$250	(2,3) \$350	(2,3) \$350	(3,2) \$400	(4,1) \$450	\$360	[\$250, \$450]
$b_2 = 2$	No. Admitted Revenue	(1,2) \$200	(2,2) \$300	(3,2) \$400	(3,2) \$400	(4,1) \$450	\$350	[\$200, \$450]
$b_2 = 1$	No. Admitted Revenue	(1,1) \$150	(2,1) \$250	(3,1) \$350	(4,1) \$450	(4,1) \$450	\$330	[\$150, \$450]
$b_2 = 0$	No. Admitted Revenue	(1,0) \$100	(2,0) \$200	(3,0) \$300	(4,0) \$400	(5,0) \$500	\$300	[\$100, \$500]

Table D.1: Two-product and 5 Units of Capacity Example, All Possible Booking Limit Controls

Appendix E

Uniformly Distributed Demand

Based on the example in [38], we begin to investigate demand distributions that are defined for a finite interval, mainly, uniformly distributed. In our first experiment, the demand for the two products is uniformly distributed between 40 and 80, with fares assumed to be 1000 and 500.

To compare with our proposed methods, which require standard deviation and mean value information, instead of uniform distributions, we assume that the input demands are normally distributed with a mean of 60 and a standard deviation of 30, where the 25th and 75th percentiles coincide with the extreme points of the uniform distribution, 40 and 80. This case is referred to as *40-80 Interval*.

For this section, we provide experiments under two real demand structures. First, we assume that the real demand is, in fact, uniformly distributed. In the second plots, we find a corresponding normal distribution from the given upper and lower demand information and assume that the real demand is distributed normally under this randomness.




Figure E.1: Percentile vs. Revenue. LBH Arrival Pattern Under 40-80 Interval.




Figure E.2: Percentile vs. Revenue. LBH Arrival Pattern Under 40-80 Interval Variant with Standard Deviations of 60 and 30.

Appendix F

On The Robustness of The Realized Revenues

We provide various metrics for 5 experimental settings, focusing on robustness of *Algorithm-NL₁* and EMSRb. For each experiment we present mean, standard deviation, skewness and kurtosis information regarding the realized revenues after 2000 simulation runs. “Skewness-2” and “Kurtosis-2” refers to quantile-based calculations as discussed in Chapter 4, section “Performance Metrics.” Below, we give a list of the experimental setups of the following Table F.1. Note that the arrival pattern is assumed to be LBH and the capacity of the resource is 100.

- *Experiment 1*: Follows 3.1
- *Experiment 2*: Follows 4.8
- *Experiment 3*: Two-products having 60 mean demand and 20 standard deviation, while the fare ratio is 0.7.
- *Experiment 4*: The same setup as the previous experiment, except the standard deviation are now 10.
- *Experiment 5*: The same setup as *Experiment 3*, where the low-fare product has a standard deviation of 5 instead of 20.

Metrics	<i>Experiment 1</i>		<i>Experiment 2</i>		<i>Experiment 3</i>		<i>Experiment 4</i>		<i>Experiment 5</i>	
	NL_1	EMSRb	NL_1	EMSRb	NL_1	EMSRb	NL_1	EMSRb	NL_1	EMSRb
Mean ($\mu, 10^3$)	59.97	59.79	73.10	79.61	79.84	80.07	82.82	84.48	80.10	80.97
Std. Dev. ($s, 10^2$)	34.92	48.60	48.10	88.13	88.04	100.9	23.08	42.45	57.81	85.93
SNR	17.17	12.30	15.20	9.03	9.07	7.93	35.87	19.89	13.85	9.42
Skewness-1	-2.26	-1.61	-0.91	-1.64	-2.29	-1.93	-3.35	-2.55	-3.60	-2.67
Skewness-2	-0.22	-0.19	-0.11	-0.17	-0.14	-0.21	0.25	-0.25	-0.09	-0.25
Kurtosis-1	6.67	2.94	2.20	2.48	6.28	3.65	26.48	7.27	15.21	7.51
Kurtosis-2	3.11	1.78	2.25	2.55	8.55	2.96	2.50	4	3.17	5.33

Table F.1: Metrics Related to The Distribution of Realized Revenues.

Appendix G

Partitioned Booking Limit Examples




Figure G.1: Percentile vs. Revenue for Table 3.1 Under LBH Arrival Pattern, $C = 100$.




Figure G.2: Percentile vs. Revenue for Table 4.8 Under LBH Arrival Pattern, $C = 100$.