



HAL
open science

Data driven optimisation models for the operations planning in industrial supply chains

Benoit Loger

► **To cite this version:**

Benoit Loger. Data driven optimisation models for the operations planning in industrial supply chains. Operations Research [math.OC]. Ecole nationale supérieure Mines-Télécom Atlantique, 2023. English. NNT : 2023IMTA0389 . tel-04497849

HAL Id: tel-04497849

<https://theses.hal.science/tel-04497849v1>

Submitted on 11 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT DE

L'ÉCOLE NATIONALE SUPÉRIEURE
MINES-TÉLÉCOM ATLANTIQUE BRETAGNE
PAYS-DE-LA-LOIRE - IMT ATLANTIQUE

ÉCOLE DOCTORALE N° 648
Sciences pour l'Ingénieur et le Numérique
Spécialité : *Informatique*

Par

Benoit LOGER

**Modèles d'optimisation basés sur les données pour l'optimisation
des opérations dans les Supply Chain industrielles**

Thèse présentée et soutenue à IMT Atlantique, Nantes, le 21/12/2023
Unité de recherche : Laboratoire des Sciences du Numérique de Nantes (LS2N)
Thèse N° 2023IMTA0389

Rapporteurs avant soutenance :

André ROSSI Professeur à Université Paris Dauphine
Céline GICQUEL Maître de Conférences à Université Paris Saclay

Composition du Jury :

Président :	Nadjib BRAHIMI	Professeur à Rennes School of Business
Examineur :	André ROSSI	Professeur à Université Paris Dauphine
	Céline GICQUEL	Maître de Conférences à Université Paris Saclay
	Stefan MINNER	Professeur à Technical University of Munich
	Evren SAHIN	Professeure à CentraleSupélec
Dir. de thèse :	Fabien LEHUÉDÉ	Professeur à IMT Atlantique
Co-encadrants :	Guillaume MASSONNET	Maître Assistant à IMT Atlantique
	Alexandre DOLGUI	Professeur à IMT Atlantique

RÉSUMÉ EN FRANÇAIS

Dans notre société, la majorité des biens de consommation est produite et distribuée par des entreprises dont l'objectif est de satisfaire efficacement (à moindre coût et avec une haute qualité de service) la demande de leurs clients. Depuis de nombreuses décennies, les logiques de compétitivité ont poussé de nombreuses entreprises à se spécialiser dans une unique tâche (de production ou de service) et à collaborer avec d'autres entités spécialisées. De ce fait, les circuits de production et de distribution sont devenus de plus en plus complexes, incluant une grande diversité d'acteurs amenés à collaborer pour satisfaire les besoins finaux. De tels réseaux d'entreprises sont appelés *chaînes d'approvisionnement* (ou *Supply Chain* en anglais). Leur diversification et leur complexification durant les dernières décennies a poussé les gestionnaires de ces réseaux à caractériser formellement leurs dynamiques afin d'en optimiser le fonctionnement. L'identification de ces besoins et problématiques, génériques comme spécifiques, ont motivé de nombreux travaux de recherches et de nombreux livres portant sur le management des chaînes d'approvisionnement (Stadtler et al. (2015), Christopher (2005), Dolgui and Proth (2010), Simchi-Levi (2008)). Ces travaux ont ainsi contribué à décrire la structure et les dynamiques de ces chaînes d'approvisionnement en détail, à identifier leurs failles et les potentiels leviers d'améliorations. Cette thèse s'inscrit dans la lignée de ces travaux et se concentre sur l'optimisation des processus au niveau opérationnel (planification de production, gestion de stock, ...).

Comme évoqué plus haut, la collaboration entre les différents maillons d'une chaîne d'approvisionnement est une des clés clairement identifiées pour améliorer sa qualité de service et diminuer les inefficiences Dolgui and Proth (2010). Alors que l'accroissement des flux de données industrielles devraient permettre à chaque entreprise de baser ses décisions sur une vision globale de la chaîne d'approvisionnement, le partage d'information se heurte bien souvent aux réticences de groupes parfois en compétition et soucieux de préserver leur avantage concurrentiel. Ainsi, on constate que de nombreux dirigeants et/ou gestionnaires ne partagent que partiellement les informations nécessaires aux autres membres d'une chaîne logistique pour adapter leurs décisions. Ce phénomène est à l'origine d'un intérêt croissant pour la conception de méthodes d'aide à la décision basées sur une connaissance partielle de l'état de la chaîne d'approvisionnement, incluant les incertitudes concernant

le comportement des autres acteurs et des clients finaux.

Contexte scientifique :

Un des problèmes d'optimisation les plus référencés dans la littérature en gestion de chaîne d'approvisionnement consiste à planifier simultanément l'approvisionnement en matières premières ou en composants et les opérations de production. Plus particulièrement, le problème de *dimensionnement de lot* (ou *lot sizing*) et ses nombreux dérivés a été l'objet de nombreuses études, dont un nombre non négligeable propose des modèles qui intègrent des incertitudes. Pendant longtemps, l'essentiel de ces travaux se concentrait sur des approches de *programmation stochastique* visant à optimiser la performance moyenne des décisions et reposant sur des hypothèses restrictives concernant la distribution des paramètres incertains. Cependant, dans un contexte applicatif, il est souvent impossible d'obtenir une caractérisation précise de ces distributions et ces approches peuvent conduire à proposer des solutions aux performances très variables, notamment lors de la réalisation de scénarios peu probables susceptibles d'entraîner des perturbations extrêmement défavorables.

Dans les deux dernières décennies, de nombreux travaux ont donc cherché à proposer des approches d'optimisation permettant d'obtenir des solutions robustes face aux aléas, possiblement sous-optimales mais dont la qualité reste satisfaisante même lorsqu'on s'affranchit des hypothèses restrictives concernant la valeur des paramètres incertains. Ce paradigme est connu sous le terme d'*Optimisation Robuste*. Si elle prend racine dans les travaux de pionniers comme Scarf (1958) ou Soyster (1973), de très nombreux modèles et méthodes de résolution ont été proposés dans la littérature depuis la fin des années 90 et se démarquent notamment de l'approche par optimisation stochastique par l'absence d'hypothèse sur la distribution des paramètres incertains inclus dans les modèles d'optimisation. Pour remplacer la description du comportement aléatoire de ses paramètres, un modèle d'optimisation robuste considère l'ensemble de leurs valeurs possibles, appelé *ensemble d'incertitude*. La définition et la formulation de tels ensembles est au centre de la recherche en optimisation robuste car le choix d'un ensemble influence à la fois la qualité des solutions obtenues et la capacité de méthodes de résolution classique telles que la programmation linéaire pour résoudre ces modèles en pratique.

Dans cette thèse, nous nous intéressons plus spécifiquement à des méthodes d'optimisation robuste dirigées par les données, un domaine de recherche particulièrement actif au cours des deux dernières décennies. Durant cette période, on assiste en effet à une

politique généralisée d'automatisation des opérations de production et de distribution, accompagnée d'une collecte généralisée des données de mesure relatives aux processus industriels. L'intégration de cette information aux outils d'aide à la décision, grâce notamment à des approches statistiques ou d'apprentissage automatisé, permet souvent d'améliorer significativement la qualité des solutions proposées. Dans cette dynamique, l'optimisation robuste s'avère être une approche particulièrement bien adaptée, à même d'intégrer directement ces données dans la caractérisation des ensembles d'incertitude considérés. Cependant, une étude de la littérature montre que l'application de ces approches aux problèmes propres à la gestion de la chaîne d'approvisionnement n'a été que très peu étudiée. En effet, la plupart de ces ensembles d'incertitude mènent à des modèles d'optimisation de grande taille dont la résolution par des approches classiques de programmation linéaire est souvent trop coûteuse en temps et en mémoire, limitant leur champs d'application à des problèmes simples n'intégrant que peu de paramètres incertains.

Les travaux de cette thèse ont donc pour objectif de proposer des approches d'optimisation robuste dirigées par les données pour des problèmes d'approvisionnement et de production. Nous nous intéressons tout particulièrement à la qualité des solutions proposées par nos modèles en comparaison des modèles déterministes (sans incertitudes), de modèles robustes plus classiques et de modèles d'optimisation stochastique. Nos contributions portent notamment sur la formulation des ensembles d'incertitude induits et sur la résolution des modèles qui en découlent.

Structure du manuscrit :

Le manuscrit est composé des 6 chapitres suivants :

Le **Chapitre 1** introduit le concept d'optimisation robuste ainsi que les modèles et les méthodes de résolution les plus utilisés dans la littérature. De nombreux travaux antérieurs à cette thèse ont proposé des ensembles d'incertitude permettant de modéliser plus précisément les incertitudes afin de les intégrer aux modèles d'optimisation robuste. Une revue de la littérature est donc faite pour présenter ces différentes approches et notamment les approches dirigées par les données. Pour introduire plus précisément le contexte applicatif de cette thèse, ce chapitre présente également une revue de littérature des applications de l'optimisation robuste pour des problèmes de gestion de la chaîne d'approvisionnement.

Le **Chapitre 2** présente les différentes contributions de cette thèse et les positionne vis-à-vis de l'état de l'art.

Le **Chapitre 3** présente une méthode d'approximation pour une classe d'ensembles d'incertitude dirigés par les données. Ces ensembles sont obtenus à l'aide d'un algorithme d'apprentissage (connu sous l'appellation *Support Vector Clustering*) et les utiliser dans des modèles d'optimisation robuste permet d'obtenir des solutions de meilleure qualité que les modèles classiques. Cependant, leur formulation complexe augmente de façon significative le temps nécessaire à l'obtention de solutions optimales et limite, de fait, leur champs d'application à des problèmes de petite taille. L'approximation introduite dans ce chapitre réduit ce temps de calcul de façon significative et permet d'obtenir des solutions optimales aux problèmes d'optimisation robuste en un temps raisonnable.

Le **Chapitre 4** est consacré à un cas d'application pratique inspiré du domaine aéronautique. Le problème consiste à planifier les opérations d'approvisionnement et de production d'une ligne d'assemblage dans le but de satisfaire les besoins en différents produits finaux. Les opérations d'approvisionnement sont soumises à des incertitudes émanant de la collaboration avec un prestataire logistique et menant à une diminution de l'efficacité du système de production. Ce chapitre introduit deux modèles d'optimisation robuste permettant de garantir un haut niveau de satisfaction de la demande en produits finis. Le premier modèle utilise une approche classique de la littérature et le second utilise l'algorithme de support vector clustering présenté au chapitre précédent, dans le but d'obtenir des solutions optimales en un temps raisonnable. Les expérimentations montrent l'intérêt des modèles robustes en comparaison à une approche déterministe, ainsi que la supériorité des solutions obtenues avec le modèle dirigé par les données.

Le **Chapitre 5** est consacré à un problème plus général d'approvisionnement face à des incertitudes sur la demande et sur le délai de livraison. Ce chapitre introduit un nouvel ensemble d'incertitude dirigé par les données qui, comparé à un modèle robuste plus classique, permet d'obtenir des solutions de meilleure qualité dans les cas de corrélation étudiés. Cet ensemble d'incertitude diffère des modèles introduits jusqu'alors. En effet, il est basé sur l'hypothèse que la demande suit un processus stochastique bien défini (mais dont les paramètres sont incertains). Nous utilisons ici un modèle ARMA (Autoregressive Moving Average), communément utilisé dans des méthodes prédictives, pour définir les bornes de notre ensemble d'incertitude. Les expérimentations montrent que cette approche permet d'obtenir des solutions moins sensibles aux scénarios de demande extrêmes en comparaison d'une approche d'optimisation stochastique. De plus, comparé à une approche d'optimisation robuste classique de la littérature, cette approche permet d'obtenir de meilleures performances en moyenne et n'induit pas d'augmentation du temps

de résolution des modèles.

Le **Chapitre 6** résume les conclusions de cette thèse et introduit diverses perspectives pour de futurs travaux de recherche.

Les travaux présentés dans cette thèse ont mené à la soumission d'un article scientifique en cours de publication dans une revue internationale ainsi qu'à des présentations dans diverses conférences nationales et internationales. Deux autres articles sont également en cours de rédaction.

ACKNOWLEDGEMENT

Je souhaite remercier chaleureusement mes directeurs Fabien Lehuédé, Guillaume Massonnet et Alexandre Dolgui pour leur soutien tout au long de ces trois années. Je tiens à leur exprimer ma plus profonde gratitude pour la confiance qu'ils m'ont accordée, pour leurs conseils, leur écoute, et plus généralement, pour tous les échanges que nous avons pu avoir (scientifique ou non).

Je remercie Céline Gicquel et André Rossi d'avoir accepté la charge d'être relecteur de cette thèse et pour leurs remarques qui ont abouti à cette version finale du manuscrit. Je remercie également Nadjib Brahimi, Evren Sahin ainsi que Stefan Minner pour leur participation en tant que membre du jury et pour nos échanges qui offrent de nouvelles perspectives pour de futurs travaux. Je remercie tout particulièrement Stefan Minner pour son accueil au sein de l'université technique de Munich et pour nos échanges lors de notre collaboration.

Je remercie également l'ensemble de l'équipe Modelis d'IMT Atlantique de m'avoir offert un cadre de travail convivial tout au long de ma thèse. En partageant avec moi un bureau, un café, une bière ou une cible d'aimant, vous avez su me remotiver lorsque je rencontrais des moments difficiles.

La rédaction de cette thèse n'aurait pas été possible sans le soutien de mes parents, Dominique et Gilles, de mes frères et sœurs, Guillaume et Hélène, et de mes plus proches amis, ils se reconnaîtront. Enfin, je remercie ma copine Mélissa, sans qui rien n'aurait été possible. Elle a été mon plus grand soutien avant et pendant la rédaction de cette thèse, ma plus grande force et ma plus grande fierté.

TABLE OF CONTENTS

Introduction	13
1 Scientific context	14
2 Overview of the thesis and notations	15
1 Literature Review	17
1 Robust Optimization	17
1.1 Solving RO models	20
1.1.1 Dual reformulation with polyhedral uncertainty set	20
1.1.2 Iterative cutting plane methods	22
1.2 Constructing Uncertainty sets	23
1.3 Data-Driven uncertainty sets	25
1.3.1 Gaussian Mixture Models	27
2 Robust inventory management	29
2.1 Lot-sizing models	29
2.2 Static robust inventory and production management	31
2.3 Dynamic robust inventory and production management	33
2 Contributions	37
3 Approximate kernel learning uncertainty set	41
1 Introduction	41
2 Support vector clustering-based uncertainty set	43
3 Approximation of the SVC-based uncertainty set	46
3.1 Initial support vectors selection by K-medoid clustering	47
3.2 Procedure $OPT-\alpha$	48
3.3 Reduction of the number of support vectors	49
3.4 Summary of the approach	50
4 Experimental results	52
4.1 Uncertainty set comparison	52
4.2 Application to optimization problems	54

4.2.1	Robust 0-1 Knapsack Problem	55
4.2.2	Robust Lot-sizing Problem	57
4.2.3	Robust Traveling Salesman Problem	61
5	Conclusion	62
4	A Data-Driven Robust Optimization approach to supply planning	65
1	Introduction	65
2	Problem statement and deterministic formulation	67
3	Robust models for uncertain setup times	69
3.1	Budget-based robust formulation	69
3.2	SVC-based robust formulation	71
4	Experimental results	73
4.1	Scalability of the SVC-based model	74
4.2	Comparing \mathcal{U}^F and \mathcal{U}^{AC}	76
5	Conclusion	78
5	Data-Driven Robust inventory management under uncertain demand and lead time	81
1	Introduction	81
2	Problem and robust solution approach	83
2.1	Deterministic model	84
2.2	Robust optimization model	85
3	Data-driven uncertainty set	86
3.1	Data driven uncertainty set based on ARMA(p,q) process	87
3.1.1	General $ARMA(p, q)$ -based uncertainty set	87
3.1.2	Examples for special cases of demand processes	88
3.2	Joint demand and lead time uncertainty set	89
3.3	A constraint generation approach	91
4	Computational experiments	92
4.1	Sample average approach	93
4.2	Uncertain correlated demand	94
4.3	Uncertain demand-dependent lead times	98
5	Conclusion and perspectives	101
6	Conclusions and Perspectives	103

Conclusion	103
Appendices	109
A Support Vector Clustering and Kernel function	109
Bibliography	111

INTRODUCTION

In today's society, the production and distribution of goods is largely carried out by large companies, which are responsible for satisfying the demand or needs of end customers. In the current context, where international trade is encouraged and the economic competitiveness of companies is a priority, more and more companies are specializing in a single activity, collaborating with other specialized companies or subcontracting some of the tasks they used to carry out. As a result, production and distribution channels have become increasingly complex, incorporating more and more organizations whose collaboration is necessary to ensure that market needs are met. Such networks of companies are often referred to as *Supply Chains (SC)*. Their diversification and expansion over the past decades has motivated numerous research studies and books focused on *Supply Chain Management (SCM)* (Stadtler et al. (2015); Christopher (2005); Dolgui and Proth (2010); Simchi-Levi (2008)), with the objective to describe them in detail, identify their bottlenecks and opportunities of improvement. By searching through the literature on SCM, we can easily see that it regroups a wide variety of problems going from the more strategic (e.g. defining the location of production sites and warehouses, selecting suppliers, etc.) to the more operational considerations (e.g. scheduling of daily production, transmission of orders to a supplier, etc.). This thesis focuses mainly on the operational level, we refer the interested reader to the wide literature on SCM for more detail on the strategic part.

Many definition of what is a supply chain co-exist because of the wide range of organizations that this term aims to describe. Christopher (2005) defines a supply chain as a « ... *network of organizations that are involved, through upstream and downstream linkages, in the different processes and activities that produce value in the form of products or services in the hand of the ultimate customer.* ». In other words, a SC is composed of several organizations who collaborate through material, financial and information exchange.

For instance, think of a production firm that collaborates with a *Logistic Service Provider (LSP)* in charge of the reception and storage of raw materials and their delivery to various locations on the production line. In this simple example, there is an upstream material flow (i.e. the LSP deliver raw material on the production line) and downstream financial and information flows. The financial flow is the price paid by the production

firm for the services provided and the information flow is the transmission of the quantity of raw materials needed on the production line.

One point of agreement over all definitions of SC is that its main objective is the satisfaction of the final customer demand. To achieve this goal, ensuring a good collaboration of all actors of a SC has been identified as a key ingredient to ensure the reliability of material flows. The quality of this collaboration may rely on information sharing, giving access to the state of the SC as a whole to each organisation in order to improve their decision making with respect to a global target, rather than limiting their knowledge to partial information from their direct upstream and downstream collaborators (Barlas and Gunduz (2011)). However, complete collaboration is hard to implement in practice because each company involved in a SC also has to maintain its competitiveness. For example, a manufacturer selling its products to several retailers may share the level of congestion on its production line to inform them about the expected lead time of their orders. Such additional knowledge may in turn lead the retailers to adapt their decisions to reduce the risk of shortage, for instance by anticipating their needs or requesting a special delivery from another producer. While this common knowledge is likely to improve the service level of the whole SC, the manufacturer may perceive it as an additional risk of losing contracts, making it harder to prioritize orders coming from a particular retailer. Such issues have lead SC managers to only partially disclose the information on their own activity to other members, raising the interest for decision support tools that are able to base their optimization procedures on a partial knowledge about the state of the SC, including uncertainties on the behavior of other actors or the end customers.

1 Scientific context

In most applications, SCM involves solving production planning problems considering simultaneously the production, procurement and distribution decisions (Pochet and Wolsey (2006)). Among the classical SCM research topics, *Lot Sizing Problems (LSP)* have been the focal point of a large part of the literature, counting numerous contributions since the pioneering works of Harris (1913) and Wagner and Whitin (1958). A rising challenge that recent works have attempted to address is related to problems in which uncertainty occurs, with the objective to mitigate the impact of uncertainties on the reliability of production plan and/or production costs.

In this context, Stochastic Programming has been one of the most popular approaches

and studies have considered a wide variety of problems and sources of uncertainties (see e.g. Brahimi et al. (2017); Metzker Soares (2022)). In order to integrate uncertainties in the LSP, stochastic programming approaches often rely on strong assumptions on the distributions of uncertain parameters. In addition, the objective of such models is to minimize the expected cost of solutions thus taking the risk to face extreme costs or to produce unreliable production plans in extreme cases of uncertain parameters. Thus, in order to reduce the risk of bad performances in extreme conditions, many decision makers seek for a "robust" solution (i.e. a solution that performs reasonably well for all realizations of uncertain parameters).

While the last two decades have witnessed the rapid development of *Robust Optimization (RO)* models for various problems, its introduction for SCM applications dates back to the work of Scarf (1958); Gallego (1992); Moon and Gallego (1994). RO models have the advantage that the approach does not require specific assumptions on the distribution of uncertain parameters, since they only consider that some of the problem parameters can take *any* value within a bounded set, often called *uncertainty set*. A solution is then said to be *robust* if it remains feasible for all values of uncertain parameters within this set.

In parallel, the volume of data collected by the actors of supply chains grows continuously, giving access to large databases of historical values of these uncertain parameters (e.g. demand, lead time, picking time) to improve the decision making. It turns out that many RO methods offer a suitable framework to incorporate directly these historical data in the formulation of optimization models (Bertsimas et al., 2017), which has led to a new research stream referred to as *Data Driven Robust Optimization (DDRO)*.

2 Overview of the thesis and notations

In this thesis, we study problems related to SCM in an uncertain context through the lens of data-driven robust optimization. Chapter 1 introduces the concept of robust optimization, giving the most common models and solution methods. It presents a structured literature review on the different approaches to model the uncertainty, going from models using basic statistical information to data-driven models built with machine learning algorithms. The second part of our survey is dedicated to the application of RO in production planning and inventory management. It gives an overview of the different problems that have been studied under the RO framework, the most common models and optimization

algorithms used in this area. In Chapter 2, we summarize the contributions presented in detail in Chapter 3, 4 and 5 of this thesis. Finally, we conclude this manuscript by a summary of the contributions developed in these chapters, giving some managerial insights and describing promising perspectives for future research.

Notations: In the rest of the manuscript we use the following notations in order to simplify mathematical expressions. All scalar parameters or variables are denoted by lower case letters such as a_{ij} . Bold lower case letters such as \mathbf{x} denote vectors of variables or parameters and bold uppercase letters are used to denote matrices such as \mathbf{A} whose j -th row is denoted by \mathbf{a}_j . We denote by $[m]$ the set of integer values going from 1 to m and $\|\cdot\|_p$ denote the l_p -norm for all $p \in \mathbb{N}_+ \cup \{\infty\}$. Sets are denoted by calligraphic upper case letters such as \mathcal{D} , \mathcal{S} or \mathcal{B} .

LITERATURE REVIEW

In this chapter, we introduce the concept of Robust Optimization and review the existing literature related to uncertainty set construction, Data Driven Robust Optimization and Robust Supply Chain Management.

1 Robust Optimization

Since the end of 1990's, *Robust Optimization (RO)* has become a popular approach to solve optimization problems involving uncertain parameters. The key aspect of RO is that it does not rely on distributional assumptions of uncertain parameters unlike other standard approaches such as *Stochastic Optimization (SO)* or *Stochastic Dynamic Programming (SDP)*. The central concept of RO is to consider that some of the problem parameters \mathbf{a}_j can take *any* value within a bounded set \mathcal{U}_j , often called *uncertainty set*, instead of assuming a fixed one. A solution is then said to be *robust* if it remains feasible for all possible values $\mathbf{a}_j \in \mathcal{U}_j$. In this section, we consider the general form of robust linear optimization problems given by:

$$\begin{aligned}
 & \max \quad \mathbf{c}^T \mathbf{x} \\
 & \text{s.t.} \quad \mathbf{a}_j^T \mathbf{x} \leq b_j \quad \forall j \in [m], \quad \forall \mathbf{a}_j \in \mathcal{U}_j \\
 & \quad \quad \mathbf{x} \in \mathbb{R}_+^n
 \end{aligned} \tag{1.1}$$

where \mathbf{x} is an n -dimensional decision variable vector, \mathbf{c} a vector of cost parameters, $\mathbf{a}_j \in \mathbb{R}^n$ are vectors of uncertain parameters of the problem for all $j \in [m]$.

The concept of (RO) was first introduced by Soyster (1973) who defines a framework in which each uncertain parameter a_{ij} belongs to the interval $[\bar{a}_{ij} - \hat{a}_{ij}, \bar{a}_{ij} + \hat{a}_{ij}]$ where \bar{a}_{ij} is called the *nominal value* and \hat{a}_{ij} represents the *maximum deviation* of a_{ij} from \bar{a}_{ij} .

Introducing additional notations $z_{ij} = \frac{a_{ij} - \bar{a}_{ij}}{\hat{a}_{ij}}$. The resulting uncertainty set is defined as:

$$\mathcal{U}_j^{\text{Box}} = \{a_{ij} = \bar{a}_{ij} + z_{ij}\hat{a}_{ij} \mid \|\mathbf{z}_j\|_\infty \leq 1\} \quad (1.2)$$

. From the definitions (1.1) and (1.2), it follows that any solution using Soyster's assumptions is protected against *all* possible values of the uncertain parameters within \mathcal{U}^{Box} . While this property may be desirable in some very specific settings, it turns out that when parameters a_{ij} are in fact random variables with support $[\bar{a}_{ij} - \hat{a}_{ij}, \bar{a}_{ij} + \hat{a}_{ij}]$, this approach often becomes overly conservative. It is well established that in most applications, it is highly unlikely that all parameters simultaneously take their extreme values, and the probability that this situation occurs decreases as the number of random parameters increases. In such cases, the solutions obtained using Soyster's modeling framework guarantee their robustness against unlikely realizations of the problem parameters, often at the expense of a significant performance reduction. Such solutions are called *over-conservative*. This issue has been a major topic for researchers over the last three decades, pursuing the right trade-off between the average performance and the protection level against bad parameters realizations.

The first attempt to reduce over-conservatism was proposed simultaneously by Ben-Tal and Nemirovski (1998) and El Ghaoui et al. (1998) who studied robust convex and semidefinite programs and their tractability. They start from the fact that, using the l_2 -norm, the cumulated deviation of uncertain parameters from their nominal values evolve as an ellipse in the parameter space. Thus, they introduce an ellipsoidal uncertainty set whose radius is controlled by a parameter Ω_j . The use of this parameter provide an natural interpretation to control the conservatism of RO models as increasing the value of Ω_j would be used to ensure a larger constraint satisfaction level but will often result in a reduction of performance. Under such uncertainty set, they propose to solve the robust problem (1.1) by replacing it by an equivalent and tractable deterministic reformulation (the *Robust Counterpart (RC)*). Among other results, they prove that the robust counterpart is a *Second Order Cone Program (SOCP)* when the initial problem is a *Linear Program (LP)*. Later on, Ben-Tal and Nemirovski (2000) extended this work and provide an upper bound on the constraint violation level for robust LP under such uncertainty sets when parameters are independent and symmetrically distributed random variables. The ellipsoidal uncertainty set is given by:

$$\mathcal{U}_j^\Omega = \left\{ a_{ij} = \bar{a}_{ij} + z_{ij}\hat{a}_{ij} \mid \|\mathbf{z}_j\|_2 \leq \Omega_j \right\} \quad (1.3)$$

Following their works, Bertsimas and Sim (2003, 2004) introduce a polyhedral uncertainty set whose size is also controlled by a parameter $\Gamma_j \in \{0, \dots, n\}$ called *budget* of uncertainty. This uncertainty set, often called budget-based uncertainty set is given by the following expression:

$$\mathcal{U}_j^\Gamma = \{a_{ij} = \bar{a}_{ij} + z_{ij}\hat{a}_{ij} \mid \|z_j\|_1 \leq \Gamma_j\} \quad (1.4)$$

As in the ellipsoidal case, Γ_j provides a natural interpretation to control the robustness of robust models as Γ_j represents the maximum number of parameters which can deviate from their nominal value \bar{a}_{ij} in the i -th constraint. In other words, if $\Gamma_j = 0$, all uncertain parameters will take their nominal value. On the contrary, if $\Gamma_j = n$ they will all take their worst value. However in this case, the use of $\|\cdot\|_1$ ensures that the robust counterpart of a LP remains linear, thus providing a significant computational advantage compared to the ellipsoidal uncertainty set (1.3). With the help of numerical experiments on a binary knapsack problem and a linear portfolio management problem, they prove that their approach reduces the conservatism of robust solutions (compared to \mathcal{U}^{Box}) while allowing large *Mixed Integer Linear Problems (MILP)* to be solved efficiently.

As an example, Table 1.1 show the performance reduction and the probability of constraint violation obtained for the binary knapsack problem under uncertain weights of Bertsimas and Sim (2003, 2004). We compare the solutions obtained by the deterministic model and the robust models using \mathcal{U}^{Box} and \mathcal{U}^Γ and we evaluate the empirical probability of constraint violation on 10000 random scenarios. On this example, we observe that using Soyster's method provide 100% of constraint satisfaction but reduce the objective value by 5.19% in comparison with the deterministic solution. On the other hand, using the approach of Bertsimas and Sim (2003, 2004) with different budget provide trade-off solutions and when $\Gamma = 20.0$ the risk of constraint violation is only 0.05% and the objective value is reduced by only 0.75%.

These three uncertainty sets are at the root of much of the work on Robust Optimization that has followed in the last two decades. Their findings introduces the central challenges of robust optimization research: producing high quality solutions while maintaining computational tractability.

Table 1.1 – Performance reduction and risk of constraint violation for the binary knapsack problem of Bertsimas and Sim (2003, 2004) under box and budget based uncertainty sets.

Method	Objective value	Performance reduction	Constraint violation
Deterministic	8135	0.0%	49.79%
$\Gamma = 0$	8135	0.0%	49.79%
$\Gamma = 5$	8123	0.14%	25.28%
$\Gamma = 10$	8105	0.35%	6.68%
$\Gamma = 15$	8085	0.61%	0.74%
$\Gamma = 20$	8074	0.75%	0.05%
\mathcal{U}^{Box}	7712	5.19%	0.0%

1.1 Solving RO models

Since RO problems as (1.1) involves an infinite number of constraints when parameters a_{ij} are continuous, using a solver directly to compute an optimal solution is inconceivable. In this section, we quickly review the most common methods encountered in the literature to tackle RO models.

1.1.1 Dual reformulation with polyhedral uncertainty set

The most common approach in the literature to solve RO models is to formulate a deterministic and tractable RC using duality theory. As an example, we describe how to derive a linear robust counterpart from the robust problem (1.1) when the uncertainty set is the budget-based uncertainty set (1.4) of Bertsimas and Sim (2003, 2004). Let us consider the following problem:

$$\begin{aligned}
 & \max \mathbf{c}^T \mathbf{x} \\
 & \text{s.t. } \max_{\mathbf{a}_j \in \mathcal{U}_j} \mathbf{a}_j^T \mathbf{x} \leq b_j, \quad \forall j \in [m] \\
 & \mathbf{x} \in \mathbb{R}_+^n
 \end{aligned} \tag{1.5}$$

This problem is equivalent to (1.1) as any feasible solution \mathbf{x} satisfying constraint $\max_{\mathbf{a}_j \in \mathcal{U}_j} \mathbf{a}_j^T \mathbf{x} \leq b_j$ remain feasible for all realization $\tilde{\mathbf{a}}_j \in \mathcal{U}_j$ such that $\tilde{\mathbf{a}}_j^T \mathbf{x} \leq \max_{\mathbf{a}_j \in \mathcal{U}_j} \mathbf{a}_j^T \mathbf{x}$.

Under the budget-based uncertainty set, this problem can be formulated as:

$$\begin{aligned}
& \max \mathbf{c}^T \mathbf{x} \\
& \text{s.t. } \sum_{i=1}^n \bar{a}_{ij} x_{ij} + \max_{\mathbf{z}_j \in [-1,1]^n, \|\mathbf{z}_j\|_1 \leq \Gamma_j} \left\{ \sum_{i=1}^n \hat{a}_{ij} x_{ij} z_{ij} \right\} \leq b_j, \forall j \in [m] \\
& \mathbf{x} \in \mathbb{R}_+^n
\end{aligned} \tag{1.6}$$

Considering an optimal solution \mathbf{x}^* to this problem, for each $j \in [m]$, the inner maximization problem $\max_{\mathbf{z}_j \in [-1,1]^n, \|\mathbf{z}_j\|_1 \leq \Gamma_j} \left\{ \sum_{i=1}^n \hat{a}_{ij} x_{ij}^* z_{ij} \right\}$ is equivalent to a linear optimization problem. As values in \mathbf{x} are all positive or null, we can restrict $\mathbf{z}_j \in [0, 1]^n$ and thus replace the l_1 -norm $\|\mathbf{z}_j\|_1$ by a simple sum $\sum_{i=1}^n z_{ij}$. This gives us the following equivalent LP:

$$\begin{aligned}
& \max \sum_{i=1}^n \hat{a}_{ij} x_{ij}^* z_{ij} \\
& \text{s.t. } \sum_{i=1}^n z_{ij} \leq \Gamma_j \\
& 0 \leq z_{ij} \leq 1, \forall i \in [n].
\end{aligned} \tag{1.7}$$

Associating the dual variables q_j and r_{ij} , $i \in [n]$ to the constraints. We now consider the dual of problem (1.7)

$$\begin{aligned}
& \min q_j \Gamma_j + \sum_{i=1}^n r_{ij} \\
& \text{s.t. } q_j + r_{ij} \geq \hat{a}_{ij} x_{ij}^* \quad \forall i \in [n] \\
& q_j \geq 0 \\
& r_{ij} \geq 0 \quad \forall i \in [n]
\end{aligned} \tag{1.8}$$

Since the primal problem (1.7) is feasible and bounded, its dual problem (1.8) is also feasible and bounded and their optimal values coincide. Therefore, we can replace the j -th inner maximisation problem of (1.6) by the objective function of (1.8), with the

associated constraints. This gives us the (linear) robust counterpart of problem (1.1):

$$\begin{aligned}
& \max \mathbf{c}^T \mathbf{x} \\
& \text{s.t. } \sum_{i=1}^n \bar{a}_{ij} x_i + q_j \Gamma_j + \sum_{i=1}^n r_{ij} \leq b_j & \forall j \in [m] \\
& q_j + r_{ij} \geq \hat{a}_{ij} x_i & \forall i \in [n], j \in [m] \\
& x_i \geq 0 & \forall i \in [n] \\
& r_{ij} \geq 0 & \forall i \in [n], j \in [m] \\
& q_j \geq 0 & \forall j \in [m]
\end{aligned} \tag{1.9}$$

which can be solved by linear programming solvers. Note that a similar approach can be applied to obtain the RC of a robust LP under any polyhedral uncertainty set. For more details on the formulation of the RC under ellipsoidal uncertainty sets, interested readers can refer to Ben-Tal and Nemirovski (1998); El Ghaoui et al. (1998); Ben-Tal and Nemirovski (2000).

1.1.2 Iterative cutting plane methods

As an alternative approach, Mutapcic and Boyd (2009) have proposed to solve RO models using an iterative cutting plane algorithm. Their method rely on solving a relaxed RO problem of the form:

$$\begin{aligned}
& \max \mathbf{c}^T \mathbf{x} \\
& \text{s.t. } \mathbf{a}_j^T \mathbf{x} \leq b_j, \mathbf{a}_j \in \Omega_j, \forall j \in [m] \\
& \mathbf{x} \in \mathbb{R}^n
\end{aligned} \tag{1.10}$$

where Ω_j is a subset of \mathcal{U}_j . Given an optimal solution \mathbf{x}^* obtained after k iterations, the algorithm solves the inner optimization problem to find the worst-case parameter values $\mathbf{a}_j^{(k)}$. If the j -th constraint is violated, a new constraint is added to the formulation (1.10). Then, if no more violating realization is found, the current solution \mathbf{x}^* is returned. If more than one constraint is subject to uncertainties the algorithm can solve one inner problem per constraint (Bertsimas et al. (2016)). If several constraints contain the same uncertain parameters, one could formulate a single inner problem corresponding to all the constraints simultaneously (Bienstock and Ozbay (2008); Thorsen and Yao (2017)).

Such approaches have been proposed to solve robust linear problems, convex prob-

lems and quadratic problems (Ben-Tal et al. (2015)). The convergence of such algorithm have been studied under ellipsoidal uncertainty set (Mínguez and Casero-Alonso (2019)) and under general polyhedral uncertainty set (Bertsimas et al. (2016)). Several papers compared the performances of such methods with a reformulation approach with ellipsoidal and budget-based uncertainty sets (Fischetti and Monaci (2012); Bertsimas et al. (2016)). To date, the latest and more complete comparison (Bertsimas et al. (2016)) claim that there is no significant evidence that one approach outperforms the other. In fact, the method to be preferred seem to depend on the considered uncertainty set (ellipsoidal, budget-based) and on the class of problem (LP, MILP). The systematic use of reformulation to solve robust problems should therefore be questioned, particularly when introducing new families of uncertainty sets. Based on the same idea, some papers also propose dedicated iterative algorithms to solve particular problems (Goerigk et al. (2015); Bienstock and Ozbay (2008); Montemanni (2006)) or study the use of approximate solution methods to solve the relaxed robust problem or to identify violating realizations (Pätzold and Schöbel (2020); Thevenin et al. (2022)).

1.2 Constructing Uncertainty sets

Uncertainties on parameters of optimization problems arise for various reasons. Parameters can model random events such as the demand of final customers of a SC or correspond to elements for which an exact measurement is not possible, leading to errors. However, in most models all types of uncertain parameters are considered as random variables and this will be the case in the rest of this manuscript. As mentioned at the beginning of this chapter, the construction of uncertainty sets is at the center of researchers considerations because of its significant role in both the quality of the obtained solutions and the computational tractability of robust models. The practical efficiency of classical uncertainty sets such as $\mathcal{U}_j^{\text{Box}}$, \mathcal{U}_j^Γ and \mathcal{U}_j^Ω have been demonstrated on numerous applications (Aharon et al. (2009); Gabrel et al. (2014); Sözüer and Thiele (2016)). However, they systematically define the set of possible values for each parameter as an interval centred on a nominal value with symmetrical and independent variations. Thus, by ignoring potential asymmetries or correlations in the distributions of uncertain parameters, using these uncertainty sets can lead to over-conservative or unreliable decisions in practice. Moreover, the probability bounds on the constraint satisfaction obtained by using such uncertainty sets are reliable if and only if uncertain parameters are indepently and symetrically distributed. For this reason, numerous papers have attempted to design alternative

uncertainty sets that are more suitable to model a wider range of behaviours for uncertain parameters.

Pachamanova (2002) and Bertsimas et al. (2004) generalize the concept of *norm-based* uncertainty set to any arbitrary norm as:

$$\mathcal{U}_j = \{\mathbf{a}_j \mid \|\mathbf{M}(\mathbf{a}_j - \bar{\mathbf{a}}_j)\| \leq \Delta_j\} \quad (1.11)$$

where $\|\cdot\|$ can be any given norm, the budget Δ_j bounds the cumulative deviation of uncertain parameters from their nominal value $\bar{\mathbf{a}}_j$ and the matrix M is an invertible matrix used to integrate statistical information in the uncertainty set. As an example, considering the l_2 -norm, if M is the identity matrix I_n then \mathcal{U}_j is a sphere. If M is a diagonal matrix whose entries are standard deviations of uncertain parameters then \mathcal{U}_j is the classic ellipsoidal uncertainty set. And if $M = \Sigma^{-1/2}$ with Σ the covariance matrix then \mathcal{U}_j is an ellipsoid stretched in the direction of greatest dispersion. This last matrix is a way to integrate correlation information in the construction of uncertainty sets. Bertsimas et al. (2004) also prove that the robust counterpart of any problem under a norm based uncertainty set can be obtained from its dual norm. They derive probability bounds on constraint violation for the general norm-base which coincides with existing results from El Ghaoui et al. (1998) and Bertsimas and Sim (2003, 2004).

The concept of general norm based uncertainty set was used later on by Chen et al. (2007) to propose an asymmetric uncertainty set based on backward and forward deviation measures. Their approach is to define $\mathbf{a}_j = \bar{\mathbf{a}}_j + (\hat{\mathbf{a}}_j^+ - \hat{\mathbf{a}}_j^-)$ with $\hat{\mathbf{a}}_j^+, \hat{\mathbf{a}}_j^- \in \mathbb{R}_+^n$ and to define an asymmetric uncertainty set:

$$\mathcal{U}_j = \left\{ \mathbf{a}_j \mid \mathbf{a}_j = \bar{\mathbf{a}}_j + (\hat{\mathbf{a}}_j^+ - \hat{\mathbf{a}}_j^-), \|\mathbf{P}\mathbf{a}_j^+ + \mathbf{Q}\mathbf{a}_j^-\| \leq \Delta_j \right\} \quad (1.12)$$

where \mathbf{P} and \mathbf{Q} are diagonal matrices built from forward and backward deviation measures.

More recently Jalilvand-Nejad et al. (2016) proposed an alternative way to integrate correlation information in the construction of uncertainty sets. They propose a polyhedral formulation based on the correlation matrix of uncertain parameters. They use experimental results on a robust job shop scheduling problem with uncertain processing times to show that their approach outperform Pachamanova (2002) under normal marginal distributions. Daneshvari and Shafaei (2021) develop an alternative formulation to integrate correlation information and improve the results on the same set of experiments.

As an other approach to reduce conservatism, Bandi and Bertsimas (2012) propose to use the central limit theorem to tune the so-called budget of uncertainty. Assuming that uncertain parameters are independent and identically distributed the corresponding set is:

$$\mathcal{U}_j = \left\{ a_{ij} = \bar{a}_{ij} + z_{ij}\sigma \mid \left| \sum_{i=1}^n z_{ij} \right| \leq \Gamma_j \sqrt{n} \right\} \quad (1.13)$$

where σ is the standard deviation of uncertain parameters.

Finally, Poss (2013) developed an other way to deal with over-conservatism. He uses the budget-based approach of Bertsimas and Sim (2003, 2004) in which the budget Γ_j is a function of the decision variables and defines an uncertainty set

$$\mathcal{U}_j^F = \{ a_{ij} = \bar{a}_{ij} + z_{ij}\hat{a}_{ij} \mid \|z_j\|_1 \leq \gamma_j(\mathbf{x}) \} \quad (1.14)$$

He derives a linear robust counterpart when the function $\gamma_j(\mathbf{x})$ is a linear or piecewise linear increasing function and proves that the above uncertainty set reduces the conservatism in problem with few non-zero variables. His approach can be applied for problems with binary and bounded integer or continuous variables. He illustrates the performance of his approach over the classic budget-based approach on the Robust 0-1 Knapsack problem and its fractional version.

The variations of uncertainty sets presented above offer promising ways to reduce the conservatism of robust models in particular cases under specific assumptions. However, they either rely on strong assumptions on the distributions of uncertain parameters or necessitate the tuning of parameters (nominal values, maximum deviations, budget parameters) that strongly influences the efficiency of the resulting robust models. Making robust optimization a tool of interest for practitioners thus requires uncertainty sets that are both as flexible as the asymmetric ones of Chen et al. (2007) and parameterized with simple and intuitive parameters such as the budget of uncertainty introduced in Bertsimas and Sim (2003, 2004).

1.3 Data-Driven uncertainty sets

Up to now, we illustrated the fact that in order to reduce conservatism, the traditional approach to include uncertain parameters in robust optimization models relied on offline procedures and/or assumptions to determine their range of admissible values or distributional characteristics. In this classical framework, upfront data analysis is sometimes

employed to define some input parameters to the problem under study. However, with the emergence of large data sets collected in a variety of applications, the last two decades have led researchers and decision makers to challenge this approach. A new trend has emerged, which allows to directly incorporate the data into a mathematical formulation, often through a combination of statistical analysis and/or dedicated algorithms, and tailor the extracted knowledge to the problem under study. Within these developments, RO has emerged as a suitable approach that enables the practitioners to both integrate raw information into optimization models while retaining good generalization properties, see e.g., Bertsimas et al. (2017). Such models are referred to as *Data-Driven Robust Optimization (DDRO)* models in the literature.

In the research literature on DDRO, most efforts focus on leveraging various statistical and machine learning approaches on historical data to build uncertainty sets that embed structural properties of the model parameters. One of the major goal of DDRO is to combine precision in the description of the uncertain parameters and simplicity of use for practitioners. In this section we review the main existing approaches from the literature to obtain exploitable data-driven uncertainty sets.

1.3.0.1 Statistical Hypothesis Test In their well-known paper Bertsimas et al. (2017) use statistical hypothesis-test (SHT) to build uncertainty sets with different assumptions on the distributions (i.e. known finite support, continuous with independent marginals, none). Each of their uncertainty sets corresponds to a single tractable robust counterpart formulation (LP, SOCP) and is related to a given bound on the probability of constraint violation (valid under independent marginals). They illustrate the use of their approach on small problems from portfolio management and queuing analysis, showing that their uncertainty sets can provide constraint satisfaction guarantees equivalent to classical models but are less conservative. To the best of our knowledge, their approach have never been used in other studies or on other application cases.

1.3.0.2 Dimensionality reduction Ning and You (2018a,b) rely on Principal Component Analysis (PCA) and Kernel Density Estimation (KDE) to build uncertainty sets that directly integrate distributional information (correlation, asymmetry) available in the historical data. They use PCA to identify correlation and KDE (as well as Robust KDE) to define upper and lower bounds in each dimension of the principle components. They illustrate the efficiency of this approach on different applications in robust model

predictive control, production batch scheduling and chemical process network design. In the same framework, Zhang et al. (2022) use coupled Partial Least Square (PLS) or Kernel PCA (KPCA) with Robust Kernel Density Estimation (RKDE) to construct uncertainty sets.

Using dimensionality reduction and distribution estimation appears as a promising approach to build uncertainty sets from historical data. However, PCA and similar techniques must be used carefully as one has to set the number of principle components to efficiently aggregate the information extracted from the data. Moreover, PCA must be applied on standardized data and ignore the difference in scale of uncertain parameters which restricts its range of application to build uncertainty sets.

1.3.1 Gaussian Mixture Models

Campbell and How (2015) use the Dirichlet Process Mixture Model (DPMM) to define uncertainty sets. This learning algorithm use historical data to build a Gaussian Mixture Model that describe the distribution of uncertain parameters. The final uncertainty set is the union of ellipsoidal uncertainty sets in which each component of the mixture model is encapsulated in an ellipsoid. This approach leads to a SOCP robust counterpart and have been applied by Zhao and You (2019) on a case study in energy production. Their result shows that using the DPMM based uncertainty set succeed in reducing operating costs against classical norm based uncertainty sets when distributions are multi-modal and correlated and when the number of uncertain parameter is small (8). Ning and You (2017) propose an alternative uncertainty set using DPMM. Their uncertainty set is the union of several polyhedral (budget-based) uncertainty sets in which each component of the mixture model is encapsulated in a bounded and convex polytope. They solve the resulting robust optimization models using a cutting plane approach such as presented in Section 1.1 in which several problems (one problem per component) are solved at each iteration to identify violating uncertain realizations.

1.3.1.1 Support Vector Clustering Shang et al. (2017) introduced an uncertainty set based on a one class *Support Vector Clustering (SVC)* algorithm. SVC is an unsupervised version of the classical (One-Class) *Support Vector Machines (SVM)*, a classification machine learning algorithm that aims at defining the best boundary in order to partition a space into two classes, based on a given set of (labeled) data points. Schölkopf et al. (2001) propose to use SVC to describe the support of the distribution of (multi-dimensional)

random variables by defining cluster that contains a given portion of data. Shang et al. (2017) build upon their approach and derive linear constraints to describe this cluster as a bounded and convex polytope. When applied to the uncertain parameters of an optimization problem, their method enables them to obtain new data-driven uncertainty sets suitable for RO. The resulting uncertainty set efficiently captures correlations and asymmetries in the distribution of uncertain parameters and the robust counterpart is tractable (potentially large scale) when the initial problem is linear.

In the recent years, the SVC-based uncertainty set has attracted the interest of researchers for various applications such as energy system optimization Shen et al. (2020), resource allocation in a cellular network Wu et al. (2022), supply chain design Mohseni and Pishvae (2020), or model predictive control Shang and You (2019); Shang et al. (2020). Several papers also extend the work of Shang et al. (2017) or propose alternative methods to build uncertainty sets based on support vectors. Mohseni and Pishvae (2020) extend the SVC-based uncertainty set with a Conservative Support Vector Clustering algorithm. Their approach provide higher satisfaction guarantees when only a small data-set is available. In the same objective, Shang and You (2018) propose to reduce the dimension of the uncertainty with PCA before using the SVC-based uncertainty set. Han et al. (2021) and Lin et al. (2022) use Multiple Kernel Learning to improve the representation accuracy of the support of uncertain parameters. The resulting uncertainty sets are similar in their formulation to the classic SVC-based one but requires much more variables and constraints, thus leading to very large robust counterparts and limiting the range of application to problems with only a small number of uncertain parameters. Finally, Goerigk and Kurtz (2023) use One-Class Deep Support Vector Data Description to formulate a SVC-based uncertainty set. This approach leads to non convex uncertainty sets which can be used to model multimodal distributions more efficiently. However, the learning process is extremely demanding and the resulting problem needs to be solved with a decomposition approach (to tackle the non-convexity).

1.3.1.2 Sampling and Cutting Plane generation An other common approach in DDRO is to consider the construction of uncertainty sets as a cutting plane generation problem whose objective is to eliminate parts of the uncertain parameters space. In that way, (Zhang et al. (2018b,a)) rely on copula theory to estimate the joint probability density function (pdf) of uncertain parameters, then use Monte-Carlo sampling and define cutting planes that eliminate a given portion of the samples. Zhang et al. (2018b) also compare

this approach with uncertainty sets based on the convex hull of the sampled realization. All the uncertainty sets introduced in these papers are polyhedral uncertainty sets but introduce a large number of cutting planes and can only be used to solve optimization problems with a small (i.e. less than 10) number of uncertain parameters.

As an alternative to cutting plane generation, (Gumte et al. (2021b,a)) use Fuzzy C-means clustering to define clusters in historical data and uniform sampling in these clusters to define a discrete and non-convex uncertainty set.

2 Robust inventory management

This thesis focuses on the management of inventory in production systems subject to uncertainties. In industrial systems, production decisions (e.g. quantity, scheduling) are made in order to satisfy final customers demand and inventory decisions (e.g. placing orders, quantity) are made to ensure the feasibility of the production plan.

It is well known that industrial systems are subject to various uncertainties (e.g. demand, yield, processing times, etc) that affect their efficiency at each step of the production process. In this context, robust optimization have been proposed as a method to reduce the impact of uncertainties on production and inventory management decisions.

Due to the dynamic aspect of inventory management, most of the robust optimization literature in this area have focused on multi-period problems. In this section we review the main existing models and techniques used in the robust inventory management and lot-sizing literature. To make this manuscript self supporting, we start with a brief introduction to lot-sizing models in Section 2.1. We then present separately static and dynamic approaches. The former, introduced in Section 2.2 corresponds to optimization problem under uncertainty in which all the decisions are made prior to the uncertainty realization (here and now decisions). Such problems can be formulated as the robust problem (1.1). On the other hand, some of the decisions in dynamic formulations are *adjustable*, that is they can be made once a portion of the uncertain parameters are realized. Section 2.3 presents applications of this principle to robust lot-sizing models.

2.1 Lot-sizing models

The Lot Sizing Problem (LSP) has been introduced by Wagner and Whitin (1958) as a way to plan the production of goods (when and in which quantity) in order to

meet customers demand over a planning horizon in a deterministic context. Since then, LSP models have been studied and extended to various types of production systems (multi-echelon, multi-item, capacitated) and uncertainties (demand, lead-time, yield). Since the last two decades, several papers proposed literature review on different aspects of LSP models. Brahimy et al. (2006, 2017) provide a survey of the single-item LSP model. Karimi and Fatemi (2003); Gicquel et al. (2008) propose reviews of capacitated lot sizing models. Stadtler et al. (2015) propose a review of LSP models applied to supply chain management and more recently, Metzker Soares et al. (2022) propose a review of the different solution methods for LSP under uncertainty, including Stochastic Programming, RO and Distributionally Robust Optimization. We refer interested readers to these references for more information on LSP models.

The static and deterministic version of the single-item uncapacitated lot sizing problem aims at defining the setups (when to produce/order) and lot sizes (quantity produced/ordered) in a finite planning horizon of T periods in order to satisfy the customers demands at minimum costs. The basic costs incurred consist in a setup cost s paid once for each order, linear ordering costs v paid for each unit ordered and a linear holding cost paid for each unit held in inventory at the end of a given period. We can formulate this problem with the following mathematical model:

$$\min \sum_{t=1}^T (sx_t + vq_t + hI_t) \quad (1.15)$$

$$\text{s.t. } I_t = \sum_{k=1}^t q_k - d_k \quad \forall t \in [T] \quad (1.16)$$

$$0 \leq q_t \leq Mx_t \quad \forall t \in [T] \quad (1.17)$$

$$x_t \in \{0, 1\} \quad \forall t \in [T] \quad (1.18)$$

$$I_t \geq 0 \quad \forall t \in [T] \quad (1.19)$$

$$(1.20)$$

in which for all $t \in [T]$, x_t is a binary setup variable equal to 1 if production/ordering occurs during period t , q_t is the lot-size decision in period t , I_t is the inventory level at the end of period t and M is a sufficiently large value (e.g. $\sum_{t=1}^T d_t$).

In order to model the large variety of supply chains and production systems, this basic model has been extended with various features in the literature. We presented the *single-echelon* model in which a product is ordered or produced in order to satisfy

the customer demand with no intermediate sub-deliveries/assemblies. There also exists a variety of *multi-echelon* models, in which several steps are necessary before meeting the final customer demand. The same distinction is made for the number of item type (*single/multi-item*).

Many robust inventory models in the literature include the possibility of a demand shortage, where unmet demand can be either lost or delayed for a penalty cost. Specifically, *lost sales* corresponds to a case in which the unmet demand is penalized but can not be satisfied in future periods. On the other hand when *backordering* or *backlogging* is allowed, demand satisfaction may be delayed when the inventory is empty to future periods one new units of product have been produced or received. Notice that some papers include both type of penalties in a mixed approach to reflect practical applications (Nabil Absi and Dauzère-Pérès (2011)). In some papers, both types of penalties are considered in a mixed approach to reflect practical applications.

In practice, resources are often constrained by a capacity (e.g. maximum supplier production, available production time, etc). When such constraints are applied, we talk about *capacitated* models while *uncapacitated* models refer to problems with unconstrained resources.

Finally, a wide variety of uncertainties can occur in supply chain and in production systems. According to multiple review (e.g. Metzker Soares (2022)) the most studied source of uncertainty is the demand. Other sources of uncertainty have also been studied such as cost uncertainties (e.g. when production or ordering costs are uncertain) (Quezada et al. (2019)), lead time uncertainties (Huang and Küçüyavuz (2008)), yield uncertainties (Metzker Soares et al. (2022)) and production and setup times uncertainties (Taş et al. (2019)).

2.2 Static robust inventory and production management

In this section, we provide an overview of existing works on robust inventory management static decision context, i.e. when decisions are made before the uncertainty is realized. The first application of Robust Optimization in inventory management is Bertsimas and Thiele (2004, 2006). The authors study a single-echelon, single-item LSP with uncertain demand, as well as extensions to multi-echelon systems, with and without capacity on the production quantity or the inventory. They use a budget-based robust model and derive a robust counterpart following the methodology introduced in Bertsimas and Sim (2003, 2004). As a by-product, they propose a closed form policy in the case of a

single echelon problem without setup costs or with known ordering periods. Compared to a dynamic programming (DP) solution, RO outperforms DP (better average and better standard deviation) when the ratio backlog/holding is high enough and DP outperforms RO in other cases. In the case of multi-echelon systems, RO outperform a myopic policy (when DP is computationally intractable).

As a matter of example, the single-echelon single-item robust lot-sizing model of Bertsimas and Thiele (2006) is given by:

$$\min \sum_{t=1}^T (sx_t + vq_t + y_t) \quad (1.21)$$

$$\text{s.t. } y_t \geq \max_{d \in \mathcal{U}_t} \left\{ h \left(\sum_{k=1}^t q_k - d_k \right) \right\} \quad \forall t \in [T] \quad (1.22)$$

$$y_t \geq \max_{d \in \mathcal{U}_t} \left\{ -b \left(\sum_{k=1}^t q_k - d_k \right) \right\} \quad \forall t \in [T] \quad (1.23)$$

$$0 \leq q_t \leq Mx_t \quad \forall t \in [T] \quad (1.24)$$

$$x_t \in \{0, 1\} \quad \forall t \in [T] \quad (1.25)$$

$$y_t \geq 0 \quad \forall t \in [T] \quad (1.26)$$

, where $\mathcal{U}_t = \{z \in [0, 1]^t \mid \|z\|_1 \leq \Gamma_t\}$ and for all $t \in [T]$ the budgets $\Gamma_t \in [0, t]$ are input parameters that control the conservatism of the solution and where b is a backlogging penalty.

Recently Hamed Mamani (2016) extends their work to more sophisticated uncertainty sets (e.g. CLT based Bandi and Bertsimas (2012)) and Metzker Soares et al. (2022) used the same methodology to LSP under yield uncertainty. These three papers derive closed form policies that are useful for practitioners because they can be easily implemented on large horizon without computational efforts. The use of such closed form policy is however restricted to cases in which the ordering or production periods are known in advance as they only compute the optimal ordering quantities.

A different solution approach was proposed in Bienstock and Ozbay (2008) for the same problem without setup costs. They point out that considering constraints (1.22) and constraints (1.23) can lead to sub-optimal robust decisions. In fact, this formulation allows the inner maximization problems to fix the demand for a same period at two different values in two distinct constraints. As a result, the worst holding/backlogging cost y_t can be overestimated and this can lead to over-conservative decisions. To overcome

this problem, Bienstock and Ozbay (2008) propose an iterative cutting plane algorithm in which a single inner maximization problem is solved at each iteration. They illustrate the efficiency of their approach by solving large instances of the problem (up to 300 periods). Their works have been extended by Thorsen and Yao (2017) to the case with demand and lead-time uncertainty, they use numerical experimentations to illustrate the conservatism of the constraint dualization approach compared to the cutting plane approach.

More complex cases have also been studied in the context using a RO approach. Aouam and Brahimi (2013) focus on an integrated production planning problem under demand uncertainty in which the producer can decide to satisfy only a part of a client demand, only incurring backlogging penalties on accepted orders. They also integrate the notion of Work In Progress to model the congestion of production systems. Their experimental results show that integrating order acceptance help to reduce congestion, increase the fulfillment rate and increase companies profitability under uncertain demands. Varas et al. (2014); Alem et al. (2018) propose a model to integrate lot-size and scheduling decision in a robust context under demand uncertainty. Their model aims to define the lot size of final product and schedule the usage of raw material with a time capacity for the scheduling part. A nearly similar approach is presented in José Alem and Morabito (2012) to combine lot sizing decisions and cutting stock problem. They face production and holding costs, backlogging penalties but also integrate costs related to cutting stock problem (e.g. trim loss). Finally, Agra et al. (2018) study a multi-product problem directly inspired from an industrial case.

Most of the existing literature uses the classical budget based approach and does not integrate more sophisticated structures of uncertainty sets such as presented in Section 1.2 and 1.3.

2.3 Dynamic robust inventory and production management

Dynamic solutions are applicable when a decision maker can adjust a part of its decisions once the uncertainty is (at least partially) revealed. In inventory management, such approaches allow to exploit the temporal dimensions of lot sizing problems and provide less conservative solutions than classic RO formulations. One solution to correct this lack of flexibility, *Adjustable Robust Optimization (ARO)* extends the RO approach to a dynamic context by making some decisions adjustable to the realization of the uncertain parameters. Multi-stage ARO suppose that the uncertainty is revealed over time and that decisions must adjust to the current knowledge of the uncertainty. In the last two decades,

there has been many studies on ARO models, their computational complexity and how to solve them. As this review focuses on papers centered on inventory or production systems management that propose significant results on this topic, we refer interested readers to the literature review presented in Daryalal et al. (2023) for a broader description of ARO.

In the multi-stage context implied by the structure of lot-sizing problems, ARO models often require to approximate the true adjustable problem to remain computationally tractable. In that objective Ben-Tal et al. (2004) introduce the concept of affinely adjustable robust solutions, in which the adjustable decision variables are replaced by affine functions of realizations of a subset of the uncertain parameteres. In the lot sizing problem for example, the quantity ordered/produced is expressed as a linear combination of past demand realizations called *linear decision rules*. Let $t \in [T]$ and assume that we are given the vector $\mathbf{d}_{t-1} = [d_1, \dots, d_{t-1}]$ of realized demands up to period $t - 1$, the decision variables q_t, y_t in period t then become:

$$q_t(\mathbf{d}_{t-1}) = q_t^0 + \sum_{k=1}^{t-1} q_t^k d_k \quad (1.27)$$

$$y_t(\mathbf{d}_{t-1}) = y_t^0 + \sum_{k=1}^{t-1} y_t^k d_k \quad (1.28)$$

Ben-Tal et al. (2004) apply this approach to solve a multi period, single item capacitated LSP with multiple factories. They minimize the sum of all time-dependent production costs in order to meet customers demand. They model the demand uncertainty with a box uncertainty set and show the efficiency of their approach with experimental results.

The same methodology has been applied in Aharon et al. (2009) to multi-echelon supply chain problem with backordering and setup costs to study the impact of information sharing. The same problem is studied in Melamed et al. (2016) under box and ellispoidal uncertainty sets. See and Sim (2010) considered the same problem and propose to use *truncated linear decision rules*. They provide models for the static case and both type of decision rules under an ellispoidal uncertainty set but without any numerical results.

Finally, Solyalı et al. (2016) compare different formulations of static and dynamic RLSP models. They compare both the quality of solutions and the tractability of models under the formulations of Bertsimas and Thiele (2006), the adjustable policy of Ben-Tal et al. (2004) and See and Sim (2010). They introduce a new formulation based on a disaggregated model (i.e. facility location based LSP model). Bertsimas and Georghiou (2015) compared the efficiency of affine, piecewise linear and piecewise constant decision

rules in a similar problem with supplier selection. At each period, an order is submitted to a supplier with either instant delivery or a delay of period. Their numerical results highlight that using several types of decision rules provides more flexibility in the nature of adjustable decisions and outperforms other methods with a single type of decision rule.

Decision rules provide an efficient way to obtain a lower bound of the true ARO problem by restricting the form of adjustable policies. In the opposite direction, Santos et al. (2018) propose a lower bound for the ARO version of the lot sizing problem with and without setup costs. They obtain a lower bound by solving a relaxed problem in which the non-anticipativity constraints are relaxed and decisions are affine functions of the demand on the whole planning horizon.

All studies considered up to now focus on uncertain demand, which represents the vast majority of the existing literature. However, recently other sources of uncertainty have been studied in the context of ARO. Thevenin et al. (2022) propose an ARO approach to the LSP with uncertain lead time (order cross-over), supplier selection and setup costs. They use an exact solution method based on a constraint and column generation approach (Zeng (2011)) and find through their numerical experiments that the problem is hard to solve. To overcome this issue, they propose a fix and optimize heuristic that leads to a significant reduction of the computation time.

Metzker Soares et al. (2023) study the dynamic lot sizing problem under yield uncertainty and backorders via affinely adjustable robust policy. They show that the robust counterpart based on affine decision rules is not tractable and compare different ways to approximate it as a LP. More complex LSPs have also been studied via the ARO approach. For example, Other works also include the relationship between multiple actors within a supply chain. For instance, Ben-Tal et al. (2005) study a retailer-supplier commitment contract problem in which the retailer makes prediction of its ordering quantities over the planning horizon before any realization of the uncertain demands. These quantities may be adjusted in a dynamic way after observing the demand, but induce additional penalties when they deviate from its predictions (in addition to the LSP costs of setup, ordering, holding and backorder). In their work, Ben-Tal et al. (2005) provide an ARO model for this problem under classical box and ellipsoidal uncertainty sets. In this context, ARO is an efficient way to obtain less conservative static decisions by reducing the amount of items ordered during the planning horizon.

CONTRIBUTIONS

In the previous chapter, we introduced the main robust optimization models used to take decisions in supply chains and production systems. In this context, most papers focus on uncertain demand while only two papers consider yield uncertainty (Metzker Soares et al., 2022, 2023) and one paper focus on uncertain lead time and demand (Thorsen and Yao, 2017). Moreover, in addition to the lack of diversity in the type of uncertainties, most works have used classical box, budget-based or ellipsoidal uncertainty sets while only three papers have used more complex uncertainty set (CLT based) and only one paper has discussed the case of correlated uncertainty (Hamed Mamani, 2016). The use of more sophisticated uncertainty sets integrating statistical information in robust models such as described in Section 1.2 remains largely unexplored in the case of lot sizing models. Moreover, one could exploit the amount of data collected by modern production systems in order to propose a set of suitable uncertainty sets in order to handle the uncertainties in a DDRO approach.

This gap in the literature may find its explanation in computational limitations. We pointed out in Section 1.3 that despite their efficiency to model the behavior of uncertain parameters and due to their more complex formulations, the use of data-driven uncertainty sets comes at the expense of computational tractability. Such sets often lead to large scale robust counterparts that are hard to solve in practice Zhang et al. (2018b); Han et al. (2021); Bertsimas et al. (2017). This scalability issue is a limiting factor to the application of DDRO in industrial problems often involving large scale mixed integer linear programs. The rest of this chapter summarizes how the contributions detailed in Chapter 3, 4 and 5 attempt to address this issue.

Chapter 3: Approximate kernel learning uncertainty set In this first contribution we study the scalability of the data driven uncertainty set based on support vector clustering. This uncertainty set has been applied successfully on several optimization problems but suffers from scalability issues (see Chapter 1, paragraph 1.3.1.1). In fact,

the number of decision variables and constraints introduced in the RC is proportional to both the number of uncertain parameters and the size of the data-set used to formulate it. In practice, it results in large scale robust formulations that are practically intractable when the initial problem is a MILP. In this chapter, we propose a two-phase method to approximate the resulting uncertainty sets and overcome these tractability issues. The main idea behind this approximation is to reduce the number of variables and constraints introduced in the model by selecting a subset of historical data to construct an approximate uncertainty set. This subset is selected by an optimization process whose objective is at first to maximize the similarity with the initial uncertainty set and then to minimize the number of data point used in the formulation.

We evaluate our approximation method on three distinct well known optimization problems, the robust binary knapsack problem, the robust lot-sizing problem as formulated by Bertsimas and Thiele (2006) and a robust traveling salesman problem. Experimental results show that the approximated uncertainty set leads to solutions that are comparable to those obtained with the classic SVC-based uncertainty set on all problems with a significant reduction of the computation time.

Chapter 4: A data driven robust optimization approach to supply planning.

In this chapter, we focus on a problem derived from a practical case encountered in the aircraft industry. We consider a single assembly line that produces several distinct item in order to satisfy the needs that are planned over a discrete and finite horizon. The components are available in an external warehouse and the picking and delivery of components to the assembly line are subcontracted to an external logistic provider following the manufacturer orders. The problem consists in planning the delivery of components on the assembly line and the quantity of each final item to produce at each period of the planning horizon. The manufacturer follows the assemble to order rule, which means that an item can not be produced earlier than its due date. In addition, the time needed by the logistic provider to pick a set of components is uncertain and the total picking time of the logistic provider during a given period is bounded by a capacity. If the planning defined by the manufacturer exceeds this capacity in a given period, a part of the components will not be delivered on time to the assembly line, possibly disrupting the production plan and affecting its reliability. On the contrary, components delivered in advance to the assembly line can reduce the efficiency of workers and the risk of incidents. The objective is then to find a compromise between guaranteeing the feasibility of the

production plan and minimizing the stock held on the assembly line.

The manufacturer can track the state of each picking operation and the flow of component from the warehouse to the assembly line leading to an historical data set of picking and delivery time. We propose a robust mixed integer model based on a capacitated lot sizing model to solve this problem. We compare the use of a classic budget based uncertainty set and the SVC based uncertainty set to integrate picking time uncertainties in the model. Building upon the contribution of Chapter 3, we use the approximation of the SVC based uncertainty set to solve the data driven robust models in a reasonable amount of time.

We conduct a numerical experiments on simulated data set inspired from the true industrial case. The results show that the robust model can significantly reduce the impact of picking time uncertainties on the assembly operations. In addition, the use of the SVC-based uncertainty set leads to a lower level of component held on the border of the assembly line while guaranteeing the reliability of the production plan.

Chapter 5: Data-driven Robust inventory management under uncertain demand and lead time In this chapter we consider a more classic single item inventory management problem under the two most common sources of uncertainty: demand and lead time. From our previous studies, it appears that non-parametric models such as SVC are particularly well suited to construct uncertainty sets when a large amount of data is available. However, such large data-set are hard to obtain for demand or lead time uncertainties as it would cover a very long time horizon during which the demand process can evolve. In this context, we propose an alternative approach to build uncertainty sets from a restricted set of historical data in which we assume that uncertain parameters follows a specific (but unknown) stochastic process. We formulate an uncertainty set for demand based on *Autoregressive Moving Average (ARMA)* process in order to provide a flexible framework that can adapt to a wide range of demand processes. We conducted a set of experiments for problems with lead time uncertainties, setup costs and autoregressive demand. We compare the solutions obtained by a classic robust model using the budget based uncertainty set, a sample average approximation model and our approach. Experimental results show that our method results in less conservative solutions than a classic budget based approach on average while maintaining a significant robustness advantage over the sample average approach.

As mentioned in the previous chapter, using a dual reformulation approach to solve

this problem can be over-conservative as the same uncertain parameters appear in several constraints simultaneously. Thus, we use a cutting plane approach such as described in Thorsen and Yao (2017). In addition to a reduction of conservatism, this solution approach provides several opportunities to improve uncertainty sets. In this context we propose two improvements of the existing approach used to model lead time in robust optimization: (i) We propose a discrete uncertainty set for lead times while the rest of the literature have only considered continuous lead time models (Thorsen and Yao (2017); Thevenin et al. (2022)), and (ii) we propose a simple approach to link uncertain lead time and demand using linear regression in order to integrate potential correlation in the variations of both types of uncertain parameters. In a second set of experiments, we study the impact of uncertain lead time on the solution proposed by our three models. Experimental results show that uncertain lead time reduces the gap of cost among all the methods while preserving the conclusions drawn from the first set of experiments.

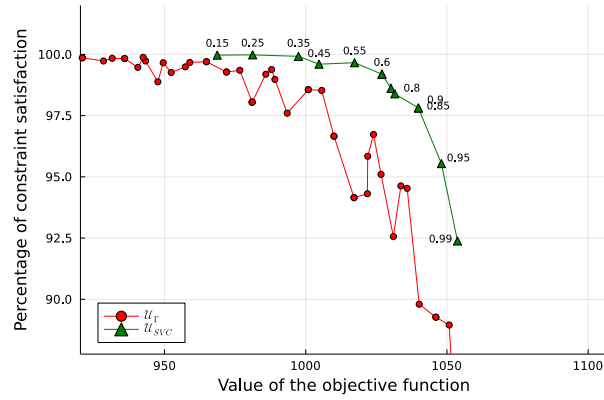
APPROXIMATE KERNEL LEARNING UNCERTAINTY SET

1 Introduction


As stated in Chapter 1, one of the main approaches to build data-driven uncertainty sets is the use of support vectors. Since the introduction of this technique by Shang et al. (2017) it has been applied in several applications (see Chapter 1). Unfortunately, all these works agree on the lack of scalability of the method with respect to the number of uncertain parameters. Indeed, so-called SVC-based uncertainty sets require the introduction of a large number of new variables and constraints in the robust counterparts of optimization models. These sets of variables and constraints grow proportionally to both the dimension of the uncertainty and the number of data samples used to build them. In many practical settings, in which we face a large optimization problem with a large data-set, this results into intractable formulations. A first attempt to overcome this tractability issue was proposed in Loger et al. (2022) to solve a robust inventory management problem under SVC-based uncertainty set. As an example, Figure 3.1b shows the increase in CPLEX CPU time to solve the robust 0-1 Knapsack problem with 30 items under uncertain weights with respect to the number of data samples.

This chapter proposes a method to reduce the number of additional variables and constraints associated with an SVC-based uncertainty set. More precisely, we develop a method to approximate the SVC-based uncertainty set introduced in Shang et al. (2017) in order to overcome the scalability issue that was pointed out. The solutions computed with the resulting approximate uncertainty set are identical or comparable to those obtained with the SVC-based uncertainty set with a significant reduction of the computational burden.

The rest of this chapter is organized as follows: Section 2 describes the uncertainty



(a)



(b)

Figure 3.1 – Performance of the SVC-based uncertainty set on a robust 0-1 Knapsack problem with uncertain weights and 30 items. Top: Trade-off between solution quality and feasibility using the SVC-based (\mathcal{U}_{SVC}) and the budget-based (\mathcal{U}_Γ) uncertainty sets. Bottom: Evolution of CPLEX CPU time in seconds using the SVC-based uncertainty set for different sizes of data sets.

set of Shang et al. (2017) and provides some explanations on its scalability issues. We introduce our approximation procedure in Section 3 and present experimental results in Section 4. The latter first focuses on the ability of the approximation to define a set that resembles the original SVC-based uncertainty set, then provides evidence of the optimization and computational performance on three different optimization problems.

2 Support vector clustering-based uncertainty set

SVC is a well-known Machine Learning algorithm, introduced by Shang et al. (2017), that describes the support of a multi-dimensional random variable with a cluster of minimal volume containing a given portion of historical realizations. Among the various advantages displayed by SVC, its usage as a ML tool to define uncertainty sets is motivated by the fact that it naturally incorporates information on mixed moments of distributions such as the covariance of random variables. In addition, depending on the chosen Kernel function, it does not necessitate any complex tuning of hyper-parameters. To make this chapter self-contained, we briefly describe below the approach developed in Shang et al. (2017).

Assume that we dispose of a set $\mathcal{D} = \{\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)}\}$ of N data samples in an m -dimensional space. The principle of SVC is to use a non-linear mapping $\phi(u) : \mathbb{R}^m \rightarrow \mathbb{R}^k$ to a high-dimensional feature space \mathcal{F} and to seek the smallest sphere that encloses at least $(1 - \nu)N$ data samples, where the parameter $\nu \in]0, 1[$ is a lower bound on the portion of data samples considered as outliers. The general idea is to express this optimization problem as a Lagrangian function that is then minimized using the KKT conditions. Shang et al. (2017) apply the so-called *kernel trick* with a suitable piecewise linear function called WGIK (Weighted General Intersection Kernel) to obtain a quadratic program (QP) that computes the best sphere with respect to \mathcal{D} . The different step used to obtain the QP and its interpretation are given in Appendix A.

The QP solution defines a subset \mathcal{S} of data points, called *support vectors* (SV), that lie outside of the boundary of the sphere. A subset $\mathcal{B} \subseteq \mathcal{S}$ of these points are exactly on the boundary of the sphere and are referred to as *boundary support vectors* (BSV). To summarize, those two sets are such that $\mathcal{B} \subseteq \mathcal{S} \subseteq \mathcal{D}$ and $|\mathcal{S} \setminus \mathcal{B}| \leq N\nu \leq |\mathcal{S}|$. In the robust optimization framework, the parameter ν is diverted from its original purpose to control the size of the uncertainty set and consequently the level of conservatism of the model. That is, selecting a larger value of ν increases the number of SV and decreases

the conservatism. Figure 3.2 represents these sets for 1000 samples following a bivariate gaussian distribution when ν is set to 0.15.




Figure 3.2 – Representation of \mathcal{S} and \mathcal{B} for a bivariate gaussian distribution when $\nu = 0.15$

From now on, we slightly abuse the notation and sometimes refer to the data point $\mathbf{u}^{(i)}$ using its index i , e.g., writing $i \in \mathcal{S}$ instead of $\mathbf{u}^{(i)} \in \mathcal{S}$, when the meaning is clear from the context. For each point $i \in \mathcal{D}$, the QP also computes a value $\alpha_i \in [0, 1/\nu N]$, where $\alpha_i = 0$ iff $i \in \mathcal{D} \setminus \mathcal{S}$, $\alpha_i = 1/\nu N$ iff $i \in \mathcal{S} \setminus \mathcal{B}$ and $0 < \alpha_i < 1/\nu N$ if $i \in \mathcal{B}$. Let $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_N]$ be the vector of obtained coefficients, Shang et al. (2017) use $\boldsymbol{\alpha}$ to define the following data-driven uncertainty set:

$$\mathcal{U}_{\text{SVC}} = \left\{ \mathbf{u} \mid \sum_{i \in \mathcal{S}} \alpha_i \|\mathbf{Q}(\mathbf{u} - \mathbf{u}^{(i)})\|_1 \leq \theta \right\} \quad (3.1)$$

where $\|\cdot\|_1$ stands for the ℓ_1 -norm, \mathbf{Q} is a weighting matrix and $\theta = \min_{i' \in \mathcal{B}} \left(\sum_{i \in \mathcal{S}} \alpha_i \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 \right)$. For simplicity, we also define for all $\bar{\boldsymbol{\alpha}} \in \mathbb{R}_+^{|\mathcal{S}|}$ and $\mathbf{u} \in \mathbb{R}_+^m$ the function:

$$\theta(\bar{\boldsymbol{\alpha}}, \mathbf{u}) = \sum_{i \in \mathcal{S}} \bar{\alpha}_i \|\mathbf{Q}(\mathbf{u} - \mathbf{u}^{(i)})\|_1 \quad (3.2)$$

which measures the weighted distance of \mathbf{u} to all points in \mathcal{S} used in definition (3.1) of \mathcal{U}_{SVC} . Note that $\mathcal{U}_{\text{SVC}} = \{\mathbf{u} \mid \theta(\boldsymbol{\alpha}, \mathbf{u}) \leq \theta\}$. In particular, there exists a data point $i' \in \mathcal{B}$ such that $\theta(\boldsymbol{\alpha}, \mathbf{u}^{(i')}) = \theta$, where θ is the threshold used to define the boundaries of \mathcal{U}_{SVC} .

In what follows, we extend definition (3.1) and denote

$$\mathcal{U}_{\bar{\alpha}} = \left\{ \mathbf{u} \left| \sum_{i \in \mathcal{S}} \bar{\alpha}_i \|\mathbf{Q}(\mathbf{u} - \mathbf{u}^{(i)})\|_1 \leq \theta \right. \right\} = \{ \mathbf{u} \mid \theta(\bar{\alpha}, \mathbf{u}) \leq \theta \} \quad (3.3)$$

for *any* vector $\bar{\alpha} \in \mathbb{R}_+^m$. In particular when $\bar{\alpha} = \alpha$, we have $\mathcal{U}_{\text{SVC}} = \mathcal{U}\alpha$. An appealing property of this class of uncertainty sets is that one can easily reformulate $\mathcal{U}_{\bar{\alpha}}$ as a polyhedron using auxiliary variables $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_{|\mathcal{S}|}]$:

$$\mathcal{U}_{\bar{\alpha}} = \left\{ \mathbf{u} \left| \begin{array}{l} \exists \mathbf{v}_i \in \mathbb{R}_+^m, i \in \mathcal{S} \text{ s.t.} \\ \sum_{i \in \mathcal{S}} \bar{\alpha}_i \mathbf{v}_i^T \mathbf{1} \leq \theta \\ -\mathbf{v}_i \leq \mathbf{Q}(\mathbf{u} - \mathbf{u}^{(i)}) \leq \mathbf{v}_i, i \in \mathcal{S} \end{array} \right. \right\} \quad (3.4)$$

which is bounded and nonempty for $0 < \nu < 1$ (Shang et al. (2017)). Based on formulation (3.4), the left-hand side of the robust linear constraint

$$\max_{\mathbf{a} \in \mathcal{U}_{\text{SVC}}} \mathbf{a}^T \mathbf{x} \leq b \quad (3.5)$$

with $\mathbf{a} \in \mathcal{U}_{\text{SVC}} \subseteq \mathbb{R}^m$ and $b \in \mathbb{R}$, is thus equivalent to the following LP:

$$\begin{aligned} & \max \mathbf{a}^T \mathbf{x} \\ & \text{s.t.} \quad \sum_{i \in \mathcal{S}} \alpha_i \mathbf{1}^T \mathbf{v}_i \leq \theta \\ & \quad \quad -\mathbf{v}_i \leq \mathbf{Q}(\mathbf{a} - \mathbf{u}^{(i)}) \leq \mathbf{v}_i \quad \forall i \in \mathcal{S} \end{aligned}$$

which is feasible and bounded since \mathcal{U}_{SVC} is nonempty and bounded for $0 < \nu < 1$. By strong duality, the dual of this problem

$$\min \sum_{i \in \mathcal{S}} (\boldsymbol{\mu}_i - \boldsymbol{\lambda}_i)^T \mathbf{Q} \mathbf{u}^{(i)} + \eta \theta \quad (3.6)$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{S}} \mathbf{Q}(\boldsymbol{\lambda}_i - \boldsymbol{\mu}_i) + \mathbf{x} = \mathbf{0} \quad (3.7)$$

$$\boldsymbol{\lambda}_i + \boldsymbol{\mu}_i = \eta \alpha_i \mathbf{1} \quad \forall i \in \mathcal{S} \quad (3.8)$$

$$\boldsymbol{\lambda}_i, \boldsymbol{\mu}_i \in \mathbb{R}_+^m \quad \forall i \in \mathcal{S} \quad (3.9)$$

$$\eta \geq 0 \quad (3.10)$$

is also feasible and bounded and their optimal value coincide. Therefore, the robust

counterpart of constraint (3.5) is given by

$$\sum_{i \in \mathcal{S}} (\boldsymbol{\mu}_i - \boldsymbol{\lambda}_i)^T \mathbf{Q} \mathbf{u}^{(i)} + \eta \theta \leq b \quad (3.11)$$

with the additional constraints (3.7)-(3.10).

As the formulation (3.11), (3.7)-(3.10) suggests, the size of the robust counterpart depends upon the number of support vectors. Specifically, consider a robust constraint j that contains m_j uncertain parameters taking their values in an SVC uncertainty set defined with support vectors \mathcal{S}_j . The robust formulation then adds $2|\mathcal{S}_j| \cdot m_j + 1$ new variables and $(|\mathcal{S}_j| + 1) \cdot m_j$ new constraints to the original one. Since the number $|\mathcal{S}_j|$ of support vectors is closely related to νN , the size of the robust counterpart is greatly influenced by both the regularization parameter ν and the size of the original dataset \mathcal{D} . In practice, increasing one of these values may greatly affect the tractability of the original program and become a significant limitation to the application of this framework. This phenomenon is particularly problematic in the context of big data where one would need to take advantage of information contained in a high volume of collected data. Moreover, it is well known that the SVC algorithm suffers from the *curse of dimensionality*, as stated by Keogh and Mueen (2017): « *For machine learning problems, a small increase in dimensionality generally requires a large increase in the numerosity of the data, in order to keep the same level of performance for regression, clustering, etc.* ». In our case, increasing the number of uncertain parameters m_j requires more data samples for the SVC algorithm to be efficient, leading in turn to a large robust counterpart unlikely to be solvable in a reasonable time.

3 Approximation of the SVC-based uncertainty set

To overcome the tractability issue pointed out in the previous section, we propose a two-step method to reduce the size of the robust program derived from the SVC-based uncertainty set. The goal is to curtail the number of support vectors that define \mathcal{U}_{SVC} , while keeping a good representation of the original cluster obtained via the SVC.

Assume that, for a given ν , coefficients $\boldsymbol{\alpha}$ and the corresponding SVC-based uncertainty set \mathcal{U}_{SVC} are computed following the methodology of Shang et al. (2017). In a nutshell, the procedure consists of computing a vector of modified weights $\hat{\boldsymbol{\alpha}} = [\hat{\alpha}_1, \dots, \hat{\alpha}_N]$ that is used to derive an approximation of the original uncertainty set. The algorithm

takes an integer K as an input parameter, which sets the maximum number of strictly positive coordinates of $\hat{\boldsymbol{\alpha}}$, thereby limiting the number of SV that shape the (modified) uncertainty set. This in turn reduces the size of the resulting LP (3.6)-(3.10) and its robust counterpart, since we only need to include dual variables for SV $\mathbf{u}^{(i)}$ with a strictly positive coefficient $\hat{\alpha}_i$ in their formulation. After selecting an initial subset of SV, the algorithm runs in two phases: (1) optimize the weights $\hat{\boldsymbol{\alpha}}$ of the selected points in order to minimize a given distance between $\mathcal{U}\hat{\boldsymbol{\alpha}}$ and \mathcal{U}_{SVC} , then (2) minimize the number of points selected without deteriorating the distance obtained during the first step. We detail the procedure below.

3.1 Initial support vectors selection by K-medoid clustering

As a strategy to reduce the size of the set \mathcal{S} , the algorithm starts by choosing a first subset $\hat{\mathcal{S}}_0$ of support vectors that are *a priori* good representatives of \mathcal{S} . By default, we include all the boundary support vectors \mathcal{B} in $\hat{\mathcal{S}}_0$, as well as K additional support vectors selected from $\mathcal{S} \setminus \mathcal{B}$. The general approach consists of partitioning $\mathcal{S} \setminus \mathcal{B}$ into K distinct subsets of data points sharing similar features and aggregate the information contained in each subset into a single data point. The underlying idea is that close samples that lie strictly outside \mathcal{U}_{SVC} have similar contributions to cluster boundaries.

We perform the partition using the so-called K -medoid clustering algorithm described in (Schubert and Rousseeuw, 2019), also known as partitioning around medoids (PAM, see Kaufman and Rousseeuw (1990)). This procedure assigns each sample in \mathcal{S} to exactly one cluster in order to minimize its distance to the center of this group. One appealing property of the K -medoid clustering compared to the more well-known K -means clustering is that it forces the center of each cluster to be an actual data point. In our case, this ensures that any selected point belongs to $\mathcal{S} \setminus \mathcal{B}$ and therefore lies outside of the initial uncertainty set \mathcal{U}_{SVC} . We compute the distances between two points $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$ using the function

$$d(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\|_2 \tag{3.12}$$

where $\|\cdot\|_2$ represents the ℓ_2 -norm. At the end of the K -medoids procedure, we obtain K clusters along with their respective centers. The union of \mathcal{B} with these centers form the initial set $\hat{\mathcal{S}}_0$, as represented in Figure 3.3b. In our numerical experiments, we rely on an existing implementation of the algorithm available in julia¹.

1. <https://juliastats.org/Clustering.jl/dev/index.html>

3.2 Procedure $OPT\text{-}\alpha$

Recall from the definition (3.1) that we use the quantity $\theta(\boldsymbol{\alpha}, \mathbf{u}^{(i)}) = \sum_{i \in \mathcal{S}} \alpha_i \|\mathbf{Q}(\mathbf{u} - \mathbf{u}^{(i)})\|_1$ to determine whether a point $\mathbf{u} \in \mathbb{R}_+^m$ belongs to \mathcal{U}_{SVC} . Following the reduction of the original set of support vectors into the smaller subset $\widehat{\mathcal{S}}_0$, our goal is now to define an alternative uncertainty set $\mathcal{U}^{\widehat{\alpha}}$ where $\widehat{\alpha}_i > 0$ iff $i \in \widehat{\mathcal{S}}_0$. A simple modification would be to set $\widehat{\alpha}_i = \alpha_i \mathbb{1}_{\{i \in \widehat{\mathcal{S}}_0\}}$, but this is likely to distort the boundaries of the resulting uncertainty set. Instead, we propose to compute modified weights $\widehat{\alpha}$ such that the uncertainty set $\mathcal{U}^{\widehat{\alpha}}$ approximate \mathcal{U}_{SVC} in a satisfactory fashion.

Our goal is then to limit as much as possible the size of the subset $(\mathcal{U}_{\text{SVC}} \setminus \mathcal{U}^{\widehat{\alpha}}) \cup (\mathcal{U}^{\widehat{\alpha}} \setminus \mathcal{U}_{\text{SVC}})$. Since a direct calculation of the latter would clearly be too cumbersome, we instead use the initial sample of data points to measure the quality of the approximation $\mathcal{U}^{\widehat{\alpha}}$.

Specifically, we rely on the empirical cumulated absolute deviation

$$\Delta = \sum_{i \in \mathcal{D}} |\theta(\widehat{\alpha}, \mathbf{u}^{(i)}) - \theta(\boldsymbol{\alpha}, \mathbf{u}^{(i)})| \quad (3.13)$$

Notice that all points $i \in \mathcal{S}$ that are discarded in $\widehat{\mathcal{S}}_0$ are not candidate to serve as support vectors. Hence, we artificially set their modified weights $\widehat{\alpha}_i$ to 0. We use Δ as a mean to evaluate the quality of the obtained approximation. We seek values $\widehat{\alpha}$ that minimize this empirical distance, which amounts to solving the optimization problem

$$\min_{\widehat{\alpha}} \left\{ \sum_{i' \in \mathcal{D}} \left| \sum_{i \in \mathcal{S}} \alpha_i \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 - \sum_{i \in \mathcal{S}} \widehat{\alpha}_i \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 \right| \right\} \quad (3.14)$$

It is straightforward that (3.14) is equivalent to the following linear program:

$$\min \quad \sum_{i' \in \mathcal{D}} \Delta_{i'} \quad (3.15)$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{S}} (\alpha_i - \widehat{\alpha}_i) \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 \leq \Delta_{i'} \quad \forall i' \in \mathcal{D} \quad (3.16)$$

$$(OPT\text{-}\alpha) \quad \sum_{i \in \mathcal{S}} (\widehat{\alpha}_i - \alpha_i) \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 \leq \Delta_{i'} \quad \forall i' \in \mathcal{D} \quad (3.17)$$

$$\widehat{\alpha}_i = 0 \quad \forall i \in \mathcal{S} \setminus \widehat{\mathcal{S}}_0 \quad (3.18)$$

$$\widehat{\alpha}_i \geq 0 \quad \forall i \in \widehat{\mathcal{S}}_0 \quad (3.19)$$

$$\Delta_{i'} \geq 0 \quad \forall i' \in \mathcal{D} \quad (3.20)$$

where constraints (3.16) and (3.17) linearize the absolute distance used in (3.14) and the variables $\Delta_{i'}$ are the absolute deviations described in (3.13) evaluated for a given point $i' \in \mathcal{D}$. The solution to this LP provides the optimal weights $\hat{\alpha}_0$ that we can then use to define the modified uncertainty set $\mathcal{U}_{\text{KM}} = \mathcal{U}\hat{\alpha}_0$. This first approximation is illustrated on Figure 3.3b where the orange triangles represent the medoids that constitute $\hat{\mathcal{S}}_0$.

3.3 Reduction of the number of support vectors

The objective of the second phase is to further reduce the number of support vectors retained in $\hat{\mathcal{S}}_0$ while keeping a good quality of the first approximation. At this stage, we relax the constraint that a support vector must belong to the initial set of medoids or to \mathcal{B} . Instead, we consider any subset of points $\hat{\mathcal{S}} \subset \mathcal{S}$ as potential support vector for the approximation. Let Δ^* be the value of the optimal solution to $(OPT-\alpha)$ computed in the first phase. We impose that the solution derived from this second phase: (1) selects at most $K + |\mathcal{B}|$ support vectors, i.e. $|\hat{\mathcal{S}}| \leq |\hat{\mathcal{S}}_0|$, and (2) defines new weights $\hat{\alpha}^*$ such that $\hat{\alpha}_i^* > 0$ iff $i \in \hat{\mathcal{S}}$ and the resulting empirical distance between \mathcal{U}_{SVC} and $\mathcal{U}\hat{\alpha}^*$ is at most Δ^* . We introduce the set of binary variables

$$y_i = \begin{cases} 1 & \text{if } \mathbf{u}^{(i)} \in \hat{\mathcal{S}} \\ 0 & \text{otherwise} \end{cases}$$

and formulate the above problem with the following MIP:

$$\min \sum_{i \in \mathcal{S}} y_i \tag{3.21}$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{S}} (\alpha_i - \hat{\alpha}_i) \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 \leq \Delta_{i'} \quad \forall i' \in \mathcal{D} \tag{3.22}$$

$$\sum_{i \in \mathcal{S}} (\hat{\alpha}_i - \alpha_i) \|\mathbf{Q}(\mathbf{u}^{(i')} - \mathbf{u}^{(i)})\|_1 \leq \Delta_{i'} \quad \forall i' \in \mathcal{D} \tag{3.23}$$

$$\sum_{i' \in \mathcal{D}} \Delta_{i'} \leq \Delta_0^* \tag{3.24}$$

$$\hat{\alpha}_i \leq y_i \quad \forall i \in \mathcal{S} \tag{3.25}$$

$$\hat{\alpha}_i \geq 0 \quad \forall i \in \mathcal{S} \tag{3.26}$$

$$\Delta_{i'} \geq 0 \quad \forall i' \in \mathcal{D} \tag{3.27}$$

$$y_i \in \{0, 1\} \quad \forall i \in \mathcal{S} \tag{3.28}$$

where (3.21) minimizes the number of SV retained in $\widehat{\mathcal{S}}$. Constraints (3.22)-(3.24) ensure that the empirical distance achieved by the selected support vectors with their modified weights is no greater than Δ^* , i.e., that the quality of the first approximation is preserved. Finally, constraints (3.25) ensure that the weight $\hat{\alpha}_i$ associated to a sample point $\mathbf{u}^{(i)}$ is positive only if the variable y_i is equal to one, that is if it belongs to $\widehat{\mathcal{S}}$.

The solution to the MIP above defines a new subset of support vectors $\widehat{\mathcal{S}}$ that we can use to approximate \mathcal{U}_{SVC} . Notice that the optimal weights $\hat{\alpha}_0$ to $(OPT-\alpha)$ directly define a feasible solution of value $K + |\mathcal{B}|$ by letting $y_i = 1$ for all $i \in \widehat{\mathcal{S}}_0 \cup \mathcal{B}$ and $y_i = 0$ otherwise. It follows that the approximation \mathcal{U}_{KM} can be used as an initial solution to the MIP (3.21)-(3.28). The corresponding initial value is thus an upper bound, which ensures that the final set of $\widehat{\mathcal{S}}$ of SV satisfies $|\widehat{\mathcal{S}}| \leq |\widehat{\mathcal{S}}_0|$. Notice that at this point, the value of the modified weights $\hat{\alpha}$ may not be optimal since this MIP primarily focuses on reducing the size of $\widehat{\mathcal{S}}$. However, running the program $(OPT-\alpha)$ on this new subset $\widehat{\mathcal{S}}$ allows us to compute optimized weights $\hat{\alpha}^*$ in order to minimize the distance Δ defined by (3.13). Let $\mathcal{U}_{\text{AC}} = \mathcal{U}\hat{\alpha}^*$ be the resulting uncertainty set. Figure 3.3c represents the uncertainty set \mathcal{U}_{AC} and the set $\widehat{\mathcal{S}} \subseteq \mathcal{S}$.

3.4 Summary of the approach


We summarize the procedure to construct this new uncertainty set $\mathcal{U}\hat{\alpha}$ in three steps:

Initialization Apply the method of Shang et al. (2017) to obtain the weights α and the associated sets \mathcal{S} and \mathcal{B} .

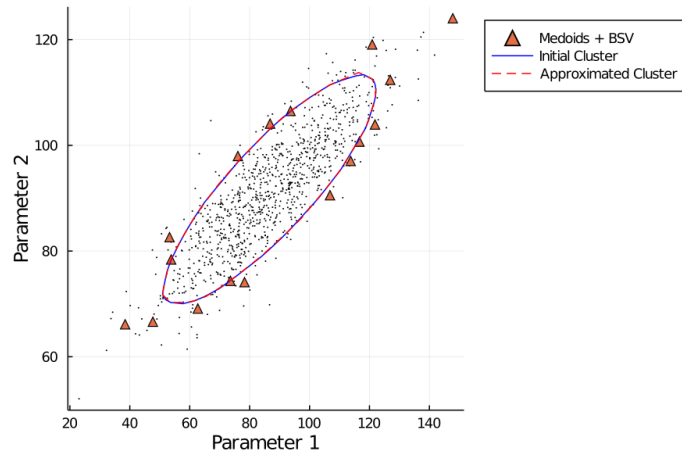
Step 1 Apply the K -medoid algorithm to obtain an initial subset $\widehat{\mathcal{S}}_0 \subseteq \mathcal{S}$, then solve $(OPT-\alpha)$ to derive the lower bound Δ^* and collect the modified the weights $\hat{\alpha}_0$ to define \mathcal{U}_{KM} .

Step 2 Solve the MIP (3.21) - (3.28) to obtain the final subset $\widehat{\mathcal{S}}$, then use $(OPT-\alpha)$ to refine the coefficients $\hat{\alpha}^*$ and obtain \mathcal{U}_{AC} .


One can then follow the method described in Section 2 to obtain the robust counterpart of a constraint with the new uncertainty set $\mathcal{U}_{\text{AC}} = \mathcal{U}\hat{\alpha}^*$. Figure 3.3 illustrates the evolution of the set of SV and the boundaries of the resulting uncertainty set throughout the different step of the procedure on a simple two-dimensional example, in which the two uncertain parameters follow a bivariate normal distribution.



(a) Initial cluster



(b) First approximation



(c) Final approximation

Figure 3.3 – Steps of uncertainty set construction

4 Experimental results

In this section we present several numerical experiments to evaluate the performance of the approximation methods introduced in this chapter.

The first part compares the approximation \mathcal{U}_{AC} derived from the procedure introduced in the previous section with the original SVC-based uncertainty set \mathcal{U}_{SVC} . The goal is to produce numerical evidence that the proposed optimization procedures lead to a satisfactory approximation in the sense that \mathcal{U}_{SVC} and \mathcal{U}_{AC} define comparable subsets of \mathbb{R}^m . In the second part, we apply the resulting uncertainty set on three distinct robust optimization problems. We find consequent performance improvements without significantly deteriorating the solution quality. Specifically, we first consider a robust knapsack problem with uncertain weights similar to the one studied in Bertsimas et al. (2004), a robust Traveling Salesman Problem (TSP) with travel time uncertainties, and the robust inventory management problem of Chapter 3 in which the demand is an autocorrelated stochastic process. For each problem we examine (1) the quality of the solutions obtained by the original SVC-based uncertainty set \mathcal{U}_{SVC} , and (2) the reduction of the computational burden incurred to solve the problem to optimality, compared to \mathcal{U}_{SVC} . As the MIP (3.21) - (3.28) is hard to solve to optimality, in what follows we restrict the time allowed to solve it to 1 hour. All models are implemented in Julia 1.9.0 with JuMP and solved with CPLEX 20.1 using a single thread on a Linux, Ubuntu 20.04.2 LTS, equipped with an Intel Xeon Gold 6230 clocked at 2.10GHz.

4.1 Uncertainty set comparison

In this section, we compare the portion of the distribution of random parameters covered by \mathcal{U}_{SVC} and the approximate uncertainty sets obtained after the second phase of the procedure. Our test bench covers different numbers of uncertain parameters and uses two types of distributions (Gamma and Gaussian). For each experiment, we generate a set \mathcal{D} of N data points that serves as input data to define the considered uncertainty sets (\mathcal{U}_{SVC} and \mathcal{U}_{AC}) and a distinct test set \mathcal{T} of 10000 data points to evaluate the out-of-sample error. Finally, we evaluate the portion of \mathcal{T} contained in \mathcal{U}_{SVC} and \mathcal{U}_{AC} . To avoid special cases and strengthen the conclusions, we repeat the above procedure 10 times. The cumulative out-of-sample error (Type I and II) is given as a percentage of $|\mathcal{T}|$. In what follows, the Type I error refers to the percentage of data points in \mathcal{T} that belongs to \mathcal{U}_{AC} but not to \mathcal{U}_{SVC} and the Type II error refers to the percentage of data points in

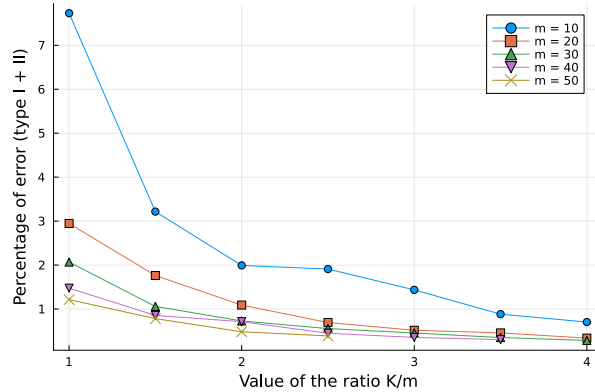


Figure 3.4 – Evolution of the out-of-sample cumulative error as a function of the ratio K/m for a multivariate Gamma distribution when $N = 1000$

\mathcal{T} that belongs to \mathcal{U}_{SVC} but not to \mathcal{U}_{AC} .

Figure 3.4 represents the cumulative error computed for \mathcal{U}_{AC} on the test sample \mathcal{T} for various values of m (number of uncertain parameters) and several ratios K/m . We represent the average error for various values of $\nu \in \{0.15, 0.25, 0.35, 0.45, 0.55, 0.60, 0.70, 0.75, 0.80, 0.85, 0.90, 0.95, 0.99\}$. It appears that the parameter K offers an efficient control over the quality of the approximation.

Figure 3.5 displays the average percentage of total error obtained as a function of ν . We find that this percentage increases as ν increases from 0.15 to 0.8 but decreases on $[0.8, 0.99]$.

Remarks:

1. As mentioned above, the MIP (3.21)-(3.28) proved to be difficult to handle by commercial solvers like CPLEX or Gurobi. In our experiments we let 1 hour to the solver to improve upon the initial solution. In comparison to this time limit, the computation time required by Step 1 (i.e. the K-medoid algorithm and solving LP ($OPT - \alpha$)) is negligible. On all the approximate uncertainty sets computed for this set of experiments, it took on average less than 1 second and always less than 3 seconds to complete Step 1.
2. We conducted other experiments with larger data-sets (up to $N = 5000$) and there was no significant difference in the quality of the approximation \mathcal{U}_{AC} .




Figure 3.5 – Evolution of the cumulative error as a function of ν for a multivariate Gamma distribution when $N = 1000, K/m = 3$

4.2 Application to optimization problems

The robust problems that we consider in this section are expressed as MIP of the form

$$\max \mathbf{c}^\top \mathbf{x} \quad (3.29)$$

$$\text{s.t. } \mathbf{u}_j^\top \mathbf{x} \leq b_j \quad \forall \mathbf{u}_j \in \mathcal{U}_j(\mathcal{D}) \quad (3.30)$$

$$\mathbf{x} \in \{0, 1\}^m \quad (3.31)$$

where \mathbf{u}_j is the vector of uncertain parameters of the j th constraint, which are restricted to belong to the uncertainty set $\mathcal{U}_j(\mathcal{D})$. In our computational experiments, we distinguish four possible implementations of $\mathcal{U}_j(\mathcal{D})$, all derived from the data available $\mathcal{D} = \{\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)}\}$. Each type of uncertainty set results in a different robust counterpart for each of the problems under study. In the following analysis, we denote by z_{SVC} and z_{AC} the final objective function value attained with the formulation based on the uncertainty set \mathcal{U}_{SVC} and \mathcal{U}_{AC} , respectively. Similarly, we let τ_{SVC} and τ_{AC} be the needed absolute CPLEX CPU-time by both formulations to converge. We compare the respective performance of the two uncertainty sets based on the relative objective value function $\rho = (z_{\text{SVC}} - z_{\text{AC}})/z_{\text{SVC}}$, while the relative CPLEX CPU-time is given by $(\tau_{\text{AC}}/\tau_{\text{SVC}}) \cdot 100\%$. In the lot sizing problem and in the TSP, the uncertainty affects the objective function, hence z_{SVC} and z_{AC} are evaluated *a posteriori* on an independent test set of 10000 new data samples. Note that for the 0-1 Knapsack problem, some realizations of the uncertain parameters may lead to the constraint violation. In that case, we also use a distinct test set of 10000 new samples to compute the achieved fraction of feasible solutions φ_{SVC} and

φ_{AC} using \mathcal{U}_{SVC} and \mathcal{U}_{AC} , respectively. We then express the relative constraint satisfaction as the ratio $(\varphi_{SVC} - \varphi_{AC})/\varphi_{SVC}$.

4.2.1 Robust 0-1 Knapsack Problem

The first problem considered in this section is the robust knapsack problem with uncertain weights.

This problem has been introduced in Bertsimas et al. (2004) and can be expressed with the following IP:

$$\max \mathbf{c}^\top \mathbf{x} \tag{3.32}$$

$$\text{s.t. } \mathbf{w}^\top \mathbf{x} \leq W \quad \forall \mathbf{w} \in \mathcal{U}(\mathcal{D}) \tag{3.33}$$

$$\mathbf{x} \in \{0, 1\}^m \tag{3.34}$$

where W is the total capacity and $\mathbf{w} \in \mathbb{R}^m$ is the vector of weights of the items considered, which is restricted to the uncertainty set $\mathcal{U}(\mathcal{D})$.

We generate instances of different size (i.e., number of items). Random weights are generated using multivariate gamma distributions with random correlations drawn uniformly in $[-1, 1]$, shape and scale parameters drawn uniformly on $[10, 13]$. As in Bertsimas et al. (2004), costs are randomly generated in $[17, 77]$. We considered instances of size $m \in \{10, 20, 30, 40\}$ and generate 10 instances of each size. Note that we solve 8 distinct models on each instance, each of them corresponding to a specific value. $\nu \in \{0.15, 0.25, 0.35, 0.55, 0.75, 0.85, 0.95, 0.99\}$. Figure 3.6 represents, respectively, the difference of objective value ρ and constraint satisfaction between solutions to models based on \mathcal{U}_{SVC} and \mathcal{U}_{AC} for different number of uncertain parameters and different value of the ratio K/m . Regarding the differences in terms of constraint satisfaction, we observe that our approximation leads to a higher level of constraint satisfaction than the SVC-based model although the difference becomes less and less significant as the ratio K/m increases. On the other hand, our approximation displays smaller values of objective function in most cases, where again the difference decreases with the ratio K/m .

Figure 3.7 represents the absolute and relative solving time of \mathcal{U}_{AC} for different values of m and of the K/m ratio. We observe that using this approximate uncertainty set leads to a reduction of the solving time by more than 50% on all models. This phenomenon diminishes as the number of uncertain parameters increases: Since the number of SV used to define \mathcal{U}_{SVC} only depends on ν and the number of data samples while our approximation




Figure 3.6 – Relative constraint satisfaction (left) and objective function (right) difference in percent for the knapsack problem




Figure 3.7 – Absolute (left) and relative (right) solving time of \mathcal{U}_{AC} as a function of m for the robust knapsack problem with different values of the ratio K/m

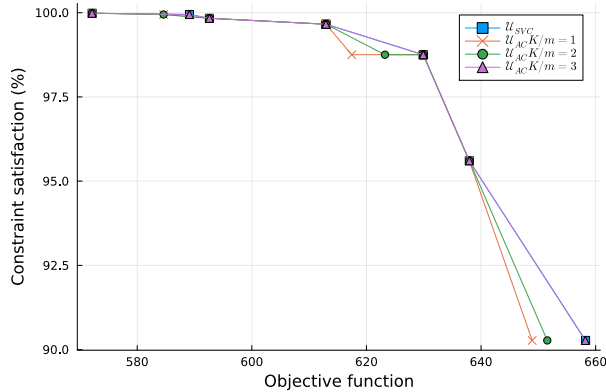


Figure 3.8 – Trade-off between objective value and constraint satisfaction of \mathcal{U}_{SVC} and \mathcal{U}_{AC} for different values of K/m

scales with the dimension m , the relative advantage of the latter is less significant when these two values compare less favorably.

Figure 3.8 represents the constraint satisfaction level in percent as a function of the objective value of the solution obtained with \mathcal{U}_{SVC} and \mathcal{U}_{AC} . The considered instance has 30 items and the different solutions correspond to different values of $\nu \in \{0.15, 0.25, 0.35, 0.45, 0.55, 0.75, 0.85, 0.95\}$. This example illustrates the typical behavior observed for solutions obtained using \mathcal{U}_{SVC} and \mathcal{U}_{AC} in our experiments. The majority of the marked points (corresponding to a specific value of ν) display comparable performance between the original robust solution and the approximated one, even when $K/m = 1$. In particular for $K/m = 3$, both the objective value and the probability of constraint satisfaction are equivalent when using \mathcal{U}_{AC} instead of \mathcal{U}_{SVC} .

Figure 3.9 represent the number of support vectors $|\widehat{\mathcal{S}}|$ used in the approximation \mathcal{U}_{AC} as a function of m and different values of the K/m ratio. When the K/m ratio is less than or equal to 2, we observe that the number of support vectors is less than m for most instances. When $K/m = 3$, the average number of support vectors is less than $1.5m$. Finally, when the ratio $K/m = 4$, the number of support vectors is less than $4m$ but the second phase does not seem to reduce that value significantly within the allowed computation time of 1 hour.

4.2.2 Robust Lot-sizing Problem

In this section we consider an uncapacitated lot-sizing problem (ULSP) similar to the one addressed in Bertsimas and Thiele (2006). The goal is to define the ordering quantities

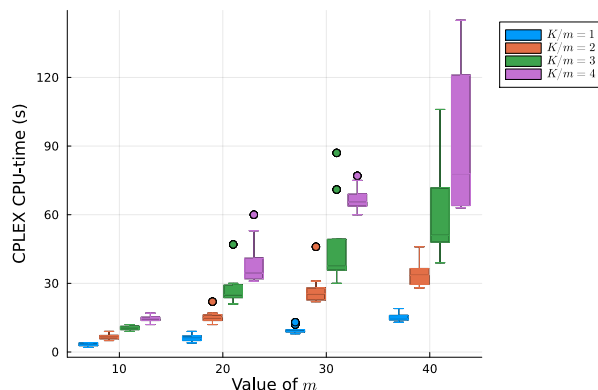


Figure 3.9 – Number of support vectors $|\widehat{\mathcal{S}}|$ used in \mathcal{U}_{AC} as a function of m and different values of K/m

of a single item over a finite and discrete planning horizon of T periods. In each period $t = 1, \dots, T$, we assume an uncertain demand d_t that can be satisfied immediately using units held in the inventory or (fully or partially) backlogged to be served in a subsequent period. Four distinct types of costs are incurred: (1) a per-unit ordering cost c , (2) a fixed ordering cost s that is paid in every period when an order is placed (regardless of the quantity), (3) a per-unit, per-period holding cost h that applies to units physically held in the inventory to satisfy future demands, and (4) a per-unit, per-period backlogging penalty cost b for every unmet unit of demand. The objective is to minimize the total cost incurred by the system of the planning horizon $[T]$.

Bertsimas and Thiele (2006) model this problem with the following MIP:

$$\min \sum_{t=1}^T (cq_t + Sx_t + y_t) \quad (3.35)$$

$$\text{s.t.} \quad y_t \geq h \left(x_0 + \sum_{i=1}^t u_i - \min_{d \in \mathcal{U}_t(\mathcal{D})} \mathbf{d} \mathbf{1}^\top \right) \quad \forall t \in [T] \quad (3.36)$$

$$y_t \geq -b \left(x_0 + \sum_{i=1}^t u_i - \max_{d_i \in \mathcal{U}_t(\mathcal{D})} \mathbf{d} \mathbf{1}^\top \right) \quad \forall t \in [T] \quad (3.37)$$

$$0 \leq q_t \leq Mx_t \quad \forall t \in [T] \quad (3.38)$$

$$x_t \in \{0, 1\} \quad \forall t \in [T] \quad (3.39)$$

$$y_t \geq 0 \quad \forall t \in [T] \quad (3.40)$$

where q_t is the quantity of items ordered in period t , x_t is a binary variable equal to 1 if $q_t > 0$, 0 otherwise, y_t correspond to the holding/backlogging cost of period t , and M is a large constant that can be set to $\sum_{t=1}^T d_t$ w.l.o.g.. In constraints (3.36)-(3.37), $\mathcal{U}_t(\mathcal{D})$ denotes one of the data-driven uncertainty set that describes the uncertain demand parameters $\mathbf{d}_t = [d_1, \dots, d_t]$ that contains the demand values over periods 1 to t .

In our experiments, we consider a horizon length of $T = 14$ and let $S = 3000$, $c = 1$ and $h = 4$. We consider random demand scenarios inspired from the ones introduced in Ben-Tal et al. (2004) that we generate as follows: for all period $t \in [T]$, we first define a seasonal nominal demand

$$\bar{d}_t = 1000 \cdot \left(1 + \frac{1}{2} \sin \left(\frac{2\pi(t-1)}{7} \right) \right)$$

and use it to generate an autocorrelated final demand process $(d_t)_{t \in [T]}$ with autocorrelation factors a_1, \dots, a_k and error term \hat{d}_t :

$$d_t = \bar{d}_t + \sum_{j=1}^k a_j d_{t-j} + \hat{d}_t.$$

The perturbations \hat{d}_t are independent and identically distributed according to either a Gamma or a Normal distribution and satisfy for all $t \in [T]$ the property $P(-0.2 \leq \hat{d}_t \leq 0.2) \geq 0.9$. For each distribution, we generate 5 different data sets and 3 different instances with varying ratio $\sigma = \frac{b}{h} \in \{0.5, 1, 3\}$. We thus obtain 30 instances of the ULSP with stochastic demand. As in the previous experiments, each instance is solved for different values of $\nu \in \{0.15, 0.25, 0.35, 0.55, 0.75, 0.85, 0.95, 0.99\}$ and each solution is evaluated on a data set containing 10,000 data points based on the ratio ρ .

Figures 3.10 show the average and standard deviation of the difference in objective value for several values of the ratios K/m and σ . Notice that we present aggregated results for the two considered types of distribution and all possible values of ν since our experiments did not reveal any significant differences when these parameters vary. For all values of σ , both the average and standard deviation (across all instances) of ρ decrease as K/m increases. Most of the solutions computed based on \mathcal{U}_{AC} are worse than those obtained with \mathcal{U}_{SVC} , although the average gap and standard deviation between the two diminishes quickly with K/m . Again, we observe that a ratio $K/m = 3$ provides a very good approximation. We also observe that the difference in objective value is less consistent than the one obtained for the Robust 0-1 Knapsack problem. This is in

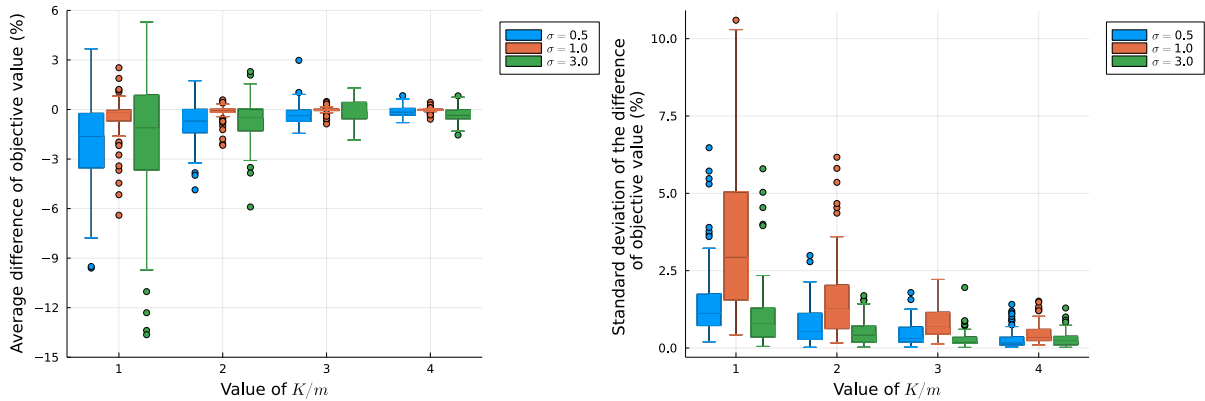


Figure 3.10 – Average (left) and standard deviation (right) of the relative difference between the objective values z_{SVC} and z_{AC} for the robust lot sizing problem.




Figure 3.11 – Solving time of \mathcal{U}_{AC} in percent of the solving time of \mathcal{U}_{SVC} as a function of K/m for different values of σ (Robust lot sizing problem)

contrast with a problem using binary or integer variables, for which a slight change of the boundaries of the uncertainty set is less likely to affect the feasibility of the optimal solution.

Figure 3.11 shows the solving time of \mathcal{U}_{AC} in percent of the solving time of \mathcal{U}_{SVC} for different values of K/m and different ratios σ . We again observe that the approximate uncertainty set succeeds in reducing the time needed to solve the robust counterpart as this ratio never exceeds 4%. For the SVC-based model, the median of CPLEX CPU-time is above 1000 seconds and goes up to 10,000 seconds in the slowest cases. In comparison, it is always shorter than 10 seconds when the robust counterpart uses \mathcal{U}_{AC} even for $K/m = 4$. In addition, we observe a slightly higher computation time for both uncertainty sets when the ratio $\sigma = \frac{b}{h}$ is equal to 0.5.

4.2.3 Robust Traveling Salesman Problem

In this section, we consider a robust traveling salesman problem with uncertain travel costs. Given a set of m cities, the objective is to compute a tour of minimal cost that visits each city once and returns to the starting city. We consider the case where some of the costs c_{ij} incurred when traveling from city i to city j are uncertain. We denote by A and \tilde{A} the sets of arcs with known and uncertain traveling costs, respectively.

For all $i, j \in [m]$, let x_{ij} be a binary variable equal to 1 if the arc (i, j) is used in the solution and $u_i \in [m]$ be an integer variable that gives the position of city i in the corresponding tour. We rely on the well-known Miller-Tucker-Zemlin (MTZ) formulation to express the problem as a MIP:

$$\min Z \tag{3.41}$$

$$\text{s.t.} \quad \sum_{(i,j) \in A} x_{ij}c_{ij} + \max_{\mathcal{c} \in \mathcal{U}(\mathcal{D})} \sum_{(i,j) \in \tilde{A}} x_{ij}c_{ij} \leq Z \tag{3.42}$$

$$\sum_{i \in [m]} x_{ij} = 1 \quad \forall j \in [m] \tag{3.43}$$

$$\sum_{j \in [m]} x_{ij} = 1 \quad \forall i \in [m] \tag{3.44}$$

$$u_i - u_j + mx_{ij} + 1 \leq m \quad 2 \leq i \neq j \leq m \tag{3.45}$$

$$1 \leq u_i \leq m \quad \forall i \in [m] \tag{3.46}$$

$$x_{ij} \in \{0, 1\} \quad \forall i, j \in [m] \tag{3.47}$$

where $\mathcal{U}(\mathcal{D})$ is one of the data-driven uncertainty sets built from historical traveling costs.

We generate instances with $m \in \{15, 20, 25\}$ cities and different partitions A and \tilde{A} . The coordinates of cities are generated randomly and drawn uniformly in the square $[1, 50]^2$. The traveling time on edges belonging to A and the average traveling cost for edges belonging to \tilde{A} are set to the Euclidean distance between the two extremities. To select the set of edges that belong to \tilde{A} , we used the following procedure:

1. Select n cities randomly in $[m]$
2. For each of the n selected cities compute the set \mathcal{C} of the 3 closest cities in terms of Euclidean distance.
3. Add edges (j, i) and (i, j) to \tilde{A} .

In this set of experiments, we used values of $n \in \{5, 7\}$. Traveling costs are such that $c_{ij} = \bar{c}_{ij} + \delta_{ij} \forall (i, j) \in \tilde{A}$ where \bar{c}_{ij} is a fixed traveling cost and δ_{ij} is a random perturbation

following a Gamma (for half of the instances) or a Normal distribution (for the other half). For each value of m , n and each type of distribution, we generate 2 instances for a total of 24 different instances. The perturbations δ_{ij} and δ_{jk} that apply to incoming and outgoing arcs to and from node j are positively correlated, with a correlation coefficient that is drawn uniformly in $[0, 1]$.

As in the previous experiments, we solve each instance for various values $\nu \in \{0.15, 0.25, 0.35, 0.55, 0.75, 0.85, 0.95\}$. We evaluate each solution on a dataset that contains 10,000 data points and compute the relative difference of objective value ρ as in the previous experiments. Following the results obtained in the two previous experiments, we fix the value K/m to 3 since it seems to result in a satisfactory trade-off between accuracy and computational performance.

On all instances considered in these experiments, both models based on \mathcal{U}_{SVC} and \mathcal{U}_{AC} return the same solution for a given input parameter ν . Figure 3.12 shows the percentage of optimal solutions found across all the instances with \mathcal{U}_{SVC} and \mathcal{U}_{AC} . In the case of \mathcal{U}_{SVC} -based robust counterparts, about 55% of the instances are solved optimally within the time limit of 10 hours. The final gap achieved on the remaining instances varies from 0.42% to 34.22%, with an average value of 9.96%. On the other hand, all instances are solved within 2 hours when the robust counterpart is defined using \mathcal{U}_{AC} , with 50% and 75% reaching the optimal solution within 2 and 7 minutes, respectively. It is worth noting that over the subset of instances solved to optimality with both uncertainty sets, the solutions found were always identical.

5 Conclusion

In this paper we presented a method to build approximations of the SVC-based uncertainty sets introduced in Shang et al. (2017) that preserves their polyhedral structure while reducing significantly the number of additional variables and constraints in the resulting robust counterpart. Our numerical experiments attest that the obtained solutions with our approximate sets are comparable but require much less computational time (from less than 1% to less than 50% in a few isolated instances) than the ones derived from the original SVC-based uncertainty sets. A direct implication is that our procedure extends the range of possible applications of the latter to problems of greater sizes. In addition, any optimization procedure that relies on the polyhedral structure of the uncertainty set may benefit from this approximation. From the practitioner standpoint, the parameter K

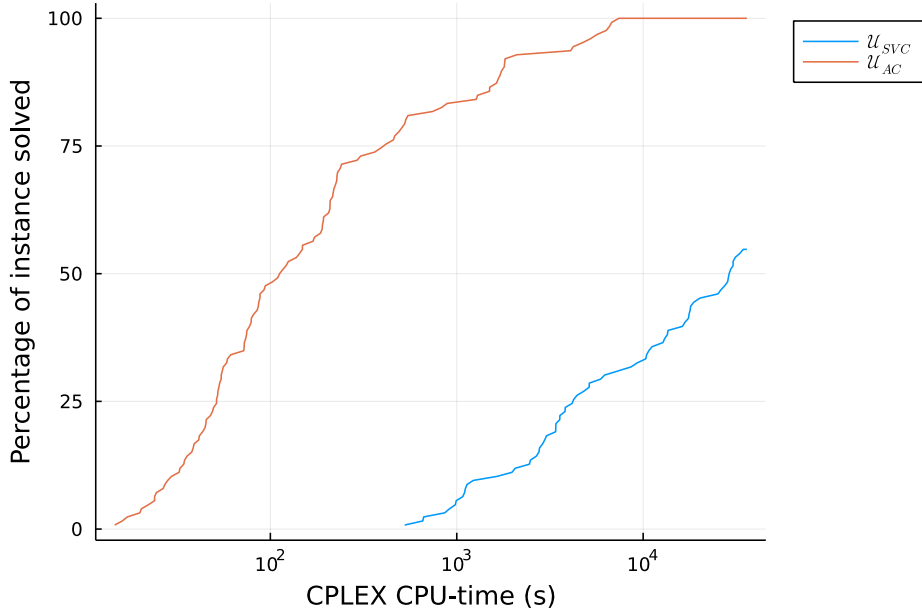


Figure 3.12 – Proportion of TSP instances solved for \mathcal{U}_{SVC} and \mathcal{U}_{AC} $K/m = 3$

used in the first phase of our method defines a trade-off between the quality of the approximation and the reduction of the computational time. We test the influence of K on these two aspects and provide some insights on appropriate values for this parameter. While tuning its value may improve the performance in some specific applications, we observe that $K \approx 3m$, where m is the number of uncertain parameters, result in a good trade-off in general. Finally, we notice that our approximation method is almost insensitive to the type of uncertainty distribution in our experiments and to the value of the parameter ν of the SVC algorithm. While our approximation does not provide the in-sample coverage guarantee of the SVC-based uncertainty set (i.e. \mathcal{U}_{SVC} contains at least a portion $(1 - \nu)$ of \mathcal{D}), our results show that for a sufficiently large value of K , the out-of-sample coverage of both uncertainty sets are similar and the approximation does not significantly impact the level of conservatism of the robust solution.

Despite these promising numerical results, we note that the use of a K -medoid procedure in its first phase makes our approach heuristic. Alternatively, one could think of jointly optimizing the position and weight coefficient of the representatives of the support vectors in a single phase. This results into large Mixed-Integer Non-Linear Programs that are essentially more challenging and require new ideas to compute solutions efficiently. Finally, notice that we rely on data points to approximate the original polyhedron ob-

tained by SVC. It will be interesting to study whether similar ideas are applicable to other data-driven uncertainty sets such as the ones developed in Bertsimas et al. (2017) or Ning and You (2018b).

A DATA-DRIVEN ROBUST OPTIMIZATION APPROACH TO SUPPLY PLANNING

1 Introduction

In this chapter, we study a problem derived from a practical case encountered in the aircraft industry. We consider the case of a single assembly line which manufactures several products in order to satisfy the needs that are planned over a discrete and finite planning horizon. Each final product has its own *Bill Of Material (BOM)* that involves different components. The components are stored in a remote warehouse and they are picked and delivered to the production line by a *Third-party Logistic Provider (TPL)* following the manufacturer's orders. In a sense, this problem is closely related to the class of assemble-to-order systems, which are often encountered in the inventory control literature. The need to define both component replenishment and production policies in the presence of uncertain factors makes assemble-to-order systems particularly challenging. Large surveys on this topic have been presented in Song and Zipkin (2003) and more recently in Atan et al. (2017).

In this chapter, we investigate the planning of picking orders when the available information related to picking times at the TPL is uncertain in a context where the time available for these picking operations in a given period is bounded by a capacity. We model this problem as an extension of the *Multi-Item Capacitated Lot Sizing Problem (CLSP)* and propose a robust optimization approach to integrate the picking time uncertainties. These picking time uncertainties are modeled as setup time uncertainties.

The capacitated lot sizing model is well-suited for production planning problems in which (at least) one of the resources used in the production process is available in limited quantity, imposing a capacity bound on the volume of goods produced during a period. In practice, this can correspond to a maximum quantity of raw material, a maximum quan-

tity of produced items or a time capacity as in our case. The CLSP has been extensively studied in the literature because of its relationship to a wide variety of practical problems arising in the industry. Several papers have reviewed the existing works on CLSP and the different extensions of these models (Karimi and Fatemi (2003); Gicquel et al. (2008)). The first application of such models in an uncertain context appears in Brandimarte (2006) who considers uncertain demand and proposes a multi-stage mixed-integer stochastic programming model. Since then, most of the existing literature on CLSP under uncertainty have focused on uncertain demand, including several RO approaches. Some extensions of this basic setting to more complex versions of CLSP exist, see e.g. Coniglio et al. (2018) who introduces two robust models for a CLSP with storage deterioration or Abdel-Aal (2019) for a problem with setup times that allows overtime decisions. The numerical experiments conducted in these studies show that the computation time needed to solve robust formulations of a CLSP increases significantly with the number of periods in the planning horizon. To the best of our knowledge however, all existing works on robust CLSP consider a budget-based uncertainty set and no study has focused on uncertain setup times yet.

The contributions of this chapter are the following:

1. We provide a MIP formulation of a practical variation of a CLSP arising in the aircraft industry.
2. We derive two robust formulations of the problem based on polyhedral uncertainty sets. The first one uses the classical budget-based approach of Bertsimas and Sim (2003, 2004) and the second one relies on the SVC-based uncertainty set of Shang et al. (2017), or the approximate version introduced in the previous chapter.
3. We evaluate and compare the performances of the three models on simulated data.

The remainder of this chapter is organized as follows: Section 2 presents the supply planning problem that motivates this study, along with a deterministic mathematical formulation. In Section 3 we derive two tractable robust formulations of the problem based on the budget-based and the SVC-based uncertainty sets, respectively. Finally, Section 4 presents the experimental protocol and the numerical results obtained with both robust formulations, while Section 5 contains our concluding remarks and some perspective for future research directions.

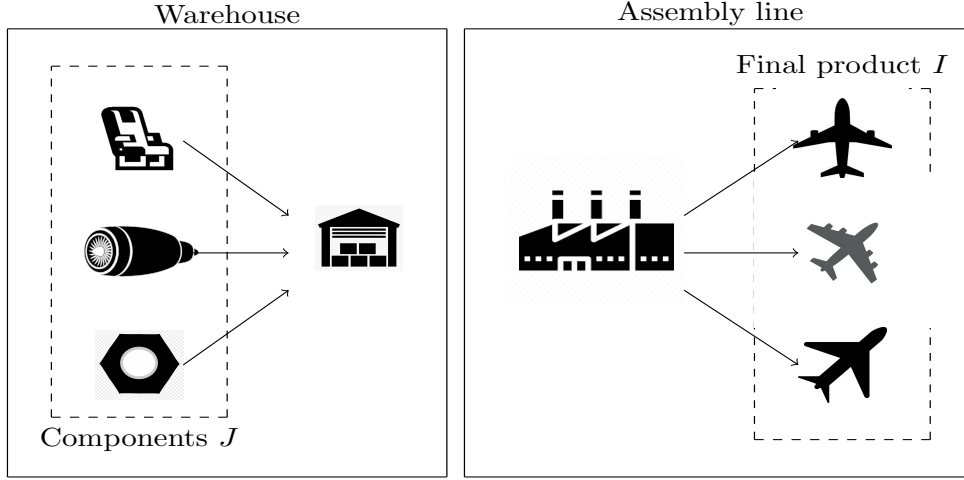


Figure 4.1 – Diagram of the problem

2 Problem statement and deterministic formulation

We study the case of a single assembly line that combines different components into several final products over a finite and discrete planning horizon of T periods. Let $[I]$ and $[J]$ the set of products and components, respectively. In each period $t = 1, \dots, T$, the system faces a demand d_{it} for product $i \in [I]$ that is assembled on demand (i.e. it is impossible to manufacture a product in advance, then store it to satisfy a future demand). For all product $i \in [I]$ and component $j \in [J]$, let r_{ij} be the number of components of type j required to produce one unit of product i .

All components are assumed to be available at all time in a remote warehouse managed by a TPL provider, who is in charge of delivering the assembly line based on its orders for components, as represented on Figure 4.1. For each picking operation for component $j \in [J]$, the quantity collected by an operator is limited to a maximum batch size of m_j units. Note that several batches of the same components can be scheduled in the same period. The picking time for a particular batch of component j of size q_j can be divided into two main parts. First, we consider a fixed time p_j that corresponds to the travel time between the shipping point and the zone where components of type j are stored, which is spent regardless of the quantity that is transported. In addition, we consider a per-unit picking time τ_j . Finally for all period $t \in [T]$, the total picking time spent by the TPL provider to collect all the components delivered to the assembly line in period t cannot exceed a maximum work capacity C_t .

Any demand for product i that is not satisfied immediately is backlogged until the

corresponding product is assembled in a subsequent period. Each such unit incurs a per-unit backlogging penalty cost b_i for each period of delay. Whenever a component j is available on the assembly line but is not immediately used to manufacture an end product, it disturbs the production process by interfering with people and other goods moving nearby. We model this situation with a per-unit, per-period obstruction cost o_j . The problem consists in planning the quantity of each component delivered to the assembly line by the TPL in each period such that the sum of the obstruction and backlogging costs is minimized. We let $q_{jt} \in \mathbb{R}_+$ be the quantity of components $j \in [J]$ that is ordered by the assembly line for a delivery in period t . These quantities induce the number of distinct picking operations of components $j \in [J]$ performed during period t represented as the integer variable $x_{jt} \in \mathbb{N}$. For any $t \in [T]$, the variable u_{it} denotes the number of final product $i \in [I]$ produced during period t and O_t corresponds to the total obstruction cost for period t . At the end of each period $t \in [T]$ the stock level s_{jt} of each component $j \in [J]$ is given by:

$$s_{jt} = s_{j0} + \sum_{k=1}^t (q_{kt} - \sum_{i \in [I]} r_{ij} u_{ik}) \quad (4.1)$$

We formulate the problem with the following MIP:

$$\min \sum_{t=1}^T \left(O_t + \sum_{i \in [I]} b_i \sum_{k=1}^t (d_{ik} - u_{ik}) \right) \quad (4.2)$$

$$\text{s.t. } O_t \geq \sum_{j \in [J]} o_j \left(s_{j0} + \sum_{k=1}^t (q_{kt} - \sum_{i \in [I]} r_{ij} u_{ik}) \right) \quad \forall t \in [T] \quad (4.3)$$

$$\sum_{k=1}^t u_{ik} \leq \sum_{k=1}^t d_{ik} \quad \forall i \in [I], \forall t \in [T] \quad (4.4)$$

$$\sum_{j \in [J]} (p_j x_{jt} + \tau_j q_{jt}) \leq C_t \quad \forall t \in [T] \quad (4.5)$$

$$x_{jt} \geq \lceil q_{jt}/m_j \rceil \quad \forall j \in [J], \forall t \in [T] \quad (4.6)$$

$$q_{jt}, O_t \in \mathbb{R}_+ \quad \forall j \in [J], \forall t \in [T] \quad (4.7)$$

$$u_{it}, x_{jt} \in \mathbb{N} \quad \forall i \in [I], \forall j \in [J], \forall t \in [T] \quad (4.8)$$

The objective (4.2) aims at minimizing the total cost incurred over the planning horizon. Note that the backlogging term is always positive due to constraints (4.4), which ensure that the production plan up to a given period cannot anticipate demands in future periods. Constraints (4.3) compute the total obstruction cost for period $t \in [T]$. Con-

straints (4.5) ensure that the capacity of the logistic provider is respected by the picking plan during a given period $t \in [T]$. Finally, constraints (4.6) link the quantity of component ordered q_{jt} and the number of batches x_{jt} , while constraints (4.7) and (4.8) define the domain of the decision variables.

3 Robust models for uncertain setup times

In practice, the manufacturer often has either incomplete or imprecise knowledge on the picking times. As a consequence, some combinations of his orders may exceed the picking capacity of the TPL provider, forcing the latter to postpone some operations to subsequent periods. This delay in the delivery of some components is susceptible to induce two types of inefficiencies on the assembly line as missing components (i) prevent the manufacturer to assemble some of the products, leading to backlogging penalty costs and (ii) leave the other components in the BOM on the border of the line, disturbing other production operations. This double negative effect strongly incentivizes the planner to protect his decisions against uncertain setup times. In what follows, we assume that we have access to a set of historical setup times $\mathcal{P} = \{\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(N)}\}$ in order to tune the different parameters of the uncertainty sets.

As the uncertain parameters of this problem are the setup picking time \mathbf{p} , its robust formulation only affects the capacity constraints (4.5). Given uncertainty set \mathcal{U}_t for a period $t \in [T]$, the robust counterpart of this constraint is expressed as follows:

$$\sum_{j \in [J]} p_j x_{jt} + \tau_j q_{jt} \leq C_t \quad \forall \mathbf{p} \in \mathcal{U}_t, \forall t \in [T] \quad (4.9)$$

which is equivalent to the following constraint (see Chapter 1):

$$\max_{\mathbf{p} \in \mathcal{U}_t} \left\{ \sum_{j \in [J]} p_j x_{jt} \right\} + \sum_{j \in [J]} \tau_j q_{jt} \leq C_t, \quad \forall t \in [T] \quad (4.10)$$

3.1 Budget-based robust formulation

As a first robust approach, we consider the classical uncertainty set of Bertsimas and Sim (2003, 2004) already defined in previous chapter and defined by :

$$\mathcal{U}_t^\Gamma = \{\mathbf{p} \mid p_{jt} = \bar{p}_j + \hat{p}_j z_{jt} \mid \|\mathbf{z}\|_1 \leq \Gamma_t, \forall j \in [J]\} \quad (4.11)$$

, where \bar{p}_j is the average value of p_j evaluated on \mathcal{P} and \hat{p}_j is a maximal deviation. The scaled deviation coefficient $z_{jt} = (p_j - \bar{p}_j)/\hat{p}_j \in [-1, 1]$ ensures that $p_{jt} \in [\bar{p}_j - \hat{p}_j, \bar{p}_j + \hat{p}_j]$ for all components $j \in [J]$ and period $t \in [T]$ and that the total scaled deviation of \mathbf{p} never exceeds the budgets of uncertainty $\mathbf{\Gamma} = (\Gamma_1, \dots, \Gamma_T)$.

Proposition 1. The robust counterpart of problem (4.2)-(4.8) based on the uncertainty sets \mathcal{U}_t^Γ , $t \in [T]$ is equivalent to the following MIP, after introducing the additional variables $\mathbf{v} \in \mathbb{R}_+^T$ and $\mathbf{w} \in \mathbb{R}_+^{J \times T}$:

$$\min \sum_{t=1}^T \left(O_t + \sum_{i \in [I]} b_i \sum_{k=1}^t d_{ik} - u_{ik} \right) \quad (4.12)$$

$$\text{s.t. (4.3), (4.4), (4.6) - (4.8)} \quad (4.13)$$

$$\sum_{j \in [J]} \bar{p}_j x_{jt} + \tau_j q_{jt} + v_t \Gamma_t + \sum_{j \in [J]} w_{jt} \leq C_t \quad \forall t \in [T] \quad (4.14)$$

$$\hat{p}_j \leq v_t + w_{jt} \quad \forall j \in [J], \forall t \in [T] \quad (4.15)$$

$$v_t, w_{jt} \in \mathbb{R}_+ \quad \forall j \in [J], \forall t \in [T] \quad (4.16)$$

Proof. We formulate the robust capacity constraint (4.10) as:

$$\sum_{j \in [J]} \bar{p}_j x_{jt} + \max_{\substack{\mathbf{z}_t \in [0, 1]^{|J|} \\ \|\mathbf{z}_t\|_1 \leq \Gamma_t}} \left\{ \sum_{j \in [J]} \hat{p}_j x_{jt} z_{jt} \right\} + \sum_{j \in [J]} \tau_j q_{jt} \leq C_t, \quad \forall t \in [T] \quad (4.17)$$

Considering an optimal solution \mathbf{x}^* to the problem, for each $t \in [T]$, the inner maximization problem is equivalent to a linear optimization problem. As values in \mathbf{x} are all positive or null, we can restrict $\mathbf{z}_t \in [0, 1]^{|J|}$ and thus replace the l_1 -norm $\|\mathbf{z}_t\|_1$ by a simple sum $\sum_{j \in [J]} z_{jt}$. This gives us the following equivalent LP:

$$\begin{aligned} & \max \sum_{j \in [J]} \hat{p}_j x_{jt}^* z_{jt} \\ & \text{s.t. } \sum_{j \in [J]} z_{jt} \leq \Gamma_t \\ & \quad 0 \leq z_{jt} \leq 1, \quad \forall j \in [J]. \end{aligned} \quad (4.18)$$

After introducing dual variables v_t and w_{jt} for $j \in [J]$, we obtain the dual formulation of

the problem:

$$\begin{aligned}
\min \quad & v_t \Gamma_t + \sum_{j \in [J]}^n w_{jt} \\
\text{s.t.} \quad & v_t + w_{jt} \geq \hat{p}_j x_{jt}^* \quad \forall j \in [J] \\
& v_t \geq 0 \\
& w_{ij} \geq 0 \quad \forall j \in [J]
\end{aligned} \tag{4.19}$$

Since the primal problem (4.18) is feasible and bounded, its dual problem (4.19) is also feasible and bounded and their optimal values coincide. Therefore, we can replace the inner maximisation problem of period $t \in [T]$ by the objective function of (4.19), with the associated constraints. This gives us the (linear) robust counterpart (4.12)-(4.16). \square

3.2 SVC-based robust formulation

As a second robust formulation, we propose to use the SVC-based uncertainty set of Shang et al. (2017) presented in the previous chapter. Let us recall that the SVC approach applied on the set of historical picking times \mathcal{P} define a subset \mathcal{S} of data points, called *support vectors* (SV), that lie on or outside of the boundary of the uncertainty set. A subset $\mathcal{B} \subseteq \mathcal{S}$ of these points, referred to as *boundary support vectors* (BSV), are located exactly on the boundary of the uncertainty set. The definition of this uncertainty set ensures that $\mathcal{B} \subseteq \mathcal{S} \subseteq \mathcal{D}$ and $|\mathcal{S} \setminus \mathcal{B}| \leq N\nu \leq |\mathcal{S}|$. In addition, the SVC algorithm also computes for each historical data point $l \in [|\mathcal{D}|]$ a value $\alpha_l \in [0, 1/\nu N]$ such that $\alpha_l = 0$ iff $l \in \mathcal{D} \setminus \mathcal{S}$, $\alpha_l = 1/\nu N$ for all $l \in \mathcal{S} \setminus \mathcal{B}$ and $0 < \alpha_l < 1/\nu N$ when $l \in \mathcal{B}$. Let $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_N]$ be the vector of obtained coefficients, Shang et al. (2017) use $\boldsymbol{\alpha}$ to define the following data-driven uncertainty set:

$$\mathcal{U}_t^{SVC} = \left\{ \mathbf{p} \mid \sum_{l \in \mathcal{S}} \alpha_l \|\mathbf{Q}(\mathbf{p} - \mathbf{p}^{(l)})\|_1 \leq \theta \right\} \tag{4.20}$$

where $\mathbf{Q} = \boldsymbol{\Sigma}^{-\frac{1}{2}}$ with $\boldsymbol{\Sigma}$ the covariance matrix of \mathbf{p} and $\theta = \min_{l' \in \mathcal{B}} \left(\sum_{l \in \mathcal{S}} \alpha_l \|\mathbf{Q}(\mathbf{u}^{(l')} - \mathbf{u}^{(l)})\|_1 \right)$.

Proposition 2. Using for all $\forall t \in [T]$ the additional variables $\boldsymbol{\lambda}_t, \boldsymbol{\mu}_t \in \mathbb{R}_+^{|J| \times |S|}$ and $\eta_t \in \mathbb{R}_+$ leads to the following MIP formulation for the robust counterpart of problem (4.2)-(4.8) based on the uncertainty sets \mathcal{U}_t^{SVC} :

$$\min \sum_{t \in [T]} \left(O_t + \sum_{i \in [I]} b_i \sum_{k=1}^t (d_{ik} - u_{ik}) \right) \quad (4.21)$$

$$\text{s.t. (4.3) - (4.8)} \quad (4.22)$$

$$\sum_{l \in \mathcal{S}} (\boldsymbol{\mu}_{tl} - \boldsymbol{\lambda}_{tl})^T \mathbf{Q} \mathbf{u}^{(l)} + \eta_t \theta + \sum_{j \in [J]} \tau_j q_{jt} \leq C_t \quad \forall t \in [T] \quad (4.23)$$

$$\sum_{l \in \mathcal{S}} \mathbf{Q} (\boldsymbol{\lambda}_{tl} - \boldsymbol{\mu}_{tl}) + \mathbf{x}_t = \mathbf{0} \quad \forall t \in [T] \quad (4.24)$$

$$\boldsymbol{\lambda}_{tl} + \boldsymbol{\mu}_{tl} = \eta_t \alpha_l \mathbf{1} \quad \forall l \in \mathcal{S}, \forall t \in [T] \quad (4.25)$$

$$\lambda_{tlj}, \mu_{tlj}, \eta_t \in \mathbb{R}_+ \quad \forall j \in [J], \forall l \in \mathcal{S}, \forall t \in [T] \quad (4.26)$$

Proof. Considering an optimal solution \mathbf{x}^* to the problem, the inner maximization problem of constraint (4.10) for a given $t \in [T]$ is equivalent to the following linear optimization problem:

$$\begin{aligned} & \max \mathbf{p}^T \mathbf{x}_t \\ & \text{s.t. } \sum_{l \in \mathcal{S}} \alpha_l \mathbf{1}^T \mathbf{v}_l \leq \theta \\ & \quad -\mathbf{v}_l \leq \mathbf{Q}(\mathbf{p} - \mathbf{p}^{(l)}) \leq \mathbf{v}_l \quad \forall l \in \mathcal{S} \end{aligned} \quad (4.27)$$

We introduce the dual variable $\boldsymbol{\lambda}_t, \boldsymbol{\mu}_t \in \mathbb{R}_+^{|J| \times |\mathcal{S}|}$ and $\eta_t \in \mathbb{R}_+$, $\forall t \in [T]$ to define the dual of the problem:

$$\begin{aligned} & \min \sum_{l \in \mathcal{S}} (\boldsymbol{\mu}_{tl} - \boldsymbol{\lambda}_{tl})^T \mathbf{Q} \mathbf{u}^{(l)} + \eta_t \theta \\ & \text{s.t. } \sum_{l \in \mathcal{S}} \mathbf{Q} (\boldsymbol{\lambda}_{tl} - \boldsymbol{\mu}_{tl}) + \mathbf{x}_t = \mathbf{0} \quad \forall t \in [T] \end{aligned} \quad (4.28)$$

$$\lambda_{tlj}, \mu_{tlj}, \eta_t \in \mathbb{R}_+ \quad \forall j \in [J], \forall l \in \mathcal{S}, \forall t \in [T]$$

Since the primal problem (4.27) is feasible and bounded, its dual problem (4.28) is also feasible and bounded and their optimal values coincide. Therefore, we can replace the inner maximisation problem of period $t \in [T]$ by the objective function of (4.28), with the associated constraints. The (linear) robust counterpart (4.21)-(4.25) follows. \square

4 Experimental results

The aim of this section is to evaluate the performances of the proposed optimization models, both in terms of their practical computational tractability and the quality of the solutions obtained. To do this, we generated a set of instances with various configuration (i.e. number of products, number of components, number of periods). In what follows, we will use the notations *Det* to denote the deterministic models or their solutions, \mathcal{U}^{Γ} to denote the budget based robust models or their solutions and \mathcal{U}^{SVC} and \mathcal{U}^{AC} to denote the SVC-based models and its approximation introduced in the previous chapter, respectively. All models are built from historical data sets \mathcal{P} of 1000 data points.

Instance generation We first describe several rules that we apply to generate the instances used to evaluate the performances of the different models presented above. The specific assumptions we consider are listed below, with the objective to reflect the original use case from the aircraft industry that motivates this work: r_{ij} is chosen randomly from an interval $[r_j, \bar{r}_j]$, the smallest interval is $[2,4]$ and the widest is $[10,25]$. Each product i requires a subset $\mathcal{J}_i \subseteq [J]$ of components. The obstruction cost o_j of component j is chosen randomly in $[0, 1]$, while $b_i = \rho \sum_{j \in [J]} r_{ij} o_j$ is a multiple of the sum of the obstruction cost of its components. Here ρ represents a ratio between the penalty incurred for backloging a demand for product i versus the maximum obstruction cost of its components. In our numerical experiments, we use $\rho = 5$, but this value can be adjusted depending on the context. In order to primarily focus on the impact of the setup times uncertainty, we disregard the linear part of picking times and set $\tau_j = 0$ and $m_j = 3 \max_{i \in [I]} r_{ij}$ for all $j \in [J]$. Demands are deterministic with values drawn from a uniform distribution whose range vary among products. We compute the capacity of the TPL C_t from the average time needed to pick the required components in each period, that we increase by a given percentage to ensure that it is not always saturated.

Setup time distribution We assume that the TPL provider organizes the storage of the components in order to optimize the picking operations. Specifically, components are stored in such a way that their accessibility improves with their order frequency. We consider three types of components and separate them based on their storage area. We assume that both the mean and the variability of the setup times decrease with component accessibility. In our instances, we thus generate setup times using Gamma distributions with three different shape parameters, one for each type of component. We also assume

that picking times are influenced by the behaviour of the operators and can be correlated if operators groups orders of components or use particular picking routes. Thus, we generate random correlations among picking time of components in the same storage area.

In practice, setting the budget parameter Γ_t is hard and the performance obtained for a given parameter value heavily depends on the characteristics of the instance. In order to compare both models, we set Γ_t to cover a certain portion of \mathcal{P} , by iteratively increasing its value until we reach the desired coverage.

Evaluating solutions As we are considering uncertain parameters, we can not compare solutions only based on the objective value of the optimization models directly. In addition, we cannot rely on a simple evaluation of the constraint satisfaction as the main objective of this problem is to increase the satisfaction of the final product demand with a reasonable increase of the amount of components held on the side of the assembly line. Thus in order to evaluate the quality of solutions obtained, we rely on a simulation process following four steps:

1. Generate a set of 10^4 setup time scenarios \mathcal{S} following the same distribution than \mathcal{P} . Each scenario is composed of one setup time for each component $j \in [J]$ and each period $t = 1, \dots, T$. This unique set will be used to evaluate all the solutions for a given instance.
2. For a given solution \mathbf{x}^* , \mathbf{q}^* simulate the picking operations at each period with each scenario in \mathcal{S} to obtain the real quantities of components delivered on the assembly line. This simulation take into account the reported quantity if the capacity of a given period is exceeded.
3. Compute the optimal production plan for the quantity delivered by the simulation.
4. Evaluate the resulting production plan, the costs, the satisfaction of the demand and the amount of components held next to the assembly line.

As in Chapter 3, all models are implemented in Julia 1.9 with JuMP and solved with CPLEX 20.1 using a single thread on a Linux, Ubuntu 20.04.2 LTS, equipped with an Intel Xeon Gold 6230 clocked at 2.10GHz.

4.1 Scalability of the SVC-based model

Recall from the previous chapter that the number of variables and constraints of the SVC-based uncertainty set is proportional to the number of uncertain parameters of

the problem and the number of data point in the training sample. This is particularly problematic in the problem considered in this chapter, which includes multiple constraints involving uncertain parameters. In this first set of experiments we consider instances with 2 final products, 5 components and a planning horizon of $T = 7$ periods. Table 4.1 shows the average and 0.95 quantile of solution costs, as well as the average computation time of both models when the uncertainty sets are tuned to cover different portion of \mathcal{P} . As we consider a set of small instances, we restricted the computation time to 1 hour for both models and report the remaining CPLEX gap (in percent) when no optimal solution is found within the time limit. In this table, we partition the results into three cases based on the desired coverage of \mathcal{P} :

- Case 1 When the coverage is 0.15 or 0.35 both models are solved to optimality and provide comparable performances in term of costs. However, using the approximate uncertainty set \mathcal{U}^{AC} reduces the computation time by more 99%.
- Case 2 When the coverage is 0.55, the original SVC-based formulation presents an average CPLEX optimality gap of 3.54% after 1 hour while our approximation converges to an optimal solution within 10.47 seconds on average. The impact of this gap on the solutions seems significant as the average cost is increased by more than 6% and the 0.95 quantile by more than 13% when using \mathcal{U}^{SVC} .
- Case 3 When the coverage is high (0.75 or 0.95) both models are extremely conservative, leading to solutions with high average costs that are nearly unaffected by the uncertainty as the difference between the average value and the 0.95 quantile suggests. Moreover, this comes with a significant increase of the CPLEX CPU time that is not influenced as much by the size of the models and affects \mathcal{U}^{AC} in a similar fashion as \mathcal{U}^{SVC} .

We restricted the CPLEX CPU-time to 1 hour as our objective is to be able to obtain good solutions in a reasonable amount of time. The results show that both models lead to comparable solutions in term of costs but with a significant reduction of the CPLEX CPU-time, in particular on small instances for which both \mathcal{U}^{SVC} and \mathcal{U}^{AC} find an optimal solution on all instances. This suggests that \mathcal{U}^{AC} may be applied to larger instances to obtain solutions in a reasonable amount of time. They illustrate that the approximate uncertainty set designed in Chapter 3 gives an opportunity to the practitioners to use the information they can extract from the data available through an unsupervised learning procedure without the computational burden that usually comes with it.

Table 4.1 – Comparison of \mathcal{U}^{SVC} and the approximation \mathcal{U}^{AC} on small instances with 2 final products, 5 components and time horizon $T = 7$ periods.

Coverage	\mathcal{U}^{SVC}				\mathcal{U}^{AC}			
	Average	Q0.95	time (s)	gap	Average	Q0.95	time (s)	gap
0.15	1487	4050	543.93	0	1485	4015	1.65	0.0
0.35	1324	3309	1112.9	0	1319	3275	6.41	0.0
0.55	1329	2727	3600.88	3.54	1236	2362	10.47	0.0
0.75	3508	3842	3600.88	42.5	3164	3612	952.41	0.0
0.95	9451	9446	3291.47	0	9451	9446	2199.95	0.0

4.2 Comparing \mathcal{U}^Γ and \mathcal{U}^{AC}

In all the robust models we present, we use the portion of data coverage as a parameter to define the degree of robustness of the solution. The impact of this parameter on the cost of solutions is represented in Figure 4.2, which compares the evolution of the average (4.2a) and 0.95-quantile (4.2b) costs obtained by both \mathcal{U}^Γ and \mathcal{U}^{AC} when the portion of \mathcal{P} covered increases. The instance includes 5 final products, 10 components and $T = 14$ periods. For both models we observe that increasing the portion of historical data covered significantly reduces the total cost in extreme cases. Figure 4.2b shows that when the coverage increases, the 0.95-quantile of costs is reduced from more than 25000 to less than 17500 for \mathcal{U}^Γ and to less than 14000 for \mathcal{U}^{AC} . On the other hand, this improvement for extreme cases comes with a deterioration of the average solution cost, as shown in Figure 4.2a. In the case of \mathcal{U}^Γ the average cost decreases from 10600 to 9900 as the coverage grows up to 0.25, but increases with the coverage afterward. While we observe the same behavior when using \mathcal{U}^{AC} instead, the decrease in average cost (from 9900 to 8000) continues until a coverage value of 0.35, but starts going up beyond that point. For all coverage values tested, \mathcal{U}^{AC} clearly outperforms \mathcal{U}^Γ by providing lower average and extreme costs. While the two 0.95-quantile average costs are comparable for coverage values between 0.2 and 0.35, this comes at the expense of a steep raise in average cost for \mathcal{U}^Γ . We observe a similar phenomenon affecting \mathcal{U}^{AC} , but starting at higher values of coverage and with milder increasing slope of average costs.

Table 4.2 details the obstruction and backloging costs obtained by both methods on a set of 10 instances with 5 final products, 10 components and a planning horizon of $T = 7$ periods when the coverage of \mathcal{P} varies from 0.1 to 0.3. The first observation

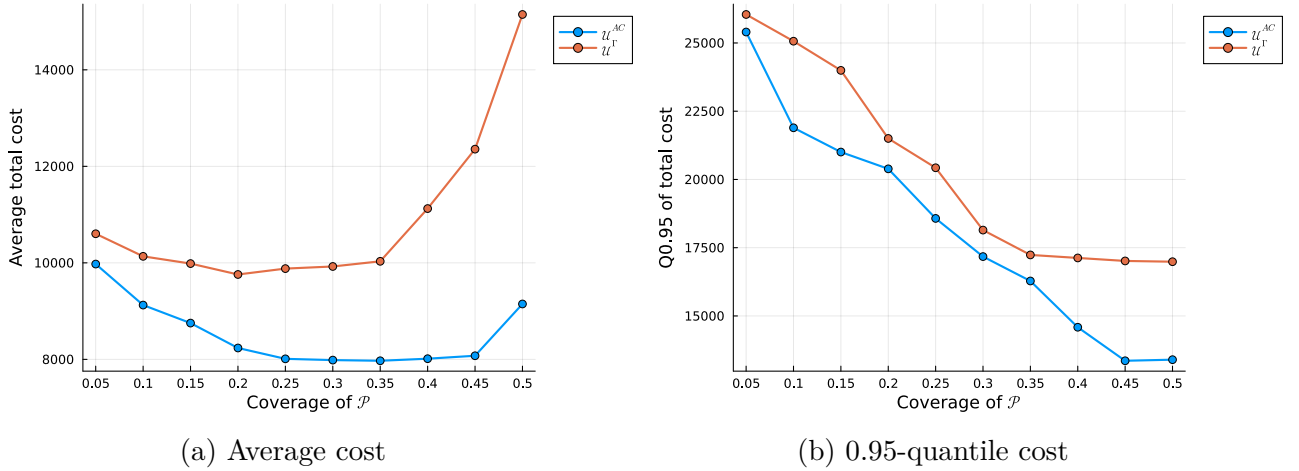


Figure 4.2 – Evolution of the average cost and 0.95-quantile obtained with \mathcal{U}^{Γ} and \mathcal{U}^{AC} on an instance with 5 components 10 products and planning horizon $T = 14$ periods.

is that \mathcal{U}^{AC} simultaneously provides lower obstruction and backlogging costs when the coverage is 0.1 and 0.2, which explains the difference observed on Figure 4.2. When the coverage is 0.3, \mathcal{U}^{Γ} achieves a greater reduction of the obstruction cost at the expense of a significant increase in backlogging costs, which suggests that its representation of possible parameter values includes unlikely scenarios which do not lead to anticipate enough the need for some components to satisfy the demand on time. In comparison, \mathcal{U}^{AC} provides a more significant reduction of backlogging costs with higher obstruction costs and lead to lower costs overall. In the last line of 4.2, when the coverage is 0.5, both approaches provides over-conservative solution increasing simultaneously the obstruction and backlogging costs. This conservatism is emphasized by the static decision context and may be reduced in a more dynamic one where the decisions can be adjusted to the realization of uncertain parameters over time. Table 4.3 shows the evolution of the average CPLEX CPU-time and the CPLEX optimality gap after 1 hour for the same set of instances as the coverage of \mathcal{P} increases. We observe that for both models, increasing the coverage increases the CPLEX CPU-time, for example the solutions of \mathcal{U}^{AC} are obtained in 2 seconds on average when the coverage is 0.1 and in 193 seconds when the coverage is 0.3. In addition, when the coverage is high (≥ 0.5) the CPLEX CPU-time increase significantly for both models (2200 seconds for \mathcal{U}^{Γ}) and the optimal solution is never obtained in 1 hour for \mathcal{U}^{AC} .

Table 4.4 compare the obstruction and backlogging costs obtained by both methods on a set of 10 instances with 10 final products and 20 components for different coverage

Table 4.2 – Comparison of \mathcal{U}^Γ and the approximation \mathcal{U}^{AC} on a set of 10 instances with 5 final products, 10 components and time horizon $T = 7$ periods.

Coverage	\mathcal{U}^Γ				\mathcal{U}^{AC}			
	Obstruction		Backlogging		Obstruction		Backlogging	
	Average	Q0.95	Average	Q0.95	Average	Q0.95	Average	Q0.95
0.1	7022	8137	4286	16852	−10.89%	−13.25%	−17.49%	−10.72%
0.2	8677	9616	3251	14600	−17.76%	−17.73%	−6.21%	−9.76%
0.3	7985	8654	6064	12502	+34.41%	+30.03	−71.82	−34.32
0.5	9125	10189	13116	13782	−10.17%	−11.31	−24.41	−22.27

Table 4.3 – Average CPLEX CPU-time and optimality gap of \mathcal{U}^Γ and the approximation \mathcal{U}^{AC} on a set of 10 instances with 5 final products, 10 components and time horizon $T = 7$ periods.

Coverage	\mathcal{U}^Γ		\mathcal{U}^{AC}	
	time (s)	gap (%)	time (s)	gap (%)
0.1	1	0	2	0
0.2	3	0	12	0
0.3	9	0	193	0
0.4	12	0	675	0
0.5	2200	2	3600	5
0.6	3600	7	3600	9

of \mathcal{P} . In line with precedent results of Table 4.2, this table show that \mathcal{U}^{AC} outperform \mathcal{U}^Γ by providing solutions with lower average and extreme costs. However, the relative difference seems to be smaller for this set of larger instances.

5 Conclusion

In this chapter, we consider a SVC-based robust optimization approach to plan the picking operations of an assembly line when the latter are subcontracted to an outside service provider. Through a set of experiment we showed that using the classic SVC-based approach of Shang et al. (2017) is not a viable option to solve this problem in practice due to its long solving time. We propose to instead rely on the approximation introduced in Chapter 3 to obtain solutions quickly. We compare these solutions with a classical budget-based approach similar to the one introduced in Bertsimas and Sim (2003, 2004). Experimental results show that the proposed method efficiently reduces the

Table 4.4 – Comparison of \mathcal{U}^{Γ} and the approximation \mathcal{U}^{AC} on a set of 10 instances with 10 final products, 20 components and time horizon $T = 7$ periods.

Coverage	\mathcal{U}^{Γ}				\mathcal{U}^{AC}			
	Obstruction		Backlogging		Obstruction		Backlogging	
	Average	Q0.95	Average	Q0.95	Average	Q0.95	Average	Q0.95
0.1	14343	17284	9674	25354	− 12.28%	− 14.17%	− 13.47%	− 9.58%
0.2	17192	20017	7191	22405	− 9.51%	− 12.18%	− 9.21%	− 10.40%
0.3	15030	18195	13011	20610	+17.42%	+15.78%	− 52.15%	− 24.89%

impact of uncertainties on the performance of the assembly line. By controlling the level of robustness of the model, good tradeoffs can be found between reducing the backlog of assembly operations and minimizing the quantity of components stored on the border of the assembly line. In comparison with a classical budget-based robust optimization model, the data-driven model yields better performance on test instances with up to 10 different final products and 20 different components. The comparisons between solutions provided by both methods shows that the data-driven approach can lead simultaneously to a higher satisfaction of the demand (smaller backlogging cost) and a lower level of obstruction of the assembly line.

Despite the computational improvement provided by the use of the approximation, using an SVC-based approach lead to a significant increase of the CPLEX CPU-time compared to the budget based method. In addition, the computation time needed to solve both types of models increases significantly with the desired level of robustness (i.e coverage of historical data). An interesting perspective to apply this approach on larger industrial instances would be to combine the proposed approximate SVC-based set with adapted heuristic approaches, e.g. the well-known Relax-and-Fix, Fix-and-Optimize one, to provide good quality solutions in a reasonable amount of time. Notice that since robust models are reformulated as a deterministic robust counterpart, existing heuristic approaches from the literature on CLSP could be applied without further modifications.

An other relevant perspective for this specific application would be to consider a dynamic decision context in which the quantity ordered to the supplier can be adjusted dynamically to take the reported picking operations into account. This would lead to a better estimation of the consequences of picking time uncertainties on the assembly line and would reduce the cost of robust solution. However, the decentralized decision-making framework we consider assumes that picking decisions are made by the logistic provider,

which makes the description of the problem dynamics particularly challenging. Whether the TPL policy can be included in the uncertainty set or incentivized through the orders placed by assembly line manager remain an open question that should be addressed if one wishes to derive a formulation that faithfully relates to practical situations.

DATA-DRIVEN ROBUST INVENTORY MANAGEMENT UNDER UNCERTAIN DEMAND AND LEAD TIME

1 Introduction

In the two previous Chapters, we studied the efficiency of a non-parametric DDRO approach (SVC-based) that can be used as a black-box on any type of uncertainty without further tuning effort when data are available. In this Chapter, we propose to formulate our uncertainty sets from historical data but we assume that random parameters are driven by a given stochastic process (ARIMA) or have specific relationships with each others (Linear correlations). Moreover, while the SVC-based uncertainty set was used to model a single type of uncertainty in the previous models (item weights, demands, travel times, picking times) in this chapter, we study a more complex problem in which two types of uncertain parameters (demand and lead time) are modeled simultaneously.

Uncertain demand and lead time are two of the main factors driving up inventory management costs. Poor estimation of demand can lead to overstocking, generating additional costs, or to stock-outs that can have negative effects on future sales and profitability. In the same way, extended lead times or delivery delays increase the risk of stock-outs if the inventory level is insufficient. One may choose to adopt a conservative approach in this case and overestimate lead times but this can significantly increase storage costs. It has been pointed out in the inventory management literature that the majority of the existing works focus on demand uncertainty while uncertainties on delivery or production lead times are seldom considered. Moreover, most contributions focus on one type of uncertainty even if they are likely to occur simultaneously in practical problems (Ben Amar et al., 2013). The majority of research papers assume that the true distribution of

uncertain parameters is either known or well approximated by fitting the first empirical moments of the sample available. They then compute optimal reorder points and order quantities based on this information. However, this situation is often unrealistic as the true distribution is generally unknown and hard to characterize in practical cases. In addition, errors in the evaluation of this assumed distribution of lead time can lead to sub-optimal ordering policies (see e.g. Eppen and Martin (1988) who show that such approaches are sub-optimal in specific cases).

To overcome this drawback, Scarf's method paved the way for distributionally robust optimization with application to inventory management (Scarf, 1958; Gallego, 1992; Moon and Gallego, 1994). More broadly, RO has emerged in recent years as an appropriate approach to address various inventory management problems, especially in a distribution-free setting (see Chapter 1 for a short review of existing references on the topic).

It is a common practice for inventory or production managers to adjust their replenishment or planning decisions based on demand forecasts, often obtained through time series models that are tuned on historical data to faithfully capture the dynamics of a demand process. The classical methodology applied both by researchers and practitioners aims at designing a model that produces the most accurate predictions before using them as input parameters for their decision support tools. Among the options available, *Auto-Regressive Moving Average (ARMA)*, or its integrated version *Auto-Regressive Integrated Moving Average (ARIMA)*, process is a popular one to produce high-quality forecasts (Udom and Phumchusri (2014); Syntetos et al. (2008)). The natural relationship that link better predictions accuracy and improved decisions is not always true, as pointed out by several studies. Most of the works that apply robust optimization to inventory management and lot-sizing problems make the standard (but restrictive) assumption of independence of uncertain parameters. However, multiple papers have shown that demand processes may display autocorrelation (Erkip et al., 1990; Disney et al., 2006), in some cases with significant impact on inventory policies (Luong, 2007; Charnes et al., 1995). In addition, existing contributions on supply chain management have stated that lead times may be correlated with the workload of suppliers (Pahl et al., 2005, 2007). Indeed, a succession of orders may saturate a supplier's production capacity and thus increase the delivery lead time and variability. This correlation could also be strengthened in a competitive context if several retailers are facing high demands in the same period. On the contrary, a sequence of periods with low demand may incentivize suppliers to group production or transportation of items, inducing an increase in lead time. One way to

reduce the negative impact of load dependent lead times on the inventory management problem is to consider a joint production and inventory problem (Boute et al., 2014) or to use the workload of the supplier as an input to estimate the delivery lead time. However, in practice, this information is not always available for retailers and the ordering policy is often based on the observed delivery lead times. In this Chapter, we consider a robust multi-period, single-item inventory management problem with uncertain demand and lead time where all ordering decisions should be made before the beginning of the time horizon (Static Uncertainty model as described by Bookbinder and Tan (1998)). To address this problem, we provide an easy to implement, distribution-free, methodology that relies on existing tools and reduces the conservatism of classic robust optimization approaches when the uncertain parameters are not identically and independently distributed (i.i.d). In addition, we extend (Thorsen and Yao, 2017) in different directions: 1. We assume that the demand is correlated over time and propose a *Data Driven Uncertainty Set* based on autoregressive processes to capture this correlation. 2. We assume that the demand is correlated with the delivery lead time and propose a joint uncertainty set for demand and lead time to handle both types of uncertainty. 3. We consider fixed ordering costs and a discrete lead time uncertainty set. We conduct a set of experiments to demonstrate that our data-driven uncertainty sets outperforms classic uncertainty sets by integrating the information retrieved from the available historical data.

The rest of the Chapter is divided as follows: Section 2 presents the mathematical model and its robust formulation. Section 3 presents the proposed Data-Driven uncertainty sets. Section 4 presents the computational experiments settings, the computation results and discussions.

2 Problem and robust solution approach

We study an inventory Lot-Sizing problem over a discrete and finite horizon of T time periods. The goal is to minimize the costs incurred to manage the inventory of a single item in order to satisfy an uncertain sequence of demands (d_1, \dots, d_T) that follows a specific (but unknown) stochastic process. At the beginning of each period $t \in [T]$, the manager may place an order for a quantity of $q_t \geq 0$ units, incurring an ordering cost of $s\mathbb{1}_{\{q_t > 0\}} + cq_t$, where s is a fixed setup cost and c is a per-unit ordering cost. We consider that an order does not replenish the inventory immediately and instead assume that the units ordered are received after a discrete lead time of $l_t \in [1, L]$ periods. When

the available inventory in period t is insufficient to cover demand d_t , the unmet demand is backlogged to be satisfied in a subsequent period. At the end of each time period, the system also incurs a holding cost h for every unit physically held in the inventory and a backlogging penalty cost b for every unit of postponed demand.

2.1 Deterministic model

In a deterministic model, both lead time and demand vectors are known. For all $k \leq t$, we model lead times using binary parameters

$$\delta_{kt} = \begin{cases} 1 & \text{if the quantity } q_k \text{ ordered in } k \text{ is delivered in or before period } t \\ 0 & \text{otherwise} \end{cases} \quad (5.1)$$

The inventory level at the end of period t is given by the following equation:

$$I_t = I_0 + \sum_{k=1}^t (\delta_{kt} q_k - d_k) \quad (5.2)$$

where I_0 is the initial inventory level. The corresponding inventory cost incurred in period t is thus $y_t = \max\{hI_t, -bI_t\}$.

We use the above observation to formulate the problem as the Mixed Integer Program (MIP) below :

$$\min_{\substack{\mathbf{x} \in \{0,1\}^T \\ \mathbf{q}, \mathbf{y} \in \mathbb{R}_+^T}} \sum_{t=1}^T s x_t + c q_t + y_t \quad (5.3)$$

$$\text{s.t. } y_t \geq h \left(I_0 + \sum_{i=1}^t (\delta_{it} q_i - d_i) \right) \quad \forall t \in [T] \quad (5.4)$$

$$y_t \geq -b \left(I_0 + \sum_{i=1}^t (\delta_{it} q_i - d_i) \right) \quad \forall t \in [T] \quad (5.5)$$

$$q_t \leq M x_t, \quad \forall t \in [T] \quad (5.6)$$

where $\mathbf{x} \in \{0,1\}^T$ are binary variable such that $x_t = 1$ if an order is placed in period t and 0 otherwise and M is a sufficiently large value defined such that constraint (5.6) is always true if $x_t = 1$, e.g. $M = \sum_{t \in [T]} d_t$.

2.2 Robust optimization model

Following the classical paradigm of robust optimization, we consider that both the demand vector \mathbf{d} and delivery matrix $\boldsymbol{\delta}$ belong to uncertainty sets \mathcal{U}_d and \mathcal{U}_δ , respectively, similar to the ones introduced by Thorsen and Yao (2017) for an inventory problem with no fixed costs. The goal is then to optimize against the worst parameters values in these sets by solving the following MIP:

$$\min_{\substack{\mathbf{x} \in \{0,1\}^T \\ \mathbf{q} \in \mathbb{R}_+^T}} \max_{\substack{\mathbf{d} \in \mathcal{U}_d \\ \boldsymbol{\delta} \in \mathcal{U}_\delta}} \min_{\mathbf{y} \in \mathbb{R}_+^T} \sum_{t=1}^T sx_t + cq_t + y_t \quad (5.7)$$

$$\text{s.t. } y_t \geq h \left(I_0 + \sum_{i=1}^t (\delta_{it}q_i - d_i) \right) \quad \forall t \in [T] \quad (5.8)$$

$$y_t \geq -b \left(I_0 + \sum_{i=1}^t (\delta_{it}q_i - d_i) \right) \quad \forall t \in [T] \quad (5.9)$$

$$q_t \leq Mx_t \quad \forall t \in [T] \quad (5.10)$$

The objective function (5.7) minimizes the maximum value of the sum of fixed/unit ordering costs and holding/backlogging costs over the planning horizon. Constraints (5.8) and (5.9) define the holding/backlogging cost y_t due to storage or stockouts at the end of period $t \in [T]$. Finally, constraints (5.10) link the binary vector \mathbf{x} with the vector of ordered quantity \mathbf{q} .

As mentioned in Chapter 1, a popular choice in robust inventory management is the budget-based uncertainty set introduced by Bertsimas and Sim (2003, 2004). It has been used in Bertsimas and Thiele (2006) to define admissible values for a demand vector from so-called budgets of uncertainty Γ :

$$\mathcal{U}_d^\Gamma = \left\{ \mathbf{d} = \bar{\mathbf{d}} + \hat{\mathbf{d}}^\top \mathbf{z} \mid \sum_{i=1}^t z_i \leq \Gamma_t \quad \forall t \in [T], \mathbf{z} \in [-1, 1]^m \right\} \quad (5.11)$$

where $\bar{\mathbf{d}}$ is the vector of nominal demand and $\hat{\mathbf{d}} = [\hat{d}_1, \dots, \hat{d}_T]$ is the vector of maximum deviation from the nominal demand. The tuning parameters Γ_t control the degree of robustness of the model by bounding the cumulative deviation of the demands from their nominal values. The value of $\bar{\mathbf{d}}$ and $\hat{\mathbf{d}}$ define how \mathcal{U}_d^Γ covers the possible values of \mathbf{d} , which may have a significant influence on the solutions obtained. Hence they must be selected carefully to include a predefined fraction of the demands, e.g. 95% (see Bertsimas and

Thiele (2006)). More recently, Thorsen and Yao (2017) and Thevenin et al. (2022) proposed a robust approach to inventory management problems under lead time uncertainty by introducing an additional uncertainty set on δ :

$$\mathcal{U}_\delta^\Gamma = \left\{ \delta \left| \begin{array}{l} \delta_{tk} \leq \delta_{tk+1}, \forall t = 1, \dots, T-1, t \leq k \leq T \\ \sum_{t=1}^T \sum_{k=t}^T (1 - \delta_{tk}) \leq \Gamma \\ \delta_{tt+L} = 1, \forall t = 1, \dots, T \\ \delta_{tk} \in [0, 1], \forall t = 1, \dots, T, t \leq k \leq T \end{array} \right. \right\} \quad (5.12)$$

where L is an upper bound on the lead time value and the budget of uncertainty Γ bounds the cumulative lead time over the planning horizon. Clearly δ must be a matrix of binary values to ensure that $l_t = \sum_{k=t}^T (1 - \delta_{tk})$ is an integer value for all $t = 1, \dots, T$. Note that \mathcal{U}_δ relaxes this constraint by considering instead that $\delta \in [0, 1]^{T^2}$ to obtain a convex uncertainty set. This models a continuous lead time, effectively allowing an order to be delivered in more than one single period. For instance, a vector $\delta_1 = [0.5, 0.5, 0.5, 1, 1]$ corresponds to a scenario in which half of the order q_1 is received in period 1 and the second half is received in period 4.

3 Data-driven uncertainty set

In this section, we apply time series analysis on available data to take advantage of the structure of classical stochastic processes used to model demands. In particular, we exploit properties of the type of processes that is often used as a demand forecasting tool to incorporate correlation information between parameters and propose an alternative uncertainty set that is less conservative than \mathcal{U}_d^Γ . This approach differs from previous data-driven uncertainty sets such as the SVC based uncertainty set studied in Chapter 3 and 4 as we assume that the demand follows a specific stochastic process. After presenting the theoretical foundations to build this uncertainty set, we provide insights on fitting the different parameters from historical data and propose an extension that incorporates the influence of past demands on order lead times (production and/or transportation).

3.1 Data driven uncertainty set based on ARMA(p,q) process

Many existing works on inventory management rely on time series analysis to produce forecasting models for demand process (see e.g. Duc et al. (2008) and included references). Among the most popular ones, *Autoregressive Moving Average (ARMA)* processes often offer enough flexibility to derive good quality predictions, with the objective to make better informed production and/or ordering decisions. Besides its linear nature that facilitates integration into MIP formulations, another key advantage is that one can easily fit demand data to such processes using one of the many existing off-the-shelf libraries dedicated to time series analysis.

3.1.1 General ARMA(p,q)-based uncertainty set

The demand $(D_t)_{t \in \mathbb{N}}$ is an *ARMA(p,q)* process of mean 0 if it satisfies the following relationship:

$$\begin{aligned} D_t &= 0 & \forall t \leq 0 \\ D_t + \beta_1 D_{t-1} + \dots + \beta_p D_{t-p} &= \epsilon_t + \theta_1 \epsilon_{t-1} + \dots + \theta_q \epsilon_{t-q} & \text{otherwise} \end{aligned}$$

where $\boldsymbol{\beta} \in \mathbb{R}^p$ and $\boldsymbol{\theta} \in \mathbb{R}^q$ are the vectors of autoregressive and moving-average coefficients and where $\boldsymbol{\epsilon}$ is a vector of i.i.d. noise of mean 0 and variance σ_ϵ^2 . If the process has mean $\mu \neq 0$, the process defined for all t as $D_t - \mu$ satisfies the above equations.

Given a vector of $N + 1$ consecutive (historical) demand values $[d_{-N}, \dots, d_0]$ that are assumed to be generated by an *ARMA(p,q)* process, one can find the Maximum Likelihood Estimators (MLE) of parameters $\mu, \sigma, \beta_1, \dots, \beta_p, \theta_1, \dots, \theta_q$ assuming a normally distributed random noise ϵ_t . One can then use classical results on *ARMA* processes to obtain the prediction interval for a given confidence level $1 - \alpha$. We define the *ARMA(p,q)* based uncertainty set as:

$$\mathcal{U}_d^{ARMA(p,q)} = \left\{ \mathbf{d} \left| \begin{array}{l} d_t = \beta_0 + \sum_{k=1}^p \beta_k d_{t-k} + \sum_{k=1}^q \theta_k \epsilon_{t-k} + \epsilon_t \quad \forall t \in [T] \\ \left| \sum_{j=1}^t \epsilon_j \right| \leq \Gamma_\epsilon \sigma_\epsilon \sqrt{t} \quad \forall t = 1, \dots, T \\ \underline{d}_t \leq d_t \leq \bar{d}_t \quad \forall t \in [T] \\ \underline{\epsilon}_t \leq \epsilon_t \leq \bar{\epsilon}_t \quad \forall t \in [T] \end{array} \right. \right\} \quad (5.13)$$

where: for all $k \in \mathbb{N}$, d_{-k} is equal to an historical value from the training sample, \underline{d}_t and \bar{d}_t are lower and upper bounds on the demand, $\underline{\epsilon}_t$ and $\bar{\epsilon}_t$ are lower and upper bounds on the gaussian noise in period t , β_0 is a constant such that $\beta_0 = \mu - \sum_{k=1}^p \beta_k \mu$, and the parameter $\Gamma_\epsilon \geq 0$ restricts the value of the noise vector ϵ . Notice that the uncertainty set restricts the maximum cumulative deviation on the errors themselves: This ensures that each demand sequence is consistent through time by linking the errors in the different time periods.

Remark 1. The interval $[\underline{\epsilon}_t, \bar{\epsilon}_t]$ can be any symmetric or asymmetric interval containing 0.

Remark 2. More advanced type of processes, such as the Autoregressive Integrated Moving Average (ARIMA) or its seasonal version, involve to fit an ARMA process on linear transformations of the original process. Embedding such structures thus only requires simple modifications of the uncertainty set (5.13) without altering its nature (i.e. bounded and convex polyhedron). For instance, if $(D_t)_{t \in [T]}$ is modeled accurately by an $ARIMA(p, 1, q)$ process, it suffices to replace all terms of the form d_j by $d_j - d_{j-1}$ in the equality constraints.

3.1.2 Examples for special cases of demand processes

We focus on a simple case where the demand for a given product follows an $AR(1)$ process, or equivalently and $ARMA(1, 0)$ process. The uncertainty set (5.13) based on the $AR(1)$ process uses the following demand relationship:

$$d_t = \beta_0 + \beta_1 d_{t-1} + \epsilon_t \quad (5.14)$$

Assume the values of β_0 , β_1 and σ_ϵ^2 are fitted using their MLE on a training sample containing $N + 1$ historical observations. Since ϵ_t has mean 0 for all t , we know from equation (5.14) that

$$\mathbb{E}(d_t) = \beta_0 + \beta_1 \mathbb{E}(d_{t-1}) \quad (5.15)$$

and we can express the variance of the demand d_t as:

$$\text{Var}(d_t) = \sigma_\epsilon^2 \frac{1 - \beta_1^{2t}}{1 - \beta_1^2} \quad (5.16)$$

We use this formula to define the bounds \underline{d}_t and \bar{d}_t as follows:

$$\mathbb{E}(d_t) - \Gamma_d \sqrt{\sigma_\epsilon^2 \frac{1 - \beta_1^{2t}}{1 - \beta_1^2}} \leq d_t \leq \mathbb{E}(d_t) + \Gamma_d \sqrt{\sigma_\epsilon^2 \frac{1 - \beta_1^{2t}}{1 - \beta_1^2}} \quad (5.17)$$

where Γ_d control the width of the interval. In what follows, we tune this parameter empirically on the vector of demand $[d_{-N}, \dots, d_0]$ in order to cover 95% of the possible values of d_t . Note that since parameter fitting assumes normally distributed errors, it is quite natural to define a symmetric interval centered on the value $\mathbb{E}(d_t)$.

With ARMA(1,1), the induction becomes:

$$D_t = \beta_0 + \beta_1 D_{t-1} + \epsilon_t + \theta_1 \epsilon_{t-1} \quad (5.18)$$

Hence, by induction, it is clear that one can express D_t as a linear combination of errors ϵ_t and data d_1, \dots, d_N . In particular, it is possible to derive a prediction interval for any future period $t > 0$ with a given confidence level $(1 - \alpha)\%$ using bootstrapping techniques to obtain approximate confidence limits. This is well-documented in the literature (see e.g. Brockwell and Davis (2016)) and readily implemented in numerous existing libraries dedicated to time series analysis (e.g. <https://www.statsmodels.org/stable/tsa.html> in python or <https://cran.r-project.org/package=tseries> in R).

3.2 Joint demand and lead time uncertainty set

As mentioned in the introduction, we assume that the lead time may be influenced by the value of past demands. Among others, *Linear Regression (LR)* is a simple way to describe the dependency between multiple parameters. In our case, it consists in fitting a linear model of the form:

$$l_t = \bar{l} + \sum_{k=0}^n b_k d_{t-k} + \zeta_t \quad (5.19)$$

where the uncertain lead time l_t of period t is described by the constant value \bar{l} , the vector of linear dependencies \mathbf{b} and the random perturbation vector $\boldsymbol{\zeta}$ lying in $[\underline{\zeta}, \bar{\zeta}]$ and with mean 0. Given a vector of $N + 1$ consecutive demand values $[d_{-N}, \dots, d_0]$ and $N + 1$ consecutive lead time values $[l_{-N}, \dots, l_0]$, one can use the *Mean Squared Error (MSE)* estimator to fit the linear model and obtain the value of \bar{l} and \mathbf{b} by minimizing the

function:

$$MSE = \frac{1}{N} \sum_{t=1}^N (\hat{l}_t - l_t)^2 \quad (5.20)$$

where \hat{l}_t is the prediction value defined by:

$$\hat{l}_t = \bar{l} + \sum_{k=0}^n b_k d_{t-k} \quad (5.21)$$

Note that we can express the value of the lead time l_t as :

$$l_t = \sum_{k=t}^T (1 - \delta_{tk}) \quad (5.22)$$

We can then define a joint uncertainty set for demand and lead time by:

$$\mathcal{U}_{\delta}^{LR} = \left\{ \boldsymbol{\delta}, \mathbf{d} \left\{ \begin{array}{l} \delta_{tk} \leq \delta_{tk+1}, \quad \forall t = 1, \dots, T-1, t \leq k \leq T \quad (a) \\ \delta_{tk} \in \{0, 1\}, \quad \forall t = 1, \dots, T, t \leq k \leq T \quad (b) \\ L_{\min} \leq \sum_{k=t}^T (1 - \delta_{tk}) \leq L_{\max}, \quad \forall t = 1, \dots, T \quad (c) \\ \sum_{k=t}^T (1 - \delta_{tk}) = \bar{l} + \sum_{k=0}^n b_k d_{t-k} + \zeta_t, \quad \forall t = 1, \dots, T \quad (d) \\ \underline{\zeta} \leq \zeta_t \leq \bar{\zeta}, \quad \forall t = 1, \dots, T \quad (e) \\ d_t = \beta_0 + \sum_{k=1}^p \beta_k d_{t-k} + \sum_{k=1}^q \theta_k \epsilon_{t-k} + \epsilon_t \quad \forall t \in [T] \quad (f) \\ \left| \sum_{j=1}^t \epsilon_j \right| \leq \Gamma_{\epsilon} \sigma_{\epsilon} \sqrt{t} \quad \forall t = 1, \dots, T \quad (g) \\ \underline{d}_t \leq d_t \leq \bar{d}_t \quad \forall t \in [T] \quad (h) \\ \underline{\epsilon}_t \leq \epsilon_t \leq \bar{\epsilon}_t \quad \forall t \in [T] \quad (i) \end{array} \right. \right\} \quad (5.23)$$

where constraints (a) ensure that the quantity received until period $t + 1$ is greater than or equal to the quantity received until period t . Constraints (b) forces the integrality of $\boldsymbol{\delta}$, which in turn implies the integrality of lead times. Constraints (c) and (d) bound the value of lead times and constraints (e) bound the possible values of the perturbations ζ_t . Note that this set of constraint allows for order crossover, (i.e. we can receive the order q_t before precedent orders q_{t-n} for all $n \in [1, L_{\max} - 1]$). Finally, constraints (f)-(i) are constraints related to the demand.

3.3 A constraint generation approach

In uncertainty set (5.12), the binary constraint on δ is relaxed to define a convex uncertainty set. This property is required in classic robust optimization to use the dual formulation of the subproblem. In the uncertainty set (5.23), the constraint (d) integrates the correlation between demand and lead time. In order to use the linear model (5.19) obtained by minimizing function (5.20) we need to ensure the integrality of the lead time values (constraint (b)). As a result, the proposed data-driven uncertainty set is a finite discrete set and solving the problem using a dual reformulation approach is not possible. However, one can adapt the constraint generation approach introduced in Bienstock and Ozbay (2008) and used in Thorsen and Yao (2017) to solve the robust formulation of the problem with the joint discrete uncertainty set. This approach rely on 1. a master problem that take decisions according to a set of uncertain scenarios and 2. a sub problem generating *worst case* scenarios to enrich the master problem. Both problems are solved iteratively by enriching a set Ω of lead time and demand scenarios until an optimal robust solution is found. For a given set of scenarios Ω , let

$$\mathcal{X}(\Omega) = \left\{ (\mathbf{x}, \mathbf{q}, \mathbf{y}) : \mathbf{x} \in \{0, 1\}^T, \mathbf{q} \in \mathbb{R}_+^T, \mathbf{y} \in \mathbb{R}_+^{T \times |\Omega|} \right\}$$

, the master problem of the inventory planning model under uncertain lead time and demand is given by:

$$\min_{\mathbf{x}, \mathbf{q}, \mathbf{y}, z} z \tag{5.24}$$

$$\text{s.t. } z \geq \sum_{t \in [T]} s x_t + c q_t + y_{t\omega} \quad \forall \omega \in \Omega \tag{5.25}$$

$$y_{t\omega} \geq h \left(I_0 + \sum_{i=1}^t (\delta_{it}^{\omega} q_i - d_i^{\omega}) \right) \quad \forall t \in [T], \forall \omega \in \Omega \tag{5.26}$$

$$y_{t\omega} \geq -b \left(I_0 + \sum_{i=1}^t (\delta_{it}^{\omega} q_i - d_i^{\omega}) \right) \quad \forall t \in [T], \forall \omega \in \Omega \tag{5.27}$$

$$q_t \leq M x_t \quad \forall t \in [T] \tag{5.28}$$

$$(\mathbf{x}, \mathbf{q}, \mathbf{y}) \in \mathcal{X}(\Omega) \tag{5.29}$$

By solving this problem, we obtain the optimal decision vectors \mathbf{x}^* and \mathbf{q}^* that are used to identify a new worst-case scenario to add to Ω . This scenario is obtained as the optimal

solution to the following MILP:

$$\max_{\delta, \mathbf{d}, H, B, p} \quad \sum_{t=1}^T sx_t^* + cq_t^* + H_t + B_t \quad (5.30)$$

$$\text{s.t.} \quad H_t \geq h \left(I_0 + \sum_{i=1}^t (\delta_{it} q_i^* - d_i) \right) \quad \forall t \in [T] \quad (5.31)$$

$$H_t \leq h \left(I_0 + \sum_{i=1}^t (\delta_{it} q_i^* - d_i) \right) + M_t(1 - p_t) \quad \forall t \in [T] \quad (5.32)$$

$$H_t \leq M_t p_t \quad \forall t \in [T] \quad (5.33)$$

$$B_t \geq -b \left(I_0 + \sum_{i=1}^t (\delta_{it} q_i^* - d_i) \right) \quad \forall t \in [T] \quad (5.34)$$

$$B_t \leq -b \left(I_0 + \sum_{i=1}^t (\delta_{it} q_i^* - d_i) \right) + M_t p_t \quad \forall t \in [T] \quad (5.35)$$

$$B_t \leq M_t(1 - p_t) \quad \forall t \in [T] \quad (5.36)$$

$$H_t, B_t \geq 0 \quad \forall t \in [T] \quad (5.37)$$

$$p_t \in \{0, 1\} \quad \forall t \in [T] \quad (5.38)$$

$$\delta, \mathbf{d} \in \mathcal{U}_\delta^{LR} \quad (5.39)$$

where M_t are sufficiently large values. This sub-problem maximize the inventory costs obtained with the optimal decisions $\mathbf{x}^*, \mathbf{q}^*$ of the master problem by selecting the worst realization of demand and lead time in the uncertainty set.

The procedure for this constraint generation approach is given by Algorithm 1. Biemstock and Ozbay (2008) proved that this algorithm converges to an optimal robust solution in a finite number of iterations when the uncertainty sets are non-empty and bounded.

4 Computational experiments

The objective of this section is to evaluate the solutions provided by the robust models introduced up to now and compare their performances against a sample average approach (presented in Section 4.1). In order to evaluate the impact of each contribution, we consider two different scenarios set of experiments. In the first set of experiments - Section 4.2), we consider that the demand follows an autoregressive process $AR(1)$ while lead times are uniformly distributed in an interval. In the second set of experiments

Algorithm 1 Adversarial Approach Bienstock and Ozbay (2008)

1. Initialization:

Select an initial set of scenarios Ω

Define tolerance value ϵ

Set $LB = 0$ and $UB = \infty$

2. Master problem:

Solve the master problem with scenario set Ω to obtain \mathbf{x}^* , \mathbf{q}^* and Z^*

Set $LB = Z^*$

3. Sub problem:

Solve adversarial sub-problem with \mathbf{x}^* and \mathbf{q}^* to obtain δ^* , \mathbf{d}^*

Set $UB = \min(UB, \text{Objective value of sub-problem})$

4. If $UB - LB > \epsilon$:

Add scenario (δ^*, \mathbf{d}^*) to Ω , and return to Step 2.

5. If $UB - LB \leq \epsilon$:

Terminate and return \mathbf{q}^* .

(Section 4.3), we study the impact of potential correlations between demand and lead time.

4.1 Sample average approach

The Sample Average Approximation (SAA) model is formulated as a two-stage model in which the first stage decisions defines the ordering periods \mathbf{x} as well as the quantities ordered \mathbf{q} , while the second stage decisions computes the holding and backlogging costs \mathbf{y} for each scenario. The model is given by the following MIP:

$$\min_{\mathbf{x}, \mathbf{q}, \mathbf{y}} \quad \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{t=1}^T sx_t + cq_t + y_{t\omega} \quad (5.40)$$

$$\text{s.t. } y_{t\omega} \geq h \left(I_0 + \sum_{i=1}^t (\delta_{it}^{\omega} q_i - d_i^{\omega}) \right) \quad \forall t \in [T], \forall \omega \in \Omega \quad (5.41)$$

$$y_{t\omega} \geq -b \left(I_0 + \sum_{i=1}^t (\delta_{it}^{\omega} q_i - d_i^{\omega}) \right) \quad \forall t \in [T], \forall \omega \in \Omega \quad (5.42)$$

$$q_t \leq Mx_t \quad \forall t \in [T] \quad (5.43)$$

$$(\mathbf{x}, \mathbf{q}, \mathbf{y}) \in \mathcal{X}(\Omega) \quad (5.44)$$

in which Ω is the set of demand and lead time scenarios. The objective function (5.40) aims at minimizing the average cost over the set of scenarios. Constraints (5.41) and (5.42) compute the costs due to storage and backlogging in each period and each scenario. Finally, constraints (5.43) link the quantity ordered q_t and the binary variable x_t .

Clearly, the quality of an SAA solution depends heavily on the set of scenarios Ω that

Table 5.1 – Costs obtained by SAA solutions with AR(1) and Normal demand scenarios for instances of 20 periods without fixed cost, ratio $\frac{b}{h} = 8$, $\sigma_\epsilon = 10$ when β_1 vary.

β_1	Average		Median		Standard deviation		Q0.99	
	Normal	AR(1)	Normal	AR(1)	Normal	AR(1)	Normal	AR(1)
0.3	695	639	551	573	488	308	2892	2055
0.5	1092	780	702	700	982	414	4955	2624
0.8	2916	1325	1631	1154	2843	864	12219	5198

is included in the formulation. In particular when the demand follows a specific process such as the ARMA described before, one may improve the performances of the formulation by generating trajectories using the information embedded in the historical data . As a matter of example, Table 5.1 shows the costs of solutions obtained with SAA models on an instance with 20 periods, no fixed cost and a ratio $\frac{b}{h} = 8$ in which the demand follows an $AR(1)$ process. We study three types of instances corresponding to different degree of auto-correlation and each line present the average values obtained over a set of 10 instances. For each type, we generate a scenario set containing 500 scenarios whose demand distribution follows either a Normal distribution or an $AR(1)$. Each parameters used to generate scenarios (e.g. β_1 , mean value, standard deviation) are estimated on historical data sets of 100 demand realizations. The solutions found by each formulation are then evaluated on 10 000 different trajectories and for each value of β_1 the process is repeated on 50 instances.

The results indicate that the SAA using scenarios generated with the $AR(1)$ process lead to better solutions than the one based on an iid demand generation, both on average and in the worst case. For example when $\beta_1 = 0.8$ the average and extreme (0.99 quantile) costs obtained with the normally distributed training sample are more than twice the ones obtained with the one generated based on an autoregressive process.

4.2 Uncertain correlated demand

The results from Table 5.1 suggest that being able to precisely characterize the demand process leads to a better scenario generation. However in practice, the only information available to define an appropriate model and tune its parameters consists of the historical demand values. Even when the nature of the underlying process is known, biased estimates of the parameters may affect scenario generation and deteriorate the performances of the

Table 5.2 – Instance configurations for the first set of experiments

Name	Definition	Value
T	time horizon	20
c	unit ordering cost	0.1
h	unit holding cost	0.3
b	unit backloging cost	2.4
s	fixed ordering costs	{0, 45, 105, 165}
TBO	time between orders	{0, 1, 2, 3}
β_1	autocorrelation factor	{0.3, 0.5, 0.8}
β_0	constant factor	$100(1 - \beta_1)$ (gives an average demand value of 100)
σ_ϵ	standard deviation of the perturbation	10

SAA. In this context, robust optimization may be seen as an alternative to prevent decision makers from paying extreme inventory costs because of these errors. In the remainder of this section, we assume that all the models use a common data set containing $N = 100$ consecutive demand (and leadtimes in Section 4.3) observations. In particular when the models have access to the specific demand process, they define its parameters as their MLE using the data available (and their possible bias). On each instance, we compare the performances of the SAA formulation against two robust models that use either a budget-based or an ARMA-based uncertainty set.

In this first set of experiment, the demand evolves according to an $AR(1)$ process with several combination of parameters as described in Table 5.2. Note that the different values of fixed cost s are computed to control the *Time Between Order (TBO)*, which corresponds to the expected number of periods between two consecutive orders in order to minimize the total costs. For each combination, we generated 50 distinct instances, i.e. LSP that share the same parameters except their training sample of historical demands. For each instance, we compare the solution obtained by the SAA described in Section 4.1 a *Sample Average Approximation (SAA)*, a robust approach with the classic uncertainty set \mathcal{U}_d^Γ described in (5.11) and the $AR(1)$ based uncertainty set described above $\mathcal{U}_d^{ARMA(1,0)} = \mathcal{U}_d^{AR(1)}$.

We evaluate the solutions obtained by each method on an independent test set of $K = 10^5$ scenarios generated using Monte Carlo simulation and we report the average, 0.95- and 0.99-quantiles of the empirical cost achieved by each solution over the test sample. More precisely, we report the absolute cost evaluated by simulation for the SAA model while for both robust approaches we report the relative cost difference from the SAA solution in %. For the classical robust model, the budget parameters Γ_t are set for all $t \in [T]$ with the objective to include the cumulative demand in the uncertainty set with

Table 5.3 – Costs obtained with SAA, \mathcal{U}_d^Γ and $\mathcal{U}_d^{AR(1)}$ for instances with fixed cost $s = 0$ when β_1 vary.

β_1	Average			Q0.95			Q0.99		
	SAA	\mathcal{U}_d^Γ	$\mathcal{U}_d^{AR(1)}$	SAA	\mathcal{U}_d^Γ	$\mathcal{U}_d^{AR(1)}$	SAA	\mathcal{U}_d^Γ	$\mathcal{U}_d^{AR(1)}$
0.3	690.28	+33.44%	+0.63%	1385.51	-15.46%	-16.48%	2266.19	-29.39%	-23.97%
0.5	859.47	+20.16%	+3.57%	1819.01	-15.99%	-16.3%	2973.26	-23.34%	-21.3%
0.8	1480.14	+11.83%	+10.29%	3473.42	-6.28%	-20.26%	5875.9	-10.64%	-27.23%

an approximate probability of 0.95, under the assumption of IID normally distributed demands. Assuming that the training sample empirical mean and standard deviation are \bar{d} and $\hat{\sigma}_d$, respectively, we denote for all $t \in [T]$ and $k = t, \dots, N + 1$: $\Delta_k(t) = \sum_{j=1}^t |d_{-k+j} - \bar{d}|$, the cumulative absolute deviation from \bar{d} of a sequence of t consecutive demands starting from index $-k$ in the training data set. We then define for all $t \in [T]$ the budget value $\Gamma_t = \min\{\gamma \geq 0 : \sum_{k=t}^N \mathbb{1}_{\{\Delta_k(t) \leq 2\gamma\hat{\sigma}_d\}} \geq 0.95(N + 1 - t)\}$. For the $AR(1)$ -based uncertainty set, we use the normality assumption on the error terms and define $\Gamma_\epsilon = 2$ so that $\mathbb{P}\left(|\sum_{k=1}^t \epsilon_t| \leq \Gamma_\epsilon \sigma_\epsilon \sqrt{t}\right) \approx 0.95$.

Table 5.3 shows the costs of solutions obtained with different degree of autocorrelation β_1 when there is no setup costs. In this situation, we observe that the SAA model always provide the lowest average cost but also the highest extreme costs. For example, when $\beta_1 = 0.3$, the 0.99-quantile of the SAA solution is more than three times its average cost. In comparison, the robust approaches provides much higher average costs but with a significant reduction of extreme costs. Using the $AR(1)$ -based uncertainty set provides the lowest average costs and 0.95-quantile among the two robust approaches and across all values of β_1 . Interestingly, the classic robust approach outperforms the $AR(1)$ -based in case of weakly auto-correlated processes ($\beta_1 \in \{0.3, 0.5\}$) but this phenomenon fades as β_1 increases. In particular when $\beta_1 = 0.8$ the solutions derived using the $AR(1)$ -based approach give the best 0.99-quantile performance. Generally speaking, it appears that the budget-based uncertainty set \mathcal{U}_d^Γ cannot lead to both robust solutions and good average performances simultaneously compared to $\mathcal{U}_d^{AR(1)}$.

Table 5.4 summarizes the costs obtained with different values of setup cost s when the autocorrelation factor $\beta_1 = 0.8$. We observe that increasing setup costs decrease the cost difference across all methods without modifying the previous observations. The SAA model still provide the lower average cost and still lead to very high 0.95- and 0.99-quantile. For example, when the setup cost $s = 45$, the $AR(1)$ based robust solution

Table 5.4 – Costs obtained with SAA, \mathcal{U}_d^Γ and $\mathcal{U}_d^{AR(1)}$ for instances with $\beta_1 = 0.8$ when the fixed cost s vary.

s	Average			Q0.95			Q0.99		
	SAA	\mathcal{U}_d^Γ	$\mathcal{U}_d^{AR(1)}$	SAA	\mathcal{U}_d^Γ	$\mathcal{U}_d^{AR(1)}$	SAA	\mathcal{U}_d^Γ	$\mathcal{U}_d^{AR(1)}$
0	1480.14	11.83%	10.29%	3473.42	-6.28%	-20.26%	5875.9	-10.64%	-27.23%
45	2022.89	4.54%	2.17%	3665.55	-9.14%	-14.58%	5785.06	-16.5%	-22.87%
105	2517.81	2.57%	-0.09%	4151.87	-5.7%	-11.95%	5980.36	-9.99%	-18.53%
165	2825.64	2.43%	0.6%	4304.56	-6.1%	-11.32%	5970.0	-11.15%	-19.74%

increase the average cost by 2.17% while it reduces the average 0.99-quantile by 22.87%.

In order to compare in more details the relative performances of the three methods, we compare the empirical *superquantile* of the cost achieved by each method over the test set. Superquantile is a well-established measure to evaluate the risk taken by a decision maker. Also known as conditional value at risk or expected shortfall in portfolio optimization, it corresponds to the average cost above a given quantile q_α of the cost distribution. For a given random variable X , the superquantile is given by:

$$S_\alpha(X) = \mathbb{E}[X|X \geq q_\alpha(U)] = \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha} \mathbb{E}[(X - \eta)^+] \right\}$$

, with its empirical counterpart for a sample $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_K)$ of size K :

$$\hat{S}_\alpha(\hat{\mathbf{x}}) = \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{(1-\alpha)K} \sum_{i=1}^K (\hat{x}_i - \eta)^+ \right\}$$

In the later, we compare the costs obtained by the SAA approach and the classic budget based robust approach with the $AR(1)$ -based robust approach by computing the relative gap between the values of their respective superquantile for various values $\alpha \in [0, 1]$. Recall that for each combination of parameters, we create 50 historical trajectories $[d_{-N}^{(i)}, \dots, d_0^{(i)}]$ that are used to obtain one solution to the corresponding LSP for each method. Each solution computed for instance i is then evaluated on K distinct out-of-sample trajectories of the process, giving three vectors containing their own realized costs $\hat{\mathbf{z}}_{\text{SAA}}^{(i)} = (\hat{z}_{\text{SAA}}^{(i,1)}, \dots, \hat{z}_{\text{SAA}}^{(i,K)})$, $\hat{\mathbf{z}}_{\Gamma}^{(i)} = (\hat{z}_{\Gamma}^{(i,1)}, \dots, \hat{z}_{\Gamma}^{(i,K)})$ and $\hat{\mathbf{z}}_{AR(1)}^{(i)} = (\hat{z}_{AR(1)}^{(i,1)}, \dots, \hat{z}_{AR(1)}^{(i,K)})$. For a given instance $i = 1, \dots, 50$, one can thus compare the relative superquantile of the three

methods by computing

$$\rho_{\text{SAA}}^{(i)}(\alpha) = \frac{\hat{S}_\alpha(\hat{\mathbf{z}}_{\text{SAA}}^{(i)}) - \hat{S}_\alpha(\hat{\mathbf{z}}_{AR(1)}^{(i)})}{\hat{S}_\alpha(\hat{\mathbf{z}}_{AR(1)}^{(i)})} \quad (5.45)$$

$$\rho_{\Gamma}^{(i)}(\alpha) = \frac{\hat{S}_\alpha(\hat{\mathbf{z}}_{\Gamma}^{(i)}) - \hat{S}_\alpha(\hat{\mathbf{z}}_{AR(1)}^{(i)})}{\hat{S}_\alpha(\hat{\mathbf{z}}_{AR(1)}^{(i)})} \quad (5.46)$$


Notice that we take as a reference the cost returned by the $AR(1)$ -based robust formulation to facilitate the evaluation of its performances compared to the other two. In Figure 5.1 and 5.2, we represent the box plot across 50 instances of the relative gaps defined in (5.45) and (5.46) on the superquantiles of the different methods when α increase. Figure 5.1 shows the evolution of this ratio for different values of β_1 with no setup costs and all other parameters fixed, while Figure 5.2 compares the same evolution for a fixed β_1 as the setup cost increases. We observe in line with Table 5.3 and 5.4 that SAA provides the best average cost in most instances while the classic robust approach provide good performances on extreme values of the cost distribution ($\alpha = 0.9$). Another important observation is that as β_1 increases, the cost difference between all methods decreases (on average over all instances) as stated by table 5.3. However both the SAA and the budget-based robust formulation compute solutions with much stronger positive deviations from the cost achieved by the $AR(1)$ -based robust solution, while improvements over the latter by one of the former is both rare and of less magnitude across all values of α and all the combinations tested.

4.3 Uncertain demand-dependent lead times


In this second set of experiments, we consider uncertain demand and lead times and assume that a positive correlation exists between these two types parameters. In the remainder, we denote by $\mathcal{U}_\delta^\Gamma$ the use of classical uncertainty sets presented in Section 2 for demand and lead time while \mathcal{U}_δ^{LR} refer to the data-driven joint uncertainty set introduced in Section 3.

Lead time generation Given a set of $N + 1$ consecutive demand $[d_{-N}, \dots, d_0]$ we generate a set of $N + 1$ consecutive lead time values $[l_{-N}, \dots, l_0]$ taking values in $[L]$ with the following formula:


$$l_t = \min\{L, \max(1, Z_t^2 * L)\} \quad (5.47)$$



(a) $s = 0, \beta_1 = 0.3$




(b) $s = 0, \beta_1 = 0.5$




(c) $s = 0, \beta_1 = 0.8$


Figure 5.1 – Evolution of the relative \hat{S}_α as α increase for solutions obtained with SAA and \mathcal{U}_d^Γ for an instance with ratio $\frac{b}{h} = 8$, $\sigma_\epsilon = 10$ and $s = 0$ for different values of β



(a) $s = 45$, $\beta_1 = 0.8$



(b) $s = 105$, $\beta_1 = 0.8$



(c) $s = 165$, $\beta_1 = 0.8$

Figure 5.2 – Evolution of the relative \hat{S}_α as α increase for solutions obtained with SAA and \mathcal{U}_d^Γ for an instance with ratio $\frac{b}{h} = 8$, $\sigma_\epsilon = 10$ and $\beta_1 = 0.8$ for different values of setup cost s

Table 5.5 – Costs obtained with SAA, $\mathcal{U}_\delta^\Gamma$ and \mathcal{U}_δ^{LR} for instances with ratio $\frac{b}{h} = 8$, $\sigma_\epsilon = 10$, $\beta_1 = 0.5$ when the fixed cost $s = 105$ and the maximum lead time L varies.

L	Average			Q0.95			Q0.99		
	SAA	$\mathcal{U}_\delta^\Gamma$	\mathcal{U}_δ^{LR}	SAA	$\mathcal{U}_\delta^\Gamma$	\mathcal{U}_δ^{LR}	SAA	$\mathcal{U}_\delta^\Gamma$	\mathcal{U}_δ^{LR}
1	1952.36	+7.71%	+1.41%	2500.95	-10.79%	-12.54%	3044.97	-18.04%	-19.02%
3	2935.96	+6.74%	+3.02%	4097.26	-7.61%	-10.04%	5351.84	-15.77%	-16.86%
5	3002.59	+4.87%	+3.13%	4132.72	-7.28%	-9.8%	5710.53	-13.79%	-13.85%

, where $Z_t = \frac{(\sum_{k=1}^n d_{t-k}) - \underline{D}}{\overline{D} - \underline{D}}$ and \overline{D} and \underline{D} are respectively the maximum and minimum value of the sum of n consecutive demands $\sum_{k=1}^n d_{t-k}$, evaluated on the training sample of $N + 1$ demands. For the current set of experiment, we set $n = 3$ periods which mean that the lead time of order depends on the level of the demand over the three previous periods.

Numerical results Table 5.5 shows the average cost obtained by the three methods applied on 50 instances sharing the same configuration. We consider a set of instances with $\beta_1 = 0.5$, $s = 105$ and for which the maximum lead time L vary from 1 to 5. The results in this table remain consistent with the previous observations as SAA formulations lead to the lowest average cost and the highest extreme costs. The average cost difference between the SAA and the classic robust solutions decreases when the maximum lead time increases, while the difference between SAA and the solutions provided by our method increases.

The joint demand and lead time uncertainty set presents clear advantages over the one used by budget-based robust model, as it is more robust and performs better on average than the latter. However in our set of experiments, the cost difference between both robust models decreases as the maximum lead time L increases. In particular when $L = 5$ the average cost difference is less than 2% between both robust solutions and the 0.95-quantiles of costs are separated by less than 3%.

5 Conclusion and perspectives

In this chapter, we study an inventory management problem in which the demand level and the lead time of orders are uncertain and can be correlated. We introduce a parametric approach to build uncertainty set in which demand is assumed to follow a stochastic

process. We use classical autoregressive and moving average models to formulate uncertainty sets that integrates information on the demand process such as autocorrelation.

We propose numerical experiments in which the demand follows an $AR(1)$ process and show that our approach exploiting relationships between consecutive demands outperforms a classic robust approach inspired Bertsimas and Sim (2003, 2004) applied directly to the demand values. When compared to a standard sample average approximation scheme, it significantly reduces the risk to incur an extremely high cost at the expense of a mild increase the average cost of solutions.

Our results suggest that lead time uncertainties affect negatively the competitiveness of using our $AR(1)$ based uncertainty set for demand uncertainty, in particular when the maximum lead time is high. Integrating a potential correlation between demand and lead time does not appear to provide a particular advantage in comparison to a classic budget-based uncertainty set when positive correlation is considered.

With this first study, we show that using statistical model such as ARMA (or, by extension, ARIMA) to integrate correlation information improves the performance of robust models when uncertain parameters are correlated. Several interesting extension of our work should be considered in future research. First of all, we restricted our experiment to an $AR(1)$ model but our framework can be applied to more general ARIMA models which may further advantage structured uncertainty sets over generic ones. An other direct extension of our work would be to consider the application of such uncertainty sets in a dynamic context in which ordering quantity are adapted as the planning horizon unroll and some uncertain parameters are realized. Finally, the application of this approach on a real case study would be particularly interesting to conclude on the practical relevance of the method in an industrial context.

CONCLUSIONS AND PERSPECTIVES

In this thesis we study the application of data-driven robust models on several optimization problems arising in supply chain management. The literature review presented in Chapter 1 highlights that this paradigm is still young and despite promising results in various application domains, the use of such data-driven models in supply chain management and production planning remains relatively scarce. The main difficulty that prevents these new data-driven robust models to translate into methods that are applicable in practice comes from their lack of scalability. Indeed data-driven robust optimization may not be a viable option in SCM considering the (potentially large) MIP formulations of the optimization problems that arise in this context.

In this manuscript, we demonstrate that using data-driven uncertainty sets to bound the behaviour of uncertain parameters of optimization problems more precisely can provide significant reduction of the conservatism of robust solutions to SCM problems. In addition, we address the scalability issue through the introduction of approximate data-driven approaches, with the objective to improve the tractability of the resulting robust optimization models. This section recalls the main conclusions drawn from our work, both in terms of scientific contributions and future research perspectives susceptible to turn these techniques into useful tools for industrial applications.

Chapter 3. Our study of the SVC-based uncertainty set exhibits the lack of scalability of this approach, despite its interest for various applications (see Chapter 1, Section 1.3). We identify the large number of addition variables and constraints introduced in SVC-based robust formulations as one of the main causes of this limitation and propose an alternative approximate model that reduces the size of the MILP formulations while maintaining a good-quality representation of uncertain parameters. We compare the solutions obtained by the SVC-based approach and our approximation on three distinct optimization problem (MILP) in order to demonstrate the efficiency of our approach.

Our computational results show that our approximation succeeds in producing solu-

tions that are comparable to the solution obtained with the SVC-based uncertainty set. For example, on a robust knapsack problem with uncertain correlated weights, our approximation provide solutions which objective values are less than 1% higher than the solutions of the SVC-based uncertainty set on most of the instances. On the other hand, the computation time needed to solve these new robust models is reduced by more than 90%. In addition, our approximation scheme uses a single parameter that scales with the dimension of uncertain parameters and efficiently controls the trade-off between the quality of the approximation and the computation time (i.e. the size of the robust counterpart).

As a result, this study extends the range of application of the SVC-based approach to more complex and larger optimization problems. In addition, as it rely on the same formulation as the initial uncertainty set, it is likely to be compatible with any other improvement of the SVC-based uncertainty set that have been proposed in the literature (see Chapter 1 Section 1.3 for examples).

Chapter 4 This chapter provides a practical application of the approximate SVC-based uncertainty set. We focus on a problem faced by a manufacturer who must plan the assembly of multiple products in order to meet a known demand over a finite planning horizon. The picking and delivery of the components necessary to obtain the end items are subcontracted to an external logistic service provider. The goal is to define the orders in order to respect the maximum time capacity dedicated to the corresponding picking operations when their duration is uncertain.

We propose a MIP formulation and develop two robust models that integrate uncertainties to solve this problem. The first one is based on the budget-based approach of Bertsimas and Sim (2003, 2004) and allows to reduce the impact of uncertain picking times on the assembly line by providing solutions with a higher demand satisfaction compared to a deterministic solution. As an alternative robust solution approach, we apply the approximate SVC-based one introduced in Chapter 3. A set of experiments on instances with up to 10 products and 20 components shows that our approximate SVC-based uncertainty set outperforms the classic budget based approach on this problem. Specifically, it simultaneously reduces significantly the risk of backloging demand and leads to a lower number of components held on the border of the assembly line. For example, on an instance with 5 components and 10 products over a planning horizon of 14 periods, our approach can reduce the average and extreme (0.95-quantile) total cost

incurred by producers by 20%.

Chapter 5 In the first two contributions, we make use and extend a data-driven robust model based on a non-parametric approach (i.e. no assumptions are made on the uncertain parameters). However, this type of approach requires a large amount of historical data to be efficient, which restricts their range of application in supply chain management. For example when the data consists of time-dependent uncertainties, e.g. in the form of a demand process, gathering large and reliable data set is often impossible as the process that drives the parameters may evolve over time (e.g. new products on the market, crisis, ...). In this context, we introduce a novel uncertainty set built upon the assumption that uncertain parameters are driven by a specific (but unknown) stochastic process that is commonly used in predictive approach.

To illustrate the efficiency of our approach, we study a simple single-item lot sizing problem under uncertain and autocorrelated demand and an extension of this problem integrating an uncertain, demand-dependent lead-time. In addition, we rely on a dedicated cutting plane method that is known to be less conservative than the dual reformulation approach to solve this problem (Thorsen and Yao (2017)).

Our experiments show that our uncertainty set results in a significant reduction of conservatism compared to a classic budget-based approach. The solutions that are found using our parametric data-driven approach are of comparable or lower extreme costs with a significant reduction of cost on average. We also compare the solutions obtained to those of a sample average approximation method whose scenarios integrate the same correlation information that our robust model. The experimental results show that the sample average approach provides the best performance on average but that it performs poorly in extreme cases. In addition, when the correlation is high, the performances of the sample average approach display more variability than our robust formulation. For some instances, the solution of the latter is better than the former both on average and in extreme cases.

In a second set of experiments, we integrate uncertain and demand-dependent lead time to the problem. The results show that this type of uncertainty significantly impacts the costs of the solutions obtained with all three methods, which all adopt a « robust » policy with regards to uncertain lead times. The conclusions drawn from the first set of experiments remain unchanged when we couple uncertain lead times with uncertain correlated demand.

Perspectives *Scalability of DDRO approaches:* The scalability issue is a major drawback of most data-driven robust optimization methods. We show in this thesis that approximating uncertainty sets may address this problem and extend the range of application of existing methods. While our proposition is only applicable to a specific type of uncertainty sets, we believe that the core idea of our method may be adapted to other families of uncertainty sets. Extending this framework may improve the range of applications for which robust counterparts remain tractable in practical settings. In addition, most existing studies on similar data-driven methods rely on the dual reformulation to obtain solutions of robust models and the existing works comparing dual reformulation and cutting plane approaches to solve robust models is restricted to classic uncertainty sets (i.e. budget based, ellipsoidal). Moreover, there is no consensus on which method should be used and the efficiency of both methods seems to vary among uncertainty sets and optimization problems (see Chapter 1 Section 1.1). Thus, solution methods based on cutting plane approaches could help in order to improve the scalability of DDRO methods.

As a small step on this path, we studied the difference of computation time to solve the robust binary knapsack problem of Chapter 1 with dual reformulation and with a cutting plane approach as described in Chapter 1 Section 1.1. Figure 6.1 shows the computation time needed to solve the binary knapsack problem with both methods. We present results obtained with the SVC-based uncertainty set (6.1a) and our approximation (6.1b). In Figure 6.1a each boxplot correspond to a set of 10 instances and the computation time evolve with the parameter ν of the SVC based approach. In Figure 6.1b each boxplot correspond to a set of 90 problems instances and the computation time evolves with the parameter K of our approximation approach. This first set of experimental results clearly shows that in this context, coupling the SVC-based uncertainty set or our approximation with a cutting plane approach greatly reduces the computation time needed to solve the resulting formulations.

Improved cutting plane algorithm for LSP: As described in Chapter 1 and Chapter 5, the dedicated cutting plane algorithm introduced by Bienstock and Ozbay (2008) for lot sizing problems solves a relaxed robust problem with only a subset of realizations selected from their uncertainty sets. At each iteration, a new scenario is added to the formulation in order to converge to the optimal solution of the true min-max robust problem. One limitation of this procedure is that given a subset of scenarios, there could exist multiple optimal solutions to the min-max problem, each with a different average objective value over the subset of scenarios. As an alternative to this pure « worst case »

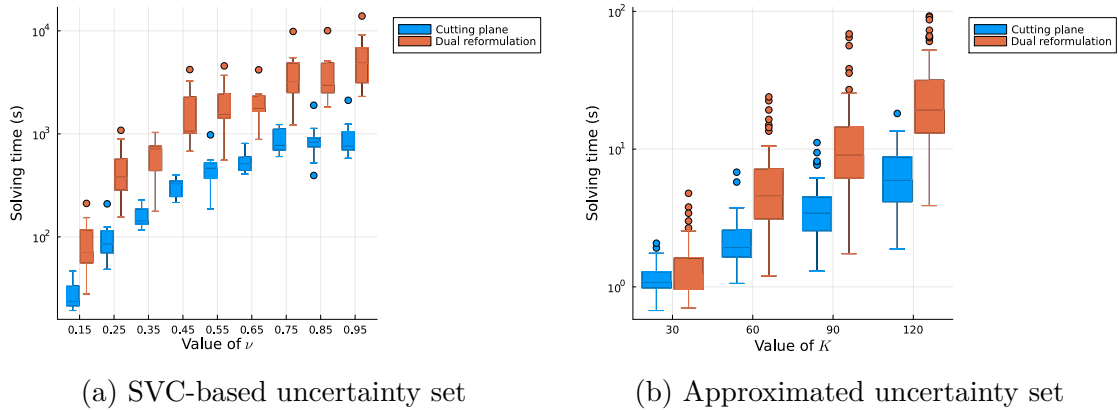


Figure 6.1 – Computation time of dual reformulation and cutting plane method for the binary knapsack problem with 30 items using the SVC-based uncertainty set and our approximation.

perspective, one could think of a bi-objective approach that focuses on the worst case performance as its first objective, then on the average cost over the subset of scenarios selected in a lexicographic order. The obtained solution would then have the same worst case performance but likely be less conservative on average. To go even further, one could also think of implementing this method into more general two stages robust problems solved with a column and constraint generation approach such as the one presented in Zeng (2011).

Adjustable data-driven robust optimization: This thesis focuss on static decision context in which all decisions are fixed prior to the realization of uncertain parameters. During the last two decades, adjustable robust models have been proposed to obtain solutions that depend on the realization of uncertain parameters (see Chapter 1 Section 2.3). Such solutions may be of interest for many applications in SCM and are known to be less conservative than static decisions making. Studies coupling such approaches with data-driven methods are still scarce and often applied to other fields. A direct extension of our work would then be to consider adjustable versions of the problem introduced in this thesis, e.g. to tackle the LSP from Chapter 5. For instance, a direct extension that consists in fixing the setup before the beginning of the planning horizon and adjusting the lot size decisions dynamically as the uncertainty is revealed would likely improve the solution quality drastically. In addition for this specific case, we expect that the more precise representation of the uncertainty developed in Chapter 5 would favor the dynamic counterparts of these static robust lot-sizing models.

Link with constraint satisfaction level: Most data-driven robust methods aims at defining an uncertainty set $\mathcal{U}(\mathcal{D})$ that covers a given portion of historical data \mathcal{D} . While this may give an accurate estimation of the probability that the realization of random parameters falls within such sets, this portion is often an underestimate of the probability of satisfaction for the corresponding robust constraint. A more interesting property would be to design sets that depend on both the uncertain parameters and the decision variables in order to define a probabilistic guarantee of constraint satisfaction instead. In other words, considering that the set \mathcal{D} describe the true distribution of uncertain parameters, a desirable property of a DDRO approach would be of the form:

$$\mathbb{P}(\mathbf{a} \in \mathcal{U}(\mathcal{D}, \mathbf{x})) \geq \mathbb{P}(\mathbf{a}^\top \mathbf{x} \leq b) \quad (6.1)$$

Existing works rely on the concept of robustness in distribution to derive ambiguity sets that allow to bound this probability, but they usually result computationally heavy models. In practice it is still difficult to estimate the true constraint satisfaction level obtained from a particular uncertainty set. Whether it is possible to design a data-driven uncertainty set that provides a tight probabilistic bound on a given constraint satisfaction thus remains an open research question. Such bounds would be of particular interest for practitioners in order to tune data-driven uncertainty set to their specific needs and avoid overconservative decisions.

APPENDICES

A Support Vector Clustering and Kernel function

In this section, we detail the formulation of the Quadratic Problem solved in the Support Vector Clustering algorithm and the kernel function introduced by Shang et al. (2017). The SVC algorithm seeks to minimize the radius R of a sphere enclosing a portion $1 - \nu$ of the available data $\{\phi(\mathbf{u}^{(i)})\}$ where $\phi(\cdot)$ is a mapping to a feature space \mathcal{F} in higher dimension. The corresponding optimization problem is given by:

$$\begin{aligned} \min \quad & R^2 + \frac{1}{N\nu} \sum_{i=1}^N \xi_i \\ \text{s.t.} \quad & \|\phi(\mathbf{u}^{(i)}) - \mathbf{p}\|^2 \leq R^2 + \xi_i \quad \mathbf{u}^{(i)} \in \mathcal{D} \\ & R \in \mathbb{R} \\ & \mathbf{p} \in \mathbb{R}^m \\ & \xi_i \geq 0 \quad i = 1, \dots, N \end{aligned}$$

where \mathbf{p} is the center of the sphere and $\boldsymbol{\xi}$ is a vector of slack variables defining if a given data-point $\mathbf{u}^{(i)} \in \mathcal{D}$ resides into the sphere (i.e. if $\xi_i > 0$ then $\mathbf{u}^{(i)}$ is outside of the sphere). Introducing lagrangian multipliers $\boldsymbol{\alpha}$ lead to the following equivalent dual quadratic problem:

$$\begin{aligned} \min \quad & \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j K(\mathbf{u}^{(i)}, \mathbf{u}^{(j)}) - \sum_{i=1}^N \alpha_i K(\mathbf{u}^{(i)}, \mathbf{u}^{(i)}) \\ \text{s.t.} \quad & 0 \leq \alpha_i \leq 1/N\nu \quad i = 1, \dots, N \\ & \sum_{i=1}^N \alpha_i = 1 \end{aligned}$$

where $K(\mathbf{u}^{(i)}, \mathbf{u}^{(j)}) = \phi(\mathbf{u}^{(i)})^T \phi(\mathbf{u}^{(j)})$ is the kernel function. In order to obtain a polyhedral description of the cluster, Shang et al. (2017) introduced the *Weighted General*

Intersection Kernel given by:

$$K(\mathbf{u}^{(i)}, \mathbf{u}^{(j)}) = \sum_{k=1}^m l_k - \|\mathbf{Q}(\mathbf{u}^{(i)} - \mathbf{u}^{(j)})\|_1 \quad (2)$$

where \mathbf{Q} is called the weighting matrix and is defined by $\mathbf{Q} = \mathbf{\Sigma}^{-\frac{1}{2}}$ with $\mathbf{\Sigma}$ the covariance matrix and where l_k is such that it satisfy the constraint $l_k > \max_{1 \leq i \leq N} \mathbf{q}_k^T \mathbf{u}^{(i)} - \min_{1 \leq i \leq N} \mathbf{q}_k^T \mathbf{u}^{(i)}$. From the optimal solution $\boldsymbol{\alpha}^*$ of this QP we identify a set of *support vectors* $\mathcal{S} = \{\mathbf{u}^{(i)} \mid \alpha_i^* > 0\}$ and a set of *boundary support vectors* $\mathcal{B} = \{\mathbf{u}^{(i)} \mid 0 < \alpha_i^* < 1/N\nu\}$ that are used to define the SVC-based uncertainty set as follows:

$$\mathcal{U}_{\text{SVC}} = \left\{ \mathbf{u} \mid \sum_{i \in \mathcal{S}} \alpha_i^* \|\mathbf{Q}(\mathbf{u} - \mathbf{u}^{(i)})\|_1 \leq \theta \right\}$$

, where $\theta = \max_{j \in \mathcal{B}} \sum_{i \in \mathcal{S}} \alpha_i^* \|\mathbf{Q}(\mathbf{u}^{(j)} - \mathbf{u}^{(i)})\|_1$.

BIBLIOGRAPHY

- Abdel-Aal, M. (2019). A robust capacitated lot sizing problem with setup times and overtime decisions with backordering allowed under demand uncertainty. *IFAC-PapersOnLine*, 52:589–594.
- Agra, A., Poss, M., and Santos, M. (2018). Optimizing make-to-stock policies through a robust lot-sizing model. *International Journal of Production Economics*, 200:302–310.
- Aharon, B.-T., Boaz, G., and Shimrit, S. (2009). Robust multi-echelon multi-period inventory control. *European Journal of Operational Research*, 199:922–935.
- Alem, D., Curcio, E., Amorim, P., and Almada-Lobo, B. (2018). A computational study of the general lot-sizing and scheduling model under demand uncertainty via robust and stochastic approaches. *Computers & Operations Research*, 90:125–141.
- Aouam, T. and Brahimi, N. (2013). Integrated production planning and order acceptance under uncertainty: A robust optimization approach. *European Journal of Operational Research*, 228(3):504–515.
- Atan, Z., Ahmadi, T., Stegehuis, C., de Kok, T., and Adan, I. (2017). Assemble-to-order systems: A review. *European Journal of Operational Research*, 261(3):866–879.
- Bandi, C. and Bertsimas, D. (2012). Tractable stochastic analysis in high dimensions via robust optimization. *Math. Program*, 134:23–70.
- Barlas, Y. and Gunduz, B. (2011). Demand forecasting and sharing strategies to reduce fluctuations and the bullwhip effect in supply chains. *The Journal of the Operational Research Society*, 62:458–473.
- Ben Ammar, O., Dolgui, A., Hnaien, F., and Louly, M. A. (2013). Supply planning and inventory control under lead time uncertainty: A review. *IFAC Proceedings Volumes*, 46(9):359–370.

- Ben-Tal, A., Golany, B., Nemirovski, A., and Vial, J.-P. (2005). Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7:248–271.
- Ben-Tal, A., Goryashko, A., and Guslitzer, E. Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99:351–376.
- Ben-Tal, A., Hazan, E., Koren, T., and Mannor, S. (2015). Oracle-based robust optimization via online learning. *Operations Research*, 63:628–638.
- Ben-Tal, A. and Nemirovski, A. (1998). Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805.
- Ben-Tal, A. and Nemirovski, A. (2000). Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming*, 88:411–424.
- Bertsimas, D., Dunning, I., and Lubin, M. (2016). Reformulation versus cutting-planes for robust optimization. *Computational Management Science*, 13:195–217.
- Bertsimas, D. and Georghiou, A. (2015). Design of near optimal decision rules in multi-stage adaptive mixed-integer optimization. *Operations Research*, 63(3):610–627.
- Bertsimas, D., Gupta, V., and Kallus, N. (2017). Data-driven robust optimization. *Mathematical Programming*, 167:235–292.
- Bertsimas, D., Pachamanova, D., and Sim, M. (2004). Robust linear optimization under general norms. *Operations Research Letters*, 32:510–516.
- Bertsimas, D. and Sim, M. (2003). Robust discrete optimization and network flows. *Mathematical programming*, 98(1-3):49–71.
- Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations Research*, 52:35–53.
- Bertsimas, D. and Thiele, A. (2004). A robust optimization approach to supply chain management. In *Integer Programming and Combinatorial Optimization*, pages 86–100, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Bertsimas, D. and Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, 54:150–168.

- Bienstock, D. and Ozbay, D. (2008). Computing robust basestock levels. *Discrete Optimization*, 5:389–414.
- Bookbinder, J. H. and Tan, J.-Y. (1998). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34:1096–1108.
- Boute, R. N., Disney, S. M., Lambrecht, M. R., and Van Houdt, B. (2014). Coordinating lead times and safety stocks under autocorrelated demand. *European Journal of Production Research*, 232:52–63.
- Brahimi, N., Absi, N., Dauzère-Pérès, S., and Nordli, A. (2017). Single-item dynamic lot-sizing problems: An updated survey. *European Journal of Operational Research*, 263(3):838–863.
- Brahimi, N., Dauzere-Peres, S., Najid, N. M., and Nordli, A. (2006). Single item lot sizing problems. *European Journal of Operational Research*, 168(1):1–16.
- Brandimarte, P. (2006). Multi-item capacitated lot-sizing with demand uncertainty. *International Journal of Production Research*, 44(15):2997–3022.
- Brockwell, P. J. and Davis, R. A. (2016). *Introduction to Time Series and Forecasting*. Springer Texts in Statistics. Springer International Publishing, Cham.
- Campbell, T. and How, J. P. (2015). Bayesian nonparametric set construction for robust optimization. In *2015 American Control Conference (ACC)*, pages 4216–4221.
- Charnes, J. M., Marmorstein, H., and Zinn, W. (1995). Safety stock determination with serially correlated demand in a periodic-review inventory system. *The Journal of the Operational Research Society*, 46(8):1006–1013.
- Chen, X., Sim, M., and Sun, P. (2007). A robust optimization perspective on stochastic programming. *Operation Research*, 55:1058–1071.
- Christopher, M. (2005). *Logistics and Supply Chain Management: Creating Value-Adding Networks (3rd ed.)*. FT Prentice Hall.
- Coniglio, S., Koster, A. M. C. A., and Spiekermann, N. (2018). Lot sizing with storage losses under demand uncertainty. *Journal of Combinatorial Optimization*, 36:763–788.

- Daneshvari, H. and Shafaei, R. (2021). A new correlated polyhedral uncertainty set for robust optimization. *Computers & Industrial Engineering*, 157:107346.
- Daryalal, M., Arslan, A. N., and Bodur, M. (2023). Two-stage and lagrangian dual decision rules for multistage adaptive robust optimization. *hal-04090602v1*. 09/19/2023.
- Disney, S., Farasyn, I., Lambrecht, M., Towill, D., and de Velde, W. V. (2006). Taming the bullwhip effect whilst watching customer service in a single supply chain echelon. *European Journal of Operational Research*, 173(1):151–172.
- Dolgui, A. and Proth, J.-M. (2010). *Supply Chain Engineering: Useful Methods and Techniques*. Springer London.
- Duc, T. T. H., Luong, H. T., and Kim, Y.-D. (2008). A measure of bullwhip effect in supply chains with a mixed autoregressive-moving average demand process. *European Journal of Operational Research*, 187(1):243–256.
- El Ghaoui, L., Oustry, F., and Lebret, H. (1998). Robust solutions to uncertain semidefinite programs. *SIAM Journal on Optimization*, 9(1):33–52.
- Eppen, G. D. and Martin, R. K. (1988). Determining safety stock in the presence of stochastic lead time and demand. *Management Science*, 34(11):1380–1390.
- Erkip, N., Hausman, W. H., and Nahmias, S. (1990). Optimal centralized ordering policies in multi-echelon inventory systems with correlated demand. *Management Science*, 36(3):381–392.
- Fischetti, M. and Monaci, M. (2012). Cutting plane versus compact formulations for uncertain (integer) linear programs. *Mathematical Programming Computation*, 4:239–273.
- Gabrel, V., Murat, C., and Thiele, A. (2014). Recent advances in robust optimization: An overview. *European Journal of Operational Research*, 235(3):471–483.
- Gallego, G. (1992). A minmax distribution free procedure for the (q, r) inventory model. *Operations Research Letters*, 11(1):55–60.
- Gicquel, C., Minoux, M., and Dallery, Y. (2008). Capacitated lot sizing models: a literature review. hal-00255830.

- Goerigk, M., Deghdak, K., and T'Kindt, V. (2015). A two-stage robustness approach to evacuation planning with buses. *Transportation Research Part B: Methodological*, 78:66–82.
- Goerigk, M. and Kurtz, J. (2023). Data-driven robust optimization using deep neural networks. *Computers & Operations Research*, 151:106087.
- Gumte, K. M., Devi Pantula, P., Miriyala, S. S., and Mitra, K. (2021a). Achieving wealth from bio-waste in a nationwide supply chain setup under uncertain environment through data driven robust optimization approach. *Journal of Cleaner Production*, 291:125702.
- Gumte, K. M., Devi Pantula, P., Miriyala, S. S., and Mitra, K. (2021b). Data driven robust optimization for handling uncertainty in supply chain planning models. *Chemical Engineering Science*, 246:116889.
- Hamed Mamani, Shima Nassiri, M. R. W. (2016). Closed-form solutions for robust inventory management. *Management Science*, 63(5):1625–1643.
- Han, B., Shang, C., and Huang, D. (2021). Multiple kernel learning-aided robust optimization: Learning algorithm, computational tractability, and usage in multi-stage decision-making. *European Journal of Operational Research*, 292:1004–1018.
- Harris, F. W. (1913). How many parts to make at once. *Factory: The Magazine of Management*, 10(2):135–136.
- Huang, K. and Küçüyavuz, S. (2008). On stochastic lot-sizing problems with random lead times. *Operations Research Letters*, 36(3):303–308.
- Jalilvand-Nejad, A., Shafaei, R., and Shahriari, H. (2016). Robust optimization under correlated polyhedral uncertainty set. *Computers & Industrial Engineering*, 92:82–94.
- José Alem, D. and Morabito, R. (2012). Production planning in furniture settings via robust optimization. *Computers & Operations Research*, 39(2):139–150.
- Karimi, B. and Fatemi, G. (2003). The capacitated lot sizing problem: a review of models and algorithms. *Omega*, 31:365–378.
- Kaufman, L. and Rousseeuw, P. J. (1990). Partitioning Around Medoids (Program PAM). In *Finding Groups in Data*, pages 68–125. John Wiley & Sons, Ltd.

- Keogh, E. and Mueen, A. (2017). Curse of Dimensionality. In Sammut, C. and Webb, G. I., editors, *Encyclopedia of Machine Learning and Data Mining*, pages 314–315. Springer US, Boston, MA.
- Lin, X., Zhao, L., Shang, C., He, W., Du, W., and Qian, F. (2022). Data-driven robust optimization for cyclic scheduling of ethylene cracking furnace system under uncertainty based on kernel learning. *Chemical Engineering Science*, 260.
- Loger, B., Dolgui, A., Lehuédé, F., and Massonnet, G. (2022). Improving the tractability of svc-based robust optimization. *IFAC-PapersOnLine*, 55(10):719–724.
- Luong, H. T. (2007). Measure of bullwhip effect in supply chains with autoregressive demand process. *European Journal of Operational Research*, 180(3):1086–1097.
- Melamed, M., Ben-Tal, A., and Golany, B. (2016). On the average performance of the adjustable ro and its use as an offline tool for multi-period production planning under uncertainty. *Computational Management Science*, 13:293–315.
- Metzker Soares, P. (2022). *L’optimisation robuste pour des problèmes de lotissement dans un contexte de rendement incertain*. PhD thesis, IMT Atlantique.
- Metzker Soares, P., Thevenin, S., Adulyasak, Y., and Dolgui, A. (2022). Robust optimization for lot sizing under yield uncertainty. *Computers and Operations Research*, 149:106025.
- Metzker Soares, P., Thevenin, S., Adulyasak, Y., and Dolgui, A. (2023). Adaptive robust optimization for lot-sizing under yield uncertainty. *European Journal of Operational Research*.
- Mohseni, S. and Pishvaei, M. S. (2020). Data-driven robust optimization for wastewater sludge-to-biodiesel supply chain design. *Computers & Industrial Engineering*, 139:105944.
- Montemanni, R. (2006). A benders decomposition approach for the robust spanning tree problem with interval data. *European Journal of Operational Research*, 174(3):1479–1490.
- Moon, I. and Gallego, G. (1994). Distribution free procedures for some inventory models. *Journal of the Operational Research Society*, 45:651–658.

- Mutapcic, A. and Boyd, S. P. (2009). Cutting-set methods for robust convex optimization with pessimizing oracles. *Optimization Methods and Software*, 24:381 – 406.
- Mínguez, R. and Casero-Alonso, V. (2019). On the convergence of cutting-plane methods for robust optimization with ellipsoidal uncertainty sets.
- Nabil Absi, S. K.-S. and Dauzère-Pérès, S. (2011). Uncapacitated lot-sizing problem with production time windows, early productions, backlogs and lost sales. *International Journal of Production Research*, 49(9):2551–2566.
- Ning, C. and You, F. (2017). Data-driven adaptive nested robust optimization: General modeling framework and efficient computational algorithm for decision making under uncertainty. *AIChE Journal*, 63(9):3790–3817.
- Ning, C. and You, F. (2018a). Data-driven adaptive robust optimization framework based on principal component analysis. In *2018 Annual American Control Conference (ACC)*, pages 3020–3025.
- Ning, C. and You, F. (2018b). Data-driven decision making under uncertainty integrating robust optimization with principal component analysis and kernel smoothing methods. *Computers & Chemical Engineering*, 112:190–210.
- Pachamanova, D. A. (2002). *A robust optimization approach to finance*. PhD thesis, Massachusetts Institute of Technology.
- Pahl, J., Voss, S., and Woodruff, D. (2007). Production planning with load dependent lead times: An update of research. *Annals of Operations Research*, 153:297–345.
- Pahl, J., Voß, S., and Woodruff, D. L. (2005). *Load Dependent Lead Times — From Empirical Evidence to Mathematical Modeling*, pages 539–554. Physica-Verlag HD.
- Pochet, Y. and Wolsey, L. A. (2006). *Production Planning by Mixed Integer Programming*. Springer New York.
- Poss, M. (2013). Robust combinatorial optimization with variable budgeted uncertainty. *4OR*, 11:75–92.
- Pätzold, J. and Schöbel, A. (2020). Approximate cutting plane approaches for exact solutions to robust optimization problems. *European Journal of Operational Research*, 284:20–30.

- Quezada, F., Gicquel, C., and Kedad-Sidhoum, S. (2019). Stochastic dual dynamic integer programming for a multi-echelon lot-sizing problem with remanufacturing and lost sales. In *IEEE International Conference on Control, Decision and Information Technologies CODIT 2019*, pages 1254–1259. IEEE.
- Santos, M. C., Poss, M., and Nace, D. (2018). A perfect information lower bound for robust lot-sizing problems. *Annals of Operations Research*, 271:887–913.
- Scarf, H. (1958). A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production*, Stanford University Press, pages 201–209.
- Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., and Williamson, R. C. (2001). Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471.
- Schubert, E. and Rousseeuw, P. J. (2019). Faster k-Medoids Clustering: Improving the PAM, CLARA, and CLARANS Algorithms. In Amato, G., Gennaro, C., Oria, V., and Radovanović, M., editors, *Similarity Search and Applications*, pages 171–187, Cham. Springer International Publishing.
- See, C.-T. and Sim, M. (2010). Robust approximation to multiperiod inventory management. *Operations Research*, 58:583–594.
- Shang, C., Chen, W.-H., Stroock, A. D., and You, F. (2020). Robust model predictive control of irrigation systems with active uncertainty learning and data analytics. *IEEE Transactions on Control Systems Technology*, 28(4):1493–1504.
- Shang, C., Huang, X., and You, F. (2017). Data-driven robust optimization based on kernel learning. *Computers & Chemical Engineering*, 106:464–479.
- Shang, C. and You, F. (2018). Robust optimization in high-dimensional data space with support vector clustering. *IFAC-PapersOnLine*, 51:19–24.
- Shang, C. and You, F. (2019). A data-driven robust optimization approach to scenario-based stochastic model predictive control. *Journal of Process Control*, 75:24–39.
- Shen, F., Zhao, L., Du, W., Zhong, W., and Qian, F. (2020). Large-scale industrial energy systems optimization under uncertainty: A data-driven robust optimization approach. *Applied Energy*, 259:114199.

- Simchi-Levi, D. (2008). *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*. McGraw-Hill/Irwin.
- Solyali, O., Cordeau, J.-F., and Laporte, G. (2016). The impact of modeling on robust inventory management under demand uncertainty. *Management Science*, 62(4):1188–1201.
- Song, J.-S. and Zipkin, P. (2003). Supply chain operations: Assemble-to-order systems. *Handbooks in Operations Research and Management Science*, 11:561–596.
- Soyster, A. L. (1973). Technical note — convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21(5):1154–1157.
- Stadtler, H., Kilger, C., and Meyr, H. (2015). *Supply Chain Management and Advanced Planning*. Springer Berlin, Heidelberg.
- Syntetos, A., Boylan, J., and Disney, S. (2008). Journal of operations research society. *50th Annual Conference of the Operational Research Society 2008, OR50*, 60:149–160.
- Sözüer, S. and Thiele, A. C. (2016). The State of Robust Optimization. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, International Series in Operations Research & Management Science, pages 89–112. Springer International Publishing.
- Taş, D., Gendreau, M., Jabali, O., and Jans, R. (2019). A capacitated lot sizing problem with stochastic setup times and overtime. *European Journal of Operational Research*, 273(1):146–159.
- Thevenin, S., Ben-Ammar, O., and Brahimi, N. (2022). Robust optimization approaches for purchase planning with supplier selection under lead time uncertainty. *European Journal of Operational Research*, 303(3):1199–1215.
- Thorsen, A. and Yao, T. (2017). Robust inventory control under demand and lead time uncertainty. *Annals of Operations Research*, 257:207–236.
- Udom, P. and Phumchusri, N. (2014). A comparison study between time series model and arima model for sales forecasting of distributor in plastic industry. *IOSR Journal of Engineering*, 4:32–38.

- Varas, M., Maturana, S., Pascual, R., Vargas, I., and Vera, J. (2014). Scheduling production for a sawmill: A robust optimization approach. *International Journal of Production Economics*, 150:37–51.
- Wagner, H. M. and Whitin, T. M. (1958). Dynamic version of the economic lot size model. *Management Science*, 5(1):89–96.
- Wu, W., Liu, R., Yang, Q., and Quek, T. Q. (2022). Learning-based robust resource allocation for d2d underlying cellular network. *IEEE Transactions on Wireless Communications*, 21(8):6731–6745.
- Zeng, B. (2011). Solving Two-stage Robust Optimization Problems by A Constraint-and-Column Generation Method. *Optimization Online*.
- Zhang, C., Wang, Z., and Wang, X. (2022). Machine learning-based data-driven robust optimization approach under uncertainty. *Journal of Process Control*, 115:1–11.
- Zhang, Y., Feng, Y., and Rong, G. (2018a). Data-driven rolling-horizon robust optimization for petrochemical scheduling using probability density contours. *Computers & Chemical Engineering*, 115:342–360.
- Zhang, Y., Jin, X., Feng, Y., and Rong, G. (2018b). Data-driven robust optimization under correlated uncertainty: A case study of production scheduling in ethylene plant. *Computers & Chemical Engineering*, 109:48–67.
- Zhao, L. and You, F. (2019). A data-driven approach for industrial utility systems optimization under uncertainty. *Energy*, 182:559–569.

PERSONAL PUBLICATIONS LIST

Submitted to an international journal

- B. Loger, A. Dolgui, F. Lehuédé, G. Massonnet. Approximate Kernel Learning Uncertainty Set for Robust Combinatorial Optimization. In press in *Informations Journal On Computing*.

International conference with proceedings

- B. Loger, A. Dolgui, F. Lehuédé, G. Massonnet. A robust data-driven approach to supply planning. APMS 2021: IFIP International Conference on Advances in Production Management Systems, IMT Atlantique, Sep 2021, Nantes, France.
- B. Loger, A. Dolgui, F. Lehuédé, G. Massonnet. Improving the tractability of SVC-based Robust Optimization. MIM 2022: 10th IFAC Conference on Manufacturing Modelling, Management and Control, Jun 2022, Nantes, France.

International conference without proceedings

- B. Loger, A. Dolgui, F. Lehuédé, G. Massonnet. Data-driven robust optimization approach to supply planning using an approximated uncertainty set. IWLS 2022: International Workshop on Lot-Sizing, Aug 2022, Oslo, Norway.

National conferences

- B. Loger, A. Dolgui, F. Lehuédé, G. Massonnet. Approximation d'un ensemble d'incertitude pour l'optimisation robuste dirigée par les données. ROADEF 2022 : 23^{ème} congrès annuel de la Société Française de Recherche Opérationnelle et d'Aide à la Décision. Villeurbanne - Lyon, France.

- B. Loger, A. Dolgui, F. Lehuédé, G. Massonnet, S. Minner, B. Sel. Robust inventory management under joint demand and lead-time uncertainty. ROADEF 2023 : 24^{ème} congrès annuel de la Société Française de Recherche Opérationnelle et d'Aide à la Décision. Rennes, France.

Titre : Modèles d'optimisation basés sur les données pour la planification des opérations dans les Supply Chain industrielles

Mot clés : Optimisation dirigées par les données, Gestion de la Supply Chain, Optimisation Robuste

Résumé : Face à la complexité croissante des chaînes logistiques, l'utilisation d'outils d'aide à la décision automatisés devient nécessaire pour appréhender les multiples sources d'incertitude susceptibles de les impacter tout en garantissant un niveau de performance élevé. Pour répondre à ces objectifs, les managers ont de plus en plus recours à des approches capables d'améliorer la résilience des chaînes logistiques en proposant des solutions robustes face aux aléas, pour garantir à la fois la qualité de service et la maîtrise des coûts liés à la production, au stockage et au transport de biens. Alors que la collecte et l'analyse des données occupe une place croissante dans la stratégie des entreprises, la bonne exploitation de ces informations afin de caractériser plus précisément ces incertitudes et leur impact sur les opérations devient un enjeu majeur pour optimiser les systèmes de production et de distribution modernes. Cette thèse se positionne au cœur de ces nouveaux dé-

fis en développant différentes méthodes d'optimisation mathématiques reposant sur l'utilisation de données historiques, dans le but de proposer des solutions robustes à plusieurs problèmes d'approvisionnement et de planification de production. Des expérimentations numériques sur des applications de natures diverses permettent de comparer ces nouvelles techniques à plusieurs autres approches classiques de la littérature pour valider leur pertinence en pratique. Les résultats obtenus démontrent l'intérêt de ces contributions, qui offrent des performances moyennes comparables tout en réduisant leur variabilité en contexte incertain. En particulier, les solutions restent satisfaisantes lorsqu'elles sont confrontées à des scénarios extrêmes, dont la probabilité d'apparition est faible. Enfin, les temps de résolution des procédures développées restent compétitifs et laissent envisager la possibilité d'une mise en œuvre sur des cas d'application à l'échelle industrielle.

Title: Data driven optimisation models for operations planning in industrial Supply chains

Keywords: Data Driven Optimization, Supply Chain Management, Robust Optimization

Abstract: With the increasing complexity of supply chains, automated decision-support tools become necessary in order to apprehend the multiple sources of uncertainty that may impact them, while maintaining a high level of performance. To meet these objectives, managers rely more and more on approaches capable of improving the resilience of supply chains by proposing robust solutions that remain valid despite uncertainty, to guarantee both a quality of service and a control of the costs induced by the production, storage and transportation of goods. As data collection and analysis become central to define the strategy of companies, a proper usage of this information to characterize more precisely these uncertainties and their impact on operations is becoming a major challenge for optimizing modern production and distribution systems. This thesis addresses these new challenges by developing different mathematical optimization methods based on historical data, with the aim of proposing robust solutions to several supply and production planning problems. To validate the practical relevance of these new techniques, numerical experiments on various applications compare them with several other classical approaches from the literature. The results obtained demonstrate the value of these contributions, which offer comparable average performance while reducing their variability in an uncertain context. In particular, the solutions remain satisfactory when confronted with extreme scenarios, whose probability of occurrence is low. Finally, the computational time of the procedures developed remain competitive, making them suitable for industrial-scale applications.