HEC MONTRÉAL École affiliée à l'Université de Montréal

Regret Minimization in Multistage Decision Problems

par Mehran Poursoltani

Thèse présentée en vue de l'obtention du grade de Ph. D. en administration (spécialisation Sciences de la décision)

août 2023

© Mehran Poursoltani, 2023

HEC MONTRÉAL École affiliée à l'Université de Montréal

Cette thèse intitulée :

Regret Minimization in Multistage Decision Problems

Présentée par :

Mehran Poursoltani

a été évaluée par un jury composé des personnes suivantes :

Okan Arslan HEC Montréal Président-rapporteur

Erick Delage HEC Montréal Directeur de recherche

Angelos Georghiou University of Cyprus Codirecteur de recherche

> Maxime Cohen McGill University Membre du jury

Merve Bodur University of Toronto Examinatrice externe

Olivier Bahn HEC Montréal Représentant du directeur de HEC Montréal

Résumé

Des études empiriques soutiennent l'idée que plusieurs décideurs sont averses au regret, i.e., qu'ils préfèrent éviter les actions qui pourraient potentiellement conduire à d'immenses sentiments de regret une fois que l'incertitude se sera dissipée et que les opportunités manquées seront connues. La minimisation du regret attire donc une attention importante dans différents domaines, allant de la psychologie et de l'économie à la recherche opérationnelle et à l'informatique. Cette thèse étudie la question de la minimisation des regrets du point de vue computationnel pour des contextes de décision à plusieurs étapes. La contribution de ce travail à la littérature existante réside dans le développement de nouveaux paradigmes de modélisation, l'avancement des méthodes de résolution et l'identification de phénomène intéressant dans les contextes de gestions d'inventaires, de trajectoire, et de portefeuilles financiers.

En termes de modélisation, nous étudions comment les problèmes de minimisation des regrets à deux étapes peuvent être représentés sous la forme de problèmes d'optimisation robuste à deux étapes. De plus, nous proposons un modèle de minimisation des regrets à plusieurs étages qui intègre deux concepts de manière innovante. Nous délaissons d'abord l'approche de minimisation du pire regret pour considérer plutôt une mesure qui intègre l'aversion au risque par l'entremise d'une mesure de risque cohérente et permettons ensuite au décideur d'imposer une structure d'information plus générique à la politique de référence, qui traditionnellement se veut clairvoyante. Enfin, nous étendons notre étude de la minisation du regret pour couvrir un contexte axé sur les données, dans lequel les décisions peuvent exploiter des informations secondaires. Nous introduisons un modèle d'optimisation stochastique conditionnel et robuste au choix de distribution qui optimise un "coefficient de prescriptivité", qui a récemment été proposé dans la littérature pour ce genre de problème.

L'optimisation numérique des problèmes de minisation du regret présente de manière générale un véritable défi computationnel. Bien qu'il soit possible d'identifier des approches tractables pour certaines classes simples du problème, jusqu'à présent, la plupart des implémentations de ce paradigme nécessitent la conception d'un schéma d'approximation spécifiquement dessiné pour le problème étudié. Malgré les efforts déployés pour améliorer la tractabilité des problèmes de minimisation des regrets, l'existence d'une approche générale pour produire des solutions exactes et approximatives de haute qualité est toujours une question ouverte dans la littérature. Cette thèse s'appuie sur des développements algorithmiques récents dans le domaine de l'optimisation robuste pour adresser cette grande famille de problèmes. Nos résultats sont appliqués à plusieurs domaines pratiques, incluant le problème du vendeur de journaux multi-éléments, le problème du transport de production, le problème de gestion des stocks multi-périodes, le problème de sélection du portefeuille et le problème de chemin le plus court. Grâce à ces applications, nous examinons le comportement des nouveaux modèles que nous proposons, l'efficacité des méthodes de résolution et la qualité des solutions obtenues.

Mots-clés

Prise de décision en plusieurs étapes; Minimisation des regrets; Optimisation robuste; Programmation stochastique, Optimisation contextuelle, Règles de décision affine, Mesures de risque.

Méthodes de recherche

Recherche quantitative; Programmation mathématique.

Abstract

Empirical studies support the idea that the decision makers are "regret-averse" in the sense that they avoid taking actions that might potentially lead to feeling regret once the uncertainties are gone and the missed opportunities are revealed. Regret minimization has drawn significant attention in different domains, ranging from psychology and economics to operations research and computer science. Broadly speaking, this thesis investigates the concept of regret minimization from a computational perspective for multistage decision contexts. The contribution of this work to the existing literature lies in the development of novel modeling paradigms, advancement of solution methods, and identification of interesting insights in the contexts of inventory, routing, and financial portfolio management.

In terms of modeling, we study how two-stage regret minimization problems can be represented as two-stage robust optimization problems. Furthermore, we propose a new multistage regret minimization model that integrates two concepts in an innovative way. We first abandon the worst-case regret minimization approach to consider a measure that integrates risk aversion through a coherent risk measure and then to allow the decision maker to impose an information structure that is more general compared to the reference policy, which is traditionally considered clairvoyant. Finally, we extend our study of regret minimization to cover a data-driven contextual decision making in which decisions can leverage side information. We introduce a conditional stochastic optimization model that is robust to the choice of distribution and optimizes a "coefficient of prescriptive-ness", which has recently been proposed in the literature for this kind of problem.

The numerical optimization of regret minimization problems is a real computational challenge in general. While it is possible to identify tractable approaches for some simple

classes of problems, thus far, most implementations of this paradigm require designing an approximation scheme that is specifically designed for the problem under study. Despite the efforts made to improve the tractability of regret minimization problems, the existence of a general approach for producing exact and approximate solutions of high quality is still an open question in the literature. This thesis builds on recent algorithmic developments in the field of robust optimization to address this large family of problems. Our findings are applied to several decision-making applications, including the multiitem newsvendor problem, production transportation problem, multi-period inventory management problem, portfolio selection problem, and shortest path problem. Thanks to these applications, we examine the behavior of the new models that we propose, the efficiency of the solution methods, and the quality of the solutions obtained.

Keywords

Multistage decision-making; Regret minimization; Robust optimization; Stochastic programming; Conditional optimization; Affine decision rules; Risk measures; Data-driven.

Research methods

Quantitative research; Mathematical programming.

Contents

R	ésum	é			iii
A	bstrad	ct			v
Li	st of '	Tables			xi
Li	st of I	Figures	5		xiii
A	cknov	wledge	ments		xv
Pı	eface				xvii
G	enera	l Intro	duction		1
	Refe	erences			6
1	Adj	ustable	e Robust (Optimization Reformulations of Two-Stage Worst-case Re-	
	gret	Minin	nization I	Problems	9
	Abs	tract .			9
	1.1	Introd	luction .		10
	1.2	Revie	w of the I	Literature on Regret Minimization	13
	1.3	Mode	rn Solutio	on Methods for Two-Stage Adjustable Robust Optimization	16
		1.3.1	The Cas	e of Fixed Recourse and Right-Hand Side Uncertainty	16
			1.3.1.1	The Column-and-Constraint Generation Method \ldots .	18
			1.3.1.2	Conservative Approximation Using Linear Decision Rules	19
			1.3.1.3	Other Solution Schemes	20
		1.3.2	The Cas	e of Objective Function Uncertainty	21

1.4	TSLRC	O Reformulations for Worst-case Absolute Regret Minimization Prob-	
	lems .		22
	1.4.1	The Case of Right-Hand Side Uncertainty	23
	1.4.2	The Case of Objective Uncertainty	25
1.5	TSLRC	O Reformulations for Worst-case Relative Regret Minimization Prob-	
	lems .		26
	1.5.1	The Case of Right-Hand Side Uncertainty	28
	1.5.2	The Case of Objective Uncertainty	29
1.6	Optim	ality of Affine Decision Rules	30
1.7	Nume	rical Results	34
	1.7.1	Multi-Item Newsvendor Problem	34
		1.7.1.1 Numerical Efficiency of AARC Compared with C&CG	37
		1.7.1.2 Value of Flexibility to Hindsight Decisions	39
		1.7.1.3 Decision Analysis	41
	1.7.2	Production-Transportation Problem	42
1.8	Apper	ndix	45
	1.8.1	Illustrative Example with Newsvendor Problem	45
	1.8.2	Summary Tables for the Literature on Regret Minimization	46
	1.8.3	Ning and You (2018)'s C&CG Approach Is an Optimistic Approxi-	
		mation	49
	1.8.4	Bertsimas and Dunning (2020)'s Conservative Approximation Is	
		Weaker than the Approximation Obtained with Problem (1.15) \ldots	49
	1.8.5	Zeng and Zhao (2013)'s Mixed-integer Linear Programming Refor-	
		mulation of C&CG's Sub-Problem	50
	1.8.6	Proof of Lemma 1.3.1	50
	1.8.7	Proof of Proposition 1.4.1	52
	1.8.8	Proof of Proposition 1.4.2	53
	1.8.9	Proof of Proposition 1.5.1	55
	1.8.10	Proof of Proposition 1.5.2	56
	1.8.11	Proof of Proposition 1.6.1	58
	1.8.12	Proof of Proposition 1.6.2	61

		1.8.13	Proof of Proposition 1.6.3	63
		1.8.14	TSLRO Reformulations for WCRRM in Cost Minimization Problems	67
			1.8.14.1 The Case of Right-Hand Side Uncertainty	68
			1.8.14.2 The Case of Objective Uncertainty	70
	Refe	erences		73
2	Risl	k-Avers	e Regret Minimization in Multistage Stochastic Programs	81
	Abs	tract .		81
	2.1	Introd	uction	81
	2.2	Literat	ture Review	84
	2.3	Δ -Reg	ret Minimization in Multistage Stochastic Programs	86
		2.3.1	The Δ -Regret	88
		2.3.2	Illustrative Example of Δ -Regret Model \ldots	89
	2.4	Proper	rties of Δ -Regret Model Under Risk Measures $\ldots \ldots \ldots \ldots$	91
		2.4.1	The Case of $\rho(\xi) = \operatorname{esssup}(\xi)$	92
		2.4.2	The Case of $ ho(\xi) = \mathbb{E}_{\mathbb{Q}}[\xi]$	95
		2.4.3	The Case of Coherent Risk Measures	96
	2.5	Nume	rical Experiments	98
	2.6	Apper	ndix	103
		2.6.1	Illustrative Example	103
		2.6.2	Proof of Lemma 2.3.1	105
		2.6.3	Proof of Lemma 2.4.1	106
		2.6.4	Proof of Theorem 2.4.1	107
		2.6.5	Proof of Proposition 2.4.1	109
		2.6.6	Proof of Corollary 2.4.1	110
		2.6.7	Proof of Proposition 2.4.2	110
		2.6.8	Proof of Corollary 2.4.2	111
		2.6.9	Proof of Proposition 2.4.3	111
		2.6.10	Column-and-Constraint Generation Algorithm	113
		2.6.11	Portfolio Management	115
			2.6.11.1 Two-Stage Portfolio Management	116

			2.6.11.2 Multistage Portfolio Management	118
		2.6.12	Parameters of Solved Instances	121
	Refe	erences	•••••••••••••••••••••••••••••••••••••••	123
3	Rob	oust Da	ta-Driven Prescriptiveness Optimization	129
	Abs	tract .		129
	3.1	Introd	luction	130
	3.2	Litera	iture Review	132
	3.3	Motiv	vation for Optimizing ${\mathcal P}$ and its Robustification $\ldots \ldots \ldots \ldots$	135
	3.4	Robus	st Data-Driven Prescriptiveness Optimization	136
		3.4.1	Convex Formulation for DRPCR	139
		3.4.2	The Nested CVaR Ambiguity Set \mathcal{D}	140
		3.4.3	A Bisection Algorithm for DRPCR	143
	3.5	Exper	iments	145
	3.6	Apper	ndix	152
	Refe	erences	• • • • • • • • • • • • • • • • • • • •	154
G	enera	l Concl	lusion	159

List of Tables

1.1	Multi-Item Newsvendor Problem - Uncorrelated Uncertainty Set	36
1.2	Multi-Item Newsvendor Problem - Correlated Uncertainty Set	38
1.3	Compare P&D and B&D Approaches - Multi-Item Newsvendor Problem with	
	Order Limits	40
1.4	Average Suboptimality of Solutions from RO, WCARM, and WCRRM with	
	Respect to RO, WCARM, and WCRRM Models Based on 120 Randomly Gen-	
	erated Instances of Three Different Sizes.	41
1.5	Production-Transportation Problem	44
1.6	WCRRM - Production-Transportation Problem with 7 Facilities and 14 Cus-	
	tomers	45
1.7	General Approaches - Linear Single-Stage Problems	47
1.8	General Approaches - Two-Stage Problems	48
2.1	Numerical Details of Example 2.1.1	83
2.2	Optimal Project Selected under Four Variants of Regret Minimization in Ex-	
	ample 2.1.1	83
2.3	Profit Function in Example 2.4.1	94
2.4	EP-WCR	103
2.5	EP-RAR	104
2.6	EA-RAR	105
2.7	Demand Realizations of Multistage Inventory Management Problem - $\pmb{\zeta}_{t-1}(\omega)$	121
2.8	Parameters of Multistage Inventory Management Problem	121
2.9	Solution Time for Multistage Inventory Management Problem	122

2.10	Price Realizations of Multistage Portfolio Management Problem - $\zeta_{t-1}(\omega)$	122
3.1	Mean Out-of-Sample Coefficient of Prescriptiveness	152

List of Figures

Comparison of Adaptation Power Between x (Beside the Timeline) and x' (on	
the Right of Each Tree) as a Function of Δ	91
Optimal Regret and Total First-Stage Production Achieved Given Different	
Levels of Conservatism (α).	99
Regret Breakdown for Different Δ -Regret Optimal Solutions	102
Optimal Regret for Different Values of Δ , α and r_f .	116
Optimal Investment Decisions in the Risk Free Asset for Δ -Regret and CVaR	
Minimization.	117
Regret and First-stage Investment in the Risk Free Account Given Different	
Levels of Conservatism (α).	119
Scenario Tree with $T = 5$, $ \Omega = 16$ and $\Delta \in \{0, 1, 2, 3, 4\}$.	122
Visualization of the Basic (Left) and Accelerated (Right) Bisection Algorithm	145
Visualization of the basic (Left) and Accelerated (Mght) Disection Augorithm.	145
Out-of-Sample Performance of Relaxed $\boldsymbol{x}(\cdot)$	150
Out-of-Sample Performance of Binary $\boldsymbol{x}(\cdot)$	151
	Comparison of Adaptation Power Between x (Beside the Timeline) and x' (on the Right of Each Tree) as a Function of Δ Optimal Regret and Total First-Stage Production Achieved Given Different Levels of Conservatism (α) Regret Breakdown for Different Δ -Regret Optimal Solutions Optimal Regret for Different Values of Δ , α and r_f Optimal Investment Decisions in the Risk Free Asset for Δ -Regret and CVaR Minimization Regret and First-stage Investment in the Risk Free Account Given Different Levels of Conservatism (α) Scenario Tree with $T = 5$, $ \Omega = 16$ and $\Delta \in \{0, 1, 2, 3, 4\}$ Visualization of the Basic (Left) and Accelerated (Right) Bisection Algorithm. Out-of-Sample Performance of Binary $x(\cdot)$

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Erick Delage. His contagious enthusiasm, thorough guidance, and constant accessibility were invaluable in guiding and inspiring me throughout my academic journey. Without his exceptional mentorship, this thesis could not have been possible. Secondly, I extend my utmost appreciation to my co-supervisor, Prof. Angelos Georghiou, whose pivotal roles as both a committee member during the earlier stages of my thesis and a precious mentor during its development have been invaluable. I will forever cherish his unwavering support and infinite patience, guiding me through intricate concepts with clarity and ease. Their collective mentorship and encouragement have been the cornerstone of my success, and I feel privileged to have had the opportunity to work with such dedicated and brilliant minds.

I extend my sincere thanks to the members of my thesis committee, Prof. Maxime Cohen, and Prof. Merve Bodur, for their precious time, thoughtful feedback, and profound expertise. Moreover, I wish to express my appreciation to the Fonds de recherche du Québec – Nature et technologies (FRQNT) for providing the funding that made this research possible.

Last but not least, my heartfelt thanks to my wife, Samira, and my parents, for their unconditional love, support, and encouragement.

Preface

This thesis consists of three articles listed as follows:

- Poursoltani, M. and Delage, E. (2022) Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems. *Operations Research*, 70(5):2906–2930.
- 2. Poursoltani, M., Delage, E., and Georghiou, A. (2023). Risk-averse regret minimization in multistage stochastic programs. *Operations Research*, 10.1287/opre.2022.2429.
- 3. Poursoltani, M., Delage, E., and Georghiou, A. (2023). Robust Data-driven Prescriptiveness Optimization. To be submitted.

General Introduction

Decision-making has always been an indispensable part of our personal and professional lives, giving rise to challenges in contexts with either large financial stakes or with large potential impact on the quality of life of individuals. This challenge is magnified when it comes to decisions that have long-term repercussions, where the information is revealed gradually, and actions can be deployed progressively as new observations are made. For-tunately, the field of Operations Research (OR) has been quite successful in the last few decades in assembling an essential toolbox of models and algorithms to support decision makers when they face irreducible uncertainties. In particular, whilst the stochastic programming paradigm exploits probability theory to optimize the trade-off between risk and returns, robust optimization has gained significant popularity by improving computation requirements through the optimization of the worst-case scenarios.

An appealing alternative to robust optimization consists in optimizing decisions using the notion of regret. Regret comes into play when the actual outcome is such that some actions achieve higher utility than the implemented one, leading individuals to experience a sense of disappointment or dissatisfaction with their choice. The introduction of regret minimization to decision theory is historically attributed to Savage (1951). Empirical evidence (Bleichrodt et al. 2010) supports that many decision-makers are "regretaverse" and corroborates the idea that optimizing regret generally leads to less conservative decisions than those produced by robust optimization (Perakis and Roels 2008, Natarajan et al. 2014, Caldentey et al. 2017). This has given rise to numerous applications of regret minimization in decision-making problems, including but not limited to portfolio selection (Lim et al. 2012), shortest path, subset selection (Natarajan et al. 2014), spanning tree, ranking problems (Audibert et al. 2014), and pricing and mechanism design (Caldentey et al. 2017, Koçyiğit et al. 2022); nevertheless, the numerical optimization of minimum regret problem is a real challenge in general. For instance, while expected value model (EVM) and robust optimization (RO) formulations are polynomially solvable in the case of a linear program with objective coefficients known to reside in a box uncertainty set, Averbakh and Lebedev (2005) proved that the alternative regret minimization problem becomes strongly NP-hard to solve. With the exception of some simple classes of the regret minimization problem, most implementations require specifically structured approximation solution methods, which in itself highlights the need for general exact and approximate solution schemes.

This concern becomes more serious when it comes to two-stage regret minimization problems, a case where the literature is quite scarce. Whereas most of the existing pieces of research (see Assavapokee et al. 2008a, Assavapokee et al. 2008b, Jiang et al. 2013, Ng 2013, Chen et al. 2014, Ning and You 2018) study two-stage worst-case regret minimization problem under specific uncertainty assumptions, their designed solution methods, in general, are based on mixed integer programs together with a column-and-constraint generation approach. On an exceptional basis, Bertsimas and Dunning (2020) exploit linear decision rules to provide a linear program that conservatively approximates the two-stage worst-case absolute regret minimization problem with right-hand side uncertainty. The lack of a general solution method for two-stage regret minimization problems, together with the availability of a rich pool of modern exact and approximate solution methods for two-stage robust optimization problems bring up the natural question of whether it is possible to reformulate the former class of problems as the latter or identify a polynomially solvable subclass of two-stage regret minimization problems. The first chapter of this thesis attempts to find an answer to these fundamental questions under both absolute and relative regret paradigms.

Most of the literature on regret minimization focuses on minimizing the worst-case regret, which is based solely on the support information of uncertain parameters. However, there have also recently been a few works that started investigating how distributional information and alternative risk measures can be used to measure regret. As a result of this extension, two distinct reformulations for regret minimization problems have implicitly emerged, namely ex-ante and ex-post regret. Natarajan et al. (2014) study ex-post regret using a worst-case Conditional Value-at-Risk (CVaR) measure. In this case, the decisionmaker benchmarks her decision against one that has access to the future realization of the unknown parameter (the optimal hindsight decision) and measures the worst-case CVaR of this difference. On the contrary, Yue et al. (2006), and Perakis and Roels (2008) use the ex-ante formulation that employs a worst-case risk measure. They employ the worst-case expected value to measure regret and rather interpret it as the expected value of distribution information (EVDI) or the maximum value of stochastic modeling (see Delage et al. 2014). In the ex-ante regret formulation, the decisions are compared against a benchmark decision, which only knows the distribution of the uncertain parameter instead of its realization.

The primary distinction between ex-ante and ex-post regret formulations is the information structure imposed on the benchmark decision. This distinction raises several interesting questions: How do these information structures influence the decisions and the optimal values of regret minimization models? How do different risk measures behave in the contexts of ex-ante and ex-post regret formulations? Can we identify additional information structures that potentially interpolate between these two extreme cases in a multistage decision making context? Chapter 2 of this thesis tackles these questions by conducting theoretical and numerical analyses of a new risk-averse multistage regret minimization model, called the Δ -regret model. More specifically, in a multistage setting, the Δ -regret model explores the idea of evaluating regret under a popular risk measure, where a benchmark policy can only benefit from foreseeing Δ steps into the future. We demonstrate how this model incorporates ex-ante and ex-post regret formulations as special cases.

Recently, the community of stochastic programming has been interested in a new family of models known as contextual optimization models, where a decision maker has access to side information that can be exploited to better predict the unknown parameters and lead, in turn, to higher quality decisions. This raises the question of how regret models should be adapted to address this contextual setting and, perhaps more importantly in a data-driven context, how they can be designed in a way that makes them robust to estimation error. Chapter 3 addresses these two questions by proposing a new distributionally robust contextual relative regret minimization model. We call this the

distributionally robust prescriptiveness competitive ratio problem as it aims at optimizing a robust form of the so-called coefficient of prescriptiveness introduced in Bertsimas and Kallus (2020). Intuitively, our prescriptiveness competitive ratio compares the difference between the performance of a proposed policy and the optimal hindsight decision, which we may refer to as our policy's ex-post regret, to the ex-post regret achieved by a reference policy that does not exploit the side-information. This optimization problem further has interesting connections to both the famous coefficient of determination used in statistics, and the Δ -regret model proposed in Chapter 2.

The main contributions of this thesis can be summarized as follows. In the first chapter, we study both two-stage worst-case absolute and relative regret minimization problems in which uncertainty impacts either the right-hand side or objective function and reformulate them as two-stage worst-case robust optimization problems. We identify conditions under which one can obtain exact and conservative approximate solutions by exploiting popular adjustable robust optimization decomposition schemes and linear decision rules. Our experiments for a multi-item newsvendor problem and a production transportation problem provide evidence to support the high quality of solutions obtained through employing linear decision rules; furthermore, we establish a subclass of regret minimization problems for which these approximate decisions are proven exact. In particular, this discovery leads to the identification of a class of two-stage worst-case regret minimization problems for which we know a general solution method with polynomial solution time.

The second chapter of this thesis contributes to the regret minimization literature in multiple directions. In a multistage stochastic programming setting with a discrete probability distribution, we explore the idea of risk-averse regret minimization. We introduce the Δ -regret model, which allows the decision-maker to impose the arbitrary information structure on the benchmark policy. In particular, in a *T*-period problem, setting $\Delta = 0$ reduces to the so-called ex-ante regret minimization problem, where the benchmark policy has similar information as the decision-maker; on the contrary, adjusting $\Delta = T - 1$ leads to the ex-post regret minimization problem in which the benchmark policy knows all the future realizations of the uncertain parameter; and for any other integer value that falls in between, the Δ -regret model interpolates between the ex-ante and ex-post mod-

els, allowing the benchmark policy get access to up to Δ stages ahead information. We provide theoretical and experimental insights on how the Δ -regret model behaves under popular risk measures for different values of Δ .

In the third chapter, we consider a data-driven context where the side information can potentially lead to more anticipative decisions. We propose a novel distributionally robust conditional stochastic optimization model where a prescriptiveness competitive ratio replaces the classical conditional stochastic programming objective. We reformulate this problem as a convex optimization problem and demonstrate how it reduces to a linear program when nested CVaR represents the ambiguity set. A bisection algorithm is proposed for solving this problem which can be further expedited through an acceleration scheme. We discuss how this novel conditional optimization model is connected to the Δ -regret model studied in Chapter 2; more specifically, we show that this model spans both $\Delta = 1$ (ex-post) and, in some sense, $\Delta = -1$ regret minimization models. Since, in the latter case, the decision maker has access to the side information, indeed she has more extensive information compared to the ex-ante regret minimization model ($\Delta = 0$). On the experimental side, we study a shortest path problem and evaluate the robustness of the resulting decisions against alternative methods when the out-of-sample dataset experiences a distribution shift compared to the in-sample dataset.

References

- Assavapokee, T., Realff, M., and Ammons, J. (2008a). Min-max regret robust optimization approach on interval data uncertainty. *Journal of Optimization Theory and Applications*, 137(2):297–316.
- Assavapokee, T., Realff, M., Ammons, J., and Hong, I. (2008b). Scenario relaxation algorithm for finite scenario-based min-max regret and min-max relative regret robust optimization. *Computers & Operations Research*, 35(6):2093–2102.
- Audibert, J., Bubeck, S., and Lugosi, G. (2014). Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45.
- Averbakh, I. and Lebedev, V. (2005). On the complexity of minmax regret linear programming. *European Journal of Operational Research*, 160(1):227–231.
- Bertsimas, D. and Dunning, I. (2020). Relative robust and adaptive optimization. *IN*-FORMS Journal on Computing, 32(2):408–427.
- Bertsimas, D. and Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044.
- Bleichrodt, H., Cillo, A., and Diecidue, E. (2010). A quantitative measurement of regret theory. *Management Science*, 56(1):161–175.
- Caldentey, R., Liu, Y., and Lobel, I. (2017). Intertemporal pricing under minimax regret. *Operations Research*, 65(1):104–129.
- Chen, B., Wang, J., Wang, L., He, Y., and Wang, Z. (2014). Robust optimization for transmission expansion planning: Minimax cost vs. minimax regret. *IEEE Transactions on Power Systems*, 29(6):3069–3077.
- Delage, E., Arroyo, S., and Ye, Y. (2014). The value of stochastic modeling in two-stage stochastic programs with cost uncertainty. *Operations Research*, 62(6):1377–1393.
- Jiang, R., Wang, J., Zhang, M., and Guan, Y. (2013). Two-stage minimax regret robust unit commitment. *IEEE Transactions on Power Systems*, 28(3):2271–2282.

- Koçyiğit, Ç., Rujeerapaiboon, N., and Kuhn, D. (2022). Robust multidimensional pricing: Separation without regret. *Mathematical Programming*, 196:841–874.
- Lim, A. E., Shanthikumar, J. G., and Vahn, G.-Y. (2012). Robust portfolio choice with learning in the framework of regret: Single-period case. *Management Science*, 58(9):1732–1746.
- Natarajan, K., Shi, D., and Toh, K.-C. (2014). A probabilistic model for minmax regret in combinatorial optimization. *Operations Research*, 62(1):160–181.
- Ng, T. (2013). Robust regret for uncertain linear programs with application to coproduction models. *European Journal of Operational Research*, 227(3):483–493.
- Ning, C. and You, F. (2018). Adaptive robust optimization with minimax regret criterion: Multiobjective optimization framework and computational algorithm for planning and scheduling under uncertainty. *Computers & Chemical Engineering*, 108:425–447.
- Perakis, G. and Roels, G. (2008). Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203.
- Savage, L. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67.
- Yue, J., Chen, B., and Wang, M.-C. (2006). Expected value of distribution information for the newsvendor problem. *Operations Research*, 54(6):1128–1136.

Chapter 1

Adjustable Robust Optimization Reformulations of Two-Stage Worst-case Regret Minimization Problems

Abstract

This chapter explores the idea that two-stage worst-case regret minimization problems with either objective or right-hand side uncertainty can be reformulated as two-stage robust optimization problems and can therefore benefit from the solution schemes and theoretical knowledge that have been developed in the last decade for this class of problems. In particular, we identify conditions under which a first-stage decision can be obtained either exactly using popular adjustable robust optimization decomposition schemes, or approximately by conservatively employing affine decision rules. Furthermore, we provide both numerical and theoretical evidence that in practice the first-stage decision obtained using affine decision rules is of high quality. Initially, this is done by establishing mild conditions under which these decisions can be proven exact, which effectively extends the space of regret minimization problems known to be solvable in polynomial time. We further evaluate both the computational efficiency of this tractable approximation scheme and the sub-optimality of the resulting policies on a multi-item newsvendor problem and a production transportation problem.

1.1 Introduction

When employing optimization in the context of uncertainty, a well-known alternative to minimizing expected value or the worst-case scenario, a.k.a. expected value model (EVM) and robust optimization (RO) respectively, consists in minimizing the regret experienced once the decision maker finds out that another action would have achieved a better performance under the realized scenario. Historically, while the paradigm of worst-case absolute regret minimization is usually attributed to Savage (1951), it became a legitimate representation of preferences through its axiomatization in Milnor (1954) and more comprehensively in Stoye (2011). Empirical studies (e.g. in Loomes and Sugden 1982 and in Bleichrodt et al. 2010) have also supported the idea that some decision makers are "regret averse" in the sense that they are inclined to abandon alternatives that might lead to large regret once they realize what would have been the best actions in hindsight. In the operations research literature, there is recently a growing number of studies that describe regret minimization models as leading to less "conservative" decisions than those produced by robust optimization (Perakis and Roels, 2008; Aissi et al., 2009; Natarajan et al., 2014; Caldentey et al., 2017). In particular, this reduced conservatism, which is often considered as the Achilles' heel of robust optimization, is achieved without requiring the assumption of knowing an underlying distribution. In support of this popular belief, we refer interested readers to Section 1.8.1 where it is shown that, in a simple newsvendor problem, orders made by a regret averse agent are always of larger magnitude than those proposed by robust optimization.

An important obstacle in the application of regret minimization models resides in the fact that they can give rise to a serious computational challenge. In particular, while both EVM and RO formulations are polynomially solvable in the case of a linear program with objective coefficients known to reside in their respective interval (a.k.a. box uncertainty), Averbakh and Lebedev (2005) demonstrated that the worst-case regret minimization problem is strongly NP-hard. While there have been extensive efforts invested in the development of exact and approximate solution schemes, most of these focus on specific applications of single-stage mixed-integer programs (e.g., shortest path, knapsack, single-period portfolio optimization). More recently, some attention was driven towards general forms of two-stage continuous/mixed-integer linear programs but, with the exception of Bertsimas and Dunning (2020) who applied affine decision rules to a facility location problem with right-hand side uncertainty, there has been no general tractable conservative approximation scheme proposed for these models. In comparison, while two-stage robust optimization problem is also known to be strongly NP-hard when uncertainty appears in the constraints (see Guslitser 2002, Minoux 2009), there has been active research in the last 10 years about deriving and analyzing tractable solution schemes for some of the most general forms of the problem (see for instance Yanikoglu et al. 2018 for a recent survey). Moreover, these efforts have led to the development of software packages (e.g., ROME in Goh and Sim 2011 and JuMPeR in Dunning et al. 2017) that facilitate the implementation of these solution schemes and certainly promoted its use in applications. Among these different schemes, there is no doubt that the most popular one, which was initially proposed in Ben-Tal et al. (2004) and will be referred to as the linear decision rule approach (as popularized in Kuhn et al. 2011), approximates the second-stage decision with a decision rule that is affine with respect to the uncertain parameters.

Generally speaking, this chapter explores both theoretically and numerically the idea that regret minimization problems can be reformulated as two-stage robust optimization problems and can therefore benefit from the tractable solution schemes and theoretical knowledge that has been developed in the last decade for this class of problems. In particular, we make the following contributions:

• We establish for the first time how, in a general two-stage linear programming setting with either objective or right-hand side uncertainty, both worst-case absolute regret minimization and worst-case relative regret minimization problems can be reformulated as a two-stage robust linear program. We also identify weak conditions on the regret minimization problems under which a tractable conservative approximation can be obtained by employing the concept of affine decision rules. Alternatively, we state conditions under which an exact solution can be obtained using the column-and-constraint generation algorithm proposed in Zeng and Zhao (2013) or in Ayoub and Poss (2016).

- We establish mild conditions on the regret minimization problem under which the theory developed in Bertsimas and Goyal (2012) and Ardestani-Jaafari and Delage (2016) can be exploited to demonstrate that the solution obtained using affine decision rules is exact. These results effectively both extend the class of regret minimization problems for which a polynomial-time solution method is known to exist and support the claim that in practice affine decision rules identify solutions of high quality.
- We present the results of numerical experiments that provide further evidence that the solutions obtained using affine decision rules are of high quality. In particular, we investigate both the computational efficiency of the solution methods and sub-optimality of such approximate first-stage decisions in multi-item newsvendor problems and production-transportation problems. We also illustrate how much improvement can be achieved in terms of worst-case regret by passing from a robust solution to a regret minimizing solution.

The rest of the chapter is composed as follows. Section 1.2 reviews the relevant literature and highlights the relevance of our proposed reformulations. Section 1.3 introduces the notation of two-stage linear programming models and summarizes some relevant results from the literature on two-stage robust optimization models. Section 1.4 proposes a two-stage robust optimization reformulation for two-stage worst-case absolute regret minimization with right-hand side uncertainty and for the one with objective uncertainty. Section 1.5 presents analogous results for the case of relative regret. Section 1.6 identifies conditions under which the use of affine decision rules in the robust optimization reformulations identifies exactly optimal first-stage decisions. Section 1.7 presents our numerical experiments. Finally, all proofs and additional materials are deferred to Section 1.8.

1.2 Review of the Literature on Regret Minimization

The computational challenges related to solving combinatorial worst-case regret minimization problems have been extensively tackled in the recent literature (see two comprehensive surveys Kouvelis and Yu (1996); Aissi et al. (2009) and references therein). In the domain of continuous decision variables, most research has focused on the singlestage version of the problem. In particular, a small number of single-stage linear regret minimization problems are known to be polynomial-time solvable. As presented in Gabrel and Murat (2010) and more recently in Bertsimas and Dunning (2020), this is the case for general linear programs with right-hand side and polyhedral uncertainty since these problems can be reformulated as equivalent linear programs. Averbakh (2004) also identifies an $O(n \log(n))$ algorithm for solving the minimum regret problem in resource allocation problems with objective and interval uncertainty. This approach is improved to linear time by Conde (2005) for the continuous knapsack problem. Nevertheless, the case of a general single-stage linear program with interval objective function uncertainty is known to be strongly NP-hard (see Averbakh and Lebedev 2005) and has motivated many algorithmic developments. First, Inuiguchi and Kume (1994), Inuiguchi and Sakawa (1995), and Inuiguchi and Sakawa (1997a) proposed to tackle the worst-case regret minimization problem by replacing the box uncertainty set with the list of its extreme points, and inserting these points progressively using a constraint generation procedure. In order to speed up the identification of violated constraints, Inuiguchi and Sakawa (1996) replace the exhaustive search with a branch-and-bound procedure that effectively solves a mixed-integer linear programming (MILP) formulation of the regret maximization subproblem. This MILP reformulation is further improved in Mausser and Laguna (1998) by exploiting the piecewise linear structure of the problem and a fast heuristic for identifying strong cuts is proposed in Mausser and Laguna (1999a), who also ported the constraint generation scheme to relative regret problems in Mausser and Laguna (1999b). The constraint generation procedure was extended for the first time to general polyhedral uncertainty in Inuiguchi and Sakawa (1997b) yet its numerical efficiency was further improved using an outer approximation scheme in Inuiguchi et al. (1999), and a cutting hyperplanes scheme in Inuiguchi and Tanino (2001). A summary of this prior work on single-stage problems is presented in Table 1.7 in Section 1.8.2.

In comparison with single-stage, the work on two-stage linear programs is rather scarce. First, in terms of application-specific methods, one might consider Vairaktarakis (2000) which proposes a linear time algorithm to solve multi-item newsvendor absolute and relative regret minimization problems with interval demand uncertainty and which proposes a dynamic programming approach for the NP-hard case of scenario-based uncertainty. Yue et al. (2006) and Perakis and Roels (2008) define closed-form solutions for the stochastic version of this problem with only one item, absolute regret, and distribution ambiguity, while Zhu et al. (2013) extend some of these results to the relative regret form. Zhang (2011) also studied a related two-stage uncapacitated lot-sizing problem with binary first-stage decisions and interval uncertainty on demands and identified a dynamic programming method that provides optimal solutions in polynomial time.

Table 1.8 (in Section 1.8.2) summarizes studies that propose general solution schemes. Specifically, Assavapokee et al. (2008b) consider two-stage worst-case absolute and relative regret minimization problems with binary first-stage decisions, continuous recourse variables, and scenario-based parametric uncertainty. The proposed approach is a precursor of the column-and-constraint generation (C&CG) algorithm in Zeng and Zhao (2013) as it relies on progressively introducing worst-case scenarios (found using an exhaustive search) in a master problem that optimizes both the first-stage decisions and recourse decisions for this subset of scenarios. This C&CG approach is extended to the case of box uncertainty set in Assavapokee et al. (2008a) where uncertainty only affects the right-hand side of constraints and the coefficients that are multiplied to first-stage decisions. This allows the authors to solve the regret maximization subproblem using two MILP reformulations that respectively generate feasibility and optimality cuts. This C&CG is further extended to the case of polyhedral uncertainty in Jiang et al. (2013) where the subproblem is solved approximately using coordinate ascent, and in Chen et al. (2014), which successfully identifies an exact MILP reformulation when uncertainty only affects the right-hand side of constraints.

Ng (2013) investigates problems that minimize the sum of linearly penalized perturbed constraint violations, which are special cases of two-stage linear worst-case absolute regret minimization problem with polyhedral uncertainty. The author proposes a
conservative *approximation* that takes the form of a two-stage robust optimization problem yet remains intractable. He employs a constraint generation scheme which involves solving a MILP at each iteration. Note that while the reformulations that we propose in Sections 1.4 and 1.5 will similarly lead to two-stage robust optimization models, our reformulations will be *exact* and available whether absolute or relative regret is considered. Furthermore, by using affine decision rules, our proposed conservative approximation models will be tractable in the sense that they can be reformulated as linear programs of comparable size.

More recently, Bertsimas and Dunning (2020) used a facility location problem to illustrate how affine decision rules can be used to conservatively approximate two-stage absolute regret minimization problems with right-hand side uncertainty. In contrast, our proposed conservative approximation will be in general tighter and applicable whether uncertainty lies in the objective function or the constraints. We further identify for the first time mild conditions under which our proposed conservative approximations and the one used in Bertsimas and Dunning (2020) are exact.

Finally, Ning and You (2018) suggested reformulating two-stage problems with righthand side polyhedral uncertainty exactly as two-stage robust optimization models yet did not extend this procedure to relative regret or to problems with objective uncertainty as we will present. The authors also mistakenly assume that worst-case scenarios always occur at extreme points of the polyhedral uncertainty set. This is in turn used to formulate a MILP that generates violated constraints in a C&CG approach effectively providing an optimistic approximation to the regret minimization problem (see Section 1.8.3 for an example). Finally, a distinguishing feature of our work will be to describe for the first time how linear decision rules can be tractably employed to obtain conservative solutions for a large family of two-stage regret minimization problems, and conditions under which such decision rules actually return exact solutions.

1.3 Modern Solution Methods for Two-Stage Adjustable Robust Optimization

In this section, we introduce our notation and present a number of modern solution methods that have appeared in recent literature concerning two-stage robust linear optimization problem. While the version of this model with right-hand side uncertainty is known to be intractable, we survey methods that either seek optimal solutions, conservative approximation, or optimistic bounds. We later present the case of objective uncertainty for which there is a tractable reformulation.

1.3.1 The Case of Fixed Recourse and Right-Hand Side Uncertainty

In this section, the focus is on the following two-stage linear robust optimization model with fixed recourse (TSLRO):

$$(TSLRO) \quad \underset{\boldsymbol{x},\boldsymbol{y}(\cdot)}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta}\in\mathcal{U}} \quad (C\boldsymbol{\zeta}+\boldsymbol{c})^T\boldsymbol{x} + \boldsymbol{d}^T\boldsymbol{y}(\boldsymbol{\zeta}) + \boldsymbol{f}^T\boldsymbol{\zeta} \tag{1.1a}$$

subject to
$$A\boldsymbol{x} + B\boldsymbol{y}(\boldsymbol{\zeta}) \leq \Psi(\boldsymbol{x})\boldsymbol{\zeta} + \boldsymbol{\psi}, \,\forall\, \boldsymbol{\zeta} \in \mathcal{U}$$
 (1.1b)

$$x \in \mathcal{X}$$
, (1.1c)

where $x \in \mathbb{R}^{n_x}$ is the first-stage decision vector implemented immediately while y: $\mathbb{R}^{n_{\zeta}} \to \mathbb{R}^{n_y}$ is a strategy for the second-stage decision vector that is implemented only once the vector of uncertain parameters $\zeta \in \mathbb{R}^{n_{\zeta}}$ has been revealed. Furthermore, we have that $C \in \mathbb{R}^{n_x \times n_{\zeta}}$, $c \in \mathbb{R}^{n_x}$, $d \in \mathbb{R}^{n_y}$, $f \in \mathbb{R}^{n_{\zeta}}$, $A \in \mathbb{R}^{m \times n_x}$ and $B \in \mathbb{R}^{m \times n_y}$, and assume that $\psi \in \mathbb{R}^m$ and $\Psi : \mathbb{R}^{n_x} \to \mathbb{R}^{m \times n_{\zeta}}$ is an affine mapping of x. Note that d and B are not affected by uncertainty which is also referred to as satisfying the fixed recourse property. Finally, we assume that both \mathcal{X} and \mathcal{U} are non-empty polyhedra such that when the latter is bounded one retrieves the more common $\min_{\zeta \in \mathcal{U}}$ notation.

A special kind of TSLRO model emerges when the uncertain vector $\boldsymbol{\zeta}$ only influences the right-hand side of constraint (1.1b) and gives rise to the following definition.

Definition 1.3.1 A TSLRO problem is considered to have "right-hand side uncertainty" when C = 0, f = 0, and $\Psi(x) = \Psi$.

TSLRO problems with right-hand side uncertainty arise for instance in a number of inventory management, and logistics problems (see Melamed et al. 2016; Kim and Chung 2017; Simchi-Levi et al. 2019).

The TSLRO problem can also be equivalently reformulated in a form where the two stages of decisions are made explicit:

$$(TSLRO) \quad \underset{\boldsymbol{x} \in \mathcal{X}}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta} \in \mathcal{U}} h(\boldsymbol{x}, \boldsymbol{\zeta}), \tag{1.2a}$$

where $h(\boldsymbol{x}, \boldsymbol{\zeta})$ is defined as:

$$h(\boldsymbol{x},\boldsymbol{\zeta}) := \sup_{\boldsymbol{y}} (C\boldsymbol{\zeta} + \boldsymbol{c})^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} + \boldsymbol{f}^T \boldsymbol{\zeta}$$
(1.3a)

s.t.
$$A\boldsymbol{x} + B\boldsymbol{y} \le \Psi(\boldsymbol{x})\boldsymbol{\zeta} + \boldsymbol{\psi}$$
. (1.3b)

In Ben-Tal et al. (2004), the authors established that the TSLRO problem is NP-hard in general due to the so-called "adversarial problem", i.e. $\inf_{\zeta \in \mathcal{U}} h(\boldsymbol{x}, \zeta)$, which reduces to the minimization of a piecewise linear concave function over an arbitrary polyhedron. Since this seminal work, a number of methods have been proposed to circumvent this issue. We present a subset of these methods in the rest of this section where it will be useful to refer to some of the following assumptions.

Assumption 1.3.1 The sets \mathcal{X} and \mathcal{U} are non-empty polyhedra of the respective form $\mathcal{X} := \{x \in \mathbb{R}^{n_x} | Wx \leq v\}$, with $W \in \mathbb{R}^{r \times n_x}$ and $v \in \mathbb{R}^r$, and $\mathcal{U} := \{\zeta \in \mathbb{R}^{n_{\zeta}} | P\zeta \leq q\}$, with $P \in \mathbb{R}^{s \times n_{\zeta}}$ and $q \in \mathbb{R}^s$. Furthermore, there exists a triplet (x, ζ, y) such that $x \in \mathcal{X}, \zeta \in \mathcal{U}$, and $Ax + By \leq \Psi(x)\zeta + \psi$.

Assumption 1.3.2 The feasible set \mathcal{X} is such that it is always possible to identify a recourse action y that will satisfy all the constraints under any realization $\zeta \in \mathcal{U}$, a property commonly referred as "relatively complete recourse". Specifically:

$$\mathcal{X} \subseteq \{ \boldsymbol{x} \in \mathbb{R}^{n_x} | \forall \boldsymbol{\zeta} \in \mathcal{U}, \exists \boldsymbol{y} \in \mathbb{R}^{n_y}, A\boldsymbol{x} + B\boldsymbol{y} \le \Psi(\boldsymbol{x})\boldsymbol{\zeta} + \boldsymbol{\psi} \}.$$
(1.4)

Assumption 1.3.3 For all $x \in \mathcal{X}$ there exists a $\zeta \in \mathcal{U}$ such that the recourse problem (1.3) is bounded. In other words, this assumes that the TSLRO problem is bounded.

1.3.1.1 The Column-and-Constraint Generation Method

A so-called column-and-constraint generation (C&CG) method was proposed in Zeng and Zhao (2013) to identify an exact solution for the TSLRO problem. Specifically, in its simplest form this method can be applied when Assumptions 1.3.1, 1.3.2, and 1.3.3 are satisfied together with the following assumption.

Assumption 1.3.4 For all feasible first-stage decisions, there is a lower bound on the worst-case profit achievable, i.e. for all $x \in \mathcal{X}$, $\inf_{\zeta \in \mathcal{U}} h(x, \zeta) > -\infty$.

In particular, the latter assumption is straightforwardly met when the uncertainty set \mathcal{U} is bounded. The C&CG method then exploits the fact that $h(\boldsymbol{x}, \boldsymbol{\zeta})$ is convex with respect to $\boldsymbol{\zeta}$ to reformulate Problem (1.2) equivalently as :

$$\underset{\boldsymbol{x} \in \mathcal{X}}{\operatorname{maximize}} \quad \min_{\boldsymbol{\zeta} \in \mathcal{U}_v} h(\boldsymbol{x}, \boldsymbol{\zeta}) \,,$$

where $U_v = {\{\bar{\zeta}_1, \bar{\zeta}_2, ..., \bar{\zeta}_K\}}$ is the set of vertices of U, i.e. $U = \text{ConvexHull}(U_v)$ when U is bounded. This allows one to approximate the TSLRO problem as a restricted master problem:

$$\max_{\boldsymbol{x}, \{\boldsymbol{y}_k\}_{k \in \mathcal{K}'}} \min_{k \in \mathcal{K}'} c(\bar{\boldsymbol{\zeta}}_k)^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y}_k + \boldsymbol{f}^T \bar{\boldsymbol{\zeta}}_k$$
(1.5a)

subject to
$$A\boldsymbol{x} + B\boldsymbol{y}_k \leq \Psi(\boldsymbol{x})\bar{\boldsymbol{\zeta}}_k + \boldsymbol{\psi}, \,\forall \, k \in \mathcal{K}'$$
 (1.5b)

$$x \in \mathcal{X}$$
, (1.5c)

where $\mathcal{K}' \subseteq \{1, 2, ..., K\}$ such that $\mathcal{U}'_v = \{\bar{\zeta}_k\}_{k \in \mathcal{K}'} \subseteq \mathcal{U}_v$, and where each $y_k \in \mathbb{R}^{n_y}$. Problem (1.5) provides an upper bound for the optimal value of the TSLRO problem. This bound can be further tightened by introducing additional vertices in \mathcal{K}' . Given any $x \in \mathcal{X}$ that is optimal with respect to Problem (1.5), one can identify an additional worstcase vertex by solving the NP-hard adversarial problem $\min_{\zeta \in \mathcal{U}_v} h(x, \zeta)$. The algorithm will converge after a number of iterations that is necessarily less or equal to the number of vertices, K. Recently, it has become common practice (see problem (15)–(20) in Zeng and Zhao 2013) to reformulate the adversarial problem as a mixed-integer linear program. We refer interested readers to Section 1.8.5 for a description of this MILP.

1.3.1.2 Conservative Approximation Using Linear Decision Rules

A common approach (initially proposed in Ben-Tal et al. 2004) for formulating a tractable approximation of the TSLRO problem consists in restricting $y(\cdot)$ to take the form of an affine policy $y(\zeta) := Y\zeta + y$, where $Y \in \mathbb{R}^{n_y \times n_\zeta}$ and $y \in \mathbb{R}^{n_y}$. This gives rise to what is commonly referred as the affinely adjustable robust counterpart (AARC) model:

$$(AARC) \max_{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{y}, \boldsymbol{Y}} \inf_{\boldsymbol{\zeta} \in \mathcal{U}} (C\boldsymbol{\zeta} + \boldsymbol{c})^T \boldsymbol{x} + \boldsymbol{d}^T (Y \boldsymbol{\zeta} + \boldsymbol{y}) + \boldsymbol{f}^T \boldsymbol{\zeta}$$
(1.6a)

subject to
$$A\boldsymbol{x} + B(Y\boldsymbol{\zeta} + \boldsymbol{y}) \leq \Psi(\boldsymbol{x})\boldsymbol{\zeta} + \boldsymbol{\psi}, \,\forall \, \boldsymbol{\zeta} \in \mathcal{U}.$$
 (1.6b)

It is said that the AARC problem conservatively approximates the TSLRO problem since it identifies a solution pair $(\hat{x}, \hat{y}(\cdot))$ that is necessarily feasible according to the TSLRO model and since its optimal value provides a lower bound on the optimal value of the TSLRO problem.

A linear programming reformulation of Problem (1.6) can be obtained by exploiting Assumption 1.3.1, which ensures that \mathcal{U} is non-empty, together with the principles of duality theory. Indeed, this gives rise to Problem (1.6)'s so called equivalent robust counterpart:

$$\underset{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{y}, \boldsymbol{Y}, \boldsymbol{\Lambda}, \boldsymbol{\lambda}}{\text{maximize}} \quad \boldsymbol{c}^{T} \boldsymbol{x} + \boldsymbol{d}^{T} \boldsymbol{y} - \boldsymbol{q}^{T} \boldsymbol{\lambda}$$
 (1.7a)

subject to
$$C^T \boldsymbol{x} + Y^T \boldsymbol{d} + \boldsymbol{f} + P^T \boldsymbol{\lambda} = 0$$
 (1.7b)

$$A\boldsymbol{x} + B\boldsymbol{y} - \boldsymbol{\psi} + \Lambda \boldsymbol{q} \le 0 \tag{1.7c}$$

$$\Psi(\boldsymbol{x}) - BY + \Lambda P = 0 \tag{1.7d}$$

$$\Lambda \ge 0, \lambda \ge 0, \tag{1.7e}$$

where $\lambda \in \mathbb{R}^s$ and $\Lambda \in \mathbb{R}^{m \times s}$ are the dual variables that arise when applying duality to the objective function (1.6a) and each constraint of (1.6b), respectively.

In the last decade, a number of theoretical and empirical arguments have reinforced a prevailing belief that linear decision rules provide high-quality solutions to TSLRO problems. One might for instance refer to Bertsimas et al. (2010b) and Ardestani-Jaafari and Delage (2016) for conditions under which this approach is exact.

1.3.1.3 Other Solution Schemes

There exists a rich pool of additional methods that have been proposed to solve TSLRO problems of the form presented in Problem (1.1). While we encourage the reader to refer to Delage and Iancu (2015) and Yanikoglu et al. (2018) for a more exhaustive description, we summarize below the main categories of approach.

In terms of exact methods, it is worth mentioning the work of Ayoub and Poss (2016), which provides a second column-and-constraint generation algorithm for deriving the exact solutions of TSLRO problems where C = 0 and d = 0. This algorithm is particularly useful for problems where Assumption 1.3.2 is violated.

In terms of approximation methods, Kuhn et al. (2011) show how linear decision rules can also be applied on a dual maximization problem associated to the TSLRO to obtain lower bounds on its optimal value. Alternatively, one can also obtain lower bounds by replacing \mathcal{U} with a finite subset of carefully selected scenarios (see Hadjiyiannis et al. 2011). Regarding conservative approximations, Chen et al. (2008) and Chen and Zhang (2009) explain how to employ piecewise linear (a.k.a. segregated) decision rules, while Ben-Tal et al. (2009) and Bertsimas et al. (2011) investigate the use of quadratic and polynomial decision rules, respectively. To improve the quality of solutions obtained using structured decision rules, Zhen et al. (2018) propose to eliminate some adjustable variables while Ardestani-Jaafari and Delage (2020) recommend reformulating an equivalent "complete recourse" problems.

Interestingly, it was recently observed in Bertsimas and de Ruiter (2016) that any TSLRO problem could be equivalently reformulated as a "dualized" TSLRO. The authors show empirically that this can improve numerical efficiency when using affine decision rules. This also allows them to obtain tighter lower bounds on TSLRO by exploiting the idea of Hadjiyiannis et al. (2011) on both versions of the TSLRO. One might also suspect that methods such as C&CG could perform differently whether they are applied on the original TSLRO or its dualized form.

Finally, an important recent methodological development consists in deriving exact copositive programming reformulations for the TSLRO problem (see Xu and Burer 2018 and Hanasusanto and Kuhn 2018). While copositive programming is known to be NP- hard in general, there are known hierarchies of tractable approximation models for these mathematical programs that will eventually identify an exactly optimal solution.

1.3.2 The Case of Objective Function Uncertainty

An alternative class of two-stage robust linear optimization problems makes the assumption that the uncertainty is limited to the objective function. This is summarized in the following formulation:

$$\max_{\boldsymbol{x}\in\mathcal{X},\boldsymbol{y}(\cdot)} \quad \inf_{\boldsymbol{\zeta}\in\mathcal{U}} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T(\boldsymbol{\zeta})\boldsymbol{y}(\boldsymbol{\zeta})$$
(1.8a)

subject to
$$A\boldsymbol{x} + B\boldsymbol{y}(\boldsymbol{\zeta}) \leq \boldsymbol{\psi}, \, \forall \, \boldsymbol{\zeta} \in \mathcal{U},$$
 (1.8b)

where $d : \mathbb{R}^{n_{\zeta}} \to \mathbb{R}^{n_y}$ is assumed to be an affine mapping of ζ , i.e. that we can characterize it in the form $d(\zeta) := D\zeta + d$, for some $D \in \mathbb{R}^{n_y \times n_{\zeta}}$ and $d \in \mathbb{R}^{n_y}$.

Remark 1.3.1 Note that Problem (1.8) can also accommodate situations where c is uncertain simply by lifting the space of second-stage decisions. Namely,

where $y_x : \mathbb{R}^{n_{\zeta}} \to \mathbb{R}^{n_x}$ and $y_y : \mathbb{R}^{n_{\zeta}} \to \mathbb{R}^{n_y}$. In this work, we adopt the more concise definition to simplify the exposition.

As for the case of TSLRO, the model can be reformulated in a format that emphasizes the dynamics:

where the recourse problem is defined as:

$$h(\boldsymbol{x},\boldsymbol{\zeta}) := \sup_{\boldsymbol{y}} c^T \boldsymbol{x} + \boldsymbol{d}^T(\boldsymbol{\zeta}) \boldsymbol{y}$$
(1.10a)

s.t. $A\boldsymbol{x} + B\boldsymbol{y} \le \boldsymbol{\psi}$. (1.10b)

From a computational perspective, it is interesting to consider the case where Assumptions 1.3.1, 1.3.2 and 1.3.3 are applicable. In particular, Assumption 1.3.2, which was referred as relatively complete recourse, simply reduces to the fact that $\mathcal{X} \subseteq \{x \in \mathbb{R}^{n_x} | \exists y \in \mathbb{R}^{n_y}, Ax + By \leq \psi\}$. Under these conditions, Problem (1.8) becomes more appealing than the TSLRO problem in (1.1) as one can easily verify that it can be reformulated as an equivalent linear program.

Lemma 1.3.1 *Given that Assumptions 1.3.1, 1.3.2 and 1.3.3 are satisfied, Problem (1.8) can be reformulated as the following equivalent linear program:*

$$\underset{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{y}, \boldsymbol{\lambda}}{\operatorname{maximize}} \quad \boldsymbol{c}^{T} \boldsymbol{x} + \boldsymbol{d}^{T} \boldsymbol{y} - \boldsymbol{q}^{T} \boldsymbol{\lambda}$$
 (1.11a)

subject to
$$A\boldsymbol{x} + B\boldsymbol{y} \le \boldsymbol{\psi}$$
 (1.11b)

$$P^T \boldsymbol{\lambda} + D^T \boldsymbol{y} = 0 \tag{1.11c}$$

$$\boldsymbol{\lambda} \ge 0, \qquad (1.11d)$$

where $\lambda \in \mathbb{R}^{s}$.

1.4 TSLRO Reformulations for Worst-case Absolute Regret Minimization Problems

As defined in Savage (1951), the worst-case absolute regret criterion aims at evaluating the performance of a decision x with respect to the so-called "worst-case regret" that might be experienced in hindsight when comparing x to the best decision that could have been made. Mathematically speaking, given a profit function $h(x, \zeta)$, which depends on both the decision and the realization of some uncertain vector of parameters ζ , one measures the regret experienced once ζ is revealed as the difference between the best profit achievable $\sup_{x' \in \mathcal{X}} h(x', \zeta)$ and the profit $h(x, \zeta)$ achieved by the decision x that was implemented. The worst-case absolute regret minimization (WCARM) problem thus takes the form:

(WCARM) minimize
$$\sup_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) - h(\boldsymbol{x},\boldsymbol{\zeta}) \right\},$$
 (1.12)

which is well-defined when one makes the assumption that the best profit achievable in hindsight never reaches infinity under any scenario for ζ .

Assumption 1.4.1 The best profit achievable is bounded, i.e., $\sup_{\zeta \in U, x \in \mathcal{X}} h(x, \zeta) < \infty$.

Assumption 1.4.1 is a natural condition to impose on the WCARM problem and implies Assumption 1.3.3. When Assumption 1.4.1 is not known to be satisfied, we will interpret the WCARM model as:

$$\min_{\boldsymbol{x} \in \mathcal{X}} \sup_{\boldsymbol{\zeta} \in \mathcal{U}} \left\{ \sup_{\boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \inf_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} \boldsymbol{c}^T(\boldsymbol{x}' - \boldsymbol{x}) + \boldsymbol{d}^T(\boldsymbol{y}' - \boldsymbol{y}) \right\} \,,$$

where $\mathcal{Y}(x, \zeta)$ is the set of feasible second-stage decisions given that x and ζ have realized, and interpret the fact that WCARM is unbounded as indicating that the optimal worst-case absolute regret is zero since there exists an $x \in \mathcal{X}$ such that for all $\zeta \in \mathcal{U}$ there is a way of reaching an arbitrarily large profit.¹ There is therefore no absolute regret under any circumstances when implementing such an x.

While we encourage interested readers to read an extensive review of the recent work regarding this problem formulation in Aissi et al. (2009), in what follows we demonstrate how the WCARM problem can be reformulated as a TSLRO problem when the profit function $h(\boldsymbol{x}, \boldsymbol{\zeta})$ captures the profit of a second-stage linear decision model with either right-hand side or objective uncertainty.

1.4.1 The Case of Right-Hand Side Uncertainty

We consider the case where $h(x, \zeta)$ takes the form presented in Problem (1.3) and where uncertainty is limited to the right-hand side as defined in Definition 1.3.1.

Proposition 1.4.1 *Given that Assumption 1.3.1 is satisfied, the WCARM problem with righthand side uncertainty is equivalent to the following TSLRO problem:*

subject to
$$A\boldsymbol{x} + B\boldsymbol{y}'(\boldsymbol{\zeta}') \leq \Psi'\boldsymbol{\zeta}' + \boldsymbol{\psi}, \,\forall\, \boldsymbol{\zeta}' \in \mathcal{U}',$$
 (1.13b)

where $\boldsymbol{\zeta}' \in \mathbb{R}^{n_{\zeta}+n_x+n_y}$, $\boldsymbol{y}' : \mathbb{R}^{n_{\zeta}+n_x+n_y} \to \mathbb{R}^{n_y}$, $\boldsymbol{f}' = [\boldsymbol{0}^T - \boldsymbol{c}^T - \boldsymbol{d}^T]^T$, and $\Psi' := \begin{bmatrix} \Psi & 0 & 0 \end{bmatrix}$, while \mathcal{U}' is defined as the new uncertainty set:

$$\mathcal{U}' := \{ \boldsymbol{\zeta}' \in \mathbb{R}^{n_{\boldsymbol{\zeta}} + n_x + n_y} | P' \boldsymbol{\zeta}' \le \boldsymbol{q}' \}$$
(1.14)

with

$$P' := \begin{bmatrix} P & 0 & 0 \\ 0 & W & 0 \\ -\Psi & A & B \end{bmatrix}, \quad and \quad q' := \begin{bmatrix} q \\ v \\ \psi \end{bmatrix}.$$

Furthermore, this TSLRO reformulation naturally satisfies Assumption 1.3.1, but also satisfies Assumptions 1.3.2 and 1.3.3 if the WCARM problem satisfies Assumption 1.3.2 and Assumptions 1.3.2 and 1.3.3 respectively, and satisfies Assumption 1.3.4 if the WCARM problem satisfies Assumptions 1.3.4 and 1.4.1.

Proposition 1.4.1 states that the WCARM model with right-hand side uncertainty can be reformulated as a TSLRO problem. This is interesting because it implies that it can benefit from the exact solution methods and conservative approximations discussed in Sections 1.3.1.1, 1.3.1.2, and 1.3.1.3. As an example, we provide below how affine decision rules can be applied to this reformulation.

Corollary 1.4.1 *Given that Assumption 1.3.1 is satisfied, the WCARM problem with right-hand side uncertainty is conservatively approximated by*

$$\min_{\boldsymbol{x} \in \mathcal{X}, \boldsymbol{y}, \boldsymbol{Y}', \boldsymbol{\Lambda}', \boldsymbol{\lambda}'} \quad -\boldsymbol{c}^T \boldsymbol{x} - \boldsymbol{d}^T \boldsymbol{y} + \boldsymbol{q}'^T \boldsymbol{\lambda}'$$
(1.15a)

subject to
$$Y'^T \boldsymbol{d} + \boldsymbol{f}' + P'^T \boldsymbol{\lambda}' = 0$$
 (1.15b)

$$A\boldsymbol{x} + B\boldsymbol{y} - \boldsymbol{\psi} + \Lambda' \boldsymbol{q}' \le 0 \tag{1.15c}$$

$$\Psi' - BY' + \Lambda' P' = 0 \tag{1.15d}$$

$$\Lambda' \ge 0, \, \boldsymbol{\lambda}' \ge 0 \,, \tag{1.15e}$$

where $Y' \in \mathbb{R}^{n_y \times n_\zeta + n_x + n_y}$, $\Lambda' \in \mathbb{R}^{m \times s + r + m}$, and $\lambda' \in \mathbb{R}^{s + r + m}$.

It is worth noting that to obtain the reformulation presented in Corollary 1.4.1, one needs to employ decision rules of the form $\mathbf{y}'(\zeta') := Y'\zeta' + \mathbf{y} = Y_{\zeta}\zeta + Y_x\mathbf{x}' + Y_y\mathbf{y}' + \mathbf{y}$, for some $Y_{\zeta} \in \mathbb{R}^{n_y \times n_{\zeta}}$, $Y_x \in \mathbb{R}^{n_y \times n_x}$, and $Y_y \in \mathbb{R}^{n_y \times n_y}$, and where $(\mathbf{x}', \mathbf{y}')$ captures the best pair of actions one would have implemented if he had a-priori information about ζ . Furthermore, one can easily show that the conservative approximation presented in (1.15) is at least as tight as the conservative approximation proposed in Bertsimas and Dunning (2020) given that the latter employs affine decision rules of the form $y'(\zeta') := Y_{\zeta}\zeta + y$. Section 1.8.4 further presents an example of two-item newsvendor problem where the bound obtained with Problem (1.15) is strictly tighter.

If one is more interested in applying an exact method for solving WCARM, then as long as the WCARM problem satisfies Assumptions 1.3.1, 1.3.2, 1.3.3, 1.3.4, and 1.4.1, based on Proposition 1.4.1 one can straightforwardly apply the column-and-constraint generation algorithm proposed in Section 1.3.1.1 to the TSLRO Problem (1.13).

1.4.2 The Case of Objective Uncertainty

We consider the case where $h(x, \zeta)$ takes the form presented in Problem (1.10).

Proposition 1.4.2 *Given that Assumptions 1.3.1 and 1.3.2 are satisfied, the WCARM problem with objective uncertainty is equivalent to the following TSLRO problem:*

$$\underset{\boldsymbol{x},\boldsymbol{y}'(\cdot)}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta}'\in\mathcal{U}'} \ (C'\boldsymbol{\zeta}'+\boldsymbol{c})^T\boldsymbol{x} + \boldsymbol{d}'^T\boldsymbol{y}'(\boldsymbol{\zeta}') + \boldsymbol{f}'^T\boldsymbol{\zeta}'$$
(1.16a)

subject to
$$A' \boldsymbol{x} + B' \boldsymbol{y}'(\boldsymbol{\zeta}') \le \Psi' \boldsymbol{\zeta}' + \boldsymbol{\psi}'$$
 (1.16b)

$$x \in \mathcal{X},$$
 (1.16c)

where $y' : \mathbb{R}^{n_{\zeta}+m} \to \mathbb{R}^{m+r}$, while \mathcal{U}' is defined as the new uncertainty set:

$$\mathcal{U}' := \{ \boldsymbol{\zeta}' \in \mathbb{R}^{n_{\boldsymbol{\zeta}} + m} | P' \boldsymbol{\zeta}' \le \boldsymbol{q}' \}$$
(1.17)

with

$$P' := \begin{bmatrix} P & 0 \\ -D & B^T \\ D & -B^T \end{bmatrix}, \quad and \quad q' := \begin{bmatrix} q \\ d \\ -d \end{bmatrix},$$

and where the matrices

$$C' := \begin{bmatrix} 0 & -A^T \end{bmatrix}, \quad d' := \begin{bmatrix} -\psi \\ -v \end{bmatrix}, \qquad f' := \begin{bmatrix} 0 \\ \psi \end{bmatrix},$$

are considered. Furthermore, the TSLRO reformulation (1.16) satisfies Assumptions 1.3.1, 1.3.2, and 1.3.3 when the WCARM also satisfies Assumptions 1.3.3 and 1.4.1, while the WCARM needs to additionally satisfy Assumption 1.3.4 for the TSLRO reformulation to satisfy Assumption 1.3.4.

Once again, Proposition 1.4.2 states that the WCARM model with objective uncertainty can be reformulated as a TSLRO problem and can therefore benefit from solution methods developed for adjustable robust optimization problems. In particular, a conservative approximation can be obtained using affine decision rules, which reduces to the linear program (1.7) when Assumptions 1.3.1, 1.3.2, 1.3.3, and 1.4.1 are satisfied by the WCARM. In order to implement the column-and-constraint generation algorithm described in Section 1.3.1.1, one needs to additionally verify that the WCARM satisfies Assumption 1.3.4.

1.5 TSLRO Reformulations for Worst-case Relative Regret Minimization Problems

An alternative form of regret minimization problem considers regret in its relative, rather than absolute, form. This approach is also equivalently measured according to a socalled "competitive ratio", which is a popular measure in the field of online optimization (Borodin and El-Yaniv, 2005). As defined in Kouvelis and Yu (1996), the worst-case relative regret criterion aims at evaluating the performance of a decision x with respect to the worst-case regret that might be experienced in hindsight relatively to the best decision that could have been made. Mathematically speaking, given a non-negative profit function $h(x, \zeta)$, which depends on both the decision and the realization of some uncertain vector of parameters ζ , one measures the relative regret experienced once ζ is revealed as the ratio of the difference between the best profit achievable $\sup_{x' \in \mathcal{X}} h(x', \zeta)$ and the profit $h(x, \zeta)$ achieved by the decision x that was implemented, over the best profit achievable. When Assumption 1.4.1 is satisfied, the worst-case relative regret minimization (WCRRM) problem thus takes the form:

(WCRRM) minimize
$$\sup_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \frac{\sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) - h(\boldsymbol{x},\boldsymbol{\zeta})}{\sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta})} \right\},$$
 (1.18)

where it is understood that the relative regret is null if $\sup_{x' \in \mathcal{X}} h(x', \zeta) = h(x, \zeta) = 0$. Mathematically speaking, we might be more accurate by defining the WCRRM problem as:

$$\underset{\boldsymbol{x} \in \mathcal{X}}{\text{minimize}} \sup_{\boldsymbol{\zeta} \in \mathcal{U}} \quad \lim_{\epsilon \to 0^+} \left\{ \frac{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})}{\epsilon + \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})} \right\} \,.$$

Besides Assumption 1.4.1, the following two assumptions will be useful in deriving TSLRO reformulations for WCRRM problems.

Assumption 1.5.1 The profit function $h(x, \zeta) \ge 0$ for all $x \in \mathcal{X}$ and all $\zeta \in \mathcal{U}$. This implies that the WCRRM problem satisfies Assumption 1.3.2 and, with Assumption 1.4.1, that the optimal value of Problem (1.18) lies in the closed interval [0, 1].

Assumption 1.5.2 It is possible to achieve a strictly positive worst-case profit, namely

$$\exists x \in \mathcal{X}, \forall \zeta \in \mathcal{U}, h(x, \zeta) > 0.$$

Together with Assumption 1.4.1, this implies that the optimal value of Problem (1.18) lies in the open interval [0, 1].

While Assumptions 1.4.1 and 1.5.1 simply formalize a hypothesis that needs to be made for the WCRRM problem to be meaningful, we argue that Assumption 1.5.2 is made without loss of generality since if it is not the case, then the WCRRM becomes trivial. Indeed, one can then simply consider any $x \in \mathcal{X}$ as an optimal solution to the WCRRM since it achieves the best possible worst-case relative regret, i.e. either 0% or 100%.

In what follows we demonstrate how the WCRRM problem can be reformulated as a TSLRO problem when the profit function $h(x, \zeta)$ captures the profit of a second-stage linear decision model with either right-hand side or objective uncertainty. Note that for completeness Section 1.8.14 presents similar TSLRO reformulation for the case where the two-stage problem is a cost minimization problem, i.e. that $h(x, \zeta)$ is non-positive.

1.5.1 The Case of Right-Hand Side Uncertainty

We consider the case where $h(x, \zeta)$ takes the form presented in Problem (1.3) and where uncertainty is limited to the right-hand side as defined in Definition 1.3.1.

Proposition 1.5.1 *Given that Assumptions 1.3.1, 1.4.1, and 1.5.1 are satisfied, the WCRRM problem with right-hand side uncertainty is equivalent to the following TSLRO problem:*

$$\begin{array}{ll} \text{maximize} & \inf_{\mathbf{x}' \in \mathcal{X}', \mathbf{y}'(\cdot)} & c'^T \mathbf{x}' \\ \end{array}$$
(1.19a)

$$\begin{array}{ll} \mathbf{x}' \in \mathcal{X}', \mathbf{y}'(\cdot) & \boldsymbol{\zeta}' \in \mathcal{U}' \\ \text{subject to} & A'\mathbf{x}' + B'\mathbf{y}'(\boldsymbol{\zeta}') \leq \Psi'(\mathbf{x}')\boldsymbol{\zeta}' + \boldsymbol{\psi}', \,\forall \, \boldsymbol{\zeta}' \in \mathcal{U}', \end{array}$$
(1.19b)

where $x' \in \mathbb{R}^{n_x+1}$, $\boldsymbol{\zeta}' \in \mathbb{R}^{n_{\boldsymbol{\zeta}}+n_x+n_y}$, $\boldsymbol{y}' : \mathbb{R}^{n_{\boldsymbol{\zeta}}+n_x+n_y} \to \mathbb{R}^{n_y}$, $\boldsymbol{c}' = \begin{bmatrix} -1 & \boldsymbol{0}^T \end{bmatrix}^T$, while $\mathcal{X}' := \{\begin{bmatrix} t & \boldsymbol{x}^T \end{bmatrix}^T \in \mathbb{R}^{n_x+1} \mid \boldsymbol{x} \in \mathcal{X}, t \in [0, 1] \}$, \mathcal{U}' is defined as in equation (1.14) and

$$A' := \begin{bmatrix} 0 & -\boldsymbol{c}^T \\ 0 & A \end{bmatrix}, \qquad \qquad B' := \begin{bmatrix} -\boldsymbol{d}^T \\ B \end{bmatrix},$$
$$\Psi'(\boldsymbol{x}') := \begin{bmatrix} \boldsymbol{0}^T & -\boldsymbol{c}^T & -\boldsymbol{d}^T \\ \Psi & 0 & 0 \end{bmatrix} + \begin{bmatrix} \boldsymbol{0}^T & \boldsymbol{c}^T & \boldsymbol{d}^T \\ 0 & 0 & 0 \end{bmatrix} x'_1, \qquad \psi' := \begin{bmatrix} \boldsymbol{0} \\ \psi \end{bmatrix}.$$

In particular, an optimal solution for the WCRRM takes the form of $x^* := x'^*_{2:n_x+1}$ and achieves a worst-case relative regret of x'^*_1 . Furthermore, this TSLRO reformulation necessarily satisfies Assumption 1.3.1 while it only satisfies Assumption 1.3.2 if all $x \in X$ achieve a worst-case regret of zero.

Proposition 1.5.1 motivates the application of solution methods developed for adjustable robust optimization problems to WCRRM problems. It is clear for instance that a conservative approximation that takes the form of the linear program (1.7) can readily be obtained by using affine decision rules. Exact methods however must be designed in a way that can handle TSLRO problems that do not satisfy relatively complete recourse. In particular, in our numerical experiments, we will make use of the method proposed in Ayoub and Poss (2016).

1.5.2 The Case of Objective Uncertainty

We consider the case where $h(x, \zeta)$ takes the form presented in Problem (1.10).

Proposition 1.5.2 Given that Assumptions 1.3.1, 1.4.1, 1.5.1, and 1.5.2 are satisfied, the WCRRM problem with objective uncertainty is equivalent to the following TSLRO problem:

$$\underset{x' \ y'(\cdot)}{\operatorname{maximize}} \quad \inf_{c' \in \mathcal{U}'} c'^T x'$$
 (1.20a)

subject to
$$A' \boldsymbol{x} + B' \boldsymbol{y}'(\boldsymbol{\zeta}') \leq \Psi'(\boldsymbol{x}') \boldsymbol{\zeta}' + \psi'$$
 (1.20b)

$$x' \in \mathcal{X}',$$
 (1.20c)

where $\boldsymbol{x}' \in \mathbb{R}^{n_x+1}$, $\boldsymbol{y}' : \mathbb{R}^{n_\zeta+m} \to \mathbb{R}^{m+r}$, while $\mathcal{X}' := \{ [u \ \boldsymbol{z}^T]^T \in \mathbb{R}^{n_x+1} \mid W \boldsymbol{z} \leq \boldsymbol{v} u, u \geq 1 \}$, \mathcal{U}' is defined as in equation (1.17). Furthermore, we have that $\mathbf{c}' := [-1 \ \mathbf{0}^T]^T$, while

$$\begin{split} A' &:= \begin{bmatrix} 0 & -c^T \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad B' &:= \begin{bmatrix} \psi^T & \psi^T \\ A^T & W^T \\ -A^T & -W^T \\ B^T & 0 \\ -B^T & 0 \\ -B^T & 0 \\ 0 & 0 \\ 0 & -I \end{bmatrix}, \\ \Psi'(x') &:= \begin{bmatrix} 0^T & \psi^T x_1' - x_{2:n_x+1}' A^T \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ -D & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad and \ \psi' &:= \begin{bmatrix} 0 \\ c \\ -c \\ d \\ -d \\ 0 \\ 0 \end{bmatrix}. \end{split}$$

In particular, an optimal solution for the WCRRM takes the form of $x^* := x_{2:n_x+1}^{\prime*}/x_1^{\prime*}$ and achieves a worst-case relative regret of $1 - 1/x_1^{\prime*}$. Finally, this TSLRO reformulation necessarily satisfies Assumption 1.3.1 while it only satisfies Assumption 1.3.2 if all $x \in \mathcal{X}$ achieve a worstcase regret of zero.

This final proposition reformulating WCRRM problems with objective uncertainty as TSLRO problems motivates once more the application of solution methods developed for adjustable robust optimization problems to this under-explored class of problems. In particular, a tractable conservative approximation can directly be obtained by using affine decision rules while to obtain an exact solution, a method such as proposed in Ayoub and Poss (2016) needs to be employed.

1.6 Optimality of Affine Decision Rules

In this section, we derive conditions under which one can establish that affine decision rules are optimal in the TSLRO reformulation of WCARM and WCRRM problems. These results will draw their arguments from similar results that have been established for two-stage robust optimization. In fact, perhaps the most famous of those results is attributed to Bertsimas and Goyal (2012) for the case where the uncertainty set takes the form of a simplex set.

Definition 1.6.1 An uncertainty set U is called a "simplex set" if it is the convex hull of $n_{\zeta} + 1$ affinely independent points in $\mathbb{R}^{n_{\zeta}}$.

One can in fact extend the known optimality of affine decisions to special classes of WCARM and WCRRM problems.

Proposition 1.6.1 If $h(x, \zeta)$ satisfies $\max_{x \in \mathcal{X}} h(x, \zeta) = \gamma^T \zeta + \overline{\gamma}$ for some $\gamma \in \mathbb{R}^{n_{\zeta}}$ and $\overline{\gamma} \in \mathbb{R}$ and \mathcal{U} is a simplex set, then affine decision rules are optimal in the TSLRO reformulation of the WCARM (under Assumption 1.3.1) and WCRRM (under Assumptions 1.3.1, 1.4.1, and 1.5.1) problems with right-hand side uncertainty, i.e. Problem (1.13) and (1.19) respectively.

Note that the condition that $\max_{\boldsymbol{x}\in\mathcal{X}} h(\boldsymbol{x},\boldsymbol{\zeta}) = \boldsymbol{\gamma}^T \boldsymbol{\zeta} + \bar{\gamma}$ is satisfied in a number of classical inventory models. For instance, this condition is satisfied for the following multiitem newsvendor problem (see Ardestani-Jaafari and Delage 2016):

$$\underset{\boldsymbol{x} \ge 0}{\text{maximize}} \quad \inf_{\boldsymbol{\zeta} \in \mathcal{U}} \sum_{i=1}^{n_y} (p_i - c_i) x_i + \min(-b_i(\zeta_i - x_i), (s_i - p_i)(x_i - \zeta_i)),$$

and where x_i is the number of units of item *i* ordered, ζ_i is the unknown demand for item *i*, p_i is the sales price for item *i*, c_i is the ordering cost, b_i is the shortage cost, and s_i is the salvage value. Exploiting a well-known epigraph formulation, the single-stage model can be reformulated using $\max_{\boldsymbol{x} \ge 0} \min_{\boldsymbol{\zeta} \in \mathcal{U}} h(\boldsymbol{x}, \boldsymbol{\zeta})$ with

$$h(\boldsymbol{x}, \boldsymbol{\zeta}) := \max_{\boldsymbol{y}} \qquad \sum_{i}^{n_{\boldsymbol{y}}} y_{i}$$

subject to
$$y_{i} \leq (p_{i} - c_{i})x_{i} + (s_{i} - p_{i})(x_{i} - \zeta_{i}), \, \forall i = 1, \dots, n_{\boldsymbol{y}}$$
$$y_{i} \leq (p_{i} - c_{i})x_{i} - b_{i}(\zeta_{i} - x_{i}), \, \forall i = 1, \dots, n_{\boldsymbol{y}}.$$

It is usually assumed that $s_i \leq c_i \leq p_i$, namely that the salvage value is smaller than the ordering cost, which is itself smaller than the retail price, so that if the demand vector $\boldsymbol{\zeta}$ was known then the optimal order would simply be $x_i^* = \zeta_i \mathbf{1}\{p_i - c_i + b_i \geq 0\}$, where $\mathbf{1}\{p_i - c_i + b_i \geq 0\}$ is the indicator function that returns 1 if $p_i - c_i + b_i \geq 0$ and 0 otherwise. Hence, we have that:

$$\max_{\boldsymbol{x} \ge 0} h(\boldsymbol{x}, \boldsymbol{\zeta}) = \sum_{i=1}^{n_y} \left(-b_i + (p_i - c_i + b_i) \mathbf{1} \{ p_i - c_i + b_i \ge 0 \} \right) \zeta_i.$$

Similarly, in a classical lot-sizing problem with backlog described as:

$$\underset{\boldsymbol{x}\geq 0}{\text{maximize}} \quad \inf_{\boldsymbol{\zeta}\in\mathcal{U}} \sum_{t=1}^{T} \left(-c_t x_t - \min\left(h_t \left(\sum_{t'=1}^{t} x_{t'} - \zeta_{t'}\right), b_t \left(\sum_{t'=1}^{t} \zeta_{t'} - x_{t'}\right)\right) \right) ,$$

where x_t is the number of units ordered for time t, ζ_t is the demand for time t, while c_t is the ordering cost, h_t the holding cost, and b_t the shortage cost. One can exploit the well-known facility location reformulation (see for instance Pochet and Wolsey 1988) to simplify the full information problem:

$$\begin{split} & \max_{\boldsymbol{x}\in\mathcal{X}} h(\boldsymbol{x},\boldsymbol{\zeta}) \\ &= \max_{X:X \ge 0, \sum_{t=1}^{T+1} X_{t,t'} = \zeta_t', \forall t'} - \left(\sum_{t=1}^T c_t \sum_{t'=1}^T X_{t,t'} + \sum_{i=1}^t \sum_{j=t+1}^T h_t X_{i,j} + \sum_{i=1}^t \sum_{j=t+1}^{T+1} b_t X_{j,i} \right) \\ &= - \left(\sum_{t=1}^T \min_{\boldsymbol{x}:\boldsymbol{x}\ge 0, \sum_{t'=1}^{T+1} x_{t'} = \zeta_t} \sum_{t'=1}^T c_{t'} x_{t'} + \sum_{t'=1}^{t-1} \sum_{t''=t'}^{t-1} h_{t''} x_{t'} + \sum_{t'=t+1}^{T+1} \sum_{t''=t}^{t'-1} b_{t''} x_{t'} \right) \\ &= - \sum_{t=1}^T \zeta_t \left(\min_{\boldsymbol{x}:\boldsymbol{x}\ge 0, \sum_{t'=1}^{T+1} x_{t'} = 1} \sum_{t'=1}^T c_{t'} x_{t'} + \sum_{t'=1}^{t-1} \sum_{t''=t'}^{t-1} h_{t''} x_{t'} + \sum_{t'=t+1}^{T+1} \sum_{t''=t}^{t'-1} b_{t''} x_{t'} \right), \end{split}$$

where $X_{t,t'}$ captures the number of units produced at time *t* to satisfy the demand at time *t'*. We see once again that the optimal value is linear with respect to ζ .

Proposition 1.6.1 also has an analog in the context of a two-stage model with objective uncertainty.

Proposition 1.6.2 If $\mathcal{Z} := \{(x, y) \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} | x \in \mathcal{X}, Ax + By \leq \psi\}$ is a simplex set, then affine decision rules are optimal in the TSLRO reformulation of the WCARM, when Assumptions 1.3.1, 1.3.2, 1.3.3, and 1.4.1 hold, and WCRRM problems, when Assumptions 1.3.1, 1.4.1, 1.5.1, and 1.5.2 hold, with objective uncertainty, i.e. Problem (1.16) and (1.20) respectively.

Interestingly, Proposition 1.6.2 provides a polynomial-time solvable reformulation for the WCARM and WCRRM versions of resource allocation problems.

Corollary 1.6.1 *The linear program obtained by employing affine decision rules on the TSLRO reformulation of the WCARM problem*

$$\min_{\boldsymbol{x} \in \mathcal{X}} \sup_{\boldsymbol{\zeta} \in \mathcal{U}} \left(\max_{\boldsymbol{x}' \in \mathcal{X}} \boldsymbol{d}(\boldsymbol{\zeta})^T \boldsymbol{x} - \boldsymbol{d}(\boldsymbol{\zeta})^T \boldsymbol{x} \right),$$

where $\mathcal{X} := \{ \boldsymbol{x} \in \mathbb{R}^{n_x}_+ | \boldsymbol{w}^T \boldsymbol{x} \leq v \}$ with $\boldsymbol{w} \in \mathbb{R}^{n_x}_+$ and $v \in \mathbb{R}_+$, is exact for all polyhedral uncertainty set \mathcal{U} , and similarly for the WCRRM version of this problem given that the assumptions described in Proposition 1.6.2 hold.

This corollary extends the result in Averbakh (2004), which identified a $O(n_x \log(n_x))$ time algorithm for the WCARM version of the continuous knapsack problem under interval uncertainty.

Following the work of Ardestani-Jaafari and Delage (2016), the result presented in Proposition 1.6.1 can be extended to other forms of uncertainty sets in the case that $h(x, \zeta)$ captures the sum of piecewise linear concave functions.

Proposition 1.6.3 If $h(x, \zeta)$ is a sum of piecewise linear concave functions of the form:

$$h(\boldsymbol{x},\boldsymbol{\zeta}) := \sum_{i=1}^{N} \min_{k=1,\dots,K} \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^{T} \boldsymbol{\zeta} + \beta_{ik}(\boldsymbol{x}) = \max_{\boldsymbol{y}} \qquad \sum_{i=1}^{n_{y}} y_{i}$$
(1.21)
s.t. $y_{i} \leq \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^{T} \boldsymbol{\zeta} + \beta_{ik}(\boldsymbol{x}), \,\forall \, i, \,\forall \, k,$

for some affine mappings $\alpha_{ik} : \mathbb{R}^{n_x} \to \mathbb{R}^{n_\zeta}$ and $\beta_{ik} : \mathbb{R}^{n_x} \to \mathbb{R}$, the uncertainty set \mathcal{U} is the budgeted uncertainty set²

$$\mathcal{U} := \{ \boldsymbol{\zeta} \in \mathbb{R}^{n_{\zeta}} \mid \exists \boldsymbol{\zeta}^{+} \in \mathbb{R}^{n_{\zeta}}_{+}, \exists \boldsymbol{\zeta}^{-} \in \mathbb{R}^{n_{\zeta}}_{+}, \boldsymbol{\zeta} = \boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}, \boldsymbol{\zeta}^{+} + \boldsymbol{\zeta}^{-} \leq 1, \sum_{i} \zeta_{i}^{+} + \zeta_{i}^{-} = \Gamma \},$$
(1.22)

with $\Gamma \in [0, n_{\zeta}]$, and the following conditions are satisfied:

- 1. Either of the following applies:
 - *i*. $\Gamma = 1$
 - *ii.* $\Gamma = n_{\zeta}$ and uncertainty is "additive": *i.e.* $\boldsymbol{\alpha}_{ik}(\boldsymbol{x}) = \bar{\alpha}_{ik}(\boldsymbol{x})(\sum_{\ell < i} \hat{\alpha}_{\ell}(\boldsymbol{x})\boldsymbol{e}_{\ell})$ for some $\bar{\alpha}_{ik} : \mathbb{R}^{n_x} \to \mathbb{R}$ for all *i* and *k* and some $\hat{\boldsymbol{\alpha}} : \mathbb{R}^{n_x} \to \mathbb{R}^{n_{\zeta}}$
 - *iii.* Γ *is integer and the objective function is "decomposable": i.e.* $\alpha_{ik}(x) = \bar{\alpha}_{ik}(x)e_i$ for some $\bar{\alpha}_{ik} : \mathbb{R}^{n_x} \to \mathbb{R}$ for all i and k
- 2. $\max_{\boldsymbol{x}\in\mathcal{X}} h(\boldsymbol{x},\boldsymbol{\zeta}) = \boldsymbol{\gamma}^T \boldsymbol{\zeta} + \bar{\gamma} \text{ for some } \boldsymbol{\gamma} \in \mathbb{R}^{n_{\zeta}} \text{ and } \bar{\gamma} \in \mathbb{R}$

Then, affine decision rules with respect to $(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-, \boldsymbol{x}', \boldsymbol{y}')$ are optimal in the TSLRO reformulation of the WCARM and WCRRM problems, i.e. Problem (1.13) and (1.19) respectively.

Propositions 1.6.1 and 1.6.3 effectively extend the set of problem classes for which a polynomial-time solution scheme is known. In particular, it extends the results of Vairaktarakis (2000) for multi-item newsvendor problems to include simplex sets and budgeted uncertainty sets with integer budget. They similarly provide a polynomial-time solution scheme for a large class of lot-sizing problems under the budgeted uncertainty set as long as $\Gamma = 1$ or n_{ζ} . Unlike in the work of Vairaktarakis (2000) and Zhang (2011), tractability does not come from exploiting specifically designed algorithms for each of these applications but is rather simply achieved by employing the general linear decision rules approach on the TSLRO reformulation. It further naturally serves as theoretical evidence of the effectiveness of such an approach for general regret minimization. Finally, it is worth noting that neither the proof of Proposition 1.6.1 nor 1.6.3 exploit the fact that the affine decision rules employed in the TSLRO reformulation were flexible with respect to (x', y'). This implies that the two propositions also hold when the simpler decision rules of the form $y(\zeta) := y + Y_{\zeta}\zeta$ are used, as was proposed in Bertsimas and Dunning (2020).

1.7 Numerical Results

In this section, we evaluate the numerical performance of exact and approximate solution schemes that are commonly used to solve two-stage linear robust optimization problems when employed to solve the TSLRO reformulations of worst-case regret minimization problems. This is done in the context of two representative applications of TSLRO, namely a multi-item newsvendor problem and a production-transportation problem, which are respectively special cases of TSLRO with right-hand side uncertainty and objective uncertainty. Our objective consists in comparing both the solution time and quality of first-stage decisions that are obtained using exact and approximate methods and provide empirical evidence regarding whether two-stage regret minimization problems are more difficult to solve than their robust optimization version.

While a number of approximation schemes from the adjustable robust optimization literature could be put to the test, we focus our analysis on the AARC approximation method described in Section 1.3.1.2. Similarly, we rely on the C&CG method presented in Section 1.3.1.1 to solve the TSLRO reformulations of WCARM problems exactly, and on the column-and-constraint generation algorithm of Ayoub and Poss (2016), called C&CG*, for WCRRM problems. A time limit of 4 hours (14, 400 seconds) and optimality tolerance of 10^{-6} are imposed on all solution schemes. The quality of the AARC approximation scheme is reported in terms of relative optimality gap (in %) in the case of a WCARM model, and absolute optimality gap for WCRRM models since the objective function is already expressed in percentage. All algorithms were implemented in MATLAB R2017b using the YALMIP toolbox and CPLEX 12.8.0 as the solver for all linear programming models.

1.7.1 Multi-Item Newsvendor Problem

The first application that we consider is the multi-item newsvendor problem, which was studied in its robust optimization form in Ardestani-Jaafari and Delage (2016) and

Ardestani-Jaafari and Delage (2020). The single-stage robust formulation of this problem is as follows:

$$\underset{\boldsymbol{x} \ge 0}{\text{maximize}} \quad \min_{\boldsymbol{\zeta} \in \mathcal{U}} \sum_{i=1}^{n_y} p_i \min(x_i, \zeta_i) - c_i x_i + s_i \max(x_i - \zeta_i, 0) - b_i \max(\zeta_i - x_i, 0), \quad (1.23)$$

where $p_i \ge 0$, $c_i \in [0, p_i]$, $s_i \in [0, c_i]$, and $b_i \ge 0$ represent sales price, ordering cost, salvage value, and shortage cost of a unit of item i, with $i = 1, ..., n_y$, respectively. Decision variable x_i is the initial ordering amount of item i. We refer the reader to Section 1.6 for a reformulation of the robust multi-item newsvendor problem as a TSLRO problem using y as epigraph variables.

We consider two forms of uncertainty sets, which respectively model the fact that the demand for each item is assumed to be correlated or not. The "uncorrelated demand" uncertainty set is defined straightforwardly in terms of the well-known budgeted set (see Bertsimas and Sim 2004):

$$\mathcal{U}(\Gamma) = \left\{ \boldsymbol{\zeta} \middle| \begin{array}{l} \boldsymbol{\delta}^{+} \geq 0, \, \boldsymbol{\delta}^{-} \geq 0 \\ \boldsymbol{\delta}^{+}_{i} + \boldsymbol{\delta}^{-}_{i} \leq 1, \, \forall i = 1, \dots, n_{y} \\ \sum_{i=1}^{n_{y}} \boldsymbol{\delta}^{+}_{i} + \boldsymbol{\delta}^{-}_{i} = \Gamma \\ \boldsymbol{\zeta}_{i} = \bar{\boldsymbol{\zeta}}_{i} + \hat{\boldsymbol{\zeta}}_{i}(\boldsymbol{\delta}^{+}_{i} - \boldsymbol{\delta}^{-}_{i}), \, \forall i = 1, \dots, n_{y} \end{array} \right\}$$

where $\bar{\zeta}_i$ and $\hat{\zeta}_i$ denote the nominal demand and the maximum demand deviation of the item *i* and where $\Gamma \in [0, n_y]$ captures a budget of maximum number of deviations from the nominal demand. We also consider a "correlated demand" uncertainty set defined as follows:

$$\tilde{\mathcal{U}}(\Gamma) = \left\{ \boldsymbol{\zeta} \middle| \exists \boldsymbol{\delta}^{+}, \boldsymbol{\delta}^{-} \in \mathbb{R}^{n_{y}}, \begin{array}{l} \boldsymbol{\delta}^{+}_{i} \geq 0, \ \boldsymbol{\delta}^{-} \geq 0 \\ \delta_{i}^{+} + \delta_{i}^{-} \leq 1, \ \forall i = 1, \dots, n_{y} \\ \sum_{i=1}^{n_{y}} \delta_{i}^{+} + \delta_{i}^{-} = \Gamma \\ \zeta_{i} = \bar{\zeta}_{i} + \hat{\zeta}_{i} (\delta_{j_{1}(i)}^{+} + \delta_{j_{2}(i)}^{+} - \delta_{j_{1}(i)}^{-} - \delta_{j_{2}(i)}^{-})/2, \ \forall i = 1, \dots, n_{y} \end{array} \right\}$$

,

where $j : \{1, ..., n_y\} \rightarrow \{1, ..., n_y\}^2$ identifies two sources of perturbation of item *i* such that items i_1 and i_2 are correlated if $j_{\ell_1}(i_1) = j_{\ell_2}(i_2)$ for some $(\ell_1, \ell_2) \in \{1, 2\}^2$. We note that for both sets, we employ a less common (but equivalent) equality representation of the budget constraint in order to be consistent with the representation used in Proposition

Problem Size	Decision	Type of		Level of Un (in % o	certainty f n_y)	ý
	Cincilon	periornance	30%	50%	70%	100%
	Worst-case Profit	Avg Rel. Gap - AARC	0.72%	0.62%	0.92%	0.00%
	(RO)	Avg CPU time (s) - AARC	1.3	1.3	1.3	1.3
	()	Avg CPU time (s) - C&CG	71.8	119.9	148.7	85.8
5 items	Worst-case	Avg Rel. Gap - AARC	2.03%	0.49%	0.14%	0.00%
5 1101115	Absolute Regret	Avg CPU time (s) - AARC	1.3	1.3	1.3	1.3
	(WCARM)	Avg CPU time (s) - C&CG	116.7	143.1	105.8	82.5
	Worst-case	Avg Abs. Gap - AARC	0.24%	0.15%	0.08%	0.00%
	Relative Regret	Avg CPU time (s) - AARC	0.2	0.2	0.3	0.3
	(WCRRM)	Avg CPU time (s) - C&CG*	142.7	154.1	166.8	118.4
	Warst case Profit	Avg Rel. Gap - AARC	0.00%	0.00%	0.00%	0.00%
	(PO)	Avg CPU time (s) - AARC	1.4	1.4	1.5	1.5
	(KO)	Avg CPU time (s) - C&CG	96.9	138.5	282.6	174.8
10 itoma	Worst-case	Avg Rel. Gap - AARC	0.00%	0.00%	0.00%	0.00%
10 nems	Absolute Regret	Avg CPU time (s) - AARC	1.5	1.5	1.5	1.5
	(WCARM)	Avg CPU time (s) - C&CG	184.0	239.4	201.8	153.1
	Worst-case	Avg Abs. Gap	0.00%	0.00%	0.00%	0.00%
	Relative Regret	Avg CPU time (s) - AARC	0.4	0.4	0.4	0.4
	(WCRRM)	Avg CPU time (s) - C&CG*	238.5	315.0	312.6	206.2
	Worst-case Profit	Avg Rel. Gap - AARC	0.00%	0.00%	0.00%	0.00%
	(RO)	Avg CPU time (s) - AARC	1.9	1.9	2.0	2.1
	(RO)	Avg CPU time (s) - C&CG	227.3	381.9	649.8	460.3
20 itoms	Worst-case	Avg Rel. Gap - AARC	0.00%	0.00%	0.00%	0.00%
20 items	Absolute Regret	Avg CPU time (s) - AARC	2.0	2.1	2.2	2.2
	(WCARM)	Avg CPU time (s) - C&CG	494.7	760.6	781.3	367.7
	Worst-case	Avg Abs. Gap - AARC	0.00%	0.00%	0.00%	0.00%
	Relative Regret	Avg CPU time (s) - AARC	1.0	1.0	1.1	1.3
	(WCRRM)	Avg CPU time (s) - C&CG*	891.1	7,528.4[6]	-[0]	5,115.3

	Table 1.1: Multi-Item	Newsvendor Problem -	 Uncorrelated Uncertaint 	v Set
--	-----------------------	----------------------	---	-------

[] indicates the number of instances solved by C&CG* algorithm within the 4 hours time limit. In this case, the average is computed on the instances that were solved to optimality within the time limit.

1.6.3. This proposition also suggests that affine decision rules should be employed on the lifted space $(\zeta^+, \zeta^-, x', y')$.

We consider three different sizes of the problem, namely $n_y \in \{5, 10, 20\}$. For each size, we generate 10 problem instances randomly according to the following procedure. Each sales price p_i is uniformly and independently generated on the interval [0.5, 1], each ordering cost c_i uniformly generated on $[0.3p_i, 0.9p_i]$, and the salvage value s_i and shortage cost b_i are drawn uniformly at random from $[0.1c_i, c_i]$. The nominal demand for each item i is $\bar{d}_i = 10$ while the maximum demand perturbation is generated uniformly on $[0.3\bar{d}_i, 0.6\bar{d}_i]$. In the case of the correlated uncertainty set $\tilde{\mathcal{U}}$, for each item i the pair $(j_1(i), j_2(i))$ is drawn randomly among all possible pairs such that $j_1(i) \neq j_2(i)$. The budget Γ is fixed among the levels $\Gamma \in \{0.3n_y, 0.5n_y, 0.7n_y, n_y\}$. In what follows, we first study the numerical efficiency and quality of solutions obtained from AARC and C&CG in the worst-case profit (RO), worst-case absolute regret (WCARM), and worst-case relative regret (WCRRM) problems. We then present a short study that focuses on the need for flexibility with respect to hindsight decisions. Finally, we investigate, from a decision analysis point of view, whether there is a real need for formulating WCARM and WCRRM problems given that RO solutions are supposed to be robust and might already be solutions that achieve low absolute and relative regret.

1.7.1.1 Numerical Efficiency of AARC Compared with C&CG

Tables 1.1 and 1.2 present the average performance of C&CG and AARC in solving the classical robust optimization, the worst-case absolute regret minimization, and the worst-case relative regret minimization formulation when accounting for the uncorrelated and correlated uncertainty sets respectively.

Looking at Table 1.1, one can remark that for the instances with $n_y = 5$ items, the average optimality gaps achieved by the AARC approach are of similar small sizes in the case of classical robust optimization as for worst-case regret minimization. The optimality gap is also surprisingly small (below 0.3%) for the WCRRM problems. Since the instances studied in this table employ an uncorrelated uncertainty set, the empirical evidence confirms the findings of Proposition 1.6.3, which states that, similarly as for the robust optimization formulation (see Ardestani-Jaafari and Delage 2016), AARC provides exact solutions for WCARM and WCRRM when Γ is an integer.

When it comes to comparing computation times, one may make three interesting observations. First, all AARC approximation models are solved in less than 3 seconds (on average), which is more than one order of magnitude faster than the time needed to solve any of these problems using C&CG. This can be explained by the well-known fact that each step of C&CG involves solving an NP-hard mixed integer linear program. Secondly, it appears to be generally true that both of the AARC and C&CG solution schemes have a similar runtime whether they are used to solve the RO model or the WCARM. This seems to support the claim that regret minimization has the same complexity as robust optimization for a two-stage linear program with right-hand side uncertainty. On the

Problem Size	Decision	Type of	L	evel of U (in %	Incertaint of n_y)	у
	Cincilon	performance	30%	50%	70%	100%
	Worst-case Profit	Avg Rel. Gap - AARC	1.46%	3.11%	2.39%	0.00%
	(RO)	Avg CPU time (s) - AARC	1.3	1.3	1.3	1.3
	(110)	Avg CPU time (s) - C&CG	81.6	86.6	93.2	73.7
5 items	Worst-case	Avg Rel. Gap - AARC	3.58%	3.68%	1.61%	0.00%
5 1101115	Absolute Regret	Avg CPU time (s) - AARC	1.3	1.3	1.3	1.3
	(WCARM)	Avg CPU time (s) - C&CG	78.9	83.0	96.8	75.1
	Worst-case	Avg Abs. Gap - AARC	0.68%	0.84%	0.69%	0.00%
	Relative Regret	Avg CPU time (s) - AARC	0.2	0.2	0.2	0.2
	(WCRRM)	Avg CPU time (s) - C&CG*	103.3	118.9	132.1	93.3
	Warst case Profit	Avg Rel. Gap - AARC	1.30%	1.62%	0.62%	0.00%
	(PO)	Avg CPU time (s) - AARC	1.4	1.4	1.4	1.4
	(KO)	Avg CPU time (s) - C&CG	115.9	145.9	178.4	126.4
10 itoms	Worst-case	Avg Rel. Gap - AARC	3.16%	0.87%	0.16%	0.00%
10 nems	Absolute Regret	Avg CPU time (s) - AARC	1.4	1.4	1.4	1.4
	(WCARM)	Avg CPU time (s) - C&CG	135.9	177.5	164.6	120.7
	Worst-case	Avg Abs. Gap - AARC	0.30%	0.13%	0.09%	0.00%
	Relative Regret	Avg CPU time (s) - AARC	0.3	0.3	0.3	0.4
	(WCRRM)	Avg CPU time (s) - C&CG*	208.3	262.0	258.9	157.8
	Worst-case Profit	Avg Rel. Gap - AARC	0.62%	0.52%	0.10%	0.00%
	(RO)	Avg CPU time (s) - AARC	1.7	1.7	1.8	1.9
	(1(0))	Avg CPU time (s) - C&CG	286.3	451.7	582.1	314.3
20 itoms	Worst-case	Avg Rel. Gap - AARC	0.66%	0.05%	0.01%	0.00%
20 1101113	Absolute Regret	Avg CPU time (s) - AARC	1.9	2.0	2.0	2.1
	(WCARM)	Avg CPU time (s) - C&CG	428.3	576.6	500.8	248.5
	Worst-case	Avg Abs. Gap - AARC	0.07%	0.05%	0.02%	0.00%
	Relative Regret	Avg CPU time (s) - AARC	0.8	0.9	0.9	1.0
	(WCRRM)	Avg CPU time (s) - C&CG*	717.1	2,681.5	6,287.9	567.8

Table 1.2: Multi-Item Newsvendor Problem - Correlated Uncertainty Set

other hand, it also appears that the C&CG^{*} approach used for WCRRM leads to longer run times than what is needed for RO models. Finally, we see that in the case of $n_y = 20$ the C&CG^{*} scheme is unable to solve a number of problem instances within the allocated time for $\Gamma = 10$ and 14. This is in sharp contrast with the AARC approach, which identifies optimal solutions in less than a couple of seconds. This evidence reinforces the idea that modern approximation methods that exist for RO models can provide highperformance algorithms for regret minimization problems.

Looking at Table 1.2 where problem instances have correlated demand, we draw similar conclusions as with Table 1.1. Namely, we observe that AARC provides an optimal solution when $\Gamma = n_y$, which might indicate that there are other conditions than those identified in Section 1.6 where affine decision rules are optimal. For other cases, the quality of the approximation is very high for all versions of the problems, presenting a maximum average gap of 3.68% and 0.84% for the WCARM and WCRRM problems respectively. In terms of the run times, the observations are also similar except for the instances where $n_y = 20$, which appear to be less challenging for the C&CG* scheme than when demand was uncorrelated. Indeed, C&CG* is able here to converge to an optimal solution within the time limit for all instances although this could simply be due to the specific structure of the 10 instances that were drawn for this part of the study. Overall, this study seems to indicate that AARC is a much more favorable approach for tackling larger scale regret minimization problems.

Remark 1.7.1 The average relative and absolute gaps presented in Tables 1.1 and 1.2 reflect the worst-case performance of a Pareto robustly optimal solutions of the AARC models, as prescribed in Iancu and Trichakis (2014). Specifically, once each AARC model is solved, we search among the robustly optimal affine decision rules for one that achieves the best objective value under a representation of the nominal scenario that lies in the relative interior of the uncertainty set.

1.7.1.2 Value of Flexibility to Hindsight Decisions

As discussed in Section 1.4.1, Bertsimas and Dunning (2020) provide a conservative approximation for the multi-stage regret minimization problems with right-hand side uncertainty, where decision rules only adapt to the realization of uncertain parameters. In contrast, our approach seeks decision rules that adapt both to the parameters and the optimal hindsight decisions, so-called x' and y'. Our initial experiments in fact indicated empirically that there was actually no value in employing the more flexible decision rules in the instances that were studied in Table 1.1 and 1.2. We suspect that this property is a consequence of the optimal hindsight profit being a linear function with respect to demand.

The difference between the two approaches already starts becoming observable when additional constraints are imposed on the size of the orders. In particular, consider the following version of multi-item newsvendor problem with order limits:

$$\max_{\boldsymbol{x} \ge 0, \{x_i \le u_i\}_{i=1}^{n_y}} \min_{\boldsymbol{\zeta} \in \mathcal{U}} \sum_{i=1}^{n_y} p_i \min(x_i, \zeta_i) - c_i x_i + s_i \max(x_i - \zeta_i, 0) - b_i \max(\zeta_i - x_i, 0), (1.24)$$

where u_i represents the maximum amount of order that can be placed for item *i*. Specifically, in this new experiment we simply modify the instances that led to Table 1.2, i.e. with correlated uncertainty, by imposing for each item *i* an order limit u_i equal to the nominal demand plus 50% of its maximum perturbation. The results of this experiment are presented in Table 1.3, where we compare our conservative approximation approach with the one proposed by Bertsimas and Dunning (2020), denoted as "P&D" and "B&D", respectively.

Problem	Decision	Type of Performance	Ι	Level of U (in %	Incertaint of n_y)	у
OILC	Cincilon	renomance	30%	50%	70%	100%
	Worst-case	Avg Gap (%) - AARC P&D	3.34%	4.67%	4.35%	2.32%
	Absolute Regret	Avg Gap (%) - AARC B&D	19.90%	40.46%	49.73%	61.42%
5 items	(WCARM)	Avg Bound Improvement (%) (B&D - P&D)	16.56%	35.79%	45.38%	59.10%
	Worst-case	Avg Absolute Gap (%) - AARC P&D	0.72%	1.56%	1.46%	0.74%
	Relative Regret	Avg Absolute Gap (%) - AARC B&D	2.80%	5.96%	6.86%	8.07%
	(WCRRM)	Avg Bound Improvement (%) (B&D - P&D)	2.07%	4.40%	5.40%	7.32%
	Worst-case	Avg Gap (%) - AARC P&D	0.68%	1.81%	2.63%	2.69%
	Absolute Regret	Avg Gap (%) - AARC B&D	33.66%	42.54%	52.00%	62.36%
10 items	(WCARM)	Avg Bound Improvement (%) (B&D - P&D)	32.98%	40.73%	49.37%	59.67%
	Worst-case	Avg Absolute Gap (%) - AARC P&D	0.29%	0.71%	0.79%	1.06%
	Relative Regret	Avg Absolute Gap (%) - AARC B&D	3.00%	4.27%	6.28%	8.15%
	(WCRRM)	Avg Bound Improvement (%) (B&D - P&D)	2.71%	3.55%	5.49%	7.08%
	Worst-case	Avg Gap (%) - AARC P&D	0.27%	1.01%	2.20%	3.19%
	Absolute Regret	Avg Gap (%) - AARC B&D	25.34%	35.23%	46.37%	55.49%
20 items	(WCARM)	Avg Bound Improvement (%) (B&D - P&D)	25.07%	34.22%	44.17%	52.30%
	Worst-case	Avg Absolute Gap (%) - AARC P&D	_	_	_	1.30%
	Relative Regret	Avg Absolute Gap (%) - AARC B&D	_	_	_	7.96%
	(WCRRM)	Avg Bound Improvement (%) (B&D - P&D)	2.48%	3.76%	5.37%	6.67%

Table 1.3: Compare P&D and B&D Approaches - Multi-Item Newsvendor Problem with Order Limits

- indicates that none of the instances were solved by C&CG* algorithm within the 4 hours time limit.

According to results presented in Table 1.3, the average optimality gap of P&D approach for the WCARM problem is less than 5% for all values of Γ and all problem sizes. Comparatively, the average gap of B&D reaches up to nearly 62%. On a case-by-case basis, we see that the average gap increases by a factor going from 6 to 90 times larger for B&D depending on the level of uncertainty and problem size. The value of hindsight flexibility also appears to increase as uncertainty is increased for the WCARM problem.

In terms of the WCRRM problem, relatively similar observations can be made. Specifically, the flexibility in P&D allows to improve the bound obtained from B&D by a factor ranging from 3 to 10 depending on the problem class that could be solved in less than 4 hours.

Overall, these results confirm a strong potential for improving the quality of the solution proposed in B&D by making the decision rules flexible with respect to optimal hindsight decisions.

1.7.1.3 Decision Analysis

We now turn to study whether the three criteria for decision making, namely worstcase profit, worst-case absolute regret, and worst-case relative regret, produce solutions that are quite different from each other. In particular, given that RO and WCARM are slightly more appealing from a computational point of view, one could ask whether there is value in solving the harder WCRRM problem. To provide some insight on this question, we evaluated the performance of each proposed solution scheme with respect to the two other criteria on the set of problem instances used for Table 1.2, with correlated demand. In details, given a two-stage problem instance, for each model $M \in$ $\mathcal{M} := \{\text{RO}, \text{WCARM}, \text{WCRRM}\}$, we compute the sub-optimality with respect to $M' \in$ $\mathcal{M} \setminus \{M\}$ of the best candidate of optimal solution set \mathcal{X}_M^* . This provides us for each model type an optimistic estimate of the sub-optimality we should expect when measuring performance with either of the two other criteria. Table 1.4 presents the average performances based on 120 problem instances, i.e. 10 instances of two-stage problems for each of 3 problem sizes and 4 uncertainty levels.³

Table 1.4: Average Suboptimality of Solutions from RO, WCARM, and WCRRM with Respect to RO, WCARM, and WCRRM Models Based on 120 Randomly Generated Instances of Three Different Sizes.

	Rel. Gap in RO	Rel. Gap in WCARM	Abs. Gap in WCRRM
$\mathcal{X}^*_{\mathrm{RO}}$	0 %	169.9%	25.2%
$\mathcal{X}^*_{\mathrm{WCARM}}$	36.3%	0 %	13.3%†
$\mathcal{X}^*_{\mathrm{WCRRM}}$	19.5%	59.2%	0 %

[†] Average is reported based on 118 instances given that two led to an infinite worst-case relative regret.

Looking at Table 1.4, we do find strong evidence of dissimilarities between the solution concepts. First, one notices that relying on the RO decisions leads to a significant average increase of 169.9% of the worst-case absolute regret performance and a 25.2% average increase in worst-case relative regret comparing to the optimal solution of these respective models. On the other hand, WCARM decisions will typically decrease the worst-case profit by 36.3%, while WCRRM decisions diminish it by a lesser 19.5%. This corroborates the conclusion from the example in Section 1.8.1 that WCRRM might be closer in spirit to RO than WCARM, especially given that WCARM actually led on two occasions to solutions that achieved infinite worst-case relative regret.

Overall, it is clear that RO models propose decisions that may be in contradiction with what leads to low absolute regret, whether it be absolute or relative. It is well known that RO decisions tend to improve worst-case profits while disregarding completely all plausible opportunities to make higher profits, which can lead to large regret in hindsight. On the other hand, WCARM and WCRRM decisions will follow a "less conservative" approach in the sense that they attempt to be well positioned to seize opportunities sacrificing to some extent the assurance of the higher possible worst-case profit.

1.7.2 Production-Transportation Problem

Our second application consists of the production-transportation problem with uncertainty in transportation cost, which was considered in Bertsimas et al. (2010a). Specifically, in this problem, one considers m facilities and n customer locations. Each facility has a production capacity of \bar{x}_i goods. The units produced at these facilities should be shipped to the customer locations in order to cover a predefined set of orders. The difficulty for the manager resides in the fact that transportation costs are unknown when production decisions are made. The corresponding TSLRO problem can be defined as follows:

$$\underset{0 \le \boldsymbol{x} \le \bar{\boldsymbol{x}}, \, y(\boldsymbol{\zeta})}{\text{minimize}} \max_{\boldsymbol{\zeta} \in \mathcal{U}} \qquad \sum_{i=1}^{m} c_i x_i + \sum_{i=1}^{m} \sum_{j=1}^{n} \zeta_{ij} y_{ij}(\boldsymbol{\zeta})$$
(1.25a)

subject to
$$\sum_{i=1}^{m} y_{ij}(\boldsymbol{\zeta}) = d_j, \, \forall j \in \mathcal{J}, \, \forall \boldsymbol{\zeta} \in \mathcal{U}$$
 (1.25b)

$$\sum_{i=1}^{n} y_{ij}(\boldsymbol{\zeta}) = x_i, \, \forall i \in \mathcal{I}, \, \forall \boldsymbol{\zeta} \in \mathcal{U}$$
(1.25c)

$$y(\boldsymbol{\zeta}) \ge 0, \, \forall \boldsymbol{\zeta} \in \mathcal{U} \,,$$
 (1.25d)

where for each facility location *i*, c_i is the production cost, while for each customer location *j*, d_j refers to the demand that needs to be covered, and ζ_{ij} is the initially unknown

transportation cost per unit from production facility i to customer location j. This problem has two-stages of decisions, namely the here-and-now production decisions x, and the wait-and-see transportation decisions y, which are made once transportation costs are observed. Finally, we define the uncertainty set as

$$\mathcal{U}(\Gamma) = \left\{ \boldsymbol{\zeta} \middle| \begin{array}{l} \boldsymbol{\delta}^{+} \geq 0, \, \boldsymbol{\delta}^{-} \geq 0 \\ \exists \boldsymbol{\delta}^{+}, \, \boldsymbol{\delta}^{-} \in \mathbb{R}^{m}, \quad \delta_{i}^{+} + \delta_{i}^{-} \leq 1, \, \forall i \in \mathcal{I} \,, \, \zeta_{ij} = \bar{\zeta}_{ij} + \hat{\zeta}_{ij}(\delta_{i}^{+} - \delta_{i}^{-}), \, \forall i \in \mathcal{I}, \, \forall j \in \mathcal{J} \\ \sum_{i} \delta_{i}^{+} + \delta_{i}^{-} = \Gamma \end{array} \right\}$$

where $\bar{\zeta}_{ij}$ and $\hat{\zeta}_{ij}$ are respectively the nominal cost and maximum cost deviations for transporting each unit of good transported from *i* to *j*. Note that in defining $\mathcal{U}(\Gamma)$, we make the uncertainty about the costs for transportation from the same facility perfectly correlated, which allows us to consider $\Gamma \in [0, m]$. Alternatively, one could easily consider each transportation cost to be independent from each other.

In our numerical experiments, we consider three different sizes of the problem, namely $(m,n) \in \{(3,6), (5,10), (7,14)\}$. In each case, we generate 10 instances randomly. To do so, we start by randomly generating m + n locations within the unit square. The nominal transportation cost per unit from facility i to customer j is set to the Euclidean distance between their locations and the maximum perturbation of this cost is supposed to be 50% of the nominal value. The production costs are uniformly and independently generated on the interval $\left[0.5 \frac{\sum_{ij} \zeta_{ij}}{mn}, 1.5 \frac{\sum_{ij} \zeta_{ij}}{mn}\right]$. We fix the production capacities \bar{x}_i to one. Given that this leads to a maximum total production of m units, the size of each order d_i is uniformly generated on the interval [0.5m/n, m/n]. The empirical performance of all solution schemes on all three forms of problems with $\Gamma \in \{0.3m, 0.5m, 0.7m, m\}$ are presented in Table 1.5. Note that in the case of the RO model, as described in Section 1.3.2, one can easily identify an optimal solution by solving the so-called robust counterpart (RC) model, which takes the form of a linear program.

Looking at Table 1.5, one remarks that the average of the optimality gaps achieved by the AARC approach for the WCARM model is always below 8% for all values of Γ and all problem sizes. This is a poorer performance than in the case of the multi-item newsvendor problem yet still makes the AARC approach attractive when comparing to the convergence time of C&CG for problems of size m = 7 and n = 14 where all AARC

Problem Size	Decision Criterion	Type of performance		Level of U (in %	Incertainty of <i>m</i>)	,
	Cincilon	performance	30%	50%	70%	100%
	Worst-case Cost (RO)	Avg CPU time (s) - RC	0.7	0.8	0.9	0.9
2 facilition	Worst-case	Avg Rel. Gap - AARC	0.55%	2.28%	2.51%	4.07%
5 facilities	Absolute Regret	Avg ČPU time (s) - AARC	1.9	1.9	2.0	2.0
6 customers	(WCARM)	Avg CPU time (s) - C&CG	8.5	10.4	10.6	11.0
	Worst-case	Avg Abs. Gap - AARC	0.02%	0.19%	0.23%	0.55%
	Relative Regret	Avg CPU time (s) - AARC	5.9	6.0	6.0	6.0
	(WCRRM)	Avg CPU time (s) - C&CG*	89.4	102.0	98.9	112.9
	Worst-case Cost (RO)	Avg CPU time (s) - RC	1.0	1.2	1.5	1.8
	Worst-case	Avg Rel. Gap - AARC	6.71%	7.21%	5.68%	4.97%
5 facilities	Absolute Regret	Avg ČPU time (s) - AARC	19.7	19.6	19.4	20.2
10 customers	(WCARM)	Avg CPU time (s) - C&CG	42.9	65.4	93.2	95.8
	Worst-case	Avg Abs. Gap - AARC	0.39%	0.66%	0.78%	0.79%
	Relative Regret	Avg CPU time (s) - AARC	23.3	24.6	23.9	23.9
	(WCRRM)	Avg CPU time (s) - C&CG*	299.7	555.4	1,000.3	1,564.0
	Worst-case Cost (RO)	Avg CPU time (s) - RC	3.0	3.9	5.0	6.0
7 fa cilition	Worst-case	Avg Rel. Gap - AARC	4.14%	4.54%	4.59%	4.21%
7 facilities	Absolute Regret	Avg ČPU time (s) - AARC	442.9	373.2	318.7	296.4
14 customers	(WCARM)	Avg CPU time (s) - C&CG	3,425.4	8,365.1	6,967.8	7,468.7
	Worst-case	Avg Abs. Gap - AARC		_	_	
	Relative Regret	Avg CPU time (s) - AARC	352.5	319.7	346.9	451.3
	(WCRRM)	Avg CPU time (s) - C&CG*	>14,400	>14,400	>14,400	>14,400

Table 1.5: Production-Transportation Problem

models were solved in less than 8 minutes while C&CG takes around 2 hours. It is also obvious that the RO model is more tractable than WCARM and WCRRM due to the fact that uncertainty is limited to the objective function. Moreover, it appears that the WCRRM model is especially difficult to solve exactly in this setting while the AARC approach once again performs surprisingly well both in terms of computation time and quality of solutions. Indeed, the average absolute gap remained under < 1% for all categories of instances where exact solutions could be identified.

In order to shed more light on the difficulties of solving WCRRM for larger size problems, we present in Table 1.6 a description of the performance of both AARC and C&CG^{*} for each of the 10 large problem instances for which C&CG^{*} was unable to converge in less than 4 hours. In particular, the table shows that when $\Gamma = 0.3m$, for 3 out of 10 instances, the C&CG^{*} algorithm is unable to provide the cuts needed to bound the minimal worst-case relative regret away from 0%. Furthermore, in instance #5, it is even unable to identify the most violated constraint in its first iteration within the allotted time. This

				Uncerta	inty Level	(in % of <i>m</i>)			
Ins.		30%			70%			100%	
	AARC UB	C&CG* LB	AARC Abs. gap	AARC UB	C&CG* LB	AARC Abs. gap	AARC UB	C&CG* LB	AARC Abs. gap
1	4.72%	0.00%	< 4.72%	7.05%	0.00%	< 7.05%	7.62%	0.00%	< 7.62%
2	4.14%	3.68%	$\stackrel{-}{\leq} 0.45\%$	5.88%	0.00%*	\leq 5.88%	6.04%	0.00%	$\stackrel{-}{\leq} 6.04\%$
3	2.78%	0.00%	$\leq 2.78\%$	4.47%	0.00%	$\leq 4.47\%$	4.77%	0.00%*	$\leq 4.77\%$
4	4.30%	2.20%	$\leq 2.10\%$	7.09%	0.00%*	$\leq 7.09\%$	7.35%	0.00%	\leq 7.35%
5	3.05%	0.00%*	$\leq 3.05\%$	4.54%	0.00%*	$\leq 4.54\%$	4.64%	0.00%*	$\leq 4.64\%$
6	3.79%	2.28%	$\leq 1.51\%$	6.15%	0.00%	$\leq 6.15\%$	6.45%	0.00%*	$\leq 6.45\%$
7	3.64%	2.61%	$\leq 1.03\%$	5.70%	0.00%*	$\leq 5.70\%$	5.93%	0.00%*	$\leq 5.93\%$
8	6.28%	1.40%	$\leq 4.88\%$	9.67%	0.00%	$\leq 9.67\%$	9.98%	0.00%	$\leq 9.98\%$
9	4.94%	1.40%	$\leq 3.54\%$	7.32%	0.00%	$\leq 7.32\%$	7.67%	0.00%*	$\leq 7.67\%$
10	2.57%	0.56%	$\leq 2.01\%$	3.91%	0.00%	$\leq 3.91\%$	4.13%	0.00%*	$\leq 4.13\%$

Table 1.6: WCRRM - Production-Transportation Problem with 7 Facilities and 14 Customers

* indicates that C&CG* was unable to identify the most violated constraint within 4 hours in its first iteration.

phenomenon becomes more frequent as Γ is increased. In the limit when $\Gamma = m$, 6 out of the 10 instances did not complete their first round of constraint generation because of the difficulty of the subproblem. For sake of completeness, we provide the bounds that can be computed on the optimality gap of AARC given the state of the C&CG* algorithm after four hours. Overall, these seem to support the idea that, for this class of problems, AARC is a valuable approximation scheme and that the design of efficient exact algorithms constitutes a promising direction for future research.

1.8 Appendix

1.8.1 Illustrative Example with Newsvendor Problem

Consider a simple newsvendor problem:

$$\max_{x\geq 0} p\min(x,\zeta) - cx\,,$$

where $x \in \mathbb{R}$ is the number of newspapers ordered, p > 0 is the sales price, c < p is the ordering cost, and $\zeta > 0$ is the demand for the newspaper only known to lie in an interval $\mathcal{U} := [\bar{\zeta} - \hat{\zeta}, \bar{\zeta} + \hat{\zeta}]$, with $\bar{\zeta} > 0$ as the nominal demand and $\hat{\zeta} < \bar{\zeta}$ as the maximum deviation. In this context, one can consider four different models. First, the so-called nominal model simply solves the newsvendor problem under the nominal demand $\bar{\zeta}$ and leads to the unique optimal solution $x_{\text{nom}}^* = \bar{\zeta}$. Second, the classical robust optimization model takes the form:

$$\max_{x \ge 0} \min_{\zeta \in \mathcal{U}} p \min(x, \zeta) - cx,$$

with its optimal solution uniquely achieved by $x_{rob}^* = \bar{\zeta} - \hat{\zeta}$, i.e. the lowest demand possible. Third, one might consider the worst-case absolute regret minimization problem:

$$\min_{x \ge 0} \max_{\zeta \in \mathcal{U}} \max_{x' \ge 0} \left(p \min(x', \zeta) - cx' \right) - \left(p \min(x, \zeta) - cx \right).$$

This problem has as unique optimal solution $x^*_{abs} = \bar{\xi} + (1 - (2c/p))\hat{\zeta}$ when $p \leq 2c$. Fourth, one could formulate the worst-case relative regret minimization problem:

$$\min_{x \ge 0} \min_{\zeta \in \mathcal{U}} \frac{\max_{x' \ge 0} (p \min(x', \zeta) - cx') - (p \min(x, \zeta) - cx)}{\max_{x' \ge 0} (p \min(x', \zeta) - cx')}$$

The unique optimal solution to this problem is $x_{\text{rel}}^* = (\bar{\zeta}^2 - \hat{\zeta}^2)/(\bar{\zeta} + (2c/p - 1)\hat{\zeta}).$

In this context, two key properties are worth discussing. First, one can show that the four different optimal solutions follow a certain order $x_{rob}^* \leq x_{rel}^* \leq x_{abs}^* \leq x_{nom}^*$, as long as $p \leq 2c$. This property provides some arguments that support the popular conclusion that regret minimizing solutions are "less conservative" than the solutions of robust optimization problem. Indeed, both x_{rel}^* and x_{abs}^* recommend submitting larger orders than x_{rob}^* .

Another interesting property is that x^*_{abs} turns out to be the optimal solution of the stochastic program

$$\max_{x\geq 0} \mathbb{E}[\min(x,\zeta) - cx],$$

when ζ is considered uniformly distributed on \mathcal{U} . This again points to the fact that worstcase absolute regret minimizers might offer a better balance between risks and returns compared to robust optimization.

1.8.2 Summary Tables for the Literature on Regret Minimization

Tables 1.7 and 1.8 respectively present a summary of the algorithmic developments of the last 25 years regarding the resolution of worst-case regret minimization problems involving single-stage and two-stage models respectively.

		Solution		Scope	
Reference	Algorithm	Туре	Regret Type	Uncertain Parameters	Uncertainty Set
Inuiguchi and Kume (1994)	Constraint Generation + Vertex Enumeration	Exact	Absolute	Obj	Box
Inuiguchi and Sakawa (1995)	Constraint Generation + Vertex Enumeration	Exact	Absolute	Obj	Box
Inuiguchi and Sakawa (1996)	Constraint Generation + MILP Reformulation	Exact	Absolute	Obj	Box
Mausser and Laguna (1998)	Constraint Generation + MILP Reformulation	Exact	Absolute	Obj	Box
Mausser and Laguna (1999a)	Constraint Generation + MILP Reformulation + Greedy Search	Exact	Absolute	Obj	Box
Inuiguchi and Sakawa (1997a)	Constraint Generation + Vertex Enumeration	Exact	Relative	Obj	Box
Mausser and Laguna (1999b)	Constraint Generation + MILP Reformulation	Exact	Relative	Obj	Box
Bertsimas and Dunning (2020)	Constraint Generation + MILP Reformulation	Exact	Absolute Relative	Obj	Budgeted
Inuiguchi and Sakawa (1997b)	Constraint Generation + MILP Reformulation	Exact	Absolute	Obj	Polyhedral
Inuiguchi et al. (1999)	Constraint Generation + Outer Approx. Scheme	Exact	Absolute	Obj	Polyhedral
Inuiguchi and Tanino (2001)	Constraint Generation + Outer Approx. Scheme + Cutting-hyperplanes scheme	Exact	Absolute	Obj	Polyhedral
Gabrel and Murat (2010)	LP Reformulation	Exact	Absolute	RHS	Box
Bertsimas and Dunning (2020)	LP Reformulation	Exact	Absolute Relative	RHS	Polyhedral

Table 1.7: General Approaches - Linear Single-Stage Problems

* RHS and Obj refer to the right-hand side and objective uncertainty, respectively.

	זמחוב זיטי	זכווכומו אן	prudutes	- 1W U-DIABE I LUDICIUS		
Dofemento	Alconithm	Solution		Scope		
INCRETENCE	unningter	Type	Regret Type	Variables (First-stage + Recourse)	Uncertain Parameters	Uncertainty Set
Assavapokee et al. (2008b)	C&CG + Exhaustive Search	Exact	Absolute Relative	Binary + Continuous	All Parameters	Discrete Scenarios
Assavapokee et al. (2008a)	C&CG + MILP Reformulation	Exact	Absolute	Binary + Continuous	RHS + First-stage Technology Matrix	Box
Jiang et al. (2013)	Constraint Generation + Coordinate Ascent	O.A.	Absolute	Binary + Continuous	RHS	Polyhedral
Ng (2013)	Constraint Generation + MILP Reformulation	C.A.	Absolute	Continuous + Continuous	RHS + Obj	Polyhedral
Chen et al. (2014)	C&CG + MILP Reformulation	Exact	Absolute	Binary + Continuous	RHS	Polyhedral
Ning and You (2018)	C&CG + MILP Reformulation	O.A.	Absolute	Continuous + Continuous	RHS	Polyhedral
Bertsimas and Dunning (2020)	LP Reformulation	C.A.	Absolute Relative	Continuous + Continuous	RHS	Polyhedral
* C.A. and O.A. stand for Conserva	itive and Optimistic Approxi	mations, resp	ectively.			

Table 1.8: General Approaches - Two-Stage Problems

48

1.8.3 Ning and You (2018)'s C&CG Approach Is an Optimistic Approximation

Consider the multi-item newsvendor problem presented in Section 1.7.1 where we let $n_x = n_y = n_{\zeta} = 2$ items, the sales price be $p_i = 1$, ordering cost $c_i = 1$, salvage value $s_i = 0$, and shortage cost $b_i = 1$. We also consider that the two items have a nominal demand of 50 and 25 with maximum deviations of 50 and 25 respectively and that the sum of absolute relative deviations must be smaller or equal to one, i.e. $\Gamma = 1$. Moreover, we consider that the maximum total number of items ordered must be smaller or equal to 100, namely that $\mathcal{X} := \{ x \in \mathbb{R}^2_+ | x_1 + x_2 \le 100 \}$. In this context, one can show numerically that the minimal worst-case absolute regret is equal to 45.833 and achieved by ordering 44.657 units of item #1 and 23.824 units of item #2. On the other hand, the C&CG approach proposed in Ning and You (2018) recommends ordering 37.5 units of item #1 and 25 units of item #2, estimating the minimal worst-case absolute regret achieved by this solution to be 37.5 when it is actually 54.167. In particular, when the solution (37.5, 25) is used, one can easily confirm that if only integer values for δ^+ and δ^- are considered in the uncertainty set, then for all possible cases the regret achieved is 37.5. However, this is an underestimation of the regret that is achieved over $\mathcal{U}(\Gamma)$ since at $\zeta = (250/3, 50/3)$ is equal to $325/6 \approx 54.167$. This confirms that the C&CG approach proposed in Ning and You (2018) solves an optimistic approximation of the WCARM problem.

1.8.4 Bertsimas and Dunning (2020)'s Conservative Approximation Is Weaker than the Approximation Obtained with Problem (1.15)

Consider again the multi-item newsvendor problem presented in Section 1.7.1 and Section 1.8.3 where we let $n_x = n_y = n_\zeta = 2$ items, the sales price be $p_i = 1$, ordering cost $c_i = 1$, salvage value $s_i = 0$, and shortage cost $b_i = 1$. We also consider that the two items have a nominal demand of 50 and 25 with maximum deviations of 50 and 25 respectively and that the sum of absolute relative deviations must be smaller or equal to one, i.e. $\Gamma = 1$. Moreover, we consider that the maximum total number of items ordered must be smaller or equal to 100, namely that $\mathcal{X} := \{ \boldsymbol{x} \in \mathbb{R}^2_+ | x_1 + x_2 \leq 100 \}$. In this context, one can show numerically that the minimal worst-case absolute regret is equal to 45.833

and achieved by ordering 44.657 units of item #1 and 23.824 units of item #2. The conservative approximation proposed in Bertsimas and Dunning (2020) recommends ordering 50 units of item #1 and 25 units of item #2, estimating the minimal worst-case absolute regret achieved by this solution to be below 50, which is actually implying that while the solution is not optimal itself, its worst-case regret bound is exact. Alternatively, the conservative approximation in Problem (1.15) recommends ordering 45.833 units of item #1 and 25 units of item #2, estimating the minimal worst-case absolute regret achieved by this solution is not optimal itself, its worst-case regret bound is exact. Alternatively, the conservative approximation in Problem (1.15) recommends ordering 45.833 units of item #1 and 25 units of item #2, estimating the minimal worst-case absolute regret achieved by this solution to be below 45.833, which is actually exact and optimal.

1.8.5 Zeng and Zhao (2013)'s Mixed-integer Linear Programming Reformulation of C&CG's Sub-Problem

In Zeng and Zhao (2013), the authors propose a column-and-constraint generation method for solving the TSLRO problem. A key step consists in solving the NP-hard adversarial problem $\min_{\boldsymbol{\zeta} \in \mathcal{U}_v} h(\boldsymbol{x}, \boldsymbol{\zeta})$ in order to identify new columns and constraints to add to Problem (1.5). They show that this can be done by reformulating this adversarial problem as the following mixed-integer linear program:

$$\min_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{y}, \boldsymbol{\lambda}, \boldsymbol{u}} \quad \boldsymbol{x}^T C \boldsymbol{\zeta} + \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} + \boldsymbol{f}^T \boldsymbol{\zeta}$$
(1.26a)

ubject to
$$A\mathbf{x} + B\mathbf{y} \le \Psi(\mathbf{x})\boldsymbol{\zeta} + \boldsymbol{\psi}$$
 (1.26b)

$$\lambda \ge 0 \tag{1.26c}$$

$$\lambda \le M \boldsymbol{u}$$
 (1.26d)

$$\Psi(\boldsymbol{x})\boldsymbol{\zeta} + \boldsymbol{\psi} - A\boldsymbol{x} - B\boldsymbol{y} \le M(1-\boldsymbol{u})$$
(1.26e)

$$\boldsymbol{d} = B^T \boldsymbol{\lambda} \tag{1.26f}$$

$$\boldsymbol{u} \in \{0,1\}^m, \tag{1.26g}$$

where $\boldsymbol{y} \in \mathbb{R}^{n_y}$, $\boldsymbol{\lambda} \in \mathbb{R}^m$, and M is some large enough constant.

1.8.6 **Proof of Lemma 1.3.1**

Based on Assumption 1.3.2, for all $x \in \mathcal{X}$ and all $\zeta \in \mathcal{U}$, there exists a y for which Problem (1.10) is feasible. Therefore, strong duality property holds for Problem (1.10) and duality
can be used to reformulate it as a minimization problem:

$$h(\boldsymbol{x},\boldsymbol{\zeta}) := \inf_{\boldsymbol{\rho}} \quad \boldsymbol{c}^{T}\boldsymbol{x} + (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho}$$
(1.27a)

s.t.
$$B^T \boldsymbol{\rho} = \boldsymbol{d}(\boldsymbol{\zeta})$$
 (1.27b)

$$\boldsymbol{\rho} \ge 0\,, \tag{1.27c}$$

where $\rho \in \mathbb{R}^m$ is the dual variable associated to constraint (1.10b). Therefore, the Problem (1.9) can be rewritten as Problem (1.28):

$$\underset{\boldsymbol{x} \in \mathcal{X}}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta} \in \mathcal{U}} h(\boldsymbol{x}, \boldsymbol{\zeta}) \equiv \underset{\boldsymbol{x} \in \mathcal{X}}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta}, \boldsymbol{\rho}} \quad \boldsymbol{c}^{T} \boldsymbol{x} + (\boldsymbol{\psi} - A \boldsymbol{x})^{T} \boldsymbol{\rho} \quad (1.28a)$$

s.t.

$$B^T \boldsymbol{\rho} = D\boldsymbol{\zeta} + \boldsymbol{d} \qquad (1.28b)$$

$$\boldsymbol{\rho} \ge 0 \tag{1.28c}$$

$$P\boldsymbol{\zeta} \leq \boldsymbol{q}$$
, (1.28d)

where we exploited the definition of $d(\zeta)$.

According to Assumption 1.3.3, for all $x \in \mathcal{X}$ there is a $\hat{\zeta} \in \mathcal{U}$ for which Problem (1.10) is bounded, and it has a finite optimal value based on Assumption 1.3.2. By the strong duality property, Problem (1.27) must also have a finite optimal value for the same $\hat{\zeta}$, hence it must have a feasible solution $\hat{\rho}$. We conclude that $(\hat{\zeta}, \hat{\rho})$ is a feasible solution for Problem (1.27). Therefore, strong duality applies for the minimization problem in (1.28) and ensures that

$$\inf_{\boldsymbol{\zeta} \in \mathcal{U}} h(\boldsymbol{x}, \boldsymbol{\zeta}) = \sup_{\boldsymbol{y}', \boldsymbol{\lambda}, \boldsymbol{\gamma}} c^T \boldsymbol{x} + d^T \boldsymbol{y}' - \boldsymbol{q}^T \boldsymbol{\lambda}$$

s.t. $A \boldsymbol{x} + B \boldsymbol{y}' + \boldsymbol{\gamma} = \boldsymbol{\psi}$
 $P^T \boldsymbol{\lambda} + D^T \boldsymbol{y}' = 0$
 $\boldsymbol{\gamma} \ge 0, \boldsymbol{\lambda} \ge 0,$

where $\boldsymbol{y} \in \mathbb{R}^{n_y}$, $\boldsymbol{\gamma} \in \mathbb{R}^m$ and $\boldsymbol{\lambda} \in \mathbb{R}^s$ are the dual variables associated with the constraints (1.28b), (1.28c), and (1.28d) respectively. This maximization problem can be reintegrated with the maximization over $\boldsymbol{x} \in \mathcal{X}$ to obtain Problem (1.11).

1.8.7 Proof of Proposition 1.4.1

By substituting Problem (1.3) in Problem (1.12) after replacing C = 0, f = 0, and $\Psi(x) = \Psi$ as prescribed by Definition 1.3.1, we can proceed with the following simple steps:

$$WCARM \equiv \min_{\boldsymbol{x} \in \mathcal{X}} \sup_{\boldsymbol{\zeta} \in \mathcal{U}} \left\{ \sup_{\boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x}' + \boldsymbol{d}^T \boldsymbol{y}' - \sup_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} \right\}$$

$$(1.29a)$$

$$\equiv \min_{\boldsymbol{x} \in \mathcal{X}} \sup_{\boldsymbol{\zeta} \in \mathcal{U}} \sup_{\boldsymbol{x}' \in \mathcal{Y}, \boldsymbol{x}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \inf_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x}' + \boldsymbol{d}^T \boldsymbol{y}' - \boldsymbol{c}^T \boldsymbol{x} - \boldsymbol{d}^T \boldsymbol{y} \quad (1.29b)$$

$$\equiv \underset{\boldsymbol{x}\in\mathcal{X}}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta}\in\mathcal{U}, \boldsymbol{x}'\in\mathcal{X}, \boldsymbol{y}'\in\mathcal{Y}(\boldsymbol{x}',\boldsymbol{\zeta})} \sup_{\boldsymbol{y}\in\mathcal{Y}(\boldsymbol{x},\boldsymbol{\zeta})} -\boldsymbol{c}^{T}\boldsymbol{x}' - \boldsymbol{d}^{T}\boldsymbol{y}' + \boldsymbol{c}^{T}\boldsymbol{x} + \boldsymbol{d}^{T}\boldsymbol{y},$$
(1.29c)

where $\mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta}) := \{ \boldsymbol{y} \in \mathbb{R}^{n_{\boldsymbol{y}}} | A\boldsymbol{x} + B\boldsymbol{y} \leq \Psi \boldsymbol{\zeta} + \boldsymbol{\psi} \}$, and where we simply regrouped the minimization and maximization operations together, and later rewrote the minimization problem as a maximization problem with the understanding that an optimal value for WCARM can be obtained by changing the sign of the optimal value returned from Problem (1.29c).

In order to formulate a TSLRO model, we simply consider a lifted uncertain vector composed as $\boldsymbol{\zeta}' := [\boldsymbol{\zeta}^T \ \boldsymbol{x'}^T \ \boldsymbol{y'}^T]^T$, which needs to realize inside the polyhedron defined as

$$\mathcal{U}' := \{ [\boldsymbol{\zeta}^T \ \boldsymbol{x}^T \ \boldsymbol{y}^T]^T \in \mathbb{R}^{n_{\boldsymbol{\zeta}} + n_{\boldsymbol{x}} + n_{\boldsymbol{y}}} \ | \ P \boldsymbol{\zeta} \leq \boldsymbol{q}, \ \boldsymbol{x} \in \mathcal{X}, \ A \boldsymbol{x} + B \boldsymbol{y} \leq \Psi \boldsymbol{\zeta} + \boldsymbol{\psi} \}$$

One also needs to consider that since ζ has been lifted to ζ' , the recourse decision y can depend on all the information revealed by ζ' . This completes the proof of how the TSLRO model presented in (1.13) is equivalent to the WCARM.

We now verify the conditions under which all four assumptions are satisfied by this new TSLRO. Firstly, given that Assumption 1.3.1 is satisfied for the WCARM problem, there must exist a triplet $(\bar{x}, \bar{\zeta}, \bar{y})$ that is such that $\bar{x} \in \mathcal{X}, \bar{\zeta} \in \mathcal{U}$, and $\bar{y} \in \mathcal{Y}(\bar{x}, \bar{\zeta})$. It is then straightforward to confirm that $\zeta' := [\bar{\zeta}^T \ \bar{x}^T \ \bar{y}^T]^T$ must be a member of \mathcal{U}' so that the triplet $(\bar{x}, \zeta', \bar{y})$ satisfies the same condition for the new TSLRO problem (1.13). We conclude from this that Assumption 1.3.1 applies. Secondly, given that the feasible set for the recourse problem is the same in WCARM and its new TSLRO reformulation, Assumption 1.3.2 carries over to the new TSLRO problem. Thirdly, one can show that Assumption 1.3.3 also carries through if Assumption 1.3.2 holds. Specifically, we start by letting $\bar{\zeta} : \mathbb{R}^{n_x} \to \mathbb{R}^{n_{\zeta}}$ be a policy that verifies that Assumption 1.3.3 holds for the WCARM problem and letting (x', y') be a feasible first-stage and recourse policy, which exists based on Assumption 1.3.2. One can construct a policy $\zeta'(x) := [\bar{\zeta}(x)^T x'^T y'^T]^T$ that will make Assumption 1.3.3 hold for the TSLRO problem. Finally, Assumption 1.3.4 carries through to the new TSLRO as long as the WCARM also satisfies Assumption 1.4.1. Indeed, when both assumptions are satisfied by the WCARM problem, we know that:

$$\begin{split} \inf_{\boldsymbol{\zeta}' \in \mathcal{U}'} h'(\boldsymbol{x}, \boldsymbol{\zeta}') &= \inf_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \sup_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} - \boldsymbol{c}^T \boldsymbol{x}' - \boldsymbol{d}^T \boldsymbol{y}' + \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} \\ &\geq \inf_{\boldsymbol{\zeta} \in \mathcal{U}} h(\boldsymbol{x}, \boldsymbol{\zeta}) - \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x}' + \boldsymbol{d}^T \boldsymbol{y}' \\ &\geq \inf_{\boldsymbol{\zeta} \in \mathcal{U}} h(\boldsymbol{x}, \boldsymbol{\zeta}) - \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) > -\infty \,, \end{split}$$

where we denoted the recourse problem that appears in the TSLRO reformulation as $h'(\boldsymbol{x}, \boldsymbol{\zeta}')$.

1.8.8 Proof of Proposition 1.4.2

Let us consider the following maximization problem, which is part of the WCARM problem with objective uncertainty:

$$\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) = \sup_{\boldsymbol{x}', \boldsymbol{y}'} \boldsymbol{c}^T \boldsymbol{x}' + \boldsymbol{d}^T(\boldsymbol{\zeta}) \boldsymbol{y}'$$
(1.30a)

s.t.
$$Ax' + By' \le \psi$$
 (1.30b)

$$W \boldsymbol{x}' \le \boldsymbol{v} \,. \tag{1.30c}$$

Based on Assumption 1.3.2, there necessarily exists a pair (x', y') that makes problem (1.30) feasible. Therefore, strong duality holds and the dual form of Problem (1.30) can be derived by introducing the dual variables $\lambda \in \mathbb{R}^m$ and $\gamma \in \mathbb{R}^r$ associated with constraints (1.30b) and (1.30c), respectively. Thus, we obtain:

$$\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) = \inf_{\boldsymbol{\lambda} \ge 0, \boldsymbol{\gamma} \ge 0} \boldsymbol{\psi}^T \boldsymbol{\lambda} + \boldsymbol{v}^T \boldsymbol{\gamma}$$
(1.31a)

s.t.
$$A^T \boldsymbol{\lambda} + W^T \boldsymbol{\gamma} = \boldsymbol{c}$$
 (1.31b)

$$B^T \boldsymbol{\lambda} = \boldsymbol{d}(\boldsymbol{\zeta}) \,. \tag{1.31c}$$

Since the strong duality property holds for both Problems (1.10) and (1.30), it is possible to rewrite the WCARM problem by substituting both $h(x, \zeta)$ and $\sup_{x' \in \mathcal{X}} h(x', \zeta)$ using their respective dual form, which results in the following reformulation:

$$WCARM \equiv \min_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) - h(\boldsymbol{x},\boldsymbol{\zeta}) \right\}$$

$$\equiv \min_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) - \inf_{\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})} \boldsymbol{c}^{T}\boldsymbol{x} + (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho} \right\}$$

$$\equiv \min_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})} \left\{ \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) - \boldsymbol{c}^{T}\boldsymbol{x} + (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho} \right\}$$

$$\equiv \min_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})} \inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})} \boldsymbol{\psi}^{T}\boldsymbol{\lambda} + \boldsymbol{v}^{T}\boldsymbol{\gamma} - \boldsymbol{c}^{T}\boldsymbol{x} - (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho}$$

$$\equiv \max_{\boldsymbol{x}\in\mathcal{X}} \inf_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})} \sup_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})} - \boldsymbol{\psi}^{T}\boldsymbol{\lambda} - \boldsymbol{v}^{T}\boldsymbol{\gamma} + \boldsymbol{c}^{T}\boldsymbol{x} + (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho},$$

(1.32)

where $\Upsilon_1(\boldsymbol{\zeta}) := \{(\boldsymbol{\lambda}, \boldsymbol{\gamma}) \in \mathbb{R}^m \times \mathbb{R}^r \mid \boldsymbol{\lambda} \ge 0, \boldsymbol{\gamma} \ge 0, \text{ (1.31b), (1.31c)}\} \text{ and } \Upsilon_2(\boldsymbol{\zeta}) := \{\boldsymbol{\rho} \in \mathbb{R}^m \mid B^T \boldsymbol{\rho} = \boldsymbol{d}(\boldsymbol{\zeta}), \ \boldsymbol{\rho} \ge 0\}.$ By using the two liftings $\boldsymbol{\zeta}' = \begin{bmatrix} \boldsymbol{\zeta} \\ \boldsymbol{\rho} \end{bmatrix} \text{ and } \boldsymbol{y}'(\boldsymbol{\zeta}) := \begin{bmatrix} \boldsymbol{\lambda}(\boldsymbol{\zeta}) \\ \boldsymbol{\gamma}(\boldsymbol{\zeta}) \end{bmatrix},$ Problem (1.32) can be rewritten in the form presented in equation (1.16).

Regarding the conditions on WCARM for the TSLRO reformulation to satisfy some of the stated assumptions, we start by considering that WCARM satisfies Assumptions 1.3.1, 1.3.2, 1.3.3, and 1.4.1. Based on Assumption 1.3.3, it is possible to identify an $\bar{x} \in \mathcal{X}$ and $\bar{\zeta} \in \mathcal{U}$ such that $h(\bar{x}, \bar{\zeta})$ is bounded. This implies by LP duality that there must be a feasible $\bar{\rho} \in \Upsilon_2(\bar{\zeta})$. Moreover, Assumption 1.4.1 implies that $\sup_{x' \in \mathcal{X}} h(x', \bar{\zeta})$ is bounded hence once again LP duality ensures that there exists a pair $(\bar{\lambda}, \bar{\gamma}) \in \Upsilon_1(\bar{\zeta})$. The TSLRO reformulation therefore satisfies Assumption 1.3.1 using the quintuplet $(\bar{x}, \bar{\zeta}, \bar{\rho}, \bar{\lambda}, \bar{\gamma})$. Next, the fact that the TSLRO reformulation satisfies Assumption 1.3.2 follows similarly from imposing Assumption 1.4.1 on WCARM since the existence of a pair $(\bar{\lambda}, \bar{\gamma}) \in \Upsilon_1(\zeta)$ holds for all $\zeta \in \mathcal{U}$. Finally, Assumption 1.3.3 implies that there exists a $\bar{\zeta}(x) \in \mathcal{U}$ such that, for all $x \in \mathcal{X}$, $h(x, \bar{\zeta}(x)) < \infty$. From this, we can conclude that:

$$\inf_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) - h(\boldsymbol{x},\boldsymbol{\zeta}) \right\} \geq \inf_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\bar{\boldsymbol{\zeta}}(\boldsymbol{x})) - h(\boldsymbol{x},\bar{\boldsymbol{\zeta}}(\boldsymbol{x})) \geq 0 > -\infty$$

The WCARM problem is therefore bounded below by zero hence the TSLRO reformulation is bounded above by zero, which demonstrates that the latter satisfies Assumption 1.3.3. Now, given that the WCARM additionally satisfies Assumption 1.3.4, we therefore have that for all $x \in \mathcal{X}$:

$$\sup_{\boldsymbol{\zeta}\in\mathcal{U}}\left\{\sup_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\zeta})-h(\boldsymbol{x},\boldsymbol{\zeta})\right\}\leq\left(\sup_{\boldsymbol{\zeta}\in\mathcal{U}}\sup_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\zeta})\right)-\left(\inf_{\boldsymbol{\zeta}\in\mathcal{U}}h(\boldsymbol{x},\boldsymbol{\zeta})\right)<\infty\,,$$

where the first term is bounded above according to Assumption 1.4.1 and the second term bounded below according to Assumption 1.3.4. We can thus conclude that for all $x \in \mathcal{X}$, the worst-case regret is bounded above, thus that for all $x \in \mathcal{X}$ the "worst-case profit" achievable in the TSLRO reformulation is bounded below, i.e. Assumption 1.3.4 is satisfied by the TSLRO reformulation.

1.8.9 Proof of Proposition 1.5.1

We first employ an epigraph form for Problem (1.18) as follows:

$$\min_{\boldsymbol{x}\in\mathcal{X},t} \quad t \tag{1.33a}$$

subject to
$$\sup_{\boldsymbol{\zeta} \in \mathcal{U}} \left\{ \frac{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})}{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})} \right\} \le t$$
(1.33b)

$$0 \le t \le 1 \,, \tag{1.33c}$$

where we impose that $t \in [0, 1]$ since Assumptions 1.4.1 and 1.5.1 ensure that the optimal value of the WCRRM problem is in [0, 1]. One can then manipulate constraint (1.33b) to show that it is equivalent to

$$\frac{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})}{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})} \le t , \, \forall \boldsymbol{\zeta} \in \mathcal{U} \,,$$

and moreover to

$$\sup_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\zeta})-h(\boldsymbol{x},\boldsymbol{\zeta})\leq t(\sup_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\zeta}))\;,\;\forall\boldsymbol{\zeta}\in\mathcal{U}\;,$$

since it is clearly the case if $\boldsymbol{\zeta}$ is such that $\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) > 0$ and otherwise would lead to the constraint that $-h(\boldsymbol{x}, \boldsymbol{\zeta}) \leq 0$, which is necessarily satisfied and is coherent with the fact that we consider regret to be equal to 0 for such a $\boldsymbol{\zeta}$. Finally, we obtain the constraint:

$$(1-t)\sup_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\zeta})-h(\boldsymbol{x},\boldsymbol{\zeta})\leq 0,\,\forall\boldsymbol{\zeta}\in\mathcal{U}\,.$$
(1.34)

By substituting Problem (1.3) in this constraint we obtain the following reformulations

$$(1.33b) \equiv (1-t) \sup_{\boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x}' + \boldsymbol{d}^T \boldsymbol{y}' - \sup_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} \leq 0, \ \forall \boldsymbol{\zeta} \in \mathcal{U}$$
$$\equiv \inf_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} (1-t) \boldsymbol{c}^T \boldsymbol{x}' + (1-t) \boldsymbol{d}^T \boldsymbol{y}' - \boldsymbol{c}^T \boldsymbol{x} - \boldsymbol{d}^T \boldsymbol{y} \leq 0, \ \forall \boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})$$

Hence the WCRRM problem reduces to:

$$\min_{\boldsymbol{x} \in \mathcal{X}, t \in [0,1]} \quad \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} h'(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{x}', \boldsymbol{y}') ,$$

where

$$\begin{aligned} h'(\boldsymbol{x},t,\boldsymbol{\zeta},\boldsymbol{x}',\boldsymbol{y}') &:= & \inf_{\boldsymbol{y}} & t \\ & \text{s.t.} & -\boldsymbol{c}^T\boldsymbol{x} - \boldsymbol{d}^T\boldsymbol{y} \leq -(1-t)\boldsymbol{c}^T\boldsymbol{x}' - (1-t)\boldsymbol{d}^T\boldsymbol{y}' \\ & & A\boldsymbol{x} + B\boldsymbol{y} \leq \Psi\boldsymbol{\zeta} + \boldsymbol{\psi} \,. \end{aligned}$$

Rewriting the minimization problem as a maximization problem, we obtain the TSLRO problem presented in equation (1.19).

Regarding the assumptions that are satisfied by this TSLRO reformulation, we can straightforwardly verify that based on Assumption 1.3.1, there must be a triplet $(\bar{x}, \bar{\zeta}, \bar{y})$ such that $\bar{x} \in \mathcal{X}, \bar{\zeta} \in \mathcal{U}$, and $\bar{y} \in \mathcal{Y}(\bar{x}, \bar{\zeta})$ and construct an assignment for $\bar{x}' := \bar{x}$ and $\bar{y}' := \bar{y}$ and $\bar{t} := 0$, which satisfies all the constraints of the new TSLRO reformulation. Unfortunately, if there exists an $x \in \mathcal{X}$ such that the worst-case relative regret is strictly greater than 0, then there clearly exists a $\bar{\tau} > 0$ and a feasible triplet $(\bar{\zeta}, \bar{x}', \bar{y}')$ for which the recourse problem $h'(x, \bar{\tau}, \bar{\zeta}, \bar{x}', \bar{y}')$ becomes infeasible, hence the new TSLRO reformulation does not satisfy Assumption 1.3.2.

1.8.10 Proof of Proposition 1.5.2

The first steps of this proof are exactly as in the proof of Proposition 1.5.1 up to equation (1.34), except for the small difference that we will consider $t \in [0, 1[$, which follows from Assumption 1.5.2. Since we are now dealing with objective uncertainty, we substitute $h(\boldsymbol{x}, \boldsymbol{\zeta})$ and $\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})$ using their respective dual form (see equations (1.27) and (1.31) respectively), where strong duality follows again from Assumption 1.3.2 implied

by Assumption 1.5.1. This leads to the following reformulation:

$$(1.33b) \equiv (1-t) \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}) \le 0, \, \forall \boldsymbol{\zeta} \in \mathcal{U}$$

$$(1.35)$$

$$\equiv (1-t) \left(\inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})} \boldsymbol{\psi}^{T}\boldsymbol{\lambda} + \boldsymbol{v}^{T}\boldsymbol{\gamma} \right) - \inf_{\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})} \{ \boldsymbol{c}^{T}\boldsymbol{x} + (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho} \} \leq 0, \forall \boldsymbol{\zeta} \in \mathcal{U}$$
(1.36)

$$\equiv \inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_1(\boldsymbol{\zeta})} (1-t)(\boldsymbol{\psi}^T\boldsymbol{\lambda} + \boldsymbol{v}^T\boldsymbol{\gamma}) - \boldsymbol{c}^T\boldsymbol{x} - (\boldsymbol{\psi} - A\boldsymbol{x})^T\boldsymbol{\rho} \le 0, \,\forall \boldsymbol{\zeta} \in \mathcal{U}, \forall \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})$$
(1.37)

$$\equiv \inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})} \boldsymbol{\psi}^{T}\boldsymbol{\lambda} + \boldsymbol{v}^{T}\boldsymbol{\gamma} - \frac{1}{1-t}\boldsymbol{c}^{T}\boldsymbol{x} - \frac{1}{1-t}(\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho} \leq 0, \forall \boldsymbol{\zeta} \in \mathcal{U}, \forall \boldsymbol{\rho} \in \Upsilon_{2}(\boldsymbol{\zeta}),$$
(1.38)

where $\Upsilon_1(\zeta)$ and $\Upsilon_2(\zeta)$ are as defined in the proof of Proposition 1.4.2. Hence the WCRRM problem reduces to:

$$\min_{\boldsymbol{x} \in \mathcal{X}, t \in [0,1[} \quad \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} h'(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{\rho}) ,$$

where

$$\begin{aligned} h'(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{\rho}) &:= & \inf_{\boldsymbol{\lambda}, \boldsymbol{\gamma}} & t \\ & \text{s.t.} & \boldsymbol{\psi}^T \boldsymbol{\lambda} + \boldsymbol{v}^T \boldsymbol{\gamma} - \frac{1}{1-t} \boldsymbol{c}^T \boldsymbol{x} - \frac{1}{1-t} (\boldsymbol{\psi} - A \boldsymbol{x})^T \boldsymbol{\rho} \leq 0 \\ & A^T \boldsymbol{\lambda} + W^T \boldsymbol{\gamma} = \boldsymbol{c} \\ & B^T \boldsymbol{\lambda} = \boldsymbol{d}(\boldsymbol{\zeta}) \\ & \boldsymbol{\lambda} \geq 0, \ \boldsymbol{\gamma} \geq 0 \,. \end{aligned}$$

Using a simple replacement of variables u := 1/(1-t) and z := (1/(1-t))x and applying a monotone transformation of the objective function $t \to 1/(1-t)$, we obtain that the WCRRM is equivalently represented as

$$\min_{u \ge 1, \boldsymbol{z}: W \boldsymbol{z} \le \boldsymbol{v} u} \quad \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} h''(\boldsymbol{z}, u, \boldsymbol{\zeta}, \boldsymbol{\rho}),$$

where

$$h''(\boldsymbol{z}, \boldsymbol{u}, \boldsymbol{\zeta}, \boldsymbol{\rho}) := \inf_{\boldsymbol{\lambda}, \boldsymbol{\gamma}} \quad u$$
(1.39a)
s.t. $\boldsymbol{\psi}^T \boldsymbol{\lambda} + \boldsymbol{v}^T \boldsymbol{\gamma} - \boldsymbol{c}^T \boldsymbol{z} - (\boldsymbol{\psi} \boldsymbol{u} - A \boldsymbol{z})^T \boldsymbol{\rho} \le 0$ (1.39b)

$$A^T \boldsymbol{\lambda} + W^T \boldsymbol{\gamma} = \boldsymbol{c} \tag{1.39c}$$

$$B^T \boldsymbol{\lambda} = \boldsymbol{d}(\boldsymbol{\zeta}) \tag{1.39d}$$

$$\boldsymbol{\lambda} \ge 0, \ \boldsymbol{\gamma} \ge 0. \tag{1.39e}$$

By converting the minimization problem into a maximization problem, this problem can be rewritten in the form presented in equation (1.20).

Regarding the assumptions that are satisfied by this TSLRO reformulation, we can straightforwardly verify that based on Assumption 1.3.1, there must be a triplet $(\bar{x}, \bar{\zeta}, \bar{y})$ such that $\bar{x} \in \mathcal{X}, \bar{\zeta} \in \mathcal{U}$, and $\bar{y} \in \mathcal{Y}(\bar{x}, \bar{\zeta})$ and construct an assignment for $\bar{x}' := \bar{x}, \bar{y}' := \bar{y},$ $\bar{z} := \bar{x}$, and $\bar{u} := 1$, which satisfy all the constraints of the TSLRO reformulation. Finally, the difficulties of satisfying Assumption 1.3.2 can be demonstrated exactly as in the proof of Proposition 1.5.1.

1.8.11 Proof of Proposition 1.6.1

Starting with the case of the WCARM problem, we let $h_1(x)$ be defined as the worst-case absolute regret achieved by x, which can be captured in the following form based on Proposition 1.4.1:

$$h_1(\boldsymbol{x}) := \inf_{\boldsymbol{\zeta}' \in \mathcal{U}'} \sup_{\boldsymbol{y} \in \mathcal{Y}'(\boldsymbol{x}, \boldsymbol{\zeta}')} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} + \boldsymbol{f}'^T \boldsymbol{\zeta}' \,,$$

where

$$\mathcal{Y}'(\boldsymbol{x},\boldsymbol{\zeta}') := \{ \boldsymbol{y} \, | \, A \boldsymbol{x} + B \boldsymbol{y} \le \Psi' \boldsymbol{\zeta}' + \boldsymbol{\psi} \}$$

Alternatively, let $h_2(x)$ denote the conservative approximation of $h_1(x)$ obtained using affine decision rules:

$$h_2(\boldsymbol{x}) := \sup_{(\boldsymbol{y}, Y_{\zeta'}) \in \mathcal{Y}'_{\mathsf{aff}}(\boldsymbol{x})} \inf_{\boldsymbol{\zeta}' \in \mathcal{U}'} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T (\boldsymbol{y} + Y_{\zeta'} \boldsymbol{\zeta}') + \boldsymbol{f}'^T \boldsymbol{\zeta}' \,,$$

with

$$\mathcal{Y}_{\rm aff}'(\boldsymbol{x}) := \{(\boldsymbol{y}, Y_{\zeta'}) \,|\, A\boldsymbol{x} + B(\boldsymbol{y} + Y_{\zeta'}\boldsymbol{\zeta}') \leq \Psi'\boldsymbol{\zeta}' + \boldsymbol{\psi}, \,\forall\, \boldsymbol{\zeta}' \in \mathcal{U}'\}\,.$$

Necessarily, we have that $h_1(x) \ge h_2(x)$ since affine decision rules provide a conservative approximation. In order to demonstrate that $h_1(x) = h_2(x)$, we are left with showing that $h_2(x) \ge h_1(x)$ and proceed as follows:

$$h_2(\boldsymbol{x}) \geq \sup_{(\boldsymbol{y}, [Y_\zeta \ 0 \ 0]) \in \mathcal{Y}_{\mathsf{aff}}'(\boldsymbol{x})} \min_{\boldsymbol{\zeta}' \in \mathcal{U}'} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T (\boldsymbol{y} + [Y_\zeta \ 0 \ 0] \boldsymbol{\zeta}') + \boldsymbol{f}'^T \boldsymbol{\zeta}'$$

$$= \sup_{(\boldsymbol{y},[Y_{\zeta} \ 0 \ 0]) \in \mathcal{Y}_{aff}'(\boldsymbol{x})} \min_{\zeta \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \zeta)} c^{T}\boldsymbol{x} + d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) - c^{T}\boldsymbol{x}' - d^{T}\boldsymbol{y}'$$

$$= \sup_{(\boldsymbol{y},[Y_{\zeta} \ 0 \ 0]) \in \mathcal{Y}_{aff}'(\boldsymbol{x})} \min_{\zeta \in \mathcal{U}} c^{T}\boldsymbol{x} + d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) - \max_{\boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \zeta)} c^{T}\boldsymbol{x}' + d^{T}\boldsymbol{y}'$$

$$= \sup_{(\boldsymbol{y},Y_{\zeta}) \in \mathcal{Y}_{aff}(\boldsymbol{x})} \min_{\zeta \in \mathcal{U}} c^{T}\boldsymbol{x} + d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) - \gamma^{T}\zeta - \bar{\gamma} \qquad (1.40)$$

$$= \sup_{t,(\boldsymbol{y},Y_{\zeta}) \in \mathcal{Y}_{aff}(\boldsymbol{x})} t \qquad (1.41)$$

$$= t \leq c^{T}\boldsymbol{x} + d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) - \gamma^{T}\zeta - \bar{\gamma}, \forall \zeta \in \mathcal{U}$$

$$= \max_{t,\boldsymbol{y}(\cdot)} t \qquad (1.42)$$

$$\text{s.t.} \quad t \leq c^{T}\boldsymbol{x} + d^{T}\boldsymbol{y}(\zeta) - \gamma^{T}\zeta - \bar{\gamma}, \forall \zeta \in \mathcal{U}$$

$$= \min_{\xi \in \mathcal{U}} \max_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x},\zeta)} c^{T}\boldsymbol{x} + d^{T}\boldsymbol{y} - \gamma^{T}\zeta - \bar{\gamma} \qquad (1.43)$$

$$= \min_{\zeta \in \mathcal{U}} \max_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x},\zeta)} c^{T}\boldsymbol{x} + d^{T}\boldsymbol{y} - \max_{\boldsymbol{x} \in \mathcal{X}} h(\boldsymbol{x},\zeta)$$

$$= h_{1}(\boldsymbol{x}),$$

where

$$\mathcal{Y}_{\text{aff}} := \{(\boldsymbol{y}, Y_{\zeta}) \,|\, A\boldsymbol{x} + B(\boldsymbol{y} + Y_{\zeta}) \leq \Psi \boldsymbol{\zeta} + \boldsymbol{\psi}, \,\forall \, \boldsymbol{\zeta} \in \mathcal{U} \} \,.$$

Detailing each step, we first obtained a lower bound by maximizing over a subset of the available affine decision rules. We then in the next three steps exploited the property that $\max_{x \in \mathcal{X}} h(x, \zeta) = \gamma^T \zeta + \bar{\gamma}$. The fourth step consists in using an epigraph representation to cast the model in a form where all the uncertainty appears on the right-hand side. The equivalence between (1.41) and (1.42) follows from the fact that affine decision rules are optimal in two-stage robust linear programs with right-hand side uncertainty when the uncertainty set is a simplex set (see Theorem 1 in Bertsimas and Goyal 2012). Finally, the steps are completed by replacing back $\gamma^T \zeta + \bar{\gamma} = \max_{x \in \mathcal{X}} h(x, \zeta)$ to obtain the expression of worst-case absolute regret, which was defined as $h_1(x)$.

In the case of WCRRM, we can follow a similar reasoning. For any fixed x and t, we can let

$$h_1(\boldsymbol{x},t) := \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \inf_{\boldsymbol{y}} t$$

s.t.
$$-\boldsymbol{c}^T \boldsymbol{x} - \boldsymbol{d}^T \boldsymbol{y} \leq -(1-t)\boldsymbol{c}^T \boldsymbol{x}' - (1-t)\boldsymbol{d}^T \boldsymbol{y}'$$

 $A\boldsymbol{x} + B\boldsymbol{y} \leq \Psi \boldsymbol{\zeta} + \boldsymbol{\psi},$

and $h_2(\boldsymbol{x}, t)$ as the upper bound obtained when applying affine decision rules of the form $\boldsymbol{y}(\boldsymbol{\zeta}, \boldsymbol{x}', \boldsymbol{y}') := \boldsymbol{y} + Y_{\boldsymbol{\zeta}} \boldsymbol{\zeta} + Y_{x'} \boldsymbol{x}' + Y_{y'} \boldsymbol{y}'$. In this context, we can show that

$$\begin{split} h_{2}(\boldsymbol{x},t) &\leq \inf_{\boldsymbol{y},Y_{\zeta}} t \\ \text{s.t.} &- c^{T}\boldsymbol{x} - d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) \leq -(1-t)c^{T}\boldsymbol{x}' - (1-t)d^{T}\boldsymbol{y}', \qquad (1.44) \\ &\quad \forall \zeta \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}',\zeta) \\ &\quad A\boldsymbol{x} + B(\boldsymbol{y} + Y_{\zeta}\zeta) \leq \Psi \zeta + \psi, \forall \zeta \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}',\zeta) \\ &= \inf_{\boldsymbol{y},Y_{\zeta}} t \\ \text{s.t.} &- c^{T}\boldsymbol{x} - d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) \leq -(1-t)(c^{T}\boldsymbol{x}' + d^{T}\boldsymbol{y}'), \forall \zeta \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}',\zeta) \\ &\quad A\boldsymbol{x} + B(\boldsymbol{y} + Y_{\zeta}\zeta) \leq \Psi \zeta + \psi, \forall \zeta \in \mathcal{U} \\ &= \inf_{\boldsymbol{y},Y_{\zeta}} t \\ \text{s.t.} &- c^{T}\boldsymbol{x} - d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) \leq -(1-t)(\bar{\gamma} + \gamma^{T}\zeta), \forall \zeta \in \mathcal{U} \\ &\quad A\boldsymbol{x} + B(\boldsymbol{y} + Y_{\zeta}\zeta) \leq \Psi \zeta + \psi, \forall \zeta \in \mathcal{U} \\ &= \begin{cases} t & \text{if } \sup_{(\boldsymbol{y},Y_{\zeta})\in\mathcal{Y}_{\text{aff}}(\boldsymbol{x}) & \min_{\zeta \in \mathcal{U}} c^{T}\boldsymbol{x} + d^{T}(\boldsymbol{y} + Y_{\zeta}\zeta) - (1-t)(\gamma^{T}\zeta - \bar{\gamma}) \geq 0 \\ \infty & \text{otherwise} \end{cases} \\ &= \begin{cases} t & \text{if } \min_{\zeta \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \zeta) & \max_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \zeta)} \end{cases} \\ (1.45) \\ &= \end{cases} \end{split}$$

 $\int \infty$ otherwise

(1.46)

$$= \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \inf_{\boldsymbol{y}} t$$

s.t. $-\boldsymbol{c}^{T}\boldsymbol{x} - \boldsymbol{d}^{T}\boldsymbol{y} \leq -(1-t)(\boldsymbol{c}^{T}\boldsymbol{x}' + \boldsymbol{d}^{T}\boldsymbol{y}')$
 $A\boldsymbol{x} + B\boldsymbol{y} \leq \Psi \boldsymbol{\zeta} + \boldsymbol{\psi}$

$$=h_1(\boldsymbol{x},t)\,,$$

where the equivalence between (1.45) and (1.46) was already demonstrated in going through equations (1.40) to (1.43).

1.8.12 Proof of Proposition 1.6.2

Considering the case of the WCARM problem, we start by establishing a second equivalent TSLRO reformulation for Problem (1.4.2). In particular, for any fixed x, we can let

$$egin{aligned} h_1(m{x}) &\coloneqq \inf_{m{\zeta}\in\mathcal{U},m{
ho}\in\Upsilon_2(m{\zeta})} &\sup_{(m{\lambda},m{\gamma})\in\Upsilon_1(m{\zeta})} &-m{\psi}^Tm{\lambda} -m{v}^Tm{\gamma} +m{c}^Tm{x} + (m{\psi} - Am{x})^Tm{
ho} \ &= \inf_{m{\zeta}\in\mathcal{U},m{
ho}\in\Upsilon_2(m{\zeta})} &\inf_{m{x}'\in\mathcal{X},m{y}'\in\mathcal{Y}'(m{x}')} &-m{c}^Tm{x}' -m{d}(m{\zeta})^Tm{y}' +m{c}^Tm{x} + (m{\psi} - Am{x})^Tm{
ho} \ &= \inf_{m{x}'\in\mathcal{X},m{y}'\in\mathcal{Y}'(m{x}')} &\sup_{m{y}\in\mathcal{Y}(m{x}),m{\lambda}\in\mathcal{L}(m{y},m{y}')} &m{c}^T(m{x} -m{x}') +m{d}^T(m{y} -m{y}') -m{q}^Tm{\lambda}\,, \end{aligned}$$

where $\mathcal{L}(\boldsymbol{y}, \boldsymbol{y}') := \{ \boldsymbol{\lambda} \in \mathbb{R}^s_+ | P^T \boldsymbol{\lambda} = D^T (\boldsymbol{y}' - \boldsymbol{y}) \}$ and where we exploited strong duality of

$$\inf_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\rho}\in\Upsilon_2(\boldsymbol{\zeta})} \ (\boldsymbol{\psi}-A\boldsymbol{x})^T\boldsymbol{\rho}-\boldsymbol{d}(\boldsymbol{\zeta})^T\boldsymbol{y}' \ = \ \sup_{\boldsymbol{y}\in\mathcal{Y}(\boldsymbol{x}),\boldsymbol{\lambda}\in\mathcal{L}(\boldsymbol{y},\boldsymbol{y}')} \ \boldsymbol{d}^T(\boldsymbol{y}-\boldsymbol{y}')-\boldsymbol{q}^T\boldsymbol{\lambda} \in\mathcal{L}(\boldsymbol{y},\boldsymbol{y}')$$

Note that strong duality follows from Assumption 1.3.3 for the same reasons as in the case of Problem (1.28) (see proof of Proposition 1.3.1). Hence, our analysis gives rise to a dual reformulation for TSLRO (1.16).

In Bertsimas and de Ruiter (2016), it was established (see Theorem 2) that the conservative approximation obtained by employing affine decision rules on a TSLRO problem is exactly equivalent to the approximation obtained by employing affine decision rules on its dual reformulation. This implies that:

$$h_{2}(\boldsymbol{x}) := \sup_{(\boldsymbol{\lambda}(\cdot),\boldsymbol{\gamma}(\cdot))\in\Upsilon_{1}^{\operatorname{aff}}} \inf_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})} -\boldsymbol{\psi}^{T}\boldsymbol{\lambda}(\boldsymbol{\zeta},\boldsymbol{\rho}) - \boldsymbol{v}^{T}\boldsymbol{\gamma}(\boldsymbol{\zeta},\boldsymbol{\rho}) + \boldsymbol{c}^{T}\boldsymbol{x} + (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho}$$

$$(1.47)$$

$$= \sup_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\gamma}(\cdot)} \min_{\boldsymbol{\zeta}\in\mathcal{U},\boldsymbol{\gamma}(\cdot)} c^{T}(\boldsymbol{x} - \boldsymbol{x}') + d^{T}(\boldsymbol{y}(\boldsymbol{x}',\boldsymbol{y}') - \boldsymbol{y}') - \boldsymbol{q}^{T}\boldsymbol{\lambda}(\boldsymbol{x}',\boldsymbol{y}')$$

$$= \sup_{\boldsymbol{y}(\cdot)\in\mathcal{Y}^{\operatorname{aff}}(\boldsymbol{x}),\boldsymbol{\lambda}(\cdot)\in\mathcal{L}^{\operatorname{aff}}(\boldsymbol{y}(\cdot))} \min_{\boldsymbol{x}'\in\mathcal{X},\boldsymbol{y}'\in\mathcal{Y}'(\boldsymbol{x}')} \boldsymbol{c}^{T}(\boldsymbol{x}-\boldsymbol{x}') + \boldsymbol{d}^{T}(\boldsymbol{y}(\boldsymbol{x}',\boldsymbol{y}')-\boldsymbol{y}') - \boldsymbol{q}^{T}\boldsymbol{\lambda}(\boldsymbol{x}',\boldsymbol{y}')$$
(1.48)

$$= \inf_{\boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}'(\boldsymbol{x})} \sup_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}), \boldsymbol{\lambda} \in \mathcal{L}(\boldsymbol{y}, \boldsymbol{y}')} c^T(\boldsymbol{x} - \boldsymbol{x}') + d^T(\boldsymbol{y} - \boldsymbol{y}') - \boldsymbol{q}^T \boldsymbol{\lambda}$$
(1.49)
= $h_1(\boldsymbol{x})$,

where Υ_1^{aff} captures the set of all affine mappings for $\lambda : \mathbb{R}^{n_{\zeta}} \times \mathbb{R}^s \to \mathbb{R}^m$ and $\gamma : \mathbb{R}^{n_{\zeta}} \times \mathbb{R}^s \to \mathbb{R}^r$ such that $(\lambda(\zeta, \rho), \gamma(\zeta, \rho)) \in \Upsilon_1(\zeta)$ for all $\zeta \in \mathcal{U}$ and $\rho \in \Upsilon_2(\zeta), \mathcal{Y}^{\text{aff}}(x)$ captures the affine mappings $y : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \to \mathbb{R}^{n_y}$ such that $y(x', y') \in \mathcal{Y}(x)$ for all $x' \in \mathcal{X}$ and $y' \in \mathcal{Y}(x')$, and $\mathcal{L}^{\text{aff}}(y(\cdot))$ captures the affine mappings $\lambda : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \to \mathbb{R}^s$ such that $\lambda(x', y') \in \mathcal{L}(y(x', y'), y')$ for all $x' \in \mathcal{X}$ and $y' \in \mathcal{Y}(x')$. Specifically, while the equivalence between expression (1.47) and (1.48) follows from Theorem 2 of Bertsimas and de Ruiter (2016), the equivalence between (1.48) and (1.49) rather follows from Bertsimas and Goyal (2012) as exploited in the proof of Proposition 1.6.1.

In the case of WCRRM, the steps are very similar to the ones used in proving Proposition 1.6.1. We first let, for any fixed feasible u and z and their associated $x := uz \in \mathcal{X}$ and t := 1 - 1/u, the operator $h_1(u, z)$ stand for

$$h_1(u, \boldsymbol{z}) := \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} h''(\boldsymbol{z}, u, \boldsymbol{\zeta}, \boldsymbol{\rho}),$$

where $h''(z, u, \zeta, \rho)$ is as defined in equation (1.39). Furthermore, we let $h_2(u, z)$ be the upper bound achieved when using affine decision rules for λ and γ . We must then have that:

$$egin{aligned} h_2(u,oldsymbol{z}) &= \inf_{egin{aligned} & \lambda(egin{aligned} & \lambda(eta,oldsymbol{
ho}) + oldsymbol{v}^Toldsymbol{\gamma}(oldsymbol{\zeta},oldsymbol{
ho}) - oldsymbol{c}^Toldsymbol{z} - (oldsymbol{\psi}u - Aoldsymbol{z})^Toldsymbol{
ho} &\leq 0\,, \ & \forall oldsymbol{\zeta} \in \mathcal{U},\,oldsymbol{
ho} \in \Upsilon_2(oldsymbol{\zeta}) \end{aligned}$$

$$= \inf_{(\lambda(\cdot),\gamma(\cdot))\in\Upsilon_{1}^{\operatorname{aff}}} u$$

s.t.
$$\frac{1}{u}\psi^{T}\lambda(\zeta,\rho) + \frac{1}{u}v^{T}\gamma(\zeta,\rho) - \frac{1}{u}c^{T}z - (\psi - \frac{1}{u}Az)^{T}\rho \leq 0,$$
$$\forall \zeta \in \mathcal{U}, \rho \in \Upsilon_{2}(\zeta)$$
$$= \begin{cases} u \quad \text{if } \inf_{(\lambda(\cdot),\gamma(\cdot))\in\Upsilon_{1}^{\operatorname{aff}} \sup_{\zeta \in \mathcal{U}, \rho \in \Upsilon_{2}(\zeta)} \frac{1}{u}\psi^{T}\lambda(\zeta,\rho) + \frac{1}{u}v^{T}\gamma(\zeta,\rho) - c^{T}x - (\psi - Ax)^{T}\rho \leq 0\\\\\infty \quad \text{otherwise} \end{cases}$$

$$= \begin{cases} u & \text{if } \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} \inf_{(\boldsymbol{\lambda}, \boldsymbol{\gamma}) \in \Upsilon_1(\boldsymbol{\zeta})} & \frac{\frac{1}{u} \boldsymbol{\psi}^T \boldsymbol{\lambda} + \frac{1}{u} \boldsymbol{v}^T \boldsymbol{\gamma} - \\ \boldsymbol{c}^T \boldsymbol{x} - (\boldsymbol{\psi} - A \boldsymbol{x})^T \boldsymbol{\rho} \leq 0 \\ \infty & \text{otherwise} \end{cases}$$
$$= \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} \inf_{(\boldsymbol{\lambda}, \boldsymbol{\gamma}) \in \Upsilon_1(\boldsymbol{\zeta})} u \\ \text{s.t.} & \boldsymbol{\psi}^T \boldsymbol{\lambda}(\boldsymbol{\zeta}, \boldsymbol{\rho}) + \boldsymbol{v}^T \boldsymbol{\gamma}(\boldsymbol{\zeta}, \boldsymbol{\rho}) - \boldsymbol{c}^T \boldsymbol{z} - (\boldsymbol{\psi} u - A \boldsymbol{z})^T \boldsymbol{\rho} \leq 0 \end{cases}$$
$$= h_1(u, \boldsymbol{z}) .$$

Note that again here we exploit the fact that affine decision rules on

$$\sup_{\boldsymbol{\zeta}\in\mathcal{U},\,\boldsymbol{\rho}\in\Upsilon_{2}(\boldsymbol{\zeta})}\inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})}\frac{1}{u}\boldsymbol{\psi}^{T}\boldsymbol{\lambda}+\frac{1}{u}\boldsymbol{v}^{T}\boldsymbol{\gamma}-\boldsymbol{c}^{T}\boldsymbol{x}-(\boldsymbol{\psi}-A\boldsymbol{x})^{T}\boldsymbol{\rho}$$

gives the same optimal value as using it on

$$\inf_{\boldsymbol{x}',\boldsymbol{y}':\boldsymbol{u}\boldsymbol{x}'\in\mathcal{X},\boldsymbol{u}\boldsymbol{y}'\in\mathcal{Y}(\boldsymbol{u}\boldsymbol{x}')} \quad \sup_{\boldsymbol{y}\in\mathcal{Y}(\boldsymbol{x}),\boldsymbol{\lambda}\in\mathcal{L}(\boldsymbol{y},\boldsymbol{y}')} \quad \boldsymbol{c}^T(\boldsymbol{x}-\boldsymbol{x}') + \boldsymbol{d}^T(\boldsymbol{y}-\boldsymbol{y}') - \boldsymbol{q}^T\boldsymbol{\lambda}$$

which is its dual reformulation and for which we can verify that the set $\{(x', y') | ux' \in \mathcal{X}, uy' \in \mathcal{Y}(ux')\}$ is a simplex set when $\{(x, y) | x \in \mathcal{X}, y \in \mathcal{Y}(x)\}$ is one. Hence, according to Theorem 2 in Bertsimas and de Ruiter (2016) and Theorem 1 in Bertsimas and Goyal (2012), affine decision rules must be optimal in both cases.

1.8.13 **Proof of Proposition 1.6.3**

The proof proceeds in two steps. The first step consists in extending Corollary 1 in Ardestani-Jaafari and Delage (2016) to the following formulation:

$$\underset{x \in \mathcal{X}}{\operatorname{maximize}} \quad \underset{(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma)}{\operatorname{min}} h(\boldsymbol{x}, \boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-) - \bar{\gamma} - \boldsymbol{\gamma}^T(\boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-), \quad (1.50)$$

where

$$\mathcal{U}_{\pm}(\Gamma) := \{ (\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathbb{R}^m_+ \times \mathbb{R}^m_+ \, | \, \boldsymbol{\zeta}^+ + \boldsymbol{\zeta}^- \le 1, \, \sum_i \boldsymbol{\zeta}^+_i + \boldsymbol{\zeta}^-_i = \Gamma \}$$

and where $h(x, \zeta)$ is a sum of piecewise linear concave functions as defined in (1.21). Namely, that affine decision rules are optimal for Problem (1.50) when $h(x, \zeta)$ and Γ satisfy one of the three conditions described in our proposition. This can then be used to demonstrate that they are optimal for Problem (1.13) and (1.19) following the same arguments as those used in the proof of Proposition 1.6.1 where the equivalence between (1.40) and (1.43), and between (1.45) and (1.46) is now supported by what was established in the first step. For the sake of conciseness, we focus on the first step.

Lemma 1.8.1 If $h(x, \zeta)$ is a sum of piecewise linear concave functions of the form presented in (1.21), the uncertainty set U is the budgeted uncertainty set defined as in (1.22), and either of the following conditions are satisfied:

- *i*. $\Gamma = 1$
- *ii.* $\Gamma = n_{\zeta}$ and uncertainty is "additive": *i.e.* $\boldsymbol{\alpha}_{ik}(\boldsymbol{x}) = \bar{\alpha}_{ik}(\boldsymbol{x})(\sum_{\ell < i} \hat{\alpha}_{\ell}(\boldsymbol{x})\boldsymbol{e}_{\ell})$ for some $\bar{\alpha}_{ik} : \mathbb{R}^{n_{x}} \to \mathbb{R}$ for all *i* and *k* and some $\hat{\boldsymbol{\alpha}} : \mathbb{R}^{n_{x}} \to \mathbb{R}^{n_{\zeta}}$
- *iii.* Γ *is integer and objective is "decomposable": i.e.* $\boldsymbol{\alpha}_{ik}(\boldsymbol{x}) = \bar{\alpha}_{ik}(\boldsymbol{x})\boldsymbol{e}_i$ for some $\bar{\alpha}_{ik} : \mathbb{R}^{n_x} \to \mathbb{R}$ for all *i* and *k*

then, affine decision rules with respect to (δ^+, δ^-) are optimal in the following two-stage linear programming formulation of maximize_{$x \in \mathcal{X}$} min_{$\zeta \in \mathcal{U}$} $h(x, \zeta) - \bar{\gamma} - \gamma^T \zeta$:

 $\begin{array}{ll} \underset{\boldsymbol{x}\in\mathcal{X},\boldsymbol{y}(\cdot,\cdot)}{\text{maximize}} & \min_{(\boldsymbol{\zeta}^+,\boldsymbol{\zeta}^-)\in\mathcal{U}_{\pm}(\Gamma)}\sum_{i=1}^{n_y}y_i(\boldsymbol{\zeta}^+,\boldsymbol{\zeta}^-)-\bar{\gamma}-\boldsymbol{\gamma}^T(\boldsymbol{\zeta}^+-\boldsymbol{\zeta}^-)\\ \text{subject to} & y_i(\boldsymbol{\zeta}^+,\boldsymbol{\zeta}^-)\leq\boldsymbol{\alpha}_{ik}(\boldsymbol{x})^T(\boldsymbol{\zeta}^+-\boldsymbol{\zeta}^-)+\beta_{ik}(\boldsymbol{x})\,,\,\forall\,(\boldsymbol{\zeta}^+,\boldsymbol{\zeta}^-)\in\mathcal{U}_{\pm}(\Gamma),\,\forall\,i,\,\forall\,k\,, \end{array}$

where $\boldsymbol{y}: \mathbb{R}^{n_{\zeta}} \times \mathbb{R}^{n_{\zeta}} \to \mathbb{R}^{n_{y}}$.

For each of the three cases, we will demonstrate that there exists a linear transformation of $y(\cdot)$ that can be used to distribute the term $\bar{\gamma} + \gamma^T (\zeta^+ - \zeta^-)$ in the constraints while preserving their respective structure. This then allows us to exploit Corollary 1 in Ardestani-Jaafari and Delage (2016) to reach our conclusion.

Condition i: Let us start by characterizing for any fixed $x \in X$, the optimal value of the adversarial problem as $h_1(x)$, namely:

$$h_1(\boldsymbol{x}) := \min_{(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma)} h(\boldsymbol{x}, \boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-) - ar{\gamma} - oldsymbol{\gamma}^T(oldsymbol{\zeta}^+ - oldsymbol{\zeta}^-)$$

and by $h_2(x)$ the lower bound on this value obtained using affine decision rules:

$$h_{2}(\boldsymbol{x}) := \max_{\bar{\boldsymbol{y}}, \{\boldsymbol{y}_{i}^{+}, \boldsymbol{y}_{i}^{-}\}_{i=1}^{n_{y}}} \min_{(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_{y}} (\bar{\boldsymbol{y}}_{i} + \boldsymbol{y}_{i}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{y}_{i}^{-T} \boldsymbol{\zeta}^{-}) - \bar{\boldsymbol{\gamma}} - \boldsymbol{\gamma}^{T} (\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-})$$
(1.51a)

s.t.
$$\bar{y}_i + y_i^{+T} \boldsymbol{\zeta}^+ + y_i^{-T} \boldsymbol{\zeta}^- \leq \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^T (\boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-) + \beta_{ik}(\boldsymbol{x}), \quad \begin{cases} \forall (\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma) \\ \forall i, \forall k. \end{cases}$$

$$(1.51b)$$

We will show that $h_2(\boldsymbol{x})$ is actually equal to $h_1(\boldsymbol{x})$. In particular, by replacing $\bar{z}_1 := \bar{y}_1 - \bar{\gamma}$, $\boldsymbol{z}_1^+ := \boldsymbol{y}_1^+ - \boldsymbol{\gamma}$, $\boldsymbol{z}_1^- := \boldsymbol{y}_1^- + \boldsymbol{\gamma}$, while $\bar{z}_i := \bar{y}_i$, $\boldsymbol{z}_i^+ := \boldsymbol{y}_i^+$, and $\boldsymbol{z}_i^- := \boldsymbol{y}_i^-$ for all $i \ge 2$, we then get that:

$$h_{2}(\boldsymbol{x}) := \max_{\boldsymbol{\bar{z}}, \{\boldsymbol{z}_{i}^{+}, \boldsymbol{z}_{i}^{-}\}_{i=1}^{n_{y}}} \min_{(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_{y}} (\bar{\boldsymbol{z}}_{i} + \boldsymbol{z}_{i}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{z}_{i}^{-T} \boldsymbol{\zeta}^{-})$$

s.t.
$$\bar{\boldsymbol{z}}_{1} + \boldsymbol{z}_{1}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{z}_{1}^{-T} \boldsymbol{\zeta}^{-} \leq (\boldsymbol{\alpha}_{1k}(\boldsymbol{x}) - \boldsymbol{\gamma})^{T} (\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}) + (\beta_{1k}(\boldsymbol{x}) - \bar{\boldsymbol{\gamma}}), \forall (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma), \forall k$$

$$\bar{\boldsymbol{z}}_{i} + \boldsymbol{z}_{i}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{z}_{i}^{-T} \boldsymbol{\zeta}^{-} \leq \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^{T} (\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}) + \beta_{ik}(\boldsymbol{x}), \forall (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma), \forall i \geq 2, \forall k.$$

One can easily recognize that this form is equivalent to the lower bound obtained when applying affine decision rules to approximate the worst-case value of a sum of piecewise linear concave functions. Following Corollary 1 in Ardestani-Jaafari and Delage (2016), since $\Gamma = 1$, we can conclude that

$$\begin{split} h_2(\boldsymbol{x}) &:= \max_{\boldsymbol{z}(\cdot)} & \min_{(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_y} z_i(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \\ & \text{s.t.} & z_1(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \leq (\boldsymbol{\alpha}_{1k}(\boldsymbol{x}) - \boldsymbol{\gamma})^T(\boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-) + (\beta_{1k}(\boldsymbol{x}) - \bar{\boldsymbol{\gamma}}), \, \forall \, (\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma), \, \forall \, k \\ & z_i(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \leq \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^T(\boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-) + \beta_{ik}(\boldsymbol{x}), \, \forall \, (\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma), \, \forall \, i \geq 2, \, \forall \, k, \end{split}$$

which once more with a replacement of variables gives us:

$$h_{2}(\boldsymbol{x}) = \max_{\boldsymbol{y}(\cdot)} \qquad \min_{(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_{y}} y_{i}(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) - \bar{\gamma} - \boldsymbol{\gamma}^{T}(\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-})$$

s.t. $z_{i}(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \leq \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^{T}(\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}) + \beta_{ik}(\boldsymbol{x}) \,\forall \, (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma), \,\forall i, \,\forall k.$

Hence, we have that $h_2(\boldsymbol{x}) = h_1(\boldsymbol{x})$.

Condition iii: The proof for Condition iii is fairly similar except that we exploit a different affine transformation for passing from y to z. In particular, now we can exploit the fact that the objective function in (1.51) can be equivalently written as:

$$\max_{\bar{\boldsymbol{y}},\{\boldsymbol{y}_i^+,\boldsymbol{y}_i^-\}_{i=1}^{n_y}} \min_{(\boldsymbol{\zeta}^+,\boldsymbol{\zeta}^-)\in\mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_y} (\bar{y}_i-\bar{\gamma}/n_y) + (\boldsymbol{y}_i^+-\gamma_i\boldsymbol{e}_i)^T \boldsymbol{\zeta}^+ + (\boldsymbol{y}_i^-+\gamma_i\boldsymbol{e}_i)^T \boldsymbol{\zeta}^- \,.$$

We can now replace $\bar{z} := \bar{y} - \bar{\gamma}/n_y$ and each $z_i^+ := y_i^+ - \gamma_i e_i$ and $z_i^- := y_i^- + \gamma_i e_i$ to get:

$$h_{2}(\boldsymbol{x}) = \max_{\bar{\boldsymbol{x}}, \{\boldsymbol{z}_{i}^{+}, \boldsymbol{z}_{i}^{-}\}_{i=1}^{n_{y}}} \min_{(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{p} (\bar{\boldsymbol{z}}_{i} + \boldsymbol{z}_{i}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{z}_{i}^{-T} \boldsymbol{\zeta}^{-})$$

s.t.
$$\forall (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)$$

$$\bar{z}_i + \boldsymbol{z}_i^{+T} \boldsymbol{\zeta}^+ + \boldsymbol{z}_i^{-T} \boldsymbol{\zeta}^- \leq (\boldsymbol{\alpha}_{ik}(\boldsymbol{x}) - \gamma_i \boldsymbol{e}_i)^T (\boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-) + \beta_{ik}(\boldsymbol{x}) - \bar{\gamma}/n_y, \quad \begin{array}{l} \forall (\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma) \\ \forall i, \forall k. \end{array}$$

One can again recognize that this form is equivalent to the lower bound obtained when applying affine decision rules to approximate the worst-case value of $h'(\boldsymbol{x}, \boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-)$, which is defined as the sum of piecewise linear concave functions using $\boldsymbol{\alpha}'_{ik}(\boldsymbol{x}) := \boldsymbol{\alpha}_{ik}(\boldsymbol{x}) - \gamma_i \boldsymbol{e}_i$ and $\beta'_{ik}(\boldsymbol{x}) := \beta_{ik}(\boldsymbol{x}) - \overline{\gamma}/n_y$. Following Corollary 1 in Ardestani-Jaafari and Delage (2016), we can conclude that

$$h_2(\boldsymbol{x}) = \min_{(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma)} h'(\boldsymbol{x}, \boldsymbol{\zeta}^+ - \boldsymbol{\zeta}^-)$$

since by Condition iii we have that:

$$\boldsymbol{\alpha}_{ik}'(\boldsymbol{x}) = \boldsymbol{\alpha}_{ik}(\boldsymbol{x}) - \gamma_i \boldsymbol{e}_i = \bar{\alpha}_{ik}(\boldsymbol{x})\boldsymbol{e}_i - \gamma_i \boldsymbol{e}_i = (\bar{\alpha}_{ik}(\boldsymbol{x}) - \gamma_i)\boldsymbol{e}_i,$$

hence Condition 3 in Ardestani-Jaafari and Delage (2016) is satisfied. We can therefore conclude that

$$h_{2}(\boldsymbol{x}) = \max_{\boldsymbol{z}(\cdot)} \qquad \min_{(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_{y}} z_{i}(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-})$$

s.t. $z_{i}(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \leq (\boldsymbol{\alpha}_{ik}(\boldsymbol{x}) - \gamma_{i}\boldsymbol{e}_{i})^{T}(\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}) + \beta_{ik}(\boldsymbol{x}) - \bar{\gamma}/n_{y}, \quad \begin{cases} \forall (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma) \\ \forall i, \forall k, \end{cases}$

which once more with a replacement of variable gives us:

$$h_{2}(\boldsymbol{x}) = \max_{\boldsymbol{y}(\cdot)} \qquad \min_{(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_{y}} y_{i}(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) - \bar{\gamma} - \boldsymbol{\gamma}^{T}(\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-})$$

s.t. $z_{i}(\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \leq \boldsymbol{\alpha}_{ik}(\boldsymbol{x})^{T}(\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}) + \beta_{ik}(\boldsymbol{x}) \,\forall \, (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma), \,\forall \, i, \,\forall \, k.$

Hence, we have that $h_2(\boldsymbol{x}) = h_1(\boldsymbol{x})$.

Condition ii: The proof for Condition ii is again entirely analogous with a new affine transformation for passing from y to z. In particular, we first assume for simplicity of exposition that $\hat{\alpha}_{\ell} \neq 0$ for all $\ell = 1, ..., n_{\zeta}$ and that $n_y = n_{\zeta} + 1$. We then exploit the fact that:

$$\boldsymbol{\gamma} = \sum_{\ell=1}^{n_{\zeta}} \gamma_{\ell} \boldsymbol{e}_{\ell} = \sum_{\ell=1}^{n_{\zeta}} \hat{\alpha}_{\ell} \boldsymbol{e}_{\ell} \left(\sum_{i=\ell}^{n_{\zeta}} \frac{\gamma_{i}}{\hat{\alpha}_{i}} - \sum_{i=\ell+1}^{n_{\zeta}} \frac{\gamma_{i}}{\hat{\alpha}_{i}} \right) = \sum_{i=1}^{n_{\zeta}-1} \left(\frac{\gamma_{i}}{\hat{\alpha}_{i}} - \frac{\gamma_{i+1}}{\hat{\alpha}_{i+1}} \right) \sum_{\ell \leq i} \hat{\alpha}_{\ell} \boldsymbol{e}_{\ell} + \frac{\gamma_{n_{\zeta}}}{\hat{\alpha}_{\zeta}} \sum_{\ell=1}^{n_{\zeta}} \hat{\alpha}_{\ell} \boldsymbol{e}_{\ell}$$

$$= \sum_{i=1}^{n_{\zeta}+1} \bar{\alpha}'_{i} \left(\sum_{\ell < i} \hat{\alpha}_{\ell} \boldsymbol{e}_{\ell} \right),$$

where

$$\bar{\alpha}'_{i} := \begin{cases} 0 & \text{if } i = 1\\ \frac{\gamma_{i-1}}{\hat{\alpha}_{i-1}} - \frac{\gamma_{i}}{\hat{\alpha}_{i}} & \text{if } i \in \{2, \dots, n_{\zeta}\}\\ \frac{\gamma_{n_{\zeta}}}{\hat{\alpha}_{\zeta}} & \text{if } i = n_{\zeta} + 1 \end{cases}$$

We, therefore, have that the objective function in (1.51) can be reformulated as

$$\max_{\bar{\boldsymbol{y}}, \{\boldsymbol{y}_i^+, \boldsymbol{y}_i^-\}_{i=1}^{n_y}} \min_{(\boldsymbol{\zeta}^+, \boldsymbol{\zeta}^-) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_y} (\bar{y}_i - \bar{\gamma}/n_y) + \left(\boldsymbol{y}_i^+ - \bar{\alpha}_i' \left(\sum_{\ell < i} \hat{\alpha}_\ell \boldsymbol{e}_\ell\right)\right)^T \boldsymbol{\zeta}^+ + \left(\boldsymbol{y}_i^- + \bar{\alpha}_i' \left(\sum_{\ell < i} \hat{\alpha}_\ell \boldsymbol{e}_\ell\right)\right)^T \boldsymbol{\zeta}^-.$$

By replacing $\bar{z}_i := \bar{y}_i - \bar{\gamma}/n_y$ as before, while replacing $\boldsymbol{z}_i^+ := \boldsymbol{y}_i^+ - \bar{\alpha}_i'(\sum_{\ell < i} \hat{\alpha}_\ell \boldsymbol{e}_\ell)$ and $\boldsymbol{z}_i^- := \boldsymbol{y}_i^- + \bar{\alpha}_i'(\sum_{\ell < i} \hat{\alpha}_\ell \boldsymbol{e}_\ell))^T \boldsymbol{\zeta}^-$, we obtain:

$$h_{2}(\boldsymbol{x}) = \max_{\boldsymbol{\bar{z}}, \{\boldsymbol{z}_{i}^{+}, \boldsymbol{z}_{i}^{-}\}_{i=1}^{n_{y}}} (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma)} \sum_{i=1}^{n_{y}} \left(\bar{z}_{i} + \boldsymbol{z}_{i}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{z}_{i}^{-T} \boldsymbol{\zeta}^{-} \right)$$

s.t.
$$\bar{z}_{i} + \boldsymbol{z}_{i}^{+T} \boldsymbol{\zeta}^{+} + \boldsymbol{z}_{i}^{-T} \boldsymbol{\zeta}^{-} \leq \left(\bar{\alpha}_{ik}(\boldsymbol{x}) - \bar{\alpha}_{i}' \right) \left(\sum_{\ell < i} \hat{\alpha}_{\ell} \boldsymbol{e}_{\ell} \right)^{T} (\boldsymbol{\zeta}^{+} - \boldsymbol{\zeta}^{-}) + \beta_{ik}(\boldsymbol{x}) - \bar{\gamma}/n_{y}, \quad \begin{cases} \forall (\boldsymbol{\zeta}^{+}, \boldsymbol{\zeta}^{-}) \in \mathcal{U}_{\pm}(\Gamma) \\ \forall i, \forall k. \end{cases}$$

Hence, once again Corollary 1 of Ardestani-Jaafari and Delage (2016) applies and allows us to complete the proof using exactly the same steps as for conditions i and iii.

1.8.14 TSLRO Reformulations for WCRRM in Cost Minimization Problems

Given a non-negative optimal second-stage cost function $f(x, \zeta)$, which depends on both the decision and the realization of some uncertain vector of parameters ζ , following the formulation presented in Mausser and Laguna (1999b), one measures the relative regret experienced once ζ is revealed as the ratio of the difference between the lowest cost achievable $\min_{x' \in \mathcal{X}} f(x', \zeta)$ and the cost $f(x, \zeta)$ achieved by the decision x that was implemented, over the lowest cost achievable. The worst-case relative regret minimization (WCRRM) problem thus takes the form:

$$\underset{\boldsymbol{x}\in\mathcal{X}}{\operatorname{minimize}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \frac{f(\boldsymbol{x},\boldsymbol{\zeta}) - \inf_{\boldsymbol{x}'\in\mathcal{X}} f(\boldsymbol{x}',\boldsymbol{\zeta})}{\inf_{\boldsymbol{x}'\in\mathcal{X}} f(\boldsymbol{x}',\boldsymbol{\zeta})} \right\},$$
(1.52)

where, when $\inf_{x' \in \mathcal{X}} f(x', \zeta) = 0$, we should interpret the relative regret as being either 0 if $f(x, \zeta) = 0$ or infinite otherwise. Equivalently, in terms of $h(x, \zeta) := -f(x, \zeta)$, we will define the WCRRM problem has:

(WCRRM) minimize
$$\sup_{\boldsymbol{x}\in\mathcal{X}} \sup_{\boldsymbol{\zeta}\in\mathcal{U}} \left\{ \frac{h(\boldsymbol{x},\boldsymbol{\zeta}) - \sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta})}{\sup_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta})} \right\}.$$
 (1.53)

As mentioned above, we make the following assumption about the profit function in this two-stage problem.

Assumption 1.8.1 The cost function $h(x, \zeta) \leq 0$ for all $x \in \mathcal{X}$ and all $\zeta \in \mathcal{U}$. This implies that Assumptions 1.3.3 and 1.4.1 are satisfied and that the optimal value of Problem (1.53) is greater or equal to zero.

In what follows we demonstrate how the WCRRM problem can be reformulated as a TSLRO when the cost function $f(x, \zeta)$ (a.k.a. $-h(\mathcal{X}, \zeta)$) captures the cost of a second-stage linear decision model with either right-hand side or objective uncertainty.

1.8.14.1 The Case of Right-Hand Side Uncertainty

We consider the case where $h(x, \zeta)$ takes the form presented in Problem (1.3) and where uncertainty is limited to the right-hand side as defined in Definition 1.3.1.

Proposition 1.8.1 *Given that Assumptions 1.3.1 and 1.8.1 are satisfied, the cost-based WCRRM problem with right-hand side uncertainty is equivalent to the following TSLRO problem:*

$$\underset{\boldsymbol{x}' \in \mathcal{X}', \boldsymbol{y}'(\cdot)}{\operatorname{maximize}} \quad \inf_{\boldsymbol{\zeta}' \in \mathcal{U}'} \boldsymbol{c}'^T \boldsymbol{x}'$$
(1.54a)

subject to
$$A'\boldsymbol{x}' + B'\boldsymbol{y}'(\boldsymbol{\zeta}') \leq \Psi'(\boldsymbol{x}')\boldsymbol{\zeta}' + \psi', \,\forall\,\boldsymbol{\zeta}' \in \mathcal{U}',$$
 (1.54b)

where $x' \in \mathbb{R}^{n_x+1}$, $\zeta' \in \mathbb{R}^{n_\zeta+n_x+n_y}$, $y' : \mathbb{R}^{n_\zeta+n_x+n_y} \to \mathbb{R}^{n_y}$, $c' = \begin{bmatrix} -1 & \mathbf{0}^T \end{bmatrix}^T$, while $\mathcal{X}' := \{[t \ \mathbf{x}^T]^T \in \mathbb{R}^{n_x+1} \mid \mathbf{x} \in \mathcal{X}, t \ge 0\}$, \mathcal{U}' is defined as in equation (1.14), and

$$A' := \begin{bmatrix} 0 & -\boldsymbol{c}^T \\ 0 & A \end{bmatrix}, \qquad \qquad B' := \begin{bmatrix} -\boldsymbol{d}^T \\ B \end{bmatrix},$$
$$\Psi'(\boldsymbol{x}') := \begin{bmatrix} \boldsymbol{0}^T & -\boldsymbol{c}^T & -\boldsymbol{d}^T \\ \Psi & 0 & 0 \end{bmatrix} + \begin{bmatrix} \boldsymbol{0}^T & -\boldsymbol{c}^T & -\boldsymbol{d}^T \\ 0 & 0 & 0 \end{bmatrix} x'_1, \qquad \psi' := \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{\psi} \end{bmatrix}.$$

In particular, a solution for the WCRRM takes the form of $x^* := x_{2:n_x+1}^{**}$ and achieves a worstcase relative regret of x_1^{\prime} . Furthermore, this TSLRO reformulation necessarily satisfies Assumption 1.3.1 while it only satisfies Assumption 1.3.2 if all $x \in \mathcal{X}$ achieve a worst-case regret of zero.

We first employ an epigraph form for Problem (1.53) as follows:

$$\min_{\boldsymbol{x}\in\mathcal{X},t} \quad t \tag{1.55a}$$

subject to
$$\sup_{\boldsymbol{\zeta} \in \mathcal{U}} \left\{ \frac{h(\boldsymbol{x}, \boldsymbol{\zeta}) - \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})}{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})} \right\} \le t$$
(1.55b)

$$0 \le t \,, \tag{1.55c}$$

where we impose that $0 \le t$ since Assumption 1.8.1 ensures that the optimal value of the WCRRM problem is greater or equal to zero. One can then manipulate constraint (1.55b) to show that it is equivalent to

$$\frac{h(\boldsymbol{x},\boldsymbol{\zeta}) - \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta})}{\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta})} \le t , \, \forall \boldsymbol{\zeta} \in \mathcal{U},$$

hence to

$$h(\boldsymbol{x},\boldsymbol{\zeta}) - \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta}) \ge t(\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\zeta})) , \, \forall \boldsymbol{\zeta} \in \mathcal{U} \,,$$

since, for a fixed ζ , either $\sup_{x' \in \mathcal{X}} h(x, \zeta) < 0$ or otherwise the new constraint becomes equivalent to $h(x, \zeta) = 0$, which captures exactly the fact that the regret is zero under this ζ scenario if $h(x, \zeta) = 0$ and otherwise infinite. Finally, we obtain the constraint:

$$(t+1) \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}) \le 0, \, \forall \boldsymbol{\zeta} \in \mathcal{U}.$$
(1.56)

By substituting Problem (1.3) in this constraint, we obtain the following reformulations

(1.55b)
$$\equiv (t+1) \sup_{\boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x}' + \boldsymbol{d}^T \boldsymbol{y}' - \sup_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x}, \boldsymbol{\zeta})} \boldsymbol{c}^T \boldsymbol{x} + \boldsymbol{d}^T \boldsymbol{y} \leq 0, \forall \boldsymbol{\zeta} \in \mathcal{U}$$

$$\equiv \min_{\boldsymbol{y} \in \mathcal{Y}(\boldsymbol{x},\boldsymbol{\zeta})} - \boldsymbol{c}^T \boldsymbol{x} - \boldsymbol{d}^T \boldsymbol{y} + (1+t)\boldsymbol{c}^T \boldsymbol{x}' + (1+t)\boldsymbol{d}^T \boldsymbol{y}' \leq \boldsymbol{0} , \forall \boldsymbol{\zeta} \in \mathcal{U}, x' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}',\boldsymbol{\zeta}) \in \mathcal{Y}(\boldsymbol{x}',\boldsymbol{\zeta})$$

Hence, the WCRRM problem reduces to:

$$\min_{\boldsymbol{x} \in \mathcal{X}, t \geq 0} \quad \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{x}' \in \mathcal{X}, \boldsymbol{y}' \in \mathcal{Y}(\boldsymbol{x}', \boldsymbol{\zeta})} h'(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{x}', \boldsymbol{y}')$$

where

$$\begin{aligned} h'(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{x}', \boldsymbol{y}') &:= & \inf_{\boldsymbol{y}} & t \\ & \text{s.t.} & -\boldsymbol{c}^T \boldsymbol{x} - \boldsymbol{d}^T \boldsymbol{y} \leq -(t+1) \boldsymbol{c}^T \boldsymbol{x}' - (t+1) \boldsymbol{d}^T \boldsymbol{y}' \\ & & A \boldsymbol{x} + B \boldsymbol{y} \leq \Psi \boldsymbol{\zeta} + \boldsymbol{\psi} \,. \end{aligned}$$

This problem can be rewritten in the form presented in equation (1.54).

Note that the arguments to support the conditions under which Assumptions 1.3.1 and 1.3.2 are satisfied are exactly the same as in the proof of Proposition 1.5.1.

1.8.14.2 The Case of Objective Uncertainty

We consider the case where $h(x, \zeta)$ takes the form presented in Problem (1.3).

Proposition 1.8.2 *Given that Assumptions 1.3.1 and 1.8.1 are satisfied, the WCRRM problem with objective uncertainty is equivalent to the following TSLRO problem:*

subject to
$$A' \boldsymbol{x}' + B' \boldsymbol{y}'(\boldsymbol{\zeta}') \leq \Psi'(\boldsymbol{x}') \boldsymbol{\zeta}' + \boldsymbol{\psi}'$$
 (1.57b)

$$x' \in \mathcal{X}'$$
, (1.57c)

where $\mathbf{x}' \in \mathbb{R}^{n_x+1}$, $\mathbf{y}' : \mathbb{R}^{n_\zeta+m} \to \mathbb{R}^{m+r}$, while $\mathcal{X}' := \{ [u \ \mathbf{z}^T]^T \in \mathbb{R}^{n_x+1} | W\mathbf{z} \ge \mathbf{v}u, -1 \le u \le 0 \}$, and \mathcal{U}' is defined as in equation (1.17). Moreover, we have that $\mathbf{c}' := [-1 \ \mathbf{0}^T]^T$, while

$$A' := \begin{bmatrix} 0 & c^T \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad B' := \begin{bmatrix} \psi^T & v^T \\ A^T & W^T \\ -A^T & -W^T \\ B^T & 0 \\ -B^T & 0 \\ -B^T & 0 \\ -I & 0 \\ 0 & -I \end{bmatrix},$$

$$\Psi'(\boldsymbol{x}') := \begin{bmatrix} \boldsymbol{0}^T & -\boldsymbol{\psi}^T \boldsymbol{x}_1' + \boldsymbol{x}_{2:n_x+1}'^T A^T \\ 0 & 0 \\ 0 & 0 \\ D & 0 \\ -D & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad \boldsymbol{\psi}' := \begin{bmatrix} 0 \\ \boldsymbol{c} \\ -\boldsymbol{c} \\ \boldsymbol{d} \\ -\boldsymbol{d} \\ \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix}$$

In particular, a solution for the WCRRM takes the form of $x^* := x_{2:n_x+1}^{*}/x_1^{*}$ and achieves a worst-case relative regret of $-1 - 1/x_1^{**}$ if $x_1^{**} < 0$ while the best worst-case relative regret should be considered infinite if $x_1^{**} = 0$. Furthermore, this TSLRO reformulation necessarily satisfies Assumption 1.3.1 while it only satisfies Assumption 1.3.2 if all $x \in X$ achieve a worst-case regret of zero.

The first steps of this proof are exactly as in the proof of Proposition 1.8.1 up to equation (1.56). The next steps are then exactly analogous to the steps followed in the proof of Proposition 1.5.2, which we repeat for completeness. Since we are now dealing with objective uncertainty, we substitute $h(\boldsymbol{x}, \boldsymbol{\zeta})$ and $\sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta})$ using their respective dual form (see equations (1.27) and (1.31) respectively). Strong duality applies since Assumption 1.8.1 implies that Assumptions 1.3.3 and 1.4.1 are satisfied, which results to the following reformulation:

$$(1.55b) \equiv (t+1) \sup_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}) \leq 0, \ \forall \boldsymbol{\zeta} \in \mathcal{U}$$

$$\equiv (t+1) \inf_{(\boldsymbol{\lambda}, \boldsymbol{\gamma}) \in \Upsilon_1(\boldsymbol{\zeta})} \psi^T \boldsymbol{\lambda} + \boldsymbol{v}^T \boldsymbol{\gamma} - \inf_{\boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} \{ \boldsymbol{c}^T \boldsymbol{x} + (\boldsymbol{\psi} - A \boldsymbol{x})^T \boldsymbol{\rho} \} \leq 0, \ \forall \boldsymbol{\zeta} \in \mathcal{U}$$

$$(1.59)$$

$$\equiv \inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})} (1+t)\boldsymbol{\psi}^{T}\boldsymbol{\lambda} + (1+t)\boldsymbol{v}^{T}\boldsymbol{\gamma} - \boldsymbol{c}^{T}\boldsymbol{x} - (\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho} \leq 0, \forall \boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_{2}(\boldsymbol{\zeta})$$
(1.60)

$$\equiv \inf_{(\boldsymbol{\lambda},\boldsymbol{\gamma})\in\Upsilon_{1}(\boldsymbol{\zeta})} \psi^{T}\boldsymbol{\lambda} + \boldsymbol{v}^{T}\boldsymbol{\gamma} - \frac{1}{1+t}\boldsymbol{c}^{T}\boldsymbol{x} - \frac{1}{1+t}(\boldsymbol{\psi} - A\boldsymbol{x})^{T}\boldsymbol{\rho} \leq 0, \forall \boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_{2}(\boldsymbol{\zeta}),$$
(1.61)

where $\Upsilon_1(\boldsymbol{\zeta})$ and $\Upsilon_2(\boldsymbol{\zeta})$ are as defined in the proof of Proposition 1.4.2. Hence the WCRRM

problem reduces to:

$$\min_{\boldsymbol{x} \in \mathcal{X}, t \geq 0} \quad \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} h''(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{\rho}) \,,$$

where

$$\begin{split} h''(\boldsymbol{x}, t, \boldsymbol{\zeta}, \boldsymbol{\rho}) &:= & \inf_{\boldsymbol{\lambda}, \boldsymbol{\gamma}} & t \\ & \text{s.t.} & \boldsymbol{\psi}^T \boldsymbol{\lambda} + \boldsymbol{v}^T \boldsymbol{\gamma} - \frac{1}{1+t} \boldsymbol{c}^T \boldsymbol{x} - \frac{1}{1+t} (\boldsymbol{\psi} - A \boldsymbol{x})^T \boldsymbol{\rho} \leq 0 \\ & A^T \boldsymbol{\lambda} + W^T \boldsymbol{\gamma} = \boldsymbol{c} \\ & B^T \boldsymbol{\lambda} = \boldsymbol{d}(\boldsymbol{\zeta}) \\ & \boldsymbol{\lambda} \geq 0, \ \boldsymbol{\gamma} \geq 0 \,. \end{split}$$

Using a simple replacement of variables $u := -\frac{1}{1+t}$ and $z := -\frac{1}{1+t}x$ and applying a monotone transformation of the objective function $t \to -\frac{1}{1+t}$, we obtain that the WCRRM is equivalently represented as

$$\min_{\substack{-1 \leq u < 0, \, \boldsymbol{z} : W \boldsymbol{z} \geq \boldsymbol{v} u}} \sup_{\boldsymbol{\zeta} \in \mathcal{U}, \boldsymbol{\rho} \in \Upsilon_2(\boldsymbol{\zeta})} h''(\boldsymbol{z}, u, \boldsymbol{\zeta}, \boldsymbol{\rho}) \,,$$

where

$$\begin{split} h''(\boldsymbol{z}, u, \boldsymbol{\zeta}, \boldsymbol{\rho}) &:= & \inf_{\boldsymbol{\lambda}, \boldsymbol{\gamma}} & u \\ \text{s.t.} & \boldsymbol{\psi}^T \boldsymbol{\lambda} + \boldsymbol{v}^T \boldsymbol{\gamma} + \boldsymbol{c}^T \boldsymbol{z} + (u \boldsymbol{\psi} - A \boldsymbol{z})^T \boldsymbol{\rho} \leq 0 \\ & A^T \boldsymbol{\lambda} + W^T \boldsymbol{\gamma} = \boldsymbol{c} \\ & B^T \boldsymbol{\lambda} = \boldsymbol{d}(\boldsymbol{\zeta}) \\ & \boldsymbol{\lambda} \geq 0, \ \boldsymbol{\gamma} \geq 0 \,. \end{split}$$

This problem can be rewritten in the form presented in equation (1.57) as long as when the optimal value of the TSLRO is 0 one concludes that best worst-case relative regret is infinite.

Note that the arguments to support the conditions under which Assumptions 1.3.1 and 1.3.2 are satisfied are similar as in the proof of Proposition 1.5.2.

Endnotes

1. Note that if WCARM is unbounded it is necessarily because such an $x \in \mathcal{X}$ exists since for any fixed x if the profit reachable under all $\zeta \in \mathcal{U}$ is finite then the regret is necessarily non-negative.

2. Note that the budgeted uncertainty set in this work follows the representation proposed in Ardestani-Jaafari and Delage (2016), i.e. with $\sum_i \zeta_i^+ + \zeta_i^- = \Gamma$ instead of $\sum_i \zeta_i^+ + \zeta_i^- \leq \Gamma$, in order for their Proposition 6 to be applicable.

3. In order to assess the average performance in terms of worst-case profit (*WC*), worstcase absolute regret (*WCARM*), and worst-case relative regret (*WCRRM*) of the different decision sets, we computed the average of $\frac{WC^*-WC}{WC^*} \times 100$, $\frac{WCARM-WCARM^*}{WCARM^*} \times 100$, and *WCRRM* – *WCRRM*^{*} measures on the 120 instances, where *WC*^{*}, *WCARM*^{*}, and *WCRRM*^{*} represent the optimal values of WC, WCARM, and WCRRM problems, respectively.

References

- Aissi, H., Bazgan, C., and Vanderpooten, D. (2009). Min-max and min-max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research*, 197(2):427–438.
- Ardestani-Jaafari, A. and Delage, E. (2016). Robust optimization of sums of piecewise linear functions with application to inventory problems. *Operations Research*, 64(2):474– 494.
- Ardestani-Jaafari, A. and Delage, E. (2020). Linearized robust counterparts of two-stage robust optimization problems with applications in operations management. *INFORMS Journal on Computing*.
- Assavapokee, T., Realff, M., and Ammons, J. (2008a). Min-max regret robust optimization approach on interval data uncertainty. *Journal of Optimization Theory and Applications*, 137(2):297–316.

- Assavapokee, T., Realff, M., Ammons, J., and Hong, I. (2008b). Scenario relaxation algorithm for finite scenario-based min-max regret and min-max relative regret robust optimization. *Computers & Operations Research*, 35(6):2093–2102.
- Averbakh, I. (2004). Minmax regret linear resource allocation problems. *Operations Research Letters*, 32(2):174–180.
- Averbakh, I. and Lebedev, V. (2005). On the complexity of minmax regret linear programming. *European Journal of Operational Research*, 160(1):227–231.
- Ayoub, J. and Poss, M. (2016). Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Computational Management Science*, 13(2):219–239.
- Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., and Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376.
- Bertsimas, D. and de Ruiter, F. (2016). Duality in two-stage adaptive linear optimization: Faster computation and stronger bounds. *INFORMS Journal on Computing*, 28(3):500–511.
- Bertsimas, D., Doan, X., Natarajan, K., and Teo, C. (2010a). Models for minimax stochastic linear optimization problems with risk aversion. *Mathematics of Operations Research*, 35(3):580–602.
- Bertsimas, D. and Dunning, I. (2020). Relative robust and adaptive optimization. *IN*-FORMS Journal on Computing, 32(2):408–427.
- Bertsimas, D. and Goyal, V. (2012). On the power and limitations of affine policies in two-stage adaptive optimization. *Mathematical Programming*, 134(2):491–531.
- Bertsimas, D., Iancu, D., and Parrilo, P. (2010b). Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research*, 35(2):363–394.

- Bertsimas, D., Iancu, D., and Parrilo, P. (2011). A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Transactions on Automatic Control*, 56(12):2809– 2824.
- Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations Research*, 52(1):35–53.
- Bleichrodt, H., Cillo, A., and Diecidue, E. (2010). A quantitative measurement of regret theory. *Management Science*, 56(1):161–175.
- Borodin, A. and El-Yaniv, R. (2005). *Online computation and competitive analysis*. Cambridge university press.
- Caldentey, R., Liu, Y., and Lobel, I. (2017). Intertemporal pricing under minimax regret. *Operations Research*, 65(1):104–129.
- Chen, B., Wang, J., Wang, L., He, Y., and Wang, Z. (2014). Robust optimization for transmission expansion planning: Minimax cost vs. minimax regret. *IEEE Transactions on Power Systems*, 29(6):3069–3077.
- Chen, X., Sim, M., Sun, P., and Zhang, J. (2008). A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56(2):344–357.
- Chen, X. and Zhang, Y. (2009). Uncertain linear programs: Extended affinely adjustable robust counterparts. *Operations Research*, 57(6):1469–1482.
- Conde, E. (2005). On the complexity of the continuous unbounded knapsack problem with uncertain coefficients. *Operations Research Letters*, 33(5):481–485.
- Delage, E. and Iancu, D. (2015). Robust multistage decision making. INFORMS TutORials in Operations Research, pages 20–46.
- Dunning, I., Huchette, J., and Lubin, M. (2017). JuMP: A modeling language for mathematical optimization. SIAM Review, 59(2):295–320.
- Gabrel, V. and Murat, C. (2010). Robustness and duality in linear programming. *Journal* of the Operational Research Society, 61(8):1288–1296.

- Goh, J. and Sim, M. (2011). Robust optimization made easy with ROME. *Operations Research*, 59(4):973–985.
- Guslitser, E. (2002). Uncertainty-immunized solutions in linear programming. Master's thesis, Technion, Israeli Institute of Technology.
- Hadjiyiannis, M. J., Goulart, P. J., and Kuhn, D. (2011). A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. In 2011 50th IEEE Conference on Decision and Control and European Control Conference, pages 7386–7391.
- Hanasusanto, G. and Kuhn, D. (2018). Conic programming reformulations of two-stage distributionally robust linear programs over wasserstein balls. *Operations Research*, 66(3):849–869.
- Iancu, D. A. and Trichakis, N. (2014). Pareto efficiency in robust optimization. *Management Science*, 60(1):130–147.
- Inuiguchi, M., Higashitani, H., and Tanino, T. (1999). On computation methods of the minimax regret solution for linear programming problems with uncertain objective function coefficients. In *IEEE SMC'99 Conference Proceedings*. 1999 *IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.99CH37028)*, volume 3, pages 979– 984.
- Inuiguchi, M. and Kume, Y. (1994). Minimax regret in linear programming problems with an interval objective function. In Tzeng, G. H., Wang, H. F., Wen, U. P., and Yu, P. L., editors, *Multiple Criteria Decision Making*, pages 65–74, New York, NY. Springer New York.
- Inuiguchi, M. and Sakawa, M. (1995). Minimax regret solution to linear programming problems with an interval objective function. *European Journal of Operational Research*, 86(3):526–536.
- Inuiguchi, M. and Sakawa, M. (1996). Maximum regret analysis in linear programs with an interval objective function. *Proceedings of International Workshop on Soft Computing in Industry* '96, pages 308–317.

- Inuiguchi, M. and Sakawa, M. (1997a). An achievement rate approach to linear programming problems with an interval objective function. *The Journal of the Operational Research Society*, 48(1):25–33.
- Inuiguchi, M. and Sakawa, M. (1997b). Minimax regret solution algorithms for linear program with convex polyhedral objective coefficients. *Proceedings of Second European Workshop on Fuzzy Decision Analysis and Neural Networks for Management, Planning and Optimization, Dortmund, Germany*, pages 116–125.
- Inuiguchi, M. and Tanino, T. (2001). On computation methods for a minimax regret solution based on outer approximation and cutting hyperplanes. *International Journal of Fuzzy Systems*, 3(4):548–557.
- Jiang, R., Wang, J., Zhang, M., and Guan, Y. (2013). Two-stage minimax regret robust unit commitment. *IEEE Transactions on Power Systems*, 28(3):2271–2282.
- Kim, B. and Chung, B. (2017). Affinely adjustable robust model for multiperiod production planning under uncertainty. *IEEE Transactions on Engineering Management*, 64(4):505–514.
- Kouvelis, P. and Yu, G. (1996). *Robust Discrete Optimization and Its Applications*. Nonconvex Optimization and Its Applications. Springer US.
- Kuhn, D., Wiesemann, W., and Georghiou, A. (2011). Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209.
- Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824.
- Mausser, H. and Laguna, M. (1998). A new mixed integer formulation for the maximum regret problem. *International Transactions in Operational Research*, 5(5):389–403.
- Mausser, H. and Laguna, M. (1999a). A heuristic to minimax absolute regret for linear programs with interval objective function coefficients. *European Journal of Operational Research*, 117(1):157–174.

- Mausser, H. E. and Laguna, M. (1999b). Minimising the maximum relative regret for linear programmes with interval objective function coefficients. *The Journal of the Oper-ational Research Society*, 50(10):1063–1070.
- 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(2):1–23.
- Milnor, J. W. (1954). Games against nature. In Thrall, R., Coombs, C., and Davis, R., editors, *Decision Processes*. New York: Wiley.
- Minoux, M. (2009). On robust maximum flow with polyhedral uncertainty sets. *Optimization letters*, 3(3):367–376.
- Natarajan, K., Shi, D., and Toh, K.-C. (2014). A probabilistic model for minmax regret in combinatorial optimization. *Operations Research*, 62(1):160–181.
- Ng, T. (2013). Robust regret for uncertain linear programs with application to coproduction models. *European Journal of Operational Research*, 227(3):483–493.
- Ning, C. and You, F. (2018). Adaptive robust optimization with minimax regret criterion: Multiobjective optimization framework and computational algorithm for planning and scheduling under uncertainty. *Computers & Chemical Engineering*, 108:425–447.
- Perakis, G. and Roels, G. (2008). Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203.
- Pochet, Y. and Wolsey, L. (1988). Lot-size models with backlogging: Strong reformulations and cutting planes. *Mathematical Programming*, 40(1-3):317–335.
- Savage, L. (1951). The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67.
- Simchi-Levi, D., Trichakis, N., and Zhang, P. (2019). Designing response supply chain against bioattacks. *Operations Research*, 67(5):1246–1268.
- Stoye, J. (2011). Statistical decisions under ambiguity. Theory and decision, 70(2):129–148.

- Vairaktarakis, G. (2000). Robust multi-item newsboy models with a budget constraint. *International Journal of Production Economics*, 66(3):213–226.
- Xu, G. and Burer, S. (2018). A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *Computational Optimization and Applications*, 70(1):33–59.
- Yanikoglu, I., Gorissen, B., and den Hertog, D. (2018). A survey of adjustable robust optimization. *European Journal of Operational Research*.
- Yue, J., Chen, B., and Wang, M. (2006). Expected value of distribution information for the newsvendor problem. *Operations Research*, 54(6):1128–1136.
- Zeng, B. and Zhao, L. (2013). Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461.
- Zhang, M. (2011). Two-stage minimax regret robust uncapacitated lot-sizing problems with demand uncertainty. *Operations Research Letters*, 39(5):342–345.
- Zhen, J., den Hertog, D., and Sim, M. (2018). Adjustable robust optimization via fouriermotzkin elimination. *Operations Research*, 66(4):1086–1100.
- Zhu, Z., Zhang, J., and Ye, Y. (2013). Newsvendor optimization with limited distribution information. *Optimization Methods & Software*, 28(3):640–667.

Chapter 2

Risk-Averse Regret Minimization in Multistage Stochastic Programs

Abstract

Within the context of optimization under uncertainty, a well-known alternative to minimizing expected value or the worst-case scenario consists in minimizing regret. In a multistage stochastic programming setting with a discrete probability distribution, we explore the idea of risk-averse regret minimization, where the benchmark policy can only benefit from foreseeing Δ steps into the future. The Δ -regret model naturally interpolates between the popular ex-ante and ex-post regret models. We provide theoretical and numerical insights about this family of models under popular coherent risk measures and shed new light on the conservatism of the Δ -regret minimizing solutions.

2.1 Introduction

The regret minimization paradigm, introduced by Savage (1951), is claimed to provide less conservative solutions compared to the ones returned by optimizing with respect to the worst-case scenario (Perakis and Roels 2008, Aissi et al. 2009, Natarajan et al. 2014, Caldentey et al. 2017). Given a profit function $h(x, \zeta)$, which depends on the decision x and an uncertain vector of parameters ζ , the regret minimization approach aims at minimizing the difference between the achieved profit and the best profit that would have been made if the realization of ζ was known before making the decision. Namely, the so-called ex-post worst-case regret minimization problem takes the form of:

(EP-WCR) minimize
$$\max_{\boldsymbol{x} \in \mathcal{X}} \left\{ \max_{\omega \in \Omega} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}, \boldsymbol{\zeta}(\omega)) \right\}$$

where \mathcal{X} is the set of admissible actions, Ω denotes the outcome space, and x' captures the decision made with full information about ω , which we will refer to as the benchmark policy.

While most of the regret minimization literature focuses on worst-case scenario analysis, there has recently been a scarce but growing interest in formulations that account for more information about the underlying potential of realization of the different outcomes. A first common approach can be referred to as the ex-post risk-averse regret minimization problem:

(EP-RAR) minimize
$$\rho\left(\max_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\zeta}(\omega))-h(\boldsymbol{x},\boldsymbol{\zeta}(\omega))\right)$$

where ρ can either be a law-invariant risk measure (see Kusuoka 2001), e.g. expected value or Conditional Value-at-Risk (CVaR), or a worst-case risk measure (see Postek et al. 2018), e.g. a worst-case expected value that accounts for incomplete distribution information. For example, Natarajan et al. (2014) proposed an ex-post regret minimization model equipped with a worst-case CVaR risk measure that accounted for information about the marginal distribution of the different terms of ζ . Indeed, having access to distributional information enables one to employ a variety of popular risk measures, which can help further control conservatism by trading off between the expected value and tail risks of the regret with respect to a fully informed decision.

A second approach departs from the traditional ex-post regret form as it instead measures regret with respect to an action x' that does not have knowledge of the realized scenario. This rather gives rise to what can be referred as the ex-ante risk-averse regret minimization problem:

(EA-RAR) minimize
$$\max_{\boldsymbol{x} \in \mathcal{X}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}, \boldsymbol{\zeta}(\omega))\right)$$
.

Such an approach was for example used in Perakis and Roels (2008), where the risk measure takes the form of a worst-case expectation. To clarify further, we illustrate the distinction between the two approaches using a simple project selection problem with partial distribution information as an example.

Example 2.1.1 A manager must choose one of the three available projects for investment (i.e. $\mathcal{X} := \{x_A, x_B, x_C\}$) and considers two possible scenarios (i.e. $\Omega := \{\omega_1, \omega_2\}$) for the projects' payoff. Although the true probability of each scenario is not known to the manager, she considers two different possibilities (i.e. $\mathcal{P} := \{\mathbb{P}^I, \mathbb{P}^{II}\}$) and employs worst-case expected value as the risk measure.¹ Table 2.1 provides the numerical details while Table 2.2 presents the optimal project selected under four different regret minimization formulations: {Ex-ante/Ex-post} {Worst-case/Risk-averse} regret minimization. The reader is referred to Section 2.6.1 for further numerical details.

Table 2.1: Numerical Details of Example 2.1.1

| Project payoffs | Probabilities

Table 2.2: Optimal Project Selected
under Four Variants of Regret Mini-
mization in Example 2.1.1

	1 loject pujono			11000	ie militee			
	x_A	x_B^-	x_C	\mathbb{P}^{I}	\mathbb{P}^{II}		Ex-anto	Ex-post
(1)1	1\$	5\$	4\$	80%	0%		LX-ante	Ex-post
ω_1	14	οφ	1 ψ	0070	0,0	Worst-case	x_C	x_C
ω_2	6\$	2\$	3\$	20%	100%	ר ית		
	1			I		Kisk-averse	x_A	$oldsymbol{x}_C$

Specifically, both ex-post models measure the regret under each outcome by comparing to the best action in hindsight: i.e., x_B and x_A under scenarios ω_1 and ω_2 respectively. However, exante model needs to consider the same action x' to compare to under all the scenarios. Namely, in the case of the risk-averse model, we have $x'^* = x_B$. Looking at Table 2 one can remark that while under a worst-case regret formulation, the optimal decision is unaffected by the use of ex-post or ex-ante regret, this is not the case anymore when using a risk-averse setting.

Example 2.1.1 raises questions such as, under what conditions are EP-RAR and EA-RAR equivalent? Do other formulations exist between ex-ante and ex-post that could fill in the gap between the two solutions (especially in a multistage setting)? And, finally, what are the implications of these formulations in terms of level of conservatism? To the best of our knowledge, this chapter investigates these questions for the first time and by presenting a new multistage regret minimization formulation that measures regret with respect to decisions that can exploit information revealed up to Δ stages into the future. This model effectively interpolates very naturally between the ex-ante (with $\Delta = 0$) and ex-post (with $\Delta = \infty$) models and effectively allows to study them under the same lens.

Overall, the contributions of this chapter can be summarized as follows:

- Theoretically, we show that EP-RAR and EA-RAR are equivalent in terms of their optimal solution in a risk neutral setting, and equivalent both in optimal solution and value when a worst-case risk measure is used if a "relatively complete recourse property" is satisfied.
- Methodologically, we introduce the Δ-regret model for multistage stochastic programming under a discrete probability space. We show how this model can be evaluated over a continuum of Δ values and can be reformulated as a special class of two-stage robust linear program that is amenable to a rich range of solution schemes when the stochastic program is linear.
- Numerically, we investigate the effect of Δ and risk aversion on the conservatism of solutions proposed by the Δ-regret model in a multistage inventory management problem. We further illustrate the effect of enforcing different information lookahead levels Δ on the experienced regret.

The rest of the chapter is composed as follows. Section 2.2 reviews the relevant literature. Section 2.3 presents the Δ -regret model, and an illustrative example involving an inventory management problem. Section 2.4 presents our theoretical contributions and proposes a solution scheme while Section 2.5 presents our numerical experiments. All proofs are deferred to Section 2.6. We finally refer interested readers to Poursoltani et al. (2021) for a discussion on how to extend our Δ -regret model to the fractional Δ setting.

2.2 Literature Review

Since the first introduction by Savage (1951), regret minimization has been used in a wide range of applications including single-period portfolio selection (Lim et al. 2012), shortest path, subset selection (Natarajan et al. 2014), spanning tree, ranking problems (Audibert et al. 2014), and in pricing and mechanism design (Caldentey et al. 2017, Koçyiğit et al. 2022) to name a few. Broadly speaking, the regret minimization models that are found in the literature can be classified based on three elements, e.g., the type of risk measure employed for regret evaluation, the length of the planning horizon, and the type of nonanticipativity constraint imposed on the benchmark policy.

In a single-stage setting, the majority of studies focus on the ex-post worst-case regret minimization problem (see for e.g. Feizollahi and Averbakh 2014, Furini et al. 2015, Park et al. 2021), perhaps because of the ease of requiring only support information about the unknown parameters. Under the assumption of partial distribution information, Natarajan et al. (2014) study ex-post regret using a worst-case Conditional Value-at-Risk measure. Chen et al. (2006) exploit a similar regret model but in the context of a *p*-median problem. In these formulations, the benchmark policy can be seen as exploiting both the information about the distribution that employs a worst-case risk measure (for an application to the newsvendor problem, see Yue et al. 2006 and Perakis and Roels 2008). Indeed, these works employ a worst-case expected value to measure regret and rather interpret it as the expected value of distribution information (EVDI) or the maximum value of stochastic modeling (see Delage et al. 2014), due to the following equivalence:

$$\max_{\boldsymbol{x}' \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}, \boldsymbol{\zeta}(\omega)) \right] = \sup_{\mathbb{P} \in \mathcal{P}} \left(\left(\max_{\boldsymbol{x}' \in \mathcal{X}} \mathbb{E}_{\mathbb{P}} \left[h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) \right] \right) - \mathbb{E}_{\mathbb{P}} [h(\boldsymbol{x}, \boldsymbol{\zeta}(\omega))] \right)$$

In the literature studying EVDI of the newsvendor problem, one can mention that Chen and Xie (2021) assume concurrent demand and supply randomness and Zhu et al. (2013) provide closed-form solutions for the relative EVDI. Other applications of EVDI can be found in portfolio optimization (see Lim et al. 2012, Benati and Conde 2022), fleet mix optimization (see Delage et al. 2014), and blood classification (see El-Amine et al. 2018).

In the multistage setting, most studies focus on a two-stage setting under an ex-post worst-case regret minimization (see Bertsimas and Dunning 2020, Poursoltani and Delage 2022 and references therein). Additionally, Xu et al. (2015) study a two-stage bidding problem in an electricity market, where perfect distribution information is assumed and different risk measures are applied on the realized ex-post regret. Similar approaches were used in Zhang et al. (2020). Lim et al. (2006) investigate ex-ante and ex-post worstcase expected regret models in a fully multistage framework involving either an inventory management or a portfolio optimization problem. The authors draw connections between regret minimization and Bayesian learning.

There has also been an interest in the economics literature to study the role that regret can play in a dynamic environment. For instance, Hayashi (2011) and Halpern and Leung (2016) study different forms of ex-post regret models and identify conditions under which regret minimizing policies are dynamically consistent. Alternatively, one can refer to the work of Krähmer and Stone (2008), which considers a two-stage setting where the decision maker optimizes a trade-off between the expected payoff and a weighted sum of the regret experienced at different point of time. Finally, Hayashi (2009) explores dynamic consistency and the role of ex-post regret in optimal stopping problems, while Strack and Viefers (2021) explores in the same application the effect of using stopping time to control the horizon over which the ex-post regret is measured.

This chapter can be viewed to contribute to multistage regret theory from an optimization point of view. Indeed, we propose for the first time an intuitive risk-averse multistage regret minimization problem where the pessimism of the benchmark policy set is controlled using a bound on the maximum amount of look-ahead. This Δ -regret model naturally interpolates between the ex-ante and ex-post regret models. Furthermore, we explore the properties of Δ -regret models under popular risk measures and provide a promising direction for numerical resolution of these models, which is based on the recent advances in two-stage robust optimization. These later results can be seen as an interesting extension of Poursoltani and Delage (2022) to the multistage setting.

2.3 \triangle -Regret Minimization in Multistage Stochastic Programs

We consider a multistage decision making setting in which at each stage $t \in \{1, ..., T\}$ a decision maker needs to make a decision $x_t \in \mathbb{R}^n$ based on the available historical information captured by $[\zeta_1 \ \zeta_2 \ \cdots \ \zeta_{t-1}]$. Focusing on a discrete probability space $(\Omega, \Sigma, \mathbb{Q})$, where \mathbb{Q} is assumed strictly positive without loss of generality, one classical decisionmaking approach formulates the following multistage stochastic program:

$$(\mathcal{MSP}) \min_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \rho(-h(\boldsymbol{x},\boldsymbol{\zeta}))$$
(2.1)
where $\boldsymbol{x} : \Omega \to \mathbb{R}^{n \times T}$ is the multistage policy, $\boldsymbol{\zeta} : \Omega \to \mathbb{R}^{m \times T-1}$ is the concatenated matrix of the random vectors observed over the whole horizon, $h(\boldsymbol{x}, \boldsymbol{\zeta})$ is the cumulative profit of implementing policy \boldsymbol{x} when $\boldsymbol{\zeta}$ is realized, ρ is a convex risk measure that maps a random cost to a risk level, $\mathcal{X} := \{\boldsymbol{x} : \Omega \to \mathbb{R}^{n \times T} | \boldsymbol{x}(\omega) \in \mathcal{X}_{\omega}, \omega \in \Omega\}$, with bounded $\mathcal{X}_{\omega} \subseteq \mathbb{R}^{n \times T}$, imposes "physical" constraints that must be satisfied by the policy under each outcome in Ω , while \mathcal{X}_{na} ensures that the policy is nonanticipative with respect to the information revealed by $\boldsymbol{\zeta}$. We formalize below some of these elements.

Definition 2.3.1 *The set of nonanticipative policies takes the form:*

 $\mathcal{X}_{na} := \left\{ \boldsymbol{x}: \Omega \to \mathbb{R}^{n \times T} \, \big| \, \boldsymbol{x}_t(\omega) = \boldsymbol{x}_t(\omega'), \; \forall \omega, \omega' \in \Omega: \boldsymbol{\zeta}^{[t-1]}(\omega) = \boldsymbol{\zeta}^{[t-1]}(\omega'), \; \forall t \in \{1, 2, ..., T\} \right\},$

where $[t] := \{1, \ldots, t\}, \zeta^{[t-1]}(\omega)$ denotes the concatenated matrix of the random vectors observed till time step t under scenario ω , and $\zeta^{[0]}(\omega) = \zeta^{[0]}(\omega')$ is interpreted as always true.

Definition 2.3.2 According to Föllmer and Schied (2002), letting $\mathcal{L} := \{\xi : \Omega \to \mathbb{R}\}$ be the space of all possible finite random financial loss², ρ is a **convex risk measure** if and only if it satisfies:

- Monotonicity: $\forall \xi^1, \xi^2 \in \mathcal{L}, \xi^1 \ge \xi^2 \text{ a.s.} \Rightarrow \rho(\xi^1) \ge \rho(\xi^2);$
- Translation invariance: $\forall \xi \in \mathcal{L}, t \in \mathbb{R}, \rho(\xi + t) = \rho(\xi) + t;$
- Convexity: $\forall \xi^1, \xi^2 \in \mathcal{L}$, and $\theta \in [0, 1]$, $\rho(\theta \xi^1 + (1 \theta)\xi^2) \leq \theta \rho(\xi^1) + (1 \theta)\rho(\xi^2)$.

Moreover, ρ is considered a **coherent risk measure** if it further satisfies:

• Scale invariance: $\forall \xi \in \mathcal{L}, \alpha \geq 0, \rho(\alpha \xi) = \alpha \rho(\xi).$

In particular, it is well known that $\rho(-h(\boldsymbol{x}, \boldsymbol{\zeta})) = \mathbb{E}_{\mathbb{Q}}[-h(\boldsymbol{x}, \boldsymbol{\zeta})]$ and Conditional Value-at-Risk (see Example 3.11 for a definition) fall in the class of coherent risk measure. Unless specified otherwise, in what follows we will assume that ρ is a convex risk measure.

To improve computational tractability, we will later (when indicated) focus on the class of problems where the constraints and the objective function are affine with respect to x.

Assumption 2.3.1 [Stochastic Linear Programming] The profit function is an affine function of x defined as

$$h(\boldsymbol{x},\boldsymbol{\zeta}) := \sum_{t=1}^{T} \boldsymbol{c}_{t}^{\top}(\boldsymbol{\zeta})\boldsymbol{x}_{t} + d(\boldsymbol{\zeta}), \qquad (2.2)$$

for some arbitrary $\mathbf{c}_t : \mathbb{R}^{m \times (T-1)} \to \mathbb{R}^n$ and $d : \mathbb{R}^{m \times (T-1)} \to \mathbb{R}$. Furthermore, for each $\omega \in \Omega$, \mathcal{X}_{ω} is a bounded polyhedron formulated as:

$$\mathcal{X}_{\omega} := \left\{ \boldsymbol{x} \in \mathbb{R}^{n \times T} \left| \sum_{t=1}^{T} \boldsymbol{a}_{jt}(\boldsymbol{\zeta}(\omega))^{\top} \boldsymbol{x}_{t}(\omega) \le b_{j}(\boldsymbol{\zeta}(\omega)), \ j = 1, 2, ..., \mathcal{J}
ight\},
ight.$$

with arbitrary $a_{jt} : \mathbb{R}^{m \times (T-1)} \to \mathbb{R}^n$ and $b_j : \mathbb{R}^{m \times (T-1)} \to \mathbb{R}$, for all j and t.

Recall that the regret models discussed in the introduction addressed a static decision model (namely with T = 2). Hence, the main difference between the ex-post and ex-ante models hinged on whether the benchmark action x' could fully anticipate or not realization ζ . A natural question to pose is therefore how the concept of regret extends in the multistage problems where we have T > 2 and where the values of ζ are progressively revealed in time. In what follows, we propose a multistage regret minimization formulation that measures regret with respect to a benchmark policy that can exploit information revealed up to Δ stages into the future, which we term Δ -regret. This model effectively interpolates very naturally between the ex-ante (with $\Delta = 0$) and ex-post (with $\Delta = \infty$) models and effectively allows to study them under the same lens. Section 2.3.1 will present the Δ -regret model, and Section 2.3.2 will present an illustrative example involving a multistage inventory management problem.

2.3.1 The \triangle -Regret

In a multistage decision making problem, a regret-averse policy maker might be interested to compare his decisions to benchmark policies that exploit shorter foresight than the total planning horizon. This gives rise to the idea of the Δ -regret model, where the benchmark policies are capable of predicting the future realizations up to Δ steps ahead in the future. As an immediate result of such setting, the benchmark policies can adapt to the information released till time step $t + \Delta$. Assuming ρ is a convex risk measure, the Δ -regret model in the multistage setting is formulated as

$$(\Delta\text{-regret}) \quad \min_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \qquad \mathcal{R}_{\Delta}(\boldsymbol{x}), \tag{2.3a}$$

where

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) := \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \quad \rho(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})), \tag{2.3b}$$

and where \mathcal{X}_{Δ} is the space of policies that violate the nonanticipativity constraints by up to a margin of Δ steps. More specifically,

$$\mathcal{X}_{\Delta} := \left\{ \boldsymbol{x} : \Omega \to \mathbb{R}^{n \times T} \middle| \begin{array}{l} \boldsymbol{x}_t(\omega) = \boldsymbol{x}_t(\omega') \\ \forall \omega, \omega' \in \Omega : \boldsymbol{\zeta}^{[t+\Delta-1]}(\omega) = \boldsymbol{\zeta}^{[t+\Delta-1]}(\omega'), \ \forall t \in \{1, 2, ..., T\} \end{array} \right\},$$

where we interpret $\boldsymbol{\zeta}^{[t]} := \boldsymbol{\zeta}$ when $t \ge T - 1$. For any $\Delta \in \{0, 1, 2, 3, ..., T - 1\}$, the Δ -regret model will evaluate the regret of the prescribed decisions as contrasted with the ones that could have been made if the uncertain parameters were revealed up to Δ steps ahead of time. Clearly, when $\Delta = 0$, \mathcal{X}_{Δ} reduces to \mathcal{X}_{na} , implying that the benchmark policy has no access to any realization beforehand. On the contrary, $\Delta = T - 1$ gives the benchmark policy full access to all the realizations of $\boldsymbol{\zeta}$ at any point of time. The Δ -regret model therefore naturally interpolates between the ex-ante and ex-post regret models. In addition, regret is a non-decreasing function of Δ . These concepts are formalized in the following lemma.

Lemma 2.3.1 The Δ -regret model, i.e. Problem (2.3), reduces to ex-ante and ex-post regret minimization when $\Delta = 0$ and $\Delta \ge T - 1$ respectively. Moreover, its optimal value is a non-decreasing function of Δ .

2.3.2 Illustrative Example of \triangle -Regret Model

We consider the multistage inventory management problem previously studied in Ben-Tal et al. (2004) and Kuhn et al. (2011). We assume that each period t consists of a day. The inventory system consists of I production facilities, which produce a single item and store it at a shared warehouse. The production cost of a single unit of the item on day t at facility i is c_{it} and the objective is to determine the optimal production level of each production facility (x_{it}) to satisfy the uncertain demand and minimize the total production cost over a planning horizon of T days. While \bar{x}_{it} indicates the production capacity of production facility i on day t, the maximum production potential over the whole planning horizon is determined by $\bar{x}_{i,tot}$. The minimum and the maximum inventory levels that should be maintained at the end of each day are denoted by \underline{x}_{wh} and \bar{x}_{wh} , respectively, and x_{wh}^0 represents the initial inventory level. If $d_t(\omega) \in \mathbb{R}$ denotes the demand of day t under scenario ω , then we let $\zeta_{t-1}(\omega) := d_t(\omega) \in \mathbb{R}$ to model the fact that the demand for day t is known when deciding of the production levels at the beginning of the day: this occurs for instance when orders for pick up need to be made at the latest one day before pickup. The MSP for this inventory problem takes the form where:

$$h(\boldsymbol{x}, \boldsymbol{\zeta}) := -\sum_{t=1}^{T} \sum_{i=1}^{I} \boldsymbol{c}_{it}^{\top} \boldsymbol{x}_{it}(\omega),$$

and

$$\mathcal{X}_{\omega} := \left\{ \boldsymbol{x} \in \mathbb{R}^{n \times T} \middle| \begin{array}{l} 0 \leq x_{it} \leq \bar{x}_{it}, \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \\ \sum_{t=1}^{T} x_{it} \leq \bar{x}_{i,tot}, \forall i \in \mathcal{I} \\ \underline{x}_{wh} \leq x_{wh}^{0} + \sum_{s=1}^{t} \sum_{i=1}^{I} x_{is} - \sum_{s=1}^{t-1} \zeta_{s}(\omega) - d_{1} \leq \bar{x}_{wh}, \forall t \in \mathcal{T} \end{array} \right\}.$$

We consider a simple instance of this problem with 2 production facilities (I = 2), 3 days planning horizon (T = 3) and 5 demand pattern scenarios ($\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$). In this setting, considering the scenario tree structure depicted in Figure 2.1, the nonanticipativity constraints are expressed as:

$$\mathcal{X}_{na} := \left\{ oldsymbol{x} : \Omega o \mathbb{R}^{2 imes 3} \left| egin{array}{l} oldsymbol{x}_1(\omega_1) = oldsymbol{x}_1(\omega_2) = oldsymbol{x}_1(\omega_3) = oldsymbol{x}_1(\omega_4) = oldsymbol{x}_1(\omega_5) \ oldsymbol{x}_2(\omega_1) = oldsymbol{x}_2(\omega_2), \ oldsymbol{x}_2(\omega_4) = oldsymbol{x}_2(\omega_5) \end{array}
ight\}.$$

When measuring regret, the policy maker might be interested in comparing her policy to one that benefits from the same information. This is an immediate implication of $\Delta = 0$ in the Δ -regret model. Setting Δ to 1 allows her to measure her regret with respect to the policy made under one stage look-ahead information. Eventually, $\Delta = 2$ compares to policies that exploit the full information. Specifically, we have the following reductions:

$$\mathcal{X}_0 = \mathcal{X}_{na}, \ \mathcal{X}_1 = \left\{ \boldsymbol{x}: \Omega \to \mathbb{R}^{2 \times 3} \left| \left| \boldsymbol{x}_1(\omega_1) = \boldsymbol{x}_1(\omega_2), \ \boldsymbol{x}_1(\omega_4) = \boldsymbol{x}_1(\omega_5) \right| \right\}, \ \mathcal{X}_2 = \left\{ \boldsymbol{x}: \Omega \to \mathbb{R}^{2 \times 3} \right\}.$$

Figure 2.1 illustrates each policy sets in a scenario tree. This example shows how increasing Δ lifts the constraints imposed on x' gradually and the full access of the realized scenario is bestowed upon x' when Δ is at its maximum value.



Figure 2.1: Comparison of Adaptation Power Between x (Beside the Timeline) and x' (on the Right of Each Tree) as a Function of Δ .

The nodes of the tree present what information is available at each point of time.

2.4 Properties of \triangle -Regret Model Under Risk Measures

In this section, we explore interesting properties that arise under different choice of risk measures for the Δ -regret model. In particular, the first two subsections initially study the properties that emerge under specific coherent risk measures, namely the worst-case $\rho(\xi) = \text{ess sup}(\xi)$ and expected value $\rho(\xi) = \mathbb{E}_{\mathbb{Q}}[\xi]$. We then consider the general class of coherent risk measures using their worst-case expectation representation, i.e. $\rho(\xi) := \sup_{\mathbb{P}\in\mathcal{P}} \mathbb{E}_{\mathbb{P}}[\xi]$ (see Artzner et al. 1999). We will show that under a worst-case risk measure, all Δ -regret models are equivalent if (and only if) a *relatively complete recourse* property is satisfied (see Theorem 2.4.1). This will also occur, yet only in terms of the optimal solution set, for models that employ an expected value. Finally, we will derive a reformulation for all coherent risk measures that take the form of a two-stage robust linear program when the stochastic program is linear and the risk measure is linear programming representable.

2.4.1 The Case of $\rho(\xi) = \operatorname{ess\,sup}(\xi)$

In this section, we consider measuring the Δ -regret using the essential supremum as the risk measure:

$$\rho(\xi) := \operatorname{ess\,sup}(\xi) = \inf \left\{ a \,|\, \mathbb{P}(\xi > a) = 0 \right\}$$

In particular, we will confirm conditions under which, the invariability of Δ -regret to Δ , observed in Example 2.1.1, holds. In order to present our main result, we first introduce an assumption about the MSP.

Assumption 2.4.1 *The multistage stochastic program satisfies the relatively complete recourse property, i.e.,*

$$\mathcal{X}_{\omega}^{[t]} = \mathcal{X}_{\omega'}^{[t]}, \quad \forall (\omega, \omega') : \zeta^{[t-1]}(\omega) = \zeta^{[t-1]}(\omega'), \forall t,$$

where

$$\mathcal{X}_{\omega}^{[t]} := \{ \boldsymbol{x} \in \mathbb{R}^{n \times t} | \exists \bar{\boldsymbol{x}} \in \mathbb{R}^{n \times T-t}, [\boldsymbol{x} \ \bar{\boldsymbol{x}}] \in \mathcal{X}_{\omega} \}$$

is a projection of \mathcal{X}_{ω} *on the space spanned by the decision vectors* $x_1, x_2, \dots x_t$ *.*

In simpler words, this assumption imposes that when looking at the set of feasible decisions x in hindsight, this set only includes candidates that had 100% chances of being feasible at the time that they were implemented. While the decision to satisfy this assumption is an important modeling choice in designing the Δ -regret model and might affect the measured regret (see Example 2.4.1 below), it is in fact always possible to modify a multistage stochastic program so that the property is satisfied.

Lemma 2.4.1 Given a MSP, one can construct \overline{MSP} that produces the same optimal value and optimal solution set as MSP while satisfying the relatively complete recourse assumption, *i.e.* Assumption 2.4.1.

We can now turn to the main result of this section, which indicates that relatively complete recourse is a necessary and sufficient condition for the Δ -regret model to be insensitive to Δ under the essential supremum risk measure.

Theorem 2.4.1 *Given that Assumption 2.4.1 holds and* $\rho(\xi) = \operatorname{ess\,sup}(\xi)$ *, for any arbitrary* $\Delta \ge 0$ *, the objective function of Problem (2.3) reduces to*

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \mathcal{R}_{T-1}(\boldsymbol{x}) = \max_{\omega \in \Omega} \max_{\boldsymbol{x}' \in \mathcal{X}_{\omega}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)).$$
(2.4)

Hence, Problem (2.3) *produces the same optimal value and solution set for all values of* $\Delta \ge 0$ *.*

At first glance, the result of Theorem 2.4.1 looks intuitive since essential supremum hedges against a single worst-case scenario. This arises from the interchangeability of the order of two maximization problems. Thus imposing a nonanticipative structure on $x' \in \mathcal{X}_{\Delta}$ will have no effect for any value of Δ . What is less intuitive is the role of Assumption 2.4.1. In this regard, the following example supports and illustrates our claims that the relatively complete recourse property is necessary to obtain this invariance, and that the \mathcal{MSP} can always be reformulated to satisfy this assumption.

Example 2.4.1 Consider a simple two-stage (i.e. T = 2) problem, where the set of first and second stage actions are defined as $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2\}$, respectively. After implementing the first stage decision, the decision maker is faced with two scenarios, ω_1 and ω_2 with 10% and 90% chances, respectively. We consider the following definition for \mathcal{X} :

$$\mathcal{X} := \left\{ \boldsymbol{x} : \left\{ \omega_1, \omega_2 \right\} \to A \times B | \boldsymbol{x}(\omega_1) \in \mathcal{X}_{\omega_1}, \ \boldsymbol{x}(\omega_2) \in \mathcal{X}_{\omega_2} \right\},\$$

with:

$$\mathcal{X}_{\omega_1} := \{(a_1, b_1), (a_3, b_1)\}$$
$$\mathcal{X}_{\omega_2} := \{(a_2, b_2), (a_3, b_2)\}$$

In words, if the decision maker chooses a_1 , he can react to ω_1 with b_1 but has no feasible recourse against ω_2 . The reverse is true for a_2 , while a_3 enables both the b_1 and b_2 actions under ω_1 and ω_2 respectively. The profit function, defined only over feasible pairs, takes the form described in Table 2.3.

In this example, $\mathcal{X}_{\omega_1}^{[1]} = \{a_1, a_3\} \neq \mathcal{X}_{\omega_2}^{[1]} = \{a_2, a_3\}$, which indicates that Assumption 2.4.1 is violated. Furthermore, there is only one feasible policy for the decision maker, i.e. $\{\bar{x}\} = \{(a_3, b_1\mathbf{1}\{\omega = \omega_1\} + b_2\mathbf{1}\{\omega = \omega_2\})\} = \mathcal{X} \cap \mathcal{X}_{na}$. Focusing on this policy, one can compute

Actions	$h(oldsymbol{x},oldsymbol{\zeta}(\omega_1))$	$h(oldsymbol{x},oldsymbol{\zeta}(\omega_2))$
(a_1, b_1)	4	-
(a_2, b_2)	-	1
(a_3, b_1)	3	-
(a_3, b_2)	-	0

Table 2.3: Profit Function in Example 2.4.1

the Δ -regret under the essential supremum measure as follows. In the case of $\Delta = 1$, the feasible space for the benchmark policy \mathbf{x}' becomes $\mathcal{X} \cap \mathcal{X}_1 = \mathcal{X} = \{\bar{\mathbf{x}}, \bar{\mathbf{x}}', \bar{\mathbf{x}}'', \bar{\mathbf{x}}'''\}$ with

$$\begin{split} \bar{x}' &:= (a_1 \mathbf{1}\{\omega = \omega_1\} + a_2 \mathbf{1}\{\omega = \omega_2\}, \ b_1 \mathbf{1}\{\omega = \omega_1\} + b_2 \mathbf{1}\{\omega = \omega_2\}).\\ \\ \bar{x}'' &:= (a_1 \mathbf{1}\{\omega = \omega_1\} + a_3 \mathbf{1}\{\omega = \omega_2\}, \ b_1 \mathbf{1}\{\omega = \omega_1\} + b_2 \mathbf{1}\{\omega = \omega_2\}).\\ \\ \\ \bar{x}''' &:= (a_3 \mathbf{1}\{\omega = \omega_1\} + a_2 \mathbf{1}\{\omega = \omega_2\}, \ b_1 \mathbf{1}\{\omega = \omega_1\} + b_2 \mathbf{1}\{\omega = \omega_2\}). \end{split}$$

We can thus conclude that:

$$\mathcal{R}_1(\bar{\boldsymbol{x}}) = \max_{\boldsymbol{x}' \in \{\bar{\boldsymbol{x}}, \bar{\boldsymbol{x}}', \bar{\boldsymbol{x}}'', \bar{\boldsymbol{x}}'''\}} \operatorname{ess\,sup}(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\bar{\boldsymbol{x}}, \boldsymbol{\zeta}))$$
$$= \max(\operatorname{ess\,sup}(0), \operatorname{ess\,sup}(1), \operatorname{ess\,sup}(\mathbf{1}\{\omega = \omega_1\}), \operatorname{ess\,sup}(\mathbf{1}\{\omega = \omega_2\})) = 1.$$

On the other hand, if $\Delta = 0$, we have that $\mathbf{x}' \in \mathcal{X} \cap \mathcal{X}_0 = \mathcal{X} \cap \mathcal{X}_{na} = \{\bar{\mathbf{x}}\}$, i.e. the benchmark decision must be chosen from among the same sets of decision as for the decision maker. This naturally leads to $\mathcal{R}_0(\bar{\mathbf{x}}) = 0$. We, therefore, showed that when Assumption 2.4.1 is violated, it is possible that $0 = \mathcal{R}_0(\mathbf{x}) \neq \mathcal{R}_1(\mathbf{x}) = 1$.

We close this example with the observation that if the MSP was modified as proposed in Lemma 2.4.1, then we would have:

$$\begin{split} \bar{\mathcal{X}}_{\omega_1} &:= \{ \boldsymbol{x} \in A \times B | x_1 \in \mathcal{X}_{\omega_1}^{[1]} \} \cap \{ \boldsymbol{x} \in A \times B | x_1 \in \mathcal{X}_{\omega_2}^{[1]} \} \cap \mathcal{X}_{\omega_1} \\ &= \{ \boldsymbol{x} \in A \times B | x_1 \in \mathcal{X}_{\omega_1}^{[1]} \cap \mathcal{X}_{\omega_2}^{[1]} \} \cap \mathcal{X}_{\omega_1} = \{ \boldsymbol{x} \in A \times B | x_1 \in \{a_3\} \} \cap \mathcal{X}_{\omega_1} = \{ (a_3, b_1) \} \end{split}$$

while

$$\bar{\mathcal{X}}_{\omega_2} := \{ \boldsymbol{x} \in A \times B | x_1 \in \mathcal{X}_{\omega_1}^{[1]} \} \cap \{ \boldsymbol{x} \in A \times B | x_1 \in \mathcal{X}_{\omega_2}^{[1]} \} \cap \mathcal{X}_{\omega_2} = \{ (a_3, b_2) \}.$$

Using $\bar{\mathcal{X}} := \{ \boldsymbol{x} : \Omega \to A \times B \, | \, \boldsymbol{x}(\omega) \in \bar{\mathcal{X}}_{\omega}, \, \omega \in \Omega \}$ instead of \mathcal{X} does not affect the solution of the \mathcal{MSP} since in any case $\mathcal{X} \cap \mathcal{X}_{na} = \{ \bar{\boldsymbol{x}} \} = \bar{\mathcal{X}} = \bar{\mathcal{X}} \cap \mathcal{X}_{na}$. Yet, using $\bar{\mathcal{X}}$ instead of \mathcal{X} does affect

the Δ -regret model given that under \bar{X} now we have that $\mathcal{R}_1(\bar{x}) = \mathcal{R}_0(\bar{x}) = 0$. Intuitively, from the regret perspective, the difference between the two models reduces to whether the benchmark policy is allowed to implement decisions that don't have a probability one guarantee to lead to long-term feasibility.

2.4.2 The Case of $\rho(\xi) = \mathbb{E}_{\mathbb{Q}}[\xi]$

For a given probability measure \mathbb{Q} , the expected value can be considered as another option among the popular risk measures. However, for any value of Δ , Problem (2.3) produces the same optimal solution as the \mathcal{MSP} . This is formalized in the following proposition.

Proposition 2.4.1 *Given that* $\rho(\xi) := \mathbb{E}_{\mathbb{Q}}[\xi]$ *, for any arbitrary* Δ *, Problems* (2.1) *and* (2.3) *have the same set of optimal solutions as the* \mathcal{MSP} *. Namely,*

$$rg\min_{oldsymbol{x}\in\mathcal{X}\cap\mathcal{X}_{na}}
hoig(-h(oldsymbol{x},oldsymbol{\zeta})ig) \ = \ rg\min_{oldsymbol{x}\in\mathcal{X}\cap\mathcal{X}_{na}} \ \mathcal{R}_{\Delta}(oldsymbol{x})$$

While Proposition 2.4.1 establishes that all Δ -regret models produce the same optimal solution under a risk neutral setting, we will see in our numerical experiments that optimal values do change for different values of Δ . Interestingly, in the case of $\Delta = 0$, one can confirm that a risk neutral decision maker never experiences regret if she acts optimally.

Corollary 2.4.1 The optimal value of Problem (2.3) with $\Delta = 0$ and $\rho(\xi) := \mathbb{E}_{\mathbb{Q}}[\xi]$ is equal to zero and achieved by the MSP solution.

In particular, Proposition 2.4.1 and Corollary 2.4.1 suggest that the regret experienced by a decision maker can be decomposed into three positive components

$$\begin{aligned} \mathcal{R}_{\Delta}(\boldsymbol{x}) &= \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right) \\ &= \left(\max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right) - \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{0}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right)\right) \\ &+ \left(\max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{0}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right) - \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{0}} \mathbb{E}_{\mathbb{Q}}\left[h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right]\right) \\ &+ \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{0}} \mathbb{E}_{\mathbb{Q}}\left[h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right] \end{aligned}$$

$$= \left(\max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right) - \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{0}} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right)\right)$$
(2.5a)

$$+ \left(\max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_0} \rho\left(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right) - \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_0} \mathbb{E}_{\mathbb{Q}}\left[h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})\right]\right)$$
(2.5b)

$$+ \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}^*, \boldsymbol{\zeta})] - \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}, \boldsymbol{\zeta})] .$$
(2.5c)

The first component (2.5a) captures the part of the regret that comes from the information that is out of the decision maker's reach. The second component (2.5b) captures a part of the regret that comes from risk aversion of the decision maker. Finally, the last component (2.5c) comes from not being optimal with respect to the risk neutral version of the MSP.

2.4.3 The Case of Coherent Risk Measures

In a more general context, it is well-known (see Artzner et al. 1999) that any coherent risk measure can be represented using a worst-case expectation formulation :

$$\rho(\xi) := \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\xi],$$

where \mathcal{P} is a non-empty convex set of probability measures that, in the distributionally robust optimization literature, is also referred to as an ambiguity set known to contain the true underlying measure \mathbb{Q} .

Definition 2.4.1 The ambiguity set \mathcal{P} is a bounded convex set which implies that

$$\sup_{\mathbb{P}\in\mathcal{P}} \mathbb{E}_{\mathbb{P}}[\xi] = \max_{\boldsymbol{p}\in\mathcal{D}\cap\mathcal{M}} \sum_{\omega\in\Omega} p_{\omega}\xi(\omega)$$

where $\xi : \Omega \to \mathbb{R}$, $\mathcal{M} \subset \mathbb{R}^{|\Omega|}$ is the simplex set, and $\mathcal{D} \subset \mathbb{R}^{|\Omega|}$ denotes a general convex and compact set.

Taking advantage of Definition 2.4.1, Problem (2.3) can be cast as a two-stage robust optimization problem. This is formalized in the following proposition.

Proposition 2.4.2 Given $\rho(\xi) := \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\xi]$, Problem (2.3) reduces to

$$\begin{array}{l} \underset{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}}{\text{min } r} & \text{min } r \\ \text{s.t. } \delta^{*}(\boldsymbol{v}|\mathcal{D}) - v_{\omega} + h(\boldsymbol{x}'(\omega), \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)) \leq r, \ \forall \omega \in \Omega, \\ \end{array}$$

$$(2.6b)$$

where $r \in \mathbb{R}$, $v \in \mathbb{R}^{|\Omega|}$ and $\delta^*(v|\mathcal{D}) := \sup_{\boldsymbol{p} \in \mathcal{D}} \boldsymbol{p}^\top v$ represents the support function of \mathcal{D} .

In general, Problem (2.6) is a non-linear two-stage robust optimization problem. However, under a number of popular ambiguity sets, \mathcal{D} is polyhedral thus the support function of $\delta^*(v|\mathcal{D})$ renders a linear programming representation, which in turn makes (2.6) a robust linear two-stage program. Such choices include the sets associated to Conditional Value-at-Risk or expectiles (see Bellini and Bernardino 2017), and, in the Distributionally Robust Optimization (DRO) literature, some type-1 Wasserstein ambiguity sets (see Mohajerin Esfahani and Kuhn 2018) or some sets based on hypothesis testing (see Bertsimas et al. 2018).

Corollary 2.4.2 Given that Assumption 2.3.1 is satisfied and that $\mathcal{D} := \{ \boldsymbol{p} \in \mathbb{R}^{|\Omega|} | \exists \boldsymbol{q} \in \mathbb{R}^{n_q}, B_p \boldsymbol{p} + B_q \boldsymbol{q} \leq \boldsymbol{b} \}$, where $B_p \in \mathbb{R}^{m \times |\Omega|}, B_q \in \mathbb{R}^{m \times n_q}, \boldsymbol{b} \in \mathbb{R}^m$, then Problem (2.6) reduces to the following robust two-stage linear optimization problem:

$$\min_{\boldsymbol{x}\in\mathcal{X}\cap\mathcal{X}_{na}} \mathcal{R}_{\Delta}(\boldsymbol{x}) \coloneqq \max_{\boldsymbol{x}'\in\mathcal{X}\cap\mathcal{X}_{\Delta}} \min_{\boldsymbol{r},\boldsymbol{v},\boldsymbol{\lambda}} \boldsymbol{r}$$
s.t. $\boldsymbol{\lambda}^{\top}\boldsymbol{b} - \boldsymbol{v}_{\omega} + \sum_{t=1}^{T} \boldsymbol{c}_{t}^{\top}(\boldsymbol{\zeta}(\omega))(\boldsymbol{x}_{t}'(\omega) - \boldsymbol{x}_{t}(\omega)) \leq \boldsymbol{r}, \ \forall \omega \in \Omega$

$$B_{p}^{\top}\boldsymbol{\lambda} = \boldsymbol{v}$$

$$B_{q}^{\top}\boldsymbol{\lambda} = 0$$

$$\boldsymbol{r} \in \mathbb{R}, \ \boldsymbol{v} \in \mathbb{R}^{|\Omega|}, \ \boldsymbol{\lambda} \in \mathbb{R}_{+}^{m}.$$
(2.7)

Example 2.4.2 Conditional Value-at-Risk (CVaR) evaluates the conditional expectation of the random variable ξ under α % worst scenarios and mathematically takes the form of

$$CVaR_{\alpha}(\xi) := \inf_{t} t + \frac{1}{1-\alpha} \mathbb{E}_{\bar{p}}[\max(0, -\xi - t)], \qquad (2.8)$$

where \bar{p} denotes the reference probability distribution. It has the following worst-case expectation representation (see Rockafellar et al. 2006):

$$CVaR_{\alpha}(\xi) := \sup_{\boldsymbol{p} \in \mathcal{D} \cap \mathcal{M}} \sum_{\omega \in \Omega} p_{\omega}\xi(\omega), \qquad (2.9)$$

where $\mathcal{D} := \{ \boldsymbol{p} \in \mathbb{R}^{|\Omega|} | \boldsymbol{p} \leq \bar{\boldsymbol{p}}/(1-\alpha) \}$, where $\bar{p}_{\omega} := \mathbb{Q}(\omega)$. Based on Corollary 2.4.2, we get the following two-stage linear program:

$$\min_{\boldsymbol{x}\in\mathcal{X}\cap\mathcal{X}_{na}} \max_{\boldsymbol{x}'\in\mathcal{X}\cap\mathcal{X}_{\Delta}} \min_{\boldsymbol{r},\boldsymbol{v}} r$$
(2.10a)

s.t.
$$\frac{\bar{\boldsymbol{p}}^{\top}\boldsymbol{v}}{1-\alpha} - v_{\omega} + \sum_{t=1}^{T} \boldsymbol{c}_{t}^{\top}(\boldsymbol{\zeta}(\omega))(\boldsymbol{x}_{t}'(\omega) - \boldsymbol{x}_{t}(\omega)) \leq r, \ \forall \omega \in \Omega$$
 (2.10b)

$$r \in \mathbb{R}, \ \boldsymbol{v} \in \mathbb{R}^{|\Omega|}_+.$$
 (2.10c)

Unfortunately, Problem (2.7) does inherit some of the NP-hardness properties from the more general class of two-stage robust linear optimization problems (see Ben-Tal et al. 2004) as demonstrated in the following proposition.

Proposition 2.4.3 Evaluating $\mathcal{R}_{\Delta}(x)$ under a coherent risk measure is NP-hard even when \mathcal{P} uses a polyhedral \mathcal{D} .

However, there is a variety of exact and approximate solution schemes that can practically solve Problem (2.7) (see Poursoltani and Delage 2022, Section 3.1). One such exact approach is the column-and-constraint generation algorithm proposed in Zeng and Zhao (2013). Such a procedure can be directly applied to Problem (2.7), which we briefly discuss in Section 2.6.10.

We finish this section by remarking that alternative two-stage robust formulations exist for Problems (2.6) and (2.7) as suggested by Bertsimas and de Ruiter (2016) for the two-stage linear programming case and de Ruiter et al. (2022) for the two-stage non-linear case. In particular, under Assumption 2.3.1, the methods in de Ruiter et al. (2022) can be used to produce a two-stage problem with linear first and second stage constraints, while non-linear constraints, caused by $\delta^*(\boldsymbol{v}|\mathcal{D})$ will now appear in the maximization problem. Such reformulation could open up new avenues for exact and approximate solution approaches for regret minimization problems.

2.5 Numerical Experiments

In this section, we describe a numerical study that provides insights on the relationship between the amount of look-ahead (Δ) allowed for the benchmark policy, the level of risk aversion on the solution quality, and the level of regret. We consider the multistage inventory management problem discussed in Section 2.3.2 and examine the structure of regret minimizing policies and the breakdown of regret as expressed in (2.5). We measure risk using Conditional Value-at-Risk as it allows to easily control the level of risk aversion with a single parameter, namely α . Solutions to all regret models are obtained using the column-and-constraint generation algorithm presented in Section 2.6.10.

We consider two instances (indicated by I and II) with three production facilities and a five-day horizon. We assume that the production cost for all facilities is the lowest in the first period and gets progressively more expensive until the end of the horizon. Furthermore, we assume that the demand information for the next day can be progressively collected on the next morning, and the production can only be planned at the beginning of each day. Overall, the outcome space includes 16 scenarios ($|\Omega| = 16$). The two instances differ in that the total demand over the planning horizon is right-skewed in instance I and left-skewed in instance II. The scenario tree that describes the evolution of the random demand, demand realizations, and other parameters of the inventory model are presented in Section 2.6.12.



Figure 2.2: Optimal Regret and Total First-Stage Production Achieved Given Different Levels of Conservatism (α).

(a) Optimal regret for instance I; (b) Total first-stage production for instance I; (c) Total first-stage production for instance II.

The first experiment compares the optimal Δ -regret values and the optimal first-stage production plans for $\Delta \in \{0, 1, 2, 3, 4\}$ at different levels of risk aversion α . The results are presented in Figure 2.2. Looking at Figure 2.2 (a), concerning instance I, one can remark that for any fixed look-ahead level, increasing the risk aversion level leads to an increased minimal Δ -regret. This originates from the fact that the Conditional Value-at-Risk only considers the worst-case $(1 - \alpha) \cdot 100\%$ of scenarios, e.g., while at $\alpha = 0\%$ it incorporates all the scenarios, at $\alpha = 100\%$ it measures the regret with respect to the worst-case scenario. On the other hand, when fixing the risk aversion level, the results demonstrate an increase in the minimal regret as the look-ahead level Δ increases from 0 to 4. This is in line with the fact that as Δ increases, we are gradually relaxing the nonanticipativity constraints imposed on the set of benchmark policies, as suggested in Lemma 2.3.1. Moreover, we can observe that as the risk aversion level reaches 100%, i.e., $\rho(\xi) = \text{CVaR}_1(\xi) = \text{ess sup}(\xi)$, all regret models achieve the same regret, thus confirming Theorem 2.4.1, which stated that all models are equivalent under the essential supremum risk measure. On the opposite side of the graph, one sees that the minimal risk for the regret model with $\Delta = 0$ converges to zero as predicted by Corollary 2.4.1. The optimal regret analysis of instance II is quite similar, and thus omitted for brevity.

Figures 2.2 (b) and 2.2 (c), associated with instances I and II, respectively, present the optimal first-stage total production from the 3 production facilities resulting from the Δ regret problem for $\Delta \in \{0, 1, 2, 3, 4\}$, together with the one offered by the corresponding CVaR minimization problem (2.1). We observe the following: (i) The production plan returned by the CVaR minimization problem stands higher than those achieved by Δ regret problem for all Δ . One should note that in this multistage inventory management problem, the production facilities must satisfy all the observed demands with the minimum production cost. Any excess inventory at the end of the planning horizon can lead to a higher level of regret, seen as a "lost opportunity". In this sense, starting the production plan with a lower number of units can be interpreted as a "more opportunistic" or "less conservative" approach. The observations in Figures 2.2 (b) and 2.2 (c) reveal that for both instances, CVaR solutions are more conservative than the ones returned by all Δ -regret models for any risk aversion level (α), i.e., CVaR solutions produce more in the first period to take advantage of the lower production costs even if later excess production is lost, while production from Δ -regret models is smaller in the first period as it takes advantage of the recourse decisions to not overproduce, a policy that results in smaller regret even if production cost is higher overall. The result validates the claims made in the literature about the regret minimization criterion producing less conservative solutions; (*ii*) The production levels from the Δ -regret models are roughly ordered based on the value of Δ , however, the order is reverse in the two instances. In instance I, we observe that $\Delta = 0$ produces, in general, the lowest, "least conservative" production and

 $\Delta = 4$ the highest "most conservative" production, while for instance II, $\Delta = 4$ induces, in general, the "least conservative" and $\Delta = 0$ the most "most conservative" production. This behavior is attributed to the difference between the total demand distributions of the two instances and can be interpreted as follows. Since demand needs to be robustly satisfied and production cost is lowest in the first period, if there is not a high production in the first period, cost savings can be made in low demand scenarios if one waits-andsees. In the left-skewed instance II most scenarios result in a high total demand, while some scenarios result in relatively low total demand. To take advantage of the low firststage production costs, a minimum production cost solution might choose to produce as much as possible in the first period, however, such a high production will cause $\Delta = 4$ to experience a very large regret in the low total demand scenarios as the fully anticipative benchmark policy will produce much less in those scenarios. Thus, the 4-regret model chooses to produce less in the first period and take advantage of the recourse decisions. In contrast, in instance I, where most of the total demand scenarios are low, there is a smaller margin of gains to be made by the benchmark policy thus, the $\Delta = 4$ policy is close to the minimum production cost policy produced by CVaR minimization. The reverse is true for $\Delta = 0$. In general, we observe that if there is a potential of a large margin of profit, a large Δ will produce less conservative solutions compared to $\Delta = 0$, and vice versa if the margin of profit is small. This is also observed in the portfolio management example in Section 2.6.11.1.

The runtimes of both instances are provided in Section 2.6.12. For any $\Delta \in \{0, 1, 2, 3\}$, for lower risk aversion levels (α), one can expect a runtime of around 10 seconds; however, as α gradually increases, we observe runtimes of more than 1000 seconds. Once approaching the high risk aversion levels, interestingly, the difficulty of the problem decreases and once again gets somewhat close to the ones it experienced at lower α levels. For $\Delta = 4$, one can show that the Δ -regret problem can simply reduce to a linear program hence resulting in runtimes of under 1 second. This is also true for the CVaR minimization problem (2.1).

The second set of results presents a breakdown of regret as this is expressed in equation (2.5). To this end, using the optimal solutions of $\Delta \in \{0, 2, 4\}$ we evaluate the three expressions in (2.5), which can be interpreted as the "look ahead regret" (2.5a), "risk aver-



Figure 2.3: Regret Breakdown for Different Δ -Regret Optimal Solutions.

sion regret" (2.5b) and "regret of being suboptimal with respect to the \mathcal{MSP} " (2.5c). Figure 2.3 presents the cumulative breakdown of the regret for instance I. As expected, for $\alpha = 0\%$, the "risk aversion regret" is zero for all Δ by definition, while for $\alpha = 100\%$ the "look ahead regret" is zero for all Δ as suggested by Theorem 2.4.1. In fact, the "look ahead regret" decreases as α increases. For $\Delta = 0$, by definition, the "look ahead regret" is zero indicating that regret is a combination of the "risk aversion regret" and "regret of being suboptimal with respect to the \mathcal{MSP} ". It is interesting to observe that the "regret of being suboptimal with respect to the \mathcal{MSP} " is relatively low in the range of $\alpha \in [0, 0.6]$ but constitutes roughly one-third of the total regret when $\alpha = 1$.

Section 2.6.11 includes further numerical studies on a two-stage (Section 2.6.11.1) and multistage (Section 2.6.11.2) portfolio management application. From these experiments, we draw insights similar to those from the inventory management application. In particular, we again observe the less conservative nature of Δ -regret solutions compared to CVaR minimization, as well as, that the value of Δ heavily affects the nature of the solutions produced. In multistage stochastic programming, time consistency refers to the property that a decision policy remains optimal as new information becomes available over time. In Section 2.6.11.2 we demonstrate numerically that the policy produced by Δ -regret is not time consistent with the effect becoming more prominent as the level of risk aversion increases.

2.6 Appendix

2.6.1 Illustrative Example

Table 2.4 provides detailed calculations of obtaining the optimal decision for the project selection problem, described in Example 2.1.1, when the ex-post worst-case regret measure is exploited. Once the decision maker chooses project A ($x = x_A$), she will face either scenario ω_1 or ω_2 , resulting in a payoff of 1\$ or 6\$, respectively. Having full access to the realized scenario, the optimal benchmark decision consists of choosing project B ($x' = x_B$) under ω_1 with 5\$ payoff and picking project A ($x' = x_A$) under ω_2 with 6\$ payoff, leading to a 4\$ regret for the decision maker in the first case and zero regret in the second one; consequently, the worst-case regret of choosing project A will be 4\$. Performing the same analysis for $x = x_B$ and $x = x_C$ brings about the worst-case regrets of 4\$ and 3\$, respectively. As a conclusion, aiming at minimizing the worst-case regret, the decision maker finds project C with 3\$ worst-case regret as her best option.

Table 2.4: EP-WCR

x	ω	$oldsymbol{x}'$	$\{\cdot\}^{\dagger}$	$\sup_{\omega\in\Omega}\{\cdot\}$					
Project A	$\omega_1 \ \omega_2$	Project B Project A	$\begin{array}{c} 4\\ 0\end{array}$	4					
Project B	$\omega_1 \ \omega_2$	Project B Project A	$\begin{array}{c} 0 \\ 4 \end{array}$	4					
$\boldsymbol{x^*} = \operatorname{Project} C$	$\omega_1 \\ \omega_2$	Project B Project A	$\frac{1}{3}$	3					
m x	3								
$\dagger \{\cdot\} \coloneqq \left\{ \max_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}, \boldsymbol{\zeta}(\omega)) \right\}$									

Putting emphasis on risk-aversion, Table 2.5 clarifies the details of getting the optimal decisions for the ex-post worst-case expected regret minimization problem. In this setting, similar to the EP-WCR case, the benchmark policy has full access to the future scenario realizations, and as an immediate result, always selects project B under ω_1 and project A under ω_2 . If the manager picks project A for investment, withdrawing the distributional information, the felt regret consists of 4\$ and 0\$ for ω_1 and ω_2 realizations, respectively. However, in contrast to EP-WCR, in EP-RAR with $\rho(\xi) := \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\xi]$ the expected regret is measured with respect to the worst \mathbb{P} from \mathcal{P} . Since the \mathbb{P}^{I} and \mathbb{P}^{II} lead to expected regrets of 3.2\$ and 0\$, the worst-case expected regret of choosing project A equals 3.2\$. Replicating the same analysis for $x = x_B$ and $x = x_C$ gives rise to worstcase expected regrets of 4\$ and 3\$. As a consequence, investment in project C with 3\$ worst-case expected regret will be the optimal choice.

$oldsymbol{x}$	\mathcal{P}	<i>x</i>	/*	$\mathbb{E}_{\mathbb{P}}\left[\cdot ight]^{\dagger}$	$\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}\left[\cdot\right]$
		ω_1	ω_2		107
Project A	\mathbb{P}^{I} \mathbb{P}^{II}	Project B Project B	Project A Project A	$\begin{array}{c} 3.2 \\ 0 \end{array}$	3.2
Project B	$\frac{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	Project B Project B	Project A Project A	0.8 4	4
$\boldsymbol{x}^* = \operatorname{Project} C$	$\frac{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	Project B Project B	Project A Project A	1.4 3	3
		3			

Table 2.5: EP-RAR

 $\dagger \mathbb{E}_{\mathbb{P}}\left[\cdot\right] := \mathbb{E}_{\mathbb{P}}\left[\max_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}, \boldsymbol{\zeta}(\omega))\right]$

An alternative to EP-RAR consists in EA-RAR, where the benchmark no longer has access to the realized scenario and only knows the true distribution. Table 2.6 summarizes the analysis for this problem. To be consistent with the previous analysis, once again, we elaborate on the details of getting the worst-case expected regret of choosing project A. In this case, the immediate payoff under ω_1 and ω_2 will be 1\$ and 6\$, respectively. Subsequently, the benchmark decision can be made after evaluating the expected regret of each of the three possible options ($x' = x_A$, x_B or x_C) under \mathbb{P}^I and \mathbb{P}^{II} and picking the one which maximizes the expected regret of decision maker's choice ($x = x_A$). Looking at Table 2.6, one remarks six expected values for $x = x_A$, representing these six settings. For instance, if $x' = x_C$ and $\mathbb{P} = \mathbb{P}^I$, the corresponding expected regret of $x = x_A$ can be derived as $\mathbb{E}_{\mathbb{P}} [h(x', \zeta) - h(x, \zeta)] = 0.8(4-1) + 0.2(3-6) = 1.8$ \$. The maximum expected regret among these six values is 2.4\$, which is associated with $x' = x_B$ under $\mathbb{P} = \mathbb{P}^I$.

Performing the same analysis for $x = x_B$ and $x = x_C$ guides the manager towards investing in the project with minimum worst-case expected regret; more specifically, the minimum value in the last column of this table is 2.4\$, indicating that the best choice for the manager is to choose project A for investment with worst-case expected regret of 2.4\$.

\boldsymbol{x}	x'	₽	$\mathbb{E}_{\mathbb{P}}\left[\cdot ight]^{\dagger}$	$\sup_{\mathbb{P}\in\mathcal{P}}\mathbb{E}_{\mathbb{P}}\left[\cdot\right]$	$\max_{\boldsymbol{x}' \in \mathcal{X}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}\left[\cdot\right]$	
	Project A	$\overset{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	0 0	0		
$\boldsymbol{x}^* = \operatorname{Project} A$	Project B	\mathbb{P}^{I} \mathbb{P}^{II}	$2.4 \\ -4$	2.4	2.4	
	Project C	$\overset{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	$1.8 \\ -3$	1.8		
	Project A	$\mathbb{P}^{I} \\ \mathbb{P}^{II}$	-2.4 4	4		
Project B	Project B	$\frac{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	0 0	0	4	
	Project C	$\frac{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	$^{-0.6}_{1}$	1		
	Project A	$\overset{\mathbb{P}^{I}}{\mathbb{P}^{II}}$	$^{-1.8}_{3}$	3		
Project C	Project B	\mathbb{P}^{I} \mathbb{P}^{II}	$0.6 \\ -1$	0.6	3	
	Project C	$\mathbb{P}^{I} \\ \mathbb{P}^{II}$	0 0	0		
	2.4					

Table 2.6: EA-RAR

 $\dagger \mathbb{E}_{\mathbb{P}}\left[\cdot
ight] := \mathbb{E}_{\mathbb{P}}\left[h(oldsymbol{x}',oldsymbol{\zeta}) - h(oldsymbol{x},oldsymbol{\zeta})
ight]$

This is in contrast with the recommended option of project C coming from EP-WCR and EP-RAR problems.

2.6.2 **Proof of Lemma 2.3.1**

Clearly when $\Delta = 0$, we have that $\mathcal{X}_{\Delta} = \mathcal{X}_0 = \mathcal{X}_{na}$. So that the 0-regret model reduces to the *ex-ante* form:

$$\min_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{na}} \rho(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})) \,,$$

and in particular to

$$\underset{\boldsymbol{x} \in \mathcal{X}}{\text{minimize}} \quad \max_{\boldsymbol{x}' \in \mathcal{X}} \rho(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})) \,,$$

when dealing with a static linear problem, i.e. T = 2 and $c_2 = a_{j2} = 0$ for all j.

Alternatively, when $\Delta \geq T - 1$, by definition we have that $\mathcal{X}_{\Delta} = \{ \boldsymbol{x} : \Omega \to \mathbb{R}^{n \times T} \}$, implying that:

$$\mathcal{R}_{\Delta}(oldsymbol{x}) = \max_{oldsymbol{x}'\in\mathcal{X}}
hoig(h(oldsymbol{x}',oldsymbol{\zeta}) - h(oldsymbol{x},oldsymbol{\zeta})ig)$$

$$= \rho \Big(\max_{\boldsymbol{x}' \in \mathcal{X}_{\omega}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)) \Big),$$

which follows from monotonicity of ρ . Specifically, we first have that for all $x' \in \mathcal{X} \cap \mathcal{X}_T$:

$$h(\boldsymbol{x}'(\omega),\boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega)) \leq \max_{\boldsymbol{x}''\in\mathcal{X}_{\omega}} h(\boldsymbol{x}'',\boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega)) \quad \forall \omega\in\Omega.$$

Hence,

$$\rho(h(\boldsymbol{x}',\boldsymbol{\zeta})-h(\boldsymbol{x},\boldsymbol{\zeta})) \leq \rho(\max_{\boldsymbol{x}''\in\mathcal{X}_{\omega}}h(\boldsymbol{x}'',\boldsymbol{\zeta}(\omega))-h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega)))\,.$$

On the other hand, we can define $\bar{x}'(\omega) \in \arg \max_{x' \in \mathcal{X}_{\omega}} h(x', \zeta(\omega)) - h(x(\omega), \zeta(\omega))$, with $\bar{x}' \in \mathcal{X} \cap \mathcal{X}_T$ to conclude that:

$$\begin{split} \rho\Big(\max_{\boldsymbol{x}'\in\mathcal{X}_{\omega}}h(\boldsymbol{x}',\boldsymbol{\zeta}(\omega))-h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega))\Big) &= \rho\Big(h(\bar{\boldsymbol{x}}',\boldsymbol{\zeta})-h(\boldsymbol{x},\boldsymbol{\zeta})\Big) \\ &\leq \max_{\boldsymbol{x}'\in\mathcal{X}\cap\mathcal{X}_{T}}\rho\Big(h(\boldsymbol{x}',\boldsymbol{\zeta})-h(\boldsymbol{x},\boldsymbol{\zeta})\Big) \\ &\leq \max_{\boldsymbol{x}'\in\mathcal{X}\cap\mathcal{X}_{T}}\rho\Big(\max_{\boldsymbol{x}''\in\mathcal{X}_{\omega}}h(\boldsymbol{x}'',\boldsymbol{\zeta}(\omega))-h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega))\Big) \\ &= \rho\Big(\max_{\boldsymbol{x}''\in\mathcal{X}_{\omega}}h(\boldsymbol{x}'',\boldsymbol{\zeta}(\omega))-h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega))\Big), \end{split}$$

where the first inequality follows since $\bar{x}' \in \mathcal{X}_T$. We can therefore conclude that the *T*-regret model reduces to the *ex-post* model:

$$\min_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \quad \rho(\max_{\boldsymbol{x}' \in \mathcal{X}_{\omega}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})) \,,$$

which takes the following form when the problem is static:

$$\min_{\boldsymbol{x} \in \mathcal{X}} \min_{\boldsymbol{x}' \in \mathcal{X}_{\omega}} h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})) \, .$$

Finally, we turn to establishing the monotonicity of the optimal value of Problem (2.3). Let $\Delta \leq \Delta'$, then $\mathcal{X}_{\Delta} \subseteq \mathcal{X}_{\Delta'}$. This implies that:

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \rho(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})) \leq \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta'}} \rho(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})) = \mathcal{R}_{\Delta'}(\boldsymbol{x}).$$

2.6.3 **Proof of Lemma 2.4.1**

Let \overline{MSP} be exactly the same as MSP except for the set of physical constraints that must be satisfied by the policy under each outcome ω , denoted by \mathcal{X}_{ω} , which is replaced with:

$$\bar{\mathcal{X}}_{\omega} := \cap_{t=1}^T \cap_{\omega': \zeta^{[t-1]}(\omega) = \zeta^{[t-1]}(\omega')} \left\{ \boldsymbol{x} \in \mathbb{R}^{n \times T} | \boldsymbol{x}_{1:t} \in \mathcal{X}_{\omega'}^{[t]} \right\}.$$

First, we can start by demonstrating that \overline{MSP} satisfies the relatively complete recourse property. Namely, for all t, if (ω, ω') is such that $\zeta^{[t-1]}(\omega) = \zeta^{[t-1]}(\omega')$, then:

$$\begin{split} \bar{\mathcal{X}}_{\omega}^{[t]} &= \{ \boldsymbol{x} \in \mathbb{R}^{n \times t} | \exists \bar{\boldsymbol{x}} \in \mathbb{R}^{n \times T-t}, [\boldsymbol{x} \ \bar{\boldsymbol{x}}] \in \bar{\mathcal{X}}_{\omega} \} \\ &= \{ \boldsymbol{x} \in \mathbb{R}^{n \times t} | \exists \bar{\boldsymbol{x}} \in \mathbb{R}^{n \times T-t}, [\boldsymbol{x} \ \bar{\boldsymbol{x}}] \in \cap_{t=1}^{T} \cap_{\omega'': \zeta^{[t-1]}(\omega) = \zeta^{[t-1]}(\omega'')} \{ \boldsymbol{x} \in \mathbb{R}^{n \times T} | \boldsymbol{x}_{1:t} \in \mathcal{X}_{\omega''}^{[t]} \} \} \\ &= \{ \boldsymbol{x} \in \mathbb{R}^{n \times t} | \exists \bar{\boldsymbol{x}} \in \mathbb{R}^{n \times T-t}, [\boldsymbol{x} \ \bar{\boldsymbol{x}}] \in \cap_{t=1}^{T} \cap_{\omega'': \zeta^{[t-1]}(\omega') = \zeta^{[t-1]}(\omega'')} \{ \boldsymbol{x} \in \mathbb{R}^{n \times T} | \boldsymbol{x}_{1:t} \in \mathcal{X}_{\omega''}^{[t]} \} \} \\ &= \{ \boldsymbol{x} \in \mathbb{R}^{n \times t} | \exists \bar{\boldsymbol{x}} \in \mathbb{R}^{n \times T-t}, [\boldsymbol{x} \ \bar{\boldsymbol{x}}] \in \bar{\mathcal{X}}_{\omega'} \} = \bar{\mathcal{X}}_{\omega'} \end{split}$$

Next, we show that \overline{MSP} produces the same set of optimal solutions and optimal value as MSP. In particular, one can show that $\mathcal{X} \cap \mathcal{X}_{na} = \overline{\mathcal{X}} \cap \mathcal{X}_{na}$ where $\overline{\mathcal{X}} := \{ \boldsymbol{x} : \Omega \to \mathbb{R}^{n \times T} \mid \boldsymbol{x}(\omega) \in \overline{\mathcal{X}}_{\omega}, \, \omega \in \Omega \}$. First, since we have that:

$$ar{\mathcal{X}}_\omega \subseteq \cap_{\omega': \zeta^{[T-1]}(\omega) = \zeta^{[T-1]}(\omega')} \{ oldsymbol{x} \in \mathbb{R}^{n imes T} | oldsymbol{x}_{1:T} \in \mathcal{X}_{\omega'}^{[T]} \} \subseteq \mathcal{X}_\omega \, ,$$

we can conclude that $\mathcal{X} \cap \mathcal{X}_{na} \supseteq \overline{\mathcal{X}} \cap \mathcal{X}_{na}$. Alternatively, we have that for all $x \in \mathcal{X} \cap \mathcal{X}_{na}$, one can confirm that $x \in \overline{\mathcal{X}}$, i.e. $x(\omega) \in \overline{\mathcal{X}}_{\omega}$ for all ω . Specifically, fixing any ω , any t, and any ω' that satisfies $\zeta^{[T-1]}(\omega) = \zeta^{[T-1]}(\omega')$, we can check that $x_{1:t}(\omega) \in \mathcal{X}_{\omega'}^{[t]}$ since

$$[\boldsymbol{x}_{1:t}(\omega) \ \boldsymbol{x}_{t+1:T}(\omega')] = [\boldsymbol{x}_{1:t}(\omega') \ \boldsymbol{x}_{t+1:T}(\omega')] \in \mathcal{X}_{\omega'},$$

where we used the fact that $x \in \mathcal{X}_{na}$ which implies that $x_{1:t}(\omega) = x_{1:t}(\omega')$.

2.6.4 Proof of Theorem 2.4.1

The argument goes as follows:

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \sup_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \operatorname{ess\,sup}(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}))$$
(2.11a)

$$\leq \sup_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{T-1}} \operatorname{ess\,sup}(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}))$$
(2.11b)

$$\leq \sup_{\boldsymbol{x}' \in \mathcal{X}} \operatorname{ess\,sup}(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}))$$
(2.11c)

$$\leq \operatorname{ess\,sup}(\sup_{\boldsymbol{x}' \in \mathcal{X}_{\omega}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)))$$
(2.11d)

$$= \max_{\omega \in \Omega} \sup_{\boldsymbol{x}' \in \mathcal{X}_{\omega}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega))$$
(2.11e)

$$= \sup_{\boldsymbol{x}' \in \mathcal{X}_{\omega^*}} h(\boldsymbol{x}', \boldsymbol{\zeta}(\omega^*)) - h(\boldsymbol{x}(\omega^*), \boldsymbol{\zeta}(\omega^*))$$
(2.11f)

$$\leq \sup_{\boldsymbol{x}' \in \mathcal{X}_{\omega^*}} \operatorname{ess\,sup}(h(\pi(\omega; \boldsymbol{x}', \omega^*), \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)))$$
(2.11g)

$$\leq \sup_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{na}} \operatorname{ess\,sup}(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}))$$
(2.11h)

$$\leq \sup_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \operatorname{ess\,sup}(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}))$$
(2.11i)

$$= \mathcal{R}_{\Delta}(\boldsymbol{x}), \qquad (2.11j)$$

where (2.11b) follows from Lemma 2.3.1 and (2.11c) results from relaxing the constraint that $\mathbf{x}' \in \mathcal{X}_{T-1}$. (2.11d) follows from the monotonicity of ess sup and the fact that $h(\mathbf{x}'(\omega), \boldsymbol{\zeta}(\omega)) - h(\mathbf{x}(\omega), \boldsymbol{\zeta}(\omega)) \leq \sup_{\mathbf{x}' \in \mathcal{X}_{\omega}} h(\mathbf{x}', \boldsymbol{\zeta}(\omega)) - h(\mathbf{x}(\omega), \boldsymbol{\zeta}(\omega))$, almost surely. (2.11e) follows from the fact that $\mathbb{Q}(\omega) \geq 0$ for all $\omega \in \Omega$ and Ω is finite. In (2.11f) we define ω^* as any maximizer of (2.11e). In (2.11g), we let $\pi(\cdot; \mathbf{x}', \omega^*)$ be a nonanticipative policy which implements \mathbf{x}' under outcome ω^* while implementing an arbitrarily chosen feasible action at each time point for all other outcomes, e.g.:

$$\pi_t(\omega; \boldsymbol{x}', \omega^*) := \begin{cases} \boldsymbol{x}'_t & \text{if } \zeta(\omega)^{[t-1]} = \zeta(\omega^*)^{[t-1]} \\ \arg \min_{\boldsymbol{\bar{x}}_t : [\pi_{[t-1]}(\omega; \boldsymbol{x}', \omega^*)^\top \ \boldsymbol{\bar{x}}_t]^\top \in \mathcal{X}_{\omega}^{[t]}} \| \boldsymbol{\bar{x}}_t \|_2 & \text{otherwise} \end{cases} \quad \forall t.$$

$$(2.12)$$

The fact that this policy exists and is in $\mathcal{X} \cap \mathcal{X}_{na}$ is due to Assumption 2.4.1 (proof below) and motivates (2.11h). Finally, (2.11i) follows from the fact that $\mathcal{X}_{na} \subseteq \mathcal{X}_{\Delta}$.

We finalize this proof by providing more details about the three facts regarding $\pi_t(\omega; x', \omega^*)$. First, this policy exists since we can construct it from t = 1, ..., T with the guarantee that the arg min in (2.12) is non-empty given that for all t and all $\omega \in \Omega$:

$$\boldsymbol{\zeta}^{[t-1]}(\omega) = \boldsymbol{\zeta}^{[t-1]}(\omega^*) \Rightarrow \boldsymbol{\zeta}^{[t'-1]}(\omega) = \boldsymbol{\zeta}^{[t'-1]}(\omega^*), \ \forall 1 \le t' \le t \Rightarrow \pi_{[t]}(\omega; \boldsymbol{x}', \omega^*) \in \mathcal{X}_{\omega}^{[t]},$$

while iteratively, from t = 2 to t = T, and for all $\omega \in \Omega$:

$$\begin{aligned} \pi_{[t-1]}(\omega; \boldsymbol{x}', \omega^*) &\in \mathcal{X}_{\omega}^{[t-1]} \cap \boldsymbol{\zeta}^{[t-1]}(\omega) \neq \boldsymbol{\zeta}^{[t-1]}(\omega^*) \\ \Rightarrow \exists \bar{\boldsymbol{x}} \in \mathbb{R}^{n \times T - t + 1}, [\pi_{[t-1]}(\omega; \boldsymbol{x}', \omega^*) \ \bar{\boldsymbol{x}}] \in \mathcal{X}_{\omega} \\ \Rightarrow \exists \bar{\boldsymbol{x}}_t \in \mathbb{R}^{n \times 1}, \ [\pi_{[t-1]}(\omega; \boldsymbol{x}', \omega^*) \ \bar{\boldsymbol{x}}_t] \in \mathcal{X}_{\omega}^{[t]} \\ \Rightarrow \arg \min_{\bar{\boldsymbol{x}}_t: [\pi_{[t-1]}(\omega; \boldsymbol{x}', \omega^*)^\top \ \bar{\boldsymbol{x}}_t]^\top \in \mathcal{X}_{\omega}^{[t]}} \| \bar{\boldsymbol{x}}_t \|_2 \in \mathcal{X}_{\omega}^{[t]} \\ \Rightarrow \pi_{[t]}(\omega; \boldsymbol{x}', \omega^*) \in \mathcal{X}_{\omega}^{[t]}, \end{aligned}$$

where we first employ the definition of $\pi_{[t-1]}(\omega; \boldsymbol{x}', \omega^*) \in \mathcal{X}_{\omega}^{[t-1]}$, and then confirmed that the first vector of the $\bar{\boldsymbol{x}}$ matrix could be used to create a member of $\mathcal{X}_{\omega}^{[t]}$.

Now regarding $\pi(\cdot; \boldsymbol{x}', \omega^*) \in \mathcal{X}$, this is necessarily the case as we just showed that $\pi(\cdot; \boldsymbol{x}', \omega^*) \in \mathcal{X}_{\omega}^{[t]}$ for all t and $\omega \in \Omega$. Hence, $\pi(\cdot; \boldsymbol{x}', \omega^*) \in \mathcal{X}_{\omega}^{[T]} = \mathcal{X}_{\omega}$ for all ω . Furthermore, $\pi(\cdot; \boldsymbol{x}', \omega^*) \in \mathcal{X}_{na}$ by construction. Namely, for any t, if $\zeta(\omega)^{[t-1]} = \zeta(\omega')^{[t-1]} = \zeta(\omega^*)^{[t-1]}$, then $\pi_t(\cdot; \boldsymbol{x}', \omega^*) = \boldsymbol{x}'_t$. Alternatively, for any (ω, ω') such that $\zeta(\omega)^{[t-1]} = \zeta(\omega')^{[t-1]} \neq \zeta(\omega^*)^{[t-1]}$, we can exploit the fact that:

$$\boldsymbol{\zeta}^{[t-1]}(\omega) = \boldsymbol{\zeta}^{[t-1]}(\omega') \Rightarrow \boldsymbol{\zeta}^{[t'-1]}(\omega) = \boldsymbol{\zeta}^{[t'-1]}(\omega'), \ \forall 1 \le t' \le t \,,$$

so that iteratively from t' = 2 to t' = t, given that:

$$\pi_{t'-1}(\omega; \boldsymbol{x}', \omega^*) = \pi_{t'-1}(\omega'; \boldsymbol{x}', \omega^*)$$

and that Assumption 2.4.1 implies that $\mathcal{X}_{\omega}^{[t']} = \mathcal{X}_{\omega}^{[t]}$, then necessarily

$$\pi_{t'}(\omega; \boldsymbol{x}', \omega^*) = \arg \min_{\boldsymbol{\bar{x}}_{t'}: [\pi_{[t'-1]}(\omega; \boldsymbol{x}', \omega^*)^\top \ \boldsymbol{\bar{x}}_t]^\top \in \mathcal{X}_{\omega}^{[t']}} \|\boldsymbol{\bar{x}}_{t'}\|_2$$

=
$$\arg \min_{\boldsymbol{\bar{x}}_{t'}: [\pi_{[t'-1]}(\omega'; \boldsymbol{x}', \omega^*)^\top \ \boldsymbol{\bar{x}}_t]^\top \in \mathcal{X}_{\omega'}^{[t']}} \|\boldsymbol{\bar{x}}_{t'}\|_2 = \pi_{t'}(\omega'; \boldsymbol{x}', \omega^*).$$

2.6.5 **Proof of Proposition 2.4.1**

The argument goes as follows:

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta})]$$
(2.13a)

$$= \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}', \boldsymbol{\zeta})] - \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}, \boldsymbol{\zeta})]$$
(2.13b)

$$= g(\Delta) - \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x},\boldsymbol{\zeta})]$$
(2.13c)

with $g(\Delta) := \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}', \boldsymbol{\zeta})]$, and where (2.13b) follows from the linearity of the risk measure. The fact that Δ only affects the constant $g(\Delta)$ in (2.13c) allows us to conclude that the optimal solution sets of Problem (2.3) is unaffected by Δ . We also observe that:

$$\min_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \mathbb{R}_{\Delta}(\boldsymbol{x}) \equiv \max_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}, \boldsymbol{\zeta})] - g(\Delta)$$
(2.14)

hence the Δ -regret model has the same optimal solution set as \mathcal{MSP} when $\rho(\xi) = \mathbb{E}_{\mathbb{Q}}[\xi]$.

2.6.6 Proof of Corollary 2.4.1

This follows from the fact that $\mathcal{X}_0 = \mathcal{X}_{na}$, hence replacing $g(\Delta) := \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{na}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}', \boldsymbol{\zeta})]$ in equation (2.14) leads to:

$$\begin{split} \min_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} & \mathcal{R}_{\Delta}(\boldsymbol{x}) = \max_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}, \boldsymbol{\zeta})] - g(\Delta) \\ & = \max_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}, \boldsymbol{\zeta})] - \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{na}} \mathbb{E}_{\mathbb{Q}}[h(\boldsymbol{x}', \boldsymbol{\zeta})] = 0 \end{split}$$

2.6.7 Proof of Proposition 2.4.2

Taking advantage of worst-case expectation risk measure, Problem (2.3) leads to

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \rho(h(\boldsymbol{x}', \boldsymbol{\zeta}) - h(\boldsymbol{x}, \boldsymbol{\zeta}))$$
$$= \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \max_{\boldsymbol{p} \in \mathcal{D} \cap \mathcal{M}} \sum_{\omega \in \Omega} p_{\omega} h(\boldsymbol{x}'(\omega), \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)).$$

Using epigraph variable r, this can be alternatively rewritten as

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{\Delta}} \min_{r} r$$
s.t.
$$\sum_{\omega \in \Omega} p_{\omega} h(\boldsymbol{x}'(\omega), \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)) \leq r, \ \forall \boldsymbol{p} \in \mathcal{D} \cap \mathcal{M}$$

$$r \in \mathbb{R}.$$
(2.15)

Letting $g(\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{p}) := \sum_{\omega \in \Omega} p_{\omega} h(\boldsymbol{x}'(\omega), \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega))$ and applying Theorem 2 in Ben-Tal et al. (2015), we get:

$$\max_{\boldsymbol{p}\in\mathcal{D}\cap\mathcal{M}} g(\boldsymbol{x},\boldsymbol{x}',\boldsymbol{p}) = \inf_{\boldsymbol{v}} \delta^*(\boldsymbol{v}|\mathcal{D}) - \inf_{\boldsymbol{p}\in\mathcal{M}} \boldsymbol{p}^\top \boldsymbol{v} - g(\boldsymbol{x},\boldsymbol{x}',\boldsymbol{p})$$
(2.16a)

$$= \inf_{\boldsymbol{v}} \delta^*(\boldsymbol{v}|\mathcal{D}) - \inf_{\boldsymbol{p}: \boldsymbol{p} \ge 0, \sum_{\omega \in \Omega} \boldsymbol{p}_{\omega} = 1} \sum_{\omega \in \Omega} \boldsymbol{p}_{\omega} \Big(v_{\omega} - h(\boldsymbol{x}'(\omega), \boldsymbol{\zeta}(\omega)) \quad (2.16b)$$

$$+h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega))\Big)$$

= $\inf_{\boldsymbol{v}} \delta^{*}(\boldsymbol{v}|\mathcal{D}) - \min_{\omega \in \Omega} v_{\omega} - h(\boldsymbol{x}'(\omega),\boldsymbol{\zeta}(\omega)) + h(\boldsymbol{x}(\omega),\boldsymbol{\zeta}(\omega))$ (2.16c)

$$= \inf_{\boldsymbol{v}} \max_{\omega \in \Omega} \, \delta^*(\boldsymbol{v}|\mathcal{D}) - v_\omega + h(\boldsymbol{x}'(\omega), \boldsymbol{\zeta}(\omega)) - h(\boldsymbol{x}(\omega), \boldsymbol{\zeta}(\omega)), \qquad (2.16d)$$

where $v \in \mathbb{R}^{|\Omega|}$ and (2.16c) follows from the fact that searching over worst-case distribution is indeed searching over the worst-case outcome. Plugging this result back into equation (2.15) leads to the two-stage optimization model presented in (2.6).

2.6.8 Proof of Corollary 2.4.2

We define

$$\mathcal{D}' := \{ oldsymbol{p}' \in \mathbb{R}^{|\Omega| + n_q} \, | \, \exists \, oldsymbol{p} \in \mathbb{R}^{|\Omega|}, \, oldsymbol{q} \in \mathbb{R}^{n_q}, oldsymbol{p}' = [oldsymbol{p}^ op oldsymbol{q}^ op]^ op, \, Boldsymbol{p}' \leq oldsymbol{b} \}$$

where $B := [B_p \ B_q]$ so that $\mathcal{D} := \{ p \in \mathbb{R}^{n_p} | \exists p' \in \mathcal{D}', p = Ap' \}$, where $A := [I \ 0]$. Since \mathcal{D} is an affine projection \mathcal{D}' , we have that

$$\delta^*(\boldsymbol{v}|\mathcal{D}) = \sup_{\boldsymbol{p}':B\boldsymbol{p}' \leq \boldsymbol{b}} \boldsymbol{v}^\top A \boldsymbol{p}' = \inf_{\boldsymbol{\lambda} \geq 0: A^\top \boldsymbol{v} = B^\top \boldsymbol{\lambda}} \boldsymbol{b}^\top \boldsymbol{\lambda}$$

where we exploited strong LP duality theory, given that D, and implicitly D', is nonempty. After replacing *B* by using its definition and reintegrating the infimum operation in constraint (2.6b) we get Problem (2.7).

2.6.9 **Proof of Proposition 2.4.3**

We start with a definition of the NP-complete 3-SAT problem.

3-SAT problem: Let *W* be a collection of disjunctive clauses $W = \{w_1, w_2, ..., w_N\}$ on a finite set of variables $V = \{v_1, v_2, ..., v_m\}$ such that $|w_i| = 3 \forall i \in \{1, ..., N\}$. Let each clause be of the form $w = v_i \lor v_j \lor \bar{v}_k$, where \bar{v} is the negation of *v*. Is there a truth assignment for *V* that satisfies all the clauses in *W*?

Consider a 3-SAT problem with N clauses of the type $v_i \lor v_j \lor \bar{v}_k$ on m variables. One can construct a multistage stochastic linear program with T = 3 with N branches on the starting node and 3 branches for each node at t = 2. Let the outcome space defined as $\Omega := \{\omega_{ij}\}_{i \in [N], j \in [3]}$. We define two sets of decision variables $x_1 \in [0, 1]^m$ and $x_3 \in \mathbb{R}$, which capture respectively here-and-now and wait-and-see decisions. In the spirit of clause 1 is given by $v_i \vee v_j \vee \bar{v}_k$ we define the sets

$$\begin{aligned} \mathcal{X}_{\omega_{11}} &:= \{ (\boldsymbol{x}_1, x_3) \in [0, 1]^m \times \mathbb{R} | x_3 \le x_{1i} \} \\ \mathcal{X}_{\omega_{12}} &:= \{ (\boldsymbol{x}_1, x_3) \in [0, 1]^m \times \mathbb{R} | x_3 \le x_{1j} \} \\ \mathcal{X}_{\omega_{13}} &:= \{ (\boldsymbol{x}_1, x_3) \in [0, 1]^m \times \mathbb{R} | x_3 \le 1 - x_{1k} \} \end{aligned}$$

with the rest of the clauses expressed though $\mathcal{X}_{\omega_{ij}}$ accordingly where $i \in \{2, ..., N\}$, $j = \{1, 2, 3\}$. For short, we will summarize the definition of \mathcal{X}_{ω} as

$$\mathcal{X}_{\omega_{ij}} := \{ (\boldsymbol{x}_1, x_3) \in [0, 1]^N \times \mathbb{R} | x_3 \le \boldsymbol{a}(\omega_{ij})^\top \boldsymbol{x}_1 + b(\omega_{ij}) \}.$$

Let $h(x, \zeta) := x_3$, and define the ambiguity set as

$$\mathcal{P} = \{ \mathbb{P} | \mathbb{P}(\omega_{i1} \cup \omega_{i2} \cup \omega_{i3}) = 1/N, \forall i \in \{1, \dots, N\} \}.$$

For $(\boldsymbol{x}_1, x_3) := (0, 0) \in \mathcal{X} \cap \mathcal{X}_{na}$ and $\Delta = 0$, we have that

$$\begin{aligned} \mathcal{R}_{\Delta}(\boldsymbol{x}) &= \max_{\boldsymbol{x}' \in \mathcal{X} \cap \mathcal{X}_{lag}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\boldsymbol{x}'_{3} - \boldsymbol{x}_{3}] \\ &= \max_{\boldsymbol{x}'_{1} \in [0,1]^{m}} \frac{1}{N} \sum_{i=1}^{N} \sup_{\boldsymbol{p} \in \mathbb{R}^{3}_{+}: \sum_{j} p_{j} = 1} \sum_{j=1}^{3} p_{j} \max_{\boldsymbol{x}'_{3}: \boldsymbol{x}'_{3} \leq \boldsymbol{a}(\omega_{ij})^{\top} \boldsymbol{x}'_{1} + b(\omega_{ij})} \boldsymbol{x}'_{3} \\ &= \max_{\boldsymbol{x}'_{1} \in [0,1]^{m}} \frac{1}{N} \sum_{i=1}^{N} \sup_{\boldsymbol{p} \in \mathbb{R}^{3}_{+}: \sum_{j} p_{j} = 1} \sum_{j=1}^{3} p_{j}(\boldsymbol{a}(\omega_{ij})^{\top} \boldsymbol{x}'_{1} + b(\omega_{ij})) \\ &= \max_{\boldsymbol{x}'_{1} \in [0,1]^{m}} \frac{1}{N} \sum_{i=1}^{N} \max_{j \in \{1,2,3\}} \boldsymbol{a}(\omega_{ij})^{\top} \boldsymbol{x}'_{1} + b(\omega_{ij}) \\ &= \max_{\boldsymbol{x}'_{1} \in [0,1]^{m}} \frac{1}{N} \left(\max(\boldsymbol{x}'_{1i}, \, \boldsymbol{x}'_{1j}, \, 1 - \boldsymbol{x}'_{1k}) + \sum_{i=2}^{N} \max_{j \in \{1,2,3\}} \boldsymbol{a}(\omega_{ij})^{\top} \boldsymbol{x}'_{1} + b(\omega_{ij}) \right). \end{aligned}$$

The second equality follows from the definition of the ambiguity set and the fact that $x_3 = 0$, while the fourth equality follows since for each scenario *i* the supremum reduces to finding the maximum amongst the 3 terms. The last equality demonstrates how the first clause $v_i \vee v_j \vee \bar{v}_k$ will appear in the problem. Notice that each term of the last equation is between 0 and 1, and reaches 1 if and only if the assignment in x_1 satisfies the associated clause. Given that $\mathcal{R}_0(0) \in [0,1]$ and reaches 1 only if there exists an assignment of x_1 that satisfies all the *N* clauses, we can conclude that verifying whether $\mathcal{R}_0(0) \geq 1$ is equivalent to answering to the 3-SAT problem. This makes the evaluation of $\mathcal{R}_0(x)$ NP-hard, thus the evaluation of $\mathcal{R}_\Delta(x)$ generally NP-hard.

2.6.10 Column-and-Constraint Generation Algorithm

The column-and-constraint generation algorithm, proposed by Zeng and Zhao (2013), is an iterative scheme that optimally solves two-stage linear robust optimization problems with right-hand-side uncertainty. Hence, we can employ it for solving Problem (2.7). Assume that $\mathcal{X} \cap \mathcal{X}_{na}$ is non-empty, and let $\{x'_1, \ldots, x'_K\}$ denote the set of policies that comprise the vertices of $\mathcal{X} \cap \mathcal{X}_{\Delta}$, i.e., $x'_k : \Omega \to \mathbb{R}^{n \times T}$ for all $k \in \{1, \ldots, K\}$. Let $\mathcal{K}' \subseteq \{1, \ldots, K\}$. The column-and-constraint generation algorithm can be viewed as a reduction to the vertex enumeration method, where at each iteration, a vertex is added to the following master problem

$$\mathcal{M}(\mathcal{K}') = \min_{\boldsymbol{x}, s, \{\boldsymbol{v}_k, \boldsymbol{\lambda}_k\}_{\forall k \in \mathcal{K}'}} s$$
s.t. $\boldsymbol{\lambda}_k^{\top} \boldsymbol{b} - \boldsymbol{v}_{\omega, k} + \sum_{t=1}^T \boldsymbol{c}_t^{\top}(\boldsymbol{\zeta}(\omega))(\boldsymbol{x}'_{t, k}(\omega) - \boldsymbol{x}_t(\omega))) \leq s,$
 $\forall \omega \in \Omega, \forall k \in \mathcal{K}'$
 $B_p^{\top} \boldsymbol{\lambda}_k = \boldsymbol{v}_k, \quad \forall k \in \mathcal{K}'$
 $B_q^{\top} \boldsymbol{\lambda}_k = 0, \quad \forall k \in \mathcal{K}'$
 $s \in \mathbb{R}, \, \boldsymbol{\lambda}_k \in \mathbb{R}_+^m, \, \boldsymbol{v}_k \in \mathbb{R}^{|\Omega|}, \quad \forall k \in \mathcal{K}'$
 $\boldsymbol{x} \in \mathcal{X} \cap \mathcal{X}_{na}.$

$$(2.17)$$

For any $\mathcal{K}' \subseteq \{1, ..., K\}$, $\mathcal{M}(\mathcal{K}')$ constitutes a lower bound on the optimal value of Problem (2.7). For a given $x \in \mathcal{X} \cap \mathcal{X}_{na}$, we can evaluate $\mathcal{R}_{\Delta}(x)$ through solving the inner maximization of Problem (2.7). Expressing the inner minimization of Problem (2.7) through its KKT conditions and merging it into the outer maximization problem yields the following bilinear optimization program

$$\mathcal{R}_{\Delta}(\boldsymbol{x}) = \max_{\boldsymbol{x}', r, \boldsymbol{v}, \boldsymbol{\lambda}, \boldsymbol{p}, \boldsymbol{q}} \quad r + \boldsymbol{\lambda}^{\top} \boldsymbol{b}$$
(2.18a)

s.t.
$$r + v_{\omega} \ge \sum_{t=1}^{T} \boldsymbol{c}_{t}^{\top}(\boldsymbol{\zeta}(\omega))(\boldsymbol{x}_{t}'(\omega) - \boldsymbol{x}_{t}(\omega)), \ \forall \omega \in \Omega$$
 (2.18b)

$$B_p^{\top} \boldsymbol{\lambda} = \boldsymbol{v} \tag{2.18c}$$

$$B_q^{\top} \boldsymbol{\lambda} = 0 \tag{2.18d}$$

$$\sum_{\omega \in \Omega} p_{\omega} = 1 \tag{2.18e}$$

$$B_p \boldsymbol{p} + B_q \boldsymbol{q} \le \boldsymbol{b} \tag{2.18f}$$

$$p_{\omega}(r + v_{\omega} - \sum_{t=1}^{T} \boldsymbol{c}_{t}^{\top}(\boldsymbol{\zeta}(\omega))(\boldsymbol{x}_{t}'(\omega) - \boldsymbol{x}_{t}(\omega))) = 0, \; \forall \omega \in \Omega \; (2.18g)$$

$$\lambda_i \boldsymbol{e}_i^{\top} (\boldsymbol{b} - B_p \boldsymbol{p} - B_q \boldsymbol{q}) = 0, \; \forall i = 1, 2, ..., m$$
(2.18h)

$$r \in \mathbb{R}, \ \boldsymbol{v} \in \mathbb{R}^{|\Omega|}, \ \boldsymbol{\lambda} \in \mathbb{R}^m_+, \ \boldsymbol{p} \in \mathbb{R}^{|\Omega|}_+, \ \boldsymbol{q} \in \mathbb{R}^{n_q}$$
 (2.18i)

$$x' \in \mathcal{X} \cap \mathcal{X}_{\Delta},$$
 (2.18j)

where $e_i \in \mathbb{R}^m$ is the *i*th column of the identity matrix. Constraints (2.18b)-(2.18d) ensure primal feasibility, (2.18e)-(2.18f) ensure dual feasibility, while the bilinear constraints (2.18g) and (2.18h) ensure complementary slackness. To make the problem amenable to efficient optimization solvers, the bilinear constraints can be linearized using McCormick inequalities, see McCormick (1983). To this end, let M denote a sufficiently large constant, typically referred to as the big-M constant in the integer programming literature. By introducing binary variables $\text{Bin}^p \in \{0,1\}^{|\Omega|}$ and $\text{Bin}^{\lambda} \in \{0,1\}^m$, constraints (2.18g) and (2.18h) can be reformulated as

$$r + v_{\omega} - \sum_{t=1}^{T} \boldsymbol{c}_{t}^{\top}(\boldsymbol{\zeta}(\omega))(\boldsymbol{x}_{t}'(\omega) - \boldsymbol{x}_{t}(\omega)) \leq M \operatorname{Bin}_{\omega}^{\boldsymbol{p}}, \ \forall \omega \in \Omega$$
(2.19a)

$$p \le 1 - \operatorname{Bin}^p \tag{2.19b}$$

$$\boldsymbol{b} - B_p \boldsymbol{p} - B_q \boldsymbol{q} \le M \mathsf{Bin}^{\boldsymbol{\lambda}} \tag{2.19c}$$

$$\lambda \le M(1 - \operatorname{Bin}^{\lambda}). \tag{2.19d}$$

Solving the resulting mixed integer linear program provides an upper bound on the optimal value of Problem (2.7). The optimal worst-case benchmark policy x' of Problem (2.18), can be added to the master problem to further strengthen the lower bound. Algorithm 1 describes the iterative process. The computational efficiency of Algorithm 1 heavily relies on the ability to evaluate efficiently $\mathcal{R}_{\Delta}(x)$ in Step 3. The choice of the big-M constant heavily influences the solution speed, i.e., choosing it too big will result to weak linear relaxation leading in longer computational times. For the special case where the risk measure is the Conditional Value-at-Risk, the matrices in \mathcal{D} reduce to $B_p = I \in \mathbb{R}^{|\Omega| \times |\Omega|}$, $B_q = 0$ and $\mathbf{b} = \bar{\mathbf{p}}/(1 - \alpha)$. Since by construction $\mathbf{p} \in [0, 1]^{|\Omega|}$, then constraint (2.19c) reduces to

$$\frac{\bar{\boldsymbol{p}}}{(1-\alpha)} - \boldsymbol{p} \leq \frac{1}{(1-\alpha)} \text{diag}(\bar{\boldsymbol{p}}) \text{Bin}^{\boldsymbol{\lambda}},$$

Algorithm 1 Column-and-constraint generation algorithm, Zeng and Zhao (2013)

1: Initialize: $lb = -\infty$, $ub = \infty$, $\mathcal{K}' = \emptyset$.

- 2: Solve Problem (2.17) and let x^* be the optimal solution. Set $lb = \mathcal{M}(\mathcal{K}')$;
- Evaluate *R*_Δ(*x*^{*}) by solving Problem (2.18) and let *x*^{'*} be the optimal solution. Set *ub* = *R*_Δ(*x*^{*});
- 4: **if** ub lb > 0 **then**
- 5: $K' = K' \cup \{i\}$ where *i* is the index of x'^* in the set of vertices $\{x'_1, \ldots, x'_K\}$, and go to Step 2
- 6: **else**
- 7: **Return:** x^* and $\mathcal{R}_{\Delta}(x^*)$.
- 8: end if

where diag(\bar{p}) is a diagonal matrix with \bar{p} appearing in the diagonal entries. In other words, the big-M constant can be set to $\bar{p}_{\omega}/(1-\alpha)$ for the ω^{th} constraint.

2.6.11 Portfolio Management

In this section, we provide additional insights on the behavior of Δ -regret using a stylized portfolio management example. We consider an investment horizon of T periods, with $\mathcal{T} := \{1 \dots, T\}$. At each period, the asset manager can either invest in a risk free asset (with the rate r_f) or in a risky asset. The constraints of the problem are as follows:

$$\mathcal{X}_{\omega} := \left\{ \begin{array}{cc} (\boldsymbol{x}^{+}, \boldsymbol{x}^{-}, \boldsymbol{y}) \\ \in \mathbb{R}_{+}^{T} \times \mathbb{R}_{+}^{T} \times \mathbb{R}^{T} \end{array} \middle| \begin{array}{c} \sum_{t'=1}^{t} x_{t'}^{+} - x_{t'}^{-} \ge 0, \quad \forall t \in \mathcal{T} \\ \sum_{t'=1}^{t} (1+r_{f})^{t-t'} y_{t'} \ge 0, \quad \forall t \in \mathcal{T} \\ w_{0} - \sum_{t'=1}^{t} y_{t'} - (x_{t'}^{+} - x_{t'}^{-}) \zeta_{t'-1} + c \zeta_{t'-1} (x_{t'}^{+} + x_{t'}^{-}) \ge 0, \quad \forall t \in \mathcal{T} \setminus T \\ \end{array} \right\},$$
(2.20)

where w_0 is the initial wealth, x_t^+ and x_t^- capture the number of assets purchased and sold at period t (incurring a proportional transaction cost of c), respectively, while y_t captures the amount invested or liquidated from the risk free account. The first constraint ensures that there is no short-selling, the second prevents the possibility of borrowing, and the last constraint ensures that there is enough liquidity to manage the portfolio. Finally, the objective function measures how much wealth can be generated at the end of the horizon:

$$h(\boldsymbol{x}^{+}, \boldsymbol{x}^{-}, \boldsymbol{y}, \boldsymbol{\zeta}) := w_{0} - \sum_{t=1}^{T} y_{t} - \left(\sum_{t=1}^{T} (x_{t}^{+} - x_{t}^{-})\zeta_{t-1} + c\zeta_{t-1}(x_{t}^{+} + x_{t}^{-})\right)$$
(2.21)

In the following, we consider a two-stage and a multistage instance of the problem.

2.6.11.1 Two-Stage Portfolio Management

The two-stage problem can be simplified by noting that at the end of the horizon, all stocks will be sold, and the risk free investment will be liquidated, thus we can also set $x_2^- = x_1^+$ and $y_2 = -(1 + r_f)y_1$. Moreover, without loss of generality, we assume that $\zeta_0 = 1$, thus the feasible set (2.20) and profit function (2.21) respectively simplify to

$$\mathcal{X}_{\omega} := \left\{ \begin{array}{c} x_{1}^{+} \ge 0 \\ (x_{1}^{+}, y_{1}) \in \mathbb{R}_{+} \times \mathbb{R} | & y_{1} \ge 0 \\ & w_{0} - y_{1} - (1 + c)x_{1}^{+} \ge 0 \end{array} \right\}$$

with

$$h(y_1, x_1^+, \zeta) = w_0 + y_1 r_f + ((1-c)\zeta_1 - (1+c))x_1^+.$$

Given that $r_f > 0$, the initial wealth is necessarily distributed among the risk free asset and the risky asset, thus, we have that $x_1^+ = (w_0 - y_1)/(1 + c)$. This allows to eliminate x_1^+ and express the investment decisions solely through the risk free investment y_1 .

In the following experiments, the random price of the risky asset is represented via 10000 scenarios taken from a normal distribution with mean $\mu = 1.02$ and standard deviation $\sigma = 0.05$. We set c = 0 and the initial wealth $w_0 = 100$. Using $\rho(\xi) = \text{CVaR}_{\alpha}(\xi)$, we solve Problem (2.3) for $\Delta = 0$ and $\Delta = 1$, as well as the corresponding CVaR minimization Problem (2.1).



Figure 2.4: Optimal Regret for Different Values of Δ , α and r_f .

To compare the investment decisions from the three performance measures, we consider instances where $r_f \in \{0.01, 0.02, 0.03\}$. Figure 2.4 plots the optimal regret, and Figure 2.5 plots the optimal investment in the risk free asset as a function of α , with the



Figure 2.5: Optimal Investment Decisions in the Risk Free Asset for Δ -Regret and CVaR Minimization.

investment in the risky asset given by $x_1^+ = (w_0 - y_1)/(1 + c)$. We observe the following: (*i*) For $\alpha = 0$, all models reduce to minimizing expected loss, thus all models produce the same optimal solution, as suggested by Proposition 2.4.1, despite that the optimal regret is different. When $r_f < (1 - c)\mathbb{E}(\zeta_1) - (1 + c)$ all models invest the initial wealth to the risky asset while if $r_f > (1 - c)\mathbb{E}(\zeta_1) - (1 + c)$ to the risk free asset. (*ii*) For $\alpha = 1$, both $\Delta = 0$ and $\Delta = 1$ produce the same optimal investments. This is not surprising as the risk measure reduces to $\rho(\xi) = \text{CVaR}_1(\xi) = \text{ess sup}(\xi)$ and from Theorem 2.4.1 we know that both regret models will achieve the same optimal regret, see Figure 2.4. Note that in general, the optimal solution might not be unique, hence it is possible that the two regret models produce different optimal solutions, however, this is not the case in this one-dimensional example. (*iii*) Although all models result in the same investment decisions for $\alpha = 0$, as the value of α increases, the CVaR minimization invests all of the wealth in the risk free asset. In contrast, Δ -regret models prefer to invest only a fraction of the wealth in the risk free asset.

If we interpret the behavior of the decision maker to be "less conservative" if the decision maker invests in the risky asset, then the results provide further evidence that regret minimization models provide less conservative solutions than CVaR minimization. Similar behavior was observed when other distributions were considered for the risky asset, such as the uniform distribution, and the symmetric and skewed Beta distributions (not presented in this chapter). (*iv*) The investment decision for both Δ models is the same for $r_f = 0.02$ since $r_f = ((1 - c)\mathbb{E}(\zeta_1) - (1 + c))$. For $r_f = 0.01$, $\Delta = 1$ invests

more in the risky asset while the reverse is true for the $r_f = 0.03$ case. This can be intuitively interpreted as follows. When $\mathbb{P}(((1-c)\zeta_1 - (1+c)) \ge r_f) > 0.5$, i.e., there is significant chance to increase profits by investing in the risky asset, to avoid large regret as α increases, the $\Delta = 1$ policy will invest more on the risky asset, while when $\mathbb{P}(((1-c)\zeta_1 - (1+c)) \ge r_f) < 0.5$ the reverse happens leading to more conservative decisions with more investment in the risk free asset. $\Delta = 0$ on the other hand, produces slightly more conservative investments than $\Delta = 1$ when $\mathbb{P}(((1-c)\zeta_1 - (1+c)) \ge r_f) > 0.5$ and less conservative when $\mathbb{P}(((1-c)\zeta_1 - (1+c)) \ge r_f) < 0.5$. This can be partly explained as the benchmark investment is chosen before the price of the risky asset realizes. (v) It is interesting to observe that for $r_f = 0.01$ and $r_f = 0.03$ both Δ -regret models have respectively the same optimal regret. This is due to the fact that the risky asset follows a symmetric distribution with $|r_f - \mathbb{E}(\zeta_1)| = 0.01$ for both $r_f = 0.01$ and $r_f = 0.03$.

We conclude this experiment by noting that the behavior of the Δ -regret solutions are different based on the situation, with the solution of $\Delta = 1$ being opportunistic in situations where there is potential for large gains, while $\Delta = 0$ is less conservative in situations where there is potential for gains but the margins are smaller. The solutions are also significantly different from CVaR minimization, which typically provides more conservative solutions.

2.6.11.2 Multistage Portfolio Management

We now turn our attention to the multistage variant of the problem, where we consider five trading periods. Similar to the multistage inventory management problem, we assume that the outcome space includes 16 scenarios ($|\Omega| = 16$). The scenario tree that describes the price evolution of the risky asset is a binomial tree. At each node of the tree, the risky asset's price may increase to u times the current price with the probability of por decrease to d times the current price with the chances of (1 - p), where u = 1/d. We assume p = 0.8 and u = 1.02, which leads to price realizations ζ_t . We set the risk free account rate $r_f = 0.0001$, the initial wealth $w_0 = 100$, the initial price of the risky asset $\zeta_0 = 1$ and assume two different settings for the transaction cost $c \in \{0.02, 0.025\}$. We consider $\rho(\xi) = \text{CVaR}_{\alpha}(\xi)$ and solve Problem (2.3) for $\Delta \in \{0, 1, 2, 3, 4\}$ as well as the corresponding CVaR minimization Problem (2.1). The experiments are conducted over the risk aversion levels of $\alpha = n/16$, where $n \in \{0, 1, ..., 16\}$. A sketch of the scenario tree, as well as the price realizations under each scenario, are depicted in Section 2.6.12.



Figure 2.6: Regret and First-stage Investment in the Risk Free Account Given Different Levels of Conservatism (α).

Figure (a) depicts the optimal Δ -regret and the regret of a shrinking-horizon (SH) policy for c = 0.025, (b) depicts the first-stage investment in the risk free account for c = 0.025 and (c) the first-stage investment in the risk free Account for c = 0.02.

The first experiment compares the optimal Δ -regret values and optimal investment in the first period for different values of $\Delta \in \{0, 1, 2, 3, 4\}$ and levels of risk aversion α . The optimal Δ -regret for c = 0.025 are depicted in Figures 2.6 (a) (solid lines), and the conclusions are qualitatively similar as in the inventory management experiment, see Figures 2.2 (a), with the only difference being that $\Delta = 3$ and $\Delta = 4$ achieve the same level of regret for all values of α . Figures 2.6 (b) and (c) depict the optimal first-stage investment in the risk free asset when the transaction cost is c = 0.025 and c = 0.02, respectively. As before, the figures depict the solutions for the Δ -regret problem for $\Delta \in \{0, 1, 2, 3, 4\}$, together with the investment corresponding to the CVaR minimization Problem (2.1). We make the following observations: (i) The investment produced by CVaR minimization favors the risk free investment in both high and low transaction cost regimes, especially as α increases. In contrast, Δ -regret solutions, in general, invest much more in risky asset. If we interpret a higher investment in the risk free asset as being "conservative", we again reach the conclusion that CVaR minimization is significantly more conservative than Δ regret, which is more risk-seeking, especially for $\alpha > 0.2$. This observation is in line with the numerical evidence from the inventory management example in Section 2.5 as well as the two-stage portfolio management example in Section 2.6.11.1. (ii) In the low transaction cost regime, and for small values of α , all models favor the risky asset as trading is less costly. As α increases, all Δ -regret models invest more in the risk free asset; however, they do that at a much lower rate than the solution returned by CVaR minimization. (*iii*) In the high transaction cost regime, the CVaR minimization invests all the wealth in the risk free asset for all risk aversion levels α , while Δ -regret chooses to also invest in the risky asset, especially for $\alpha > 0.2$. (*iv*) Unlike the inventory management and the two-stage portfolio examples where we can see an ordering in the levels of the solutions as Δ increases, in this example, we cannot observe a clear, interpretable pattern in the behavior of the first stage investments. Nevertheless, we observe that different values of Δ return quite different first-stage portfolio investments.

In our second experiment, we examine the time consistency properties of Δ -regret using c = 0.025. To this end, for each Δ , we solve the problem in a shrinking-horizon manner, i.e., for each node of the scenario tree, we fix the policy for all preceding nodes and re-optimize for the remaining horizon. Figure 2.6 (a) reports the optimal regrets obtained from the Δ -regret model (solid lines) and the optimal regrets resulting from solving the problem using shrinking-horizon (dotted lines). We make the following observations: (*i*) For $\alpha > 0$, Δ -regret is in general not time consistent, as we observe empirically that the value of regret differs in the static versus shrinking-horizon evaluation. This is to be expected as risk-averse multistage stochastic programs are not considered time-consistent, except in certain cases, e.g., exploiting expected value or so-called nested dynamic risk measures, see Shapiro 2009, Ruszczyński 2010, and Pichler et al. 2022. (*ii*) For $\alpha = 0$, Δ regret is time consistent since the CVaR reduces to the expected value risk measure, and by Proposition 2.4.1 the set of the optimal solution of Δ -regret coincides with the set of optimal solutions of MSP. (*iii*) As α increases, we observe an increase in the difference of the regret achieved by the static versus the shrinking-horizon implementation, with the largest difference occurring when $\alpha = 1$.

We summarize this section by noting that Δ -regret models produce less conservative solutions than CVaR minimization, which mirrors the results from previous sections. In addition, we verify numerically that Δ -regret is, in general, not time-consistent except when $\rho(\xi) = \mathbb{E}_{\mathcal{Q}}(\xi)$ in which case all Δ -regret models reduce to \mathcal{MSP} which is known to be time-consistent.

2.6.12 Parameters of Solved Instances

		ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ω_7	ω_8	ω_9	ω_{10}	ω_{11}	ω_{12}	ω_{13}	ω_{14}	ω_{15}	ω_{16}
	1	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15	15
	2	0	0	0	0	0	0	0	0	4	4	4	4	4	4	4	4
t	3	6	6	6	6	3	3	3	3	11	11	11	11	12	12	12	12
	4	10	10	6	6	13	13	4	4	0	0	15	15	9	9	14	14
	5	4	1	7	14	2	4	9	12	15	2	12	8	13	4	11	3
То	tal	35	32	34	41	33	35	31	34	45	32	57	53	53	44	56	48
Instance I																	
		ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ω_7	ω_8	ω_9	ω_{10}	ω_{11}	ω_{12}	ω_{13}	ω_{14}	ω_{15}	ω_{16}
	1	ω_1	ω_2 15	ω_3 15	ω_4 15	ω_5 15	ω_6 15	ω ₇ 15	ω_8 15	ω_9 15	ω_{10} 15	$\frac{\omega_{11}}{15}$	$\frac{\omega_{12}}{15}$	ω_{13} 15	$\frac{\omega_{14}}{15}$	ω_{15} 15	$\frac{\omega_{16}}{15}$
	1 2	$\begin{array}{c} \omega_1 \\ 15 \\ 15 \end{array}$	$\frac{\omega_2}{15}$			$\frac{\omega_5}{15}$	$\frac{\omega_6}{15}$	ω ₇ 15 15	$\frac{\omega_8}{15}$	$\frac{\omega_9}{15}$	$\frac{\omega_{10}}{15}$	$\frac{\omega_{11}}{15}$	ω_{12} 15 11	ω_{13} 15 11	$\frac{\omega_{14}}{15}$	ω_{15} 15 11	$\frac{\omega_{16}}{15}$ 11
t	1 2 3	ω_1 15 15 9	$\frac{\omega_2}{15}$ 15 9	$\frac{\omega_3}{15}$ 15 9	$ \frac{\omega_4}{15} 15 9 $	ω_5 15 15 12	$ \frac{\omega_6}{15} 15 12 $	ω ₇ 15 15 12				ω_{11} 15 11 4	ω_{12} 15 11 4	ω_{13} 15 11 3		ω_{15} 15 11 3	
t	1 2 3 4	$\begin{array}{c c} \omega_1 \\ 15 \\ 15 \\ 9 \\ 5 \end{array}$	$ \frac{\omega_2}{15} 15 9 5 $		ω_4 15 15 9 9	ω_5 15 15 12 2		ω_7 15 15 12 11		$ \frac{\omega_9}{15} 11 4 15 $	ω_{10} 15 11 4 15	ω_{11} 15 11 4 0	ω_{12} 15 11 4 0	ω_{13} 15 11 3 6		ω_{15} 15 11 3 1	$ \frac{\omega_{16}}{15} 11 3 1 $
t	1 2 3 4 5	$\begin{array}{c c} \omega_1 \\ 15 \\ 15 \\ 9 \\ 5 \\ 11 \end{array}$	ω_2 15 15 9 5 14		$ \frac{\omega_4}{15} 15 9 9 1 $	ω_5 15 15 12 2 13		$ $			ω_{10} 15 11 4 15 13	ω_{11} 15 11 4 0 3	ω_{12} 15 11 4 0 7	ω_{13} 15 11 3 6 2	ω_{14} 15 11 3 6 11	ω_{15} 15 11 3 1 4	ω_{16} 15 11 3 1 12
t	1 2 3 4 5	$\begin{array}{c c} \omega_1 \\ 15 \\ 15 \\ 9 \\ 5 \\ 11 \end{array}$	ω_2 15 15 9 5 14	ω_3 15 15 9 9 8		ω_5 15 15 12 2 13		ω_7 15 15 12 11 6	ω_8 15 15 12 11 3		ω_{10} 15 11 4 15 13	ω_{11} 15 11 4 0 3	ω_{12} 15 11 4 0 7	ω_{13} 15 11 3 6 2	ω_{14} 15 11 3 6 11	ω_{15} 15 11 3 1 4	ω_{16} 15 11 3 1 12
t To	1 2 3 4 5 tal	$\begin{array}{c c} \omega_1 \\ 15 \\ 15 \\ 9 \\ 5 \\ 11 \\ 55 \end{array}$	ω_2 15 15 9 5 14 58			ω_5 15 12 2 13 57					ω_{10} 15 11 4 15 13 58	ω_{11} 15 11 4 0 3 33	ω_{12} 15 11 4 0 7 37	ω_{13} 15 11 3 6 2 37	ω_{14} 15 11 3 6 11 46	ω_{15} 15 11 3 1 4 34	ω_{16} 15 11 3 1 12 42

Table 2.7: Demand Realizations of Multistage Inventory Management Problem - $\zeta_{t-1}(\omega)$

Table 2.8: Parameters of Multistage Inventory Management Problem

$oldsymbol{c}_{it}$	i = 1	i = 2	i = 3
t = 1	14	12	20
t = 2	18	16	22
t = 3	24	18	28
t = 4	26	20	30
t = 5	28	24	32
	$\bar{x}_{1t} = 10$ $\bar{x}_{2t} = 10$ $\bar{x}_{3t} = 20$ $\bar{x}_{1,tot}$ $\bar{x}_{2,tot}$ $\bar{x}_{3,tot}$	$\begin{array}{l} \forall t \in \mathcal{T} \\ 0, \ \forall t \in \mathcal{T} \\ 0, \ \forall t \in \mathcal{T} \\ = 20 \\ = 20 \\ = 40 \end{array}$	-
$x_{wh}^{0} =$	= 0, <u>x</u> _{wh}	$=0, \bar{x}_{wh}$	$_{n} = 30$

~			Insta	ince I		Instance II						
ά	$\Delta =$	$\Delta =$	Δ =2	$\Delta =$	$\Delta =$	CVaR	$\Delta =$	CVaR				
	0	1		3	4		0	1	2	3	4	
0%	10.0	9.4	9.3	9.4	1.4	1.3	10.2	9.9	10.0	10.2	1.2	1.2
6.25%	30.6	15.1	15.5	22.4	1.2	1.3	16.4	15.9	16.1	16.5	1.2	1.2
12.5%	73.3	38.0	30.0	39.3	1.3	1.2	31.2	75.8	40.5	32.1	1.2	1.2
18.75%	32.4	79.1	69.1	56.5	1.2	1.3	43.9	55.8	120.2	56.1	1.2	1.2
25%	42.6	140.3	174.0	78.3	1.3	1.2	33.9	213.3	224.5	183.4	1.2	1.2
31.25%	79.4	219.0	224.9	285.9	1.2	1.2	81.5	198.2	254.5	316.9	1.2	1.2
37.5%	188.7	627.4	360.5	276.1	1.2	1.3	178.6	363.7	603.4	207.8	1.2	1.2
43.75%	245.6	795.9	550.3	234.7	1.2	1.3	177.9	558.7	1031.2	379.5	1.2	1.2
50%	209.6	623.9	440.9	119.4	1.2	1.3	93.0	306.8	622.0	154.0	1.2	1.2
56.25%	205.7	262.6	379.8	87.6	1.3	1.3	81.4	332.5	880.1	181.4	1.4	1.2
62.5%	228.7	291.0	271.8	194.5	1.2	1.3	138.0	109.5	668.4	275.9	1.2	1.2
68.75%	92.7	340.4	120.0	33.7	1.3	1.2	95.7	385.9	334.7	131.6	1.2	1.3
75%	130.7	339.4	158.7	60.9	1.2	1.2	70.9	315.5	330.7	94.6	1.2	1.2
81.25%	203.1	232.8	208.3	23.4	1.3	1.3	108.4	387.0	241.9	81.6	1.2	1.2
87.5%	157.8	74.1	121.8	22.7	1.3	1.2	106.2	206.2	77.7	42.9	1.2	1.2
93.75%	48.4	49.0	49.7	40.0	1.2	1.3	50.7	40.7	32.3	24.4	1.2	1.2
100%	48.7	88.3	39.0	30.5	1.4	1.4	51.0	40.3	52.5	42.0	1.2	1.2

Table 2.9: Solution Time for Multistage Inventory Management Problem

Runtime in seconds.

Table 2.10: Price Realizations of Multistage Portfolio Management Problem - $\zeta_{t-1}(\omega)$

		ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ω_7	ω_8	ω_9	ω_{10}	ω_{11}	ω_{12}	ω_{13}	ω_{14}	ω_{15}	ω_{16}
	$\begin{vmatrix} 1 \\ 2 \end{vmatrix}$	1.00	1.00 1.02	1.00 1.02	1.00	1.00 1.02	1.00 1.02	1.00 1.02	1.00 1.02	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
t	3	1.02	1.02	1.02	1.02	1.02	1.02	1.02	1.02	1.00	1.00	1.00	1.00	0.96	0.96	0.96	0.96
	4 5	$1.06 \\ 1.08$	$1.06 \\ 1.04$	1.02 1.04	1.02	1.02 1.04	1.02 1.00	0.98 1.00	0.98 0.96	$1.02 \\ 1.04$	1.02	0.98 1.00	0.98 0.96	0.98 1.00	0.98 0.96	0.94 0.96	0.94 0.92


Endnotes

1. Note that the selected values of \mathbb{P}^{I} and \mathbb{P}^{II} are such that $\max_{\mathbb{P}\in\mathcal{P}} \mathbb{E}_{\mathbb{P}}[X]$ can be reinterpreted as the 75%-conditional value-at-risk of X when using the probability measure $\mathbb{P}(\omega_{1}) = 1 - \mathbb{P}(\omega_{2}) = 20\%$.

2. We note that in the original definition of convex risk measures, *X* is considered a random financial gain.

References

- Aissi, H., Bazgan, C., and Vanderpooten, D. (2009). Min-max and min-max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research*, 197(2):427–438.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
- Audibert, J., Bubeck, S., and Lugosi, G. (2014). Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45.
- Bellini, F. and Bernardino, E. D. (2017). Risk management with expectiles. *The European Journal of Finance*, 23(6):487–506.
- Ben-Tal, A., den Hertog, D., and Vial, J.-P. (2015). Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming*, 149(1-2):265–299.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., and Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376.
- Benati, S. and Conde, E. (2022). A relative robust approach on expected returns with bounded cvar for portfolio selection. *European Journal of Operational Research*, 296(1):332–352.
- Bertsimas, D. and de Ruiter, F. (2016). Duality in two-stage adaptive linear optimization: Faster computation and stronger bounds. *INFORMS Journal on Computing*, 28(3):500–511.

- Bertsimas, D. and Dunning, I. (2020). Relative robust and adaptive optimization. *IN*-FORMS Journal on Computing, 32(2):408–427.
- Bertsimas, D., Gupta, V., and Kallus, N. (2018). Robust sample average approximation. *Mathematical Programming*, 171(1):217–282.
- Caldentey, R., Liu, Y., and Lobel, I. (2017). Intertemporal pricing under minimax regret. *Operations Research*, 65(1):104–129.
- Chen, G., Daskin, M. S., Shen, Z.-J. M., and Uryasev, S. (2006). The α -reliable meanexcess regret model for stochastic facility location modeling. *Naval Research Logistics*, 53(7):617–626.
- Chen, Z. and Xie, W. (2021). Regret in the newsvendor model with demand and yield randomness. *Production and Operations Management*, 30(11):4176–4197.
- de Ruiter, F. J. C. T., Zhen, J., and den Hertog, D. (2022). Dual approach for two-stage robust nonlinear optimization. *Operations Research*, Advance online publication.
- Delage, E., Arroyo, S., and Ye, Y. (2014). The value of stochastic modeling in two-stage stochastic programs with cost uncertainty. *Operations Research*, 62(6):1377–1393.
- El-Amine, H., Bish, E. K., and Bish, D. R. (2018). Robust postdonation blood screening under prevalence rate uncertainty. *Operations Research*, 66(1):1–17.
- Feizollahi, M. J. and Averbakh, I. (2014). The robust (minmax regret) quadratic assignment problem with interval flows. *INFORMS Journal on Computing*, 26(2):321–335.
- Föllmer, H. and Schied, A. (2002). Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447.
- Furini, F., Iori, M., Martello, S., and Yagiura, M. (2015). Heuristic and exact algorithms for the interval min-max regret knapsack problem. *INFORMS Journal on Computing*, 27(2):392–405.
- Halpern, J. Y. and Leung, S. (2016). Minimizing regret in dynamic decision problems. *Theory and Decision*, 81(1):123–151.

- Hayashi, T. (2009). Stopping with anticipated regret. *Journal of Mathematical Economics*, 45(7-8):479–490.
- Hayashi, T. (2011). Context dependence and consistency in dynamic choice under uncertainty: the case of anticipated regret. *Theory and decision*, 70(4):399–430.
- Koçyiğit, Ç., Rujeerapaiboon, N., and Kuhn, D. (2022). Robust multidimensional pricing: Separation without regret. *Mathematical Programming*, 196:841–874.
- Krähmer, D. and Stone, R. (2008). Regret in dynamic decision problems. Technical report.
- Kuhn, D., Wiesemann, W., and Georghiou, A. (2011). Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209.
- Kusuoka, S. (2001). On law invariant coherent risk measures. In *Advances in Mathematical Economics*, pages 83–95. Springer Japan, Tokyo.
- Lim, A. E., Shanthikumar, J. G., and Vahn, G.-Y. (2012). Robust portfolio choice with learning in the framework of regret: Single-period case. *Management Science*, 58(9):1732–1746.
- Lim, A. E. B., Shanthikumar, J. G., and Shen, Z. J. M. (2006). Model uncertainty, robust optimization, and learning. In *Models, Methods, and Applications for Innovative Decision Making*, pages 66–94. INFORMS TutORials in Operations Research (INFORMS, Cantonsville, MD).
- McCormick, G. P. (1983). Nonlinear programming; theory, algorithms, and applications. Technical report.
- Mohajerin Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166.
- Natarajan, K., Shi, D., and Toh, K.-C. (2014). A probabilistic model for minmax regret in combinatorial optimization. *Operations Research*, 62(1):160–181.

- Park, J., El-Amine, H., and Mutlu, N. (2021). An exact algorithm for large-scale continuous nonlinear resource allocation problems with minimax regret objectives. *INFORMS Journal on Computing*, 33(3):1213–1228.
- Perakis, G. and Roels, G. (2008). Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203.
- Pichler, A., Liu, R. P., and Shapiro, A. (2022). Risk-averse stochastic programming: Time consistency and optimal stopping. *Operations Research*, 70(4):2439–2455.
- Postek, K., Ben-Tal, A., den Hertog, D., and Melenberg, B. (2018). Robust optimization with ambiguous stochastic constraints under mean and dispersion information. *Operations Research*, 66(3):814–833.
- Poursoltani, M. and Delage, E. (2022). Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems. *Operations Research*, 70(5):2906– 2930.
- Poursoltani, M., Delage, E., and Georghiou, A. (2021). Risk-averse regret minimization in multi-stage stochastic programs. Technical report, Groupe d'études et de recherche en analyse des décisions, GERAD, Montréal QC H3T 2A7, Canada.
- Rockafellar, R. T., Uryasev, S., and Zabarankin, M. (2006). Generalized deviations in risk analysis. *Finance and Stochastics*, 10(1):51–74.
- Ruszczyński, A. (2010). Risk-averse dynamic programming for markov decision processes. *Mathematical Programming*, 125(2):235–261.
- Savage, L. (1951). The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67.
- Shapiro, A. (2009). On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*, 37(3):143–147.
- Strack, P. and Viefers, P. (2021). Too proud to stop: Regret in dynamic decisions. *Journal* of the European Economic Association, 19(1):165–199.

- Xu, Z., Hu, Z., Song, Y., and Wang, J. (2015). Risk-averse optimal bidding strategy for demand-side resource aggregators in day-ahead electricity markets under uncertainty. *IEEE Transactions on Smart Grid*, 8(1):96–105.
- Yue, J., Chen, B., and Wang, M.-C. (2006). Expected value of distribution information for the newsvendor problem. *Operations Research*, 54(6):1128–1136.
- Zeng, B. and Zhao, L. (2013). Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461.
- Zhang, R., Li, G., Jiang, T., Chen, H., Li, X., Pei, W., and Xiao, H. (2020). Incorporating production task scheduling in energy management of an industrial microgrid: A regret-based stochastic programming approach. *IEEE Transactions on Power Systems*, 36(3):2663–2673.
- Zhu, Z., Zhang, J., and Ye, Y. (2013). Newsvendor optimization with limited distribution information. *Optimization Methods & Software*, 28(3):640–667.

Chapter 3

Robust Data-Driven Prescriptiveness Optimization

Abstract

The abundance of data has led to the emergence of a variety of optimization techniques, attempting to leverage the available side information to provide more anticipative decisions. Recent developments span a wide range of methods in the context of conditional optimization; on the other hand, the necessity of the existence of a universal unitless measure for the evaluation of different optimization schemes has given rise to the introduction of the coefficient of prescriptiveness, a two-folded metric for quantification of the quality of a data-driven decision compared to a reference decision as well as the prescriptiveness content of the side information. This chapter introduces a distributionally robust conditional stochastic optimization model where the coefficient of prescriptiveness substitutes for the classical empirical risk minimization objective. We provide a convex optimization reformulation for this problem, demonstrate how it reduces to a linear program when a nested Conditional Value at Risk represents the ambiguity set, and provide a bisection method together with an acceleration scheme for tackling it. Studying a shortest path problem, we evaluate the robustness of the resulting policies against alternative methods when the out-of-sample dataset experiences a distribution shift.

3.1 Introduction

Stochastic programming is perceived as one of the fundamental methods devised for decision-making under uncertainty (see Shapiro et al. 2021, Birge and Louveaux 2011). Given a cost function $h(x, \xi)$ that depends on a decision $x \in \mathbb{R}^{n_x}$ and a random vector $\xi \in \mathbb{R}^{n_{\xi}}$, the stochastic programming (SP) problem is defined as

$$(SP) \quad \boldsymbol{x}^* \in \arg\min_{\boldsymbol{x} \in \mathcal{X}} \quad \mathbb{E}_F[h(\boldsymbol{x}, \boldsymbol{\xi})], \tag{3.1}$$

where \mathcal{X} is a convex feasible set, $h(x, \xi)$ is a cost function that is assumed convex in x for all ξ , and ξ is assumed to be drawn from the distribution F. The solution methods for this problem mainly rely on either assuming a priori distribution for F or exploiting a set of independent and identically distributed observations. In the latter case, a set of i.i.d. observations of the random vector ξ denoted by $\mathcal{S} := {\xi_i}_{i=1}^N$ can be used to formulate the following sample average approximation problem:

$$(SAA) \quad \boldsymbol{x}^* \in \arg\min_{\boldsymbol{x} \in \mathcal{X}} \quad \frac{1}{N} \sum_{i=1}^N h(\boldsymbol{x}, \boldsymbol{\xi}_i), \tag{3.2}$$

where we assume a uniform distribution over the observed data. Recently, the availability of large datasets has played a critical role in redirecting the optimization methods devised for decision-making under uncertainty towards taking advantage of so-called "side information" or "covariates". This paradigm encourages decision-makers to benefit from the available data beyond the desired random variables to make more anticipative decisions. For instance, a portfolio manager who optimizes her investments in the stock market may consider a variety of available micro and macroeconomic indicators as side information to make more anticipative decisions (see Brandt et al. 2009, Bazier-Matte and Delage 2020). This gives rise to the following conditional stochastic optimization (CSO) problem:

(CSO)
$$\boldsymbol{x}^*(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x}\in\mathcal{X}} \mathbb{E}_F[h(\boldsymbol{x},\boldsymbol{\xi})|\boldsymbol{\zeta}],$$
 (3.3)

where $\zeta \in \mathbb{R}^{n_{\zeta}}$ denotes the given vector of "covariates", or so-called "features". In this case, any observed random vector ξ_i is accompanied by a vector of covariates $\zeta_i \in \mathbb{R}^{n_{\zeta}}$. The difficulty of this problem shows up when the conditional probability distribution

function $F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}$ is unknown, and only a set of i.i.d. observations $\mathcal{T} := \{(\boldsymbol{\zeta}_i, \boldsymbol{\xi}_i)\}_{i=1}^N$ is available. In this case, a data-driven variant of the CSO problem can be written as:

$$\boldsymbol{x}^{*}(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x}\in\mathcal{X}} \mathbb{E}_{\hat{F}_{\xi|\zeta}}[h(\boldsymbol{x},\boldsymbol{\xi})],$$
 (3.4)

where $\hat{F}_{\xi|\zeta}$ is a conditional probability model for $\boldsymbol{\xi}$ given $\boldsymbol{\zeta}$ inferred from the available data, e.g. by training a random forest (Breiman 2001), or estimated via kernel density estimation (Ban and Rudin 2019).

In order to compare the performance of different CSO approaches, Bertsimas and Kallus (2020) introduce the "coefficient of prescriptiveness" as:

$$\mathcal{P}_{F}(\boldsymbol{x}(\cdot)) := 1 - \frac{\mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}), \boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}{\mathbb{E}_{F}[h(\hat{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]},$$
(3.5)

where $\hat{x} \in \operatorname{argmin}_{x} \mathbb{E}_{\hat{F}}[h(x, \xi)]$ with \hat{F} as the empirical distribution that puts equal weights on each observed data point $\{\xi_i\}_{i=1}^N$ (i.e. the solution of SAA). The idea behind the coefficient of prescriptiveness is that it measures the performance of a given policy $x(\zeta)$ relative to the constant decision \hat{x} which is agnostic to the side information ζ , and to the fully anticipative policy which achieves the progressive optimal value of $\mathbb{E}_F[\min_{x'\in\mathcal{X}} h(x',\xi)]$. It is easy to see that a high value of \mathcal{P}_F indicates that the policy can leverage the contextual information of ζ with $\mathcal{P}_F = 1$ indicating that the policy is achieving the fully anticipative performance in terms of ξ . In contrast, a low value of \mathcal{P}_F indicates that the policy is not able to exploit (or even be misled by) the available information.

Following the introduction of the coefficient of prescriptiveness, this metric has been employed in several pieces of research to demonstrate the potential of proposed datadriven policies for leveraging the available side information. One can refer to Bertsimas et al. (2016) for such a comparison in the context of inventory management, Stratigakos et al. (2022) for energy trading, Notz and Pibernik (2022) for flexible capacity planning, and Kallus and Mao (2023) for shortest path and portfolio optimization problems. We note that, in the current literature, \mathcal{P}_F is only used as a benchmark metric for assessing the performance of policies computed using different approaches, e.g., in Bertsimas and Kallus (2020), the metric compares policies computed (amongst others) using CSO where the conditional probability is estimated by random forests and kernel density estimation. Given its prevalence as a performance measure, it is natural to question whether it is possible and useful to directly optimize the coefficient of prescriptiveness.

While one can show that maximizing \mathcal{P}_F reduces to solving CSO problem, one may wonder how the \mathcal{P}_F measure should be robustified in order to improve out-of-sample performance. In this work, we introduce for the first time a distributionally robust version of \mathcal{P}_F . We establish connections to other models in the literature and present an efficient algorithm to maximize it when the conditional probability model is discrete (such as with a random forest or with a Kernel density estimator). The rest of the chapter is organized as follows. Section 3.2 reviews the literature. Section 3.3 motivates the optimization of the coefficient of prescriptiveness by its relationship to the coefficient of determination in the field of statistics. Section 3.4 introduces a robust data-driven prescriptiveness optimization model that can be used to maximize a distributionally robust version of the coefficient of prescriptiveness. We reformulate this problem as a convex optimization problem that can reduce to a linear program when the ambiguity set takes the form of a so-called "nested Conditional Value-at-Risk (CVaR) set". A bisection method is proposed to solve the latter, as well as an acceleration scheme; finally, Section 3.5 presents the numerical experiments, where we evaluate the robustness of the resulting policies against benchmark ones in a shortest path problem when the out-of-sample dataset confronts a distribution shift.

3.2 Literature Review

Conditional Stochastic Optimization (CSO) In an attempt to use the side-information in a single-item newsvendor problem, Ban and Rudin (2019) apply decision rules to represent the decision x as an affine function of covariates ζ ; alternatively, they exploit Nadaraya-Watson Kernel regression (Nadaraya 1964, Watson 1964) to estimate the conditional probabilities $P(\boldsymbol{\xi}|\boldsymbol{\zeta})$ in (3.4). A prominent stream of research tackling Problem (3.4) falls under the category of the second case, i.e. "predict-then-optimize" methods. These practices seek an efficient predictive method to estimate the conditional probabilities $P(\boldsymbol{\xi}|\boldsymbol{\zeta})$ from the training data and then optimize the decisions accordingly. Hannah et al. (2010) use two nonparametric density estimators for conditional probabilities, including Nadaraya-Watson Kernel regression and Dirichlet process mixture models. In addition to Kernel methods, Bertsimas and Kallus (2020) employ k-nearest-neighbors regression, local linear regression, regression trees, and random forests for the same purpose.

Distributionally Robust Conditional Stochastic Optimization (DRCSO) All aforementioned methods deal with the estimation of conditional probability distributions given some empirical observations; consequently, the concerns provoking the emergence of Distributionally Robust Optimization (DRO) formulation proposed by Scarf (1958) potentially apply to the conditional stochastic optimization problem as well. One can refer to Delage and Ye (2010) and Wiesemann et al. (2014) for moment-based DRO, and see Mohajerin Esfahani and Kuhn (2018) and Gao and Kleywegt (2022) for distance-based DRO. A recent stream of research studies applying the DRO framework to conditional stochastic programs. Leveraging statistical bootstrap, Bertsimas and Van Parys (2022) extends the conditional stochastic optimization to the distributionally robust optimization setting to avoid overfitting in the presence of limited data and improve the out-of-sample performance. Ho and Hanasusanto (2019) propose out-of-sample performance bounds when using the Nadaraya-Watson Kernel regression estimator in the conditional stochastic optimization problem. Analyzing the derived bound, they suggest a regularization term to improve the out-of-sample performance and finally reformulate the problem as a distributionally robust optimization problem. Wang et al. (2021) use the Wasserstein ambiguity set to formalize a distributionally robust conditional optimization problem, where the conditional probabilities are predicted by Nadaraya-Watson Kernel regression estimator. DRCSO, in general, takes the following form

$$(DRCSO) \ \boldsymbol{x}^{*}(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sup_{F \in \mathcal{D}} \ \mathbb{E}_{F}[h(\boldsymbol{x}, \boldsymbol{\xi}) | \boldsymbol{\zeta}]$$
(3.6)

where \mathcal{D} is the ambiguity set containing the set of admissible distributions.

Coefficient of Predictive Prescriptiveness Concurrent with the advancement of conditional optimization methods, Bertsimas and Kallus (2020) developed the notion of the coefficient of prescriptiveness as a unitless measure to evaluate the quality of data-driven policies as well as the value of the prescriptive content of the data. One can refer to Bertsimas and Kallus (2020) for a comparison of coefficients of prescriptiveness resulting from different estimation methods for a two-stage shipping planning problem with synthetic data and a real-world problem faced by the distribution arm of an international media company. Notz and Pibernik (2022) compare their distribution-free Kerneralized empirical risk minimization approach against the alternatives using the coefficient of prescriptiveness. Kallus and Mao (2023) propose end-to-end solution methods for solving Problem (3.4) and consider the coefficient of prescriptiveness to compare their methods with their benchmarks within a shortest path problem. The exploitation of the coefficient of predictive prescriptiveness as a metric to measure the efficiency of different data-driven prescription methods proposed for an OR/MS problem with a specific dataset brings up the question of whether optimization of this metric itself in the DRCSO context has any potential for improvement.

Distribution Shifts Since we are in a data-driven setting and we are combining Machining Learning (ML) methods and optimization tools for predictive prescription, one might ask about the strength of these methods in capturing the effect of distribution shifts. Distribution shifts can be specific to the covariates vector (ζ), the random vector ($\boldsymbol{\xi}$), or both. Assume that the algorithm learns the conditional probabilities with the weekday dataset for a shortest path problem, and it is supposed to determine the optimal route for the weekend. This is a critical and common question in machine learning tasks, as it appears in many applications. One can refer to Schrouff et al. (2022) for a discussion about the relationship of robustness to distribution shift and fairness in healthcare applications. In some cases, the algorithms are designed to detect out-of-distribution scenarios, so they get rejected or the user receives an alert. Hsu et al. (2020), and Lee et al. (2018) build on the previous work to propose an efficient detection of out-of-distribution images; however, rejection of out-of-distribution samples cannot always be an applicable strategy. One can assume an autonomous driver which needs to make online decisions with minimum probabilities of catastrophic consequences according to the scenes. Filos et al. (2020) propose a robust method so that the autonomous driver can adapt to the distribution shifts. Although different methods are devised to mitigate the distribution shift effects in supervised learning, one might be curious about what could be done on the prescription side of the existing predict-then-optimize methods to achieve a higher coefficient of prescriptiveness.

3.3 Motivation for Optimizing \mathcal{P} and its Robustification

As argued in Bertsimas and Kallus (2020), in the context of predictive models, where one wishes to predict the value of $\xi \in \mathbb{R}$ based on a list of covariates ζ using a statistical model $f : \mathbb{R}^{n_{\zeta}} \to \mathbb{R}$, one popular metric that is employed takes the form of the so-called "coefficient of determination":

$$R^{2}(f(\cdot)) := 1 - \frac{\mathbb{E}_{\hat{F}}[(f(\boldsymbol{\zeta}) - \xi)^{2}]}{\mathbb{E}_{\hat{F}}[(\hat{\xi} - \xi)^{2}]},$$

where $\hat{\xi} := \mathbb{E}_{\hat{F}}[\xi]$ is the empirical mean of ξ in the data set. The popularity of R^2 compared to mean squared error as a measure of performance can be partially attributed to being unitless. It is upper bounded by 1, with a value closer to 1, indicating that most of the variation of ξ can be modeled using $f(\cdot)$. On the flip side, when strictly smaller than 0, its absolute value measures the percentage of additional variations that are introduced by the predictive model, thus indicating a degradation of predictive power when compared to the simple sample average $\hat{\xi}$.

The coefficient of prescriptiveness can be viewed as an attempt to introduce an analogous measure in the conditional optimization setting. More specifically, it reduces to R^2 when $n_x = 1$ and $h(x, \xi) := (x - \xi)^2$, namely:

$$\mathcal{P}_{\hat{F}}(x(\cdot)) = 1 - \frac{\mathbb{E}_{\hat{F}}[(x(\boldsymbol{\zeta}) - \xi)^2)] - \mathbb{E}_{\hat{F}}[\min_{x'}(x' - \xi)^2]}{\mathbb{E}_{\hat{F}}[(\hat{x} - \xi)^2] - \mathbb{E}_{\hat{F}}[\min_{x'}(x' - \xi)^2]} = R^2(x(\cdot)),$$

since $\mathbb{E}_{\hat{F}}[\min_{x'}(x'-\xi)^2] = 0$ and $\hat{x} \in \operatorname{argmin}_x \mathbb{E}_{\hat{F}}[(x-\xi)^2] = \hat{\xi}$. Hence, the coefficient of prescriptiveness has a similar interpretation as R^2 . Namely, \mathcal{P}_F is upper bounded by 1, and as it gets closer to 1, it indicates how successful the data-driven policy has been in closing the gap between the SAA solution that makes no use of covariate information and a hypothetical policy that would have access to full information about ξ .

One can also find traces in the literature of attempts to measure $R^2(x(\cdot))$ out-ofsample. Namely, Campbell and Thompson (2008) studies whether excess stock return predictors can outperform historical averages in terms of out-of-sample explanatory power of such predictors. This measure can be captured using

$$R_F^2(f(\cdot), \hat{F}) := 1 - \frac{\mathbb{E}_F[(f(\boldsymbol{\zeta}) - \xi)^2]}{\mathbb{E}_F[(\hat{\boldsymbol{\zeta}} - \xi)^2]} = \mathcal{P}_F(f(\cdot)),$$

which naturally leads to the question of whether $R^2(f(\cdot))$ is a good approximation for $R_F^2(f(\cdot), \hat{F})$ in a data-driven environment (potentially susceptible to distribution shifts). If not, then one must turn to employ more robust estimation methods.

3.4 Robust Data-Driven Prescriptiveness Optimization

In order to tackle the robustification and optimization of \mathcal{P} , we consider a more general version of this measure, which relaxes the assumption that the benchmark is the solution to (3.2) and widens the scope of our analysis. To this end, we define the prescriptiveness competitive ratio (PCR) of a policy $\boldsymbol{x}(\cdot)$ with respect to a reference policy $\bar{\boldsymbol{x}}$ as:

$$\mathcal{V}_{F}(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}) := \begin{cases} 1 - \frac{\mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}), \boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}{\mathbb{E}_{F}[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]} & \text{if } \mathbb{E}_{F}[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})] > 0\\ 1 & \text{if } \mathbb{E}_{F}[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] = \mathbb{E}_{F}[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})] = \mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}), \boldsymbol{\xi})] \cdot\\ -\infty & \text{otherwise} \end{cases}$$

$$(3.7)$$

Indeed, the coefficient of prescriptiveness can be considered a special case when $\bar{x} := \hat{x}$:

$$\mathcal{V}_F(\boldsymbol{x}(\cdot), \hat{\boldsymbol{x}}) = 1 - \frac{\mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}), \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\hat{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]} = \mathcal{P}_F(\boldsymbol{x}(\cdot)).$$

when $\mathbb{E}_F[h(\hat{x}, \xi)] - \mathbb{E}_F[\min_{x' \in \mathcal{X}} h(x', \xi)] > 0$, while the two other cases follow from the natural extension of the definition of $\mathcal{P}_F(x(\cdot))$. In contrast to $\mathcal{P}_F(x(\cdot))$, which benchmarks policy $x(\cdot)$ only to the SAA solution, the definition of \mathcal{V}_F allows to benchmark against any other static policy. This allows our model to accommodate situations where more sophisticated statistical tools might be used to obtain the reference decision (e.g. regularized or distributionally robust SAA).¹

In a finite sample regime, where \hat{F} might fail to capture the true underlying distribution, or in a situation where we expect distribution shifts, one should be interested in a distributionally robust estimation of the PCR (or equivalently of the coefficient of prescriptiveness), which takes the form of:

$$\mathcal{V}_{\mathcal{D}}(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}) := \inf_{F \in \mathcal{D}} \mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}) = \inf_{F \in \mathcal{D}} \mathcal{P}_F(\boldsymbol{x}(\cdot)) ext{ when } \bar{\boldsymbol{x}} := \hat{\boldsymbol{x}}$$

where \mathcal{D} is a set of distribution over the joint space (ζ, ξ) , and the notation $\mathcal{V}_{\mathcal{D}}$ is overloaded to denote the distributional robust PCR measure. Furthermore, one might be interested in identifying the policy that maximizes the PCR in the form of the following distributionally robust optimization problem:

$$(DRPCR) \quad \max_{oldsymbol{x}(\cdot)\in\mathcal{H}}\mathcal{V}_{\mathcal{D}}(oldsymbol{x}(\cdot),oldsymbol{ar{x}})$$

where $\mathcal{H} \subseteq \{ x : \mathbb{R}^{n_{\zeta}} \to \mathcal{X} \}$. The following lemma provides interpretable bounds for the value of $\mathcal{V}_{\mathcal{D}}$.

Lemma 3.4.1 If $\bar{x} \in \mathcal{H}$, then the optimal value of DRPCR is necessarily in the interval [0, 1].

Proof. This follows simply from $\mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}})$ being bounded above by 1 for any policy $x(\cdot)$ and any distribution F due to:

$$\begin{aligned} \mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}) &= 1 - \frac{\mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}), \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]} \\ &\leq 1 - \frac{\mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]} = 1 \end{aligned}$$

when $\mathbb{E}_F[h(\hat{x}, \xi)] - \mathbb{E}_F[\min_{x' \in \mathcal{X}} h(x', \xi)] > 0$, and otherwise equal to 1 or $-\infty$ both bounded above by 1. Hence,

$$\max_{\boldsymbol{x}(\cdot)\in\mathcal{H}}\inf_{F\in\mathcal{D}}\mathcal{V}_F(\boldsymbol{x}(\cdot),\bar{\boldsymbol{x}})\leq 1.$$

Moreover, if $\bar{x} \in \mathcal{H}$, then we have that

$$\max_{\boldsymbol{x}(\cdot)\in\mathcal{H}}\mathcal{V}_{\mathcal{D}}(\boldsymbol{x}(\cdot),\bar{\boldsymbol{x}}) \geq \mathcal{V}_{\mathcal{D}}(\bar{\boldsymbol{x}},\bar{\boldsymbol{x}}) = \begin{cases} 0 & \text{if } \mathbb{E}_{F}[h(\bar{\boldsymbol{x}},\boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\xi})] > 0\\ 1 & \text{otherwise.} \end{cases}$$

Lemma 3.4.1 can be interpreted as follows. First, if $x(\zeta)$ achieves a $\mathcal{V}_{\mathcal{D}}(x(\cdot), \bar{x}) = 1$, then the policy is guaranteed to exploit ζ just as efficiently as if it had full information about $\boldsymbol{\xi}$ (namely achieves the fully anticipative performance). On the other end of the spectrum, $\mathcal{V}_{\mathcal{D}}(x(\cdot), \bar{x}) = 0$ indicates that the policy can potentially fail to exploit any of the information present in $\boldsymbol{\xi}$. When $\bar{x} \in \mathcal{H}$, one can always prevent negative PCR by falling back to the benchmark policy \bar{x} .

Next, we show that in an environment where the distribution is known, the optimal policy obtained from CSO is an optimal solution to DRPCR. Before proceeding, we first make the following assumption.

Assumption 3.4.1 *The policy set* \mathcal{H} *contains all possible mappings, i.e.* $\mathcal{H} := \{ x : \mathbb{R}^{n_{\zeta}} \to \mathcal{X} \}.$

Lemma 3.4.2 Given that Assumption 3.4.1 is satisfied, if the distribution set is a singleton, i.e. $\mathcal{D} = \{\bar{F}\}\$, then the optimal policy obtained from the CSO problem that employs \bar{F} maximizes DRPCR.

Proof. Let $\tilde{x}(\cdot)$ be a CSO optimal policy, then necessarily $\tilde{x}(\cdot) \in \mathcal{H}$ since $\tilde{x}(\zeta) \in \mathcal{X}$ for all ζ . This confirms that $\tilde{x}(\cdot)$ is feasible in DRPCR. Next, we can demonstrate optimality through:

$$\mathcal{V}_{\mathcal{D}}(ilde{m{x}}(\cdot), ar{m{x}}) = \mathcal{V}_{ar{F}}(ilde{m{x}}(\cdot), ar{m{x}}) \geq \max_{m{x}(\cdot) \in \mathcal{H}} \mathcal{V}_{ar{F}}(m{x}(\cdot), ar{m{x}}) = \max_{m{x}(\cdot) \in \mathcal{H}} \mathcal{V}_{\mathcal{D}}(m{x}(\cdot), ar{m{x}}),$$

since for all $\boldsymbol{x}(\cdot) \in \mathcal{H}$, we have that $\mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})|\boldsymbol{\zeta}] \geq \min_{\boldsymbol{x}(\cdot)\in\mathcal{H}} \mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})|\boldsymbol{\zeta}] = \mathbb{E}_{F}[h(\tilde{\boldsymbol{x}}(\boldsymbol{\zeta}),\boldsymbol{\xi})|\boldsymbol{\zeta}]$ for all $\boldsymbol{\zeta}$, which we can show implies that $\mathcal{V}_{F}(\tilde{\boldsymbol{x}}(\cdot),\bar{\boldsymbol{x}}) \geq \mathcal{V}_{F}(\boldsymbol{x}(\cdot),\bar{\boldsymbol{x}})$. More specifically, if $\mathbb{E}_{F}[h(\hat{\boldsymbol{x}},\boldsymbol{\xi})] = \mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi})]$, then either $\mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})] = \mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi})]$ thus

$$\mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\xi})] = \mathbb{E}_{F}[h(\tilde{\boldsymbol{x}}(\boldsymbol{\zeta}),\boldsymbol{\xi})] \leq \mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})] = \mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\xi})]$$

meaning that $\mathcal{V}_F(\tilde{\boldsymbol{x}}(\cdot), \bar{\boldsymbol{x}}) = \mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}) = 1$, or $\mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}), \boldsymbol{\xi})] > \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]$ thus $\mathcal{V}_F(\tilde{\boldsymbol{x}}(\cdot), \bar{\boldsymbol{x}}) \geq -\infty = \mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}})$. Alternatively, the case where $\mathbb{E}_F[h(\hat{\boldsymbol{x}}, \boldsymbol{\xi})] > \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]$ is straightforward as the function

$$f(y) := 1 - \frac{y - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}{\mathbb{E}_F[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})]}$$

is strictly decreasing.

While Lemma 3.4.2 implies that DRPCR reduces to CSO when the distribution is known, thus making the question of PCR optimization and performance irrelevant, this is not the case anymore for larger ambiguity sets D.

In this section, we first present a convex reformulation of DRPCR and then provide a reformulation of the problem for the nested CVaR ambiguity set. Finally, we propose a decomposition algorithm for solving the problem based on a bisection algorithm.

3.4.1 Convex Formulation for DRPCR

The following proposition provides a convex reformulation of DRPCR.

Proposition 3.4.1 *Given that* $\bar{x} \in H$ *, DRPCR is equivalent to*

$$\max_{\boldsymbol{x}(\cdot)\in\mathcal{H},\gamma} \gamma \tag{3.8a}$$

subject to
$$Q(\boldsymbol{x}(\cdot), \gamma) \leq 0$$
 (3.8b)

$$0 \le \gamma \le 1. \tag{3.8c}$$

where

$$Q(\boldsymbol{x}(\cdot),\gamma) := \sup_{F \in \mathcal{D}} \mathbb{E}_F \Big[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi}) - \Big((1-\gamma)h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) + \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi}) \Big) \Big]$$

is a convex increasing function of γ . Moreover, Problem (3.8) is a convex optimization problem when \mathcal{H} is convex.

Proof. We first present the DRPCR in epigraph form:

$$\max_{\gamma, \boldsymbol{x}(\cdot) \in \mathcal{H}} \gamma \tag{3.9a}$$

subject to
$$\mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}) \ge \gamma, \ \forall F \in \mathcal{D}$$
 (3.9b)

$$0 \le \gamma \le 1 \tag{3.9c}$$

where we added the redundant constraint $\gamma \in [0, 1]$ since Lemma 3.4.1 ensures that the optimal value of DRPCR is in this interval.

Focusing on constraint (3.9b), we can then consider two cases for the definition of $\mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}})$. In the case that $\mathbb{E}_F[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi})] - \mathbb{E}_F[\min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})] > 0$, one can multiply both sides of the inequality to equivalently obtain:

$$\mathbb{E}_{F}[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\xi})] \le (1-\gamma) \left(\mathbb{E}_{F}[h(\bar{\boldsymbol{x}},\boldsymbol{\xi})] - \mathbb{E}_{F}[\min_{\boldsymbol{x}'\in\mathcal{X}}h(\boldsymbol{x}',\boldsymbol{\xi})]\right)$$

which is equivalent, when rearranging the terms, to:

$$\mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi}) - (1-\gamma)h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) - \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi})] \le 0.$$
(3.10)

In the second case where $\mathbb{E}_F[h(\bar{x}, \xi)] = \mathbb{E}_F[\min_{x' \in \mathcal{X}} h(x', \xi)]$, then constraint (3.9b) is equivalent to:

$$\mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})] = \mathbb{E}_F[\min_{\boldsymbol{x}'\in\mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi})] \quad \& \quad \gamma \leq 1,$$

yet $\gamma \leq 1$ is redundant while the former condition can equivalently be posed as (3.10). We are left with

$$\mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi}) - (1-\gamma)h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) - \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi})] \leq 0, \ \forall F \in \mathcal{D},$$

which can equivalently be described by $Q(\boldsymbol{x}(\cdot), \gamma) \leq 0$. One can further conclude that $Q(\boldsymbol{x}(\cdot),\gamma) \leq 0$ is convex and increasing in γ given that it is the supremum of a set of affine increasing functions:

$$Q(\boldsymbol{x}(\cdot),\gamma) = \sup_{F \in \mathcal{D}} \mathbb{E}_F[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi}) - h(\bar{\boldsymbol{x}},\boldsymbol{\xi})] + \gamma(\mathbb{E}_F[h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) - \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi})]),$$
$$h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) \ge \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi}) \text{ for all } \boldsymbol{\xi} \text{ since } \bar{\boldsymbol{x}} \in \mathcal{X}.$$

with $h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) \geq \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})$ for all $\boldsymbol{\xi}$ since $\bar{\boldsymbol{x}} \in \mathcal{X}$.

From the reformulation (3.8) one can draw interesting insides between the connection of DRPCR and risk-averse regret minimization, see Poursoltani et al. (2023). For $\gamma = 1$, the problem reduces to the ex-post risk-averse regret minimization problem. In contrast, for $\gamma = 0$, one can interpret the problem as regretting the performance of the policy compared to a policy with less information. In the notation of Poursoltani et al. (2023), this will lead to a risk-averse regret problem with $\Delta = -1$.

3.4.2 The Nested CVaR Ambiguity Set D

In the following, we consider a discrete empirical distribution \overline{F} and restrict \mathcal{D} to be a nested CVaR ambiguity set. This ambiguity set is motivated by the works on nested dynamic risk measures (see Riedel 2004, Detlefsen and Scandolo 2005, Ruszczyński and Shapiro 2006) as will be explained shortly. We formalize our approach through the following assumption.

Assumption 3.4.2 There is a discrete distribution \overline{F} , with $\{\zeta_{\omega}\}_{\omega\in\Omega_{\mathcal{C}}}$ and $\{\xi_{\omega}\}_{\omega\in\Omega_{\mathcal{E}}}$ as the set of distinct scenarios for ζ and ξ respectively, such that the distribution set \mathcal{D} takes the form of the "nested CVaR ambiguity set" with respect to $\mathbb{P}_{\bar{F}}$ and defined as

$$\bar{\mathcal{D}}(\bar{F},\alpha) := \left\{ F \in \mathcal{M}(\Omega_{\zeta} \times \Omega_{\xi}) \middle| \begin{array}{l} \mathbb{P}_{F}(\boldsymbol{\zeta} = \boldsymbol{\zeta}_{\omega}) = \mathbb{P}_{\bar{F}}(\boldsymbol{\zeta} = \boldsymbol{\zeta}_{\omega}) \ \forall \omega \in \Omega_{\zeta}, \\ \mathbb{P}_{F}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \boldsymbol{\zeta}_{\omega}) \leq (1/(1-\alpha)) \mathbb{P}_{\bar{F}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \boldsymbol{\zeta}_{\omega}) \\ \forall \omega \in \Omega_{\zeta}, \omega' \in \Omega_{\xi} \end{array} \right\}.$$
(3.11)

where $\mathcal{M}(\Omega_{\zeta} \times \Omega_{\xi})$ is the set of all distributions supported on over the joint space $\{\zeta_{\omega}\}_{\omega \in \Omega_{\zeta}} \times \{\xi_{\omega}\}_{\omega \in \Omega_{\xi}}$.

The structure of $\overline{D}(\overline{F}, \alpha)$ implies that there is no ambiguity in the marginal distribution of the observed random variable ζ . Rather, the ambiguity is solely on the unobserved random variable ξ and is sized using the parameter α . The nested CVaR ambiguity set owes its name from Ruszczyński and Shapiro (2006) and the fact that for any function $g(x, \xi)$:

$$\begin{split} \sup_{F\in\bar{\mathcal{D}}(\bar{F},\alpha)} & \mathbb{E}_{F}\left[g(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})\right] = \sup_{F\in\bar{\mathcal{D}}(\bar{F},\alpha)} \sum_{\omega\in\Omega_{\boldsymbol{\zeta}}} \sum_{\omega'\in\Omega_{\boldsymbol{\xi}}} \mathbb{P}_{F}(\boldsymbol{\zeta}=\boldsymbol{\zeta}_{\omega})\mathbb{P}_{F}(\boldsymbol{\xi}=\boldsymbol{\xi}_{\omega'}|\boldsymbol{\zeta}=\boldsymbol{\zeta}_{\omega})g(\boldsymbol{x}(\boldsymbol{\zeta}_{\omega}),\boldsymbol{\xi}_{\omega'}) \\ & = \sum_{\omega\in\Omega_{\boldsymbol{\zeta}}} \mathbb{P}_{\bar{F}}(\boldsymbol{\zeta}=\boldsymbol{\zeta}_{\omega}) \sup_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}\in\bar{\mathcal{D}}(\bar{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}\omega},\alpha)} \sum_{\omega'\in\Omega_{\boldsymbol{\xi}}} \mathbb{P}_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}(\boldsymbol{\xi}=\boldsymbol{\xi}_{\omega'})g(\boldsymbol{x}(\boldsymbol{\zeta}_{\omega}),\boldsymbol{\xi}_{\omega'}) \\ & = \mathbb{E}_{\bar{F}}\left[\operatorname{CVaR}_{\bar{F}}^{\alpha}\left(g(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})|\boldsymbol{\zeta}\right)\right], \end{split}$$

where $\bar{F}_{\xi|\zeta_{\omega}}$ is the conditional distribution of \bar{F} given ζ_{ω} and where we overload the notation of \bar{D} letting:

$$\bar{\mathcal{D}}(\bar{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}},\alpha) := \{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}} \in \mathcal{M}(\Omega_{\boldsymbol{\xi}}) : \mathbb{P}_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'}) \leq (1/(1-\alpha))\mathbb{P}_{\bar{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'}) \forall \, \omega' \in \Omega_{\boldsymbol{\xi}}\}.$$

For $\alpha = 0$, the problem reduces to

$$\min_{\boldsymbol{x}(\cdot)\in\mathcal{H}} \mathbb{E}_{\bar{F}}[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi})],$$

effectively recovering the CSO policy. On the other spectrum, for $\alpha = 1$ the problem reduces to

$$\min_{\boldsymbol{x}(\cdot)\in\mathcal{H}} \mathbb{E}_{\bar{F}}[\max_{\omega:P_{\bar{F}}(\boldsymbol{\xi}=\boldsymbol{\xi}_{\omega}|\boldsymbol{\zeta})>0} h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi}_{\omega})],$$

which implies that for each realization of ζ_{ω} the decision $x(\zeta_{\omega})$ is robust against all admissible realizations of ξ given ζ_{ω} .

The nested CVaR representation and full policy space Assumption 3.4.1 can be exploited to optimize $Q(\boldsymbol{x}(\cdot), \gamma)$. Namely, letting

$$g(\boldsymbol{x},\boldsymbol{\xi},\gamma) := h(\boldsymbol{x},\boldsymbol{\xi}) - \left((1-\gamma)h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) + \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi}) \right),$$

we have that

$$\psi(\gamma) := \min_{\boldsymbol{x}(\cdot) \in \mathcal{H}} Q(\boldsymbol{x}(\cdot), \gamma)$$

$$= \min_{\boldsymbol{x}(\cdot)\in\mathcal{H}} \sup_{F\in\bar{\mathcal{D}}(\bar{F},\alpha)} \mathbb{E}_{F} \Big[g(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi},\gamma) \Big]$$

$$= \min_{\boldsymbol{x}(\cdot)\in\mathcal{H}} \mathbb{E}_{\bar{F}} \Big[\operatorname{CVaR}_{\bar{F}}^{\alpha} \Big(g(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi},\gamma) | \boldsymbol{\zeta} \Big) \Big]$$

$$= \min_{\boldsymbol{x}(\cdot)\in\mathcal{H}} \mathbb{E}_{\bar{F}} \Big[\inf_{t} t + \frac{1}{1-\alpha} \mathbb{E}_{\bar{F}} \Big[\max \Big(0, g(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi},\gamma) - t \Big) | \boldsymbol{\zeta} \Big] \Big]$$

$$= \mathbb{E}_{\bar{F}} \Big[\inf_{\boldsymbol{x}\in\mathcal{X},t} t + \frac{1}{1-\alpha} \mathbb{E}_{\bar{F}} \Big[\max \Big(0, g(\boldsymbol{x},\boldsymbol{\xi},\gamma) - t \Big) | \boldsymbol{\zeta} \Big] \Big],$$

where we exploit the infimum representation of CVaR and the interchangeability property of expected value operators (see Shapiro 2017 and reference therein).

Given that \overline{F} is a discrete distribution as described in Assumption 3.4.2, one can compute $\psi(\gamma)$ by solving for each scenario ζ_{ω} with $\omega \in \Omega_{\zeta}$ the following optimization problem:

$$\phi_{\omega}(\gamma) := \min_{\boldsymbol{x} \in \mathcal{X}, t, \boldsymbol{s}} \quad t + \frac{1}{1 - \alpha} \sum_{\omega' \in \Omega_{\xi}} \mathbb{P}_{\bar{F}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \boldsymbol{\zeta} = \boldsymbol{\zeta}_{\omega}) s_{\omega'}$$
(3.12a)
subject to $s_{\omega'} \ge h(\boldsymbol{x}, \boldsymbol{\xi}_{\omega'}) - \left((1 - \gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}_{\omega'}) + \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}_{\omega'}) \right)$ $-t, \forall \omega' \in \Omega_{\xi}$ (3.12b)
 $s_{\omega'} \ge 0, \forall \omega' \in \Omega_{\xi}.$ (3.12c)

Based on the solution of Problem (3.12) for each $\omega \in \Omega_{\zeta}$, one can obtain $\psi(\gamma) := \sum_{\omega \in \Omega_{\zeta}} \mathbb{P}_{\bar{F}}(\zeta = \zeta_{\omega})\phi_{\omega}(\gamma)$ together with a potentially feasible policy $\boldsymbol{x}(\zeta) := \boldsymbol{x}^*_{\omega(\zeta)}$, where $\omega(\zeta) \in \operatorname{argmin}_{\omega \in \Omega_{\zeta}} \|\zeta - \zeta_{\omega}\|$ and $\boldsymbol{x}^*_{\omega}$ refers to the minimizer of Problem (3.12). We further note that Problem (3.12) can be reduced to a linear program when \mathcal{X} is polyhedral and $h(\boldsymbol{x}, \boldsymbol{\xi}_{\omega'})$ is linear programming representable for all $\omega' \in \Omega_{\xi}$.

Problem (3.8) can thus be reformulated as

$$\max_{\gamma} \gamma \qquad (3.13a)$$

subject to
$$\sum_{\omega \in \Omega_{\mathcal{L}}} \mathbb{P}_{\bar{F}}(\boldsymbol{\zeta} = \boldsymbol{\zeta}_{\omega}) \phi_{\omega}(\gamma) \le 0$$
(3.13b)

$$0 \le \gamma \le 1 \,, \tag{3.13c}$$

which can be reduced to a linear program when \mathcal{X} is polyhedral and $h(\boldsymbol{x}, \boldsymbol{\xi}_{\omega'})$ is linear programming representable. Whether the problem is reduceable to a linear program or

more generally a convex optimization model, its size scales with $|\Omega_{\zeta}| \cdot |\Omega_{\xi}|$, which can be computationally challenging. In the following, we present a decomposition algorithm that allows solving the problem efficiently.

3.4.3 A Bisection Algorithm for DRPCR

From the definition of $\psi(\gamma)$ and $\phi_{\omega}(\gamma)$, we observe that for fixed γ one can evaluate $\psi(\gamma)$ by solving $|\Omega_{\zeta}|$ distinct Problem (3.12) for each $\omega \in \Omega_{\zeta}$. Moreover, Proposition 3.4.1 states that $\psi(\gamma)$ is an increasing convex function of γ . Hence, one can design a bisection algorithm on γ to solve the DRPCR Problem (3.8). Namely, each step consists in identifying the mid-point $\tilde{\gamma}$ of an interval known to contain the optimal value of γ , and verifying whether $\tilde{\gamma}$ is feasible by solving $\min_{\boldsymbol{x}(\cdot) \in \mathcal{H}} Q(\boldsymbol{x}(\cdot), \tilde{\gamma})$ to decide which of the two sub-interval below or above $\tilde{\gamma}$ contains γ^* , see Figure 3.1 (left). The details of this algorithm are presented in Algorithm 2. Its efficiency relies on the difficulty of executing step 5, i.e. solving $\min_{\boldsymbol{x}(\cdot) \in \mathcal{H}} Q(\boldsymbol{x}(\cdot), \gamma)$.

Algorithm 2 Bisection algorithm for DRPCR

1: Input: Tolerance $\epsilon > 0$ 2: Set $\gamma^{-} := 0, \gamma^{+} := 1, x^{*}(\zeta) := \bar{x}$ for all ζ 3: while $\gamma^+ - \gamma^- > \epsilon$ do Set $\tilde{\gamma} := (\gamma^+ + \gamma^-)/2$ 4: Solve $\min_{\boldsymbol{x}(\cdot) \in \mathcal{H}} Q(\boldsymbol{x}(\cdot), \gamma)$ to get optimal policy $\tilde{\boldsymbol{x}}(\cdot)$ and optimal value ψ 5: if $\tilde{\psi} \leq 0$ then 6: 7: Set $\gamma^- := \tilde{\gamma}$ and $\boldsymbol{x}^*(\cdot) := \tilde{\boldsymbol{x}}(\cdot)$ else 8: Set $\gamma^+ := \tilde{\gamma}$ 9: end if 10: 11: end while 12: Return $\gamma^* := \gamma^-$ and $x^*(\cdot)$

One can possibly accelerate the convergence rate on the bisection algorithm by exploiting the fact that $\psi(\cdot)$ is an increasing convex function when \mathcal{X} is convex. Indeed, for the current interval $[\gamma^-, \gamma^+]$, $\psi(\gamma)$ can be under- and over-estimated, see Figure 3.1 (right). The procedure can be described as follows. First, we construct a line that will underestimate ψ by identifying a subgradient of the function at $\tilde{\gamma}$. This can be computed

analytically since

$$\begin{split} \psi(\gamma) &:= \mathbb{E}_{\bar{F}} \left[\min_{x \in \mathcal{X}} \operatorname{CVaR}_{\alpha} \left(h(\boldsymbol{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) + \gamma \min_{x' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}) \right) \Big| \zeta \right) \right] \\ &= \mathbb{E}_{\bar{F}} \left[\min_{x \in \mathcal{X}} \sup_{F \in \bar{\mathcal{D}}(\bar{F}, \alpha)} \mathbb{E}_{F} \left[h(\boldsymbol{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) + \gamma \min_{x' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}) \right) \Big| \zeta \right] \right] \\ &\geq \mathbb{E}_{\bar{F}} \left[\sup_{F \in \bar{\mathcal{D}}(\bar{F}, \alpha)} \min_{x \in \mathcal{X}} \mathbb{E}_{F} \left[h(\boldsymbol{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) + \gamma \min_{x' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}) \right) \Big| \zeta \right] \right] \\ &\geq \mathbb{E}_{\bar{F}} \left[\min_{x \in \mathcal{X}} \mathbb{E}_{F_{\xi|\zeta}} \left[h(\boldsymbol{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) + \gamma \min_{x' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}) \right) \Big| \zeta \right] \right] \\ &= \mathbb{E}_{\bar{F}} \left[\min_{x \in \mathcal{X}} \mathbb{E}_{F_{\xi|\zeta}} \left[h(\boldsymbol{x}, \boldsymbol{\xi}) - \left((1 - \gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) + \gamma \min_{x' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}) \right) \Big| \zeta \right] \right] \\ &= \underbrace{\mathbb{E}_{\bar{F}} \left[\min_{x \in \mathcal{X}} \mathbb{E}_{F_{\xi|\zeta}} \left[h(\boldsymbol{x}, \boldsymbol{\xi}) - h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) \right] \right]}_{\text{offset "a"}} + \gamma \underbrace{\mathbb{E}_{\bar{F}} \left[\mathbb{E}_{F_{\xi|\zeta}} \left[h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) - \min_{x' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi}) \right] \right]}_{\text{slope "b"}}, \end{split}$$

where $F_{\xi|\zeta}^*$ is the conditional probability given ζ of any member (hopefully a maximizer) of $\overline{\mathcal{D}}(\overline{F}, \alpha)$. Note that the first inequality is tight based on Sion's minimax theorem (see Sion 1958) given that $\overline{\mathcal{D}}(\overline{F}, \alpha)$ is compact, while the second is tight as long as $F_{\xi|\zeta}^*$ achieves the supremum. Such a maximizer can be identified using:

$$F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}^{*} \in \underset{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}} \in \mathcal{M}(\Omega_{\boldsymbol{\zeta}}): \\ \mathbb{P}_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}(\boldsymbol{\xi}) \leq (1-\alpha)^{-1} \mathbb{P}_{\bar{F}}(\boldsymbol{\xi}|\boldsymbol{\zeta}), \forall \boldsymbol{\xi}}{\operatorname{argmax}} \mathbb{E}_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}\left[h(\boldsymbol{x}_{\gamma}^{*}(\boldsymbol{\zeta}), \boldsymbol{\xi}) - \left((1-\gamma)h(\bar{\boldsymbol{x}}, \boldsymbol{\xi}) + \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}', \boldsymbol{\xi})\right)\right]$$

where $\boldsymbol{x}_{\gamma}^{*}(\boldsymbol{\zeta})$ is the minimizer of (3.12) with $\boldsymbol{\zeta}_{\omega} = \boldsymbol{\zeta}$ since $(\boldsymbol{x}_{\gamma}^{*}(\cdot), F^{*})$, with F^{*} as the composition of \bar{F} marginalized on $\boldsymbol{\zeta}$ and $F_{\boldsymbol{\xi}|\boldsymbol{\zeta}'}^{*}$ is a saddle point of:

$$g(\boldsymbol{x}(\cdot),F) := \mathbb{E}_F \Big[h(\boldsymbol{x}(\boldsymbol{\zeta}),\boldsymbol{\xi}) - \Big((1-\gamma)h(\bar{\boldsymbol{x}},\boldsymbol{\xi}) + \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} h(\boldsymbol{x}',\boldsymbol{\xi}) \Big) \Big].$$

Such a $F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}^*$ can be obtained as a side product of solving Problem (3.12) using the optimal dual variables associated with constraint (3.12b). If we denote by $\gamma_u := a/b$, then the right bound of the interval can be updated to $\gamma^{+'} := \min(\gamma^+, \gamma_u)$.

The second step is to construct an overestimator. If $\psi(\tilde{\gamma}) > 0$, then we evaluate $\psi(\gamma^{-})$ and construct the line that passes through $(\gamma^{-}, \psi(\gamma^{-}))$ and $(\tilde{\gamma}, \psi(\tilde{\gamma}))$. If $\psi(\tilde{\gamma}) < 0$ then we evaluate $\psi(\gamma^{+})$ and construct the line that passes through $(\gamma^{+}, \psi(\gamma^{+}))$ and $(\tilde{\gamma}, \psi(\tilde{\gamma}))$. We denote the point for which the line evaluates to zero as γ_{o} , and update the left bound of the interval to $\gamma^{-'} := \max(\gamma^{-}, \gamma_{o})$. Hence, the new interval is given by $[\gamma^{-'}, \gamma^{+'}] \subseteq$ $[\gamma^{-}, \gamma^{+}]$, which would potentially significantly reduce the search space.



Figure 3.1: Visualization of the Basic (Left) and Accelerated (Right) Bisection Algorithm. The blue squared brackets indicate the current estimated interval containing the optimal γ^* , and the red squared brackets indicate the interval in the next iterations. The right graph also visualizes the over and under estimators of $\psi(\gamma)$.

We conclude this section by commenting that the accelerated bisection algorithm could require up to two evaluations of the ψ function at each iteration instead of a single one as described in the original algorithm.

3.5 Experiments

In this section, we describe a numerical study that compares the performance of DRPCR against three other data-driven benchmark methods to evaluate its robustness to perturbations of data generating process; more specifically, we compare the performance of the corresponding data-driven policies in terms of the coefficient of prescriptiveness over an out-of-sample dataset. Indeed, we demonstrate how these models react to the situation where we face a distribution shift for $\boldsymbol{\xi}$. In a vehicle routing problem with travel time uncertainties, this can be interpreted as a shift in the distribution of the travel times, for instance, when a special event is happening in the town. Alternatively, one can think of an inventory management problem where the manager faces a shift in the distribution, e.g., an unforeseen increase in demand for disinfectants during the first days of the COVID-19 pandemic. In general, this may contain many cases where one may confront so-called "disruptions" or "extreme cases" in supply chain management as an immediate result of unexpected changes in distributions of the uncertain parameters.

The application that we consider for our numerical experiments is a shortest path problem described in Kallus and Mao (2023). A directed graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where \mathcal{V} denotes the set of nodes and $\mathcal{A} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of arcs, i.e., ordered pairs (i, j) of nodes describing the existence of a directed path from node *i* to node *j*. The corresponding travel time of such an arc is assumed to be $\xi_{(i,j)}$. The objective of this problem is to identify the shortest path from an origin (node *o*) to a destination (node *d*). Moving away from an ideal world of known parameters gives rise to a stochastic version of this problem. In this setting, the traveling times along the arcs $\boldsymbol{\xi} \in \mathbb{R}^{|\mathcal{A}|}$ are uncertain; however, one might still have access to side information or observed covariates. In this case, aiming at minimizing the expected travel time leads to the following CSO problem:

$$\boldsymbol{x}^{*}(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x}\in\mathcal{X}} \mathbb{E}_{\hat{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}[\boldsymbol{x}^{\top}\boldsymbol{\xi}],$$
 (3.14)

where

$$\mathcal{X} = \left\{ \boldsymbol{x} \in \mathbb{R}^{|\mathcal{A}|} \middle| \begin{array}{l} \boldsymbol{x} \ge 0 & \forall (i,j) \in \mathcal{A} \\ \sum_{j:(i,j) \in \mathcal{A}} x_{(i,j)} - \sum_{j:(j,i) \in \mathcal{A}} x_{(j,i)} = 1 & \text{if } i = o \\ \sum_{j:(i,j) \in \mathcal{A}} x_{(i,j)} - \sum_{j:(j,i) \in \mathcal{A}} x_{(j,i)} = -1 & \text{if } i = d \\ \sum_{j:(i,j) \in \mathcal{A}} x_{(i,j)} - \sum_{j:(j,i) \in \mathcal{A}} x_{(j,i)} = 0 & \forall i \in \mathcal{V} \setminus \{o, d\} \end{array} \right\}$$

and $x_{(i,j)} = 1$ if we decide to travel from node *i* to node *j* and $x_{(i,j)} = 0$ otherwise. Similar to Kallus and Mao (2023), we do not force integrality constraints. Furthermore, $\hat{F}_{\xi|\zeta}$ denotes the conditional distribution inferred from the training dataset.

As discussed in Section 3.2, DRCSO is a method proposed for robustifying the policies against distributional uncertainties in the data-driven context. Consequently, one can consider DRCSO, as an alternative to CSO, for solving this shortest-path problem. Using the nested CVaR ambiguity set introduced in Assumption 3.4.2 as the ambiguity set of DRCSO, one gets the model below:

$$(DRCSO) \qquad \boldsymbol{x}^{*}(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sup_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}} \in \bar{\mathcal{D}}(\hat{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}}, \alpha)} \quad \mathbb{E}_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}[\boldsymbol{x}^{\top}\boldsymbol{\xi}], \tag{3.15}$$

where

$$\bar{\mathcal{D}}(\hat{F}_{\xi|\zeta},\alpha) := \{F_{\xi|\zeta} \in \mathcal{M}(\Omega_{\xi}) : \mathbb{P}_{F_{\xi|\zeta}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'}) \le (1/(1-\alpha))\mathbb{P}_{\hat{F}_{\xi|\zeta}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'}) \forall \, \omega' \in \Omega_{\xi}\},$$

and α is the control parameter for the size of the ambiguity set. Staying in the DRCSO context, one can exploit a worst-case regret minimization approach instead of worst-case expected travel time. In our experiments, we look into the optimal solutions arising from an ex-post regret minimization setting, introduced as a $\Delta = 1$ regret minimization model in Poursoltani et al. (2023). This leads to the following distributionally robust conditional regret optimization (DRCRO) problem:

$$(DRCRO) \ \boldsymbol{x}^{*}(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x} \in \mathcal{X}} \sup_{F_{\xi|\boldsymbol{\zeta}} \in \bar{\mathcal{D}}(\hat{F}_{\xi|\boldsymbol{\zeta}},\alpha)} \ \mathbb{E}_{F_{\xi|\boldsymbol{\zeta}}}[\boldsymbol{x}^{\top}\boldsymbol{\xi} - \min_{\boldsymbol{x}' \in \mathcal{X}} \boldsymbol{x}'^{\top}\boldsymbol{\xi}].$$
(3.16)

In this case, the decision maker compares her travel time to the one resulting from a benchmark decision that knows the future realization of $\boldsymbol{\xi}$. The ultimate goal is to minimize the worst-case expectation of this gap, so-called "worst-case expected regret", where the ambiguity set is nested CVaR. Finally, we solve our introduced DRPCR problem under nested CVaR ambiguity set, where the $Q(\boldsymbol{x}(\cdot), \gamma)$ function takes the form of:

$$Q(\boldsymbol{x}(\cdot),\gamma) := \sup_{F \in \bar{\mathcal{D}}(\tilde{F},\alpha)} \mathbb{E}_F \Big[\boldsymbol{x}(\boldsymbol{\zeta})^\top \boldsymbol{\xi} - \Big((1-\gamma) \hat{\boldsymbol{x}}^\top \boldsymbol{\xi} + \gamma \min_{\boldsymbol{x}' \in \mathcal{X}} \boldsymbol{x}'^\top \boldsymbol{\xi} \Big) \Big],$$
(3.17)

where \tilde{F} denotes the distribution derived from the training dataset, composed of the empirical distribution \hat{F}_{ζ} of ζ and the inferred conditional distribution $\hat{F}_{\xi|\zeta}$, while $\hat{x} \in \operatorname{argmin}_{\boldsymbol{x}} \mathbb{E}_{\hat{F}}[h(\boldsymbol{x}, \boldsymbol{\xi})]$ with \hat{F} that puts equal weights on each observed data point $\{\boldsymbol{\xi}_i\}_{i=1}^N$ (i.e. the SAA solution). Based on an optimal solution γ^* for the DRPCR problem, one can retrieve an optimal policy using:

$$\boldsymbol{x}^{*}(\boldsymbol{\zeta}) \in \arg\min_{\boldsymbol{x}\in\mathcal{X}} \sup_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}\in\bar{\mathcal{D}}(\hat{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}},\alpha)} \mathbb{E}_{F_{\boldsymbol{\xi}|\boldsymbol{\zeta}}} \left[\boldsymbol{x}^{\top}\boldsymbol{\xi} - \left((1-\gamma^{*})\hat{\boldsymbol{x}}^{\top}\boldsymbol{\xi} + \gamma^{*}\min_{\boldsymbol{x}'\in\mathcal{X}} \boldsymbol{x}'^{\top}\boldsymbol{\xi}\right) \right], \quad (3.18)$$

which can be obtained by solving (3.12) with γ^* and replacing $\mathbb{P}_{\bar{F}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \boldsymbol{\zeta} = \boldsymbol{\zeta}_{\omega})$ with $\mathbb{P}_{\hat{F}_{\xi|\zeta}}(\boldsymbol{\xi}_{\omega'})$.

We adapt our numerical experiments to the graph (\mathcal{G}) structure employed in Kallus and Mao (2023) with the same origin (o) and destination (d); therefore, we study a graph with the size of 45 nodes ($|\mathcal{V}| = 45$) and 97 arcs ($|\mathcal{A}| = 97$). We assume there exist 200 covariates ($n_{\zeta} = 200$) and the vector composed of travel times $\boldsymbol{\xi}$ and covariates $\boldsymbol{\zeta}$ follows a multivariate normal distribution. Specifically, each covariate ζ_i follows a normal distribution with a mean of zero and standard deviation of one (i.e. $\zeta_i \sim \mathcal{N}(0, 1)$). Similarly, each travel time $\xi_{(i,j)}$ is normal with a standard deviation that matches the deviation present in Kallus and Mao (2023)'s dataset, yet both the correlation and mean vector are treated differently. Starting with correlation, we introduce a new correlation structure for $(\boldsymbol{\zeta}, \boldsymbol{\xi})^3$ by instantiating a random correlation matrix (see Appendix 3.6 for details).

Our treatment of the mean of $\boldsymbol{\xi}$ embodies our objective to study robustness to distribution shifts. Namely, while the data generating process for the training set employs the same mean vector as in Kallus and Mao (2023), our validation data set and out-ofsample test set will measure the performance of proposed policies on generating processes where the mean of $\boldsymbol{\xi}$ as been perturbed, i.e. $\mathbb{E}[\xi_{(i,j)}] := (1 + \delta_{(i,j)})\mu_{(i,j)}$. Six tests were conducted for different levels of mean perturbations: no distribution shift $\delta_{(i,j)} = 0$, which does not allow for any perturbation, along with tests that take into account shifts with $\delta_{(i,j)}$ generated i.i.d. according to a uniform distribution on [0%, m], where $m \in \mathcal{M} := \{20\%, 30\%, 40\%, 50\%, 60\%\}$ represents the maximum possible perturbation. Furthermore, the perturbation experienced in the validation set is independent of the test set. This is to simulate situations where the level of robustness would be calibrated on a data set where a distribution shift of similar size is observed as the shift experienced out-of-sample.

Experiments for each perturbation range contain 50 instances generated by resampling the training, validation, and test data sets. Both the training and validation datasets consist of 400 data points, while the test set contains 1000 data points and is used to measure the "out-of-sample" performance. The training dataset is used for learning purposes, which allows us to infer the conditional probabilities of $\hat{F}_{\xi|\zeta}$ once a new covariate vector $\boldsymbol{\zeta}$ is observed. From a wide range of existing predictive tools for inference of $\hat{F}_{\xi|\zeta}$, Bertsimas and Kallus (2020) compare methods such as k-nearest-neighbors regression (Hastie et al. 2001), local linear regressions (Cleveland and Devlin 1988), classification and regression trees (CART; Breiman et al. 1984), and random forests (RF; Breiman 2001). In their experiments, the best coefficient of prescriptiveness belongs to random forests. We exploit the code provided in Kallus and Mao (2023) to train random forests over our training datasets and then use it as the conditional distribution estimator $\hat{F}_{\xi|\zeta}$ for our validation and out-of-sample data points. The validation dataset is used to calibrate the size of the ambiguity set (α) for the DRCSO, DRCRO, and DRPCR models. The procedure

of calibrating α and the associated optimal γ for the DRPCR model is described in Algorithm 3. A similar process in Algorithm 4 clarifies how we calibrate α for the DRCSO and DRCRO models. We define the set of discretized α values as $\mathcal{A} := \mathcal{A}_1 \cap \mathcal{A}_2$, where \mathcal{A}_1 includes 20 logarithmically spaced values in [0.01, 0.99] and \mathcal{A}_2 includes 20 evenly spaced values in [0, 1). For CSO, Algorithm 4 can also be used with $\mathcal{A} = \{0\}$.

Algorithm 3 Algorithm for calibrating the size of the ambiguity set (α) for DRPCR

- 1: Input: Training dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$ and validation dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$ and $\mathcal{A} := \{\alpha_i\}_{i=1}^n \subset [0, 1]$
- 2: Train a random forest model $\hat{F}_{\xi|\zeta}$ on $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$
- 3: Let \tilde{F} be the composition of $\hat{F}_{\xi|\zeta}$ with empirical distribution \hat{F}_{ζ} of ζ in the training set $\{\zeta_j\}_{j=1}^{N_{train}}$
- 4: for i = 1, ..., n do
- 5: //Construct $\hat{x}_i^*(\cdot)$ with α_i and F
- 6: Solve DRPCR with α_i and \tilde{F} to get γ_i^*
- 7: //Evaluate $\hat{x}_i^*(\cdot)$ on empirical distribution of realizations in $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$
- 8: **for** $j = 1, \ldots, N_{validation}$ **do**
- 9: Solve (3.12) with γ_i^* , α_i , and replacing $\mathbb{P}_{\bar{F}}(\boldsymbol{\xi} = \boldsymbol{\xi}_{\omega'} | \boldsymbol{\zeta} = \boldsymbol{\zeta}_{\omega})$ with $\mathbb{P}_{\hat{F}_{\boldsymbol{\xi}|\boldsymbol{\zeta}_j}}(\boldsymbol{\xi}_{\omega'})$ to get optimal \boldsymbol{x}_i^*
- 10: Let $\hat{x}_i^*(\zeta_j) := x_j^*$
- 11: end for
- 12: Set $s^i := \mathcal{P}^{\alpha}_{\hat{F}}(\hat{x}^*_i(\cdot))$ for empirical distribution \hat{F} on $\{\boldsymbol{\zeta}_j, \, \boldsymbol{\xi}_j\}_{i=1}^{N_{validation}}$
- 13: end for

14: Let $i^* \in \operatorname{argmax}_i s^i$ and set $\alpha^* := \alpha_{i^*}, \gamma^* := \gamma_{i^*}$, and $\boldsymbol{x}^*(\cdot) := \hat{\boldsymbol{x}}^*_{i^*}(\cdot)$

```
15: Return (\alpha^*, \gamma^*, \boldsymbol{x}^*(\cdot))
```

Figure 3.2 reports the coefficients of prescriptiveness $\mathcal{V}_F(\mathbf{x}^*(\cdot), \hat{\mathbf{x}})$, where F is the test dataset, for the four policies and perturbation levels. We observe the following: (*i*) When considering a particular optimization model, the coefficient of prescriptiveness decreases as the magnitude of the distribution shift increases. Indeed, these policies face a more serious robustness challenge as they approach more extreme scenarios beyond what was seen in the train dataset. (*ii*) When the test set follows the same distribution as the train set, all four policies roughly demonstrate similar performance; however, when this set experiences a distribution shift, DRPCR policies differentiate their performance compared to the alternative ones. (*iii*) Imposing a more severe distribution shift accentuates this differentiation. For instance, when the mean travel times across the edges are perturbed up to 50% in the test set, DRPCR policies provide a positive coefficient of prescriptive-

Algorithm 4 Algorithm of calibrating the size of the ambiguity set (α) for CVaR-loss/CVaR-regret

- 1: Input: Training dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$ and validation dataset $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$ and $\mathcal{A} := \{\alpha_i\}_{i=1}^n \subset [0, 1]$
- 2: Train a random forest model $\hat{F}_{\xi|\zeta}$ on $\{\zeta_j, \xi_j\}_{j=1}^{N_{train}}$
- 3: for i = 1, ..., n do
- 4: //Evaluate $\hat{x}_i^*(\cdot)$ on empirical distribution of realizations in $\{\zeta_j, \xi_j\}_{j=1}^{N_{validation}}$
- 5: **for** $j = 1, \ldots, N_{validation}$ **do**
- 6: Solve (3.15)/(3.16) with α_i for $\boldsymbol{\zeta} := \boldsymbol{\zeta}_j$ in validation set to get optimal \boldsymbol{x}_j^*
- 7: Let $\hat{x}_{i}^{*}(\zeta_{j}) := x_{j}^{*}$
- 8: end for
- 9: Set $s^i := \mathcal{P}^{\alpha}_{\hat{F}}(\hat{x}^*_i(\cdot))$ for empirical distribution \hat{F} on $\{\zeta_j, \xi_j\}_{i=1}^{N_{validation}}$

```
10: end for
```

- 11: Let $i^* \in \operatorname{argmax}_i s^i$ and set $\alpha^* := \alpha_{i^*}$ and $\boldsymbol{x}^*(\cdot) := \hat{\boldsymbol{x}}^*_{i^*}(\cdot)$
- 12: Return $(\alpha^*, \boldsymbol{x}^*(\cdot))$



Figure 3.2: Out-of-Sample Performance of Relaxed $x(\cdot)$



Figure 3.3: Out-of-Sample Performance of Binary $x(\cdot)$

ness, at least over 75% of instances. On the contrary, the alternative policies fail to reach a positive ratio over almost a similar number of instances. This observation is further amplified in the case of 60% perturbation. In this scenario, while CSO, DRCSO, and DRCRO policies fail to return a positive out-of-sample coefficient of prescriptiveness, DRPCR still can reach a positive median of 2% which can go up to 11% at its best.

While our first set of experiments considered a relaxed version of the shortest path problem to be closer to the real-world application, we also conduct a second set of experiments where $x(\cdot)$ represents binary variables and leads to implementable trajectories. Figure 3.3 illustrates the coefficients of prescriptiveness obtained from optimal binary policies. These results, in general, are aligned with the ones spotted in Figure 3.2; however, one remarks the following. Firstly, the results derived from CSO remain exactly the same as the relaxed case. This stems from the fact that optimal relaxed CSO decisions are

Problem Type	Method	Level of Perturbation					
		0%	20%	30%	40%	50%	60%
Relaxed $x(\cdot)$	CSO	0.45	0.30	0.19	0.04	-0.13	-0.31
	DRCSO	0.45	0.30	0.18	0.04	-0.13	-0.31
	DRCRO	0.45	0.30	0.18	0.04	-0.13	-0.32
	DRPCR	0.45	0.31	0.23	0.13	0.05	0.01
Binary $oldsymbol{x}(\cdot)$	CSO	0.45	0.30	0.19	0.04	-0.13	-0.31
	DRCSO	0.44	0.30	0.19	0.06	-0.09	-0.25
	DRCRO	0.44	0.30	0.19	0.05	-0.11	-0.28
	DRPCR	0.44	0.32	0.24	0.15	0.07	0.02

Table 3.1: Mean Out-of-Sample Coefficient of Prescriptiveness

known to be integral for the stochastic shortest path problems; conversely, this is not the case for DRCSO, DRCRO, and DRPCR, where robustness breaks the linearity of the objective. Secondly, forcing DRCSO and DRCRO to propose binary policies enhances their out-of-sample performance, surpassing those of CSO. Indeed, this setting seems to provide these two approaches the chance to better exploit the information about potential distribution shifts; however, despite their enhanced performance, the highest degree of robustness to distribution shift remains associated with DRPCR policies. Thirdly, Figure 3.3 presents counter-intuitive empirical evidence that out-of-sample performance might be slightly improved when imposing integrality constraints on the three robust models. We hypothesize that this might be caused by the additional flexibility of the relaxed models, which makes them more susceptible to overfitting their assumed stochastic models. Finally, one should note that the additional price of obtaining DRPCR policies, compared to the alternative ones, mainly consists in the computations embedded in Step 6 of Algorithm 3. The average runtime of this step across all instances and perturbation levels is 27 minutes for the relaxed version of the experiments and 33 minutes for the non-relaxed problem.

3.6 Appendix

A random covariance matrix for the random vector of (ζ, ξ) with arbitrary variances is generated based on a two-step procedure that follows. The first step consists in generating a random symmetric positive-definite matrix described in Algorithm 5, a method implemented in the sklearn.datasets.make_spd_matrix function of scikit-learn machine learning library in Python.

Algorithm 5 Algorithm for generating random symmetric positive-definite matrix

- 1: Input: Dimension of the square matrix $n_{\zeta} + n_{\xi}$
- 2: Generate random square matrix $A_{n_{\zeta}+n_{\xi}}$ sampling from the uniform distribution $\mathcal{U}_{[0,1]}$
- 3: Construct the symmetric matrix $M = A^{\top}A$
- 4: Decompose M with Singular Value Decomposition (SVD) method as $M = U\Sigma V^{\top}$
- 5: Generate random diagonal matrix S sampling from the uniform distribution $\mathcal{U}_{[0,1]}$
- 6: Construct $\Sigma' = S + J$ where J is the square matrix of ones with the size of $n_{\zeta} + n_{\xi}$
- 7: Get the symmetric positive-definite matrix as $M' = U\Sigma' V^{\top}$
- 8: Return M'

Given the vector of standard deviations for (ζ, ξ) denoted by $[\sigma_{\zeta}^{\top} \sigma_{\xi}^{\top}]^{\top}$ and also a random symmetric positive-definite matrix generated by Algorithm 5, one can implement the second stage described in Algorithm 6 to get a random covariance matrix with arbitrary standard deviations of $[\sigma_{\zeta}^{\top} \sigma_{\xi}^{\top}]^{\top}$.

Algorithm 6 Algorithm for generating random covariance matrix with arbitrary standard deviations

- Input: Random symmetric positive-definite matrix (*M*) and vector of standard deviations [σ^T_ζ σ^T_ξ]^T
- 2: Convert matrix *M* into its associated correlation matrix Corr = $\left(\operatorname{diag}(M)\right)^{-\frac{1}{2}} M\left(\operatorname{diag}(M)\right)^{-\frac{1}{2}}$

3: Get the arbitrary covariance matrix of $\text{Cov} = \text{diag}\left(\begin{bmatrix} \sigma_{\zeta} \\ \sigma_{\xi} \end{bmatrix}\right) (\text{Corr}) \text{diag}\left(\begin{bmatrix} \sigma_{\zeta} \\ \sigma_{\xi} \end{bmatrix}\right)$

4: Return Cov

Endnotes

1. In fact, one can go a step further and define $\mathcal{V}_F(\boldsymbol{x}(\cdot), \bar{\boldsymbol{x}}(\cdot))$ were $\bar{\boldsymbol{x}}(\cdot)$ is not a static policy. For example, $\bar{\boldsymbol{x}}(\cdot)$ could be a simple rule-based policy such as the order-up-to policy in inventory control. For ease of exposition, we treat the benchmark policy $\bar{\boldsymbol{x}}$ as a static policy for the remainder of the chapter.

2. Namely, $P_{F^*}(\boldsymbol{\xi}) = P_{\bar{F}}(\boldsymbol{\xi})$ and $P_{F^*}(\boldsymbol{\xi}|\boldsymbol{\zeta}) = P_{F^*_{\boldsymbol{\xi}|\boldsymbol{\zeta}}}(\boldsymbol{\xi})$ for all $\boldsymbol{\zeta}$.

3. This was done after observing that with Kallus and Mao (2023)'s dataset, the optimal uninformed decisions produced nearly the same performance as the optimal hindsight decisions that exploited full information about realized travel cost.

References

- Ban, G.-Y. and Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108.
- Bazier-Matte, T. and Delage, E. (2020). Generalization bounds for regularized portfolio selection with market side information. *INFOR: Information Systems and Operational Research*, 58(2):374–401.
- Bertsimas, D. and Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044.
- Bertsimas, D., Kallus, N., and Hussain, A. (2016). Inventory management in the era of big data. *Production and Operations Management*, 25(12):2006–2009.
- Bertsimas, D. and Van Parys, B. (2022). Bootstrap robust prescriptive analytics. *Mathematical Programming*, 195(1-2):39–78.
- Birge, J. R. and Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.
- Brandt, M. W., Santa-Clara, P., and Valkanov, R. (2009). Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns. *The Review of Financial Studies*, 22(9):3411–3447.

Breiman, L. (2001). Random forests. *Machine learning*, 45:5–32.

- Breiman, L., Friedman, J., Stone, C., and Olshen, R. (1984). Classification algorithms and regression trees. *Classification and regression trees*, 15(2):246.
- Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies*, 21(4):1509– 1531.
- Cleveland, W. S. and Devlin, S. J. (1988). Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83(403):596–610.
- Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612.
- Detlefsen, K. and Scandolo, G. (2005). Conditional and dynamic convex risk measures. *Finance and Stochastics*, 9(4):539–561.
- Filos, A., Tigkas, P., McAllister, R., Rhinehart, N., Levine, S., and Gal, Y. (2020). Can autonomous vehicles identify, recover from, and adapt to distribution shifts? In *International Conference on Machine Learning*, pages 3145–3153. PMLR.
- Gao, R. and Kleywegt, A. (2022). Distributionally robust stochastic optimization with wasserstein distance. *Mathematics of Operations Research*.
- Hannah, L., Powell, W., and Blei, D. (2010). Nonparametric density estimation for stochastic optimization with an observable state variable. *Advances in Neural Information Processing Systems*, 23.
- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA.
- Ho, C. P. and Hanasusanto, G. A. (2019). On data-driven prescriptive analytics with side information: A regularized nadaraya-watson approach. Working Paper.

- Hsu, Y.-C., Shen, Y., Jin, H., and Kira, Z. (2020). Generalized odin: Detecting out-ofdistribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10951–10960.
- Kallus, N. and Mao, X. (2023). Stochastic optimization forests. *Management Science*, 69(4):1975–1994.
- Lee, K., Lee, K., Lee, H., and Shin, J. (2018). A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31.
- Mohajerin Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166.
- Nadaraya, E. A. (1964). On estimating regression. *Theory of Probability & Its Applications*, 9(1):141–142.
- Notz, P. M. and Pibernik, R. (2022). Prescriptive analytics for flexible capacity management. *Management Science*, 68(3):1756–1775.
- Poursoltani, M., Delage, E., and Georghiou, A. (2023). Risk-averse regret minimization in multistage stochastic programs. *Operations Research*, https://doi.org/10.1287/opre.2022.2429.
- Riedel, F. (2004). Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112(2):185–200.
- Ruszczyński, A. and Shapiro, A. (2006). Conditional risk mappings. *Mathematics of Operations Research*, 31(3):544–561.
- Scarf, H. (1958). A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production*.
- Schrouff, J., Harris, N., Koyejo, O., Alabdulmohsin, I., Schnider, E., Opsahl-Ong, K., Brown, A., Roy, S., Mincu, D., Chen, C., et al. (2022). Maintaining fairness across dis-

tribution shift: do we have viable solutions for real-world applications? *arXiv preprint arXiv:2202.01034*.

- Shapiro, A. (2017). Interchangeability principle and dynamic equations in risk averse stochastic programming. *Operations Research Letters*, 45(4):377–381.
- Shapiro, A., Dentcheva, D., and Ruszczynski, A. (2021). *Lectures on stochastic programming: modeling and theory*. SIAM.
- Sion, M. (1958). On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176.
- Stratigakos, A., Camal, S., Michiorri, A., and Kariniotakis, G. (2022). Prescriptive trees for integrated forecasting and optimization applied in trading of renewable energy. *IEEE Transactions on Power Systems*, 37(6):4696–4708.
- Wang, T., Chen, N., and Wang, C. (2021). Distributionally robust prescriptive analytics with Wasserstein distance. *arXiv preprint arXiv:2106.05724*.
- Watson, G. S. (1964). Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 359–372.
- Wiesemann, W., Kuhn, D., and Sim, M. (2014). Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376.
General Conclusion

Regret minimization is a key concept in decision theory, game theory, economics, artificial intelligence, and machine learning that provides a powerful framework for making better decisions in the face of uncertainty and ambiguity. The main contributions of this thesis are mainly attributed to the development of mathematical models and optimization techniques to obtain optimal "regret-averse" decisions, as well as the subsequent analysis of these decisions to examine their behaviors and outcomes.

In the first chapter, we cast the two-stage worst-case regret minimization problems as two-stage robust optimization problems. This discovery allows us to identify a vast array of advanced solution methods that already exist within the adjustable robust optimization literature for this challenging and complex class of problems; besides, we identified subclasses of two-stage worst-case regret minimization problems that are polynomially solvable. The second chapter explored risk-averse regret minimization in multistage stochastic programs. We introduced the Δ -regret model, which provides the policymaker with an extra tool to deal with the uncertainty, i.e., imposing desired information structure on the benchmark policy. We studied this model under popular risk measures and demonstrated how it reduces to special cases or programs which can be solved via existing solution schemes from robust optimization. This model interpolates between the ex-ante and ex-post regret minimization models, the concepts introduced implicitly in single-stage regret minimization problems. The third chapter introduced a novel distributionally robust conditional stochastic optimization problem that optimizes the coefficient of prescriptiveness. Given a data-driven decision, this coefficient compares the difference between the outcomes of this decision and the hindsight decision against the gap arising from the outcomes of a reference decision made without any access to the side information and the hindsight policy. We formulated this problem as a linear program when nested CVaR represented the ambiguity set and tested the robustness of its optimal decisions under distribution shift.

This thesis opens new opportunities for future research in regret minimization. Twostage regret minimization problems investigated in Chapter 1 can be explored under different assumptions. For instance, one may consider both objective and right-hand side uncertainties simultaneously or study the case where the technology matrix is uncertain. In Chapter 2, we assumed a discrete outcome space which allowed us to reformulate the risk-averse multi-stage regret minimization problem as a two-stage robust optimization problem. To capture more general cases, this study can be further extended by investigating the setting where the outcome space is continuous; besides, one could explore the idea of how other alternative risk measures can affect regret-averse policies. Regarding the model proposed in Chapter 3, one could study whether optimizing the coefficient of prescriptiveness under ambiguities sets other than nested CVaR could identify solutions with a better edge over alternative methods in the presence of distribution shift. Throughout the thesis, we exploited miscellaneous applications to provide experimental evidence, e.g., multi-item newsvendor problem, multi-period inventory management problem, production transportation problem, shortest path problem, and portfolio selection problem; however, we still believe that the potentials of the proposed models and solution schemes for enhancement of the best practices in real-world applications remain largely unexplored.