

Diss. ETH No. 21706

Scenario-Based Optimization for Multi-Stage Stochastic Decision Problems

A dissertation submitted to
ETH ZURICH

for the degree of
Doctor of Sciences

presented by

GEORG SCHILDBACH

Dipl.-Ing., Darmstadt University of Technology

Dipl.-Wirtsch.-Ing., Darmstadt University of Technology

born July 21st, 1981
citizen of Germany

accepted on the recommendation of

Prof. Dr. Manfred Morari, examiner

Prof. Dr. Marco Campi, co-examiner

Prof. Dr. Francesco Borrelli, co-examiner

2014

Scenario-Based Optimization for Multi-Stage Stochastic Decision Problems

Georg Schildbach

Automatic Control Laboratory
ETH Zurich
Zurich, 2014

Automatic Control Laboratory
ETH Zurich
Switzerland

© 2014 Georg Schildbach. All rights reserved.

To my parents, for their eternal and unconditional love and support,
at all times and in all situations of my life.

Acknowledgments

“There is nothing as practical as a good theory.”

— Kurt Lewin, 1890–1947 (attributed)

First and foremost, I thank my advisor Prof. Manfred Morari for supervising my dissertation. He has given me a great deal of freedom and support in exploring new topics and in pursuing research directions that fit my personal interests. The wisdom and clarity of his guidance, and the example he sets for his students, became an invaluable part of my research experience.

Moreover, I thank Prof. Marco Campi and Prof. Francesco Borrelli for agreeing to act as co-examiners of my dissertation. In a variety of ways—through lectures, papers, discussions, and reviews—both of them have taught me important lessons that go well beyond their feedback on my dissertation.

It was also great experience for me to be part of the Automatic Control Laboratory (IfA) at ETH Zurich. Without a doubt, Prof. Manfred Morari and Prof. John Lygeros have created the best research group in control engineering imaginable. The group brings together many outstanding, interesting, and diverse people and it offers an inspiring and collaborative atmosphere. I have greatly enjoyed many planned and spontaneous discussions—on control and totally different topics.

The most important collaborators during my doctorate were the post-doctoral researchers Colin Jones, Lorenzo Fagiano, and Paul Goulart. Colin Jones has introduced me to the algorithms of Model Predictive Control (MPC) and the means and ways to make them work at “ludicrous speed”. Lorenzo Fagiano has brought the “scenario approach” to my attention and worked with me on its integration into MPC. Paul Goulart has helped me in exploring some of the depths in robust and stochastic control and semidefinite programming. I am grateful to all of them for their advice on my research, and for sharing an invaluable amount of knowledge, insights, and ideas with me.

It was also a privilege to have worked with various remarkable colleagues and students during my time at IfA. I have enjoyed collaborating with Frauke Oldewurtel, David Sturzenegger, Xiaojing Zhang, Stefan Deml, and Robin Franz on new approaches for a more comfortable and energy efficient building climate control. I have gained many new insights from the meetings with Melanie Zeilinger and Alexander Domahidi on fast optimization and real-time MPC. It was a pleasure working with Joe Warrington and Julius Pfommer on optimizing incentive schemes to induce a favorable behavior of customers in shared transportation networks.

Acknowledgments

My teaching duties over the course of eight semesters turned out to be a truly rewarding experience. The work among the assistant teams of Control Systems I (Prof. Manfred Morari, Alexander Domahidi, Dave Ochsenbein, Joe Warrington, David Sturzenegger, Christian Conte, Manfred Quack, Andreas Hempel, Xiaojing Zhang, Giampaolo Torrisi, Marcello Colombino) and Control Systems II (Prof. Roy Smith, Thomas Besselmann, Robert Nguyen, Peyman Mohajerin, Alexander Fuchs, Marko Tanaskovic, Marcello Colombino, Xiajing Zhang) was productive and cheerful at all times. All courses gave me the chance of interacting with many talented and energetic students, who ended up teaching me more than they have learned from me.

Furthermore, I thank my office mates Stefan Almér, Tobias Sutter, Peyman Mohajerin, Robin Vujanic and other lab members Claudia Fischer, Joe Warrington, Manfred Quack, David Sturzenegger, Stefan Richter, Bart van Parys, Sergio Grammatico, Xiaojing Zhang, Alexander Domahidi, Alexander Liniger, Tyler Summers, Andreas Hempel, Alexander Fuchs, Martin Herceg, Christian Conte, Aldo Zraggen, Stephan Huck, Robert Nguyen, Alberto Busetto, Maryam Kamgarpour, Jakob Ruess, Marcello Colombino, Marko Tanaskovic, Tony Wood, and Sean Summers for making IfA into the inspiring and exciting place that it is!

It was a pleasure for me to participate in the EMBOCON project, where I have met so many interesting and visionary people. Moritz Diehl, Sebastian Sager, Eric Kerrigan, Paul Goulart, Ion Necoara, and Colin Jones deserve a particular thanks for pulling the project together. It was also very pleasant to meet and collaborate with Christian Kirches, Janick Frasch, Joachim Ferreau, Juan Jerez, Stefano Longo, Leonard Wirsching, Michael Engelhart, Milan Vukov, Peter Kühn, Christian Schoppmeyer, Sergio Lucia, Stijn de Bruyne, and Valentin Nedelcu at various places, from Leuven over Heidelberg to Bucharest.

But, most of all, I am indebted to my family. The biggest thanks go to Frauke for always being there for me, and to my parents, my brother, and my grandparents for their support during my entire life. My parents taught me the most essential things in life and they stood behind me in all of my endeavors. They have always shown the greatest respect for my personal way of life. Therefore I gratefully dedicate this dissertation to them.

Financial Support

My doctorate has received funding from the European Seventh Framework Programme FP7/2007-2013 under grant agreement number FP7-ICT-2009-4 248940 (project EMBOCON).

Contents

A. Introduction and Background	3
1. Introduction	5
1.1 Motivation	5
1.2 Outline and Contributions	6
1.3 Publications	8
2. Background	9
2.1 The Scenario-Based Optimization Approach	9
2.2 Multi-Stage Stochastic Programs	17
2.3 Model Predictive Control	24
2.4 Uncertainty in Model Predictive Control	28
References	40
 B. Structure in Multi-Stage Stochastic Programs	 47
Paper I. Randomized Solutions to Convex Programs with Multiple Chance	
Constraints	49
1. Introduction	50
2. Problem Formulation	53
3. Structural Properties of the Constraints	57
4. Feasibility of the Scenario Solution	65
5. The Sampling-and-Discarding Approach	68
6. Example: Minimal Diameter Cuboid	72
Acknowledgments	76
References	76
 C. Scenario-Based Model Predictive Control	 79
Paper II. The Scenario Approach for Stochastic Model Predictive Control with Bounds on Closed-Loop Constraint Violations	81

1.	Introduction	82
2.	Optimal Control Problem	83
3.	Scenario-Based Model Predictive Control	85
4.	Problem Structure and Sample Complexity	90
5.	Numerical Example	97
6.	Conclusion	99
A.	Proof of Lemmas C.1 and C.2	100
	Acknowledgments	102
	References	102
	D. Application to Supply Chain Management	105
	Paper III. Scenario-Based Model Predictive Control for Multi-Echelon Supply Chain Management	107
1.	Introduction	108
2.	Supply Chain Model	111
3.	Scenario-Based Model Predictive Control	115
4.	Case Study	122
5.	Conclusion	127
	References	127
	E. Risk-Averse Two-Stage Stochastic Programs	131
	Paper IV. A Scenario Approach for Two-Level Stochastic Programs with Expected Shortfall Probability	133
1.	Introduction	134
2.	A Review of the Classic Scenario Approach	137
3.	The Scenario Approach for Two-Stage Stochastic Programs	141
4.	A Farmer's Problem	145
5.	Conclusion	148
A.	Proof of Theorems E.3 and E.4	149
	References	151
	Curriculum Vitae	155

Abstract

Technological advancements over the past decades have increased the availability of ever more powerful and inexpensive hardware. This development has caused a substantial shift of research focus from classic control theory to advanced, optimization-based control methods. The latter are characterized by the fact that the decisions about the control actions are obtained by solving a numerical optimization program. In particular, *model predictive control* (MPC) offers an effective approach for handling multivariable control problems with a defined stage cost criterion and constraints on the inputs, states, and outputs.

The main contribution of this dissertation is the development of a novel method of *scenario-based MPC* (SCMPC) for handling *multi-stage stochastic decision problems* in a receding-horizon fashion. Indeed, MPC originally assumes that an exact model of the control system is available and there are no unknown disturbances, so it can accurately predict the system's state trajectory. However, uncertainty in these predictions can lead to substantial constraint violations and a significant performance degradation (in terms of stage costs) for the system in closed-loop operation.

Various approaches to cope with uncertainty in MPC have previously been proposed. *Robust MPC* (RMPC) considers uncertainties contained inside a pre-fitted uncertainty set. For systems with stochastic disturbances, however, RMPC may result in a sub-optimal performance. The reason is that this uncertainty model contains no probabilistic information and the decisions of RMPC are often based on extreme and unlikely disturbance realizations. *Stochastic MPC* (SMPC) approaches account for a probability distribution of the uncertainty. The constraints are typically relaxed in a probabilistic sense (e.g., as *chance constraints*), in exchange for an improved performance. This turns out to be a reasonable choice for many practical applications, where performance is critical. General distribution functions, however, are not amenable to numerical computations. Therefore many SMPC approaches are either computationally very demanding, or they are specialized to uncertainties of a particular distribution type (e.g., a normal distribution).

The novel SCMPC method provides an alternative to SMPC, using sampled uncertainty scenarios of an arbitrary stochastic model (as opposed to explicit probability distributions). The number of scenarios is determined a priori, such that controller satisfies a given set of chance constraints on the system state. Compared to similar approaches that

have previously been proposed, the novel SCMPC method requires a significantly lower number of scenarios. This reduces the computational complexity and improves the performance of the controller. Moreover, examples show that the desired level of constraint violations can accurately be achieved.

The development of SCMPC in this dissertation is the result of multiple contributions. First, existing results in *scenario-based optimization* have established a direct link between the number of scenarios and bounds on the probability of constraint violations. These results are extended to problems with multiple chance constraints. Moreover, the existing bounds on the probability of constraint violations are improved in cases where a chance constraint has a limited *support rank*. The support rank is a novel concept defined in this thesis. The presented theory is applicable to very general stochastic optimization problems, in particular arising from multi-stage stochastic decision problems. Moreover, it potentially leads to a significant reduction in the number of scenarios, as compared to the previous theory.

Second, the theory for a novel SCMPC method is introduced, with a focus on its mathematical properties. The theory builds on the results of the first contribution, and it additionally provides a new framework for analyzing the behavior of the closed loop under SCMPC. In contrast to previous SCMPC approaches, this framework allows for the chance constraints to be interpreted as the time-average of state constraint violations, rather than a joint probability over an open-loop prediction horizon. This leads to a potentially massive reduction in the number of required scenarios. Furthermore, the novel SCMPC approach features the possibility of sample removal (as known from scenario-based optimization), and it is compatible with previously considered methods of disturbance feedback for closed-loop predictions.

Third, a possible implementation of SCMPC is examined in an extensive case study. The case study considers a networked supply chain distribution system with multiple products and uncertainty in the demands. It is shown that SCMPC is able to keep the prescribed service level constraints and significantly reduces the inventory holding costs. At the same time, SCMPC is computationally efficient for large-scale, complex problems with high-dimensional, correlated uncertainties. This type of problem can often not adequately be handled by means robust or stochastic optimization.

Fourth, a new implementation of the scenario approach is presented for risk averse solutions to two-stage stochastic decision problems. Instead of a conventional mean-risk optimization, the new approach optimizes the stochastic objective function value with respect to a maximal shortfall probability. The advantage of this approach is its ability to handle high-dimensional uncertainties of a very general nature in a computationally efficient manner. Its application is demonstrated for a particular version of the farmer's problem.

Zusammenfassung

Der technologische Fortschritt der letzten Jahrzehnte hat zu einer erhöhten Verfügbarkeit leistungsfähiger und kostengünstiger Hardware geführt. Diese Entwicklung hat zu einer erheblichen Verschiebung des Forschungsschwerpunktes in der Regelungstechnik beigetragen, weg von der klassischen Theorie und hin zu modernen, optimierungsbasierten Regelungsmethoden. Letztere sind dadurch gekennzeichnet, daß die Entscheidungen über die Steuergrößen auf der Lösung eines numerischen Optimierungsproblems basieren. Insbesondere die *modellprädiktive Regelung* (engl. *model predictive control*, MPC) bietet einen effektiven Ansatz zur Regelung von Mehrgrößensystemen mit einem definierten Gütekriterium und unter Beschränkungen der Steuergröße, des Zustands sowie der Regelgröße.

Der Hauptbeitrag dieser Dissertation ist die Entwicklung einer neuen Methode der *Szenario-basierten modellprädiktiven Regelung* (engl. *scenario-based MPC*, SCMPC) zur Anwendung auf *mehrstufige stochastische Entscheidungsprobleme* mit rollierendem Horizont. MPC basiert nämlich auf der Annahme, daß ein exaktes Modell der Regelstrecke zur Verfügung steht und keine zufälligen Störungen von außen auftreten, sodaß die Zustandstrajektorie des betrachteten Systems sicher und genau vorhergesagt werden kann. Jedoch können Unsicherheiten in der Vorhersage zu substantiellen Verletzungen der Beschränkungen und einer signifikanten Verschlechterung des Gütekriteriums im geschlossenen Regelkreis führen.

Es existieren bereits verschiedene Ansätze zum Umgang mit diesen Unsicherheiten in MPC. *Robuste modellprädiktive Regelung* (engl. *robust MPC*, RMPC) betrachtet Unsicherheiten aus einer zuvor festgelegten Unsicherheitsmenge. Für Systeme mit stochastischen Störungen kann RMPC jedoch zu einer schlechten Regelgüte führen. Der Grund hierfür ist, daß dieses Unsicherheitsmodell keine probabilistischen Informationen enthält und sich die Entscheidungen von RMPC häufig an extremen und unwahrscheinlichen Ausprägungen der Störungen orientieren. Diese treten allerdings in der Realität nur mit einer geringen Wahrscheinlichkeit auf. *Stochastische modellprädiktive Regelung* (engl. *stochastic MPC*, SMPC) bezieht die Wahrscheinlichkeitsverteilungen mit in die Entscheidungen ein. Die Beschränkungen werden typischerweise in einem stochastischen Sinne relaxiert (z.B. in Form von *chance constraints*), um dadurch eine bessere Regelgüte zu erzielen. Dies stellt sich als ein sinnvoller Ansatz für eine ganze Reihe von praktischen Anwendungen heraus, in denen die Regelgüte von entscheidender Bedeutung ist. Allgemeine

Verteilungsfunktionen sind jedoch für numerische Berechnungen schlecht geeignet. Daher sind die meisten Methoden des SMPC entweder sehr rechenaufwändig, oder aber sie sind spezialisiert auf Unsicherheiten mit spezieller Wahrscheinlichkeitsverteilung (z.B. einer Normalverteilung).

Die neu entwickelte SCMPC Methode bietet eine Alternative zu SMPC, indem sie eine Stichprobe möglicher Szenarien betrachtet. Diese Stichprobe kann durch ein beliebiges stochastisches Modell der Unsicherheit erzeugt werden (und benötigt also keine explizite Wahrscheinlichkeitsverteilung). Die genaue Größe der Stichprobe wird so bestimmt, daß der Regler eine gegebene Menge von stochastischen Beschränkungen („chance constraints“) einhält. Im Vergleich zu anderen Ansätzen dieser Art benötigt die neu entwickelte SCMPC Methode dazu eine erheblich geringere Anzahl an Szenarien. Dies führt zum einen zu einer beträchtlichen Reduktion des Rechenaufwands und zum anderen einer verbesserten Regelgüte. Anhand von Beispielen kann zudem gezeigt werden, daß die stochastischen Beschränkungen voll ausgeschöpft werden können.

Die Entwicklung der neuen SCMPC Methode in dieser Dissertation basiert auf verschiedenen Forschungsergebnissen. Erstens, vorherige Ergebnisse auf dem Gebiet der *Szenario-basierten Optimierung* haben eine mathematische Verbindung zwischen der Stichprobengröße (für die Szenarien) und der Verletzungswahrscheinlichkeit einer Beschränkung hergeleitet. Diese Ergebnisse werden nun erweitert auf den Fall mit mehreren stochastischen Beschränkungen. Weiterhin können die vorherigen Formeln verbessert werden für den Fall, daß die stochastischen Beschränkungen einen begrenzten *Stützrang* (engl. *support rank*) besitzen. Der Stützrang ist ein neues Konzept, das in dieser Arbeit definiert wird. Die dadurch entwickelte Theorie ist sehr weitreichend in der stochastischen Optimierung anwendbar, insbesondere aber für *mehrstufige stochastische Entscheidungsprozesse* (engl. *multi-stage stochastic decision problems*). Sie kann zu einer signifikanten Reduktion der Stichprobengröße, verglichen mit der vorherigen Theorie, führen.

Zweitens wird die grundlegende Theorie der neuen SCMPC Methode beschrieben, mit Fokus auf ihre mathematischen Eigenschaften. Die Vorgehensweise stützt sich dabei auf die obigen Ergebnisse und schafft zudem einen neuen Rahmen für die Analyse des geschlossenen Regelkreises unter SCMPC. Im Gegensatz zu früheren Methoden können die stochastischen Beschränkungen im Sinne der durchschnittlichen Häufigkeit von Verletzungen der Zustandsbeschränkungen interpretiert werden, anstatt als kumulierte Wahrscheinlichkeit über einen Planungshorizont. Dies führt zu einer potentiell massiven Reduktion der benötigten Anzahl an Szenarien. Des weiteren beinhaltet die neue SCMPC Methode die Möglichkeit des Ausschlusses von Szenarien (engl. *sample removal*) und sie ist vollständig kompatibel mit existierenden Verfahren zur Vorhersage unter Störgrößenrückführung (engl. *disturbance feedback*).

Drittens wird eine mögliche Anwendungsform der neuen SCMPC Methode in einer ausgiebigen Fallstudie untersucht. Die Fallstudie betrachtet die Steuerung des Verteilungsnetzwerks einer Logistikkette (engl. *supply chain*), mit mehreren Produkten und Nachfrageunsicherheit. Es zeigt sich, daß SCMPC den vorgeschriebenen Liefergrad einhalten kann und die Lagerhaltungskosten signifikant reduziert. Gleichzeitig bietet SCMPC

recheneffiziente Lösungen für große und komplexe Systeme mit hochdimensionalen, korrelierten Unsicherheiten. Diese Probleme können oftmals nicht adäquat mittels robuster oder stochastischer Optimierung gelöst werden.

Viertens wird ein neues Verfahren der Szenario-basierten Optimierung vorgestellt, welches zur Lösung von zweistufigen stochastischen Entscheidungsproblemen unter Risikoaversion eingesetzt werden kann. Statt der herkömmlichen Methode der Einbeziehung eines Riskomaßes in die Zielfunktion, optimiert der neue Ansatz die stochastische Zielfunktion mit einer bestimmten Unterschreitungswahrscheinlichkeit. Der Vorteil dieses Ansatzes liegt in der Behandlung von Problemen mit sehr allgemeinen, hochdimensionalen Unsicherheiten mittels einer recheneffizienten Prozedur. Die Anwendung der Methode wird anhand einer speziellen Version des *Farmer's Problem* demonstriert.

Part A

Introduction and Background

1

Introduction

1.1 Motivation

Over the past decades, technological advancements in computational hardware have led to the increasing availability of computation power at lower costs. Furthermore, available algorithms for numerical optimization have become faster and better at handling problems of large scale. These developments have caused a natural shift in research focus from classic control theory to *optimization-based control methods*, where the control actions are computed by numerical optimization [72]. In particular, the concept of *model predictive control* (MPC) has received increasing attention, and its theory has matured considerably over the past two decades [26, 67, 69, 83].

MPC is an effective approach for multivariable control problems in which performance is measured by a *stage cost criterion* and *constraints* must be observed on the control inputs, the system states, and the regulated outputs. While originally designed for chemical processes, MPC has now extended its scope to a wide range of applications [80]. The fundamental idea of MPC is to formulate a *finite-horizon optimal control problem* (FHOC), based on the stage cost criterion and the constraints. The control actions over the *prediction horizon* are computed by solving the FHOC numerically on-line. Only the first computed control action is implemented; the FHOC is then updated and re-solved in the next sampling time step. Hence, MPC represents a (generally non-linear) state feedback law for the system.

Leveraging the increasingly powerful hardware platforms and optimization algorithms, it is reasonable to predict that the use of optimization-based control methods will expand even further in the future. Particularly attractive targets are large-scale and complex systems where the stage cost criterion translates directly into physical or monetary costs; for instance, energy consumption in buildings [76] or inventory cost in supply chains [95]. Due to the complexity of these applications, conventional control methods are difficult to tune and often have a sub-optimal performance. In fact, these types of systems are the guiding motivation for the theory of this dissertation.

Many open questions remain about optimization-based control algorithms from the viewpoint of control theory. One key issue is the proper handling of uncertainty in the

underlying optimization program; i.e., the FHOC. This uncertainty may result from a variety of sources, such as inaccurate models, unknown parameters, random disturbances, or incomplete state information. In general, the proper representation of uncertainty in optimization programs is a non-trivial issue and subject to intensive, ongoing research. In the context of MPC, uncertainty poses an even greater challenge because the FHOC must be solved within a small fraction of the sampling time.

The use of particles (or “*scenarios*”) is a promising approach for integrating uncertainty into stochastic decision problems. Particle-based algorithms have been successfully implemented, for example, in approximations to multi-stage stochastic programs [21, 57, 99] and in recursive state estimation [44].

In this thesis, recent breakthroughs in *scenario-based optimization* [29, 34, 35] (also known as the “*scenario approach*”) are cast into a novel method for multi-stage stochastic decision problems. The main contribution of this dissertation is the development of a new *scenario-based MPC* (SCMPC) approach, where this novel method is applied in a receding-horizon fashion.

Depending on the particular control problem, the SCMPC approach may offer a variety of advantages over other MPC approaches. Compared to approaches of *robust MPC* (RMPC), where the uncertainty is bounded inside a pre-fitted uncertainty set, SCMPC accounts for the likelihood with which uncertainty scenarios occur. Compared to approaches of *stochastic MPC* (SMPC), where the uncertainty model is given by a probability distribution, SCMPC is purely data-based. Hence the numerical computations are simplified, in all but a very few exceptions (for example, when the joint uncertainty distribution is normal). And finally, SCMPC solves a convex optimization program whose scale is only modestly higher than that of *certainty-equivalent MPC*, where a fixed value for the uncertainty is assumed.

An exemplary application of the new SCMPC method is presented in the case study of a typical networked supply chain distribution system. SCMPC shows a good performance in the operation of the system, whose scale (several thousand decision variables) and complexity (time-varying, correlated, high-dimensional uncertainties) preclude a straightforward application of robust or stochastic optimization.

1.2 Outline and Contributions

The topics in this dissertation are clustered around multiple steps in extending the *scenario-based optimization approach* to *multi-stage stochastic decision problems*. The remainder of Chapter 1 outlines the main contributions and provides a list of scientific publications that I have prepared over the course of my doctorate. Chapter 2 covers some essential background material, introducing the basic concepts and methods from related research fields.

The following parts of this dissertation contain four main publications. The ordering goes from the abstract, general results to the more specific, concrete applications.

Part B presents new theoretical extensions to the classic scenario-based optimization approach (*scenario approach*) [29,34,35]. New results are obtained for uncertain problems with *multiple* chance constraints as they appear, for example, in multi-stage stochastic decision problems. Due to the problem structure, each chance constraint often has a limited *support rank*—a new concept that is introduced in the paper. This structural property allows for a reduction in the corresponding sample size, improving on the bounds previously derived for the scenario approach.

Part C leverages the results from Part B to develop a novel predictive control algorithm for stochastic linear systems, called *scenario-based MPC* (SCMPC). The method is based on random samples of the uncertainty over the prediction horizon (“*scenarios*”), and provides bounds on the probability of constraint violations (“*chance constraints*”). Compared to previous approaches of combining MPC with the scenario approach [32,68,77,89], the new SCMPC algorithm considers the constraint violations by the closed-loop system as an average over time, as opposed to a joint violation probability for all constraints in the FHOCP. Hence the sample size is significantly reduced, thereby lowering the computational efforts and improving the controller performance. Furthermore, the SCMPC method features the possibility of scenario constraint removals, as known from scenario-based optimization. This allows for further improvements of the controller performance, in exchange for a higher computational complexity. SCMPC is also compatible with previously considered methods of disturbance feedback for closed-loop predictions.

Part D applies the SCMPC algorithm towards the operation of a typical supply chain distribution system. As shown in this case study, SCMPC can handle supply chains with stochastic planning uncertainties of a very general type (demands, lead times, prices, etc.) and nature (distributions, correlations, etc.), while guaranteeing a specified customer service level. Moreover, SCMPC is applicable to problems of a similar scale as manageable by deterministic optimization. Hence it may be applied to problem instances that can not be adequately treated by robust or stochastic optimization. Many of the practical advantages and drawbacks of the SCMPC can be observed for this particular case study, hinting at other potential application areas for SCMPC.

Part E considers the special class of *two-stage stochastic decision problems*. Existing approaches of stochastic programming are able to efficiently optimize the first-stage decision in terms of the expected payoff. In many practical applications, however, some robustness of the solution to downside risk is a desirable feature. Further details and potential applications are provided in the paper. Previous approaches have therefore tried to include a risk measure into the objective function, which may seriously complicate its solution. It is shown that the scenario approach can be applied quite naturally for producing high quality first-stage decisions with a specified level of a shortfall probability. The effectiveness of this approach is demonstrated for a particular version of the Farmer’s Problem [21, Sec. 1.1].

1.3 Publications

The results presented in this thesis were obtained in close cooperation with various colleagues. They are mostly contained in the following publications.

The introductory part is partly based on the following conference publications:

- [96] G. Schildbach, M.N. Zeilinger, M. Morari, and C.N. Jones. Input-to-state stabilization of low-complexity model predictive controllers for linear systems. In *50th IEEE Conference on Decision and Control*, Orlando (FL), United States, 2011.
- [89] G. Schildbach, G.C. Calafiore, L. Fagiano, and C.N. Jones. Randomized model predictive control for stochastic linear systems. In *American Control Conference*, Montréal, Canada, 2012.

Part B has been published as the article:

- [91] G. Schildbach, L. Fagiano, and M. Morari. Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4):2479–2501, 2013.

Part C has been provisionally accepted for publication as the article:

- [90] G. Schildbach, L. Fagiano, C. Frei, and M. Morari. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, (under review).

Part D has been submitted as the manuscript:

- [95] G. Schildbach and M. Morari. Scenario-based model predictive control for multi-echelon supply chain management. *European Journal of Operational Research*, (submitted).

Part E has been submitted as the manuscript:

- [94] G. Schildbach and M. Morari. The scenario approach for two-level stochastic programs with expected shortfall probability. *International Journal of Production Economics*, (submitted).

The following publications have been prepared during my doctorate, but are not a part of this dissertation:

- [92] G. Schildbach, P. Goulart, and M. Morari. The linear quadratic regulator with chance constraints. In *12th European Control Conference*, Zürich, Switzerland, 2013.
- [93] G. Schildbach, P. Goulart, and M. Morari. Linear controller design for chance constrained systems. *Automatica*, (submitted).

2

Background

This chapter reviews essential background material from areas that are closely related to the contents of this dissertation. In particular, the chapter covers the scenario-based optimization approach (Section 2.1), multi-stage stochastic programming (Section 2.2), model predictive control (Section 2.3), and uncertainty in model predictive control (Section 2.4).

It is beyond the scope of this chapter to review all of these areas rigorously and in all detail. The goal is rather to build up a basic structure and framework, which can not be provided within the research results (Parts B,C,D,E), due to the usual brevity restrictions. The reader shall be equipped with the most fundamental methods and results, and pointed to the relevant literature for further details.

2.1 The Scenario-Based Optimization Approach

This section contains a brief review of the *scenario-based optimization approach* (scenario approach). Research on the scenario approach started with the paper of Calafiore and Campi [29], and achieved a breakthrough with the work of Campi and Garatti [34, 35]. The reader is referred to the above references for further details.

Uncertain Convex Program

The scenario approach considers an *uncertain convex program* $\text{UP}[\varepsilon]$ of the form

$$\text{UP}[\varepsilon] : \quad \min_y \quad c^T y \quad (\text{A.1a})$$

$$\text{s.t.} \quad \mathbf{P}[f(y, \delta) \leq 0] \geq 1 - \varepsilon, \quad (\text{A.1b})$$

$$y \in \Omega. \quad (\text{A.1c})$$

Here $y \in \mathbb{R}^d$ denotes the *decision vector* and $c \in \mathbb{R}^d$ is a vector defining a linear *objective function*. The decision vector must be chosen optimally from a compact¹ and convex²

¹*Compactness* is fundamentally a topological property [71, 104]; in a Euclidean space, it is equivalent to *closedness* and *boundedness* (Heine-Borel Theorem).

²*Convexity* is fundamentally a vector space property [46, 65, 107].

domain $\Omega \subset \mathbb{R}^d$.

The variable δ comprises all uncertain quantities in $\text{UP}[\varepsilon]$, whose *sample space* Δ is of an entirely generic nature (e.g., a vector space). The constraint (A.1b) is formulated as a *chance constraint*³, containing a *constraint function* $f : \mathbb{R}^d \times \Delta \rightarrow \mathbb{R}$ of both the decision variable y and the uncertainty δ . It must be kept with a *probability level* of at least $1 - \varepsilon$, where $\varepsilon \in (0, 1)$. The following assumptions about the nature of the uncertainty δ and the constraint function f are made throughout.

ASSUMPTION 1—UNCERTAINTY (a) There exists a probability measure \mathbf{P} on Δ ; i.e., δ is a *random variable*.⁴ (b) The probability measure \mathbf{P} (alternatively, the distribution of δ) may be unknown, but a sufficient number (to be made precise later on) of independent random samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$ are available. ■

ASSUMPTION 2—CONSTRAINT FUNCTION The constraint function $f(\cdot, \delta)$ is a convex function for almost every uncertainty $\delta \in \Delta$ (i.e., except on some subset of Δ with zero probability measure).⁵ ■

The short-hand notation for the probability measure of subsets of Δ ,

$$\mathbf{P}[f(y, \delta) \leq 0] := \mathbf{P}\{\delta \in \Delta \mid f(y, \delta) \leq 0\} , \quad (\text{A.2})$$

as used in (A.1b), is common and shall be used in the sequel. In other words, for a decision $y \in \Omega$ to be feasible in the sense of the chance constraint, the condition “ $f(y, \delta) \leq 0$ ” has to characterize a subset of Δ whose probability measure is at least $(1 - \varepsilon)$. Throughout the thesis, the complications arising from a σ -algebra of *measurable sets* are ignored; that is, any stated *event* (i.e., a subset of Δ) is assumed to be measurable.⁶

The use of “min” instead of “inf” in (A.1a) is justified by the fact that the feasible set of a chance constraint can be shown to be closed for very general cases [57, Thm. 2.1]. The feasible set of (A.1) is hence compact, by the presence of Ω , and any infimum is indeed attained; otherwise “min” can always be replaced by “inf” for finding an optimal point in the closure of the feasible set.

Assumption 1 is very general, since it requires only sufficient data as a knowledge of δ . Note that the formulation of $\text{UP}[\varepsilon]$ comprises all uncertain optimization programs that become convex if the uncertainty variable δ were known and fixed. In particular, the objective function may be an arbitrary convex function [34, 91] and the constraint function covers *joint chance constraints* (i.e., the maximum of multiple linear constraint functions) [21, 57, 99].

³Chance constraints are a fundamental concept in the field of *stochastic programming*; see [21, 57, 99] for more details.

⁴Familiarity of the reader with basic *probability theory* is assumed; excellent monographs are, e.g., [20, 100].

⁵A real-valued function defined on a vector space (more precisely, a *functional*) is said to be *convex* if its *epigraph* is a convex set [28, 65, 86].

⁶In uncountable sample spaces, the basic axioms for a *measure* may inevitably lead to the existence of *non-measurable sets* [2–4, 24].

Despite of the convexity of the constraint function $f(\cdot, \delta)$, the feasible set of chance constraint (A.1b) is generally non-convex [21, 57, 99]. Therefore, $\text{UP}[\varepsilon]$ is difficult to solve, even if the distribution of δ is known.⁷

The Scenario Program

The scenario approach provides a computationally efficient approximation to $\text{UP}[\varepsilon]$, based on the optimal solution to the *scenario program* $\text{SP}[\omega^{(\mathcal{K})}]$,

$$\text{SP}[\omega^{(\mathcal{K})}] : \quad \min_w \quad c^T y \quad (\text{A.3a})$$

$$\text{s.t.} \quad f(y, \delta^{(k)}) \leq 0 \quad \forall k = 1, 2, \dots, K, \quad (\text{A.3b})$$

$$y \in \Omega. \quad (\text{A.3c})$$

In $\text{SP}[\omega^{(\mathcal{K})}]$, the chance constraint of $\text{UP}[\varepsilon]$ has been replaced by $\mathcal{K} := \{1, \dots, K\}$ fixed constraints, namely by substituting the samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$ of the uncertainty into the constraint function $f(y, \cdot)$. For notational convenience, the samples are also denoted as a *multi-sample* $\omega^{(\mathcal{K})} := \{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}\}$. They can be interpreted as *training samples* or *scenarios* for the solution of $\text{SP}[\omega^{(\mathcal{K})}]$, which is called the *scenario solution* and denoted $y^*(\omega^{(\mathcal{K})})$.

Existence and uniqueness of $y^*(\omega^{(\mathcal{K})})$ is assumed without any loss of generality [34, Sec. 2 (5)]; a generalization of the presented theory accounting for infeasibility can be developed as in [31].

ASSUMPTION 3—EXISTENCE AND UNIQUENESS (a) The $\text{SP}[\omega^{(\mathcal{K})}]$ admits a feasible point almost surely. By the compactness of Ω , this implies that the problem has at least one optimal point. (b) If there exist multiple optimal points of $\text{SP}[\omega^{(\mathcal{K})}]$, a unique one is selected by the use of a tie-break rule. ■

A *tie-break rule* is any decision rule that uniquely selects a single point from a compact subset of \mathbb{R}^d ; e.g., the *lexicographic order*. Indeed, in practical applications it does not matter much which optimal solution is selected, as long as the objective function is minimized. However, a tie-break rule may be difficult to implement in practice, because most algorithms do not provide the entire set of optimal solutions. Instead, small perturbations can be applied to the problem [56].

In practical applications, the scenario solution $y^*(\omega^{(\mathcal{K})})$ is obtained after the outcomes of the training samples $\omega^{(\mathcal{K})}$ are observed. Hence $\text{SP}[\omega^{(\mathcal{K})}]$ is a *deterministic* convex optimization program of a pre-defined type; e.g., a linear or quadratic program. Efficient numerical algorithms exist for the solution of such problems, even in high dimensions [16, 28, 74].

⁷Familiarity of the reader with the basic theory of *optimization*, in particular *convex optimization*, is assumed; excellent monographs are, e.g., [15, 28, 66, 74].

The Sampling Theorem

The generalization properties of the scenario solution $y^*(\omega^{(\mathcal{K})})$, with respect to the chance constraint of the original UP $[\varepsilon]$, lead into the theory of the scenario approach. In particular, the theory establishes a link between the *sample size* K and the probability of $y^*(\omega^{(\mathcal{K})})$ violating the chance constraint (A.1b):

$$v(\omega^{(\mathcal{K})}) := \mathbf{P}[f(y^*(\omega^{(\mathcal{K})}), \delta) > 0] . \quad (\text{A.4})$$

Note that, for the purposes of this analysis, the scenario solution $y^*(\omega^{(\mathcal{K})})$ and the *violation probability* $v(\omega^{(\mathcal{K})})$ are considered as (unknown) functions of the random multi-sample $\omega^{(\mathcal{K})}$. Therefore, two levels of probability have to be considered for the theory of the scenario approach: The first is introduced by the random training samples $\omega^{(\mathcal{K})}$, affecting the choice of $y^*(\omega^{(\mathcal{K})})$. The second is the actual uncertainty δ , which determines whether $y^*(\omega^{(\mathcal{K})})$ satisfies the original chance constraint.

To highlight the two probability levels more clearly, suppose for the moment that the multi-sample has already been observed. Let $\bar{\omega}^{(\mathcal{K})}$ denote its outcome, and $\bar{y} := y^*(\bar{\omega}^{(\mathcal{K})})$ the corresponding scenario solution. Then the *a posteriori violation probability* $\bar{v} := v(\bar{\omega}^{(\mathcal{K})})$ is a deterministic, albeit unknown, value in the interval $[0, 1]$:

$$\bar{v} := \mathbf{P}[f(\bar{y}, \delta) > 0] . \quad (\text{A.5})$$

Now suppose that the multi-sample has not yet been observed. Then the *a priori violation probability* $v(\omega^{(\mathcal{K})})$, as defined in (A.4), is itself a random variable with support $[0, 1]$. It is defined on the probability space (Δ^K, \mathbf{P}^K) , where Δ^K and \mathbf{P}^K denote the K -th product space of Δ and the K -th product measure of \mathbf{P} , respectively.

The following fundamental result is due to Campi and Garatti [34, Thm. 2.4].

THEOREM 1—DISTRIBUTION BOUND Suppose Assumptions 1 and 3 hold. Then the probability distribution of the violation probability of SP $[\omega^{(\mathcal{K})}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(\mathcal{K})}) > \nu] \leq B(\nu; K, d - 1) , \quad (\text{A.6})$$

for any $\nu \in [0, 1]$, where

$$B(\nu; K, d - 1) := \sum_{j=0}^{d-1} \binom{K}{j} \nu^j (1 - \nu)^{K-j} \quad (\text{A.7})$$

denotes the *Beta Distribution Function* (see the last part this section), with parameters d (the dimension of the decision variable) and K (the sample size). ■

Theorem 1 shows that there exists an upper bound for all quantiles of the probability distribution of $v(\omega^{(\mathcal{K})})$. Moreover, this upper bound is tight over the entire support $[0, 1]$; i.e., there exists a class of uncertain convex programs (those which are *fully supported*

almost surely, [34, Def. 2.3]) for which (A.6) holds with equality [34, Sec. 2.1].

Given the upper bound for the cumulative distribution (A.6), the upper tail of $v(\omega^{(\mathcal{K})})$ can be bounded by a simple *regula falsi* procedure. It is also possible to use Chernoff bounds [41] for bounding the upper tail of $v(\omega^{(\mathcal{K})})$; see [30, Rem. 2.3] for details.

COROLLARY 1—EXPLICIT TAIL BOUND Suppose Assumptions 1 and 3 hold. If the sample size is selected according to

$$K \geq \frac{2}{\varepsilon} \left(\log \frac{1}{\theta} + d - 1 \right) , \quad (\text{A.8})$$

where $\log(\cdot)$ denotes the natural logarithm and $\varepsilon \in (0, 1)$, $\theta \in (0, 1)$ are fixed, then the violation probability satisfies

$$\mathbf{P}^K [v(\omega^{(\mathcal{K})}) \geq \varepsilon] \leq \theta . \quad (\text{A.9})$$

■

Furthermore, Theorem 1 allows to compute an upper bound on its expectation by integrating the upper bound of the distribution function (A.6):

$$\begin{aligned} \mathbf{E}[v(\omega^{(\mathcal{K})})] &= \int_0^1 \mathbf{P}^K [v(\omega^{(\mathcal{K})}) > \nu] d\nu \\ &\leq \int_0^1 B(\nu; K, d - 1) d\nu = \frac{d}{K + 1} . \end{aligned} \quad (\text{A.10})$$

This yields the main result of Calafiore and Campi [29, Thm. 1] as a simple corollary of Theorem 1.

COROLLARY 2—EXPECTATION BOUND Suppose Assumptions 1 and 3 hold. Then the expected violation probability of $\text{SP}[\omega^{(\mathcal{K})}]$ satisfies

$$\mathbf{E}^K [v(\omega^{(\mathcal{K})})] \leq \frac{d}{K + 1} . \quad (\text{A.11})$$

■

The Sample and Removal Theorem

Considering the randomized nature of the scenario solution, it may be desirable to reduce the exposure of the scenario solution to extreme outliers in the samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$. To this end, it is possible to deliberately increase the sample size K above its minimal value of Theorem 1, in exchange for being allowed to remove R samples *a posteriori* (i.e., after the sample values have been observed). The samples must be removed by a valid *removal procedure*, as defined below; cf. [35, Ass. 2.2].

DEFINITION 1—REMOVAL PROCEDURE A removal procedure is an algorithm $\mathcal{A}_{K,R} : \Delta^K \rightarrow \Delta^{K-R}$ that selects R of the K samples, $\mathcal{R} \subset \mathcal{K}$, to be removed a posteriori. The removal procedure is opportunistic in the sense that for the remaining samples $\omega^{(\mathcal{K} \setminus \mathcal{R})}$, the scenario solution $y^*(\omega^{(\mathcal{K} \setminus \mathcal{R})})$ violates all of the removed constraints. ■

Particular removal procedures can be based on *optimal*, *greedy*, or *marginal algorithms*; see e.g., [31, Sec. 5.1] for more details. Similar to Theorem 1, an upper bound on the distribution of the violation probability $v(\omega^{(K,R)})$ after sample removal has been established by Campi and Garatti [35, Thm. 2.1].

THEOREM 2—DISTRIBUTION BOUND WITH SAMPLE REMOVAL Suppose Assumptions 1 and 3 hold. Let $R < K$ sampled constraints be removed from $\text{SP}[\omega^{(\mathcal{K})}]$ by a removal procedure according to Definition 1. Then the distribution of the violation probability of $\text{SP}[\omega^{(\mathcal{K} \setminus \mathcal{R})}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(\mathcal{K} \setminus \mathcal{R})}) > \nu] \leq u_d^{(K,R)}(\nu) , \quad (\text{A.12})$$

for any $\nu \in ([0, 1])$, where

$$u_d^{(K,R)}(\nu) := \min\left\{1, \binom{R+d-1}{d-1} \text{B}(\nu; K, R+d-1)\right\} \quad (\text{A.13})$$

and $\text{B}(\cdot; \cdot, \cdot)$ denotes the Beta distribution function (see the last part this section), with parameters d (the dimension of the decision variable) and K (the sample size). ■

The upper bound on the distribution (A.13) equals the one of Campi and Garatti [35, Thm. 2.1], saturated at 1. The saturation is justified by the fact that $v(\omega^{(\mathcal{K} \setminus \mathcal{R})})$ is itself a probability and can hence be no larger than 1. As for the non-removal case, the upper tail of $v(\omega^{(\mathcal{K} \setminus \mathcal{R})})$ can be bounded by a *regula falsi* procedure, and an explicit bound can be derived as in [31, Sec. 5].

COROLLARY 3—EXPLICIT TAIL BOUND Suppose Assumptions 1 and 3 hold. Let $R < K$ sampled constraints be removed from $\text{SP}[\omega^{(\mathcal{K})}]$ by a removal procedure according to Definition 1. If the sample size K is selected according to

$$K \geq \frac{2}{\varepsilon} \left(\log \frac{1}{\theta} \right) + \frac{4}{\varepsilon} (R+d-1) , \quad (\text{A.14})$$

where $\varepsilon \in (0, 1)$ and $\theta \in (0, 1)$ are fixed, then the violation probability of $\text{SP}[\omega^{(\mathcal{K} \setminus \mathcal{R})}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(\mathcal{K} \setminus \mathcal{R})}) \geq \varepsilon] \leq \theta . \quad (\text{A.15})$$

■

Furthermore, Theorem 2 allows to compute an upper bound on its expectation by integrating the upper bound of the distribution function, analogously to (A.10).

COROLLARY 4—EXPECTATION BOUND Suppose Assumptions 1 and 3 hold. Let $R < K$ sampled constraints be removed from $\text{SP}[\omega^{(K)}]$ by a removal procedure according to Definition 1. Then the expected violation probability of $\text{SP}[\omega^{(K \setminus R)}]$ satisfies

$$\mathbf{E}^K[v(\omega^{(K)})] \leq \int_0^1 u_d^{(K,R)}(\nu) d\nu . \quad (\text{A.16})$$

■

While the expectation bound without sample removal (A.11) comes as a nice explicit formula, the expectation bound with sample removal (A.16) is a one-dimensional integral. Both tail bounds and expectations bounds, however, can be efficiently evaluated by numerical integration, given values for K and R . In order to find appropriate values of K and R , the number of removed constraints R is usually fixed. Then K is computed by a bisection procedure, evaluating the bounds repeatedly for different values of K , and observing that the violation probability monotonically decreases with K . Alternatively, K can be fixed and a bisection procedure can yield the corresponding value of R , observing that the violation probability monotonically increases with R .

Related Probability Distributions

Throughout this thesis, several probability-related functions are used in conjunction with the scenario approach. The *Binomial Distribution Function* [1, Sec. 26.1.20]

$$\Phi(d; K, \nu) := \sum_{j=0}^d \binom{K}{j} \nu^j (1 - \nu)^{K-j} \quad (\text{A.17})$$

expresses the probability of seeing at most $d \in \{0, 1, \dots, K\}$ successes in K independent Bernoulli trials, where the probability of success is $\nu \in [0, 1]$ per trial. If the random variable is not the number of successes, but instead the parameter $\nu \in [0, 1]$, then the same function is called the *Beta Distribution Function*

$$\text{B}(\nu; K, d) := \Phi(d; K, \nu) = \sum_{j=0}^d \binom{K}{j} \nu^j (1 - \nu)^{K-j} . \quad (\text{A.18})$$

The (real) *Beta Function* [1, Sec. 6.2.1]

$$\beta(a, b) := \int_0^1 \xi^{a-1} (1 - \xi)^{b-1} d\xi \quad (\text{A.19})$$

is defined for any parameters $a, b \in \mathbb{R}_+$, and $\xi \in (0, 1)$; it also satisfies the identity [1,

Sec. 6.2.2]

$$\beta(a, b) = \beta(b, a) = \frac{\gamma(a)\gamma(b)}{\gamma(a+b)} , \quad (\text{A.20})$$

where $\gamma : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denotes the (real) *Gamma Function* with $\gamma(n+1) = n!$ for any $n \in \mathbb{N}_0^\infty$ [1, Sec. 6.1.5]. The corresponding *Incomplete Beta Function* [1, Sec. 6.6.1] is then given by

$$\iota(\nu; a, b) := \int_0^\nu \xi^{a-1}(1-\xi)^{b-1} d\xi = \int_{1-\nu}^1 \xi^{b-1}(1-\xi)^{a-1} d\xi , \quad (\text{A.21})$$

where the last equality follows by a simple substitution. An important identity is obtained from [1, Sec. 3.1.1, 6.6.2, 26.5.7],

$$\iota(\varepsilon; a, b) = \beta(a, b) \sum_{j=a}^{a+b-1} \binom{a+b-1}{j} \nu^j (1-\nu)^{a+b-1-j} , \quad (\text{A.22})$$

which can be written more compactly by use of the binomial distribution (A.17), see for instance [31, p. 3437]:

$$\iota(\nu; a, b) = \frac{1}{b} \binom{a+b-1}{b}^{-1} B(1-\nu; a+b-1, b-1) . \quad (\text{A.23})$$

Conclusion

The fundamental results on the scenario approach have been discovered recently and have sparked a great deal of research attention. From an engineering point of view, the scenario approach offers several key advantages. First, the range of optimization models covered by the formulation of the uncertain convex program $\text{UP}[\varepsilon]$ is quite large. Second, the uncertainty model requires no knowledge about the support set or the probability distribution, since the method is purely data-based. Third, the computational effort is restricted to solving a convex optimization program, even for otherwise difficult instances of $\text{UP}[\varepsilon]$.

In contrast to $\text{UP}[\varepsilon]$, many problems require a sequence of decisions for an uncertain system, while new information becomes available as time progresses. In particular, problems in the field of control require a sequence of *control actions* for steering a *dynamical system*, which is subject to uncertain disturbances and whose sensor data is affected by noise. Therefore the goal of this thesis is to extend the fundamental results of the scenario approach to *multi-stage decision problems*.

2.2 Multi-Stage Stochastic Programs

In the setting of $\text{UP}[\varepsilon]$, a decision is made once under uncertainty δ , and then the final outcome is observed. In the literature on stochastic programming, this is considered as a *single-stage stochastic program*. More common in practice, however, are *multi-stage decision problems*. They require a sequence of multiple, interrelated decisions based on different sets of available information; see Figure A.1 for an illustration.

Multi-stage decision problems have been studied extensively in the past. This section provides a brief review of the available theory from the field of stochastic programming. Excellent monographs on the subject, with far more details, are Birge and Louveaux [21], Kall and Mayer [57], Prékopa [78], and Shapiro et al. [99].




FIGURE A.1 Schematic overview of stochastic programming. For (a) two-stage stochastic programs and (b) multi-stage stochastic programs, the figure illustrates the sequential (not necessarily the temporal) order of the decisions and the observations of uncertainty variables.

Overview and Applications

In a *two-stage stochastic program*, as in Figure A.1(a), the *first-stage decision* $y_1 \in \mathbb{R}^{d_1}$ is made before and the *second-stage decision* $y_2 \in \mathbb{R}^{d_2}$ is made after the uncertainty δ_1 is observed. In other words, the first-stage decision must be made *here-and-now*, while for the second-stage decision one can *wait-and-see* to take corrective action depending on which value of the uncertainty δ_1 actually occurs. Therefore the second-stage decision can be considered as a function $y_2(\delta_1)$, which is called the *recourse action*.

In a *multi-stage stochastic program*, as in Figure A.1(b), the first-stage decision y_1 is made here-and-now and then the *first-stage uncertainty* δ_1 is observed. Before each of the other stages $i \in \{2, \dots, N\}$, the outcome of all past uncertainties $\delta_1, \dots, \delta_{i-1}$ is revealed for the purposes of the i -th stage decision $y_i(\delta_1, \dots, \delta_{i-1})$. Unlike for sampled dynamical systems, the stages in stochastic programming are sequential, but not necessarily uniform time periods [21, Sec. 2.4].

Two-stage stochastic programs have arisen from various practical applications. A prominent example is the *news vendor problem*, going back to Edgeworth [45], where a news vendor has to decide on how many newspapers to buy in the morning. He does not know the exact demand for these newspapers over the course of the day, but he has some options of reacting to weak or strong demands (e.g., by adjusting his prices). Similarly, problems of system design (e.g., of telecommunication networks, production facilities,

etc.) require first a decision about capacities, facing uncertain demands in the future. Thereafter, adjustments can be made in the system operation according to the actual demands.

Many practical problems also fit naturally into the framework of multi-stage stochastic programs. Analogous to the news vendor problem, an *inventory management problem* over multiple time periods may be considered [99, Sec. 1.2.3]. In each time period, decisions about sales and replenishments have to be made, facing uncertain demands in future periods. Furthermore, problems of *multi-period portfolio optimization* or *asset-liability management*, where a portfolio of assets with uncertain returns can be modified (i.e., by buying or selling decisions) over multiple time periods, can often be cast into multi-stage stochastic programs [103].

Two-Stage Stochastic Programs

Most of the literature focuses on a linear *two-stage stochastic program* (TSP) in the following standard form:

$$\min_{y_1, y_2(\delta_1)} \quad c_1^T y_1 + \mathbf{E}_1 [c_2(\delta_1)^T y_2(\delta_1)] \quad (\text{A.24a})$$

$$\text{s.t.} \quad G y_1 = h, \quad y_1 \geq 0, \quad (\text{A.24b})$$

$$T(\delta_1) y_1 + S(\delta_1) y_2(\delta_1) = r(\delta_1), \quad y_2(\delta_1) \geq 0. \quad (\text{A.24c})$$

Here $c_1 \in \mathbb{R}^{d_1}$ is the cost vector for the first-stage decision variables, and G and h are a matrix and a vector (of appropriate dimensions) defining linear constraints for the first-stage decision problem. Moreover, $c_2(\delta_1) \in \mathbb{R}^{d_2}$ is the cost vector for the second-stage decision variables, and $T(\delta_1)$, $S(\delta_1)$ and $r(\delta_1)$ are matrices and a vector (of appropriate dimensions), defining linear constraints for the second-stage decision problem. $T(\delta_1)$ and $S(\delta_1)$ are also called *technology matrix* and *recourse matrix*, respectively. $\mathbf{E}_1[\cdot]$ denotes the expectation operator with respect to δ_1 , i.e., on the probability space (Δ_1, \mathbf{P}_1) .

The data of the second-stage problem involves the uncertainty δ_1 , and its solution depends on the first-stage decision y_1 as well. Note that the values of both variables are known by the time the wait-and-see decision $y_2(\delta_1)$ must be fixed. Hence the second-stage problem is easy to solve. The difficulty of the TSP lies in finding the optimal here-and-now decision y_1 .

Define $\mathcal{F}_1 \subseteq \mathbb{R}^{d_1}$ and $\mathcal{F}_2(\delta_1) \subseteq \mathbb{R}^{d_2}$ as the feasible set of the first-stage and the second-stage, respectively:

$$\mathcal{F}_1 := \{y_1 \in \mathbb{R}^{d_1} \mid G y_1 = h, y_1 \geq 0\}, \quad (\text{A.25a})$$

$$\mathcal{F}_2(\delta_1) := \{y_2 \in \mathbb{R}^{d_2} \mid \exists y_1 \in \mathcal{F}_1 : T(\delta_1) y_1 + S(\delta_1) y_2 = r(\delta_1), y_2 \geq 0\}. \quad (\text{A.25b})$$

Moreover, define \mathcal{F}_2 as the set of first-stage decisions that leave a feasible recourse decision almost surely:

$$\mathcal{F}_2 := \bigcap_{\delta_1 \in \Delta_1} \mathcal{F}_2(\delta_1) , \quad (\text{A.26})$$

where any subset of Δ_1 with probability measure zero can be removed from the intersection in (A.26).

DEFINITION 2—RELATIVELY COMPLETE, COMPLETE, AND SIMPLE RECOURSE (a) The TSP is said to have *relatively complete recourse* if for any first-stage feasible decision y_1 there exists a feasible second stage decision y_2 almost surely, i.e., if $\mathcal{F}_1 \subseteq \mathcal{F}_2$. (b) The TSP is said to have *complete recourse* if $S(\delta_1)y_2$ can take on any vector value for an appropriate choice of $y_2 \geq 0$. (c) The TSP is said to have *simple recourse* if $S(\delta_1) = [I \ -I]$, where $I \in \mathbb{R}^{d_2 \times d_2}$ denotes the identity matrix. ■

PROPOSITION 1 (a) If δ_1 is a discrete random variable (i.e., the sample space Δ_1 is countable) then \mathcal{F}_2 is a closed convex subset of \mathbb{R}^{d_1} . (b) If, moreover, Δ_1 is finite, then \mathcal{F}_2 is a polyhedron in \mathbb{R}^{d_1} .⁸ ■

The proof of Proposition 1 is straightforward, by virtue of the facts that (a) the infinite intersection of convex sets is convex [28, Sec.2.3.1] and (b) the finite intersection of polyhedrons is a polyhedron [109, Ch.0]. Some properties of \mathcal{F}_2 in cases where the random variable δ_1 is not discrete are derived in Birge and Louveaux [21, Sec.3.1(c)].

The (*expected*) *value function* $J_2 : \mathbb{R}^{d_1} \rightarrow \mathbb{R} \cup \{\pm\infty\}$ is defined as the optimal expected cost of the second-stage problem

$$J_2(y_1) = \begin{cases} \min_{y_2} \mathbf{E}_1 \left[\{c_2(\delta_1)^T y_2 \mid T(\delta_1)y_1 + S(\delta_1)y_2 = r(\delta_1), y_2 \geq 0\} \right] & \text{if } y_1 \in \mathcal{F}_2 \\ +\infty & \text{if } y_1 \notin \mathcal{F}_2 \end{cases} . \quad (\text{A.27})$$

Note that any first-stage decision that does not have a feasible recourse action almost surely is assigned an infinite value function. The (*expected*) value function is also known as the *cost-to-go function* in the field of *dynamic programming* [17].⁹

With the definition of the value function (A.27), the TSP for finding the first-stage decision can be equivalently expressed as

$$\min_{y_1} \quad c_1^T y_1 + J_2(y_1) \quad (\text{A.28a})$$

$$\text{s.t.} \quad Gy_1 = h , \quad y_1 \geq 0 . \quad (\text{A.28b})$$

⁸A *polyhedron* is the intersection of a finite number of closed half-spaces of a vector space, and a *polytope* is a bounded polyhedron [49, 109].

⁹Dynamic programming, in its modern sense, is based on the work of Bellman [9]. For a comprehensive introduction see, e.g., the monographs of Bertsekas [18, 19].

For any fixed value of $\delta_1 \in \Delta_1$, the problem of finding the optimal second-stage decision y_2 is a *multi-parametric linear program* [50], whose parameter is y_1 .¹⁰ Hence in this case $J_2(\cdot)$ is known to be a convex function which is continuous and *piecewise affine on polyhedral sets* (p.w.a.). Because the positively weighted sum of convex and p.w.a. functions is again convex and p.w.a., the following result is immediate [21, Sec. 3.1(b)].

PROPOSITION 2 If δ_1 is a finite discrete random variable, then the value function $J_2(\cdot)$ is convex and p.w.a., and problem (A.28a) can be represented as a linear program. ■

Proposition 2 means that, if δ_1 is a finite discrete random variable, then the TSP can be solved as a standard linear program. If δ_1 is a more general random variable, then the TSP is generally not a linear program. However, it may exhibit some convenient properties, as in the following result, that make it amenable to the solution by numerical algorithms [57, Sec. 3.2].

THEOREM 3 (a) If the recourse matrix is fixed (i.e., $S(\delta_1) \equiv S$) and δ_1 has finite second moments, then $J_2(\cdot)$ is a convex and Lipschitz-continuous function on \mathcal{F}_2 . (b) If, moreover, the distribution of δ_1 is absolutely continuous, then $J_2(\cdot)$ is differentiable on the relative interior of \mathcal{F}_2 . ■

A large number of numerical algorithms are proposed in the literature for solving specific instances of the TSP. In many practical instances, the TSP becomes a large-scale linear program with a particular structure. Therefore, a variety of decomposition methods have been developed for an efficient solution, most prominently the *L-shaped method* of van Slyke and Wets [105], which is based on *Bender's decomposition* [14]. See Birge and Louveaux [21, Part III,IV] or Kall and Mayer [57, Ch. 4] for a comprehensive overview.

Multi-Stage Stochastic Programs

The linear TSP extends to the linear *multi-stage stochastic program* (MSP) in a straightforward way:

$$\begin{aligned}
 \min_{y_1, \dots, y_N(\omega_{N-1})} \quad & c_1^T y_1 + \mathbf{E}_1 [c_2(\omega_1)^T y_2(\omega_1)] + \dots + \mathbf{E}_{N-1} [c_N(\omega_{N-1})^T y_N(\omega_{N-1})] \\
 \text{s.t.} \quad & S_1 y_1 = p_1, \quad y_1 \geq 0, \\
 & T_2(\omega_1) z_1 + S_2(\omega_1) y_2(\omega_1) = p_2(\omega_1), \quad y_2(\omega_1) \geq 0, \\
 & T_3(\omega_2) z_2(\omega_1) + S_3(\omega_2) y_3(\omega_2) = p_3(\omega_2), \quad y_3(\omega_2) \geq 0, \\
 & \dots \\
 & T_N(\omega_{N-1}) z_{N-1}(\omega_{N-2}) + S_N(\omega_{N-1}) y_N(\omega_{N-1}) = r_N(\omega_{N-1}), \\
 & y_N(\omega_{N-1}) \geq 0.
 \end{aligned} \tag{A.29}$$

¹⁰The research on multi-parametric linear programs goes back to Gal and Nedoma [50]; see Bank et al. [7] for an excellent introduction to general parametric programming. The concept of multi-parametric linear and quadratic programming also appears as the explicit policy of linear model predictive control (cf. Section 2.3) [11, 13, 25, 56]. Many state-of-the-art numerical algorithms have been implemented in the Multi-Parametric Toolbox for Matlab, developed by Herceg et al. [55].

Here $\omega_i := \{\delta_1, \delta_2, \dots, \delta_i\}$ is the collection of all uncertainties observed up to stage $i + 1 \in \{2, 3, \dots, N\}$, i.e., before deciding on the variable $y_{i+1}(\omega_i)$. The probability $\mathbf{P}_i[\cdot]$ and the expectation $\mathbf{E}_i[\cdot]$ refer to the random variable ω_i . Moreover, $z_i(\omega_{i-1}) := [y_1^T \ y_2^T(\omega_1) \ \dots \ y_i^T(\omega_{i-1})]^T$ is the column vector in $\mathbb{R}^{(d_1 + \dots + d_i)}$ of all stacked-up decision vectors up to stage $i \in \{1, 2, \dots, N\}$.

The MSP can again be considered as a dynamic program, and its solution can be described by a backwards recursion; see [99, Ch. 3]. Analogous to (A.25) and (A.26), the feasible sets \mathcal{F}_i of stages $i = N, N - 1, \dots, 2$ are defined as

$$\mathcal{F}_i := \bigcap_{\delta_{i-1} \in \Delta_{i-1}} \{z_{i-1}(\omega_{i-2}) \mid \exists y_i \in \mathbb{R}^{d_i} : \\ T_i(\omega_{i-1})z_{i-1}(\omega_{i-2}) + S_i(\omega_{i-1})y_i = r(\omega_{i-1}), y_i \geq 0\} \quad (\text{A.30})$$

And analogous to (A.27), the cost-to-go $J_i : \mathbb{R}^{d_1 + \dots + d_{i-1}} \times (\Delta_1 \times \dots \times \Delta_{i-2}) \rightarrow \mathbb{R}$ becomes for the final stage $i = N$

$$J_N(z_{N-1}, \omega_{N-2}) := \min_{y_N} \mathbf{E}_{N-1} [c_N(\omega_{N-1})^T y_N : \\ T_N(\omega_{N-1})z_{N-1} + S(\omega_{N-1})y_N = r_N(\omega_{N-1}), y_N \geq 0 \mid \omega_{N-2}] \quad (\text{A.31a})$$

if $z_{N-1}(\omega_{N-2}) \in \mathcal{F}_N$, otherwise $J_N(z_{N-1}, \omega_{N-2}) = +\infty$; and for the other stages $i = N - 1, N - 2, \dots, 2$

$$J_i(z_{i-1}, \omega_{i-2}) := \min_{y_i} \mathbf{E}_{i-1} [c_i(\omega_{i-1})^T y_i + J_{i+1}(z_i, \omega_{i-1}) : \\ T_i(\omega_{i-1})z_{i-1} + S(\omega_{i-1})y_i = r_i(\omega_{i-1}), y_i \geq 0 \mid \omega_{i-2}] \quad (\text{A.31b})$$

if $z_{i-1}(\omega_{i-2}) \in \mathcal{F}_i$, otherwise $J_i(z_{i-1}, \omega_{i-2}) = +\infty$. Note that the expectations in (A.31) are conditional on the past uncertainties, accounting for the fact that the random uncertainties may be statistically dependent between different stages.

Dynamic programming, although theoretically appealing, is usually an impractical method for solving multi-stage stochastic programs. In general, for the applicability of dynamic programming, the cost-to-go functions $J_i(z_{i-1}, \omega_{i-2})$ must have a parameterization of a tractable size. The parameterization, however, tends to grow exponentially with the dimensions of the domain of J_i (*curse of dimensionality*). A particular case where dynamic programming is applicable is when the set of possible decisions is finite. Moreover, in some cases, the past decisions z_{i-1} and uncertainties ω_{i-2} have a joint representation as a low-dimensional *system state* [99, p. 68]. This essentially leads into the field of dynamical systems and control theory.

The fundamental approach of stochastic programming for solving the MSP is via a *scenario tree*, as shown in Figure A.2 [99, Sec. 3.1.3]. If the distributions of $\delta_1, \dots, \delta_{N-1}$ are non-finite (or finite with too many possible outcomes), they are usually approximated

by selecting a discrete support set of small cardinality [57, Sec. 4.8.2].




FIGURE A.2 Example of a scenario tree of an N -stage stochastic program, with *tree levels* ranging from 1 (*root node*) to N (*leaf nodes*). Each path from the root node to a leaf node is called a *scenario*. Note that different nodes (e.g., of δ_2) may correspond to the same value (e.g., if the preceding value of δ_1 is different).

Each leaf node $k \in \{1, 2, \dots, K\}$ is associated with one particular uncertainty *scenario* $\{\delta_1^{(k)}, \delta_2^{(k)}, \dots, \delta_{N-1}^{(k)}\}$, whose probability $p_k \in [0, 1]$ is straightforward to compute. Moreover, one particular set of decision variables $\{y_1^{(k)}, y_2^{(k)}, \dots, y_N^{(k)}\}$ is introduced for each scenario k . Since a future decision $y_i^{(k)}$ may depend only on the preceding uncertainties $\{\delta_1^{(k)}, \delta_2^{(k)}, \dots, \delta_{i-1}^{(k)}\}$, but not the successive uncertainties $\{\delta_i^{(k)}, \delta_{i+1}^{(k)}, \dots, \delta_{N-1}^{(k)}\}$, the following *non-anticipativity constraints* must be respected:

$$y_i^{(k_1)} = y_i^{(k_2)} \quad \forall i \leq b(k_1, k_2), \quad \forall k_1, k_2 \in \{1, 2, \dots, K\}, \quad k_1 \neq k_2. \quad (\text{A.32})$$

Here $b(k_1, k_2) \in \{1, \dots, N-1\}$ denotes the level of the *branch node* of scenarios k_1 and k_2 (i.e., the tree level up to which the paths of scenarios k_1 and k_2 are identical).

The MSP (A.29) can hence be reformulated as the *equivalent deterministic program*

$$\begin{aligned}
 \min_{y_1^{(k)}, \dots, y_N^{(k)}} \quad & \sum_{k=1}^K p_k [c_1^T y_1^{(k)} + c_2(\omega_1^{(k)})^T y_2^{(k)} + \dots + c_N(\omega_{N-1}^{(k)})^T y_N^{(k)}] \\
 \text{s.t.} \quad & S_1 y_1^{(k)} = p_1, \quad y_1^{(k)} \geq 0, \\
 & T_2(\omega_1^{(k)}) z_1^{(k)} + S_2(\omega_1^{(k)}) y_2^{(k)} = r_2(\omega_1^{(k)}), \quad y_2^{(k)} \geq 0, \\
 & T_3(\omega_2^{(k)}) z_2^{(k)} + S_3(\omega_2^{(k)}) y_3^{(k)} = r_3(\omega_2^{(k)}), \quad y_3^{(k)} \geq 0, \\
 & \dots \\
 & T_N(\omega_{N-1}^{(k)}) z_{N-1}^{(k)} + S_N(\omega_{N-1}^{(k)}) y_N^{(k)} = r_N(\omega_{N-1}^{(k)}), \quad y_N^{(k)} \geq 0, \\
 & y_i^{(k_1)} = y_i^{(k_2)} \quad \forall i \leq b(k_1, k_2),
 \end{aligned} \tag{A.33}$$

where the constraints hold for all $k, k_1, k_2 \in \{1, \dots, K\}$ with $k_1 \neq k_2$. Note that $\omega_1^{(k)} := \{\delta_1^{(k)}, \delta_2^{(k)}, \dots, \delta_i^{(k)}\}$ and $z_i^{(k)} := [y_1^{(k)T} y_2^{(k)T} \dots y_i^{(k)T}]^T$.

Problem (A.33) is a large scale linear program of a particular characteristic structure. Efficient numerical algorithms, similar to those for two-stage stochastic programs, can therefore be developed for computing exact solutions. The majority of such algorithms is based on the *nested decomposition method* [57, Sec. 4.8]. Since the dimension of (A.33) is often enormous (in fact, it grows exponentially with the number of stages N), the problem is commonly tackled with approximation methods; see [21, Ch. 6] for an overview.

Conclusion

Stochastic programming is concerned with solution methods for multi-stage stochastic decision problems. A large variety of algorithms have been proposed, based on solving a finite sequence of single-stage, deterministic optimization programs. Many of these algorithms exploit the particular multi-stage structure of the problem, by employing various decomposition techniques [21, Part III], [57, Ch. 4]. Hence these algorithms are applicable to problems with a relatively low number of decision variables and constraints (as compared to deterministic optimization), and a moderate size of the scenario tree. Dynamic programming is applicable in special circumstances, or in very small dimensions.

Problems with a high number of stages tend to exceed the manageable computational complexity, as a result of the combinatorial growth of the scenario tree. Probability distributions, if known exactly, must usually be approximated with a small number of discrete outcomes per stage. In order to reduce the computations, one may resort to approximation methods, such as *sample-average approximation* [21, Part IV], [99, Ch. 5].

From the viewpoint of control theory, where multi-stage stochastic decision problems have to be solved on-line, their computational complexity easily becomes prohibitive. On the other hand, the available approximation methods often cannot effectively handle constraints on the system state and/or provide violation bounds. Moreover, no analysis of the closed control loop (i.e., with a receding horizon implementation) is available.

2.3 Model Predictive Control

Model Predictive Control (MPC) is an advanced control method, whose origins can be traced back into the 1960s [69]. Compared to approaches of classical control theory [48, 63, 64, 73, 102], MPC can handle multivariable control problems and constraints on the inputs, states, and outputs in a natural way.

Significant advances since the 1980s have created a solid theoretical basis for MPC. Furthermore, from its origins in the process industry [51, 84], MPC has now been successfully tested and applied in numerous industrial applications [80]. The potentials of MPC, however, are not yet fully uncovered. Many theoretical questions remain open to further investigations—in particular, about the handling of uncertainty in the underlying model, forecast errors, and disturbances. Moreover, for a large range of control applications the performance may be improved by an MPC implementation.

This section provides a brief introduction to the fundamental concepts and basic terminology of MPC. Further background can be found in the excellent monographs of Maciejowski [67], Rossiter [87], Rawlings and Mayne [83], and Borrelli et al. [26], as well as several survey papers on this subject [33, 47, 51, 69, 82].

Basic Assumptions

The basic principles of MPC are introduced for a linear time-invariant (LTI) system in discrete time $t \in \mathbb{N}$,

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad x_0 \in \mathbb{X}. \quad (\text{A.34})$$

Here $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ are the *system matrix* and *input matrix*, respectively. The trajectory of *states* $x_t \in \mathbb{R}^n$ must be kept inside the constrained set $\mathbb{X} \subseteq \mathbb{R}^n$ for all times $t \in \mathbb{N}$, where $x_0 \in \mathbb{R}^n$ is a given *initial condition*. The external *inputs* $u_t \in \mathbb{R}^m$ can be selected from a set $\mathbb{U} \subseteq \mathbb{R}^m$ to control the system. The external disturbances are modeled by an additive term $w_t \in \mathbb{W} \subset \mathbb{R}^n$.

It should be mentioned that a comprehensive theory exists also for systems (A.34) with nonlinear dynamics [33, 47, 54, 83]. However, the theory of *Nonlinear MPC* (NMPC) is beyond the scope of this thesis.

ASSUMPTION 4—SYSTEM DYNAMICS (a) The pair of matrices A, B is stabilizable. (b) The state x_t is measured at every step t . (c) The *state constraint set* \mathbb{X} is convex and contains the origin in its interior. (d) The *set of admissible controls* \mathbb{U} is compact, convex, and contains the origin in its interior. ■

For computational reasons, the constraint sets \mathbb{X} and \mathbb{U} are typically chosen as polytopes or polyhedrons (cf. footnote 8), even though this assumption is not required for the following theory.

The objective is to find a *feedback policy* $\kappa : \mathbb{X} \rightarrow \mathbb{U}$ that maps any measured state $x_t \in \mathbb{X}$ to an admissible control $u_t = \kappa(x_t) \in \mathbb{U}$, in a way that meets two key requirements: (a) the system state remains within the state constraint set \mathbb{X} at all times; and (b) the

trajectory shows good performance according to specified criterion and, in particular, it is stable (i.e., it converges to the origin).

Nominal MPC Feedback Policy

Throughout this section, the following additional assumption are made, leading to the basic theory of *nominal MPC*.

ASSUMPTION 5—NOMINAL MPC (a) An accurate *model* of system (A.34) is available for controller design; i.e., the system and input matrices A and B are deterministic and known exactly. (b) There are no disturbances; i.e., $w_t = 0$ for all times $t \in \mathbb{N}$. ■

Based on Assumptions 4 and 5, at any time $t \in \mathbb{N}$ and starting from the current state $x_{t|t} = x_t$, the state trajectory $\{x_{t|t}, x_{t+1|t}, \dots, x_{t+N|t}\}$ can be exactly predicted based on any selection $\{u_{t|t}, u_{t+1|t}, \dots, u_{t+N-1|t}\}$ of inputs. Here N is a finite number of time steps, called the *prediction horizon*. Moreover, a particular sequence of inputs $\{u_{t|t}^*, u_{t+1|t}^*, \dots, u_{t+N-1|t}^*\}$ can be computed as the solution to the following *Finite-Horizon Optimal Control Problem* (FHOCp):

$$\min_{u_{t|t}, \dots, u_{t+N-1|t}} \sum_{i=0}^{N-1} \ell(u_{t+i|t}, x_{t+i|t}) + \ell_f(x_{t+N|t}) \quad (\text{A.35a})$$

$$\text{s.t.} \quad x_{t+i+1|t} = Ax_{t+i|t} + Bu_{t+i|t}, \quad x_{t|t} = x_t \quad \forall i = 0, \dots, N-1, \quad (\text{A.35b})$$

$$u_{t+i|t} \in \mathbb{U} \quad \forall i = 0, \dots, N-1, \quad (\text{A.35c})$$

$$x_{t+i|t} \in \mathbb{X} \quad \forall i = 1, \dots, N-1, \quad (\text{A.35d})$$

$$x_{t+N|t} \in \mathcal{X}_f. \quad (\text{A.35e})$$

The *cost function* (A.35a) to be minimized in the FHOCp is assumed to be decomposable as the sum of *stage costs* $\ell : \mathbb{U} \times \mathbb{X} \rightarrow \mathbb{R}_{0+}$ and a *terminal cost* $\ell_f : \mathbb{X} \rightarrow \mathbb{R}_{0+}$. Constraint (A.35e) requires the final predicted state to be driven into a *terminal set* $\mathcal{X}_f \subseteq \mathbb{X} \subseteq \mathbb{R}^n$. It should be emphasized that the doubly-indexed states $x_{t+1|t}, x_{t+2|t}, \dots$ represent *predictions*—as opposed to the (*actual*) states x_{t+1}, x_{t+2}, \dots of system (A.34).

ASSUMPTION 6—COST FUNCTION (a) The stage cost $\ell : \mathbb{U} \times \mathbb{X} \rightarrow \mathbb{R}_{0+}$ is convex, continuous, $\ell(0, 0) = 0$, and there exists a lower-bounding K_∞ -function¹¹ $\alpha_l : \mathbb{R}_{0+} \rightarrow \mathbb{R}_{0+}$, i.e., $\ell(\xi, v) \geq \alpha_l(\|\xi\|)$ for all $\xi \in \mathbb{X}$ and $v \in \mathbb{U}$. (b) The terminal cost $\ell_f : \mathcal{X}_f \rightarrow \mathbb{R}_{0+}$ is convex, continuous, $\ell_f(0) = 0$, and there exists an upper-bounding K_∞ -function¹¹, i.e., $\ell_f(\xi) \leq \alpha_u(\|\xi\|)$ for all $\xi \in \mathcal{X}_f$; cf. [83, Ass. 2.16]. ■

¹¹A function $\alpha : \mathbb{R}_{0+} \rightarrow \mathbb{R}_{0+}$ is a *K-function* if it is continuous, strictly monotonically increasing, and $\alpha(0) = 0$; it is a *K_∞-function* if, in addition, $\alpha(r) \rightarrow \infty$ as $r \rightarrow \infty$; see [60, Def. 4.2]. These definitions appear in the context of Lyapunov stability theory; see [60, 88, 106] for further details.

Typical choices for the stage cost $\ell(\cdot, \cdot)$ include a so-called *linear stage costs* (A.36a,b) or a *quadratic stage cost* (A.36c) [26],

$$\ell(\xi, v) := \|Q_\ell \xi\|_1 + \|R_\ell v\|_1, \quad (\text{A.36a})$$

$$\ell(\xi, v) := \|Q_\ell \xi\|_\infty + \|R_\ell v\|_\infty, \quad (\text{A.36b})$$

$$\text{or } \ell(\xi, v) := \|Q_\ell \xi\|_2^2 + \|R_\ell v\|_2^2. \quad (\text{A.36c})$$

Here $Q_\ell \in \mathbb{R}^{n \times n}$ represents a positive definite, and $R_\ell \in \mathbb{R}^{m \times m}$ a positive semi-definite weighting matrix. Note that the terminal cost in (A.35a) may be trivial, i.e., $\ell_f \equiv 0$.

The *MPC feedback policy* $\kappa_{\text{MPC}} : \mathbb{X} \rightarrow \mathbb{U}$ is defined as follows. The current state x_t is substituted into the FHOCF via constraint (A.35b); then the FHOCF is solved for the optimal input sequence $\{u_{t|t}^*, u_{t+1|t}^*, \dots, u_{t+N-1|t}^*\}$ over the prediction horizon; finally, only the first entry $u_t = \kappa_{\text{MPC}}(x_t) := u_{t|t}^*$ is applied as a control input to the system at time t .

Recursive Feasibility and Stability

For nominal MPC, various conditions have been proposed that guarantee *recursive feasibility* and *stability*; see Mayne et al. [69]. The common approach is also the most relevant one for this thesis; it is briefly described below.

Let $\mathcal{X}_N \subseteq \mathbb{X}$ be the N -step (backward) reachable set of the terminal set \mathcal{X}_f [59],

$$\mathcal{X}_N := \{x_0 \in \mathbb{X} \mid \exists u_{0|0}, \dots, u_{N-1|0} \in \mathbb{U} : x_{1|0}, \dots, x_{N|0} \in \mathbb{X}, x_{N|0} \in \mathcal{X}_f\}. \quad (\text{A.37})$$

Note that \mathcal{X}_N is exactly the set of all initial states x_t for which the FHOCF (A.35) is feasible. The nominal MPC policy $u_t = \kappa_{\text{MPC}}(x_t)$ aims (a) at keeping the state x_t inside \mathcal{X}_N at all times $t \in \mathbb{N}$ and (b) at stabilizing the state trajectory of system (A.34) about the origin.

Both conditions can be achieved by a properly selected *terminal condition*; i.e., an appropriate combination of the terminal set \mathcal{X}_f with a terminal cost $\ell_f(\cdot)$. The set \mathcal{X}_N can be regarded as the *region of attraction* of the model predictive controller.

DEFINITION 3—CONTROL INVARIANT SETS (a) A set $\mathcal{X} \subset \mathbb{X}$ is called a *control invariant* (CI) set of system (A.34) if for any $\xi \in \mathcal{X}$ there exists a $v \in \mathbb{U}$ such that $A\xi + Bv \in \mathcal{X}$ [23, Def. 2.3]. (b) A set $\mathcal{X} \subset \mathbb{X}$ is called a *robust control invariant* (RCI) set of system (A.34) if for any $\xi \in \mathcal{X}$ there exists a $v \in \mathbb{U}$ such that $A\xi + Bv + \omega \in \mathcal{X}$ for all $\omega \in \mathbb{W}$ [23, Def. 2.3]. ■

Details on the computation of CI and RCI sets can be found in Kerrigan [59, Ch. 2].

DEFINITION 4—CONTROL LYAPUNOV FUNCTION Let $l : \mathcal{X} \rightarrow \mathbb{R}_{0+}$ be a continuous function defined on a CI set $\mathcal{X} \subset \mathbb{X}$ of system (A.34), $l(\cdot)$ be upper bounded by a

K_∞ -function, and $l(0) = 0$. If $l(\cdot)$ satisfies the condition

$$\min_{v \in \mathbb{U}, A\xi + Bv \in \mathcal{X}} [l(A\xi + Bv) - l(\xi) + \ell(\xi, v)] \leq 0 \quad \forall \xi \in \mathcal{X}_0, \quad (\text{A.38})$$

then it is called a *control Lyapunov function* for system (A.34) [83, Ass. 2.12]. ■

THEOREM 4—NOMINAL MPC Let Assumptions 4, 5, 6 hold and $\ell_f(\cdot)$ be a control Lyapunov function on a CI set \mathcal{X}_f for system (A.34). If $x_0 \in \mathcal{X}_N$, then the MPC policy $u_t = \kappa_{\text{MPC}}(x_t)$ (a) keeps the state trajectory $\{x_t\}_{t \in \mathbb{N}}$ inside \mathcal{X}_N and (b) drives the state trajectory $\{x_t\}_{t \in \mathbb{N}}$ asymptotically to the origin. ■

A proof of Theorem 4 is standard and therefore only a short outline is given below; see Rawlings and Mayne [83, Sec. 2.4] for more details. Since $\kappa_{\text{MPC}}(\cdot)$ is a nonlinear controller (in fact, it is *piecewise affine on polyhedral sets* (p.w.a.) [11, 13]), the proof is based on Lyapunov stability theory.¹²

The *value function* $J_N^* : \mathcal{X}_N \rightarrow \mathbb{R}_{0+}$, returning the minimal objective function value $J_N^*(x_t)$ of the FHOC for any $x_t \in \mathcal{X}_N$, can be employed as a Lyapunov function for the closed-loop system. Let $\{u_{t|t}^*, u_{t+1|t}^*, \dots, u_{t+N-1|t}^*\}$ be the optimal input sequence and $\{x_{t|t}^*, x_{t+1|t}^*, \dots, x_{t+N|t}^*\}$ be the optimal state sequence of the FHOC for the current state $x_t \in \mathcal{X}_N$. Under the assumptions of Theorem 4, the shifted input sequence

$$\{u_{t+1|t}^*, u_{t+2|t}^*, \dots, u_{t+N-1|t}^*, v^*\} \quad (\text{A.39})$$

is feasible in the FHOC for the subsequent state $x_{t+1} := Ax_t + Bu_{t|t}^*$. Here v^* is the optimal input according to (A.38) with $\xi := x_{t+N|t}^*$. Moreover, the shifted input sequence (A.39) induces a lower cost than $J_N^*(x_t)$ and therefore $J_N^*(x_{t+1}) < J_N^*(x_t)$.

REMARK 1—TERMINAL CONDITION A variety of approaches are known for computing a stabilizing terminal condition \mathcal{X}_f , $\ell_f(\cdot)$; see Mayne et al. [69, Sec. 3.7] for an overview. The earliest and simplest variant is $\mathcal{X}_f := \{0\}$ and $\ell_f(\cdot) \equiv 0$ [58]. For the case of quadratic costs, a better approach is to compute \mathcal{X}_f as the maximum *positively invariant* set (see Definition 5 below) of the corresponding *linear quadratic regulator* [5]. ■

DEFINITION 5—POSITIVELY INVARIANT SETS Let $F \in \mathbb{R}^{m \times n}$ be a linear feedback gain. (a) A set $\mathcal{X} \subset \mathbb{X}$ is called a *positively invariant* (PI) set of system (A.34) under the control law $u_t = Fx_t \in \mathbb{U}$ if for any $\xi \in \mathcal{X}$ it holds that $(A + BF)\xi \in \mathcal{X}$ [23, Def. 2.1]. (b) A set $\mathcal{X} \subset \mathbb{X}$ is called a *robust positively invariant* (RPI) set of system (A.34) under the control law $u_t = Fx_t \in \mathbb{U}$ if for any $\xi \in \mathcal{X}$ it holds that $(A + BF)\xi + \omega \in \mathcal{X}$ for all $\omega \in \mathbb{W}$ [23, Def. 2.2]. ■

Details on the computation of PI and RPI sets can be found in Kerrigan [59, Ch. 2].

¹²See Khalil [60], Sastry [88], or Vidyasagar [106] for the basic concepts of stability, as well as methods of stability verification for nonlinear systems.

Conclusion

MPC is a powerful concept for handling multivariable control problems with constraints on the inputs, states, and outputs. For nominal MPC, a theoretical framework has been developed that guarantees recursive feasibility and stability of the MPC policy (i.e., when the solution to the FHOCP is applied in a receding horizon fashion).

For some control systems, a highly accurate linear model is available and the disturbances are small. Then nominal MPC has proven to be a very effective control strategy in many cases. Dynamic feedback is introduced from re-solving the FHOCP in each time step.

For other control systems, the assumptions of no model uncertainty and no external disturbances are less reasonable. In these cases, the theoretical guarantees are lost and infeasible instances of the FHOCP may occur. This may have a detrimental effect on the performance of MPC. The latter type of systems are the main motivation for this thesis.

2.4 Uncertainty in Model Predictive Control

In the presence of uncertainty, nominal MPC is also referred to as *certainty-equivalent MPC*. This means that uncertain quantities in the predictions are simply replaced by a nominal value (e.g., the most likely or the average value). For certainty-equivalent MPC, all theoretical guarantees (on recursive feasibility, stability, etc.) are lost and the controller performance (constraint satisfaction, closed-loop cost, etc.) may actually be poor.

In these cases, the MPC approach often benefits significantly from the inclusion of an explicit *uncertainty model*. This model can be based on set membership (as in *robust MPC*, RMPC), on a probability distribution (as in *stochastic MPC*, SMPC), or on uncertainty scenarios (as in *scenario-based MPC*, SCMPC). Each of these approaches is briefly reviewed in this section, as they are especially relevant for Parts C and D of this thesis.

Disturbance Feedback

Firstly in this section, let Assumption 5 be replaced with the following Assumption 7.

ASSUMPTION 7—ROBUST MPC (a) The system and input matrices A and B are exactly known for the controller design. (b) The *disturbance set* \mathbb{W} is compact and contains the origin. ■

For the purposes of this introduction, Assumption 7 allows only for an additive uncertainty. Nonetheless, there exist several RMPC approaches (e.g., [12], [37]) that consider also uncertainty in the controller model. This also holds for the SCMPC algorithm presented in Part C of this thesis.

If the predictions by the system model are uncertain, it is important to distinguish between *open-loop* and *closed-loop predictions* [83, Sec.3.1]. Open-loop predictions refer to a sequence of fixed control actions $\{u_{t|t}, u_{t+1|t}, \dots, u_{t+N-1|t}\}$. Closed-loop

predictions include the possibility for *recourse actions* based on future observations of the uncertainties; cf. Section 2.2. For closed-loop predictions, the predicted control actions $\{u_{t|t}, \mu_{t+1|t}(w_{t|t}), \dots, \mu_{t+N-1|t}(w_0, \dots, w_{t+N-2|t})\}$ are in fact feedback policies $\mu_{t+i|t} : \mathbb{W}^i \rightarrow \mathbb{U}$, as opposed to constants, for any $i = 1, 2, \dots, N-1$.

The advantage of closed-loop over open-loop predictions is a more accurate representation of the future system trajectory. This generally leads to an improvement of the MPC performance, in particular because the open-loop predictions have to be more conservative [69, Sec. 4.6.1]. The drawback of closed-loop predictions is that they involve general feedback policies—i.e., multi-dimensional nonlinear maps. Hence they cannot be solved for, in general. These policies either have to be determined offline (before solving the FHOCP), or they have to be approximated in the FHOCP by a suitable parameterization (one that keeps the computations tractable).

Researchers in control theory prefer to think of the predicted control actions in terms of *state feedback policies* $\kappa_{t+i|t} : \mathbb{X} \rightarrow \mathbb{U}$, rather than *disturbance feedback policies* $\mu_{t+i|t} : \mathbb{W}^i \rightarrow \mathbb{U}$ as researchers in stochastic programming; cf. Section 2.2. The following result of Skaf and Boyd [101, Sec. IV] shows that both are indeed equivalent.

PROPOSITION 3—STATE AND DISTURBANCE FEEDBACK For system (A.34) under Assumption 7(a), there exists a bijection between the set of all disturbance feedback policies $\mu_{t+i|t} : \mathbb{W}^i \rightarrow \mathbb{U}$ and the set of all extended state feedback policies $\tilde{\kappa}_{t+i|t} : \mathbb{X}^{i+1} \rightarrow \mathbb{U}$, for any $i = 1, 2, \dots, N-1$. ■

Proof. Denote the nominal state predictions $\{x_{t|t}, x_{t+1|t}, \dots, x_{t+N|t}\}$ and define the disturbed state predictions $\{\tilde{x}_{t|t}, \tilde{x}_{t+1|t}, \dots, \tilde{x}_{t+N|t}\}$ as

$$\tilde{x}_{t+i+1|t} := A\tilde{x}_{t+i|t} + Bu_{t+i|t} + w_{t+i|t}, \quad \tilde{x}_{t|t} := x_t \quad \forall i = 0, 1, \dots, N-1. \quad (\text{A.40})$$

First, suppose that all future control inputs are given by an extended state feedback policy,

$$u_{t+i|t} := \tilde{\kappa}_{t+i|t}(\tilde{x}_{t+1|t}, \tilde{x}_{t+2|t}, \dots, \tilde{x}_{t+i|t}) \quad \forall i = 1, 2, \dots, N-1. \quad (\text{A.41})$$

Then the prediction errors can be defined as a function of the past disturbances,

$$e_{t+i|t}(w_{t|t}, w_{t+1|t}, \dots, w_{t+i-1|t}) := \tilde{x}_{t+i|t} - x_{t+i|t} \quad \forall i = 1, 2, \dots, N-1. \quad (\text{A.42})$$

Substituting (A.42) into (A.41) yields an explicit expression for the equivalent disturbance feedback policies:

$$\begin{aligned} \mu_{t+i|t}(w_{t|t}, w_{t+1|t}, \dots, w_{t+i-1|t}) &:= \\ \tilde{\kappa}_{t+i|t}(x_{t+1|t} + e_{t+1|t}(w_t), x_{t+2|t} + e_{t+2|t}(w_{t|t}, w_{t+1|t}), \dots) &\quad \forall i = 1, 2, \dots, N-1. \end{aligned}$$

Second, suppose that all future control inputs are given by a disturbance feedback policy,

$$u_{t+i|t} := \mu_{t+i|t}(w_{t|t}, w_{t+1|t}, \dots, w_{t+i-1|t}) \quad \forall i = 1, 2, \dots, N-1. \quad (\text{A.43})$$

Then explicit expressions for the equivalent extended state feedback policies can be obtained recursively. For the first step,

$$w_{t|t} = \tilde{x}_{t+1|t} - x_{t+1|t} \implies \tilde{\kappa}_{t+1|t}(\tilde{x}_{t+1|t}) := \mu_{t+1|t}(\tilde{x}_{t+1|t} - x_{t+1|t}). \quad (\text{A.44})$$

For all subsequent steps $i = 2, \dots, N-1$, suppose that for $\mu_{t+i-1|t}(w_{t|t}, w_{t+1|t}, \dots, w_{t+i-2|t})$ (abbreviated as $\mu_{t+i-1|t}(\cdot)$) an equivalent expression $\tilde{\kappa}_{t+i-1|t}(\tilde{x}_{t+1|t}, \tilde{x}_{t+2|t}, \dots, \tilde{x}_{t+i-1|t})$, (abbreviated as $\tilde{\kappa}_{t+i-1|t}(\cdot)$) has already been determined. Then

$$w_{i-1} = \tilde{x}_{t+i|t} - (A\tilde{x}_{t+i-1|t} + B \underbrace{\mu_{t+i-1|t}(\cdot)}_{=\tilde{\kappa}_{t+i-1|t}(\cdot)}) \implies \tilde{\kappa}_{t+i|t}(\cdot) := \mu_{t+i|t}(\cdot) \quad (\text{A.45})$$

yields all equivalent extended state feedback policies. \square

By virtue of Proposition 3, disturbance feedback can be assumed without any loss of generality. Note that any state feedback policy $\kappa_{t+i|t} : \mathbb{X} \rightarrow \mathbb{U}$ is, in particular, also an *extended state feedback policy* $\tilde{\kappa}_{t+i|t} : \mathbb{X}^{i+1} \rightarrow \mathbb{U}$; hence it has an equivalent disturbance feedback policy.

The disturbance feedback policies are, essentially, multi-dimensional nonlinear maps. Therefore they cannot be numerically determined, or even stored, exactly. The first approximation that comes to mind is by a linear or affine parameterization. For this special case, Goulart et al. [52] have proved a similar result to Proposition 3.

PROPOSITION 4—AFFINE DISTURBANCE FEEDBACK There exists a bijection between the class of *affine extended state feedback* (ASF) policies

$$\begin{aligned} \tilde{\kappa}_{t+i|t}(\tilde{x}_{t+1|t}, \tilde{x}_{t+2|t}, \dots, \tilde{x}_{t+i|t}) := \\ F_{i,1}\tilde{x}_{t+1|t} + F_{i,2}\tilde{x}_{t+2|t} + \dots + F_{i,i}\tilde{x}_{t+i|t} + \tilde{c}_{t+i|t} \quad \forall i = 1, 2, \dots, N-1, \end{aligned}$$

and the class of *affine disturbance feedback* (ADF) policies

$$\begin{aligned} \mu_{t+i|t}(w_{t|t}, w_{t+1|t}, \dots, w_{t+i-1|t}) := \\ F_{i,1}w_t + F_{i,2}w_{t+1} + \dots + F_{i,i}w_{t+i-1} + c_{t+i|t} \quad \forall i = 1, 2, \dots, N-1. \end{aligned}$$

Moreover, the feedback matrices $F_{i,j} \in \mathbb{R}^{m \times n}$ for all $i = 1, 2, \dots, N-1$ and $j = 1, 2, \dots, i$ are the same in both parameterizations and only the constant terms $\tilde{c}_{t+i|t}, c_{t+i|t} \in \mathbb{R}^m$ differ. \blacksquare

Proof. The argument can proceed similarly to that of Proposition 3. In fact, define

$$\tilde{c}_{t+i|t} := c_{t+i|t} - F_{i,1}x_{t+1|t} - F_{i,2}x_{t+2|t} - \cdots - F_{i,i}x_{t+i|t} \quad \forall i = 1, 2, \dots, N-1$$

to obtain the result. \square

Robust MPC (RMPC)

Recall that Robust MPC (RMPC) considers a linear system (A.34) with persistent additive disturbances w_t inside a stationary uncertainty set \mathbb{W} (Assumption 7). The main control task is to keep the (closed-loop) states x_t and inputs u_t within their respective constraint sets \mathbb{X} and \mathbb{U} at all times $t \in \mathbb{N}$.

The literature of contributions to RMPC is vast. The existing approaches can be classified into three main categories [108, Sec. 5.3]: *constraint tightening RMPC*, *min-max RMPC*, and *tube-based RMPC*. Their main ideas are briefly outlined in the following.

(a) Constraint tightening RMPC. For the constraint tightening RMPC approach see, e.g., Chisci et al. [42] and the references therein. Constraint tightening RMPC methods use the nominal cost function and predictions (A.35a,b), while tightening the constraints (A.35c,d) gradually over the prediction horizon [85].

In order to mitigate excessive conservatism in the state predictions, it is common practice to *pre-stabilize* the system by a linear feedback $F \in \mathbb{R}^{m \times n}$, which is computed offline [10]. This means that the inputs are transformed via

$$u_{t|t} = c_{t|t} \quad , \quad u_{t+i|t} = Fw_{t+i-1|t} + c_{t+i|t} \quad \forall i = 1, 2, \dots, N-1 \quad , \quad (\text{A.46})$$

where the *corrective control inputs* $\{c_{t|t}, c_{t+1|t}, \dots, c_{t+N-1|t}\}$ become the new decision variables. The modified FHOC of *constraint tightening RMPC* reads as follows:

$$\min_{c_{t|t}, \dots, c_{t+N-1|t}} \sum_{i=0}^{N-1} \ell(c_{t+i|t}, x_{t+i|t}) + \ell_f(x_{t+N|t}) \quad (\text{A.47a})$$

$$\text{s.t.} \quad x_{t+i+1|t} = Ax_{t+i|t} + Bc_{t+i|t}, \quad x_{t|t} = x_t \quad \forall i = 0, \dots, N-1 \quad , \quad (\text{A.47b})$$

$$c_{t+i|t} \in \mathbb{C}_i \quad \forall i = 0, \dots, N-1 \quad , \quad (\text{A.47c})$$

$$x_{t+i|t} \in \mathbb{X}_i \quad \forall i = 1, \dots, N-1 \quad , \quad (\text{A.47d})$$

$$x_{t+N|t} \in \mathcal{X}_f \quad . \quad (\text{A.47e})$$

Note that the state predictions $\{x_{t+1|t}, x_{t+2|t}, \dots, x_{t+N|t}\}$ in (A.47b) are deterministic and based on the corrective control inputs only. The reason is that the future states x_{t+1}, x_{t+2}, \dots , are considered to be uncertain at time t , due to the disturbances. In order to account for this uncertainty over the prediction horizon, the input and state constraints

(A.46c,d) are tightened according to

$$\mathbb{C}_0 := \mathbb{U} , \quad \mathbb{C}_i := \mathbb{C}_{i-1} \ominus F(A + BF)^{i-1}\mathbb{W} , \quad (\text{A.48a})$$

$$\mathbb{X}_0 := \mathbb{X} , \quad \mathbb{X}_i := \mathbb{X}_{i-1} \ominus (A + BF)^{i-1}\mathbb{W} , \quad (\text{A.48b})$$

where $i = 1, \dots, N - 1$. Here ‘ \ominus ’ and ‘ \oplus ’ denote the *Minkowski sum* and *Pontryagin difference*, respectively,¹³ $F\mathbb{W}_i$ represents a *multiplicative set mapping*.¹³

The terminal set (A.46e) must be selected as a *robust positively invariant* (RPI) set \mathcal{X}_f under the terminal feedback law $u_t = Fx_t$. More concretely, it must be a subset of \mathbb{X}_N , for inputs in \mathbb{C}_N , and disturbances in $(A + BF)^N\mathbb{W}$ (cf. Definition 5, p. 27).

The above terminal condition implies recursive feasibility of the FHOCP and satisfaction of the input and state constraints in closed-loop (see Theorem 5 below). Due to the persistent disturbances, convergence of the trajectory to the origin cannot be ensured in general. Instead, the trajectory remains bounded in an RCI set [42, Thm. 8] and the closed-loop system is *input-to-state stable* [52, Thm. 23].

THEOREM 5—CONSTRAINT TIGHTENING RMPC Let Assumptions 4 and 7 hold. If the modified FHOCP (A.47) is feasible for the initial state x_0 , it remains feasible at all times $t \in \mathbb{N}$ and keeps all states x_t and inputs u_t inside the constraint sets \mathbb{X} and \mathbb{U} . ■

Proof. Suppose (A.47) at time $t \in \mathbb{N}$ has a sequence of inputs $\{c_{t|t}^*, c_{t+1|t}^*, \dots, c_{t+N-1|t}^*\}$ and a sequence of predicted states $\{x_{t|t}^*, x_{t+1|t}^*, \dots, x_{t+N|t}^*\}$ which is feasible with respect to (A.47c,d,e). It suffices to prove that (A.47) at time $t + 1$ has a feasible solution for any $w_t \in \mathbb{W}$, i.e., for the updated state

$$x_{t+1|t+1} = x_{t+1} = \underbrace{Ax_t + Bc_{t|t}^*}_{=x_{t+1|t}} + w_t = x_{t+1|t} + w_t .$$

To see this, define the following sequence of corrective inputs

$$\begin{aligned} c_{t+1|t+1} &:= c_{t+1|t}^* + Fw_t \in \mathbb{C}_1 \oplus F\mathbb{W} = \mathbb{C}_0 , \\ c_{t+2|t+1} &:= c_{t+2|t}^* + (A + BF)Fw_t \in \mathbb{C}_2 \oplus (A + BF)F\mathbb{W} = \mathbb{C}_1 , \\ &\dots \\ c_{t+N-1|t+1} &:= c_{t+1|t}^* + (A + BF)^{N-2}Fw_t \in \mathbb{C}_{N-1} \oplus (A + BF)^{N-2}F\mathbb{W} = \mathbb{C}_{N-2} , \\ c_{t+N|t+1} &:= \underbrace{Fx_{N|t}^*}_{\in \mathbb{C}_N} + (A + BF)^{N-1}Fw_t \in \mathbb{C}_N \oplus (A + BF)^{N-1}F\mathbb{W} = \mathbb{C}_{N-1} , \end{aligned}$$

¹³Let $\mathcal{P}_1, \mathcal{P}_2$ be two subsets of \mathbb{R}^d . The *Minkowsky sum* of \mathcal{P}_1 and \mathcal{P}_2 is defined as $\mathcal{P}_1 \oplus \mathcal{P}_2 := \{p_1 + p_2 \in \mathbb{R}^d \mid p_1 \in \mathcal{P}_1, p_2 \in \mathcal{P}_2\}$, and the *Pontryagin difference* of \mathcal{P}_1 and \mathcal{P}_2 is defined as $\mathcal{P}_1 \ominus \mathcal{P}_2 := \{p \in \mathbb{R}^d \mid p + p_2 \in \mathcal{P}_1 \forall p_2 \in \mathcal{P}_2\}$. The *multiplicative set mapping* $M\mathcal{P}_1 \subset \mathbb{R}^{\bar{d}}$ of a matrix $M \in \mathbb{R}^{\bar{d} \times d}$ is defined as $M\mathcal{P}_1 := \{Mp \in \mathbb{R}^{\bar{d}} \mid p \in \mathcal{P}_1\}$. If $\mathcal{P}_1, \mathcal{P}_2$ are polyhedral sets, then their Minkowski sum and their Pontryagin difference are also polyhedral sets, and so is the image of a multiplicative set mapping by any matrix; cf. [49, 109].

and the corresponding sequence of states

$$\begin{aligned}
 x_{t+1|t+1} &:= x_{t+1|t}^* + w_t \in \mathbb{X}_1 \oplus \mathbb{W} = \mathbb{X}_0 , \\
 x_{t+2|t+1} &:= x_{t+2|t}^* + (A + BF)w_t \in \mathbb{X}_2 \oplus (A + BF)\mathbb{W} = \mathbb{X}_1 , \\
 &\dots \\
 x_{t+N|t+1} &:= x_{t+N|t}^* + (A + BF)^{N-1}w_t \in \mathbb{X}_N \oplus (A + BF)^{N-1}\mathbb{W} = \mathbb{X}_{N-1} , \\
 x_{t+N+1|t+1} &:= (A + BF) \underbrace{x_{t+N|t}^*}_{\in \mathcal{X}_f} + (A + BF)^N w_t \in \mathcal{X}_f .
 \end{aligned}$$

Note that both sequences are feasible with respect to (A.47c,d,e). \square

REMARK 2—AFFINE DISTURBANCE FEEDBACK It is possible to solve for the feedback gain $F \in \mathbb{R}^{m \times n}$, together with the corrective control inputs, by means of convex optimization; see Goulart et al. [52]. More generally, the optimization problem over all ADF policies can be expressed as a convex optimization program (cf. Proposition 4, p. 30). Using ADF for constraint tightening RMPC leads to an improved value of the cost function, while it increases the computational burden of the FHOCP. \blacksquare

(b) Min-max RMPC. Sokaert and Mayne [98] describe a min-max RMPC approach, which considers a scenario tree of extremal disturbance realizations (cf. Figure A.2, p. 22). In particular, suppose that \mathbb{W} is a polytope and let Δ be the (finite) set of its vertices. Then exactly $\text{card}(\Delta)$ branches spawn from each node of the scenario tree, and there is a total of $K = \text{card}(\Delta)^{N-1}$ leaf nodes, or scenarios.¹⁴ Hence the scenario tree is closely related to that of multi-stage stochastic programs, except that no probabilities are assigned to the scenarios.

The modified FHOCP of *min-max RMPC* reads as follows:

$$\min_{u_{t|t}^{(k)}, \dots, u_{t+N-1|t}^{(k)}} \max_{k=1, \dots, K} \sum_{i=0}^{N-1} \ell(u_{t+i|t}^{(k)}, x_{t+i|t}^{(k)}) + \ell_f(x_{t+N|t}^{(k)}) \quad (\text{A.49a})$$

$$\text{s.t.} \quad x_{t+i+1|t}^{(k)} = Ax_{t+i|t}^{(k)} + Bu_{t+i|t}^{(k)} + \delta_{t+i|t}^{(k)}, \quad x_{t|t}^{(k)} = x_t \quad \forall i = 0, \dots, N-1, \quad (\text{A.49b})$$

$$u_{t+i|t}^{(k)} \in \mathbb{U} \quad \forall i = 0, \dots, N-1, \quad (\text{A.49c})$$

$$x_{t+i|t}^{(k)} \in \mathbb{X} \quad \forall i = 0, \dots, N-1, \quad (\text{A.49d})$$

$$x_{t+N|t}^{(k)} \in \mathcal{X}_f, \quad (\text{A.49e})$$

$$u_{t+i|t}^{(k_1)} = u_{t+i|t}^{(k_2)} \quad \forall i \leq b(k_1, k_2), \quad (\text{A.49f})$$

where the constraints must hold for all $k, k_1, k_2 \in \{1, 2, \dots, K\}$ with $k_1 \neq k_2$. Here $\delta_{t+i|t}^{(k)}$

¹⁴Here $\text{card}(\cdot)$ denotes the *cardinality* of a set, i.e., the number of its elements.

denotes the vertex of \mathbb{W} that is chosen by scenario $k \in \{1, 2, \dots, K\}$ at time $t + i$ (stage $i + 1$ of the scenario tree in Figure A.2, p. 22).

Because the system model is linear, the input/state trajectory can be kept inside the convex hull of the predicted trajectories (A.49b), for all possible disturbance sequences $\{w_{t|t}, w_{t+1|t}, \dots, w_{t+N-1|t}\} \in \mathbb{W}^N$. Hence the input and state constraints (A.49c,d,e) are robustly satisfied. Moreover, by convexity, (A.49a) minimizes the maximal cost over all disturbance sequences $w_{t|t}, w_{t+1|t}, \dots, w_{t+N-1|t} \in \mathbb{W}$, by considering only the extremal cases. The terminal set \mathcal{X}_f in (A.49e) is the maximum RPI set under a predetermined terminal feedback law $u_t = Fx_t$ (cf. Definition 5, p. 27). The non-anticipativity constraints (A.49f) require the inputs of any two scenarios k_1, k_2 to be identical up to their branch node $b(k_1, k_2)$; see p. 22.

While min-max RMPC yields a high control performance, it is computationally intractable for most problems, due to the combinatorial growth of the scenario tree.

(c) Tube-based RMPC. Tube-based RMPC has been developed by Mayne et al. [70]; see also [83, Sec. 3.4] and [108, Sec. 5.3.1] for an overview. The basic idea is to split the control inputs into two parts,

$$u_{t+i|t} = F(\tilde{x}_{t+i|t} - x_{t+i|t}) + c_{t+i|t} \quad \forall i = 1, 2, \dots, N-1 . \quad (\text{A.50})$$

The first term in (A.50) keeps the actual state $\tilde{x}_{t+i|t}$ close to a nominal state $x_{t+i|t}$ by means of a linear feedback gain $F \in \mathbb{R}^{n \times m}$. In fact, $\tilde{x}_{t+i|t}$ is inside $x_{t+i|t} \oplus \mathcal{Z}$, where $\mathcal{Z} \subset \mathbb{R}^n$ is an RPI set under the linear feedback $u_t = Fx_t$, called the *tube cross section*, and $x_{t+i|t}$ is the *tube center*. The second term in (A.50) constitutes the part of the control inputs that is used for steering the tube center.

Ideally, the tube cross section \mathcal{Z} is chosen as small as possible to avoid excessive conservatism [70, Sec. 2]. The *minimal* RPI (mRPI) set is not necessarily polytopic; however, it can always be approximated by a polytopic RPI set [81].

The modified FHOC of *tube-based MPC* reads as follows:

$$\min_{c_{t|t}, \dots, c_{t+N-1|t}, x_{t|t}, \dots, x_{t+N|t}} \sum_{i=0}^{N-1} \ell(c_{t+i|t}, x_{t+i|t}) + \ell_f(x_{t+N|t}) \quad (\text{A.51a})$$

$$\text{s.t.} \quad x_{t+i+1|t} = Ax_{t+i|t} + Bc_{t+i|t} \quad \forall i = 0, \dots, N-1 , \quad (\text{A.51b})$$

$$x_t \in x_{t|t} \oplus \mathcal{Z} , \quad (\text{A.51c})$$

$$c_{t+i|t} \in \mathbb{U} \ominus F\mathcal{Z} \quad \forall i = 0, \dots, N-1 , \quad (\text{A.51d})$$

$$x_{t+i|t} \in \mathbb{X} \ominus \mathcal{Z} \quad \forall i = 0, \dots, N-1 , \quad (\text{A.51e})$$

$$x_{t+N|t} \in \mathcal{X}_f . \quad (\text{A.51f})$$

For tube-based MPC, the state predictions $\{x_{t|t}, x_{t+1|t}, \dots, x_{t+N|t}\}$ represent the tube centers. As the initial tube center $x_{t|t}$ is not fixed (in particular, it need not be equal to

x_t), the entire sequence of tube centers are optimization variables, subject to (A.51b,c). The input and state constraints must be tightened as shown in (A.51d,e); however, this tightening is uniform over the prediction horizon, unlike for the constraint tightening RMPC approach. The terminal condition for the tube centers must involve a control Lyapunov function $\ell_f(\cdot)$, on a *control invariant* terminal set \mathcal{X}_f with respect to the tightened constraints $\mathbb{X} \ominus \mathcal{Z}$ and $\mathbb{U} \ominus F\mathcal{Z}$ (cf. Definitions 3 and 4, p. 26).

The following key result for tube-based RMPC has been shown by Mayne et al. [70, Thm. 1].

THEOREM 6—TUBE-BASED RMPC Let Assumptions 4 and 7 hold, \mathcal{Z} be an RPI set of system (A.34) under $u_t = Fx_t$, and the modified FHOC (A.51) be feasible for the initial state x_0 . Then the modified FHOC (A.51) remains feasible at all times $t \in \mathbb{N}$ and the tube-based RMPC keeps the states x_t and inputs u_t inside the constraint sets \mathbb{X} and \mathbb{U} . Moreover, the state trajectory $\{x_t\}_{t \in \mathbb{N}}$ converges exponentially to the set $\mathcal{Z} \subset \mathbb{R}^n$. ■

Stochastic MPC (SMPC)

In contrast to RMPC, SMPC assumes that a probability distribution is available over the disturbance set \mathbb{W} . Assumption 7 is hence replaced with the following Assumption 8.

ASSUMPTION 8—STOCHASTIC MPC (a) The system and input matrices A and B are exactly known for controller design. (b) There exists a probability measure $\mathbf{P}[\cdot]$ on the *disturbance set* \mathbb{W} . ■

There exist some SMPC approaches that also handle uncertainty in the controller model; e.g., [39], [79]. However, they are beyond the discussion of this introduction.

Knowledge of the probability measure provides an algorithm with the additional information of how likely it is for specific disturbances to occur. Extreme disturbance realizations often have a very low associated probability, and this feature is not accounted for by RMPC. As a consequence, the performance of RMPC (in terms of *cumulative closed-loop stage costs*) can possibly be improved by accepting a very low level of transgressions for the state constraints [8]. This goal is the general motivation for the approaches of SMPC.

In many practical applications, the cost function represents a relevant physical or monetary quantity. For example, in building climate control, the energy consumption is minimized while observing constraints on the room temperature levels [76]; or in supply chain management, the inventory cost is minimized while customer demands must be satisfied [95]. In some of these applications, the constraints can be relaxed to so-called *chance constraints* (cf. Section 2.1). This means that they may be violated occasionally, but not too frequently.

Three different categories of SMPC approaches can be distinguished: *recursively feasible SMPC*, *soft constrained SMPC*, and *probabilistically constrained SMPC*. Their main ideas are briefly outlined in the following.

(a) Recursively feasible SMPC. Recursively feasible SMPC approaches resemble those of RMPC closely. In particular, they assume that, in addition to Assumption 8(b), the disturbance set \mathbb{W} is known and compact.

Kouvaritakis et al. [61] introduce an SMPC approach similar to constraint tightening RMPC, except that in (A.47) the first state constraint $x_{t+1|t} \in \mathbb{X}_1$ is exchanged for a chance constraint:

$$\mathbf{P}[x_{t+1|t} \in \mathbb{X}] \geq 1 - \varepsilon . \quad (\text{A.52})$$

The argument of Theorem 5 (p. 32) for recursive feasibility carries over almost analogously, since all other constraints are maintained in their robust versions and \mathbb{X}_1 is a subset of the feasible set of (A.52). The main difference is that now $x_{t+1|t+1} = x_{t+1}$ may violate the constraint set $\mathbb{X}_0 = \mathbb{X}$ with a probability of at most $\varepsilon \in [0, 1]$.

Cannon et al. [36, 38] propose a similar approach based on *stochastic tubes*. The key idea is that the error $(\tilde{x}_{t+i|t} - x_{t+i|t})$ between the actual state $\tilde{x}_{t+i|t}$ and the nominal state $x_{t+i|t}$ is bounded in an ellipsoidal set, with the specified probability $\varepsilon \in [0, 1]$. Analogously to tube-based RMPC, these ellipsoidal sets for all $i = 0, 1, \dots, N$ form a tube and must be steered to a CI terminal set. The algorithm then achieves recursive feasibility and constraint satisfaction in closed-loop operation.

In general, recursively feasible SMPC approaches only mildly reduce the cumulative closed-loop stage costs, as compared to their RMPC counterparts. Furthermore, modeling a support set for the stochastic disturbances is often a delicate task in practice: If \mathbb{W} is chosen too small, it may not contain all disturbances almost surely; if \mathbb{W} is chosen too large, the SMPC suffers from a performance loss.

(b) Soft constrained SMPC. Soft constrained MPC approaches circumvent the problem of state constraint feasibility by moving these constraints into the cost function with a penalty term. This construction is called a “*soft constraint*”. Note that, independently of the kind of uncertainty in MPC, input constraints can always be enforced as hard constraints.

Batina et al. [8] have proposed an early approach of this type. Their basic idea is to introduce a convex penalty function for state constraint violations, $\psi : \mathbb{R}^n \rightarrow \mathbb{R}_{0+}$ with $\psi(\xi) = 0$ for all $\xi \in \mathbb{X}$. The FHOCP is then solved as a large-scale convex program, which is constructed via sample approximation.

The algorithm of Chatterjee et al. [40] includes no cost penalty for state constraint violations other than the generic state cost weight. The inputs are modeled as the weighted sum of bounded basis functions of the past disturbances. Hence they satisfy hard input constraints by appropriate constraints on the sum of the coefficients. The disturbances are assumed to be i.i.d., known only in terms of their first and second order moments. In other words, the approach is robust with respect to all distributions that are compatible with the given second-order moment characteristics [110]. The FHOCP is cast as a *semi-definite program* (SDP), for which efficient algorithms are available based on convex optimization [27, 53].

(c) Probabilistically constrained SMPC. Probabilistically constrained SMPC bounds the state constraint violations in terms of some *risk measure*. Chance constraints, as in (A.1b), are just one example. Other examples are *expectation constraints*, such as

$$\mathbf{E}[g^T x_{t+i|t}] \leq h \quad \forall i = 1, 2, \dots, N \quad (\text{A.53})$$

where $g \in \mathbb{R}^d$, $h \in \mathbb{R}_{0+}$, or *integrated chance constraints*

$$\mathbf{E}[(g^T x_{t+i|t} - h)_{0+}] \leq \alpha \quad \forall i = 1, 2, \dots, N \quad (\text{A.54})$$

where $(\cdot)_{0+} := \max\{\cdot, 0\}$ and $\alpha \in \mathbb{R}_{0+}$.¹⁵ Since these approaches usually consider disturbances with unbounded support, naturally they can not guarantee recursive feasibility in the sense of nominal or robust MPC [83, p.187], [79, Sec. IV.A].

Schwarm and Nikolaou [97] and Li et al. [62] have developed an early approach of this type, with particular focus on applications in process control. Their assumption is for disturbances to be normally distributed and correlated over the prediction horizon, with a known mean and covariance matrix. The FHOCP includes a joint chance constraint on the outputs of all stages over the prediction horizon. As they show, the stochastic FHOCP with open loop predictions can be transformed into an equivalent convex program and hence be solved efficiently.

Oldewurtel et al. [75] extend this SMPC approach to closed-loop predictions, with particular focus on applications in building climate control. The additive state disturbances are assumed to be i.i.d. between different time steps, as well as normally distributed with known mean and covariance matrix. In contrast to [62, 97], the chance constraints are enforced individually per stage. The SMPC algorithm uses closed-loop predictions based on an ADF parameterization, for which the FHOCP remains a convex program.

Primbs and Sung [79] consider an approach for systems with state and input multiplicative disturbances; i.e., not additive as in (A.34). The disturbances are i.i.d. in time, with known mean and covariance, and the SMPC algorithm features closed-loop predictions via ADF. The FHOCP is formulated and solved as an SDP [27, 53]. The linear and quadratic constraints on the states and inputs are satisfied probabilistically in closed-loop operation.

Probabilistically constrained SMPC approaches perform well for some systems with stochastic uncertainties. However, they place strong assumptions on the random disturbances, such as independence in time or a normal distribution. Moreover, most of the aforementioned approaches involve a heavy computational burden—which is potentially prohibitive for systems of higher dimensions.

¹⁵In the above sense, probabilistic constraints aim at limiting the *tail risk* of a scalar random variable, which is defined as a function of the state. The range of possible risk measures is vast; see Artzner et al. [6] for an overview. For example, chance constraints correspond to the “*value-at-risk*” (VaR) and integrated chance constraints to the “*conditional value-at-risk*” (CVaR) risk measures.

Scenario-Based MPC (SCMPC)

SMPC approaches have to rely on probability distribution functions (or classes of probability distributions) which are fully characterized by a finite parameterization, e.g., by their mean and variance. More general probability distribution functions, however, appear in many practical applications. They have to be approximated by a finite set of samples (or “*particles*”) in order to be amenable for numerical computation.

Sample-based algorithms have already been successfully implemented in many areas of engineering, such as filtering problems [44]. These algorithms are often easy to implement and they can also handle *implicit* uncertainty models—i.e., models for which a probability distribution is not given explicitly, but for which samples can be obtained empirically or by simulation (e.g., a Markov chain model).

The main idea of SCMPC is to set up the FHOCP based on full sample paths for the uncertainty (or “*scenarios*”) over the prediction horizon. In this way, SCMPC does not require any probabilistic computations, unlike SMPC, but it accounts (implicitly) for the probability of the uncertainty scenarios, unlike RMPC. The drawback of SCMPC is that its policy becomes *randomized*, because it relies on randomly extracted scenarios.

Blackmore et al. [22] have shown how scenarios can, in principle, be used for approximating almost any stochastic expression in the FHOCP. For example, expectation constraints can be approximated by the average over scenarios, chance constraints by a fraction of scenarios, etc. [22, Tab. I]. In theory, the approximations become exact as the sample size goes to infinity, by the law of large numbers [100]. In practice, the quality of the approximation of chance constraints becomes almost unchanged above a certain sample size [22, Sec. VII]. This threshold, however, can generally be quite high and it is not known a priori.

Therefore, subsequent contributions have examined the use of the scenario approach for a rigorous analysis of chance constraints under finite sample sizes. The basic idea is to formulate the stochastic FHOCP as a relaxed version of the robust counterpart, where all constraints over the prediction horizon must hold with a specified probability, as in the approach of Calafiore and Fagiano [32]:

$$\mathbf{P}[u_{t+i|t} \in \mathbb{U}, x_{t+i|t} \in \mathbb{X}, x_{t+N|t} \in \mathcal{X}_f \ \forall i = 0, 1, \dots, N-1] \geq 1 - \varepsilon . \quad (\text{A.55})$$

As (A.55) represents a single chance constraint, the FHOCP turns into the form of UP[ε], and can hence be solved with the classic scenario approach (see Section 2.1).

Prandini et al. [77] show how closed-loop predictions, in the form of ADF, can be integrated into this framework. Deori et al. [43] consider closed-loop predictions by augmenting the corrective control inputs with a weighted sum of bounded basis functions of past disturbances, similar to [40]. Hence the input constraints remain satisfied almost surely. The algorithm of Matuško and Borrelli [68] provides an extension for the automatic removal of constraints, included in the solution to the FHOCP.

Note that all contributions in this framework have used a single chance constraint

formulation. Hence the corresponding sample size (due to Theorem 1) is based on the full dimension of the FHOCP, $d = Nm$. This means that the FHOCP potentially contains a large number of constraints, in particular for a long prediction horizon N . Furthermore, all of the existing approaches are limited to solving the FHOCP as a chance-constrained optimization problem and they do not consider the properties of the closed-loop system. These issues lead to the main contributions of this dissertation, described in the following Parts B,C,D,E.

Conclusion

MPC is a powerful concept for multivariable control problems with constraints. The underlying theory has matured over the past two decades and it has been successfully implemented in a wide range of applications. Systems without a reasonably accurate nominal model, however, remain open to theoretical investigations. In general, uncertain predictions are much more difficult to handle computationally than the certainty-equivalent case—yet they are critical to the system performance in many practical cases.

Several paradigms for the inclusion of uncertainty in the model-based predictions compete: set bounds (RMPC), probability distributions (SMPC), and scenarios (SCMPC). Besides accuracy of the predictions, computational complexity is an important factor to consider, because the FHOCP must be solved on-line and within a prescribed time interval.

References

- [1] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions*. Dover Publications, New York, 9th edition, 1970.
- [2] H. Amann and J. Escher. *Analysis*, volume I. Birkhäuser Verlag, Basel et al., 3rd edition, 2005.
- [3] H. Amann and J. Escher. *Analysis*, volume II. Birkhäuser Verlag, Basel et al., 2nd edition, 2008.
- [4] H. Amann and J. Escher. *Analysis*, volume III. Birkhäuser Verlag, Basel et al., 2nd edition, 2009.
- [5] B.D.O. Anderson and J.B. Moore. *Optimal Control – Linear Quadratic Methods*. Dover, Mineola (NY), United States, 2007.
- [6] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [7] B. Bank, J. Guddat, D. Klatte, B. Kummer, and K. Tammer. *Non-Linear Parametric Optimization*. Birkhäuser Verlag, Basel et al., 1983.
- [8] I. Batina, A.A. Stoorvogel, and S. Weiland. Optimal control of linear, stochastic systems with state and input constraints. In *41st IEEE Conference on Decision and Control*, Las Vegas (NV), United States, 2002.
- [9] R.E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, 1957. Dover 2003 reprint edition.
- [10] A. Bemporad. Reducing the conservativeness in predictive control of constrained systems with disturbances. In *37th IEEE Conference on Decision and Control*, Tampa (FL), United States, 1998.
- [11] A. Bemporad, F. Borrelli, and M. Morari. The explicit solution of constrained lp-based receding horizon control. In *39th IEEE Conference on Decision and Control*, Sydney, Australia, 2000.
- [12] A. Bemporad, F. Borrelli, and M. Morari. Minmax control of constrained uncertain discrete-time linear systems. *IEEE Transactions on Automatic Control*, 48(9):1600–1606, 2003.

- [13] A. Bemporad, M. Morari, V. Dua, and E. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38:3–20, 2002.
- [14] J.F. Benders. Partitioning procedure for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252, 1962.
- [15] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, Nashua (NH), 2nd edition, 1999.
- [16] D. Bertsekas. *Convex Analysis and Optimization*. Athena Scientific, Nashua (NH), 2003.
- [17] D. Bertsekas. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11:310–334, 2005.
- [18] D.P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, Belmont (MA), 3rd edition, 2005.
- [19] D.P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, Belmont (MA), 4th edition, 2012.
- [20] P. Billingsley. *Probability and Measure*. John Wiley & Sons, New York et al., 3rd edition, 1995.
- [21] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, New York, 2nd edition, 2011.
- [22] L. Blackmore, M. Ono, A. Bektassov, and B. Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Transactions on Robotics*, 26(3):502–516, 2010.
- [23] F. Blanchini. Set invariance in control. *Automatica*, 35:1747–1767, 1999.
- [24] V.I. Bogachev. *Measure Theory*. Springer, Heidelberg et al., 2007.
- [25] F. Borrelli, A. Bemporad, and M. Morari. Geometric algorithm for multiparametric linear programming. *Journal of Optimization Theory and Applications*, 118(3):515–540, 2003.
- [26] F. Borrelli, A. Bemporad, and M. Morari. Predictive control for linear and hybrid systems. Lecture notes ME290J, University of California at Berkeley, Berkeley (CA), United States, 2011. available at <http://www.mpc.berkeley.edu/mpc-course-material>.
- [27] S. Boyd and L. Vandenberghe. Semidefinite programming. *SIAM Review*, 38-1:49–95, 1996.
- [28] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.
- [29] G. Calafiore and M.C. Campi. Uncertain convex programs: Randomized solutions and confidence levels. *Mathematical Programming, Series A*, 102-1:25–46, 2005.
- [30] G.C. Calafiore. On the expected probability of constraint violation in sampled convex programs. *Journal of Optimization Theory and Applications*, 143:405–412, 2009.

- [31] G.C. Calafiore. Random convex programs. *SIAM Journal on Optimization*, 20(6):3427–3464, 2010.
- [32] G.C. Calafiore and L. Fagiano. Robust model predictive control via scenario optimization. *IEEE Transactions on Automatic Control*, 58(1):219–224, 2013.
- [33] E.F. Camacho and C. Bordons. Nonlinear model predictive control: An introductory review. In R. Findeisen et al., editor, *Assesment and Future Directions*, pages 1–16. Springer, Berlin et al., 2007.
- [34] M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19:1211–1230, 2008.
- [35] M.C. Campi and S. Garatti. A sampling and discarding approach to chance-constrained optimization: Feasibility and optimality. *Journal of Optimization Theory and Applications*, 148:257–280, 2011.
- [36] M. Cannon, Q. Cheng, B. Kouvaritakis, and S.V. Raković. Stochastic tube MPC with state estimation. *Automatica*, 48(3):536–541, 2012.
- [37] M. Cannon and B. Kouvaritakis. Optimizing prediction dynamics for robust MPC. *IEEE Transactions on Automatic Control*, 50(11):1892–1897, 2005.
- [38] M. Cannon, B. Kouvaritakis, S.V. Raković, and Q. Cheng. Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Transactions on Automatic Control*, 56(1):194–200, 2011.
- [39] M. Cannon, B. Kouvaritakis, and X. Wu. Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty. *IEEE Transactions on Automatic Control*, 54(7):1626–1632, 2009.
- [40] D. Chatterjee, P. Hokayem, and J. Lygeros. Stochastic receding horizon control with bounded control inputs: A vector space approach. *IEEE Transactions on Automatic Control*, 56(11):2704–2710, 2011.
- [41] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952.
- [42] L. Chisci, J.A. Rossiter, and G. Zappa. Systems with persistent disturbances: Predictive control with restricted constraints. *Automatica*, 37:1019–1028, 2001.
- [43] L. Deori, S. Garatti, and M. Prandini. Stochastic constrained control: Trading performance for state constraint feasibility. In *12th European Control Conference*, Zurich, Switzerland, 2013.
- [44] B. Doucet and A.M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan and D. Rozovskiĭ, editors, *The Oxford Handbook of Nonlinear Filtering*, pages 656–704. Oxford University Press, Oxford, 2011.
- [45] F.Y. Edgeworth. The mathematical theory of banking. *Journal of the Royal Statistical Society*, 51(1):113–127, 1888.
- [46] M. Einsiedler and T. Ward. Functional analysis I & II. Lecture notes, Eidgenössische Technische Hochschule (ETH) Zürich, Zürich, Switzerland, 2012.

- [47] R. Findeisen and F. Allgöwer. An introduction to nonlinear model predictive control. In *21st Benelux Meeting on Systems and Control*, Veldhoven, Netherlands, 2002.
- [48] G.F. Franklin, J.D. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Prentice Hall, Upper Saddle River (NJ), 5th edition, 2009.
- [49] K. Fukuda. Polyhedral computation. Lecture notes, Eidgenössische Technische Hochschule (ETH) Zürich, Zürich, Switzerland, 2011.
- [50] T. Gal and J. Nedoma. Multiparametric linear programming. *Management Science*, 18(7):406–422, 1972.
- [51] C.E. García, D.M. Prett, and M. Morari. Model predictive control: Theory and practice. *Automatica*, 23(3):335–348, 1989.
- [52] P.J. Goulart, E.C. Kerrigan, and J.M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42:522–533, 2006.
- [53] M. Grant and S. Boyd. *CVX Users’ Guide, version 1.22*. Palo Alto (CA), 2012.
- [54] L. Grüne and J. Pannek. *Nonlinear Model Predictive Control*. Springer, London et al., 2011.
- [55] M. Herceg, M. Kvasnica, C.N. Jones, and M. Morari. Multi-Parametric Toolbox 3.0. In *12th European Control Conference*, pages 502–510, Zürich, Switzerland, 2013. available at <http://control.ee.ethz.ch/~mpt>.
- [56] C.N. Jones, E.C. Kerrigan, and J.M. Maciejowski. Lexicographic perturbation for multiparametric linear programming with applications to control. *Automatica*, 43(10):1808–1816, 2007.
- [57] P. Kall and J. Mayer. *Stochastic Linear Programming*. Springer, New York et al., 2nd edition, 2011.
- [58] S.S. Keerthi and E.G. Gilbert. Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations. *Journal of Optimization Theory and Applications*, 57(2):265–293, 1988.
- [59] E.C. Kerrigan. *Robust Constraint Satisfaction: Invariant Sets and Predictive Control*. Ph.D. thesis, University of Cambridge, Cambridge, United Kingdom, 2000.
- [60] H.K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, 3rd edition, 2002.
- [61] B. Kouvaritakis, M. Cannon, S.V. Raković, and Q. Cheng. Explicit use of probabilistic distributions in linear predictive control. *Automatica*, 46:1719–1724, 2010.
- [62] P. Li, M. Wendt, and G. Wozny. A probabilistically constrained model predictive controller. *Automatica*, 38:1171–1176, 2002.
- [63] G. Ludyk. *Theoretische Regelungstechnik*, volume 1. Springer, Berlin et al., 1995.
- [64] G. Ludyk. *Theoretische Regelungstechnik*, volume 2. Springer, Berlin et al., 1995.

References

- [65] D. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, New York et al., 1969.
- [66] D. Luenberger and Y. Ye. *Linear and Nonlinear Programming*. Springer, Berlin et al., 3rd edition, 2008.
- [67] J.M. Maciejowski. *Predictive Control with Constraints*. Pearson Education, Harlow, 2002.
- [68] J. Matuško and F. Borrelli. Scenario-based approach to stochastic linear predictive control. In *51st IEEE Conference on Decision and Control*, Maui (HI), United States, 2012.
- [69] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.
- [70] D.Q. Mayne, M.M. Seron, and S.V. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41:219–224, 2005.
- [71] J. Munkres. *Topology*. Prentice Hall, Upper Saddle River, second edition edition, 2000.
- [72] R.M. Murray, K.J. Åström, S.P. Boyd, R.W. Brockett, and G. Stein. Future directions in control in an information rich world. *IEEE Control Systems Magazine*, 23(2):20–33, 2003.
- [73] N.S. Nise. *Control Systems Engineering*. John Wiley & Sons, Hoboken (NJ), 4th edition, 2004.
- [74] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.
- [75] F. Oldewurtel, C.N. Jones, A. Parisio, and M. Morari. Stochastic model predictive control for energy efficient building climate control. *IEEE Transactions on Control Systems Technology*, 22(3):1198–1205, 2014.
- [76] F. Oldewurtel, A. Parisio, C.N. Jones, D. Gyliastras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27, 2012.
- [77] M. Prandini, S. Garatti, and J. Lygeros. A randomized approach to stochastic model predictive control. In *51st IEEE Conference on Decision and Control*, Maui (HI), United States, 2012.
- [78] A. Prékopa. *Stochastic Programming*. Kluwer, Dordrecht et al., 1995.
- [79] J.A. Primbs and C.H. Sung. Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Transactions on Automatic Control*, 54(2):221–230, 2012.
- [80] S.J. Qin and T.A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11:733–764, 2003.

- [81] S.V. Raković, E.C. Kerrigan, K.I. Kouramas, and D.Q. Mayne. Invariant approximations of the minimal robust positively invariant set. *IEEE Transactions on Automatic Control*, 50(3):406–410, 2005.
- [82] J.B. Rawlings. A tutorial overview of model predictive control. *IEEE Control Systems Magazine*, 20(3):38–52, 2000.
- [83] J.B. Rawlings and D.Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, Madison (WI), 2009.
- [84] J. Richalet, A. Rault, J.L. Testud, and J. Papon. Model predictive heuristic control: Applications to industrial processes. *Automatica*, 14:413–428, 1978.
- [85] A. Richards and J. How. Robust stable Model Predictive Control with constraint tightening. In *American Control Conference*, Minneapolis (MN), United States, 2006.
- [86] R.T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.
- [87] J.A. Rossiter. *Model-Based Predictive Control: A Practical Approach*. CRC Press, Boca Raton (FL), 2004.
- [88] S. Sastry. *Nonlinear Systems: Analysis, Stability, and Control*. Springer, New York et al., 1999.
- [89] G. Schildbach, G.C. Calafiore, L. Fagiano, and M. Morari. Randomized model predictive control for stochastic linear systems. In *American Control Conference*, Montréal, Canada, 2012.
- [90] G. Schildbach, L. Fagiano, C. Frei, and M. Morari. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, (under review).
- [91] G. Schildbach, L. Fagiano, and M. Morari. Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4):2479–2501, 2013.
- [92] G. Schildbach, P. Goulart, and M. Morari. The linear quadratic regulator with chance constraints. In *12th European Control Conference*, Zurich, Switzerland, 2013.
- [93] G. Schildbach, P. Goulart, and M. Morari. Linear controller design for chance constrained systems. *Automatica*, (submitted).
- [94] G. Schildbach and M. Morari. The scenario approach for two-level stochastic programs with expected shortfall probability. *International Journal of Production Economics*, (submitted).
- [95] G. Schildbach and M. Morari. Scenario-based model predictive control for multi-echelon supply chain management. *European Journal of Operational Research*, (submitted).
- [96] G. Schildbach, M.N. Zeilinger, M. Morari, and C.N. Jones. Input-to-state stabilization of low-complexity model predictive controllers for linear systems. In *50th IEEE Conference on Decision and Control*, Orlando (FL), United States, 2011.

References

- [97] A.T. Schwarm and M. Nikolaou. Chance-constrained model predictive control. *AIChE Journal*, 45(8):1743–1752, 1999.
- [98] P.O.M. Scokaert and D.Q. Mayne. Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43(8):1136–1142, 1998.
- [99] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia, 2009.
- [100] A.N. Shiryaev. *Probability*. Springer, New York et al., 2nd edition, 1996.
- [101] J. Skaf and S. Boyd. Nonlinear q-design for convex stochastic control. *IEEE Transactions on Automatic Control*, 54(10):2426–2430, 2009.
- [102] S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control - Analysis and Design*. John Wiley & Sons, Hoboken (NJ), 2nd edition, 2005.
- [103] M.S. Sohdi and C.S. Tang. *Managing Supply Chain Risk*. Springer, New York et al., 2012.
- [104] L.A. Steen and J.A. Seebach. *Counterexamples in Topology*. Holt, Rinehart and Winston, New York et al., 1970.
- [105] R.M. van Slyke and R. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17(4):638–663, 1969.
- [106] M. Vidyasagar. *Nonlinear Systems Analysis*. Prentice Hall, Englewood Cliffs (NJ), 2nd edition, 1993.
- [107] D. Werner. *Funktionalanalysis*. Springer, Berlin et al., 6th edition, 2007.
- [108] M.N. Zeilinger. *Real-time Model Predictive Control*. Ph.D. thesis, Eidgenössische Technische Hochschule Zürich, Zürich, Switzerland, 2011.
- [109] G.M. Ziegler. *Lectures on Polytopes*. Springer, New York et al., 1st edition, 2007.
- [110] S. Zymler, D. Kuhn, and B. Rustem. Distributionally robust joint chance constraints with second-order moment information. *Mathematical Programming, Series A*, 137:167–198, 2013.

Part B

Structure in Multi-Stage Stochastic Programs

Paper I

Randomized Solutions to Convex Programs with Multiple Chance Constraints

Georg Schildbach · Lorenzo Fagiano · Manfred Morari

Abstract

The scenario-based optimization approach (“scenario approach”) provides an intuitive way of approximating the solution to chance-constrained optimization programs, based on finding the optimal solution under a finite number of sampled outcomes of the uncertainty (“scenarios”). A key merit of this approach is that it neither requires explicit knowledge of the uncertainty set, as in robust optimization, nor of its probability distribution, as in stochastic optimization. The scenario approach is also computationally efficient because it only requires the solution to a convex optimization program, even if the original chance-constrained problem is non-convex. Recent research has obtained a rigorous foundation for the scenario approach, by establishing a direct link between the number of scenarios and bounds on the constraint violation probability. These bounds are tight in the general case of an uncertain optimization problem with a single chance constraint.

This paper shows that the bounds can be improved in situations where the chance constraints have a limited “support rank”, meaning that they leave a linear subspace unconstrained. Moreover, it shows that also a combination of multiple chance constraints, each with individual probability level, is admissible. As a consequence of these results, the number of scenarios can be reduced from that prescribed by the existing theory for problems with the indicated structural property. This leads to an improvement in the objective value and a reduction in the computational complexity of the scenario approach. The proposed extensions have many practical applications, in particular high-dimensional problems such as multi-stage uncertain decision problems or design problems of large-scale systems.

This article has been published in the SIAM Journal on Optimization.

DOI: <http://dx.doi.org/10.1137/120878719>

©2013 SIAM Journal on Optimization.

1. Introduction

Optimization is ubiquitous in modern problems found in engineering, logistics, and other sciences. A common pattern is that a decision or design variable $x \in \mathbb{R}^d$ has to be selected from a subset of \mathbb{R}^d , as described by constraints $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, and its quality is measured against some objective or cost function $f_0 : \mathbb{R}^d \rightarrow \mathbb{R}$:

$$\min_{x \in \mathbb{R}^d} f_0(x) , \quad (\text{B.1a})$$

$$\text{s.t. } f_i(x) \leq 0 \quad \forall i = 1, 2, \dots, N . \quad (\text{B.1b})$$

1.1 Chance-Constrained Optimization

Unfortunately, in many practical applications the underlying problem data is uncertain. This uncertainty shall be represented with an abstract variable $\delta \in \Delta$, where Δ is an uncertainty set whose nature is not specified. The uncertainty may affect the objective function f_0 and/or the constraints f_i . Thus for a particular decision x it becomes uncertain what objective value is achieved and/or whether the constraints are indeed satisfied. The second situation represents a particular challenge, as good solutions are usually located on the boundary of the feasible set.

This gives rise to a trade-off problem between the (uncertain) objective value and the robustness of the chosen decision to a constraint violation. A large variety of approaches addressing this issue have been proposed in the areas of robust and stochastic optimization [3–5, 14, 15, 17, 19, 21], with the preferred method of choice depending on the requirements of the application at hand.

In many practical applications, δ can be assumed to be of a stochastic nature. In this case, the formulation of *chance constraints*, where the decision variable x has to be feasible with a least probability $(1 - \varepsilon)$ for $\varepsilon \in (0, 1)$, has proven to be an appropriate concept for handling the uncertainty in the constraints. However, chance-constrained optimization problems are usually very difficult to solve. The *scenario approach*, as explained below, represents an attractive method for finding an ‘approximate solution’ to stochastic programs, since it is both intuitive and computationally efficient.

1.2 The Scenario Approach

Recent contributions [8–12] have revealed the theoretical links between the scenario approach and the solution to an optimization problem with a linear objective function and a single chance constraint (SCP):

$$\min_{x \in \mathbb{X}} c^T x , \quad (\text{B.2a})$$

$$\text{s.t. } \Pr[f(x, \delta) \leq 0] \geq (1 - \varepsilon) . \quad (\text{B.2b})$$

Here $\mathbb{X} \subset \mathbb{R}^d$ is a compact and convex set, c^T denotes the transpose of a vector $c \in \mathbb{R}^d$, $\Pr[\cdot]$ is the probability measure on the uncertainty set Δ , $f : \mathbb{R}^d \times \Delta \rightarrow \mathbb{R}$ is a convex function in its first argument $x \in \mathbb{R}^d$ for \Pr -almost every uncertainty $\delta \in \Delta$, and ε is some value in the open real interval $(0, 1)$.

The chance constraint (B.2b) is interpreted as follows. For any given $x \in \mathbb{R}^d$, the left-hand side represents the probability of the event that x indeed belongs to the feasible set. Written more properly,

$$\Pr[f(x, \delta) \leq 0] := \Pr\{\delta \in \Delta \mid f(x, \delta) \leq 0\} , \quad (\text{B.3})$$

however, the left-hand side notation is kept throughout for brevity. Note that x is considered to be a *feasible point* of the chance constraint (B.2b) if this probability is at least $(1 - \varepsilon)$.

REMARK B.1—PROBLEM FORMULATION The formulation of the SCP encompasses a vast range of problems, namely any uncertain optimization problem that becomes convex if the value of δ were fixed. (a) Any uncertain convex objective function $f(\cdot, \delta)$ can be included by an epigraph reformulation, with the new objective being a scalar and hence linear [7, Sec. 3.1.7]. (b) Joint chance constraints, where x must satisfy multiple convex constraints simultaneously with probability $(1 - \varepsilon)$, are covered since the intersection of convex sets is convex. (c) Additional deterministic, convex constraints can be included by intersection with the compact set \mathbb{X} . ■

The characterization of the feasible set of a chance constraint requires exact knowledge of the probability distribution of δ . Moreover, the feasible set is non-convex and difficult to express explicitly, except for very special cases [5, 14, 19, 21]. This makes the SCP, in full generality and especially in higher dimensions d , an extremely difficult problem to solve.

The scenario approach can be used to find an *approximate solution* to the SCP, which is considered to be any point in \mathbb{X} that is feasible for the chance constraint with some given (very high) *confidence* $(1 - \theta) \in (0, 1)$. This problem is usually not as hard, if an approximate solution is chosen in a low-violation region of the decision space (with high confidence). However, then the resulting objective value may be poor, in which case the approximate solution shall be called “*conservative*”. Clearly, it is of major interest to find approximate solutions that are the least conservative (i.e., with an objective value as low as possible), and this is the goal of the scenario approach.

The basic idea of the scenario approach is to draw a specific number $K \in \mathbb{N}$ of samples (“*scenarios*”) from the uncertainty δ , and to take the optimal solution that is feasible under all of these scenarios (“*scenario solution*”) as an approximate solution. Computing the scenario solution involves a deterministic optimization program (“*scenario program*”), which is obtained by replacing the chance constraint (B.2b) with the K sampled deterministic constraints.

By construction, the scenario program is a deterministic, convex optimization program

that can be solved efficiently by standard algorithms [7, 16, 18]. Moreover, the scenario approach is distribution-free in the sense that it does not rely on a particular mathematical model for the distribution of δ , or even its support set Δ . In fact, both may be unknown; the only requirements are stated in the following assumption.

ASSUMPTION B.1—UNCERTAINTY (a) The uncertainty δ is a random variable with (possibly unknown) probability measure \Pr and support set Δ . (b) A sufficient number of independent random samples from δ can be obtained. ■

Note that Assumption B.1 is fairly general. It could even be argued that the scenario approach is at the heart of any robust and stochastic optimization method, because either the uncertainty set Δ or the probability distribution of δ are usually constructed based on some (necessarily finite) experience of the uncertainty.

Tight bounds for the proper choice of the sample size K are established by [10, 11], when linking it directly to the probability with which the scenario solution violates the chance constraint (B.2b). Moreover, [10, 12] show that the theory can be extended to the case where $R \leq K$ sampled constraints are discarded *a posteriori*, that is after observing the outcomes of the K samples. While this increases the complexity of the scenario approach (in terms of data requirement and computation), it can be used to improve the objective value achieved by the scenario solution. In fact, the scenario solution can be shown to converge to the exact solution of (B.2) when the number of discarded constraints are increased, given that some mild technical assumptions hold, cf. [12, Sec. 4.4]

1.3 Novel Contributions

From a practical point of view, the strongest appeal of the scenario approach is the facility of its application and the low computational complexity. It becomes particularly attractive for uncertain optimization problems in higher dimensions, as these occur frequently in fields such as engineering or logistics. In these cases, an uncertain constraint will often not involve all decision variables simultaneously, as allowed by the general case of (B.2b). Instead, multiple uncertain constraints may be present, each of them involving only a subset of the decision variables.

EXAMPLE B.1—MULTI-STAGE DECISION PROBLEMS An important example are uncertain *multi-stage decision problems* [5, Cha. 7], [14, Cha. 8] [19, Cha. 13] [21, Cha. 3], which occur in many fields such as production planning, portfolio optimization, or control theory. The basic setting is that some *decision* (e.g., on production quantities, buy/sell orders, or control inputs) has to be taken repeatedly at a finite number of time steps. Each decision affects the *state* of the system (e.g., inventory level, portfolio, or state variable) at the subsequent time step. Besides the decision, the state is also subject to uncertain influences (e.g., customer demand, price fluctuations, or dynamic disturbances). If constraints on the state variables are present (e.g., service levels, value at risk, or safety regions), this adds multiple uncertain constraints (one for the state of each time step) to the overall decision problem. Further deterministic constraints may hold

for the decision variables, for example. The special structure of such a problem is that a constraint on the state at some time step involves only the decisions made prior to this time step, while the decisions afterwards are not involved. ■

This paper extends the theory of the scenario approach for problems where a single (or multiple) chance constraint(s) are present that involve only a subset of the decision variables. More precisely, the chance constraint(s) may affect only a certain subspace of the decision space, whose dimension will be called its “*support rank*”. Other constraints, either deterministic or uncertain, cover the directions that are left unconstrained, so that the solution remains bounded.

The main result of this paper is that an uncertain constraint with a lower support rank can only supply a lower number of *support constraints* [8, 10, 11], and therefore its associated sample size can be reduced. This leads to a subtle shift from the idea of a “*problem dimension*” in the existing theory to that of a “*support dimension*” of a particular chance constraint. Moreover, it requires an extension of the existing theory to cope with multiple chance constraints in the uncertain optimization program. Finally, the approach of constraint removal *a posteriori* is carried over almost analogously to this extended setting.

From a practical point of view, these extensions improve on the merits of the scenario approach for problems that have a structure described above. In particular, the lower sample sizes reduce the computational complexity of the scenario approach and simultaneously improve the objective value of the scenario solution. At the same time, the feasibility guarantees for the scenario solution remain as strong as before. Hence the extensions of this paper, when applicable, offer only advantages over the existing results on the scenario approach.

1.4 Organization of the Paper

Section 2 contains the problem statement. Section 3 introduces some background on its properties, and states the rigorous definitions for the “*support dimension*” and the “*support rank*” of a chance constraint. Section 4 contains the main results of this paper, which give the improved sample bounds in the presence of a single (or multiple) chance constraint(s) of limited support rank. Section 5 extends this theory to the sampling-and-discarding procedure, which can be used to improve the objective value of the scenario solution, at the price of larger data requirements and an increased computational complexity. Section 6 presents a brief numerical example that demonstrates the application of the presented theory, as well as its potential benefits when compared to existing results.

2. Problem Formulation

This section introduces the generalized problem formulation with multiple chance constraints, the corresponding scenario program, and some basic terminology.

2.1 Stochastic Program with Multiple Chance Constraints

Consider the following extension of the SCP to an optimization problem with linear objective function and multiple chance constraints (MCP):

$$\min_{x \in \mathbb{X}} c^\top x, \quad (\text{B.4a})$$

$$\text{s.t.} \quad \Pr[f_i(x, \delta) \leq 0] \geq (1 - \varepsilon_i) \quad \forall i \in \mathbb{N}_1^N, \quad (\text{B.4b})$$

where i is the chance constraint index in $\mathbb{N}_1^N := \{1, 2, \dots, N\}$. The remarks for the SCP in Section 1.2 apply analogously; in particular the following key assumption is made.

ASSUMPTION B.2—CONVEXITY The constraint functions $f_i : \mathbb{R}^d \times \Delta \rightarrow \mathbb{R}$ of all chance constraints $i \in \mathbb{N}_1^N := \{1, \dots, N\}$ are convex in their first argument $x \in \mathbb{R}^d$ for Pr-almost every $\delta \in \Delta$. ■

Other than Assumption B.2, the dependence of the functions $f_i(x, \delta)$ on the uncertainty δ is completely generic.

The use of “min” instead of “inf” in (B.4a) is justified by the fact that the feasible set of a single chance constraint is closed under fairly general assumptions [14, Thm. 2.1]. This implies that the feasible set of the MCP is compact, due to the presence of \mathbb{X} , and the infimum is indeed attained.

It remains a standing assumption that the σ -algebra of Pr-measurable sets in Δ is large enough to contain all sets whose probability is measured in this paper, like the ones in (B.4b), cf. [11, p. 4].

In order to avoid technical issues, which are of little relevance for most practical applications, the following is assumed, cf. [11, Ass. 1].

ASSUMPTION B.3—EXISTENCE AND UNIQUENESS (a) Problem (B.4) admits at least one feasible point. By the compactness of \mathbb{X} , this implies that there exists at least one optimal point of (B.4). (b) If there are multiple optimal points of (B.4), a unique one is selected by the help of a *tie-break rule* (e.g., the lexicographic order on \mathbb{R}^d). ■

In principle, an approximate solution to the MCP can be obtained by the classic scenario approach. Namely, a SCP can be set up with the same objective function (B.2a) as the MCP, and a chance constraint (B.2b) defined by

$$f(x, \delta) := \max\{f_1(x, \delta), \dots, f_N(x, \delta)\} \quad \text{and} \quad \varepsilon := \min\{\varepsilon_1, \varepsilon_2, \dots, \varepsilon_N\}. \quad (\text{B.5})$$

Note that $f(x, \delta)$ is convex in x for almost every δ , since the pointwise maximum of convex functions is convex. Any feasible point of this SCP is also a feasible point of the MCP, and hence an approximate solution to the SCP with confidence $(1 - \theta)$ is also an approximate solution to the MCP with confidence $(1 - \theta)$.

However, this procedure introduces a considerable amount of conservatism, because it requires the scenario solution to simultaneously satisfy *all* constraints $i = 1, \dots, N$ with

the *highest* of all probabilities $(1 - \varepsilon_i)$. Clearly, this conservatism becomes more severe if the number of chance constraints N is large and there is a great variation in the values of ε_i .

2.2 The Extended Scenario Approach

The extended scenario approach of this paper can be used to compute an approximate solution of the MCP, which is a feasible point of every chance constraint $i = 1, \dots, N$ with a given confidence probability of $(1 - \theta_i)$. The key difference from the classic scenario approach is that each chance constraint $i \in \mathbb{N}_1^N$ is sampled separately, and with an individual sample size $K_i \in \mathbb{N}$.

Let the *random samples* pertaining to constraint i be denoted $\delta^{(i, \kappa_i)}$, where $\kappa_i \in \{1, \dots, K_i\}$, and for brevity also as the collective *multi-sample* $\omega^{(i)} := \{\delta^{(i, 1)}, \dots, \delta^{(i, K_i)}\}$. The collection of all samples is combined in an overall multi-sample $\omega := \{\omega^{(1)}, \dots, \omega^{(N)}\}$, with the total number of samples given by $K := \sum_{i=1}^N K_i$. All of these samples can be considered “identical copies” of the random uncertainty δ , in the sense that they are themselves random variables and satisfy the following key assumption.

ASSUMPTION B.4—INDEPENDENCE AND IDENTICAL DISTRIBUTION The sampling procedure is designed such that the set of all random samples, together with the actual random uncertainty,

$$\bigcup_{i \in \mathbb{N}_1^N} \{\delta^{(i, 1)}, \dots, \delta^{(i, K_i)}\} \cup \{\delta\}$$

form a set of *independent and identically distributed (i.i.d.)* random variables. ■

The multi-sample ω is an element of Δ^K , the K -th product of the uncertainty set Δ , and it is distributed according to \mathbf{P}^K , the K -th product of the measure \mathbf{P} . The scenario program for multiple chance constraints ($\text{MSP}[\omega^{(1)}, \dots, \omega^{(N)}]$) is constructed as follows:

$$\min_{x \in \mathbb{X}} c^T x, \tag{B.6a}$$

$$\text{s.t. } f_i(x, \delta^{(i, \kappa_i)}) \leq 0 \quad \forall \kappa_i \in \mathbb{N}_1^{K_i}, \forall i \in \mathbb{N}_1^N. \tag{B.6b}$$

In problem (B.6), the objective function of the MCP is minimized, while forcing x to lie inside the constrained sets for all samples $\delta^{(i, \kappa_i)}$ substituted into the corresponding constraint $i \in \mathbb{N}_1^N$. Clearly, the solution to problem (B.6) is itself a random variable, as it depends on the random multi-sample ω . For this reason, the scenario approach is a *randomized method* for finding an approximate solution to the MCP.

Of course, the MSP is actually solved for the observations of the random samples, leading to its deterministic instance ($\overline{\text{MSP}}[\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)}]$):

$$\min_{x \in \mathbb{X}} c^T x, \quad (\text{B.7a})$$

$$\text{s.t. } f_i(x, \bar{\delta}^{(i, \kappa_i)}) \leq 0 \quad \forall \kappa_i \in \mathbb{N}_1^{K_i}, \forall i \in \mathbb{N}_1^N. \quad (\text{B.7b})$$

Note that (B.7) arises from (B.6) by replacing the *(random) samples* $\delta^{(i, \kappa_i)}$, $\omega^{(i)}$, ω with their *(deterministic) outcomes* $\bar{\delta}^{(i, \kappa_i)}$, $\bar{\omega}^{(i)}$, $\bar{\omega}$. Throughout the paper, these outcomes are indicated by a bar, to distinguish them from the corresponding random variables. By Assumption (B.2), $\overline{\text{MSP}}$ constitutes a convex program that can be solved efficiently by a suitable algorithm for convex optimization, cf. [7, 16, 18].

Note that (B.6) remains important for analyzing the (probabilistic) properties of the (random) scenario solution. In fact, the subsequent theory is mainly concerned with showing that, with a very high confidence, the scenario solution is a feasible point of the chance constraints (B.4b), provided that the sample sizes K_1, \dots, K_N are appropriately selected.

2.3 Randomized Solution and Violation Probability

In order to avoid unnecessary complications, the following technical assumption ensures that there always exists a feasible solution to the MSP, cf. [11, p. 3].

ASSUMPTION B.5—FEASIBILITY (a) For any number of samples K_1, \dots, K_N , the MSP admits a feasible solution almost surely. (b) For the sake of notational simplicity, any Pr-null set for which (a) may not hold is assumed to be removed from Δ . ■

Assumption B.5 can be taken for granted in the majority of practical problems. When it does not hold in a particular case, a generalization of the presented theory accounting for the infeasible case can be developed along the lines of [10].

Hence the existence of a solution to $\overline{\text{MSP}}$ is ensured, and uniqueness holds by Assumption B.2 and by carry-over of the tie-break rule of Assumption B.3(b), see [20, Thm. 10.1, 7.1]. Therefore the *solution map*

$$\bar{x}^* : \Delta^K \rightarrow \mathbb{X} \quad (\text{B.8})$$

is well-defined, returning the unique optimal point $\bar{x}^*(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)})$ of the $\overline{\text{MSP}}$ for a given outcome of the multi-samples $\{\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)}\} \in \Delta^K$. The solution map can also be applied to the MSP, for which it is denoted by $x^* : \Delta^K \rightarrow \mathbb{X}$. Now $x^*(\omega^{(1)}, \dots, \omega^{(N)})$ represents a random vector of unknown probability distribution, which is also referred to as the *scenario solution*. In fact, its distribution is a complicated function of the geometry and the parameters of the problem.

Note that there are two levels of randomness present in the analysis. The first is introduced by the random samples in ω , which affect the choice of the scenario solution.

The second is the actual random uncertainty δ , which determines whether or not the scenario solution is feasible with respect to the chance constraints (B.6b). For this reason, the scenario approach presented here is also called a *double-level-of-probability approach* [9, Rem. 2.3].

To highlight the two probability levels more clearly, suppose first that the multi-sample $\bar{\omega}$ has already been observed, so that the scenario solution $\bar{x}^*(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)})$ is fixed. Then for each chance constraint $i = 1, \dots, N$ in (B.4b), the *a posteriori violation probability* $\bar{V}_i(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)})$ is given by

$$\bar{V}_i(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)}) := \mathbf{P}[f_i(\bar{x}^*(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)}), \delta) > 0] . \quad (\text{B.9})$$

In particular, each \bar{V}_i has a deterministic, yet generally unknown, value in $[0, 1]$. If the multi-sample ω has not yet been observed, the scenario solution $x^*(\omega^{(1)}, \dots, \omega^{(N)})$ is a random vector and so the *a priori violation probability*

$$V_i(\omega^{(1)}, \dots, \omega^{(N)}) := \mathbf{P}[f_i(x^*(\omega^{(1)}, \dots, \omega^{(N)}), \delta) > 0] \quad (\text{B.10})$$

becomes itself a random variable on (Δ^K, \mathbf{P}^K) , with support $[0, 1]$. Hence the goal is to choose appropriate sample sizes K_1, \dots, K_N which ensure that $V_i(\omega^{(1)}, \dots, \omega^{(N)}) \leq \varepsilon_i$ for all $i = 1, \dots, N$, with a sufficiently high confidence $(1 - \theta_i)$. Before these results are derived however, some structural properties of scenario programs and technical lemmas ought to be discussed.

3. Structural Properties of the Constraints

In this section, a structural property of a chance constraint is introduced which yields a reduction in the number of samples below the levels given by the existing theory [8, 10, 11]. This property relates to the new concept of the *support dimension* or, in a form that is more easily checked for many practical instances, the *support rank*.

3.1 Support Constraints

The concept of a *support constraint* carries over from the SCP case, cf. [8, Def. 4]. An illustration is given in Figure B.1.

DEFINITION B.1—SUPPORT CONSTRAINT Consider the $\overline{\text{MSP}}$ for some outcome of the multi-sample $\bar{\omega}$. (a) For some $i \in \mathbb{N}_1^N$ and $\kappa_i \in \mathbb{N}_1^{K_i}$, constraint $f_i(x, \bar{\delta}^{(i, \kappa_i)}) \leq 0$ is a *support constraint* of (B.7) if its removal from the problem entails a change in the optimal solution:

$$\bar{x}^*(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(N)}) \neq \bar{x}^*(\bar{\omega}^{(1)}, \dots, \bar{\omega}^{(i-1)}, \bar{\omega}^{(i)} \setminus \{\bar{\delta}^{(i, \kappa_i)}\}, \bar{\omega}^{(i+1)}, \dots, \bar{\omega}^{(N)}) .$$




FIGURE B.1. Illustration of Definition B.1 in \mathbb{R}^2 . The arrow indicates the optimization direction, the bold lines are the *support constraints* of the respective configuration.

In this case the sample $\bar{\delta}^{(i, \kappa_i)}$ is also said “to *generate* this support constraint.” (b) For each $i \in \mathbb{N}_1^N$, the indices κ_i of all samples that generate a support constraint of the $\overline{\text{MSP}}$ are included in the set $\overline{\text{Sc}}_i$. Moreover, the tuples (i, κ_i) of all support constraints of the $\overline{\text{MSP}}$ are collected in the *support (constraint) set* $\overline{\text{Sc}}$. With some abuse of this notation, $\overline{\text{Sc}} = \bigcup_{i=1}^N \overline{\text{Sc}}_i$. ■

Definition B.1(a) can be stated equivalently in terms of the objective function: a sampled constraint is a support constraint if and only if the optimal objective function value (or its preference by the tie-break rule) is strictly larger than when the constraint were removed. To be more precise, Definition B.1(b), $\overline{\text{Sc}}$ may also account for the set \mathbb{X} as an additional support constraint. This minor subtlety is tacitly understood in the sequel.

In the stochastic setting of the $\text{MSP}[\omega^{(1)}, \dots, \omega^{(N)}]$, whether or not a particular random sample $\delta^{(i, \kappa_i)}$ generates a support constraint becomes a random event, which can be associated with a certain probability. Similarly, the support constraint set Sc , and its subsets $\text{Sc}_1, \dots, \text{Sc}_N$ contributed by the various chance constraints, are naturally random sets.

3.2 Support Dimension

The link between the sample sizes K_1, \dots, K_N and the corresponding violation probability of the scenario solution depends decisively on the “dimensions” of the problem. The following lower bounds represent a mild technical condition, cf. [10, Thm. 3.3] and [11, Def. 2.3].

ASSUMPTION B.6 The sample sizes satisfy $K_1, \dots, K_N \geq d$. ■

In the existing literature, the dimension of the SCP has been characterized by *Helly’s dimension*, cf. [10, Def. 3.1]. In this paper, there is a subtle shift from the problem dimension to the dimension of chance constraint i in the MCP, embodied by its *support dimension*.

DEFINITION B.2—SUPPORT DIMENSION (a) Denote by $|\text{Sc}|$ the (random) cardinality of the set Sc . *Helly's dimension* is the smallest integer ζ that satisfies

$$\text{ess sup}_{\omega \in \Delta^K} |\text{Sc}| \leq \zeta .$$

(b) The *support dimension* of a chance constraint $i \in \mathbb{N}_1^N$ in the MSP is the smallest integer ζ_i that satisfies

$$\text{ess sup}_{\omega \in \Delta^K} |\text{Sc}_i| \leq \zeta_i .$$

■

From a basic argument using Helly's Theorem, the number of support constraints $|\text{Sc}|$ of any (feasible) convex optimization problem in \mathbb{R}^d is upper bounded by the dimension of the decision space d , cf. [8, Thm. 2]. This result implies that finite integers ζ and ζ_1, \dots, ζ_N matching Definition B.2 always exist, so that the concepts of “Helly's dimension” and “support dimension” are indeed well-defined. Moreover, the result provides immediate upper bounds on the support dimension of each chance constraint $i \in \mathbb{N}_1^N$ in (B.6), namely $\zeta_i \leq \zeta \leq d$.

It turns out that the support dimension ζ_i directly relates to the minimum sample size K_i that is required for a given violation level ε_i and residual probability θ_i . The basic mechanism shall be illustrated by the proposition below, for the simpler case of a *single-level of probability* problem, cf. [8, Thm. 1].

PROPOSITION B.1—PROBABILITY BOUND Consider a particular constraint $i \in \mathbb{N}_1^N$ in the MSP $[\omega^{(1)}, \dots, \omega^{(N)}]$ with some fixed sample size K_i , and let $\hat{\zeta}_i$ be an upper bound for its support dimension ζ_i . Then the following holds:

$$\mathbf{P}^{K+1}[f_i(x^*(\omega^{(1)}, \dots, \omega^{(N)}), \delta) > 0] \leq \frac{\hat{\zeta}_i}{K_i + 1} . \quad (\text{B.11})$$

■

Proof. Consider $\text{MSP}' := \text{MSP}[\omega^{(1)}, \dots, \omega^{(i-1)}, \omega^{(i)} \cup \{\delta\}, \omega^{(i+1)}, \dots, \omega^{(N)}]$ and let $\text{Sc}'_i \subset \{1, \dots, K_i, K_i + 1\}$ denote the set of support constraints generated by samples from $\omega^{(i)} \cup \{\delta\}$, where $(K_i + 1) \in \text{Sc}'_i$ stands for δ generating a support constraint. Note that the event where $f_i(x^*(\omega^{(1)}, \dots, \omega^{(N)}), \delta) > 0$ can be equivalently expressed as δ generating a support constraint of MSP' . Hence condition (B.11) can be reformulated as

$$\mathbf{P}^{K+1}[(K_i + 1) \in \text{Sc}'_i] \leq \frac{\hat{\zeta}_i}{K_i + 1} . \quad (\text{B.12})$$

To analyze the event $(K_i + 1) \in \text{Sc}'_i$, observe that by Assumption B.4 all samples in $\omega^{(i)} \cup \{\delta\}$ are i.i.d., whence all sampled instances of constraint i in (B.6b) along with

“ $f_i(\cdot, \delta) \leq 0$ ” are probabilistically identical. In particular, they are all equally likely to become a support constraint of MSP’. Hence if the number of support constraints $|\text{Sc}'_i|$ were known, then

$$\mathbf{P}^{K+1}[(K_i + 1) \in \text{Sc}'_i] = \frac{|\text{Sc}'_i|}{K_i + 1} .$$

Even though $|\text{Sc}'_i|$ is a random variable, by Definition B.2(b) $|\text{Sc}'_i| \leq \zeta_i$ almost surely, and by assumption $\zeta_i \leq \hat{\zeta}_i$. This immediately yields (B.11). \square

3.3 The Support Rank

In many practical cases, the support dimension ζ_i of a chance constraint $i \in \mathbb{N}_1^N$ in the MSP is not known exactly. Then it has to be replaced by some upper bound. As argued above, the existing upper bound is given by the dimension d of the decision space. However, this bound may not be tight in the case where the constraints satisfy a certain structural property, namely when they have a limited *support rank*.

Intuitively speaking, the support rank is the dimension d of the decision space less the maximal dimension of an (almost surely) *unconstrained subspace*. The latter is understood as a linear subspace of \mathbb{R}^d that cannot be constrained by the sampled instances of constraint i , for almost every value of the multi-sample $\omega^{(i)}$.

Before the support rank is introduced in a rigorous manner, three examples of constraint classes with bounded support rank are described, in order to equip the reader with the necessary intuition behind this concept. They also show that very common constraint classes possess this property, and that in practical problems it can often be spotted easily.

EXAMPLE B.2 For each of the following cases, a visual illustration can be found in Figure B.2.

(a) *Single Linear Constraint*. Suppose some chance constraint $i \in \mathbb{N}_1^N$ of (B.4b) takes the linear form

$$f_i(x, \delta) \equiv a^\text{T}x - b(\delta) , \tag{B.13}$$

where $a \in \mathbb{R}^d$, and $b : \Delta \rightarrow \mathbb{R}$ is a scalar depending on the uncertainty in a generic way. Note that these constraints in the MSP are unable to constrain any direction in the subspace orthogonal to the span of a , $\text{span}\{a\}^\perp$, regardless of the outcome of the multi-sample $\omega^{(i)}$. Hence the support rank α of the chance constraint (B.13) is equal to 1.

(b) *Multiple Linear Constraints*. As a generalization of case (a), suppose that some chance constraint $i \in \mathbb{N}_1^N$ of (B.4b) is given by

$$f_i(x, \delta) \equiv A(\delta)x - b(\delta) , \tag{B.14}$$

where $A : \Delta \rightarrow \mathbb{R}^{r \times d}$ and $b : \Delta \rightarrow \mathbb{R}^r$ represent a matrix and a vector that depend on the uncertainty δ . Moreover, suppose that the uncertainty enters the matrix $A(\delta)$ in such a

way that the dimension of the linear span of its rows $A_{j,\cdot}(\delta)$, for $j = 1, \dots, r$, satisfies

$$\dim \text{span}\{A_{j,\cdot}(\delta) \mid j \in \mathbb{N}_1^r, \delta \in \Delta\} \leq \beta < d .$$

Note that these constraints in the MSP are unable to constrain any direction in $\text{span}\{A_{j,\cdot}(\delta) \mid j \in \mathbb{N}_1^r, \delta \in \Delta\}^\perp$, regardless of the outcome of the multi-sample $\omega^{(i)}$. Hence the support rank of the chance constraint (B.14) is equal to β .


(c) *Quadratic Constraint.* For a nonlinear example, consider the case where some chance constraint $i \in \mathbb{N}_1^N$ of (B.4b) is given by

$$f_i(x, \delta) \equiv (x - x_c(\delta))^T Q (x - x_c(\delta)) - r(\delta) , \quad (\text{B.15})$$

where $Q \in \mathbb{R}^{d \times d}$ is positive semi-definite with $\text{rank } Q = \gamma < d$, and $x_c : \Delta \rightarrow \mathbb{R}^d$, $r : \Delta \rightarrow \mathbb{R}_+$ represent a vector and scalar that depend on the uncertainty. Note that these constraints in the MSP are unable to constrain any direction in the null space of the matrix Q , regardless of the outcome of the multi-sample $\omega^{(i)}$. Since this null space has dimension $d - \gamma$, the support rank of the chance constraint (B.15) is equal to γ . ■

To introduce the support rank in a rigorous manner, pick a chance constraint $i \in \mathbb{N}_1^N$ of the MCP. For each point $x \in \mathbb{X}$ and each uncertainty $\delta \in \Delta$, denote the corresponding level set of $f_i : \mathbb{R}^d \times \Delta \rightarrow \mathbb{R}$ by

$$F_i(x, \delta) := \{\xi \in \mathbb{R}^d \mid f_i(x + \xi, \delta) = f_i(x, \delta)\} . \quad (\text{B.16})$$



(a) Single Linear Constraint. (b) Multiple Linear Constraints. (c) Quadratic Constraint.

FIGURE B.2. Illustration of Example B.2 in \mathbb{R}^3 . The arrows indicate the dimension of the *unconstrained subspace*, equal to 3 minus the respective *support rank* α , β , or γ .

Let \mathcal{L} be the collection of all linear subspaces in \mathbb{R}^d . In order to be unconstrained, select only those subspaces that are contained in almost all level sets $F_i(x, \delta)$:

$$\mathcal{L}_i := \bigcap_{\delta \in \Delta} \bigcap_{x \in \mathbb{R}^d} \{L \in \mathcal{L} \mid L \subset F_i(x, \delta)\} . \quad (\text{B.17})$$

Introduce “ \preceq ” as the partial order on \mathcal{L}_i defined by set inclusion; i.e., for any two subspaces $L, L' \in \mathcal{L}_i$, $L \preceq L'$ if and only if $L \subseteq L'$. Then the following concepts are well-defined, as shown in Proposition B.2 below.

DEFINITION B.3—UNCONSTRAINED SUBSPACE, SUPPORT RANK (a) The *unconstrained subspace* L_i of chance constraint $i \in \mathbb{N}_1^N$ is the unique maximal element in \mathcal{L}_i , in the sense that $L \preceq L_i$ for all $L \in \mathcal{L}_i$. (b) The *support rank* $\rho_i \in \mathbb{N}_0^d$ of chance constraint $i \in \mathbb{N}_1^N$ equals to d minus the dimension of L_i ,

$$\rho_i := d - \dim L_i .$$

■

It is a minor technicality in Definition B.3 that any \mathbf{P} -null set that adversely influences the dimension of the unconstrained subspace can be removed from Δ ; this is tacitly understood.

Observe that if \mathcal{L}_i contains only the trivial subspace, then the support rank is actually equal to Helly’s dimension d . On the other hand, if \mathcal{L}_i contains more than the trivial subspace, then the support rank becomes strictly less than d .

PROPOSITION B.2—WELL-DEFINEDNESS OF UNCONSTRAINED SUBSPACE The collection \mathcal{L}_i contains a unique maximal element L_i in the set-inclusion sense, i.e., L_i contains all other elements of \mathcal{L}_i as subsets. ■

Proof. First, note that \mathcal{L}_i is always non-empty, because for every $x \in \mathbb{X}$ and every $\delta \in \Delta$ the level set $F_i(x, \delta)$ includes the origin by its definition in (B.16). Therefore \mathcal{L}_i contains (at least) the trivial subspace $\{0\}$.

Second, since every chain in \mathcal{L}_i has an upper bound (namely \mathbb{R}^d), *Zorn’s Lemma* (or the *Axiom of Choice*, cf. [6, p. 50]) implies that \mathcal{L}_i has at least one maximal element in the “ \preceq ”-sense.

Third, in order to prove that the maximal element is unique, suppose that $L_i^{(1)}, L_i^{(2)}$ are two maximal elements of \mathcal{L}_i . It will be shown that their direct sum $L_i^{(1)} \oplus L_i^{(2)} \in \mathcal{L}_i$, so that $L_i^{(1)} \neq L_i^{(2)}$ would contradict their maximality. According to (B.17), it must be shown that $L_i^{(1)} \oplus L_i^{(2)} \subset F_i(x, \delta)$ for any fixed values $x \in \mathbb{X}$ and $\delta \in \Delta$. To see this, pick

$$\xi \in L_i^{(1)} \oplus L_i^{(2)} \implies \xi = \xi^{(1)} + \xi^{(2)} \quad \text{for } \xi^{(1)} \in L_i^{(1)}, \xi^{(2)} \in L_i^{(2)} .$$

Then apply (B.16) twice to obtain

$$f_i(x + \xi^{(1)} + \xi^{(2)}, \delta) = f_i(x + \xi^{(1)}, \delta) = f_i(x, \delta) ,$$

because $\xi^{(2)} \in L_i^{(2)}$ and $\xi^{(1)} \in L_i^{(1)}$. □

3.4 The Support Rank Lemma

The following lemma provides the link between the support rank of a chance constraint and its support dimension.

LEMMA B.1—SUPPORT RANK Suppose that a chance constraint $i \in \mathbb{N}_1^N$ has the support rank $\rho_i \in \mathbb{N}_1^d$. Then its support dimension in the MSP is bounded by $\zeta_i \leq \rho_i$. ■

Proof. Without loss of generality, the proof is given for the first chance constraint $i = 1$. Pick any random multi-sample $\bar{\omega} \in \Delta^K$ (less any \Pr^K -null set for which the support rank condition may not hold).

By the assumption, there exists a linear subspace $L_1 \subset \mathbb{R}^d$ of dimension $d - \rho_1$ for which

$$f_1(x + \xi) = f_1(x) \quad \forall x \in \mathbb{X}, \quad \forall \xi \in L_1 .$$

The orthogonal complement of L_1 , L_1^\perp , is also a linear subspace of \mathbb{R}^d with dimension ρ_1 , and every vector in \mathbb{R}^d can be uniquely written as the orthogonal sum of vectors in L_1 and L_1^\perp , cf. [6, p.135].

For the sake of a contradiction, suppose that $i = 1$ contributes more than ρ_1 support constraints to the resulting $\overline{\text{MSP}}$, i.e., $|\overline{\text{Sc}}_1| \geq \rho_1 + 1$. For any $\kappa_1 \in \overline{\text{Sc}}_1$, let

$$\bar{x}_{\kappa_1}^* := \bar{x}^*(\bar{\omega}^{(1)} \setminus \{\bar{\delta}^{(1, \kappa_1)}\}, \bar{\omega}^{(2)}, \dots, \bar{\omega}^{(N)})$$

be the solution obtained if this support constraint is omitted. By Definition B.1, if a support constraint is omitted from $\overline{\text{MSP}}$, its solution moves away from \bar{x}_0^* , i.e., $\bar{x}_0^* \neq \bar{x}_{\kappa_1}^*$ for all $\kappa_1 \in \overline{\text{Sc}}_1$. Denote the collection of all solutions by

$$X := \{\bar{x}_{\kappa_1}^* \mid \kappa_1 \in \overline{\text{Sc}}_1\} \cup \{\bar{x}_0^*\} ,$$

so that $|X| \geq \rho_1 + 2$. Observe that each $\bar{x}_{\kappa_1}^*$ is feasible with respect to all constraints of the $\overline{\text{MSP}}$, except for the one generated by $\delta^{(1, \kappa_1)}$, which is necessarily violated according to Definition B.1.

Since \mathbb{R}^d is the orthogonal direct sum of L_1 and L_1^\perp , for each point in X there is a unique orthogonal decomposition of

$$\bar{x}_{\kappa_1}^* = v_{\kappa_1} + w_{\kappa_1} , \quad \text{where } v_{\kappa_1} \in L_1, \quad w_{\kappa_1} \in L_1^\perp ,$$

where $\kappa_1 \in \overline{\text{Sc}}_1 \cup \{0\}$. Consider the set

$$W := \{w_{\kappa_1} \mid \kappa_1 \in \overline{\text{Sc}}_1 \cup \{0\}\} .$$

By the hypothesis, W contains at least $\rho_1 + 2$ distinct points in the ρ_1 -dimensional subspace L_1^\perp . According to Radon's Theorem [23, p. 151], W can be split into two disjoint subsets, W_A and W_B , such that there exists a point \tilde{w} in the intersection of their convex hulls:

$$\tilde{w} \in \text{conv}\{W_A\} \cap \text{conv}\{W_B\} . \quad (\text{B.18})$$

Split the indices in $\overline{\text{Sc}}_1 \cup \{0\}$ correspondingly into I_A and I_B , and observe that every $w_A \in W_A$ satisfies the constraints in I_B :

$$f_1(w_A, \bar{\delta}^{(1, \kappa_1)}) \leq 0 \quad \forall \kappa_1 \in I_B \quad \implies \quad f_1(\tilde{w}, \bar{\delta}^{(1, \kappa_1)}) \leq 0 \quad \forall \kappa_1 \in I_B .$$

The last implication follows because $\tilde{w} \in \text{conv}\{W_A\}$ and $f_1(\cdot, \bar{\delta}^{(1, \kappa_1)})$ is convex. Similarly, every point $w_B \in W_B$ satisfies the constraints in I_A :

$$f_1(w_B, \bar{\delta}^{(1, \kappa_1)}) \leq 0 \quad \forall \kappa_1 \in I_A \quad \implies \quad f_1(\tilde{w}, \bar{\delta}^{(1, \kappa_1)}) \leq 0 \quad \forall \kappa_1 \in I_A .$$

Combining both statements thus yields

$$f_1(\tilde{w}, \bar{\delta}^{(1, \kappa_1)}) \leq 0 \quad \forall \kappa_1 \in \overline{\text{Sc}}_1 . \quad (\text{B.19})$$

According to (B.18), \tilde{w} can be expressed as a convex combination of elements in W_A or W_B . Splitting the points in X into X_A and X_B correspondingly and applying the same convex combination yields some

$$\tilde{x} \in \text{conv}\{X_A\} \cap \text{conv}\{X_B\} , \quad (\text{B.20})$$

and thereby also some $\tilde{v} \in L_1$ with $\tilde{x} = \tilde{v} + \tilde{w}$.

To establish the contradiction two things remain to be verified: first that \tilde{x} is feasible with respect to all constraints, and second that it has a lower cost (or a better tie-break value) than \bar{x}_0^* . For the first, $\tilde{x} \in \mathbb{X}$ because all points of X lie in \mathbb{X} and $\tilde{x} \in \text{conv}\{X\}$. Moreover, thanks to (B.19),

$$f_1(\tilde{x}, \bar{\delta}^{(1, \kappa_1)}) = f_1(\tilde{w}, \bar{\delta}^{(1, \kappa_1)}) \leq 0 \quad \forall \kappa_1 \in \overline{\text{Sc}}_1 .$$

For the second, pick the set from X_A and X_B that does not contain \bar{x}_0^* ; without loss of generality, say this is X_A . By construction, all elements of X_A have a strictly lower objective function value (or at least a better tie-break value) than \bar{x}_0^* . By linearity this also holds for all points in $\text{conv}\{X_A\}$, where \tilde{x} lies according to (B.20). \square

REMARK B.2—SUPPORT RANK VERSUS SUPPORT DIMENSION While the support rank ρ_i is a property of chance constraint i alone, the support dimension ζ_i may depend on the overall setup of the MSP. The support dimension ζ_i constitutes the relevant basis for selecting the sample size K_i . However, it may be difficult to determine for practical problems, as it may depend on the interactions of multiple chance constraints (see Example B.3 below). The support rank ρ_i provides an easier-to-handle upper bound to ζ_i , which can be used in place of ζ_i for selecting K_i . ■

EXAMPLE B.3—UPPER BOUNDING OF SUPPORT DIMENSION To illustrate the statements in Remark B.2, consider a small example of (B.4) in dimension $d = 3$. Let $\mathbb{X} = [-1, 1]^3$ be the unit cube, $c^T = [0 \ 1 \ 1]$ with a lexicographic tie-break rule, and two chance constraints $i = 1, 2$. Both constraints affect only the first and second coordinates x_1 and x_2 , leaving the choice of $x_3 = -1$ for the third coordinate. For $i = 1$, the constraints are parallel hyperplanes constraining x_1 from below, where the lower bound is given by the first uncertainty δ_1 :

$$f_1(x, \delta) = -x_1 + \delta_1 \ .$$

For $i = 2$, the constraints are V-shaped, with the vertex located at $x_1 = -\delta_2$ and $x_2 = -1$:

$$f_2(x, \delta) = |x_1 + \delta_2| - x_2 - 1 \ .$$

Both uncertainties $\delta := \{\delta_1, \delta_2\}$ are uniformly distributed on the interval $[0, 1]$. The setup is illustrated in Figure B.3.

In this case, the support dimensions are $\zeta_1 = 1$, $\zeta_2 = 1$ and the support ranks are $\rho_1 = 1$, $\rho_2 = 2$ for the constraints $i = 1, 2$. Notice that for $i = 2$ the support rank is strictly greater than its support dimension, due to the presence of constraint 1. Hence there is some conservatism in the upper bound, although both bounds are better than the existing upper bound by the dimension of the decision space $d = 3$ [8, Thm. 2]. ■

4. Feasibility of the Scenario Solution

In the first part of this section, it is shown that for a proper choice of the sample sizes K_1, \dots, K_N the scenario solution $x^*(\omega^{(1)}, \dots, \omega^{(N)})$ is an approximate solution of the MCP (i.e., it is a feasible point of each chance constraint $i = 1, \dots, N$ in (B.4b) with a high confidence $(1 - \theta_i)$). In the second part of this section, an explicit formula for computing the sample sizes K_1, \dots, K_N for given residual probabilities θ_i is provided.




FIGURE B.3. Illustration of Example B.3. The plot shows a projection on the x_1, x_2 -plane for $x_3 = -1$. The unit box \mathbb{X} is depicted by a dotted line. Two (possible) samples are shown for the linear constraint $i = 1$ ($x_1 \geq \delta_1$) and for the V-shaped constraint $i = 2$ ($x_2 \geq |x_1 + \delta_2| - 1$).

4.1 The Sampling Theorem

Denote by $B(\cdot; \cdot, \cdot)$ the beta distribution function, cf. [1, p. 26.5.3, 26.5.7]:

$$B(\varepsilon; n, K) := \sum_{j=0}^n \binom{K}{j} \varepsilon^j (1 - \varepsilon)^{K-j} . \quad (\text{B.21})$$

THEOREM B.1—SAMPLING THEOREM Consider problem (B.6) under Assumptions B.2, B.3, B.4, B.5, B.6. Then

$$\mathbf{P}^K [V_i(\omega^{(1)}, \dots, \omega^{(N)}) > \varepsilon_i] \leq B(\varepsilon_i; \rho_i - 1, K_i) , \quad (\text{B.22})$$

for each chance constraint $i \in \mathbb{N}_1^N$, whose support rank is ρ_i . ■

Proof. The result is an extension of [11, Thm.2.4] for the classic scenario approach, which is also used as a basis for this proof.¹⁶

Without loss of generality, consider the first chance constraint $i = 1$; the result for the other chance constraints $i = 2, \dots, N$ follows analogously. Consider the conditional probability

$$\mathbf{P}^K [V_1(\omega^{(1)}, \dots, \omega^{(N)}) > \varepsilon_1 \mid \omega^{(2)}, \dots, \omega^{(N)}] , \quad (\text{B.23})$$

i.e., the probability of drawing $\omega^{(1)}$ such that $x^*(\omega^{(1)}, \dots, \omega^{(N)})$ has a probability of violating “ $f_1(\cdot, \delta) \leq 0$ ” that is higher than ε_1 , given fixed values for the other samples $\omega^{(2)}, \dots, \omega^{(N)}$.

¹⁶The authors thank an anonymous reviewer for his/her helpful suggestions on simplifying the proof.

Clearly, the quantity in (B.23) generally depends on the multi-samples $\omega^{(2)}, \dots, \omega^{(N)}$. However, for $\mathbf{P}^{K_2+\dots+K_N}$ -almost every value of these multi-samples (B.23) can be bounded by

$$\mathbf{P}^K[V_1(\omega^{(1)}, \dots, \omega^{(N)}) > \varepsilon_1 \mid \omega^{(2)}, \dots, \omega^{(N)}] \leq B(\varepsilon_1; \rho_1 - 1, K_1) . \quad (\text{B.24})$$

Indeed, by Assumption B.2, for $\mathbf{P}^{K_2+\dots+K_N}$ -almost every $\omega^{(2)}, \dots, \omega^{(N)}$ the function $\tilde{f} : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by

$$\tilde{f}(x) \equiv \max_{i \in \mathbb{N}_2^N} \max_{\kappa_i \in \mathbb{N}_1^{K_i}} f_i(x, \delta^{(i, \kappa_i)})$$

is convex, as it is the point-wise maximum of convex functions. Then all sampled constraints of $i = 2, \dots, N$ can be expressed as the deterministic convex constraint “ $\tilde{f}(x) \leq 0$ ”, which can be considered as part of the convex set \mathbb{X} . Thus for $\mathbf{P}^{K_2+\dots+K_N}$ -almost every $\omega^{(2)}, \dots, \omega^{(N)}$ the problem takes the form of a classic SCP, to which the results of [11] apply. In particular, [11, Thm. 2.4] yields (B.24) for $\mathbf{P}^{K_2+\dots+K_N}$ -almost every $\omega^{(2)}, \dots, \omega^{(N)}$.

The difference from using the support rank ρ_1 in place of the optimization dimension d in [11, Thm. 2.4] is minor. The key fact is that ρ_1 provides an upper bound for the number of support constraints contributed by constraint 1, according to Lemma B.1, and hence it can replace d in [11, Prop. 2.2] and all subsequent results.

The final result is obtained by deconditioning the probability in (B.23):

$$\begin{aligned} \mathbf{P}^K[V_1(\omega^{(1)}, \dots, \omega^{(N)}) > \varepsilon_1] &= \\ &= \int_{\omega^{(2)}, \dots, \omega^{(N)}} \mathbf{P}^K[V_1(\omega^{(1)}, \dots, \omega^{(N)}) > \varepsilon_1 \mid \omega^{(2)}, \dots, \omega^{(N)}] \mathbf{P}^{K_2}[\mathrm{d} \omega^{(2)}] \dots \mathbf{P}^{K_N}[\mathrm{d} \omega^{(N)}] \\ &\leq \int_{\omega^{(2)}, \dots, \omega^{(N)}} \Phi(\varepsilon_1; \rho_1 - 1, K_1) \mathbf{P}^{K_2}[\mathrm{d} \omega^{(2)}] \dots \mathbf{P}^{K_N}[\mathrm{d} \omega^{(N)}] \\ &= \Phi(\varepsilon_1; \rho_1 - 1, K_1) , \end{aligned}$$

based on [22, pp. 183, 222], where the third line uses (B.24). \square

4.2 Explicit Bounds on the Sample Sizes

Formula (B.22) in Theorem B.1 ensures that with a *confidence level* of $1 - B(\varepsilon_i; \rho_i - 1, K_i)$, the violation probability $V_i(\omega^{(1)}, \dots, \omega^{(N)}) \leq \varepsilon_i$. However, in practical applications a given confidence level $(1 - \theta_i) \in (0, 1)$ is often imposed, while an appropriate sample size K_i has to be identified.

The most accurate way of finding this sample size is by observing that $B(\varepsilon_i; \rho_i - 1, K_i)$ is a monotonically decreasing function in K_i and applying a numerical procedure (e.g., *regula falsi*) for computing the smallest sample size that ensures $B(\varepsilon_i; \rho_i - 1, K_i) \leq \theta_i$. The resulting K_i shall be referred to as the *implicit bound* on the sample size.

For a qualitative analysis of the behavior of this implicit bound as ε_i and θ_i vary (and also for a good initialization of the *regula falsi* procedure), it is useful to derive an *explicit bound* on the sample size K_i . Since formula (B.22) cannot be readily inverted,

the beta distribution function must first be controlled by some upper bound, which is then inverted.

A straightforward approach is to use a Chernoff bound [13], as shown in [9, Rem. 2.3] and [10, Sec. 5]. This provides a simple explicit formula for K_i :

$$K_i \geq \frac{2}{\varepsilon_i} \left[\log\left(\frac{1}{\theta_i}\right) + \rho_i - 1 \right] , \quad (\text{B.26})$$

where $\log(\cdot)$ denotes the natural logarithm. As shown in [2, Cor. 1], this can be further improved to a better, albeit more complicated bound for K_i :

$$K_i \geq \frac{1}{\varepsilon_i} \left[\log\left(\frac{1}{\theta_i}\right) + \sqrt{2(\rho_i - 1) \log\left(\frac{1}{\theta_i}\right)} + \rho_i - 1 \right] . \quad (\text{B.27})$$

5. The Sampling-and-Discarding Approach

The sampling-and-discarding approach has previously been proposed for the classic scenario approach [10, 12]; this section describes its extension to problems with multiple chance constraints.

The fundamental goal is to reduce the objective value of the scenario solution, while maintaining the same confidence levels for feasibility with respect to the chance constraints (see Section 1.2). To this end, the sample sizes K_i are deliberately increased above the bounds derived in Section 4, in exchange for allowing a certain number of R_i sampled constraints to be discarded *a posteriori*, i.e., after the outcomes of the samples have been observed.

In this section, first the possible procedures for discarding constraints are recalled. Second, the main result on the sampling-and-discarding approach for the MCP is stated. It provides an implicit formula for the selection of appropriate sample-and-discarding pairs (K_i, R_i) , which may again vary for different chance constraints $i = 1, \dots, N$. Third, explicit bounds for the choice of pairs (K_i, R_i) are provided.

5.1 Constraint Discarding Procedure

For each chance constraint of the MCP, if $R_i \geq 0$ sampled constraints are to be discarded a posteriori, the discarding procedure is performed by a pre-defined *(sample) removal algorithm*.

DEFINITION B.4—REMOVAL ALGORITHM For each chance constraint $i = 1, \dots, N$, the *(sample) removal algorithm* $\mathcal{A}_i^{(K_i, R_i)} : \Delta^K \rightarrow \Delta^{K_i - R_i}$ is a deterministic function on the overall multi-sample $\omega \in \Delta^K$. It returns a subset of samples $\tilde{\omega}^{(i)} \in \Delta^{K_i - R_i}$, in which R_i out of the K_i samples in $\omega^{(i)} \in \Delta^{K_i}$ have been removed. ■

Obviously, the algorithm should aim at improving the objective value from $\text{MSP}[\omega^{(1)}, \dots, \omega^{(N)}]$ to $\text{MSP}[\tilde{\omega}^{(1)}, \dots, \tilde{\omega}^{(N)}]$ as much as possible. Various possible removal algorithms are described in [10, Sec. 5.1], and further references are found in [12, Sec. 2]. Brief descriptions of the most important removal algorithms are listed below.

EXAMPLE B.4 (a) *Optimal Constraint Removal.* The best improvement of the objective function value is achieved by solving the reduced problem for all possible ways of removing R_i of the K_i samples. However, a major drawback of this removal algorithm is its combinatorial complexity. Therefore the algorithm becomes computationally intractable for larger values of R_i , in particular when samples have to be removed for multiple constraints.

(b) *Greedy Constraint Removal.* Starting by solving the $\text{MSP}[\omega^{(1)}, \dots, \omega^{(N)}]$ for all K_i samples, the R_i samples are removed in R_i sequentially steps. In each step, a single sample is removed by the optimal constraint removal procedure. Between multiple constraints i , the removal algorithm can either proceed in a fixed order or again greedy-based. For most practical problems this algorithm can be expected to work almost as good as (a), while carrying a much lower computational burden.

(c) *Marginal Constraint Removal.* The R_i samples are removed in R_i sequential steps, where the removed sample in each step is selected according to the highest Lagrange multiplier. Compared to the greedy constraint removal, the decision is thus based on the highest marginal cost improvement [7, Cha. 5]), instead of the highest total cost improvement. In the case of multiple constraints i , the removal algorithm can either handle them all together, or proceed sequentially. ■

The existing theory for the SCP [10, Sec. 4.1.1] and [12, Ass. 2.2] assumes that all of the removed constraints are violated by the relaxed scenario solution.

ASSUMPTION B.7—VIOLATION OF DISCARDED CONSTRAINTS Every chance constraint $i \in \mathbb{N}_1^N$ with $R_i > 0$ satisfies the following condition: for almost every $\omega \in \Delta^K$, each of the constraints discarded by the removal algorithm $\mathcal{A}_i^{(K_i, R_i)}(\omega)$ is violated by the solution of the reduced problem, i.e.,

$$f_i(x^*(\tilde{\omega}^{(1)}, \dots, \tilde{\omega}^{(N)}), \delta^{(i, \kappa_i)}) > 0 \quad \forall \delta^{(i, \kappa_i)} \in (\omega \setminus \tilde{\omega}) . \quad (\text{B.28})$$

■

While Assumption B.7 is sufficient for the MCP as well, it may turn out to be too restrictive for some problem instances. In fact, due to the interplay of multiple chance constraints, it may not be possible to find R_i constraints that are violated by the relaxed scenario solution (this situation may also occur for a single chance constraint, in the presence of a deterministic constraint set \mathbb{X}). In this case, the *monotonicity property*, as introduced below, provides a possible alternative.

DEFINITION B.5—MONOTONICITY PROPERTY A chance constraint $i \in \mathbb{N}_1^N$ is called *monotonic* if for all $K_i \in \mathbb{N}$ and almost every $\omega^{(i)} \in \Delta^{K_i}$ the following condition holds: Every point in the feasible set of sampled instances of chance constraint i ,

$$\mathbb{X}_i(\omega^{(i)}) := \{\xi \in \overline{\mathbb{R}}^d \mid f_i(\xi, \delta^{(i, \kappa_i)}) \leq 0 \quad \forall \kappa_i \in \mathbb{N}_1^{K_i}\} , \quad (\text{B.29})$$

where $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$, is violated by a new sampled constraint only if also the optimal point in $\mathbb{X}_i(\omega^{(i)})$,

$$x_i^*(\omega^{(i)}) := \arg \min \{c^T \xi \mid \xi \in \mathbb{X}_i(\omega^{(i)})\} \quad (\text{B.30})$$

is violated. In other words, for every $\xi \in \mathbb{X}_i(\omega^{(i)})$ and almost every $\delta \in \Delta$,

$$f_i(\xi, \delta) > 0 \quad \implies \quad f_i(x_i^*(\omega^{(i)}), \delta) > 0 . \quad (\text{B.31})$$

■

ASSUMPTION B.8—MONOTONICITY OF CHANCE CONSTRAINTS Every chance constraint $i \in \mathbb{N}_1^N$ enjoys the *monotonicity property*. ■

Definition B.5 is easy to check for most practical problems, without involving any calculations. The following example illustrates the intuition behind this concept.

EXAMPLE B.5—MONOTONIC CHANCE CONSTRAINTS Consider an MSP in $d = 2$ dimensions, where $\mathbb{X} = [-100, 100]^2 \subset \mathbb{R}^2$ and $c = [0 \ 1]^T$, $\delta = [\delta_1 \ \delta_2 \ \delta_3]$ belongs to $\Delta = \{-1, 1\} \times [-1, 1] \times [-1, 1]$, and there are $N = 2$ chance constraints.

(a) *Monotonic Chance Constraint.* Let the first chance constraint $i = 1$ be of the linear form

$$[\delta_1^{(1, \kappa_1)} \quad 1]x - \delta_2^{(1, \kappa_1)} \leq 0 \quad \forall \kappa_1 = 1, \dots, K_1 .$$

Observe that for any number $K_1 \in \mathbb{N}$ and every possible sample values $\omega^{(1)}$, an additional sample δ either cuts off no point from $\mathbb{X}_1(\omega^{(1)})$, or the the point $x_1^*(\omega^{(1)})$ becomes infeasible. This fact is illustrated in Figure B.4(a). Therefore chance constraint $i = 1$ enjoys the monotonicity property.

(b) *Non-Monotonic Chance Constraint.* Let the second chance constraint $i = 2$ be of the linear form

$$[\delta_2^{(2, \kappa_2)} \quad 1]x - \delta_3^{(2, \kappa_2)} \leq 0 \quad \forall \kappa_2 = 1, \dots, K_2 .$$

Observe that for any number K_2 there exist sample values $\omega^{(2)}$ that make it possible for a new sample δ to cut off some previously feasible point from $\mathbb{X}_2(\omega^{(2)})$, without rendering the point $x_2^*(\omega^{(2)})$ infeasible. A possible configuration of this type is depicted in Figure B.4(b). Therefore chance constraint $i = 2$ does not enjoy the monotonicity property. ■




FIGURE B.4. Illustration of Example B.5. Non-bold constraints are generated by the multi-sample $\omega^{(i)} \in \Delta^{K_i}$ of chance constraint $i = 1, 2$; bold constraints are generated by the uncertainty $\delta \in \Delta$. In (b) a feasible point is made infeasible without affecting the optimum, which is not possible in the case of (a).

The usefulness of the monotonicity property is based on the following result, whose proof is an straightforward consequence of Definition B.5 and therefore omitted.

LEMMA B.2 Let $K_i \in \mathbb{N}$ and $R_i \leq K_i$. Suppose chance constraint $i \in \mathbb{N}_1^N$ of MCP is monotonic and the removal algorithm $\mathcal{A}_i^{(K_i, R_i)}$ is sequential. Then for almost every $\omega^{(i)} \in \Delta^{K_i}$ the following holds:

(a) With probability one every point ξ in the set $\mathbb{X}_i(\omega^{(i)})$ has a violation probability less than or equal to that of the cost-minimal point $x_i^*(\omega^{(i)})$:

$$\mathbf{P}[f_i(\xi, \delta) > 0] \leq \mathbf{P}[f_i(x_i^*(\omega^{(i)}), \delta) > 0] \quad \forall \xi \in \mathbb{X}_i(\omega^{(i)}) . \quad (\text{B.32})$$

(b) The solution $x_i^*(\tilde{\omega}^{(i)})$, with $\tilde{\omega}^{(i)} = \mathcal{A}_i^{(K_i, R_i)}(\omega_i)$, violates all R_i removed constraints. ■

5.2 The Discarding Theorem

For the sampling-and-discarding approach, the following result holds for the MCP.

THEOREM B.2—DISCARDING THEOREM Consider the problem (B.4) under Assumptions B.2, B.3, B.4, B.5, B.6, and either B.7 or B.8. Let $\mathcal{A}_i^{(K_i, R_i)}$ be sample removal algorithms for each of its chance constraints $i = 1, \dots, N$, some of which may be trivial (i.e., $R_i = 0$). Then it holds that

$$\mathbf{P}^K[V_i(\tilde{\omega}^{(1)}, \dots, \tilde{\omega}^{(N)}) > \varepsilon_i] \leq \binom{R_i + \rho_i - 1}{R_i} B(\varepsilon_i; R_i + \rho_i - 1, K_i) \quad (\text{B.33})$$

ρ_i is the support rank of chance constraint i and $B(\cdot; \cdot, \cdot)$ the beta distribution (B.21). ■

Proof. Here the MCP case is reduced to the SCP case, for which a detailed proof is available in [12, Sec. 5.1].

First, suppose that Assumption B.7 holds. The proof in [12, Sec. 5.1] works analogously for an arbitrary chance constraint $i \in \mathbb{N}_1^N$, given that an upper bound of the violation distribution is readily available from Theorem B.1.

Second, suppose that Assumption B.8 holds. In this case the proof in [12, Sec. 5.1] can be applied directly to the SCP which arises from the MCP if all chance constraints other than a particular $i \in \mathbb{N}_1^N$ are omitted (and also \mathbb{X} is omitted). In particular, (B.33) holds for the scenario solution of this SCP, using Lemma B.2(b). Given that the chance constraint is monotonic and by virtue of Lemma B.2(a), (B.33) also holds for any point in $\mathbb{X}_i(\omega^{(i)})$, in particular for the scenario solution of the MCP. \square

The work of [12] already provides an excellent account of the merits of the sampling-and-discarding approach, which does not require a restatement here. However, it should be emphasized that the scenario solution converges to the true solution of the MCP as the number of discarded constraints increases, provided that the constraints are removed by the optimal procedure of Example B.4(a).

5.3 Explicit Bounds on the Sample-and-Discarding Pairs

Similar to Section 4, explicit bounds on the sample size K_i can also be derived for the sampling-and-discarding approach, assuming the number of discarded constraints R_i to be fixed. The technical details, using Chernoff bounds [13], are worked out in [10, Sec. 5]. The resulting explicit bound is indicated here for the sake of completeness,

$$K_i \geq \frac{2}{\varepsilon_i} \log\left(\frac{1}{\theta_i}\right) + \frac{4}{\varepsilon_i} (R_i + \rho_i - 1) , \quad (\text{B.34})$$

where $\log(\cdot)$ denotes the natural logarithm.

Similarly, explicit bounds on the number of discarded constraints R_i can be obtained, assuming the sample size K_i to be fixed:

$$R_i \leq \varepsilon_i K_i - \rho_i + 1 - \sqrt{2\varepsilon_i K_i \log\left(\frac{(\varepsilon_i K_i)^{\rho_i-1}}{\theta_i}\right)} . \quad (\text{B.35})$$

The technical details of this are found in [12, Sec. 4.3].

6. Example: Minimal Diameter Cuboid

The following academic example has been selected to highlight the strengths of the extensions to the scenario approach presented in this paper.

6.1 Problem Statement

Let δ be a random point in $\Delta \subset \mathbb{R}^n$, whose distribution and support set are unknown, but sampled values can be obtained. The objective in this example is to construct the Cartesian product C of closed intervals in \mathbb{R}^n (“ n -cuboid”) of minimal n -diameter W , which is large enough to contain the point δ in its i -th coordinate with probability $(1 - \varepsilon_i)$. The setting is illustrated in Figure B.5.

Let $z \in \mathbb{R}^n$ denote the center point of the cuboid and $t \in \mathbb{R}_+^n$ the interval widths in each dimension, so that

$$C = \{\xi \in \mathbb{R}^n \mid |\xi_i - z_i| \leq t_i/2\} . \quad (\text{B.36})$$

Then the corresponding stochastic program reads as follows:

$$\min_{z \in \mathbb{R}^n, t \in \mathbb{R}_+^n} \|t\|_2 , \quad (\text{B.37a})$$

$$\text{s.t.} \quad \Pr[z_i - t_i/2 \leq \delta_i \leq z_i + t_i/2] \geq (1 - \varepsilon_i) \quad \forall i \in \mathbb{N}_1^n . \quad (\text{B.37b})$$

Since the objective function is not linear, (B.37) has to be reformulated (see Remark B.1(a)) as

$$\min_{z \in \mathbb{R}^n, t \in \mathbb{R}_+^n, T \in \mathbb{R}} T , \quad (\text{B.38a})$$

$$\text{s.t.} \quad \|t\|_2 \leq T , \quad (\text{B.38b})$$

$$\Pr\left[\max\{z_i - t_i/2 - \delta_i, -z_i - t_i/2 + \delta_i\} \leq 0\right] \geq (1 - \varepsilon_i) \quad \forall i \in \mathbb{N}_1^n . \quad (\text{B.38c})$$

Note that (B.38) takes the form of a MCP, for a $d = 2n + 1$ dimensional search space and $N = n$ chance constraints: the objective function (B.38a) is linear; constraint (B.38b) is deterministic and convex; and each of the chance constraints in (B.38c) is convex in z, t for any fixed value of the uncertainty $\delta \in \Delta$.

Here each of the chance constraints $i = 1, \dots, n$ depends on exactly two decision variables z_i and t_i , which is a special case of involving $[z; t; T] \in \mathbb{R}^{2n+1}$ (see Remark B.1(c)). The convex and compact set \mathbb{X} is constructed from the positivity constraints on t , the deterministic and convex constraint (B.38b), and some artificial bounds assumed on all variables. Existence of a feasible solution, and hence Assumption B.3, holds automatically from the problem setup.

6.2 Solution via Scenario Approach

By inspection, each of the chance constraints $i = 1, \dots, n$ has support rank $\rho_i = 2$, because it only involves the two variables z_i and t_i . For a fixed confidence level, e.g., $\theta = 10^{-6}$, the implicit sample sizes K_1, \dots, K_n in (B.22) can be computed for given values of n and




FIGURE B.5. Illustration of the numerical example for $n = 2$. The point $\delta \in \Delta$ appears at random in \mathbb{R}^2 , according to some unknown distribution; the points drawn here are 166 i.i.d. samples of δ . The objective is to construct the smallest product of two closed intervals (“2-cuboid”), drawn here as the shaded rectangle, such that the probability of failing to contain the realization of δ is smaller than ε_1 and ε_2 in dimension 1 and 2, respectively.

$\varepsilon_1, \dots, \varepsilon_n \in (0, 1)$ by a bisection-based algorithm (see Section 4.2). For simplicity, all $\varepsilon_1 = \dots = \varepsilon_n$ are selected as equal, and since $\rho_1 = \dots = \rho_N = 2$, the implicit sample sizes $K_1 = \dots = K_n$ are also identical.

Given the outcomes of all multi-samples, the $\overline{\text{MSP}}$ is easily solved by the smallest n -cuboid that contains all sampled points; see also Figure B.5. In other words, here the $\overline{\text{MSP}}$ has an analytic solution.

Table B.1(a) summarizes the implicit sample sizes required for guaranteeing various chance constraint levels ε_i in various dimensions n (all with $\theta = 10^{-6}$). These sample sizes are also compared to those from the classic scenario approach, based on a reformulation of (B.38) as an SCP according to the procedure outlined in Section 2.1.

Observe from Table B.1 that the SCP-based sample sizes are always larger than those using the extensions of the MCP theory. This effect increases, in particular, as the dimension n of the optimization space grows larger. The reason is that the support dimension of each chance constraint remains constant for all n , whereas Helly’s dimension grows as it equals to n . The marginal growth of the sample size of the MCP, despite the support rank $\rho_i = 2$ being constant, is the result of adjusting the confidence level θ to be (evenly) distributed among the chance constraints, i.e., $\theta_i = \theta/n$ for all $i = 1, \dots, n$.

sample size K_i	cuboid dimension $n =$						
	2	3	5	10	50	100	500
$\varepsilon_i =$	1%	1,734	1,777	1,831	1,903	2,072	2,144
	5%	341	349	360	374	407	421
	10%	166	170	176	182	199	205
	25%	62	63	65	67	73	76

(a) MCP-based Scenario Approach.

sample size K_i	cuboid dimension $n =$						
	2	3	5	10	50	100	500
$\varepsilon_i =$	1%	2,334	2,722	3,431	5,020	15,588	27,535
	5%	459	536	677	992	3,095	5,477
	10%	225	263	332	488	1,533	2,719
	25%	84	99	125	186	595	1,063

(b) SCP-based Scenario Approach.

TABLE B.1. Implicit sample sizes $K_1 = \dots = K_n$ for the MCP-based and the SCP-based scenario approach, assuming a confidence level of $\theta = 10^{-6}$, for varying problem dimension n and chance constraint levels $\varepsilon_1 = \dots = \varepsilon_n$.

The larger sample size of the SCP-based approach, as compared to the MCP-based approach, implies higher data requirements and higher computational efforts, but it also increases the conservatism of the scenario solution. The latter effect is quantified in Table B.2, showing the relative excess of the (average) objective function values of the SCP-based solutions over those of the MCP-based solutions. Note that the objective values achieved by the SCP-based approach are always higher than those achieved by the MCP-based approach, with the effect becoming increasingly significant as the dimension n of the decision space grows larger.

relative obj. value	cuboid dimension $n =$						
	2	3	5	10	50	100	500
$\varepsilon_i =$	1%	2.4%	3.4%	5.0%	7.5%	14.8%	18.4%
	5%	3.3%	4.6%	6.6%	9.8%	19.3%	23.8%
	10%	3.9%	5.4%	7.6%	11.5%	22.2%	27.4%
	25%	5.0%	7.2%	10.1%	15.1%	28.5%	34.7%

TABLE B.2. Objective function value of SCP-based scenario solution as a percentage increase over the MCP-based scenario solution, based on the sample sizes in Table B.1 and a multivariate standard normal distribution for δ . Each of the indicated values represents an average over one million simulation runs.

Acknowledgments

The authors would like to thank two anonymous reviewers for their helpful and productive comments, and Joe Warrington for his carefully proofreading of the manuscript.

The research of L. Fagiano has received funding from the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement number PIOF-GA-2009-252284, Marie Curie project “Innovative Control, Identification and Estimation Methodologies for Sustainable Energy Technologies”.

References

- [1] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions*. Dover Publications, New York, 9th edition, 1970.
- [2] T. Alamo, R. Tempo, and A. Luque. On the sample complexity of probabilistic analysis and design methods. In J.C Willems et al., editor, *Perspectives in Mathematical System Theory, Control, and Signal Processing*, pages 39–50. Springer, Berlin et al., 2010.
- [3] D. Bai, T. Carpenter, and J. Mulvey. Making a case for robust optimization models. *Management Science*, 43(7):895–907, 1997.
- [4] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.
- [5] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, New York, 1997.
- [6] B. Bollobás. *Linear Analysis*. Cambridge University Press, Cambridge et al., 2nd edition, 1999.
- [7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.
- [8] G. Calafiore and M.C. Campi. Uncertain convex programs: Randomized solutions and confidence levels. *Mathematical Programming, Series A*, 102-1:25–46, 2005.
- [9] G.C. Calafiore. On the expected probability of constraint violation in sampled convex programs. *Journal of Optimization Theory and Applications*, 143:405–412, 2009.
- [10] G.C. Calafiore. Random convex programs. *SIAM Journal on Optimization*, 20(6):3427–3464, 2010.
- [11] M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19:1211–1230, 2008.

- [12] M.C. Campi and S. Garatti. A sampling and discarding approach to chance-constrained optimization: Feasibility and optimality. *Journal of Optimization Theory and Applications*, 148:257–280, 2011.
- [13] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952.
- [14] P. Kall and J. Mayer. *Stochastic Linear Programming*. Springer, New York et al., 2nd edition, 2011.
- [15] P. Kouvelis and G. Yu. *Robust Discrete Optimization and Its Applications*. Kluwer Academic Publishers, Dordrecht, 1997.
- [16] D. Luenberger and Y. Ye. *Linear and Nonlinear Programming*. Springer, Berlin et al., 3rd edition, 2008.
- [17] J.M. Mulvey and R.J. Vanderbei. Robust optimization of large-scale systems. *Operations Research*, 43(2):264–281, 1995.
- [18] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.
- [19] A. Prékopa. *Stochastic Programming*. Kluwer, Dordrecht et al., 1995.
- [20] R.T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.
- [21] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming, Modeling and Theory*. SIAM, Philadelphia, 2009.
- [22] A.N. Shiryaev. *Probability*. Springer, New York et al., 2nd edition, 1996.
- [23] G.M. Ziegler. *Lectures on Polytopes*. Springer, New York et al., 1st edition, 2007.

Part C

Scenario-Based Model Predictive Control

Paper II

The Scenario Approach for Stochastic Model Predictive Control with Bounds on Closed-Loop Constraint Violations

Georg Schildbach · Lorenzo Fagiano · Christoph Frei · Manfred Morari

Abstract

Many practical applications in control require that constraints on the inputs and states of the system be respected, while optimizing some performance criterion. In the presence of model uncertainties or disturbances, for many control applications it suffices to keep the state constraints for at least a prescribed share of the time, as e.g., in building climate control or load mitigation for wind turbines. For such systems, a new control method of Scenario-Based Model Predictive Control (SCMPC) is presented in this paper. It optimizes the control inputs over a finite horizon, subject to robust constraint satisfaction under a finite number of random scenarios of the uncertainty and/or disturbances. While previous approaches have shown to be conservative (i.e., to stay far below the specified rate of constraint violations), the new method is the first to account for the special structure of the MPC problem in order to significantly reduce the number of scenarios. In combination with a new framework for interpreting the probabilistic constraints as average-in-time, rather than pointwise-in-time, the conservatism is eliminated. The presented method retains the essential advantages of SCMPC, namely the reduced computational complexity and the handling of arbitrary probability distributions. It also allows for adopting sample-and-remove strategies, in order to trade performance against computational complexity.

This manuscript is currently under review for publication in *Automatica*.

©2014 by the authors.

1. Introduction

Model Predictive Control (MPC) is a powerful approach for handling multi-variable control problems with constraints on the states and inputs. Its feedback control law can also incorporate feedforward information, e.g., about the future course of references and/or disturbances, and the optimization of a performance criterion of interest.

Over the past two decades, the theory of linear and robust MPC has matured considerably [22]. There are also widespread practical applications in diverse fields [26]. Yet many potentials of MPC are still not fully uncovered.

One active line of research is Stochastic MPC (SMPC), where the system dynamics are of a stochastic nature. They may be affected by additive disturbances [3, 10, 13, 14, 18, 19], by random uncertainty in the system matrices [11], or both [12, 15, 25, 30]. In this framework, a common objective is to minimize a cost function, while the system state is subject to chance constraints, i.e., constraints that have to be satisfied only with a given probability.

Stochastic systems with chance constraints arise naturally in some applications, such as building climate control [23], wind turbine control [12], or network traffic control [34]. Alternatively, they can be considered as relaxations of robust control problems, in which the robust satisfaction of state constraints can be traded for an improved cost performance.

A major challenge in SMPC is the solution to chance-constrained finite-horizon optimal control problems (FHOCs) in each sample time step. These correspond to non-convex stochastic programs, for which finding an exact solution is computationally intractable, except for very special cases [17, 31]. Moreover, due to the multi-stage nature of these problems, it generally involves the computation of multi-variate convolution integrals [10].

In order to obtain a tractable solution, various sample-based approximation approaches have been considered, e.g., [2, 4, 32]. They share the significant advantage of coping with generic probability distributions, as long as a sufficient number of random samples (or “scenarios”) can be obtained. The open-loop control laws can be approximated by sums of basis functions, as in the Q-design procedure proposed by [32]. However, these early approaches of Scenario-Based MPC (SCMPC) remain computationally demanding [2] and/or of a heuristic nature, i.e., without specific guarantees on the satisfaction of the chance constraints [4, 32].

More recent approaches [6, 7, 21, 24, 28, 33] are based on advances in the field of scenario-based optimization. However, these approaches share the drawback of being *conservative* when applied in a receding horizon fashion, i.e., the focus is either on obtaining a robust solution [6, 7, 33] or the chance constraints are over-satisfied by the closed loop system [21, 24, 28].

This conservatism of SCMPC represents a major practical issue, that is resolved by the contributions of this paper. In contrast to the previous results, the novel approach interprets the chance constraints as a time average, rather than pointwise-in-time with a

high confidence, which is much less restrictive. Furthermore, the sample size is reduced by exploiting the structural properties of the finite-horizon optimal control problem [29]. The approach also allows for the presence of multiple simultaneous chance constraints on the state, and an a-posteriori removal of adverse samples for improving the controller performance [21].

In the most general setting, this paper considers linear systems with stochastic additive disturbances and uncertainty in the system matrices, which may only be known through a sufficient number of random samples. The computational complexity can be traded against performance of the controller by removing samples a-posteriori, starting from a simple convex linear or quadratic program and converging to the optimal SMPC solution in the limit.

The paper is organized as follows: Section 2 presents a rigorous formulation of the optimal control problem that one would like to solve; Section 3 describes how an approximated solution is obtained by SCMPC; Section 4 develops the theoretical details, including the technical background and closed-loop properties; Section 5 demonstrates the application of the method to a numerical example; and Section 6 presents the main conclusions.

2. Optimal Control Problem

Consider a discrete-time control system with a linear stochastic transition map

$$x_{t+1} = A(\delta_t)x_t + B(\delta_t)u_t + w(\delta_t) , \quad x_0 = \bar{x}_0 , \quad (\text{C.1})$$

for some fixed initial condition $\bar{x}_0 \in \mathbb{R}^n$. The *system matrix* $A(\delta_t) \in \mathbb{R}^{n \times n}$ and the *input matrix* $B(\delta_t) \in \mathbb{R}^{n \times m}$ as well as the additive disturbance $w(\delta_t) \in \mathbb{R}^n$ are random, as they are (known) functions of a primal uncertainty δ_t . For notational simplicity, δ_t comprises all uncertain influences on the system at time t .

ASSUMPTION C.1—UNCERTAINTY (a) The uncertainties $\{\delta_0, \delta_1, \dots\}$, are independent and identically distributed (i.i.d.) random variables on a probability space (Δ, \mathbf{P}) . (b) A “sufficient number” of i.i.d. samples from δ_t can be obtained, either empirically or by a random number generator. ■

The support set Δ of δ_t and the probability measure \mathbf{P} on Δ are entirely generic. In fact, Δ and \mathbf{P} need not be known explicitly. The “sufficient number” of samples, which is required instead, will become concrete in later sections of the paper. Note that any issues arising from the definition of a σ -algebra on (Δ, \mathbf{P}) are glossed over in this paper, as they are unnecessarily technical. Instead, every relevant subset of Δ is assumed to be measurable.

The system (C.1) can be controlled by inputs $\{u_0, u_1, \dots\}$, to be chosen from a set of feasible inputs $\mathbb{U} \subset \mathbb{R}^m$. Since the future evolution of the system (C.1) is uncertain, it is

generally impractical to indicate all future inputs explicitly. Instead, each u_t should be determined by a static feedback law

$$\psi : \mathbb{R}^n \rightarrow \mathbb{U} \quad \text{with} \quad u_t = \psi(x_t) ,$$

based only on the current state of the system.

The optimal state feedback law ψ should be determined in order to minimize the time-average of expected stage costs $\ell : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}_{0+}$,

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbf{E}[\ell(x_t, u_t)] . \quad (\text{C.2})$$

Each stage cost is taken in expectation $\mathbf{E}[\cdot]$, since its arguments x_t and u_t are random variables, being functions of $\{\delta_0, \dots, \delta_{t-1}\}$. The time horizon T is considered to be very large, yet it may not be precisely known at the point of the controller design.

The minimization of the cost is subject to keeping the state inside a state constraint set \mathbb{X} for a given fraction of all time steps. For many applications, the robust satisfaction of the state constraint (i.e., $x_t \in \mathbb{X}$ at all times t) is too restrictive for the choice of ψ , and results in a poor performance in terms of the cost function. This is especially true in cases where the lowest values of the cost function are achieved close to the boundary of \mathbb{X} . Moreover, it may be impossible to enforce if the support of $w(\delta_t)$ is unknown and possibly unbounded.

In order to make this more precise, let $M_t := \mathbf{1}_{\mathbb{X}^c}(x_{t+1})$ denote the random variable indicating that $x_{t+1} \notin \mathbb{X}$, i.e., $\mathbf{1}_{\mathbb{X}^c} : \mathbb{R}^n \rightarrow \{0, 1\}$ is the indicator function on the complement \mathbb{X}^c of \mathbb{X} . The expected time-average of constraint violations should be upper bounded by some $\varepsilon \in (0, 1)$,

$$\mathbf{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} M_t\right] \leq \varepsilon . \quad (\text{C.3})$$

ASSUMPTION C.2—CONTROL PROBLEM (a) The state of the system can be measured at each time step t . (b) The set of *feasible inputs* \mathbb{U} is bounded and convex. (c) The *state constrained set* \mathbb{X} is convex. (d) The stage cost $\ell(\cdot, \cdot)$ is a convex function. ■

Assumption C.2(b) holds for most practical applications, and very large artificial bounds can always be introduced for input channels without natural bounds. Typical choices for the stage cost ℓ include

$$\ell(\xi, v) := \|Q_\ell \xi\|_1 + \|R_\ell v\|_1 , \quad (\text{C.4a})$$

$$\text{or } \ell(\xi, v) := \|Q_\ell \xi\|_\infty + \|R_\ell v\|_\infty , \quad (\text{C.4b})$$

$$\text{or } \ell(\xi, v) := \|Q_\ell \xi\|_2^2 + \|R_\ell v\|_2^2 , \quad (\text{C.4c})$$

where $Q_\ell \in \mathbb{R}^{n \times n}$ and $R_\ell \in \mathbb{R}^{m \times m}$ are positive semi-definite weighting matrices. Typical choices for the constraints \mathbb{U} and \mathbb{X} are polytopic or ellipsoidal sets.

Combining the previous discussions, the *optimal control problem (OCP)* can be stated as follows:

$$\min_{\psi} \quad \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{E}[\ell(x_t, u_t)] \quad , \quad (\text{C.5a})$$

$$\text{s.t.} \quad x_{t+1} = A(\delta_t)x_t + B(\delta_t)u_t + w(\delta_t) \quad , \quad x_0 = \bar{x}_0 \quad \forall t = 0, \dots, T-1 \quad , \quad (\text{C.5b})$$

$$\mathbf{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbf{1}_{\mathbb{X}^c}(x_t)\right] \leq \varepsilon \quad , \quad (\text{C.5c})$$

$$u_t = \psi(x_t) \quad \forall t = 0, \dots, T-1 \quad . \quad (\text{C.5d})$$

The equality constraints (C.5b) are understood to be substituted recursively to eliminate all state variables x_0, x_1, \dots, x_{T-1} from the problem. Thus only the state feedback law ψ remains as a free variable in (C.5).

REMARK C.1—ALTERNATIVE FORMULATIONS (a) Instead of the sum of expected values, the cost function (C.5a) can also be defined as a desired quantile of the sum of discounted stage costs. Then the problem formulation corresponds to a minimization of the “value-at-risk”, see e.g., [31]. (b) Multiple chance constraints on the state \mathbb{X}_j , each with an individual probability level ε_j , can be included without further complications. A single chance constraint is considered here for notational simplicity. ■

Many practical control problems can be cast in the general form of (C.5). For example in building climate control [23], the energy consumption of a building should be minimized, while its internal climate is subject to uncertain weather conditions and the occupancy of the building. The comfort range for the room temperatures may occasionally be violated without major harm to the system. Another example is wind turbine control [12], where the power efficiency of a wind turbine should be maximized, while its dynamics are subject to uncertain wind conditions. High stress levels in the blades must not occur too often, in order to achieve a desired fatigue life of the turbine.

3. Scenario-Based Model Predictive Control

The OCP is generally intractable, as it involves an infinite-dimensional decision variable ψ (the state feedback law) and a large number of constraints (growing with T). Therefore it is common to approximate it by various approaches, such as *Model Predictive Control (MPC)*.

3.1 Stochastic Model Predictive Control (SMPC)

The basic concept of MPC is to solve a tractable counterpart of (C.5) over a small horizon

N repeatedly at each time step. Only the first input of this solution is applied to the system (C.1). In Stochastic MPC (SMPC), a *Finite Horizon Optimal Control Problem (FHOC)* is formulated by introducing chance constraints on the state:

$$\min_{u_{0|t}, \dots, u_{N-1|t}} \sum_{t=0}^{N-1} \mathbf{E}[\ell(x_{i|t}, u_{i|t})] , \quad (\text{C.6a})$$

$$\text{s.t. } x_{i+1|t} = A(\delta_{t+i})x_{i|t} + B(\delta_{t+i})u_{i|t} + w(\delta_{t+i}), \quad x_{0|t} = x_t \quad \forall i = 0, \dots, N-1 , \quad (\text{C.6b})$$

$$\mathbf{P}[x_{i+1|t} \notin \mathbb{X}] \leq \varepsilon_i \quad \forall i = 0, \dots, N-1 , \quad (\text{C.6c})$$

$$u_{i|t} \in \mathbb{U} \quad \forall i = 0, \dots, N-1 . \quad (\text{C.6d})$$

Here $x_{i|t}$ and $u_{i|t}$ denote predictions and plans of the state and input variables made at time t , for i steps into the future. The current measured state x_t is introduced as an initial condition for the dynamics. The predicted states $x_{1|t}, \dots, x_{N|t}$ are understood to be eliminated by recursive substitution of (C.6b). Note that the predicted states are random by the influence of the uncertainties $\delta_t, \dots, \delta_{t+N-1}$.

The *probability levels* ε_i in the *chance constraints* (C.6c) usually coincide with ε from the OCP [14, 23, 30], but they may generally differ [34]. Some formulations also involve chance constraints over the entire horizon [12, 19], or as a combination with robust constraints [10, 18]. Other alternatives of SMPC consider integrated chance constraints [13], or constraints on the expectation of the state [25].

REMARK C.2—TERMINAL COST An optional (convex) terminal cost $\ell_f : \mathbb{R}^n \rightarrow \mathbb{R}_{0+}$ can be included in the FHOC [20, 27]. In this case the term

$$\mathbf{E}[\ell_f(x_{N|t})]$$

would be added to the cost function (C.6a). ■

The state feedback law provided by SMPC is given by a receding horizon policy: the current state x_t is substituted into (C.6b), then the FHOC is solved for an input sequence $\{u_{0|t}^*, \dots, u_{N-1|t}^*\}$, and the current input is set to $u_t := u_{0|t}^*$. This means that the FHOC must be solved online at each time step t , using the current measurement of the state x_t .

However, the FHOC is a stochastic program that remains difficult to solve, except for very special cases. In particular, the feasible set described by chance constraints is generally non-convex, despite of the convexity of \mathbb{X} , and hard to determine explicitly. Hence a further approximation shall be made by scenario-based optimization.

3.2 Scenario-Based Model Predictive Control (SCMPC)

The basic idea of Scenario-Based MPC (SCMPC) is to compute an optimal finite-horizon input trajectory $\{u'_{0|t}, \dots, u'_{N-1|t}\}$ that is feasible under K of sampled “scenarios” of the uncertainty. Clearly, the scenario number K has to be selected carefully in order to

attain the desired properties of the controller. In this section, the basic setup of SCMPC is discussed, while the selection of a value for K is deferred until Section 4.

More concretely, let $\delta_{i|t}^{(1)}, \dots, \delta_{i|t}^{(K)}$ be i.i.d. samples of δ_{t+i} , drawn at time $t \in \mathbb{N}$ for the prediction steps $i = 0, \dots, N-1$. For convenience, they are combined into *full-horizon samples* $\omega_t^{(k)} := \{\delta_{0|t}^{(k)}, \dots, \delta_{N-1|t}^{(k)}\}$, also called *scenarios*. The *Finite-Horizon Scenario Program (FHSCP)* then reads as follows:

$$\min_{u_{0|t}, \dots, u_{N-1|t}} \sum_{k=1}^K \sum_{i=0}^{N-1} \ell(x_{i|t}^{(k)}, u_{i|t}) , \quad (\text{C.7a})$$

$$\text{s.t. } x_{i+1|t}^{(k)} = A(\delta_{i|t}^{(k)})x_{i|t}^{(k)} + B(\delta_{i|t}^{(k)})u_{i|t} + w(\delta_{i|t}^{(k)}) ,$$

$$x_{0|t}^{(k)} = x_t \quad \forall i = 0, \dots, N-1, \quad \forall k = 1, \dots, K , \quad (\text{C.7b})$$

$$x_{i+1|t}^{(k)} \in \mathbb{X} \quad \forall i = 1, \dots, N-1, \quad \forall k = 1, \dots, K , \quad (\text{C.7c})$$

$$u_{i|t} \in \mathbb{U} \quad \forall i = 0, \dots, N-1 . \quad (\text{C.7d})$$

The dynamics (C.7b) provide K different state trajectories over the prediction horizon, each corresponding to one sequence of affine transition maps defined by a particular scenario $\omega_t^{(k)}$. Note that these K state trajectories are not fixed, as they are still subject to the inputs $u_{0|t}, \dots, u_{N-1|t}$. The cost function (C.7a) approximates (C.6a) as an average over all K scenarios. The state constraints (C.7c) are required to hold for K sampled state trajectories over the prediction horizon.

Applying a receding horizon policy, the SCMPC feedback law is defined as follows (see also Figure C.1, for $R = 0$). At each time step $t \in \mathbb{N}$ the current state measurement x_t is substituted into (C.7b), and the current input $u_t := u'_{0|t}$ is set to the first of the optimal FHSCP solution $\{u'_{0|t}, \dots, u'_{N-1|t}\}$, which is called the *scenario solution*.

Unlike many MPC approaches, SCMPC does not have an inherent guarantee of *recursive feasibility*, in the sense of [22, Sec. 4]. Hence for a proper analysis of the closed-loop system, the following is assumed.

ASSUMPTION C.3—RESOLVABILITY Under the SCMPC regime, each FHSCP admits a feasible solution at every time step t almost surely. ■

While Assumption C.3 appears to be restrictive from a theoretical point of view, it is often reasonable from a practical point of view. For some applications, such as buildings [23], recursive feasibility may hold by intuition, or it may be ensured by the use of *soft constraints* [26, Sec. 2]. All in all, MPC remains a useful tool in practice, even for difficult stochastic systems (C.1) without the possibility of an explicit guarantee of recursive feasibility.

The following are possible alternatives and also convex formulations of (C.7). The reasoning in each case is based on the theory in [29] and omitted for brevity.

REMARK C.3—ALTERNATIVE FORMULATIONS (a) Instead of the average cost in (C.7a), the minimization may concern the cost of a nominal trajectory, as e.g., in [24, 28]; or the average may be taken over any sample size other than K . (b) The inclusion of additional chance constraints into (C.7), as mentioned in Remark C.1(b), is straightforward. The number of scenarios K_j may generally differ between multiple chance constraints. (c) In case of a value-at-risk formulation, as in Remark C.1(a), the average cost in (C.7a) is replaced by the maximum:

$$\text{“} \sum_{k=1}^K \text{”} \longrightarrow \text{“} \max_{k=1, \dots, K} \text{”} ,$$

where the sample size K must be selected according to the desired risk level. ■

REMARK C.4—DISTURBANCE FEEDBACK In the FHSCP, the predicted control inputs $u_{0|t}, \dots, u_{N-1|t}$ may also be parameterized as a weighted sum of basis functions of the disturbances, as proposed in [32, 33]. In particular, for each time step $i = 1, \dots, N$, let $q_{i|t}^{(j)} : \Delta^{i-1} \rightarrow \mathbb{R}^m$ be a finite set $j \in \{1, \dots, J_i\}$ of pre-selected basis functions. The terms

$$\begin{aligned} u_{0|t} &:= c_{0|t} , \\ u_{i|t} &:= c_{i|t} + \sum_{j=1}^{J_i} \phi_i^{(j)} q_{i|t}^{(j)} (\delta_{0|t}^{(k)}, \dots, \delta_{i-1|t}^{(k)}) \quad \forall i = 1, \dots, N-1 , \end{aligned}$$

can be substituted into problem (C.7). The corrective control inputs $c_{0|t}, \dots, c_{N-1|t} \in \mathbb{R}^m$ become the new decision variables, and the weights $\phi_i^{(j)} \in \mathbb{R}$ for $i = 0, \dots, N-1$ and $j = 1, \dots, J_i$ can be determined on-line or off-line. ■

A disturbance feedback parameterization with an increasing number of basis functions J_1, \dots, J_{N-1} generally improves the quality of the SCMPC feedback, while increasing the number of decision variables and hence the computational complexity; see [32, 33] for more details.

Given the sampled scenarios, (C.7) is a convex optimization program for which efficient solution algorithms exist, depending on its structure [5]. In particular, if \mathbb{X} and \mathbb{U} are polytopic (respectively ellipsoidal) sets, then the FHSCP has linear (second-order cone) constraints. If the stage cost is either (C.4a,b), then the FHSCP has a reformulation with a linear objective function, using auxiliary variables. If the stage cost is (C.4c), then the FHSCP can be expressed as a quadratic program. More details on these formulation procedures are found in [20, pp. 154f.].

3.3 A-Posteriori Scenario Removal

A key merit of SCMPC is that it renders the uncertain control system (C.6b) into multiple deterministic affine systems (C.7b) by substituting particular scenarios. This significantly

simplifies the solution to the FHSCP, as compared to the FHOCPC. However, by introducing these random scenarios, a randomizing element is added to the SCMPC feedback law. In particular, the closed-loop system may occasionally show an erratic behavior due to highly unlikely outliers in the sampled scenarios.

This effect can be mitigated by a-posteriori scenario removal, see [9]. This allows for the *state constraints* (C.7c) corresponding to $R > 0$ scenarios to be removed *after* the outcomes of all samples have been observed. In exchange, the original sample size K must be (appropriately) increased over its value for $R = 0$. Any appropriate combination (K, R) is called a *sample-removal pair*. The choice of appropriate values for K and R is deferred to Section 4. The selection of removed scenarios is performed by a (*scenario*) *removal algorithm* [9, Def. 2.1].

DEFINITION C.1—REMOVAL ALGORITHM (a) For each $\xi \in \mathbb{R}^n$, the (*scenario*) *removal algorithm* $\mathcal{A}_\xi : \Delta^{NK} \rightarrow \Delta^{N(K-R)}$ is a deterministic function selecting $(K - R)$ out of K scenarios $\{\omega_t^{(1)}, \dots, \omega_t^{(K)}\}$. (b) The selected scenarios at time step t shall be denoted by

$$\Omega_t := \mathcal{A}_{x_t}(\omega_t^{(1)}, \dots, \omega_t^{(K)}) .$$

■

Definition C.1 is very general, in the sense that it covers a great variety of possible scenario removal algorithms. However, the most common and practical algorithms are described below:

Optimal Removal: The FHSCP is solved for all possible combinations of choosing R out of K scenarios. Then the combination that yields the lowest cost function value of all the solutions is selected. This requires the solution to K choose R instances of the FHSCP, a complexity that is usually prohibitive for larger values of R .

Greedy Removal: The FHSCP is first solved with all K scenarios. Then, in each of R consecutive steps, the state constraints of a single scenario are removed that yields the biggest improvement, either in the total cost or in the first stage cost. Thus the procedure terminates after solving $KR - R(R-1)/2$ instances of FHSCP.

Marginal Removal: The FHSCP is first solved with the state constraints of all K scenarios. Then, in each of R consecutive steps, the state constraints of a single scenario are removed based on the highest Lagrange multiplier. Hence the procedure requires the solution to K instances of FHSCP.

Figure C.1 depicts an algorithmic overview of SCMPC, for the general case with scenario removal $R > 0$. For the case without scenario removal, consider $R = 0$ and the selected scenarios $\Omega_t := \{\omega_t^{(1)}, \dots, \omega_t^{(K)}\}$.

At every time step t , perform the following steps:

1. Measure current state x_t .
2. Extract K scenarios $\omega_t^{(1)}, \dots, \omega_t^{(K)}$.
3. Remove R scenarios via \mathcal{A}_{x_t} , and solve FHSCP with only the state constraints of the remaining scenarios Ω_t .
4. Apply the first input of the scenario solution $u_t := u'_{0|t}$ to the system.

FIGURE C.1. Schematic overview of the SCMPC algorithm, for the case with scenario removal ($R > 0$) and without scenario removal ($R = 0$).

4. Problem Structure and Sample Complexity

For the SCMPC algorithm described in Section 3, the sample-removal pair (K, R) remains to be specified. Appropriate values for K and R are theoretically derived in this section. Their values generally depend on the control system and the constraints, and K is referred to as the *sample complexity* of the SCMPC problem.

For some intuition about this problem, suppose that $R \geq 0$ is fixed and the sample size K is increased. This means that the solution to the FHSCP becomes robust to more scenarios, with the following consequences. First, the average-in-time state constraint violations (C.3) decrease, in general. Therefore the state constraint will translate into a lower bound on K . Second, the computational complexity increases as well as the average-in-time closed-loop cost (C.2), in general. Therefore the objective is to choose K as small as possible, and ideally equal to its lower bound.

The higher the number of removed constraints $R \geq 0$, the higher will be the lower bound on K , in order for the state constraints (C.3) to be satisfied. Now consider pairs (R, K) of removed constraints R together with their corresponding lower bounds K , which equally satisfy the state constraints (C.3). For the intuition, suppose R is increased, so K increases as well. Then the computational complexity grows, due to more constraints in the FHSCP and the removal algorithm. At the same time, the solution quality of the FHSCP improves, in general, and hence the average-in-time closed-loop cost (C.2) decreases. Therefore R is usually fixed to a value that is as high as admitted by the available computational resources.

4.1 Support Rank

According to the classic scenario approach [8, 9], the relevant quantity for determining the sample size K for a single chance constraint (with a fixed R) is the number of *support constraints* [8, Def. 2.1]. In fact, K grows with the (unknown) number of support constraints, so the goal is to obtain a tight upper bound. For the classic scenario approach,

this upper bound is given by the dimension of the decision space [8, Prop.2.2], i.e., Nm in the case of the FHSCP.

The FHSCP is a multi-stage stochastic program, with multiple chance constraints (namely N , one per stage). This requires an extension to the classic scenario approach; the reader is referred to [29] for more details. Now each chance constraint contributes an individual number of support constraints, to which an upper bound must be obtained. These individual upper bounds are provided by the *support rank* of each chance constraint [29, Def. 3.6].

DEFINITION C.2—SUPPORT RANK (a) The *unconstrained subspace* \mathcal{L}_i of a constraint $i \in \{0, \dots, N - 1\}$ in (C.7c) is the largest (in the set inclusion sense) linear subspace of the search space \mathbb{R}^{Nm} that remains unconstrained by all sampled instances of i , almost surely. (b) The *support rank* of a constraint $i \in \{0, \dots, N - 1\}$ in (C.7c) is

$$\rho_i := Nm - \dim \mathcal{L}_i ,$$

where $\dim \mathcal{L}_i$ represents the dimension of the unconstrained subspace \mathcal{L}_i . ■

Note that the support rank is an inherent property of a particular chance constraint and it is not affected by the simultaneous presence of other constraints. Hence the set of constraints of the FHSCP may change, for instance, due to the reformulations of Remark C.1.

Besides the extension to multiple chance constraints, the support rank has the merit of a significant reduction of the upper bound on the number of support constraints. Indeed, the following two lemmas replace the classic upper bound Nm with much lower values, such as $l \leq n$ or m , depending on the problem structure.

For systems affected by *additive* disturbances only, the support rank of any state constraint in the FHSCP is given by the support rank $l \leq n$ of \mathbb{X} in \mathbb{R}^n (i.e., the co-dimension of the largest linear subspace that is unconstrained by \mathbb{X}).

LEMMA C.1—PURE ADDITIVE DISTURBANCES Let $l \leq n$ be the support rank of \mathbb{X} and suppose that $A(\delta_{i|t}^{(k)}) \equiv A$ and $B(\delta_{i|t}^{(k)}) \equiv B$ are constant and the control is not parameterized (as in Remark C.4). Then the support rank of any state constraint $i \in \{0, \dots, N - 1\}$ in (C.7c) is at most l . ■

For systems affected by *additive and multiplicative* disturbances, Lemma C.1 no longer holds. However, it will be seen that for the desired closed-loop properties, the relevant quantity for selecting the sample size K is the support rank ρ_1 of the state constraint on $x_{1|t}$ only. For this first predicted step, the support rank is restricted to at most m , under both additive and multiplicative disturbances.

LEMMA C.2—ADDITIVE AND MULTIPLICATIVE DISTURBANCES The support rank ρ_1 of constraint $i = 1$ in (C.7c) is at most m . ■

For the sake of readability, the proofs of Lemmas C.1 and C.2 are deferred to Appendix A. They effectively decouple the support rank, and hence the sample size K , from the horizon length N .

Note that the result of Lemma C.2 holds also for the parameterized control laws of Remark C.4. In this case, it decouples the sample size K from the number of basis functions J_i for all stages $i = 1, \dots, N - 1$.

Tighter bounds of ρ_1 than those in Lemmas C.1 and C.2 may exist, resulting from a special structure of the system (C.1) and/or the state constraint set \mathbb{X} . The basic insights to exploit this can be found in the Appendix A and [29].

4.2 Sample Complexity

This section describes the selection of the sample-removal pair (K, R) , based on a bound of the support rank ρ_1 . Throughout this section, the initial state x_t is considered to be fixed to an arbitrary value.

Let $V_t|x_t$ denote the (*first step*) *violation probability*, i.e., the probability with which the first predicted state falls outside of \mathbb{X} :

$$V_t|x_t := \mathbf{P}[A(\delta_t)x_t + B(\delta_t)u'_{0|t} + w(\delta_t) \notin \mathbb{X} | x_t] . \quad (\text{C.8})$$

Recall that $u'_{0|t}$ denotes the first input of the scenario solution $\{u'_{0|t}, \dots, u'_{N-1|t}\}$. Clearly, $u'_{0|t}$ and $V_t|x_t$ depend on the scenarios Ω_t that are substituted into the FHSCP at time t . The notation $u'_{0|t}(\Omega_t)$ and $V_t|x_t(\Omega_t)$ shall be used occasionally to emphasize this fact.

The violation probability $V_t|x_t(\Omega_t)$ can be considered as a random variable on the probability space $(\Delta^{KN}, \mathbf{P}^{KN})$, with support in $[0, 1]$. Here Δ^{KN} and \mathbf{P}^{KN} denote the KN -th product of the set Δ and the measure \mathbf{P} , respectively. For distinction, the expectation operator on (Δ, \mathbf{P}) is denoted \mathbf{E} , and that on $(\Delta^{KN}, \mathbf{P}^{KN})$ is denoted \mathbf{E}^{KN} .

The distribution of $V_t|x_t(\Omega_t)$ is unknown, being a complicated function of the entire control problem (C.6) and the removal algorithm \mathcal{A}_{x_t} . However, it is possible to derive the following upper bound on this distribution.

LEMMA C.3—UPPER BOUND ON DISTRIBUTION Let Assumptions C.1, C.2, C.3 hold and $x_t \in \mathbb{R}^n$ be an arbitrary initial state. For any violation level $\nu \in [0, 1]$,

$$\mathbf{P}^{KN}[V_t|x_t(\Omega_t) > \nu] \leq U_{K,R,\rho_1}(\nu) , \quad (\text{C.9a})$$

$$U_{K,R,\rho_1}(\nu) := \min \left\{ 1, \binom{R + \rho_1 - 1}{R} B(\nu; K, R + \rho_1 - 1) \right\} , \quad (\text{C.9b})$$

where $B(\cdot; \cdot, \cdot)$ represents the beta distribution function [1, frm. 26.5.3, 26.5.7],

$$B(\nu; K, R + \rho_1 - 1) := \sum_{j=0}^{R+\rho_1-1} \binom{K}{j} \nu^j (1 - \nu)^{K-j} .$$

■

Proof. The proof is a straightforward extension of [29, Thm. 6.7], where the bound on $V_t|x_t(\Omega_t)$ is saturated at 1. \square

This paper exploits the result of Lemma C.3 to obtain an upper bound on the expectation

$$\mathbf{E}^{KN}[V_t | x_t] := \int_{\Delta^{KN}} V_t|x_t(\Omega_t) d\mathbf{P}^{KN} . \quad (\text{C.10})$$

A reformulation via the indicator function $\mathbf{1} : \Delta^{KN} \rightarrow \{0, 1\}$ yields that

$$\begin{aligned} \mathbf{E}^{KN}[V_t | x_t] &= \int_{[0,1]} \int_{\Delta^{KN}} \mathbf{1}(V_t|x_t(\Omega_t) > \nu) d\mathbf{P}^{KN} d\nu \\ &= \int_{[0,1]} \mathbf{P}^{KN}[V_t|x_t(\Omega_t) > \nu] d\nu \\ &\leq \int_{[0,1]} U_{K,R,\rho_1}(\nu) d\nu . \end{aligned} \quad (\text{C.11})$$

DEFINITION C.3—ADMISSIBLE SAMPLE-REMOVAL PAIR A sample-removal pair (K, R) is *admissible* if its substitution into (C.11) yields $\mathbf{E}^{KN}[V_t | x_t] \leq \varepsilon$. \blacksquare

Whether a given sample-removal pair (K, R) is admissible can be tested by performing the one-dimensional numerical integration (C.11). It can easily be seen that the integral value (C.11) monotonically decreases with K and monotonically increases with R . Hence, if either K or R is fixed, an admissible sample-removal pair (K, R) can be determined e.g., by a bisection method. Moreover, if R is fixed, there always exist K large enough to generate an admissible pair (K, R) .

REMARK C.5—NO SCENARIO REMOVAL If $R = 0$, the integration (C.11) can be replaced by the exact analytic formula

$$\mathbf{E}^{KN}[V_t | x_t] \leq \frac{\rho_1}{K+1} . \quad (\text{C.12})$$

\blacksquare

Figure C.2 illustrates the monotonic relationship of the upper bound (C.11) in K and R . Supposing that $R = 0, 30, 100$ is fixed, the corresponding admissible pair (K, R) can be found by moving along the graphs until the desired violation level ε is reached. The solid and the dashed line correspond to different support dimensions $\rho_1 = 2$ and $\rho_1 = 5$.




FIGURE C.2. Upper bound on the expected violation probability $\mathbf{E}^{KN}[V_t | x_t]$, as a function of the sample size K , for different scenario removals R and support dimensions $\rho_1 = 2$ (solid lines) and $\rho_1 = 5$ (dashed lines).

4.3 Closed-Loop Properties

This section analyzes the closed-loop properties of the control system under the SCMPC law for an admissible sample-removal pair (K, R) . To this end, the underlying stochastic process is first described. Recall that

- x_0, \dots, x_{T-1} is the closed-loop trajectory, where x_t depends on all past uncertainties $\delta_0, \dots, \delta_{t-1}$ as well as all past scenarios $\Omega_0, \dots, \Omega_{t-1}$;
- V_0, \dots, V_{T-1} are the violation probabilities, where V_t depends on x_t and Ω_t , and hence on $\Omega_0, \dots, \Omega_t$ and $\delta_0, \dots, \delta_{t-1}$;
- M_0, \dots, M_{T-1} indicate the actual violation of the constraints, where M_t depends on x_{t+1} , and hence on $\Omega_0, \dots, \Omega_t$ and $\delta_0, \dots, \delta_t$.

At each time step t , there are a total of $D := (KN + 1)$ random variables, namely the scenarios together with the disturbance $\{\delta_t, \Omega_t\} \in \Delta^{(KN+1)} = \Delta^D$. In order to simplify notations, define

$$\mathcal{F}_t := \{\delta_0, \Omega_0, \dots, \delta_t, \Omega_t\} \in \Delta^{(t+1)D},$$

for any $t \in \{0, \dots, T-1\}$. These auxiliary variables allow for the random variables $x_t(\mathcal{F}_{t-1})$, $V_t(\mathcal{F}_{t-1}, \Omega_t)$, $M_t(\mathcal{F}_t)$ to be expressed in terms of their elementary uncertainties. Moreover, let $\mathbf{P}^{(t+1)D}$ denote the probability measure and $\mathbf{E}^{(t+1)D}$ the expectation

operator on $\Delta^{(t+1)D}$, for any $t \in \{0, \dots, T-1\}$.

Observe that $M_t \in \{0, 1\}$ is a Bernoulli random variable with (random) parameter V_t , because

$$\begin{aligned} \mathbf{E}[M_t | \mathcal{F}_{t-1}, \Omega_t] &= \int_{\Delta} M_t(\mathcal{F}_t) d\mathbf{P}(\delta_t) \\ &= V_t(\mathcal{F}_{t-1}, \Omega_t) \end{aligned} \quad (\text{C.13})$$

for any values of $\mathcal{F}_{t-1}, \Omega_t$.

THEOREM C.1 Let Assumptions C.1, C.2, C.3 hold and (K, R) be an admissible sample-removal pair. Then the expected time-average of closed-loop constraint violations (C.3) remains below the specified level ε ,

$$\mathbf{E}^{TD} \left[\frac{1}{T} \sum_{t=0}^{T-1} M_t \right] \leq \varepsilon. \quad (\text{C.14})$$

for any $T \in \mathbb{N}$. ■

Proof. By linearity of the expectation operator,

$$\begin{aligned} &\mathbf{E}^{TD} \left[\frac{1}{T} (M_0 + M_1 + \dots + M_{T-1}) \right] \\ &= \frac{1}{T} (\mathbf{E}^D[M_0] + \mathbf{E}^{2D}[M_1] + \dots + \mathbf{E}^{TD}[M_{T-1}]) \\ &= \frac{1}{T} (\mathbf{E}^{D-1}[V_0] + \mathbf{E}^{2D-1}[V_1] + \dots + \mathbf{E}^{TD-1}[V_{T-1}]), \end{aligned}$$

by virtue of (C.13). Moreover, for any $t \in \{0, \dots, T-1\}$,

$$\mathbf{E}^{(t+1)D-1}[V_t] = \int_{\Delta^{tD}} \underbrace{\mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}]}_{\leq \varepsilon} d\mathbf{P}^{tD} \leq \varepsilon,$$

where the integrand is pointwise upper bounded by ε because (K, R) is an admissible sample-removal pair. □

Theorem C.1 shows that the chance constraints of the OCP can be expected to be satisfied over any finite time horizon T . The next Lemma C.4 sets the stage for an even stronger result, Theorem C.2, showing that the chance constraint are satisfied almost surely as $T \rightarrow \infty$.

LEMMA C.4 If Assumptions C.1, C.2, C.3 hold, then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(M_t - \mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}] \right) = 0 \quad (\text{C.15})$$

almost surely. ■

Proof. For any $t \in \mathbb{N}$, define $Z_t := M_t - \mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}]$ and observe that

$$\mathbf{E}^D[Z_t | \mathcal{F}_{t-1}] \tag{C.16}$$

$$\begin{aligned} &= \mathbf{E}^D[M_t | \mathcal{F}_{t-1}] - \mathbf{E}^D[\mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}] | \mathcal{F}_{t-1}] \\ &= \mathbf{E}^D[M_t | \mathcal{F}_{t-1}] - \mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}] \\ &= 0, \end{aligned} \tag{C.17}$$

by virtue of (C.13). In probabilistic terms, this says that $\{Z_t\}_{t \in \mathbb{N}}$ is a sequence of martingale differences. Moreover,

$$\sum_{t=0}^{\infty} \frac{1}{(t+1)^2} \mathbf{E}^D[Z_t^2 | \mathcal{F}_{t-1}] < \infty \tag{C.18}$$

almost surely, because $|Z_t| \leq 1$ is bounded for $t \in \mathbb{N}$. Therefore [16, Thm. 2.17] can be applied, which yields that

$$\sum_{t=0}^{T-1} \frac{1}{t+1} Z_t \tag{C.19}$$

converges almost surely as $T \rightarrow \infty$. The result (C.15) now follows by use of Kronecker's Lemma, [16, p. 31]. \square

Note that Lemma C.4 does not imply that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} M_t = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}] \tag{C.20}$$

almost surely, because it is not clear that the right-hand side converges almost surely. However, if it converges almost surely, then (C.20) holds.

THEOREM C.2 Let Assumptions C.1, C.2, C.3 hold and (K, R) be an admissible sample-removal pair. Then

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} M_t \leq \varepsilon \tag{C.21}$$

almost surely. \blacksquare

Proof. From Lemma C.4,

$$\begin{aligned}
0 &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(M_t - \mathbf{E}^{D-1}[V_t | \mathcal{F}_{t-1}] \right) \\
&\geq \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (M_t - \varepsilon) \\
&= \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} M_t - \varepsilon
\end{aligned} \tag{C.22}$$

almost surely, where the second line follows from Definition C.3. \square

5. Numerical Example

5.1 System Data

Consider the stochastic linear system

$$x_{t+1} = \begin{bmatrix} 0.7 & -0.1(2 + \theta_t) \\ -0.1(3 + 2\theta_t) & 0.9 \end{bmatrix} x_t + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} u_t + \begin{bmatrix} w_t^{(1)} \\ w_t^{(2)} \end{bmatrix},$$

where $x_0 = [1 \ 1]^T$. Here $\theta_t \sim \mathcal{U}([0, 1])$ is uniformly distributed on the interval $[0, 1]$ and $w_t^{(1)}, w_t^{(2)} \sim \mathcal{N}(0, 0.1)$ are normally distributed with mean 0 and variance 0.1. The inputs are confined to

$$\mathbb{U} := \{v \in \mathbb{R}^2 \mid |v^{(1)}| \leq 5 \wedge |v^{(2)}| \leq 5\},$$

and two state constraints are considered:

$$\mathbb{X}_1 := \{\xi \in \mathbb{R}^2 \mid \xi^{(1)} \geq 1\}, \quad \mathbb{X}_2 := \{\xi \in \mathbb{R}^2 \mid \xi^{(2)} \geq 1\},$$

either individually or in combination $\mathbb{X} := \mathbb{X}_1 \cap \mathbb{X}_2$. The stage cost function is chosen to be of the quadratic form (C.4c), with the weights $Q_\ell := I$ and $R_\ell := I$. The MPC horizon is set to $N := 5$.

5.2 Joint Chance Constraint

The support rank of the joint chance constraint \mathbb{X} is bounded by $\rho_1 = 2$. Figure C.3 depicts a phase plot of the closed-loop system trajectory, for two admissible sample-removal pairs (a) (19, 0) and (b) (1295, 100), corresponding to $\varepsilon = 10\%$. Instances in which the state trajectory leaves \mathbb{X} are indicated in red. Note that the distributions are centered around a similar mean in both cases, however the case $R = 0$ features stronger outliers than $R = 100$.




FIGURE C.3. Phase plot of closed-loop system trajectory (red: violating states; black: other states). The axis lines mark the boundary of the feasible set \mathbb{X} .

Table C.1 shows the empirical results of a simulation of the closed-loop system over $T = 10,000$ time steps. Note that there is essentially no conservatism in the case of no removals ($R = 0$). Some minor conservatism is present for small removal sizes, disappearing asymptotically as $R \rightarrow \infty$. At the same time, the reduction of the average closed-loop cost ℓ_{avg} is minor for this example, while the standard deviation ℓ_{std} is affected significantly.

To highlight the impact of the presented SCMPC approach, the results of Table C.1 can be compared to those of previous SCMPC approaches [6, 28]. The sample size is 19 (compared to about 400), and the empirical share of constraint violations in closed-loop is 9.87% (compared to about 0.05%). These figures become even worse when longer horizons are considered; e.g., for $N = 20$, previous approaches require about 900 samples and yield about 0.2% violations.

$\varepsilon = 10\%$	$R = 0$	$R = 50$	$R = 100$	$R = 500$
K	19	702	1,295	5,723
V_{avg}	9.87%	7.37%	8.06%	8.74%
ℓ_{avg}	3.78	3.75	3.72	3.68
ℓ_{std}	0.54	0.44	0.42	0.37

TABLE C.1. Joint chance constraint: closed-loop results for mean violations V_{avg} , mean stage cost ℓ_{avg} , and standard deviation of stage costs ℓ_{std} .

5.3 Individual Chance Constraints

For the same example, the two chance constraints \mathbb{X}_1 and \mathbb{X}_2 are now considered separately, with the individual probability levels $\varepsilon_1 = 5\%$ and $\varepsilon_2 = 10\%$. Each support rank is bounded by $\rho_1 = 1$. Figure C.4 depicts a phase plot of the closed-loop system trajectory,

for the admissible sample-removal pairs (a) $(19, 0)$, $(9, 0)$ and (b) $(2020, 100)$, $(1010, 100)$.

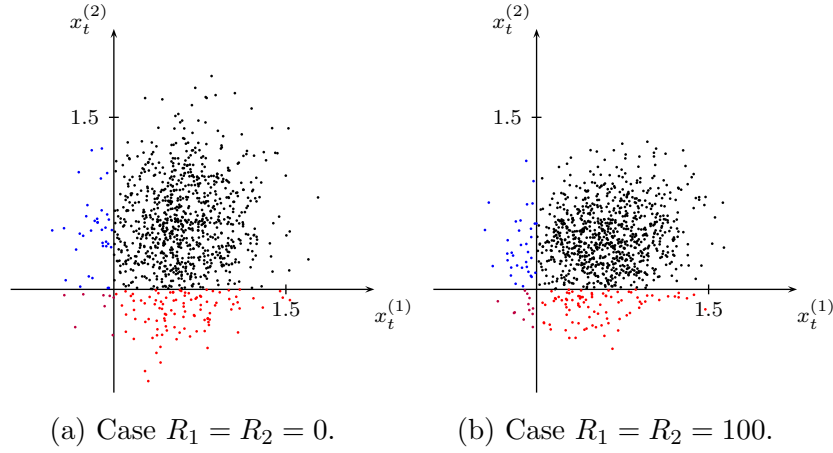


FIGURE C.4. Phase plot of closed-loop system trajectory (blue, red, purple: violating states of \mathbb{X}_1 , \mathbb{X}_2 , \mathbb{X}_1 and \mathbb{X}_2 ; black: other states). The axis lines mark the boundaries of the feasible sets \mathbb{X}_1 and \mathbb{X}_2 , respectively.

Table C.2 shows the empirical results of a simulation of the closed-loop system over $T = 10,000$ time steps. Note that there is very little conservatism in all cases. As in the previous example, the reduction of the average closed-loop cost ℓ_{avg} is minor, while the standard deviation ℓ_{std} is affected significantly.

$\varepsilon_1 = 5\%$, $\varepsilon_2 = 10\%$	$R_1 = R_2 = 0$	$R_1 = R_2 = 50$	$R_1 = R_2 = 100$
K_1	19	1,020	2,020
K_2	9	510	1,010
$V_{\text{avg},1}$	5.14%	4.84%	4.95%
$V_{\text{avg},2}$	9.94%	9.81%	9.93%
ℓ_{avg}	3.67	3.62	3.51
ℓ_{std}	0.54	0.46	0.42

TABLE C.2. Single chance constraint: closed-loop results for mean violations $V_{\text{avg},1}$ and $V_{\text{avg},2}$ of \mathbb{X}_1 and \mathbb{X}_2 , mean stage cost ℓ_{avg} , and standard deviation of stage costs ℓ_{std} .

6. Conclusion

The paper has presented new results on Scenario-Based Model Predictive Control (SCMPC). By focusing on the average-in-time probability of constraint violations and by exploiting the multi-stage structure of the finite-horizon optimal control problem

(FHOCP), the number of scenarios has been greatly reduced compared to previous approaches. Moreover, the possibility to adopt a-posteriori constraint removal strategies is also accommodated. Due to its computational efficiency, the presented approach paves the way for a tractable application of Stochastic Model Predictive Control (SMPC) to large-scale problems with hundreds of decision variables.

A. Proof of Lemmas C.1 and C.2

The particular bounding arguments follow rather easily after some general observations on the support rank. Pick any state constraint $i \in \{1, \dots, N\}$ from (C.7c). Recursively substituting the dynamics (C.7b), the constrained state can be expressed as

$$x_{i|t}^{(k)} = (A_{i|t}^{(k)} \cdot \dots \cdot A_{0|t}^{(k)})x_t + \bar{A}_{i|t}^{(k)} \bar{B}_{i|t}^{(k)} \begin{bmatrix} u_{0|t} \\ \vdots \\ u_{N-1|t} \end{bmatrix} + \bar{A}_{i|t}^{(k)} \begin{bmatrix} w_{0|t}^{(k)} \\ \vdots \\ w_{i-1|t}^{(k)} \end{bmatrix}, \quad (\text{C.23a})$$

$$\bar{A}_{i|t}^{(k)} := \begin{bmatrix} A_{i|t}^{(k)} \cdot \dots \cdot A_{1|t}^{(k)} \\ \vdots \\ A_{1|t}^{(k)} \\ I \end{bmatrix}^T, \quad (\text{C.23b})$$

$$\bar{B}_{i|t}^{(k)} := \begin{bmatrix} B_{0|t}^{(k)} & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & B_{1|t}^{(k)} & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & 0 & \dots & 0 \\ 0 & 0 & \dots & B_{i|t}^{(k)} & 0 & \dots & 0 \end{bmatrix}, \quad (\text{C.23c})$$

where $I \in \mathbb{R}^{n \times n}$ denotes the identity matrix, and for any $i = 0, \dots, N-1$ the following abbreviations are used:

$$A_{i|t}^{(k)} := A(\delta_{i|t}^{(k)}), \quad B_{i|t}^{(k)} := B(\delta_{i|t}^{(k)}), \quad w_{i|t}^{(k)} := w(\delta_{i|t}^{(k)}).$$

Let $l \leq n$ be the support rank of \mathbb{X} , i.e., the co-dimension of the largest linear subspace that is unconstrained by \mathbb{X} . Then there exists a projection matrix $P \in \mathbb{R}^{l \times n}$ such that for each $x \in \mathbb{R}^n$

$$x \in \mathbb{X} \iff Px \in P\mathbb{X} := \{P\xi \mid \xi \in \mathbb{X}\}.$$

For example, if the state constraint concerns only the first two elements of the state vector, then $l = 2$ and $P \in \mathbb{R}^{2 \times n}$ may contain the first two unit vectors $e_1, e_2 \in \mathbb{R}^n$ as its rows.

Proof of Lemma C.1

If $A(\delta_{i|t}^{(k)}) \equiv A$ and $B(\delta_{i|t}^{(k)}) \equiv B$ are constant for all $i \in \{0, \dots, N-1\}$, then (C.23a) reduces to

$$\underbrace{\begin{bmatrix} PA^{i-1}B & \dots & P & 0 & \dots \end{bmatrix}}_{\text{rank}(\cdot) \leq l} \begin{bmatrix} u_{0|t} \\ \vdots \\ u_{N-1|t} \end{bmatrix} + PA^i x_t + \begin{bmatrix} PA^{i-1}B & \dots & P \end{bmatrix} \begin{bmatrix} w_{0|t}^{(k)} \\ \vdots \\ w_{i-1|t}^{(k)} \end{bmatrix} \in P\mathbb{X} , \quad (\text{C.24})$$

for any $i \in \{1, \dots, N\}$. The rank of the first matrix of dimension $l \times Nm$ can be at most l , and therefore it has a null space of dimension at least $Nm - l$. The disturbance has no effect on this null space, because it enters only through the third, additive term in (C.24). Hence this null space is clearly an unconstrained subspace of the constraint and $\rho_i \leq l \leq n$ for all $i \in \{1, \dots, N\}$, proving Lemma C.1.

Proof of Lemma C.2

Consider the first state constraint $i = 1$ of (C.7c). Here (C.23a) reduces to

$$\underbrace{\begin{bmatrix} P\bar{B}_{0|t}^{(k)} & 0 & \dots & 0 \end{bmatrix}}_{\text{rank}(\cdot) \leq m} \begin{bmatrix} u_{0|t} \\ \vdots \\ u_{N-1|t} \end{bmatrix} + PA_{0|t}^{(k)} x_t + Pw_{0|t}^{(k)} \in P\mathbb{X} . \quad (\text{C.25})$$

The rank of the first matrix can here be at most m for all outcomes of $\bar{B}_{0|t}^{(k)}$, because the last $(N-1)m$ variables in the decision vector are always in its null space. Hence $\rho_1 \leq m$ in all cases, proving Lemma C.2.

Disturbance Feedback

For the case of parameterized control laws as in Remark C.4, it will be shown that the argument of Lemma C.2 continues to apply. Define for any $i = 1, \dots, N-1$

$$Q_{i|t}^{(k)} := \underbrace{\begin{bmatrix} q_{i|t}^{(1)} & q_{i|t}^{(2)} & \dots & q_{i|t}^{(J_i)} \end{bmatrix}}_{\in \mathbb{R}^{m \times J_i}} , \quad \Phi_{i|t} := \underbrace{\begin{bmatrix} \phi_{i|t}^{(1)} \\ \vdots \\ \phi_{i|t}^{(J_i)} \end{bmatrix}}_{\in \mathbb{R}^{J_i}} ,$$

where $q_{i|t}^{(j)} := q_{i|t}^{(j)}(\delta_{0|t}^{(k)}, \dots, \delta_{i|t}^{(k)})$ is used as an abbreviation. Then the vector of control

inputs under scenario $k = 1, \dots, K$ can be put into the affine expression

$$\begin{bmatrix} u_{0|t} \\ u_{1|t}^{(k)} \\ \vdots \\ u_{N-1|t}^{(k)} \end{bmatrix} = \begin{bmatrix} c_{0|t} \\ c_{1|t} \\ \vdots \\ c_{N-1|t} \end{bmatrix} + \underbrace{\begin{bmatrix} 0 & 0 & \dots & 0 \\ 0 & Q_{1|t}^{(k)} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & Q_{N-1|t}^{(k)} \end{bmatrix}}_{=:\bar{Q}_t^{(k)}} \underbrace{\begin{bmatrix} \Phi_{0|t} \\ \Phi_{1|t} \\ \vdots \\ \Phi_{N-1|t} \end{bmatrix}}_{=:\bar{\Phi}_t}.$$

Substitute this for the original decision vector into (C.25). In the case of off-line optimization, where $\bar{\Phi}_t$ is fixed, and in the case of on-line optimization, where $\bar{\Phi}_t$ is part of the decision variables, the same rank argument as before applies.

Acknowledgments

Research leading to these results has received funding from the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement number FP7-ICT-2009-4 248940, and under grant agreement number PEOF-GA-2009-252284. Christoph Frei gratefully acknowledges financial support by the Natural Sciences and Engineering Research Council of Canada.

References

- [1] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions*. Dover Publications, New York, 9th edition, 1970.
- [2] I. Batina. *Model Predictive Control for Stochastic Systems by Randomized Algorithms*. Ph.d. thesis, Technische Universiteit Eindhoven, Eindhoven, Netherlands, 2004.
- [3] I. Batina, A.A. Stoorvogel, and S. Weiland. Optimal control of linear, stochastic systems with state and input constraints. In *41st IEEE Conference on Decision and Control*, Las Vegas (NV), United States, 2002.
- [4] L. Blackmore, M. Ono, A. Bektassov, and B. Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Transactions on Robotics*, 26(3):502–516, 2010.
- [5] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.
- [6] G.C. Calafiore and L. Fagiano. Robust model predictive control via scenario optimization. *IEEE Transactions on Automatic Control*, 58(1):219–224, 2013.

- [7] G.C. Calafiore and L. Fagiano. Stochastic model predictive control of LPV systems via scenario optimization. *Automatica*, 49(6):1861–1866, 2013.
- [8] M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19:1211–1230, 2008.
- [9] M.C. Campi and S. Garatti. A sampling and discarding approach to chance-constrained optimization: Feasibility and optimality. *Journal of Optimization Theory and Applications*, 148:257–280, 2011.
- [10] M. Cannon, B. Kouvaritakis, S.V. Raković, and Q. Cheng. Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Transactions on Automatic Control*, 56(1):194–200, 2011.
- [11] M. Cannon, B. Kouvaritakis, and X. Wu. Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints. *Automatica*, 45(1):167–172, 2009.
- [12] M. Cannon, B. Kouvaritakis, and X. Wu. Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty. *IEEE Transactions on Automatic Control*, 54(7):1626–1632, 2009.
- [13] D. Chatterjee, P. Hokayem, and J. Lygeros. Stochastic receding horizon control with bounded control inputs: A vector space approach. *IEEE Transactions on Automatic Control*, 56(11):2704–2710, 2011.
- [14] E. Cinquemani, M. Agarwal, D. Chatterjee, and J. Lygeros. Convexity and convex approximations of discrete-time stochastic control problems with constraints. *Automatica*, 47(9):2082–2087, 2011.
- [15] D. Muñoz de la Peña, A. Bemporad, and T. Alamo. Stochastic programming applied to model predictive control. In *44th IEEE Conference on Decision and Control*, Seville, Spain, 2005.
- [16] P. Hall and C.C. Heyde. *Martingale Limit Theory and Its Application*. Academic Press, New York et al., 1980.
- [17] P. Kall and J. Mayer. *Stochastic Linear Programming*. Springer, New York et al., 2nd edition, 2011.
- [18] B. Kouvaritakis, M. Cannon, S.V. Raković, and Q. Cheng. Explicit use of probabilistic distributions in linear predictive control. *Automatica*, 46:1719–1724, 2010.
- [19] P. Li, M. Wendt, and G. Wozny. A probabilistically constrained model predictive controller. *Automatica*, 38:1171–1176, 2002.
- [20] J.M. Maciejowski. *Predictive Control with Constraints*. Pearson Education, Harlow, 2002.
- [21] J. Matuško and F. Borrelli. Scenario-based approach to stochastic linear predictive control. In *51st IEEE Conference on Decision and Control*, Maui (HI), United States, 2012.

- [22] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.
- [23] F. Oldewurtel, A. Parisio, C.N. Jones, D. Gyliastras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27, 2012.
- [24] M. Prandini, S. Garatti, and J. Lygeros. A randomized approach to stochastic model predictive control. In *51st IEEE Conference on Decision and Control*, Maui (HI), United States, 2012.
- [25] J.A. Primbs and C.H. Sung. Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Transactions on Automatic Control*, 54(2):221–230, 2012.
- [26] S.J. Qin and T.A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11:733–764, 2003.
- [27] J.B. Rawlings and D.Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, Madison (WI), 2009.
- [28] G. Schildbach, G.C. Calafiore, L. Fagiano, and M. Morari. Randomized model predictive control for stochastic linear systems. In *American Control Conference*, Montréal, Canada, 2012.
- [29] G. Schildbach, L. Fagiano, and M. Morari. Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4):2479–2501, 2013.
- [30] A.T. Schwarm and M. Nikolaou. Chance-constrained model predictive control. *AIChE Journal*, 45(8):1743–1752, 1999.
- [31] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Stochastic Programming, Modeling and Theory*. SIAM, Philadelphia, 2009.
- [32] J. Skaf and S. Boyd. Nonlinear q-design for convex stochastic control. *IEEE Transactions on Automatic Control*, 54(10):2426–2430, 2009.
- [33] P. Vayanos, D. Kuhn, and B. Rustem. A constraint sampling approach for multi-stage robust optimization. *Automatica*, 48:459–471, 2012.
- [34] J. Yan and R.R. Bitmead. Incorporating state estimation into model predictive control and its application to network traffic control. *Automatica*, 41:595–604, 2005.

Part D

Application to Supply Chain Management

Paper III

Scenario-Based Model Predictive Control for Multi-Echelon Supply Chain Management

Georg Schildbach · Manfred Morari

Abstract

Policies for managing multi-echelon supply chains can mathematically be considered as large-scale dynamic programs, affected by uncertainty and incomplete information. Except for a few special cases, optimal solutions are computationally intractable for systems of realistic size. This paper proposes a novel approximation scheme based on scenario-based model predictive control (SCMPC), using recent results in scenario-based optimization. The presented SCMPC approach can handle supply chains with stochastic planning uncertainties from various sources (demands, lead times, prices, etc.) and of a very general nature (distributions, correlations, etc.). Moreover, it guarantees a specified customer service level, when applied in a rolling horizon fashion. At the same time, SCMPC is computationally efficient and able to tackle problems of a similar scale as manageable by deterministic optimization. For a large class of supply chain models, SCMPC may therefore offer substantial advantages over robust or stochastic optimization.

This manuscript has been submitted for publication to European Journal of Operational Research.

©2014 by the authors.

1. Introduction

Methods for the automated management of supply chains continue to attract significant research attention [28]. The main reason is that even small improvements carry an enormous potential for cost savings and performance enhancements in actual supply chain systems [14]. However, the mathematical formulations of multi-echelon supply chain problems poses great challenges for modern theories of optimization and control. In essence, these formulations represent large-scale dynamic optimization problems, affected by uncertainty and incomplete information.

The difficulty of these problems precludes the exact computation of stochastic optimal control policies, except for very special cases [11, 20, 28]. A variety of approximation approaches have therefore been proposed, and shown to be effective for certain problem types [28]. One promising direction are methods related to *model predictive control* (MPC) [13, 21]. Various studies have demonstrated the potential of MPC to reduce the operation costs of supply chains; see Sarimveis et al. [28] for an excellent overview.

MPC assumes that a *model* of the actual supply chain is available to the decision maker. This model is used to predict the state of the system (inventories, service level, etc.) over a fixed time horizon, and to optimize the operational decisions (production, shipments, etc.) through an optimization program. The model, however, may contain *inaccurate parameters* (capacities, etc.) or *uncertain quantities* (demands, lead times, etc.). MPC is known to have good robustness properties with respect to mismatches between model and system parameters; cf. Braun et al. [6, 7]. Another advantage of MPC is that available forecasts on uncertain quantities can easily be integrated into the planning problem; cf. Bose and Pekny [5].

The MPC theory is concerned with the *dynamic implementation*—i.e., the long-term behavior of the supply chain under a rolling horizon implementation of the numerical optimization. In particular, the objective of MPC is to minimize the long-term costs of operation, while respecting the capacity constraints and maintaining a specific customer service level [5]. In this respect, it is well known that an explicit *uncertainty model* generally improves the performance of MPC over the deterministic-equivalent approach [18, 20, 37]. This is because the uncertainty model conveys important information about the variability of uncertain quantities, as compared to substituting them by their expected or the most likely values.

Two uncertainty models have previously been considered for integration into MPC: *robust models* and *stochastic models* [2, 37]. Robust models consider uncertainties to be set-bounded, so the resulting decisions are made irrespective of the probability with which various scenarios may occur; e.g., Li and Marlin [20]. Stochastic models have difficulties in handling multiple types of uncertainties with general correlations, or they are computationally intractable for realistic systems; e.g., Yıldırım et al. [39].

This paper presents a novel way of handling uncertainty in MPC, namely by considering sampled *uncertainty scenarios*. Hence the approach is called *scenario-based model predictive control* (SCMPC). As the sampled scenarios are eventually fixed, SCMPC can

handle uncertainties from various sources (demands, lead times, prices, etc.) and of a very general nature (distributions, correlations, etc.). The underlying theory configures SCMPC to satisfy a specified customer service level in its rolling horizon implementation. At the same time, SCMPC is computationally efficient. It is able to tackle problems of a similar scale as manageable by deterministic optimization.

The paper is organized as follows. The remainder of Section 1 provides a brief review on existing approaches for supply chain management and an introduction to the proposed method. Section 2 introduces an exemplary supply chain model, which is considered as a case study throughout the paper. The novel method described in Section 3 is indeed far more general, since it can be applied to a large variety of different models. For a demonstration of its properties, Section 4 presents some results obtained from numerical experiments. Finally, Section 5 provides a brief summary and conclusion.

1.1 Supply Chain Management

Since the seminal work of Clark and Scarf [10], a great variety of quantitative methods have been proposed for the inventory management in multi-echelon supply chains. Two fundamental approaches can be distinguished [11]: *centralized approaches* and *decentralized approaches*. Both approaches are discussed briefly below.

Decentralized Approaches. In a decentralized approach, each node of the supply chain network manages its inventory locally, typically as a *pull system*. The local replenishment decisions are commonly based on a *base-stock policy* or a (s, S) -*policy* [34]. The placement of safety stocks constitutes a strategic and centralized decision of *supply chain design*, with a critical effect on the overall performance of the system [15].

A key advantage of the decentralized approach is that—after the base stock or order levels have been determined—the *supply chain operation* is simple and requires little computations. Therefore, well-tuned decentralized base-stock policies have been successfully deployed in several industrial applications [15]. The base-stock policies can be adjusted to cope with non-stationary demands, as shown by Graves and Willems [16]. Moreover, they can also account for capacity constraints, see Schoenmeyr and Graves [31], and for evolving demand forecasts, see Schoenmeyr and Graves [32].

Decentralized approaches face potential drawbacks in practical applications. First, multiple products may share a single capacity constraints for processing or transportation. Then it is unclear how these products should be prioritized or how long they should be procured in advance. Second, base-stock policies are unable to choose between multiple sources for their supply. Therefore they are limited to simple network topologies. Third, the base-stock levels may have to be adapted in a dynamic manner. These adjustments are known to cause the undesirable *bullwhip effect* [19, 28]. In fact, higher demands lead to higher base stocks levels at all stages of the supply chain, and hence a disproportionate wave of replenishments propagates upwards through the supply chain. Fourth, the manufacturing stages may not have the flexibility to adjust their production as an instantaneous reaction to customer demand. For example, Sohdi [36] argues that for the consumer electronics industry the production schedule must usually be fixed several

weeks in advance.

Some of these shortcomings have been addressed by *decentralized MPC*, where separate MPC policies are designed for the local use in multiple parts of the supply chain [5, 12, 25]. The design typically assumes a *cooperative management* (as opposed to a *non-cooperative management*); cf. Subramanian et al. [38]. This means that the local decisions are made with the intent to optimize the overall performance of the supply chain. The goal of decentralized MPC is then to recover, as closely as possible, the performance of centralized MPC, while lowering its computational burden [5, 25].

Centralized Approaches. Centralized approaches can address the above issues by optimizing the decisions for the entire supply chain. They are often organized as *push systems*. A fundamental requirement is the existence of a central coordinator, who has access to all local information and the authority to make all relevant decisions throughout the network [25].

For centralized approaches, the decisions of each node are usually determined by solving one (or several) comprehensive mathematical program(s) for the entire supply chain. To this end, a large variety of mathematical models have been proposed; cf. the survey of Mula et al. [22].

Centralized approaches may reap the benefits from coordinated action of all nodes in the supply chain network. They are able to allocate limited capacities to multiple products and also to handle multiple suppliers with possibly different prices [9]. Moreover, they can cope with very general network topologies and the integration of forecasts is straightforward [27, 36]. They have also been successfully implemented in industrial applications [18, 27, 36].

The main bottleneck of centralized approaches lies in the numerical solution to a typically large-scale mathematical program. In particular, the problem dimensions grow with the time horizon, the number of products, and the network size. Furthermore, the problem becomes distinctly more difficult if it also accounts for uncertain parameters [37]. Therefore many existing contributions have chosen a *certainty-equivalence approach*, where uncertain quantities are substituted by fixed values (e.g., their expected or most likely values) [22]. However, this may lead to a considerable performance degradation [18, 20, 37].

1.2 Novel Contributions

This paper presents a novel centralized approach for multi-echelon supply chain management. Its core strength is the inclusion of uncertain parameters of very general nature, while it retains a computationally complexity that is similar to deterministic approaches. The presented method can be used with a great variety of different supply chain models. It builds on recent results of Campi and Garatti [8] on the *scenario-based optimization approach* (scenario approach) and of Schildbach et al. [29] on *scenario-based model predictive control* (SCMPC).

SCMPC solves, repeatedly, an optimal planning problem over a finite time horizon. Unlike most centralized approaches, SCMPC analyzes its rolling horizon implementation

rather than just the finite horizon optimal solution. In the long run, SCMPC keeps capacity constraints as well as pre-specified *customer service levels*, which are defined as the fraction of time steps in which customer demands are met from available inventory [3, 4]. At the same time it aims at keeping the long-term operational costs low.

SCMPC can handle supply chains with multiple *operational uncertainties* over the planning horizon—such as demand quantities, lead times, production yields, perishable inventories, backlog carry-overs, prices, etc. [23]. For notational convenience, they are combined into a large random uncertainty vector $\delta_t \in \Delta$. SCMPC can readily incorporate also forecasts on these uncertainties. In this case, δ_t represents the random deviation from the nominal forecast that is available. The distribution of δ_t is completely arbitrary—as long as a sufficient number K of independent samples (i.e., the *scenarios*) can be obtained in each time step t . In particular, there are no restrictions on the dimension of δ_t , nor on the type of distribution or on the correlations. The distribution of δ_t may also be time-varying and correlated in time.

SCMPC bases its decisions on a particular number K of sampled scenarios. Since the scenarios are sampled at random, SCMPC is essentially a *randomized algorithm*. A key advantage of SCMPC is its computational efficiency, which allows it to tackle stochastic problems with a complicated uncertainty structure and of a relatively large scale. The number of scenarios K can be shown to stand in a precise inverse relationship to the desired service level. Moreover, K turns out to be relatively small for most problems and does not grow with the horizon or problem dimensions. Compared to *multi-stage stochastic optimization*, SCMPC thus avoids the combinatorial complexity arising e.g., from the exhaustive sampling from a (binary) scenario tree [37]. Compared to *robust optimization*, SCMPC accounts for the probability of scenarios and needs to handle only a finite uncertainty set in the computations [2].

2. Supply Chain Model

The SCMPC method proposed in this paper can be used with a great variety of supply chain models [29]. Here a concrete case study is introduced for illustration purposes, which is comparable in size and type to the deterministic model in [25]. Its basic structure is depicted in Figure D.1.

This supply chain model is selected as a typical example of a variety of models that appear across several industries [14], such as consumer electronics [36], chemical products [24], or retail goods [33, Cha. 2]. The time period T over which the system has to be operated is very long, much longer than a sensible planning horizon.

2.1 Model Structure

The supply chain model has a tree-like structure with simple loops, three supply chain *stages* (or *echelons*) and multiple facilities per stage. The first echelon consists of $L = 3$ *suppliers* or *production facilities* $l \in \mathcal{L} := \{1, \dots, L\}$; the second echelon consists of




FIGURE D.1. Schematic overview of the exemplary supply chain model. The model includes $J = 2$ goods. There are a total number of $L = 3$ production facilities, $M = 5$ warehouses, and $N = 15$ retail stores. The arrows indicate the available shipment routes.

$M = 5$ warehouses or distribution centers $m \in \mathcal{M} := \{1, \dots, M\}$; and the third echelon consists of $N = 15$ customers or retail stores $n \in \mathcal{N} := \{1, \dots, N\}$.

Products can be shipped along the pre-defined routes, as indicated by arrows in Figure D.1. Here mainly downstream shipments are considered, with the additional possibility for a re-distribution of products from the retailers over the warehouse level. For notational convenience, the sets of routes

$$\mathcal{S}_{l,m} \subseteq \mathcal{L} \times \mathcal{M} , \quad \mathcal{S}_{m,m} \subseteq \mathcal{M} \times \mathcal{M} , \quad \mathcal{S}_{m,n} \subseteq \mathcal{M} \times \mathcal{N} , \quad \mathcal{S}_{n,m} \subseteq \mathcal{N} \times \mathcal{M} ,$$

are defined as subsets of ordered tuples of an origin and a target destination. The set of all shipment routes is denoted $\mathcal{S} := \mathcal{S}_{l,m} \cup \mathcal{S}_{m,m} \cup \mathcal{S}_{m,n} \cup \mathcal{S}_{n,m}$.

There are multiple $J = 2$ products $j \in \mathcal{J} := \{1, \dots, J\}$ that can be manufactured by the production facilities (or sourced from suppliers) $l \in \mathcal{L}$, at given production costs $\gamma_{j,l}$. Each production facility $l \in \mathcal{L}$ has a limited production capacity of C_l , of which one product $j \in \mathcal{J}$ consumes $c_{j,l}$. The finished products $j \in \mathcal{J}$ can then be shipped along the routes $s \in \mathcal{S}$, incurring shipment costs of $\sigma_{j,s}$. They can also be stored at each node, up to its storage capacity P_l (for production facilities $l \in \mathcal{L}$), W_m (for warehouses $m \in \mathcal{M}$), or R_n (for retail stores $n \in \mathcal{N}$), where each product $j \in \mathcal{J}$ has the storage coefficient $p_{j,l}$, $w_{j,m}$, or $r_{j,n}$, respectively. The storage cost for product $j \in \mathcal{J}$ amounts to $\phi_{j,l}$ (for production facilities $l \in \mathcal{L}$), $\omega_{j,m}$ (for warehouses $m \in \mathcal{M}$), or $\rho_{j,n}$ (for retail stores $n \in \mathcal{N}$) per time step.

Time is separated into discrete periods $t \in \mathcal{T} := \{1, \dots, T\}$. Before the end of each period $t \in \mathcal{T}$ and for each product $j \in \mathcal{J}$, first the decisions about the production $y_{t,j,l}$ at $l \in \mathcal{L}$ and the shipments $x_{t,j,s}$ along $s \in \mathcal{S}$ are taken. Shipped products leave the respective inventory $I_{t-1,j,l}$, $I_{t-1,j,m}$, or $I_{t-1,j,n}$ of their origin s_1 immediately.

The shipments $x_{t,j,s}$ arrive at their destination s_2 after a transportation lead time of

$\lambda_{t,s}$ for route $s \in \mathcal{S}$. The productions $y_{t,j,l}$ have a lead time $\lambda_{t,l}$ at $l \in \mathcal{L}$. Moreover, the production schedule must be determined T_{prod} time steps in advance; cf. [36]. This means that at each time t , the next T_{prod} production quantities for $t+1, \dots, t+T_{\text{prod}}$ are already fixed by previous decisions, while previous transportation decisions from all production facilities $l \in \mathcal{L}$ can still be changed.

After the shipments for the period $t \in \mathcal{T}$ are dispatched, the shipments and produced quantities from previous periods are received, and then the *demand quantities* $d_{t,j,n}$ for product $j \in \mathcal{J}$ at the retail store $n \in \mathcal{N}$ are observed. Any demand that is higher than the corresponding inventory level $I_{t,j,n}$ is backordered. However, more involved constructions of handling unmet demand, such as random backlog carry-over [36], could alternatively be incorporated.

2.2 Uncertainty and Forecasts

In this case study, the demand quantities $\delta_t := \{d_{t,j,n} \mid j \in \mathcal{J}, n \in \mathcal{N}\}$ are considered as the main source of uncertainty. In many practical cases, however, SCM has to cope also with different uncertainties; cf. [33] and [37]. They can be incorporated into the SCMPC approach under some conditions that are outlined below.

First, SCMPC is generally more suited for high-probability low-impact *operational risks*—such as demand quantities, lead times, production yields, perishable inventories, backlog carry-overs, prices etc.—rather than low-probability high-impact *disruption risks* [23]. Second, without loss of generality, joint *forecasts* for all uncertain variables in the form of scenarios are assumed to exist. They may be inexact [33], but nonetheless contain some information that is valuable for operational planning.

More precisely, let δ_t be the set of all uncertainties at time $t \in \mathcal{T}$ (i.e., not only the demands, in general). Let $\delta_{t+\tau|t-1}$ be the forecast for the uncertainty over the T_{plan} next time steps, $\tau = 1, \dots, T_{\text{plan}}$, conditional on the information available at time $t-1$ when the decisions for periods $t, t+1, \dots$ have to be made. Here $T_{\text{plan}} \ll T$ represents a rolling *planning horizon*, which is much shorter than the total *operation time* T for the supply chain.

ASSUMPTION D.1—FORECASTS At each time step t , $(t+1) \in \mathcal{T}$, a ‘sufficient number’ $k = 1, \dots, K$ of random samples $\delta_{t+\tau|t}^{(k)}$ of the uncertainty $\delta_{t+\tau|t}$ is available, over the entire planning horizon $\tau = 1, \dots, T_{\text{plan}}$. In particular, combining $\delta_{t+\tau|t}^{(k)}$ for all $k \in \{1, \dots, K\}$ together with $\delta_{t+\tau|t}$ forms a set of independent and identically distributed (i.i.d.) random variables, for all $\tau \in \{1, \dots, T_{\text{plan}}\}$. ■

For notational convenience, the forecasts over one planning horizon are combined into *multi-samples* (or *scenarios*):

$$\Omega_t^{(k)} := \{\delta_{t+1|t}^{(k)}, \delta_{t+2|t}^{(k)}, \dots, \delta_{t+T_{\text{plan}}|t}^{(k)}\} \quad \forall k \in \mathcal{K} := \{1, \dots, K\}.$$

Assumption D.1 poses no restriction on the joint distribution of each $\delta_{t+\tau|t}^{(k)}$. In fact,

it is quite plausible that future demands at different retail stores $n \in \mathcal{N}$ and between different products $j \in \mathcal{J}$ are (positively or negatively) correlated. Similarly, the forecasted uncertainties in the scenarios $\Omega_t^{(k)}$ can be correlated in time. Furthermore, the distributions of the forecasts are allowed to change over time; e.g., the distribution of $\delta_{t+2|t}$ need not be the same as that of $\delta_{t+2|t+1}$. In particular, the variance of forecasted demands may decrease due to new market information becoming available over time.

Assumption D.1 also does not require exact knowledge of the distributions of any future uncertainties. It suffices to have a ‘sufficient number’ of scenarios available (the exact number will be derived in Section 3). These scenarios may be obtained, for instance, from a stochastic prediction model or based on historical data.

2.3 Objective and Constraints

For SCM, the supply chain system has to be operated over the time steps $t \in \{1, \dots, T\}$ by a *control policy*, that is a function mapping all available information (past demands, forecasts, etc.) into a possible decision at time $t - 1$ (about production and shipments). The objective is to minimize the expected sum of stage costs

$$\underbrace{\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \sum_{l \in \mathcal{L}} \gamma_{j,l} y_{t,j,l}}_{\text{production costs}} + \underbrace{\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \sum_{s \in \mathcal{S}} \sigma_{t,j,s} x_{t,j,s}}_{\text{shipment costs}} + \underbrace{\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \left(\sum_{l \in \mathcal{L}} \phi_{j,l} I_{t,j,l} + \sum_{m \in \mathcal{M}} \omega_{j,l} I_{t,j,m} + \sum_{n \in \mathcal{N}} \rho_{j,l} I_{t,j,n} \right)}_{\text{inventory holding costs}} . \quad (\text{D.1})$$

The production and shipments costs result directly from the decision variables. For the holding costs, the inventory quantities in (D.1) are recursively computed as

$$I_{t,j,l} = I_{t-1,j,l} - \underbrace{\sum_{s \in \mathcal{S}, s_1=l} x_{t,j,s}}_{\text{shipments outbound}} + \underbrace{y_{t-\lambda_{t,j,l}}}_{\text{production}} \quad \forall l \in \mathcal{L} , \quad (\text{D.2a})$$

$$I_{t,j,m} = I_{t-1,j,m} - \underbrace{\sum_{s \in \mathcal{S}, s_1=m} x_{t,j,s}}_{\text{shipments outbound}} + \underbrace{\sum_{s \in \mathcal{S}, s_2=m} x_{t-\lambda_{t,s,j},s}}_{\text{shipments inbound}} \quad \forall m \in \mathcal{M} , \quad (\text{D.2b})$$

$$I_{t,j,n} = I_{t-1,j,n} - \underbrace{\sum_{s \in \mathcal{S}, s_1=n} x_{t,j,s}}_{\text{shipments outbound}} + \underbrace{\sum_{s \in \mathcal{S}, s_2=n} x_{t-\lambda_{t,s,j},s}}_{\text{shipments inbound}} - \underbrace{d_{t,j,n}}_{\text{demand}} \quad \forall n \in \mathcal{N} , \quad (\text{D.2c})$$

for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$.

The control policy has to respect constraints on the production quantities and the

shipment quantities

$$0 \leq y_{t,j,l} , \quad \sum_{j \in \mathcal{J}} c_j y_{t,j,l} \leq C_l \quad \forall l \in \mathcal{L} , \quad (\text{D.3a})$$

$$0 \leq x_{t,j,s} \quad \forall s \in \mathcal{S} , \quad (\text{D.3b})$$

for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$. Analogous to the maximum production capacities in (D.3a), constraints on the maximum shipment quantities in (D.3b) could be incorporated. However, they are not considered for this case study.

Finally, the following constraints on the inventory holdings at the production sites and the warehouses must be observed,

$$0 \leq I_{t,j,l} , \quad \sum_{j \in \mathcal{J}} p_j I_{t,j,l} \leq P_l \quad \forall l \in \mathcal{L} , \quad (\text{D.4a})$$

$$0 \leq I_{t,j,m} , \quad \sum_{j \in \mathcal{J}} w_j I_{t,j,m} \leq W_m \quad \forall m \in \mathcal{M} , \quad (\text{D.4b})$$

$$\sum_{j \in \mathcal{J}} r_j I_{t,j,n} \leq R_n \quad \forall n \in \mathcal{N} , \quad (\text{D.4c})$$

for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$. Despite of the uncertainty about the inventory levels of the retail stores, the capacity constraints can be kept by careful order placements.

Unless reliable upper bounds on the demands $d_{t,j,n}$ are known, a robust satisfaction of the demands at all times cannot be ensured. Moreover, in practice, such a policy is often too *conservative* [29], as it entails very high inventory holding costs. Instead, constraints on minimum *service levels* of the supply chain shall be imposed:

$$\frac{1}{T} \sum_{t \in \mathcal{T}} \mathbf{1}(I_{t,j,n} < 0) \leq \varepsilon_{j,n} \quad \forall j \in \mathcal{J}, n \in \mathcal{N} . \quad (\text{D.5})$$

Here $\mathbf{1}(\cdot)$ denotes an indicator function that equals to 1 if its argument condition is true and to 0 if it is false. The values of $\varepsilon_{j,n} \in (0, 1)$ represent the prescribed service levels of product $j \in \mathcal{J}$ at store $n \in \mathcal{N}$ over the whole operation time $t \in \mathcal{T}$.

Note that (D.5) essentially corresponds to the concept of an α -service level [3], [4], except that it is considered as the *time-average* over a long period T rather than a forward-looking *probability*.

3. Scenario-Based Model Predictive Control

The generic supply chain model described in Section 2 can be considered as a linear control system; see [29]. The level of all inventories at time $t \in \mathcal{T}$ represent the *state*

of the system, the decision variables the *input* of the system, and the inventory update equations (D.2) the (uncertain) *dynamics*.

3.1 Rolling Horizon Planning Policy

The SCMPC policy operates the supply chain in a rolling horizon planning framework, as illustrated in Figure D.2. Besides the *planning horizon* T_{plan} , it involves a *decision horizon* $T_{\text{dec}} \ll T_{\text{plan}}$ (usually $T_{\text{dec}} = 1$) and a *fixed production schedule* $T_{\text{prod}} \ll T_{\text{plan}}$.

In each *decision step* $t = (i - 1)T_{\text{dec}}$ where $i \in \{1, 2, \dots\}$, it computes the optimal decisions for the next T_{plan} time steps, which satisfy the constraints and guarantee demand satisfaction under all scenarios $k \in \mathcal{K}$.

The corresponding Rolling Horizon Scenario Program (RHSCP) reads as follows:

$$\min \frac{1}{K} \sum_{k \in \mathcal{K}} \left[\sum_{\tau=1}^{T_{\text{plan}}} \sum_{j \in \mathcal{J}} \sum_{l \in \mathcal{L}} \gamma_{j,l} y_{t+\tau|t,j,l} + \sum_{\tau=1}^{T_{\text{plan}}} \sum_{j \in \mathcal{J}} \sum_{s \in \mathcal{S}} \sigma_{t+\tau,j,s} x_{t+\tau|t,j,s} + \right. \\ \left. + \sum_{\tau=1}^{T_{\text{plan}}} \sum_{j \in \mathcal{J}} \left(\sum_{l \in \mathcal{L}} \phi_{j,l} I_{t+\tau|t,j,l} + \sum_{m \in \mathcal{M}} \omega_{j,l} I_{t+\tau|t,j,m} + \sum_{n \in \mathcal{N}} \rho_{j,l} I_{t+\tau|t,j,n}^{(k)} \right) \right], \quad (\text{D.6a})$$

$$\text{s.t.} \quad I_{t+\tau|t,j,l} = I_{t+\tau-1|t,j,l} - \sum_{s \in \mathcal{S}, s_1=l} x_{t+\tau|t,j,s} + y_{t+\tau-\lambda_t|t,j,l} \quad \forall l \in \mathcal{L}, \quad (\text{D.6b})$$

$$I_{t+\tau|t,j,m} = I_{t+\tau-1|t,j,m} - \sum_{s \in \mathcal{S}, s_1=m} x_{t+\tau|t,j,s} + \sum_{s \in \mathcal{S}, s_2=m} x_{t+\tau-\lambda_{t,s}|t,j,s} \quad \forall m \in \mathcal{M}, \quad (\text{D.6c})$$

$$I_{t+\tau|t,j,n}^{(k)} = I_{t+\tau-1|t,j,n}^{(k)} - \sum_{s \in \mathcal{S}, s_1=n} x_{t+\tau|t,j,s} + \sum_{s \in \mathcal{S}, s_2=n} x_{t+\tau-\lambda_{t,s}|t,j,s} - d_{t+\tau|t,j,n}^{(k)} \quad \forall n \in \mathcal{N}, \quad (\text{D.6d})$$

$$0 \leq y_{t+\tau|t,j,l}, \quad \sum_{j \in \mathcal{J}} c_j y_{t+\tau|t,j,l} \leq C_l \quad \forall l \in \mathcal{L}, \quad (\text{D.6e})$$

$$0 \leq x_{t+\tau|t,j,s} \quad \forall s \in \mathcal{S}, \quad (\text{D.6f})$$

$$0 \leq I_{t+\tau|t,j,l}, \quad \sum_{j \in \mathcal{J}} p_j I_{t+\tau|t,j,l} \leq P_l \quad \forall l \in \mathcal{L}, \quad (\text{D.6g})$$

$$0 \leq I_{t+\tau|t,j,m}, \quad \sum_{j \in \mathcal{J}} w_j I_{t+\tau|t,j,m} \leq W_m \quad \forall m \in \mathcal{M}, \quad (\text{D.6h})$$

$$0 \leq I_{t+\tau|t,j,n}^{(k)}, \quad \sum_{j \in \mathcal{J}} r_j I_{t+\tau|t,j,n}^{(k)} \leq R_n \quad \forall n \in \mathcal{N}, \quad (\text{D.6i})$$

where the constraints (D.6b-i) must hold for all $k \in \mathcal{K}$, $j \in \mathcal{J}$, and $\tau \in \{1, \dots, T_{\text{plan}}\}$. The

minimization refers to the decisions of shipments and productions,

$$x_{t+\tau|t,j,s} \quad \text{and} \quad y_{t+\tau|t,j,l} \quad \forall \tau \in \{1, \dots, T_{\text{plan}}\}, j \in \mathcal{J}, s \in \mathcal{S}, l \in \mathcal{L}, \quad (\text{D.7})$$

except that the productions $y_{t+\tau|t,j,s}$ for $\tau = 1, \dots, T_{\text{prod}}$ are already fixed. Moreover, the initial inventory levels $I_{t|t,j,l}$, $I_{t|t,j,m}$, $I_{t|t,j,n}^{(k)}$ are given from previous steps.

In (D.6), the $k = 1, \dots, K$ uncertainty scenarios $\Omega_t^{(k)}$ have been substituted to obtain a deterministic and convex optimization program. The cost function (D.6a) intends to approximate the expectation in (D.1) by taking the average costs over all K scenarios.

In this particular case study, the demand scenarios $d_{t+\tau|t,j,n}^{(k)}$ lead to scenario-specific inventory levels $I_{t+\tau|t,j,n}^{(k)}$ for the retail stores. An artificial scenario $K + 1$, in which all demands $d_{t+\tau|t,j,n}^{(K+1)}$ are zero, can be introduced to ensure the satisfaction of the retail storage capacity constraints with certainty. If other uncertainties are present—such as lead times, production yields, perishable inventories, or backlog carry-overs—they must also be included in the RHSCP, by adding a scenario-specific superscript to all scenario-dependent variables.




FIGURE D.2. Rolling horizon planning by SCMPC. The lower bars represent the transportation decisions and the upper bars the production decisions. The crosshatched area is the fixed production schedule (over period T_{prod}). The widely hatched decisions are computed only for planning purposes (over the planning horizon T_{plan}). The densely hatched decisions (over the decision horizon T_{dec}) are actually implemented.

The overall algorithm of SCMPC can be summarized as follows.

For each decision step $t := (i - 1)T_{\text{dec}}$ where $i \in \{1, 2, \dots\}$, perform the following steps:

1. Extract K scenarios $\Omega_t^{(1)}, \dots, \Omega_t^{(K)}$ of the uncertainty over the planning horizon.
2. Substitute these scenarios into the RHSCP, and solve for the optimal shipment decisions $x_{t+\tau|t,j,s}$, where $\tau = 1, \dots, T_{\text{plan}}$, and production decisions $y_{t+\tau|t,j,s}$, where $\tau = T_{\text{prod}} + 1, \dots, T_{\text{plan}}$.
3. Implement these decisions only for $\tau = 1, \dots, T_{\text{dec}}$ (shipments) and $\tau = T_{\text{prod}} + 1, \dots, T_{\text{prod}} + T_{\text{dec}}$ (production), respectively.

As the number K of scenarios increases, the operation of the supply chain by SCMPC becomes more *conservative*—i.e., the service levels are expected to improve while the total storage costs are expected to increase. A quantitative analysis of this intuitive relationship is provided below. In particular, a lower bound on K is derived such that the desired service level constraints (D.5) are satisfied. In order to minimize the long-term operational costs, and also to keep the computational efforts low, K should be selected equal to that lower bound.

REMARK D.1—FEEDBACK POLICY In the RHSCP, the future decisions $x_{t+\tau|t,j,s}$ for $\tau \in \{2, \dots, T_{\text{plan}}\}$ and $y_{t+\tau|t,j,l}$ for $\tau \in \{T_{\text{prod}} + 1, \dots, T_{\text{plan}}\}$ are assumed as ‘here-and-now’ decisions. They may also be implemented as ‘wait-and-see’ decisions; i.e., by reacting to the information that becomes available until time $t + \tau - 1$, see [35]. In fact, they can be modeled as the sum of (arbitrary) basis functions of the past uncertainties $\delta_{t+1}, \dots, \delta_{t+\tau-1}$, see [29, Rem. 7]. ■

3.2 Long-Term Properties of SCMPC

The subsequent analysis examines the behavior of the SCMPC policy over the entire operation time $t \in \mathcal{T}$. The theory draws heavily on the results of [8] and [30], as well as on the SCMPC approach presented in [29]. Note that the mathematical statements are kept brief, and the reader is referred to the previous work for further details. The following basic assumption is made throughout.

ASSUMPTION D.2—RESOLVABILITY The RHSCP remains a feasible optimization program at each decision step $t = (i - 1)T_{\text{dec}}$, $i \in \{1, 2, \dots\}$ [29, Ass. 5]. ■

Assumption D.2 appears to be restrictive from a theoretical point of view. However, it is quite reasonable for many practical applications. It presumes that the supply chain

is ‘well designed’ with regards to the available capacities. Moreover, it requires that a ‘back-up decision rule’ is available for cases when the demands cannot be fulfilled; e.g., by replacing the lower inventory constraints with high penalties on stockout quantities. This *soft constrained* approach shall be used for the simulation results in Section 4 [26].

DEFINITION D.1—SCENARIO PROGRAM (a) An *uncertain convex program* is a general optimization program with one (or multiple) uncertain constraint(s), that becomes convex if any possible value of the uncertainty were fixed. (b) The *scenario program* derives from the uncertain convex program by replacing the uncertain constraint with a finite number of sampled versions of this constraint, which are obtained by substituting fixed samples of the uncertainty [8, Sec. 1]. ■

Observe that the RHSCP is a scenario program, with the retail inventory levels (D.6i) being the sampled constraints. [8] have shown that the solution of a scenario program (*scenario solution*) has certain generalization properties with respect to the original uncertain convex program. Only some of these aspects, however, are relevant for the theory of this paper.

DEFINITION D.2—SUPPORT RANK (a) The *unconstrained subspace* of a sampled constraint in a scenario program is the largest linear subspace (in a set inclusion sense) of the decision space that remains unconstrained by the sampled instances of this constraint. (b) The *support rank* of a sampled constraint in a scenario program is the co-dimension of its unconstrained subspace [30, Def. 3.6]. ■

LEMMA D.1—SUPPORT RANK Let $\zeta_{t+\tau|t,j,n}$ be the support rank of the sampled retail inventory constraint (D.6i) of the RHSCP, for $\tau \in \{1, \dots, T_{\text{dec}}\}$ and any $t = (i-1)T_{\text{dec}}$, $i \in \{1, 2, \dots\}$, $j \in \mathcal{J}$, $n \in \mathcal{N}$. Then, for any value of these indices, $\zeta_{t+\tau|t,j,n} = 1$. ■

Proof. Consider the sampled retail inventory constraint (D.6i) of the RHSCP, for a fixed $\tau \in \{1, \dots, T_{\text{dec}}\}$ and any $t = (i-1)T_{\text{dec}}$, $i \in \{1, 2, \dots\}$, $j \in \mathcal{J}$, $n \in \mathcal{N}$. Of all the decision variables (D.7), the constraint directly affects only the subsets

$$\{-x_{t+\alpha|t,j,s} \mid \alpha \leq \tau, s_1 = n\} \quad \text{and} \quad \{x_{t+\alpha|t,j,s} \mid \alpha + \lambda_{t+\alpha,s} \leq \tau, s_2 = n\} . \quad (\text{D.8})$$

Moreover, the sampled constraints require that the sum of all decision variables in (D.8) must be greater than or equal to the quantities

$$\underbrace{\sum_{i=1}^{\tau} d_{t+i|t,j,n}^{(k)}}_{\text{sum of demands}} - \underbrace{I_{t|t,j,n}^{(k)}}_{\text{initial inventory}} \quad \text{for all samples } k \in \mathcal{K} . \quad (\text{D.9})$$

For all $k \in \mathcal{K}$, the sum of all decision variables in (D.8) is equivalent to a vector product of the entire decision vector with a fixed vector (of zeros, ones, and negative ones). This fixed vector has a null space of co-dimension one [30, Ex. 3.5], and hence the result follows. □

DEFINITION D.3—VIOLATION PROBABILITY Let $t = (i-1)T_{\text{dec}}$, $i \in \{1, 2, \dots\}$, $j \in \mathcal{J}$, $n \in \mathcal{N}$ be fixed. For any $\tau \in \{1, \dots, T_{\text{dec}}\}$, the *violation probability* $V_{t+\tau|t,j,n} \in [0, 1]$ is the probability with the inventory at $t + \tau$ becomes negative if the scenario solution of the RHSCP at time step t is applied to the supply chain [8, Def. 1.1]. ■

For the first part of the analysis, fix an arbitrary state of the supply chain at time t ; see Figure D.3. Note that each violation probability $V_{t+\tau|t,j,n}$ is then a function of the scenarios $\Omega_t^{(k)}$ for $k \in \mathcal{K}$ extracted at time t . Since the scenarios are random, $V_{t+\tau|t,j,n}$ is itself a random variable, see [29, Sec. 4.2].




FIGURE D.3. Overview of theoretical analysis. Theorem 1 considers each decision period individually for an arbitrary state of the supply chain. Theorem 2 concerns a sequence of decision periods.

THEOREM D.1—VIOLATION PROBABILITY Let Assumptions D.1 and D.2 hold and an arbitrary state of the supply chain at time $t = (i-1)T_{\text{dec}}$, where $i \in \{1, 2, \dots\}$, be fixed. Then for any $j \in \mathcal{J}$, $n \in \mathcal{N}$, and $\tau \in \{1, \dots, T_{\text{dec}}\}$ the random violation probability satisfies

$$\mathbf{E}_{\Omega_t^{(k)}}[V_{t+\tau|t,j,n}] \leq \frac{\zeta_{t+\tau|t,j,n}}{K+1} = \frac{1}{K+1} , \quad (\text{D.10})$$

where the expectation $\mathbf{E}_{\Omega_t^{(k)}}[\cdot]$ refers to the random scenarios $\Omega_t^{(k)}$, $k \in \mathcal{K}$. ■

Proof. By Lemma D.1, for any $t = (i-1)T_{\text{dec}}$, $i \in \{1, 2, \dots\}$, $j \in \mathcal{J}$, $n \in \mathcal{N}$, $\tau \in \{1, \dots, T_{\text{dec}}\}$ the support rank of the corresponding sampled constraint in (D.6i) is $\zeta_{t+\tau|t,j,n} = 1$. By [30, Thm. 4.1], the distribution of $V_{t+\tau|t,j,n}$ is therefore upper bounded by

$$\mathbf{P}_{\Omega_t^{(k)}}[V_{t+\tau|t,j,n} \geq \nu] \leq B(\nu; 0, K) , \quad (\text{D.11})$$

where $B(\cdot; \cdot, \cdot)$ denotes the *beta distribution* [1, Sec. 26.5.3, 26.5.7]. An integration of this distribution yields the desired result. □

As a result of Theorem D.1, the lower bound on K that is required for the service level constraint (D.5) is given by

$$\frac{1}{K+1} \leq \varepsilon \quad \Longleftrightarrow \quad K \geq \frac{1}{\varepsilon} - 1 . \quad (\text{D.12})$$

For the remainder of the section, it is therefore assumed that the sample size K is chosen according to (D.12).

After the scenarios $\Omega_t^{(k)}$ at $t = (i-1)T_{\text{dec}}$ are fixed, the actual violations $\mathbf{1}(I_{t+\tau,j,n} < 0)$ over the decision horizon $\tau = 1, \dots, T_{\text{dec}}$ are random variables with (now fixed) parameters $V_{t+\tau|t,j,n}$, respectively. Their outcome is determined by the actual uncertainties $\delta_{t+\tau}$ over $\tau = 1, \dots, T_{\text{dec}}$ and they satisfy

$$\mathbf{E}_{\Omega_t^{(k)}} [\mathbf{E}_{\delta_{t+\tau}} [\mathbf{1}(I_{t+\tau,j,n} < 0)]] = \underbrace{\mathbf{E}_{\Omega_t^{(k)}} [V_{t+\tau|t,j,n}]}_{\leq \varepsilon}, \quad (\text{D.13})$$

where $\mathbf{E}_{\delta_{t+\tau}}[\cdot]$ refers to the expectation with respect to the uncertainties $\delta_{t+1}, \dots, \delta_{t+T_{\text{dec}}}$. Note that (D.13) holds because $\mathbf{1}(I_{t+\tau,j,n} < 0)$ is a generic binomial random variable with parameter $V_{t+\tau|t,j,n}$, once $V_{t+\tau|t,j,n}$ is fixed.

In general, the random variables $\mathbf{1}(I_{t+\tau,j,n} < 0)$ for $\tau \in \{1, \dots, T_{\text{dec}}\}$, $j \in \mathcal{J}$, $n \in \mathcal{N}$ are correlated. Nonetheless, by linearity of expectation, they satisfy

$$\begin{aligned} \mathbf{E}_{\Omega_t^{(k)}} [\mathbf{E}_{\delta_{t+\tau}} [\mathbf{1}(I_{t+1,j,n} < 0) + \dots + \mathbf{1}(I_{t+T_{\text{dec}},j,n} < 0)]] \\ = \mathbf{E}_{\Omega_t^{(k)}} [V_{t+1|t,j,n}] + \dots + \mathbf{E}_{\Omega_t^{(k)}} [V_{t+T_{\text{dec}}|t,j,n}] \\ \leq T_{\text{dec}} \cdot \varepsilon. \end{aligned} \quad (\text{D.14})$$

With the above observations, it is possible to analyze the sequence of identical decision periods, see Figure D.3. Hence the following fundamental result for the long-term properties of SCMP is obtained.

THEOREM D.2—CONSTRAINT VIOLATIONS Let Assumptions D.1 and D.2 hold and consider the supply chain over a sequence $i \in \{1, 2, \dots, T/T_{\text{dec}}\}$ of decision periods. Then

$$\limsup_{i \rightarrow \infty} \frac{1}{i} \left(\mathbf{1}(I_{t+1,j,n} < 0) + \dots + \mathbf{1}(I_{t+T_{\text{dec}},j,n} < 0) \right) \leq T_{\text{dec}} \cdot \varepsilon \quad (\text{D.15})$$

for any $j \in \mathcal{J}$ and $n \in \mathcal{N}$, almost surely. ■

Note that Theorem D.2 is essentially the same as

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t \in \mathcal{T}} \mathbf{1}(I_{t,j,n} < 0) \leq \varepsilon \quad (\text{D.16})$$

almost surely, except for the minor requirement that the time length T be divisible by T_{dec} . The limit supremum in (D.15) and (D.16) refers to the sequence of *time-average constraint violations*. Therefore (D.16) means satisfaction of the service level constraint (D.5) in the strongest sense possible: namely, SCMP keeps the service level of $1 - \varepsilon$ almost surely as time goes to infinity.

Proof. The violations of each decision period $i \in \{1, 2, \dots, T/T_{\text{dec}}\}$ are generally dependent on the decisions and uncertainty outcomes of the previous time periods. However, they are independent conditionally on a fixed state of the system in the time period $(i - 1)T_{\text{dec}}$, respectively.

Theorem D.1 holds identically for all states at the beginning of each time period $(i - 1)T_{\text{dec}}$, under Assumptions D.1 and D.2. Therefore, the actual violations (D.14) over the decision periods $i = 1, 2, \dots$ form a sequence of martingale differences [17]. The result now follows by noting that the left-hand side of (D.15) is the sum of martingale differences, see [29, Lem. 16 and Thm. 17]. \square

REMARK D.2 In general, the service levels $\varepsilon_{j,n} \in (0, 1)$ can be chosen individually for each product $j \in \mathcal{J}$ and retail store $n \in \mathcal{N}$. However, then individual sample sizes $K_{j,n}$ must be selected accordingly, see [29]. Henceforth, a universal service level $\varepsilon_{j,n} = \varepsilon$ is assumed in this paper, for all $j \in \mathcal{J}$ and $n \in \mathcal{N}$. \blacksquare

4. Case Study

This case study considers the supply chain model introduced in Section 2 under the SCMPC algorithm described in Section 3. First, the numerical parameters of the model are specified. Thereafter, the results of the simulation experiments are presented and discussed.

4.1 Numerical Specifications


The numerical values of the simulation model are given in Table D.1. The costs are assumed to be in any unit currency, and the time steps $t \in \mathcal{T}$ can be considered to be weeks. The planning horizon is $T_{\text{plan}} = 26$ weeks, the decision horizon is $T_{\text{dec}} = 1$ week, and the production schedule is fixed $T_{\text{prod}} = 4$ weeks in advance. The simulations are run for a total operation time of two years ($T = 104$ weeks).

For practical purposes, the sampled constraints (D.6i) are implemented as *soft constraints* [26]. This means that additional slack variables are introduced, which are heavily penalized in the objective function so that they become non-zero only if the problem is otherwise infeasible. This construction essentially guarantees Assumption D.2.

The demand quantities $d_{t,1,n}$ for product $j = 1$ are modeled as the sum of a constant, a seasonal, and a random component. The random component is given by an autoregressive model with stationary, uniformly distributed innovations. In contrast, product $j = 2$ is assumed to have a finite life cycle, such that the demand quantities $d_{t,2,n}$ peak around $t = 40$ to $t = 60$. However, the size and the time of the peak are uncertain. For an illustration of the demand patterns, Figure D.4 depicts a few random sample paths.

DESCRIPTION	PARAMETER
production capacities	$C_1 = 60,000, C_2 = 80,000, C_3 = 60,000$
production capacity coefficients	$c_{1,l} = 1, c_{2,l} = 2 \forall l \in \mathcal{L}$
storage capacities	$P_l = 100,000 \forall l \in \mathcal{L}$ $W_m = 500,000 \forall m \in \mathcal{M}$ $R_n = 40,000 \forall n \in \mathcal{N}$
storage capacity coefficients	$p_{1,l} = 1, p_{2,l} = 2 \forall l \in \mathcal{L}$ $w_{1,m} = 1, w_{2,m} = 2 \forall m \in \mathcal{M}$ $r_{1,n} = 1, r_{2,n} = 2 \forall n \in \mathcal{N}$
production costs	$\gamma_{1,l} = 5, \gamma_{2,l} = 8 \forall l \in \mathcal{L}$
shipment costs	$\sigma_{1,s} = 0.2, \sigma_{2,s} = 0.3 \forall s \in \mathcal{S}_{l,m} \cup \mathcal{S}_{m,m}$ $\sigma_{1,s} = 0.3, \sigma_{2,s} = 0.4 \forall s \in \mathcal{S}_{m,n} \cup \mathcal{S}_{n,m}$
storage costs	$\phi_{1,l} = 0.08, \phi_{2,l} = 0.2 \forall l \in \mathcal{L}$ $\omega_{1,m} = 0.04, \omega_{2,m} = 0.1 \forall m \in \mathcal{M}$ $\rho_{1,n} = 0.08, \rho_{2,n} = 0.2 \forall n \in \mathcal{N}$
production lead times	$\lambda_{t,l} = 1 \forall t \in \mathcal{T}, l \in \mathcal{L}$
transportation lead times	$\lambda_{t,s} = 1 \forall t \in \mathcal{T}, s \in \mathcal{S}$

TABLE D.1. Parameter values used in the numerical simulation.

FIGURE D.4. Random sample paths of the demands for product $j = 1, 2$.

4.2 Simulation Results

Production. Consider the case where the service level of $\varepsilon = 5\%$ is fixed. Figure D.5 depicts a typical trajectory for the productions by the facilities $l = 1, 2, 3$. Observe that during the weeks around the peak demand of product $j = 2$, its production (long dashed lines) utilizes almost all of the available capacities (solid lines).




FIGURE D.5. Production capacity utilization of production facilities $l = 1, 2, 3$ by product $j = 1$ (short dash), $j = 2$ (long dash), and total (solid line), for one simulation run.

Inventory. Figure D.6 depicts the inventories of the warehouses $m \in \mathcal{M}$ and retail stores $n \in \mathcal{N}$ during the simulation period. Observe that for product $j = 1$ the retail inventory levels are fairly constant, and for product $j = 2$ the retail inventory levels are distinctly higher during the peak demand period; see Figure D.6(c,d). The explanation is that the retail inventories serve as a buffer against the demand *uncertainty*, while they are unrelated to the demand *level*. In particular, for product $j = 1$ the demand uncertainty is constant over time, while there are seasonal fluctuations in the demand level. For product $j = 2$, the highest demand uncertainty is around the peak of the life cycle; cf. Figure D.4.

Some safety stocks are also kept on the warehouse level, where inventory holding costs are significantly lower; see Figure D.6(a,b). However, an additional peak can be observed in the inventory, in particular of product $j = 1$, around weeks 40 to 50. This is due the anticipated bottleneck in the production capacity during weeks 40 to 60; cf. Figure D.5. Since the inventory holding cost per production capacity coefficient is cheaper for product $j = 1$, the peak demand for product $j = 2$ induces a pre-production of product $j = 1$.

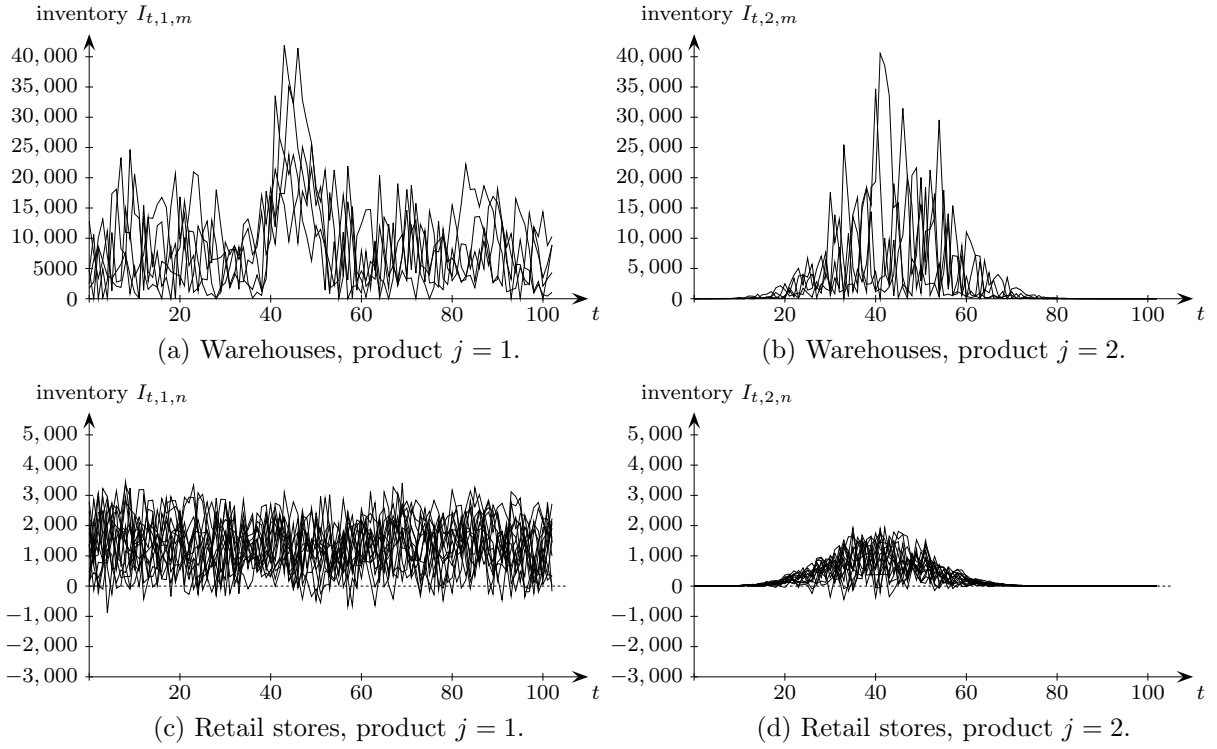


FIGURE D.6. Inventory of product $j = 1$ and $j = 2$ at the warehouses $m \in \mathcal{M}$ and the retail stores $n \in \mathcal{N}$ for one simulation run.

Service Levels. For the simulation run in Figure D.6, the average service levels of all retail stores $n \in \mathcal{N}$ (i.e., the fraction of time steps in with stockout) are $\bar{\varepsilon}_{1,n} = 4.68\%$ for product $j = 1$ and $\bar{\varepsilon}_{2,n} = 4.04\%$ for product $j = 2$. Hence they satisfy the desired service level of $\varepsilon_{1,n}, \varepsilon_{2,n} \leq 5\%$.

The results of a more extensive study is summarized in Table D.2. For a variety of sample sizes K and its corresponding guaranteed service levels $\varepsilon_{j,n}$, the empirical service levels $\bar{\varepsilon}_{j,n}$ are evaluated based on one hundred simulation runs. Observe that in all cases the theoretical services level bounds, according to Section 3.2, are fairly accurately matched by the empirical service levels.

Performance. A comparison of the performance of SCMPC to other approaches is not straightforward, because they are difficult to apply in this setting. Indeed, the uncertainty consists of 30 demand quantities $d_{t,j,n}$ (for 2 products at 15 stores) at each time step over the planning horizon $\tau = 1, \dots, 26$. Moreover, the demand quantities evolve in time by an autoregressive process—therefore, they are correlated in time and an explicit distribution is not available.

Stochastic optimization would have to handle a general multi-variate (30-dimensional), continuous, time-varying demand distribution at each stage $\tau = 1, \dots, 26$ that is conditionally dependent on the uncertainty of previous stages. Since this is generally impossible, an approximation by a much simpler stochastic model, e.g., a binary scenario tree,

SAMPLE SIZE	THEORETICAL SERVICE LEVEL		EMPIRICAL SERVICE LEVEL	
	$j = 1$	$j = 2$	$j = 1$	$j = 2$
$K = 4$	$\varepsilon_{1,n} \leq 20\%$	$\varepsilon_{2,n} \leq 20\%$	$\bar{\varepsilon}_{1,n} = 20.58\%$	$\bar{\varepsilon}_{2,n} = 19.87\%$
$K = 9$	$\varepsilon_{1,n} \leq 10\%$	$\varepsilon_{2,n} \leq 10\%$	$\bar{\varepsilon}_{1,n} = 10.37\%$	$\bar{\varepsilon}_{2,n} = 9.81\%$
$K = 19$	$\varepsilon_{1,n} \leq 5\%$	$\varepsilon_{2,n} \leq 5\%$	$\bar{\varepsilon}_{1,n} = 5.06\%$	$\bar{\varepsilon}_{2,n} = 4.78\%$
$K = 99$	$\varepsilon_{1,n} \leq 1\%$	$\varepsilon_{2,n} \leq 1\%$	$\bar{\varepsilon}_{1,n} = 1.03\%$	$\bar{\varepsilon}_{2,n} = 0.89\%$

TABLE D.2. Comparison of target service levels and empirical service levels, averaged over all retail stores $n = 1, \dots, 15$. The indicated values represent an average over one hundred simulation runs.

is required. Yet even this simple model would consist of $2^{26} \approx 67.1$ million scenarios, so the computations become intractable; cf. [37].

Robust optimization would have to (optimally) fit uncertainty sets for 30-dimensional demand quantities, in particular such that these sets contain $(1 - \varepsilon)$ of the multi-variate probability mass. In general, this problem is extremely hard, and it must be solved repeatedly for the time-varying distribution; cf. [2].

However, robust optimization can make a simple, conservative approximation (ROPT). Each uncertain demand quantity $d_{t,j,n}$ is independently set to its $(1 - \varepsilon)$ *quantile*; that is, the value that it will not exceed with a probability of more than ε . This procedure is simple and tractable, e.g., by a Monte Carlo simulation. However, it is conservative because, for instance, the $(1 - \varepsilon)$ quantile of the sum of demands (over some stores $n \in \mathcal{N}$ or periods $t \in \mathcal{T}$) is much lower than the sum of their $(1 - \varepsilon)$ quantiles (except if they are perfectly correlated).

Table D.3 compares the performances of SCMPC as compared to the simplified robust optimization scheme (ROPT), regarding the total transportation costs H_{tran} and total

SAMPLE SIZE	SERVICE LEVEL	SCMPC	ROPT
$K = 4$	$\varepsilon_{1,n} \leq 20\%$	$H_{\text{tran}} = 4.59$	$H_{\text{tran}} = 4.59$
		$H_{\text{stor}} = 0.44$	$H_{\text{stor}} = 1.90$
$K = 9$	$\varepsilon_{1,n} \leq 10\%$	$H_{\text{tran}} = 4.58$	$H_{\text{tran}} = 4.58$
		$H_{\text{stor}} = 0.58$	$H_{\text{stor}} = 2.56$
$K = 19$	$\varepsilon_{1,n} \leq 5\%$	$H_{\text{tran}} = 4.60$	$H_{\text{tran}} = 4.60$
		$H_{\text{stor}} = 0.67$	$H_{\text{stor}} = 2.96$
$K = 99$	$\varepsilon_{1,n} \leq 1\%$	$H_{\text{tran}} = 4.57$	$H_{\text{tran}} = 4.58$
		$H_{\text{stor}} = 0.87$	$H_{\text{stor}} = 3.26$

TABLE D.3. Comparison of total transportation costs H_{tran} and storage costs H_{stor} (in millions) for the SCMPC and ROPT policies. The indicated values represent an average over one hundred simulation runs.

storage costs H_{stor} . Observe that the transportation costs remain fairly constant between SCMPC and ROPT, and across all service levels. However, the storage costs rise for an increasing service level, both within SCMPC and ROPT. Finally, SCMPC shows significant savings in the storage cost compared to ROBT for all service levels.

5. Conclusion

This paper has presented a new scenario-based model predictive control (SCMPC) approach for the coordinated management of a large variety of supply chains. SCMPC can handle very general stochastic uncertainties in the system and guarantees a pre-specified customer service level. It is computationally efficient, as it requires only a few sample scenarios in each time step. Therefore, SCMPC may offer substantial advantages over existing approaches based on stochastic or robust optimization.

The effectiveness of the SCMPC approach has been demonstrated in a case study. It has been shown that the decisions of the method are reasonable and keep the desired service level constraints accurately, in the long run. Moreover, SCMPC has achieved substantial cost savings as compared to a computationally tractable robust optimization approach.

References

- [1] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions*. Dover Publications, New York, 9th edition, 1970.
- [2] D. Bertsimas and A. Thiele. A robust optimization approach to supply chain management. In D. Bienstock and G.Nemhauser, editors, *Integer Programming and Combinatorial Optimization*, volume 10, pages 86–100. Springer, Berlin et al., 2004.
- [3] G.R. Bitran and H.H. Yanasse. Deterministic approximations to stochastic production problems. *Operations Research*, 32(5):999–1018, 1984.
- [4] J.H. Bookbinder and J.-Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34(9):1096–1108, 1988.
- [5] S. Bose and J.F. Pekny. A model predictive framework for planning and scheduling problems: A case study of consumer goods supply chain. *Computers and Chemical Engineering*, 24:329–335, 2000.
- [6] M.W. Braun, D.E. Rivera, M.W. Carlyle, and K.G. Kempf. Application of Model Predictive Control to robust management of multiechelon demand networks in semiconductor manufacturing. *Simulation*, 79(3):139–156, 2003.

- [7] M.W. Braun, D.E. Rivera, M.E. Flores, M.W. Carlyle, and K.G. Kempf. A model predictive control framework for robust management of multi-product, multi-echelon demand networks. *Annual Reviews in Control*, 27:229–245, 2003.
- [8] M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19:1211–1230, 2008.
- [9] M. Chen and W. Wang. A linear programming model for integrated steel production and distribution planning. *International Journal of Operations & Production Management*, 17(6):592–610, 1997.
- [10] A.J. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Science*, 6(4):475–490, 1960.
- [11] A. Federgruen. Centralized planning models for multi-echelon inventory systems under uncertainty. In S.C. Graves, A.H.G. Rinnooy Kan, and P.H. Zipkin, editors, *Handbooks in Operations Research and Management Science*, volume 4, pages 133–173. Elsevier, Amsterdam et al., 1993.
- [12] D. Fu, C.M. Ionescu, E. Aghezzaf, and R. De Keyser. Decentralized and centralized model predictive control to reduce the bullwhip effect in supply chain management. *Computers & Industrial Engineering*, 73:21–31, 2014.
- [13] C.E. García, D.M. Prett, and M. Morari. Model predictive control: Theory and practice. *Automatica*, 23(3):335–348, 1989.
- [14] A.M. Geoffrion and R.F. Powers. Twenty years of strategic distribution system design: An evolutionary perspective. *Interfaces*, 25:105–127, 1995.
- [15] S.C. Graves and S.P. Willems. Supply chain design: Safety stock placement and supply chain configuration. In S.C. Graves and A.G. de Kok, editors, *Handbooks in Operations Research and Management Science*, volume 11, pages 95–132. Elsevier, Amsterdam et al., 2003.
- [16] S.C. Graves and S.P. Willems. Strategic inventory placement in supply chains: Non-stationary demand. *Manufacturing & Service Operations Management*, 10(2):278–287, 2008.
- [17] P. Hall and C.C. Heyde. *Martingale Limit Theory and Its Application*. Academic Press, New York et al., 1980.
- [18] H.M.S. Lababidi, M.A. Ahmed, I.M. Alatiqi, and A.F. Al-Enzi. Optimizing the supply chain of a petrochemical company under uncertain operating and economic conditions. *Industrial & Engineering Chemistry Research*, 43:43–63, 2004.
- [19] H.L. Lee, V. Padmanabhan, and S. Whang. The bullwhip effect in supply chains. *Sloan Management Review*, 38(3):93–102, 1997.
- [20] X. Li and T.E. Marlin. Robust supply chain performance via model predictive control. *Computers & Chemical Engineering*, 33:2134–2143, 2009.
- [21] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.

- [22] J. Mula, D. Peidro, M. Díaz-Madroñero, and E. Vicens. Mathematical programming models for supply chain production and transportation planning. *European Journal of Operations Research*, 204:377–390, 2010.
- [23] A. Oke and M. Gopalakrishnan. Managing disruptions in supply chains: A case study of a retail supply chain. *International Journal of Production Economics*, 118:168–174, 2009.
- [24] E. Perea, I. Grossmann, E. Ydstie, and T. Tahmassebi. Dynamic modeling and classical control theory for supply chain management. *Computers and Chemical Engineering*, 24:1143–1149, 2000.
- [25] E. Perea-López, B.E. Ydstie, and I.E. Grossmann. A model predictive control strategy for supply chain optimization. *Computers and Chemical Engineering*, 27:1201–1218, 2003.
- [26] S.J. Qin and T.A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11:733–764, 2003.
- [27] F. Rømo, A. Tomasgard, L. Hellemo, M. Fodstad, B.H. Eidesen, and B. Pedersen. Optimizing the norwegian natural gas production and transport. *Interfaces*, 39(1):46–56, 2009.
- [28] H. Sarimveis, P. Patrinos, C.D. Tarantilis, and C.T. Kiranoudis. Dynamic modeling and control of supply chain systems. *Computers & Operations Research*, 35:3530–3561, 2008.
- [29] G. Schildbach, L. Fagiano, C. Frei, and M. Morari. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, (under review).
- [30] G. Schildbach, L. Fagiano, and M. Morari. Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4):2479–2501, 2013.
- [31] T. Schoenmeyr and S.C. Graves. Strategic safety stocks in supply chains with capacity constraints. working paper, Massachusetts Institute of Technology, Cambridge (MA), United States of America, 2009.
- [32] T. Schoenmeyr and S.C. Graves. Strategic safety stocks in supply chains with evolving forecasts. *Manufacturing & Service Operations Management*, 11(4):657–673, 2009.
- [33] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi. *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*. McGraw-Hill, Boston et al., 2008.
- [34] K.F. Simpson. In-process inventories. *Operations Research*, 6(6):863–873, 1958.
- [35] J. Skaf and S. Boyd. Nonlinear q-design for convex stochastic control. *IEEE Transactions on Automatic Control*, 54(10):2426–2430, 2009.

- [36] M.S. Sohdi. Managing demand risk in tactical supply chain planning for a global consumer electronics company. *Production and Operations Management*, 14(1):69–79, 2005.
- [37] M.S. Sohdi and C.S. Tang. *Managing Supply Chain Risk*. Springer, New York et al., 2012.
- [38] K. Subramanian, J.B. Rawlings, C.T. Maravelias, J. Flores-Cerrillo, and L. Megan. Integration of control theory and scheduling methods for supply chain management. *Computers & Chemical Engineering*, 51:4–20, 2013.
- [39] I. Yildirim, B. Tan, and F. Karaesmen. A multiperiod stochastic production planning and sourcing problem with service level constraints. In G. Liberopoulos, C.T. Papadopoulos, B. Tan, J.M. Smith, and S.B. Gershwin, editors, *Stochastic Modeling of Manufacturing Systems*, pages 345–363. Springer, Berlin et al., 2006.

Part E

Risk-Averse Two-Stage Stochastic Programs

Paper IV

A Scenario Approach for Two-Level Stochastic Programs with Expected Shortfall Probability

Georg Schildbach · Manfred Morari

Abstract

This paper presents a novel approach for risk averse two-stage stochastic programming. Instead of a traditional mean-risk combination, this approach optimizes the stochastic objective function value with respect to a maximal shortfall probability ε . More precisely, the first stage decision is made such that the objective function value takes the best possible value that can be guaranteed, for after the second stage, with a probability of $(1 - \varepsilon)$. A computationally efficient approximation method is presented for this problem, based on recent results in scenario-based optimization. In particular, the method optimizes the first stage decision for the worst case of a finite number of random scenarios. A theoretical analysis then provides a generalization property of this scenario-based decision, with respect to its shortfall probability. Hence the proposed method is an intuitive way of including risk aversion into two-stage stochastic programming. Moreover, it is computationally efficient, as the sample size is typically low and independent of the dimensions of the uncertainty variable. The approach has a wide range of potential applications, as discussed in the paper, allowing also for binary variables in the first stage decision.

This manuscript has been submitted for publication to the International Journal of Production Economics.

©2014 by the authors.

1. Introduction

1.1 Two-Stage Stochastic Integer Program

Consider the standard formulation of a *two-stage stochastic program* with binary variables in the first stage,

$$\min_{x,y,z} \mathbf{E}[g(x,y,z,\delta)] \quad (\text{E.1a})$$

$$\text{s.t. } x \in \{0,1\}^l, \quad y \in \mathcal{Y}(x), \quad (\text{E.1b})$$

$$z \in \mathcal{Z}(x,y,\delta). \quad (\text{E.1c})$$

The abstract random variable $\delta \in \Delta$ includes all uncertain quantities affecting the optimization problem. The *first-stage* variables are split into $l \geq 0$ binary variables $x \in \{0,1\}^l$ and $m > 0$ continuous variables $y \in \mathcal{Y}(x) \subseteq \mathbb{R}^m$. They must be decided on “here-and-now”, that is *before* the value of δ is observed. For deciding on the $n \geq 0$ *second-stage* variables $z \in \mathcal{Z}(x,y,\delta) \subseteq \mathbb{R}^n$, one can “wait-and-see” until *after* the value of δ is observed. All variables count towards an *objective function* $g(x,y,z,\delta)$, which is also subject to the uncertainty δ and must be minimized. The following assumptions are made throughout.

ASSUMPTION E.1 (a) For any choice of $\bar{x} \in \{0,1\}^l$, the feasible set $\mathcal{Y}(\bar{x})$ of the first stage continuous variables is non-empty and convex. (b) For any choice of $\bar{x} \in \{0,1\}^l$, $\bar{y} \in \mathcal{Y}(\bar{x})$ and almost every $\bar{\delta} \in \Delta$, the second stage feasible set $\mathcal{Z}(\bar{x}, \bar{y}, \bar{\delta})$ is non-empty. (c) For any choice of $\bar{x} \in \{0,1\}^l$ and almost every $\bar{\delta} \in \Delta$, the second stage feasible constraint $z \in \mathcal{Z}(\bar{x}, y, \bar{\delta})$ is jointly convex in y and z . (d) For any choice of $\bar{x} \in \{0,1\}^l$ and almost every $\bar{\delta} \in \Delta$, the objective function $g(\bar{x}, y, z, \bar{\delta})$ is jointly convex in y and z . ■

Assumption E.1 is satisfied for many practical applications. Its two main aspects can be summarized as follows: problem (E.1) has a feasible solution under almost every scenario $\delta \in \Delta$, and if the uncertainty $\delta \in \Delta$ were known already in the first stage and the binary decision is fixed, then (E.1) becomes a convex optimization problem. Assumption E.1(b) is also known as the property of *relatively complete recourse* [26, Sec. 2.1.3]. Assumption E.1(c) means, in particular, that for any $\bar{x} \in \{0,1\}^l$, $\bar{\delta} \in \Delta$, $\bar{y}_1, \bar{y}_2 \in \mathcal{Y}(\bar{x})$, and $\lambda \in [0,1]$,

$$\bar{z}_1 \in \mathcal{Z}(\bar{x}, \bar{y}_1, \bar{\delta}), \bar{z}_2 \in \mathcal{Z}(\bar{x}, \bar{y}_2, \bar{\delta}) \implies \lambda \bar{z}_1 + (1-\lambda) \bar{z}_2 \in \mathcal{Z}(\bar{x}, \lambda \bar{y}_1 + (1-\lambda) \bar{y}_2, \bar{\delta}).$$

For the presented theory in its most general form, the objective function in (E.1a) needs not be linear, nor separable into additive terms for the first and the second stage variables; cf. [5, 15, 26]. Moreover, the constraint sets $\mathcal{Y}(x)$ and $\mathcal{Z}(x,y,\delta)$ in (E.1b,c) are not necessarily polytopes, nor is the problem required to have a *fixed recourse matrix*; e.g., [26, Sec. 2.1.3]. However, additional assumptions such as these may be required for

efficient computation, depending on the problem at hand. In particular, standard numerical solution techniques easily become intractable in high dimensions, unless tailored algorithms are available to exploit the problem structure [16, 23, 27]. The focus of this paper, however, is not on algorithmic development; instead the theory is presented in its most general form.

1.2 Risk Averse Formulation

The standard formulation of a two-stage stochastic program minimizes the expected value of the objective function, as in (E.1a). This may be justified in cases where the problem is solved repeatedly for many times, and the interest is in the average performance [5, 15, 26].

However, the fundamental results of von Neumann and Morgenstern [28] suggest that in many practical situations the utility lost from a less-than-average outcome is much more severe than the utility gained from a higher-than-average outcome (of the same magnitude). This fact has spurred significant interest in incorporating *risk aversion* into the objective function (E.1a), cf. [2, 13, 19, 24].

One way is to optimize the first-stage decision for the worst-case of all possible uncertainty scenarios [5, Sec. 2.9]. This leads into the field of robust optimization, which has been successful in some applications [4]. However, the approach has two potential drawbacks. First, for many applications it is overly conservative, because the worst case is both unrealistic and highly unlikely. Second, it requires the uncertainty set Δ to be bounded and precisely known.

The main avenue of research on stochastic optimization has therefore focused on *mean-risk models*, see e.g., [2]. The basic idea is to add a (*downside*) *risk measure* to the objective function, weighted by a parameter for adjusting the level of risk aversion. A variety of risk measures are available for this purpose [3], such as the *variance* or the *conditional value-at-risk* (CVaR); see Shapiro et al. [26, Sec. 6.2].

In this paper, a novel approach is proposed for minimizing the ε -quantile of the objective function. The goal is to find the best first-stage solution whose objective function value is exceeded with a probability of no more than ε . Here $\varepsilon \in (0, 1)$ a risk parameter, called the *shortfall probability*. In financial applications, it is also known as the $(1 - \varepsilon)$ *value-at-risk* (VaR) [12].

It should be emphasized that VaR is a popular risk measure in many practical applications [14]. However, it does not satisfy all the desirable properties of a *coherent risk measure*, as shown by Artzner et al. [3]. In particular, it lacks the property of convexity, which makes it difficult to handle by numerical optimization. The method presented in this paper provides an approximate solution to such problems in a computationally efficient manner.

1.3 Contributions and Outline

The presented method is based on the scenario-based optimization approach, or simply the “*scenario approach*”, as developed by Calafiore and Campi [7], Campi and Garatti

[9, 10], and Calafiore [8] for single-stage stochastic programs. The novel extension can therefore be referred to as the *two-stage scenario approach* (TSA).

The basic idea is to optimize the first-stage decision for the worst-case of a finite number K of sampled uncertainty scenarios. This decision shall be referred to as the *scenario solution*. A theoretical underpinning for this approach is presented, concerning the generalization properties of the scenario solution, in terms of its shortfall probability, by an appropriate choice of K .

The TSA offers a variety of practical advantages. It is an intuitive and computationally efficient way of finding risk-averse solutions to two-stage stochastic programs. The required sample size K is finite and relatively small, depending on the problem dimensions. In particular, K is independent of the dimensions of the uncertain quantity δ , whose probability measure \mathbf{P} and support set Δ can be completely arbitrary. In fact, \mathbf{P} and Δ need not even be known explicitly, as long as a sufficient number of independent samples of δ are available.

Moreover, the TSA has several theoretical extensions compared to the classic scenario approach [7–10]. It allows for the inclusion of scenario-dependent decision variables in the second stage of the optimization program. Furthermore, binary variables can be included in the first-stage decision, which represents a common feature for many practical applications.

The paper is organized as follows. The introduction in Section 1 is completed with a short review of potential applications. Section 2 introduces the basic concepts of the classic scenario approach, as far as they are needed by Section 3, where the new results on the TSA are derived. Section 4 demonstrates the application of the TSA for an illustrative example. Section 5 states the final conclusions.

1.4 Applications

The number of actual and potential applications of *two-stage stochastic programming* (TSP) is vast, and risk aversion is a common feature in many of these applications. A selection of potential applications is discussed below, as a motivation for the presented theory.

TSP is one option for modeling problems of *supply chain design*, as described e.g., by Santoso et al. [21]. The decision to be made here-and-now may include the location and procurement of machines (binary variables) or the capacity expansions of production and logistic facilities (continuous variables). Typically there is a substantial uncertainty about future demands and prices, and potentially about other factors such as resource or transportation capacities. The wait-and-see decision consists of processing and transportation of products to customers, based on the available capacities and in optimal fashion. The objective is to minimize the cost of supply, or to maximize the profits.

TSP can also be applied to problems of *network design*, for example private communication networks, as shown by Sen et al. [25]. The here-and-now decision concerns the capacity expansion of the links in the network, under uncertainty about the future demand for data transfer on each link of the network. The capacity extensions may have

fixed costs (leading to binary first-stage variables) and costs proportional to the added capacity (leading to continuous first-stage variables). The wait-and-see decision is about the optimal routing of data between the nodes of the expanded network, after the actual demands are observed. In [25], this is formulated as a multi-commodity network flow problem. The objective is to minimize the number of unserved data transfer requests.

TSP is also a potential tool for *disaster management*, as proposed by Rawls and Turnquist [20] and Noyan [19]. By the here-and-now decision the response locations (binary variables) for the placement of emergency supplies, such as food, shelter, medicine, etc. are determined, as well as their quantities (continuous variables). The uncertainty is about if, when, where, and in what extent a natural disaster, e.g., a hurricane, will occur. The wait-and-see decision is then about the optimal shipment of the supplies, given the particular emergency incidence. The optimization is with respect to the cost for disaster management, with penalties for not meeting the demands on time.

TSP can also be used for *airline revenue management*, as shown by Chen and Homem-de-Mello [11]. The here-and-now decision is to allocate the available seats on particular itineraries of aircrafts (continuous variables) to a variety of ticket categories (binary variables), ranging e.g., from very early bookings at cheap fares up to late bookings at expensive fares (*not* travel classes). The demands are revealed only in the second stage of the problem, where possible recourse actions can be taken. Note that airlines must solve this problem on a regular basis, and therefore risk aversion plays a minor role. However, Birge and Louveaux [5, p. 67] suggest that similar problems may occur for one-time events, such as a soccer championship.

2. A Review of the Classic Scenario Approach

This section reviews the existing work on the (classic) scenario approach that is relevant for the subsequent theory of this paper. More details are found, in particular, in the groundbreaking work of Calafiore and Campi [7], Campi and Garatti [9,10], and Calafiore [8].

2.1 The Scenario Program

The (classic) scenario approach, considers a single-stage uncertain program $\text{UP}[\varepsilon]$ of the form

$$\text{UP}[\varepsilon] : \quad \min_w \quad c^T w \quad (\text{E.2a})$$

$$\text{s.t.} \quad \mathbf{P}[f(w, \delta) \leq 0] \geq 1 - \varepsilon, \quad (\text{E.2b})$$

$$w \in \Omega. \quad (\text{E.2c})$$

Here w denotes the decision variable, which must be chosen optimally from a convex, compact domain $\Omega \subset \mathbb{R}^d$. The variable δ represents any uncertain quantity in $\text{UP}[\varepsilon]$,

whose support Δ is of an entirely generic nature (e.g., a vector space). The constraint (E.2b) is formulated as a *chance constraint* [5, 15, 26]. Hence it must be kept with a probability level of at least $1 - \varepsilon$, where $\varepsilon \in (0, 0.5)$. The following assumption about the nature of the uncertainty δ is made throughout.

ASSUMPTION E.2—UNCERTAINTY (a) There exists a probability measure \mathbf{P} on Δ ; i.e., δ is a random variable. (b) The measure \mathbf{P} , i.e., the distribution of δ , may be unknown, but a sufficient number of independent random samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$ are available. ■

Assumption E.2 is very general, since it requires only sufficient data as a knowledge of δ . Note that the formulation of $\text{UP}[\varepsilon]$ comprises all uncertain optimization programs that become convex if the uncertain quantity δ were fixed and known [8, 9]. Moreover, the approach also extends to problems with multiple chance constraints [22].

If the distribution of δ is known, then $\text{UP}[\varepsilon]$ represents a stochastic program. However, it remains difficult to solve in the general case [5, 15, 26]. The main reason is that the feasible set of a chance constraint is non-convex and hard to express in explicit terms, except for very special cases.

The scenario approach provides an approximation method to $\text{UP}[\varepsilon]$, based on the optimal solution to the *scenario program*

$$\text{SP}[\omega^{(K)}] : \quad \min_w \quad c^T w \quad (\text{E.3a})$$

$$\text{s.t.} \quad f(w, \delta^{(k)}) \leq 0 \quad \forall k = 1, 2, \dots, K, \quad (\text{E.3b})$$

$$w \in \Omega. \quad (\text{E.3c})$$

In $\text{SP}[\omega^{(K)}]$, the chance constraint of $\text{UP}[\varepsilon]$ has been replaced by K fixed constraints, namely by substituting the samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$ of the uncertainty into the constraint function $f(w, \cdot)$. For notational convenience, the samples are also denoted as a *multi-sample* $\omega^{(K)} := \{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}\}$. They can be interpreted as *training samples* or *scenarios* for the solution of $\text{SP}[\omega^{(K)}]$, which is called the *scenario solution* and denoted $w^*(\omega^{(K)})$. Existence and uniqueness of $w^*(\omega^{(K)})$ can be assumed without any loss of generality [9, Sec. 2 (5)].

In practical application, the scenario solution $w^*(\omega^{(K)})$ is obtained after the outcomes of the training samples $\omega^{(K)}$ are observed. Hence $\text{SP}[\omega^{(K)}]$ is a *deterministic convex optimization program* of a given type (e.g., a linear or quadratic program), for which efficient numerical algorithms exist, even in high dimensions; cf. [6, 18].

2.2 Theoretical Properties of the Scenario Solution

The theoretical properties of scenario solution involve a deep mathematical theory. It establishes a link between the sample size K and the probability of $w^*(\omega^{(K)})$ violating the chance constraints:

$$v(\omega^{(K)}) := \mathbf{P}[f(w^*(\omega^{(K)}), \delta) > 0] \quad (\text{E.4})$$

In particular, for the purpose of theoretical analysis, the scenario solution $w^*(\omega^{(K)})$ and the violation probability $v(\omega^{(K)})$ are considered as (unknown) functions of the random multi-sample $\omega^{(K)}$.

Hence there are two levels of probability present in the scenario approach. The first is introduced by the random training sample $\omega^{(K)}$, affecting the choice of $w^*(\omega^{(K)})$. The second is the random quantity δ , which determines whether $w^*(\omega^{(K)})$ actually satisfies the constraint.

To highlight the two probability levels more clearly, suppose for a moment that the multi-sample has already been observed. Let $\bar{\omega}^{(K)}$ denote its outcome, and $\bar{w} := w^*(\bar{\omega}^{(K)})$ the corresponding scenario solution. Then the *a posteriori violation probability* $\bar{v} := v(\bar{\omega}^{(K)})$ is a deterministic, albeit unknown, value in the interval $[0, 1]$:

$$\bar{v} := \mathbf{P}[f(\bar{w}, \delta) > 0] . \quad (\text{E.5})$$

Now suppose that the multi-sample has not yet been observed. Then the *a priori violation probability* $v(\omega^{(K)})$, as defined in (E.4), is itself a random variable on the probability space (Δ^K, \mathbf{P}^K) , where Δ^K and \mathbf{P}^K are the K -th product space of Δ and the K -th product measure of \mathbf{P} , respectively. Note that $v(\omega^{(K)})$ has support $[0, 1]$ and the following result holds for its distribution, according to Campi and Garatti [9, Thm. 2.4].

THEOREM E.1—DISTRIBUTION BOUND The distribution of the violation probability $v(\omega^{(K)})$ of $\text{SP}[\omega^{(K)}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(K)}) > \nu] \leq \text{B}(\nu; K, d) , \quad (\text{E.6})$$

for any $\nu \in [0, 1]$, where

$$\text{B}(\nu; K, d) := \sum_{j=0}^{d-1} \binom{K}{j} \nu^j (1 - \nu)^{K-j} \quad (\text{E.7})$$

denotes the Beta distribution function [1, pp. 26.5.3, 26.5.7], with parameters d (the dimension of the decision variable) and K (the sample size). ■

Given the upper bound for the cumulative distribution of $v(\omega^{(K)})$, it is possible to compute an upper bound on its expectation by integrating the distribution function:

$$\begin{aligned} \mathbf{E}[v(\omega^{(K)})] &= \int_0^1 \mathbf{P}^K[v(\omega^{(K)}) > \nu] d\nu \\ &\leq \int_0^1 \text{B}(\nu; K, d) d\nu = \frac{d}{K+1} . \end{aligned} \quad (\text{E.8})$$

This leads to a simple corollary for selecting the sample size K , based on the dimension d of the decision space and the probability level ε of the chance constraint.

COROLLARY E.1—PROBABILITY BOUND The scenario solution $w^*(\omega^{(K)})$ of $\text{SP}[\omega^{(K)}]$ satisfies the chance constraint (E.2b) in expectation, i.e., $\mathbf{E}^K[v(\omega^{(K)})] \leq \varepsilon$, if K is selected large enough such that

$$\frac{d}{K+1} \leq \varepsilon . \quad (\text{E.9})$$

■

In other words, for K selected as in (E.9), the violation probability $v(\omega^{(K)})$ is *expected* to be lower than ε , so chance constraint (E.2b) is satisfied *on average*.

In most cases, one is interested in the *minimal* sample size K that satisfies (E.9). The reason for that is twofold: First, for a higher K the probability level of chance constraint (E.2b) is even lower than ε ; hence the scenario solution is feasible, but *conservative*. Second, the computational complexity increases when solving $\text{SP}[\omega^{(K)}]$ with higher samples sizes.

2.3 A Posteriori Sample Removal

The bound of Corollary E.1 is tight in the sense that there exists a class of optimization problems (namely, those that are “fully-supported”) for which (E.9) is exact [9, Sec. 2.1,(1)]. Hence this bound on the sample size cannot be improved, in general. However, it may be desirable to reduce the dependence of the scenario solution on extreme outliers in the samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$.

To this end, it is possible to deliberately increase the sample size K above its minimal value of Corollary E.1, in exchange for being allowed to remove R samples *a posteriori* (i.e., after the sample values have been observed). The samples are assumed to be removed by a valid *removal procedure*, as defined below; cf. [10, Ass. 2.2].

Let $\omega^{(K,R)}$ denote the remaining samples, after R of the K samples have been removed. Consequently, $w^*(\omega^{(K,R)})$ is the scenario solution of $\text{SP}[\omega^{(K,R)}]$ after the sample removal, and $v(\omega^{(K,R)})$ is its violation probability, as in (E.4)

DEFINITION E.1—REMOVAL PROCEDURE A removal procedure is an algorithm that selects R of the K samples to be removed from $\omega^{(K)}$ a posteriori. The removal procedure is opportunistic in the sense that the scenario solution $w^*(\omega^{(K,R)})$ violates all of the removed constraints. ■

Particular removal procedures are based on optimal, greedy, and marginal constraint removal; see e.g., [8, Sec. 5.1] for more details. Similar to Theorem E.1, an upper bound on the distribution of $v(\omega^{(K,R)})$ can be established for the case of sample removal.

THEOREM E.2—DISTRIBUTION BOUND WITH SAMPLE REMOVAL Let $\omega^{(K,R)}$ be the remaining samples after applying a removal procedure to $\omega^{(K)}$. Then the distribution of

the violation probability $v(\omega^{(K,R)})$ of $\text{SP}[\omega^{(K,R)}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(K,R)}) > \nu] \leq u_d^{(K,R)}(\nu) , \quad (\text{E.10a})$$

$$u_d^{(K,R)}(\nu) := \min\left\{1, \binom{R+d-1}{d-1} \text{B}(\nu; K, R+d)\right\} , \quad (\text{E.10b})$$

for any $\nu \in ([0, 1])$, where $\text{B}(\cdot; \cdot, \cdot)$ denotes the Beta distribution function as defined in (E.7). ■

Note that the upper bound on the distribution (E.10b) equals to that of Campi and Garatti [10, Thm. 2.1], except that it is saturated at 1. The saturation is justified by the fact that $v(\omega^{(K,R)})$ is itself a probability can hence be no larger than 1.

As for the no-removal case, the expected violation probability will be lower than ε if the *sample-removal pair* (K, R) is selected appropriately.

COROLLARY E.2—PROBABILITY BOUND WITH SAMPLE REMOVAL The scenario solution $w^*(\omega^{(K,R)})$ of $\text{SP}[\omega^{(K,R)}]$ satisfies the chance constraint (E.2b) in expectation, i.e., $\mathbf{E}^K[v(\omega^{(K,R)})] \leq \varepsilon$, if K and R are selected such that

$$\int_0^1 u_d^{(K,R)}(\nu) d\nu \leq \varepsilon . \quad (\text{E.11})$$

■

While the bound without sample removal (E.9) has a nice explicit form, the bound with sample removal (E.11) comes as a one-dimensional integral. However, it can be solved efficiently by numerical integration, provided that values of K and R are given. In order to find appropriate values of K and R , the number of removed constraints R is usually fixed. Then K is computed by a bi-section procedure, solving (E.11) repeatedly for different values of K , observing that the left-hand side is monotonically decreasing with K . Alternatively, K can be fixed and a bi-section procedure can yield the corresponding value of R , observing that the left-hand side of (E.11) is monotonically increasing with R .

3. The Scenario Approach for Two-Stage Stochastic Programs

In this section, the theory of the (classic) scenario approach is extended to the two-stage scenario approach (TSA) for two-stage stochastic programs. It shall be used to find risk-averse solutions, in the sense that the expected shortfall probability of the objective value is bounded.

3.1 Two-Stage Uncertain Optimization Program

Consider a *two-stage uncertain program* of the general form

$$\text{TUP}[\varepsilon] : \min_{q,x,y,z} q, \quad (\text{E.12a})$$

$$\text{s.t.} \quad \mathbf{P}[g(x,y,z,\delta) \leq q] \geq 1 - \varepsilon, \quad (\text{E.12b})$$

$$x \in \{0,1\}^l, \quad y \in \mathcal{Y}(x), \quad (\text{E.12c})$$

$$z \in \mathcal{Z}(x,y,\delta). \quad (\text{E.12d})$$

Assumption E.2 continues to hold for the uncertainty δ . The goal of $\text{TUP}[\varepsilon]$ is to find the optimal decision for the first stage variables x and y with the best objective function value q , that has a shortfall probability of no more than ε . Hence q is the ε -quantile of the (random) two-stage optimal cost $g(x,y,z,\delta)$, also known as the $(1 - \varepsilon)$ *value-at-risk* [12, 14]. The parameter $\varepsilon \in (0, 0.5)$ can be considered as a tuning knob for the desired risk level.

REMARK E.1—CHANCE CONSTRAINTS Further chance constraints on the first-stage variables x and y can be included in (E.12) by the theory in Schildbach et al. [22]. ■

3.2 Two-Stage Scenario Program

In analogy to classic scenario approach, a *two-stage scenario program* is obtained by substituting the random samples $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$,

$$\text{TSP}[\omega^{(K)}] : \min_{q,x,y,z^{(1)},\dots,z^{(K)}} q, \quad (\text{E.13a})$$

$$\text{s.t.} \quad g(x,y,z^{(k)},\delta^{(k)}) \leq q \quad \forall k = 1, 2, \dots, K, \quad (\text{E.13b})$$

$$x \in \{0,1\}^l, \quad y \in \mathcal{Y}(x), \quad (\text{E.13c})$$

$$z^{(k)} \in \mathcal{Z}(x,y,\delta^{(k)}) \quad \forall k = 1, 2, \dots, K. \quad (\text{E.13d})$$

The first stage decision variables $q^*(\omega^{(K)})$, $x^*(\omega^{(K)})$, $y^*(\omega^{(K)})$ are selected jointly for all scenarios (*non-anticipativity*). The second-stage decision variables $z^{(k)}$, on the other hand, are generally dependent on the scenarios $\delta^{(k)}$. By substituting the sample values, $\text{TSP}[\omega^{(K)}]$ becomes a deterministic mixed-integer optimization program. For the purposes of this paper, it is assumed to be efficiently solvable [17], for instance by exploiting the special structure of the problem [16, 23, 27].

REMARK E.2—SAMPLE AVERAGE APPROXIMATION The key difference of the TSA to the widely used *sample average approximation* (SAA) methods [5, 15, 26] is that the SAA methods minimize the *average* cost over all samples, while the TSA minimizes the cost of the *worst-case* sample. ■

Let $z^*(\bar{x}, \bar{y}, \delta)$ denote the *recourse function* of the TSP $[\omega^{(K)}]$, i.e., the optimal solution to the second-stage problem when $\bar{x} \in \{0, 1\}^l$ and $\bar{y} \in \mathcal{Y}(\bar{x})$ have been selected as the first-stage decisions [5, 15, 26] (see the appendix for more details). Analogously to (E.4), for the TSP $[\omega^{(K)}]$ the (*a priori*) *violation probability* $v(\omega^{(K)})$ can then be defined as

$$v(\omega^{(K)}) := \mathbf{P}[g(x^*(\omega^{(K)}), y^*(\omega^{(K)}), z^*(x^*(\omega^{(K)}), y^*(\omega^{(K)}), \delta), \delta) > q^*(\omega^{(K)})] .$$

Intuitively speaking, $v(\omega^{(K)})$ can be expected to decrease (i.e., the first stage solution becomes “more risk averse”) if the number of scenarios K increases. A precise relationship, depending on the problem dimensions l and m , is given by the following theorem.

THEOREM E.3—DISTRIBUTION BOUND The distribution of the violation probability $v(\omega^{(K)})$ of the TSP $[\omega^{(K)}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(K)}) > \nu] \leq u_{l,m}^{(K)}(\nu) , \quad (\text{E.14a})$$

$$u_{l,m}^{(K)}(\nu) := \min\{1, 2^l \text{B}(\nu; K, m)\} , \quad (\text{E.14b})$$

for any $\nu \in [0, 1]$, where $\text{B}(\cdot; \cdot, \cdot)$ denotes the Beta distribution function as defined in (E.7). ■

For the sake of readability, the proof of Theorem E.3 has been moved to the appendix. The following result is a straightforward consequence.

COROLLARY E.3—PROBABILITY BOUND The scenario solution $x^*(\omega^{(K)})$ of the TSP $[\omega^{(K)}]$ keeps the shortfall probability (E.2b) in expectation, i.e., $\mathbf{E}^K[v(\omega^{(K)})] \leq \varepsilon$, if K is selected large enough such that

$$\int_0^1 u_{l,m}^{(K)}(\nu) d\nu \leq \varepsilon . \quad (\text{E.15})$$

■

The concept of optimizing for the worst-case scenario, rather than the average over all scenarios, has received little attention in the research on stochastic programming. Corollary E.3 provides a theoretical underpinning for this approach, by linking the sample size K to the *expected shortfall probability* of the scenario solution.

Remarkably, for this novel approach the sample size is generally small and independent of the dimension of the uncertainty space Δ . Exemplary sample sizes K are shown in Figure E.1, for two probability levels $\varepsilon = 10\%$ and $\varepsilon = 20\%$, and for varying numbers of first-stage binary variables l and continuous variables m .

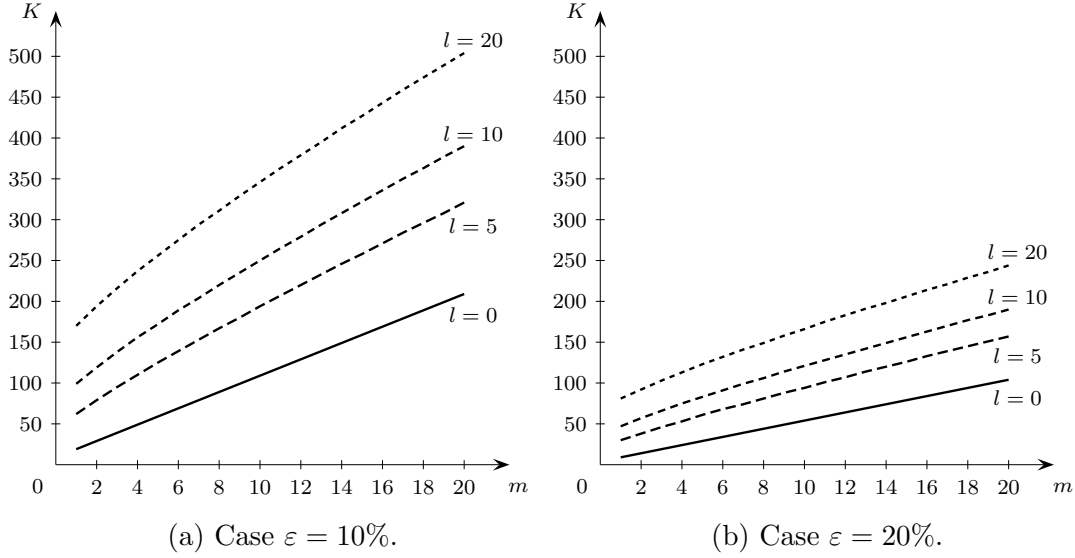


FIGURE E.1. Illustration of the minimal sample sizes K , according to Corollary E.3, for different numbers of first-stage binary variables l and continuous variables m .

3.3 Two-Stage Sample Removal

Similar to the classic scenario approach, the theory of the TSA can be extended to the case where R of K samples are removed after their observation. The purpose is again to reduce the effect of outliers in the samples $\omega^{(K)}$ on the scenario solution $q^*(\omega^{(K)})$, $x^*(\omega^{(K)})$, $y^*(\omega^{(K)})$.

The removal procedure is assumed to be as in Definition E.1. In particular, a greedy removal procedure seems an effective choice for the two-stage scenario program; that is removing the samples one-by-one, according to which has the worst associated objective function value [8, Sec. 5.1].

THEOREM E.4—DISTRIBUTION BOUND WITH SAMPLE REMOVAL Let $\omega^{(K,R)}$ be the remaining samples after applying a removal procedure to $\omega^{(K)}$. Then the distribution of the violation probability $v(\omega^{(K)})$ of TSP $[\omega^{(K)}]$ satisfies

$$\mathbf{P}^K[v(\omega^{(K)}) > \nu] \leq u_{l,m}^{(K,R)}(\nu) , \quad (\text{E.16a})$$

$$u_{l,m}^{(K,R)}(\nu) := \min\left\{1, 2^l \binom{R+m}{m} B(\nu; K, R+m+1)\right\} , \quad (\text{E.16b})$$

for any $\nu \in [0, 1]$, where $B(\cdot; \cdot, \cdot)$ denotes the Beta distribution function as defined in (E.7). ■

The proof of Theorem E.4 can be found in the appendix. The following corollary can be deduced.

COROLLARY E.4—PROBABILITY BOUND WITH SAMPLE REMOVAL The scenario solution $x^*(\omega^{(K)})$ of TSP $[\omega^{(K)}]$ keeps the shortfall probability (E.2b) in expectation, i.e.,

$\mathbf{E}^K[v(\omega^{(K)})] \leq \varepsilon$, if K and R are selected such that

$$\int_0^1 u_{l,m}^{(K,R)}(\nu) d\nu \leq \varepsilon . \quad (\text{E.17})$$

■

Figure E.2 shows some exemplary sample sizes K , for the cases of $R = 10$ and $R = 20$, assuming a risk level of $\varepsilon = 10\%$, and varying the problem dimensions. For comparison, the case of $R = 0$ can be found in Figure E.1(a).

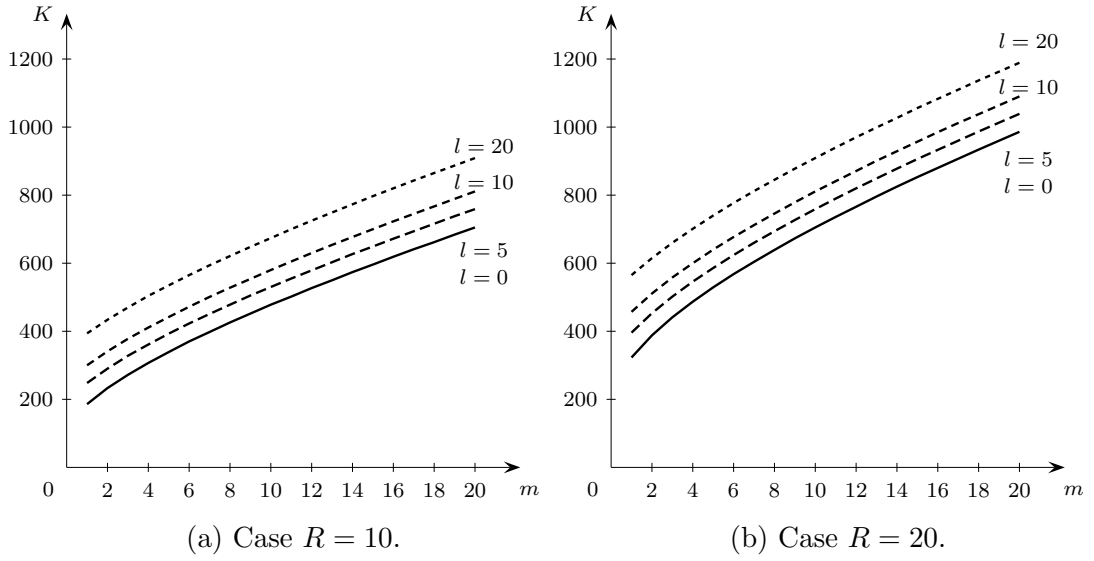


FIGURE E.2. Illustration of the minimal sample sizes K , according to Corollary E.4, for $\varepsilon = 10\%$ and different numbers of first-stage binary variables l and continuous variables m .

4. A Farmer's Problem

This section demonstrates the application of the presented concepts to a modified version of the *farmer's problem* [5, Sec. 1.1], which is similar to the classic *newsvendor's problem*.

4.1 Problem Description

A farmer can raise five different crops $i \in \{b, c, p, s, w\}$, namely barley (b), corn (c), potatoes (p), sugar beets (s), and wheat (w). He can arbitrarily distribute these on $C = 500$ ha (hectares) of land. Each crop bears a fixed cost of ρ_i , as well as a variable cost per hectare of γ_i , respectively. The first-stage decision includes the binary variables $x := [x_b \ x_c \ x_p \ x_s \ x_w]^T$ in $\{0, 1\}^5$ of which crop to grow, and the continuous variables $y := [y_b \ y_c \ y_p \ y_s \ y_w]^T$ in \mathbb{R}^5 of how much area to dedicate to each crop.

The total harvest is proportional to the dedicated area for each crop, multiplied by an uncertain yield factor η_i . The yields depend on the weather conditions over the growing seasons as well as other uncertain events, which are either specific to each individual crop or may affect all crops in the same manner.

The (uncertain) harvest is sold on the global market after the growing season, at the uncertain prices less some transportation costs, ϕ_i^g . However, the farmer is able to sell a small quantity of β_i^l of each crop i on the local market at a premium of 50% to the global price; and another quantity of β_i^n to preferred customers on the national market at a premium of 20%. The demand quantities θ_i^l and θ_i^n of the local and national markets are again uncertain, and correlated with the crop-specific yields and the weather conditions. The parameter values for this example are summarized in Table E.1.

The second-stage decision variables comprise the quantities to be sold on the local (l), national (n), and global market (g): $z_j := [z_b^j \ z_c^j \ z_p^j \ z_s^j \ z_w^j]^T$, where $j \in \{l, n, g\}$. The overall second-stage decision vector is denoted $z := [z_l^T \ z_n^T \ z_g^T]^T$.

All random quantities in this example $\delta = \{\phi_l, \phi_n, \phi_g, \eta, \theta_l, \theta_n\}$ are generally correlated. The exact data generation process is rather involved and not central to the results—in fact, *any* stochastic model will work. Without describing the details, note that barley and wheat are considered as the most risky crops in terms of yields, followed by corn, while potatoes and sugar beets have the most steady yields.

4.2 Mathematical Formulation

Based on the problem description and the data in Table E.1, the corresponding scenario program $\text{TSP}[\omega^{(K)}]$ reads as follows:

$$\min_{q, x, y, z^{(1)}, \dots, z^{(K)}} q, \quad (\text{E.18a})$$

$$\text{s.t. } \phi_l^{(k)T} z_l^{(k)} + \phi_n^{(k)T} z_n^{(k)} + \phi_g^{(k)T} z_g^{(k)} - \rho^T x - \gamma^T y \geq -q \quad \forall k = 1, 2, \dots, K, \quad (\text{E.18b})$$

$$e^T x \leq C, \quad (\text{E.18c})$$

$$y \leq Cx, \quad (\text{E.18d})$$

$$x \in \{0, 1\}^5, \quad y \geq 0, \quad (\text{E.18e})$$

$$z_l^{(k)} + z_n^{(k)} + z_g^{(k)} \leq \eta^{(k)} \odot x \quad \forall k = 1, 2, \dots, K, \quad (\text{E.18f})$$

$$z_l^{(k)} \leq \theta_l^{(k)}, \quad z_n^{(k)} \leq \theta_n^{(k)} \quad \forall k = 1, 2, \dots, K, \quad (\text{E.18g})$$

$$z_l^{(k)} \geq 0, \quad z_n^{(k)} \geq 0, \quad z_g^{(k)} \geq 0 \quad \forall k = 1, 2, \dots, K. \quad (\text{E.18h})$$

All vector inequalities are understood as element-wise, $e \in \mathbb{R}^5$ denotes the vector of ones, and “ \odot ” represents the element-wise vector product.

Constraint (E.18b) states that the profit, i.e., the revenues minus costs, must be at least $-q$ in all scenarios k , where q is minimized. Constraints (E.18c) ensure that the given land area is not exceeded, and (E.18d) that the fixed cost is paid for the growing of each crop. Constraints (E.18f) restrict the total sales on the local, national, and global

DESCRIPTION	PARAMETER	AVERAGE VALUE
land area	C	500 ha
fixed cost	$\rho := [\rho_b \ \rho_c \ \rho_p \ \rho_s \ \rho_w]^T$	10,000 €
variable cost	$\gamma := [\gamma_b \ \gamma_c \ \gamma_p \ \gamma_s \ \gamma_w]^T$	500 €/ha
yield factors	$\eta := [\eta_b \ \eta_c \ \eta_p \ \eta_s \ \eta_w]^T$	$[6.7 \ 6.1 \ 40.0 \ 65.0 \ 7.3]^T$ t/ha
global prices	$\phi_g := [\phi_b^g \ \phi_c^g \ \phi_p^g \ \phi_s^g \ \phi_w^g]^T$	$[280 \ 270 \ 46 \ 30 \ 280]^T$ €/t
national prices	$\phi_n := [\phi_b^n \ \phi_c^n \ \phi_p^n \ \phi_s^n \ \phi_w^n]^T$	$[336 \ 324 \ 55 \ 36 \ 336]^T$ €/t
local prices	$\phi_l := [\phi_b^l \ \phi_c^l \ \phi_p^l \ \phi_s^l \ \phi_w^l]^T$	$[420 \ 405 \ 69 \ 45 \ 420]^T$ €/t
national demand	$\theta_n := [\theta_b^n \ \theta_c^n \ \theta_p^n \ \theta_s^n \ \theta_w^n]^T$	$[100 \ 100 \ 200 \ 200 \ 100]^T$ t
local demand	$\theta_l := [\theta_b^l \ \theta_c^l \ \theta_p^l \ \theta_s^l \ \theta_w^l]^T$	$[100 \ 100 \ 200 \ 200 \ 100]^T$ t

TABLE E.1. Parameters and numerical values of the Farmer's Problem (b=barley, c=corn, p=potatoes, s=sugar beets, w=wheat, l=local market, n=national market, g=global market).

market to the total harvest for each crop, and (E.18g) observe the maximum demands on the local and national markets.

Note that the farmer's problem satisfies Assumption E.1, in particular (E.18) is feasible for any outcome of the scenarios $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(K)}$.

REMARK E.3—COMPUTATION (a) An outer loop for binary variables x can be based on a suitable integer solver, e.g., branch and bound [17]. (b) The inner problem, where the binary variables are either fixed or relaxed, can be solved as a standard linear program. (c) To increase the efficiency, it is possible to employ a tailored solver that exploits the special structure of the linear program, cf. [16, 23, 27]. ■

4.3 Discussion

The results of the two-stage scenario approach (TSA) can be compared with the traditional sample average approximation (SAA). The risk level is first set to $\varepsilon = 10\%$, resulting in a sample size of $K = 255$ according to Theorem E.3, used for both approaches.

For an empirical comparison, 100 sampled instances of the respective two-stage stochastic programs are solved, for the first-stage decision of the TSA and the SAA. Then each solution is tested under a total of 1,000 scenarios of the uncertainty δ .

As expected, the shortfall risk is much higher for the SAA (51.3%) than for the TSA (2.4%). However, note that there is some *conservatism* towards the required level of 10.0%. On the other hand, the SAA has an average profit of 0.79 m€, which is distinctly higher than the TSA with 0.71 m€.

A better picture is obtained when comparing the corresponding empirical distributions of the profit, as shown in Figure E.3. The distribution of the TSA has a lower mean, but also a lower variance than the distribution of the SAA.




FIGURE E.3. Distribution of objective function values (in million euros, m€), for the Two-Stage Scenario Approach (TSA) and Sample Average Approximation (SAA). The abscissa labels refer to the least profit required for a scenario to fall into a bin.

The empirical distribution in Figure E.3(a) is changed only invisibly if constraint removals are introduced, even though the computational requirements increase substantially. For example, for $R = 5$ the shortfall risk is 2.6% and the average profit is 0.72m€ and for $R = 10$ the shortfall risk is 2.8% and the average profit is 0.72m€. This suggests that the benefits of sample removal are quite limited for the two-stage scenario approach.

Finally, consider the TSA for different risk levels. For $\varepsilon = 1\%$ (sample size $K = 1,143$), the average profit is 0.70m€ and the actual violations are 0.22%. For $\varepsilon = 40\%$ (sample size $K = 25$), the average profit is 0.72m€ and the actual violations are 9.74%. Again, note that there is some conservatism in the risks for this particular case. The corresponding empirical distributions of the profit are shown in Figure E.4.




FIGURE E.4. Distribution of objective function values (in million euros, m€), for the Two-Stage Scenario Approach with different risk levels $\varepsilon = 1\%, 40\%$. The abscissa labels refer to the least profit required for a scenario to fall into a bin.

5. Conclusion

This paper has presented a novel two-stage scenario approach (TSA) for *risk averse two-stage stochastic programming*. The method is based on recent results in scenario-

based optimization [7–10], where a *scenario solution* is computed which is optimal for the worst case of a finite number of sampled scenarios. This scenario solution has a generalization property with respect to the *expected shortfall probability* of the original two-stage stochastic program.

The TSA is an intuitive and computationally efficient method, as the required sample size is small and independent of the uncertainty dimensions. Moreover, the method can be used in a variety of practical applications, as it allows for binary variables in the first stage decision. The properties and the effectiveness of the method have been demonstrated in a small numerical example.

A. Proof of Theorems E.3 and E.4

A.1 Recourse Function

For any $x \in \{0, 1\}^l$, $y \in \mathcal{Y}(x)$, and $\delta \in \Delta$, define $z^*(x, y, \delta)$ as the solution map to the second-stage problem (also called the *recourse function*),

$$z^*(x, y, \delta) := \arg \min_{z \in \mathcal{Z}(x, y, \delta)} g(x, y, z, \delta) . \quad (\text{E.19})$$

This solution map is well defined, as Assumption E.1(a) guarantees the existence of a solution, and uniqueness can be ensured by invoking a tie-break rule [9, Sec. 2,(5)].

A.2 Convexity of Shortfall Constraint

First, suppose that the binary decision variable x is fixed to any value $\bar{x} \in \{0, 1\}^l$. Substituting (E.19) into the TSP $[\omega^{(K)}]$ yields

$$\min_{y, q} \quad q , \quad (\text{E.20a})$$

$$\text{s.t.} \quad \mathbf{P}[g(\bar{x}, y, z^*(\bar{x}, y, \delta), \delta) \leq q] \geq 1 - \varepsilon , \quad (\text{E.20b})$$

$$y \in \mathcal{Y}(\bar{x}) . \quad (\text{E.20c})$$

It is now shown that (E.20) is of the shape of UP $[\varepsilon]$, for the continuous decision variables y and q . To this end, the convexity of $g(\bar{x}, y, z^*(\bar{x}, y, \bar{\delta}), \bar{\delta})$ must be verified, for any fixed value of the uncertainty $\bar{\delta} \in \Delta$. Using the standard condition [6, Equ. (3.1)] for convexity, pick arbitrary points $\bar{y}_1, \bar{y}_2 \in \mathcal{Y}(\bar{x})$ and a real $\lambda \in [0, 1]$:

$$\begin{aligned} & g(\bar{x}, \lambda \bar{y}_1 + (1 - \lambda) \bar{y}_2, z^*(\bar{x}, \lambda \bar{y}_1 + (1 - \lambda) \bar{y}_2, \bar{\delta}), \bar{\delta}) \\ & \leq g(\bar{x}, \lambda \bar{y}_1 + (1 - \lambda) \bar{y}_2, \lambda z^*(\bar{x}, \bar{y}_1, \bar{\delta}) + (1 - \lambda) z^*(\bar{x}, \bar{y}_2, \bar{\delta}), \bar{\delta}) \\ & \leq \lambda g(\bar{x}, \bar{y}_1, z^*(\bar{x}, \bar{y}_1, \bar{\delta}), \bar{\delta}) + (1 - \lambda) g(\bar{x}, \bar{y}_2, z^*(\bar{x}, \bar{y}_2, \bar{\delta}), \bar{\delta}) . \end{aligned}$$

The first inequality follows since by Assumption E.1(c)

$$\lambda z^*(\bar{x}, \bar{y}_1, \bar{\delta}) + (1 - \lambda) z^*(\bar{x}, \bar{y}_2, \bar{\delta}) \in \mathcal{Z}(\bar{x}, \bar{y}_1 + (1 - \lambda)\bar{y}_2, \bar{\delta}) ,$$

and because of the optimality property of the solution map (E.19). The second inequality is due to Assumption E.1(d). Note that the convexity of the recourse function is a classic result in two-stage stochastic programming [29].

A.3 Fixed Binary Variables

Since (E.20) has been shown to have the form of UP[ε], the result of the classic scenario approach can be applied. Let $y_{\bar{x}}^*(\omega^{(K)})$ and $q_{\bar{x}}^*(\omega^{(K)})$ denote the scenario solution of the SP[ε] that corresponds (E.20). Analogously to (E.4), define the violation probability

$$v_{\bar{x}}(\omega^{(K)}) := \mathbf{P}[g(\bar{x}, y_{\bar{x}}^*(\omega^{(K)}), z^*(\bar{x}, y_{\bar{x}}^*(\omega^{(K)}), \delta)) > q_{\bar{x}}^*(\omega^{(K)})] . \quad (\text{E.21})$$

Then Theorem E.1 yields that for any $\nu \in (0, 1)$,

$$\mathbf{P}^K[v_{\bar{y}}(\omega^{(K)}) \geq \nu] \leq \text{B}(\nu; K, m) . \quad (\text{E.22})$$

A.4 Free Binary Variables

Note that (E.22) holds for any fixed \bar{x} , as chosen *before* the observation of the scenarios. The desired result, however, must hold for $x^*(\omega^{(K)})$, the optimal choice in the scenario program *after* the observation of the scenarios.

This is resolved by a union bound over the set of all possible binary variables:

$$\mathbf{P}^K[v_{x^*(\omega^{(K)})}(\omega^{(K)}) \geq \nu] \leq \mathbf{P}^K[v_{\bar{x}}(\omega^{(K)}) \geq \nu \forall \bar{x} \in \{0, 1\}^l] . \quad (\text{E.23})$$

As the cardinality of the binary decision set $\{0, 1\}^l$ is 2^l ,

$$\mathbf{P}^K[v_{\bar{x}}(\omega^{(K)}) \geq \nu \forall \bar{x} \in \{0, 1\}^l] \leq 2^l \text{B}(\nu; K, m) . \quad (\text{E.24})$$

Theorem E.3 now follows from the combination of (E.23) and (E.24).

A.5 The Case with Sample Removal

For Theorem E.4, the steps in A.3 and A.4 may be repeated analogously for the case with sample removal. Indeed, for any fixed binary variable $\bar{x} \in \{0, 1\}^l$, define the violation probability of the scenario solution $y_{\bar{x}}^*(\omega^{(K,R)})$, $q_{\bar{x}}^*(\omega^{(K,R)})$ of (E.20), after removing R of the K samples, as

$$v_{\bar{x}}(\omega^{(K,R)}) := \mathbf{P}[g(\bar{x}, y_{\bar{x}}^*(\omega^{(K,R)}), z^*(\bar{x}, y_{\bar{x}}^*(\omega^{(K,R)}), \delta)) > q_{\bar{x}}^*(\omega^{(K,R)})] . \quad (\text{E.25})$$

Then Theorem E.2 yields that for any $\nu \in (0, 1)$,

$$\mathbf{P}^K[v_{\bar{x}}(\omega^{(K,R)}) \geq \nu] \leq \binom{R+m}{m} B(\nu; K, R+m+1) . \quad (\text{E.26})$$

The bound for the optimal binary variable $x^*(\omega^{(K,R)})$ follows from a union bound over the set of all possible binary variables,

$$\begin{aligned} \mathbf{P}^K[v_{x^*(\omega^{(K,R)})}(\omega^{(K,R)}) \geq \nu] &\leq \mathbf{P}^K[v_{\bar{x}}(\omega^{(K,R)}) \geq \nu \forall \bar{x} \in \{0, 1\}^l] \\ &\leq 2^l \binom{R+m}{m} B(\nu; K, R+m+1) , \end{aligned} \quad (\text{E.27})$$

and hence Theorem E.4.

References

- [1] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions*. Dover Publications, New York, 9th edition, 1970.
- [2] S. Ahmed. Convexity and decomposition of mean-risk stochastic programs. *Mathematical Programming, Series A*, 106:433–446, 2006.
- [3] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [4] D. Bertsimas, D.B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.
- [5] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, New York, 2nd edition, 2011.
- [6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.
- [7] G. Calafiore and M.C. Campi. Uncertain convex programs: Randomized solutions and confidence levels. *Mathematical Programming, Series A*, 102-1:25–46, 2005.
- [8] G.C. Calafiore. Random convex programs. *SIAM Journal on Optimization*, 20(6):3427–3464, 2010.
- [9] M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19:1211–1230, 2008.
- [10] M.C. Campi and S. Garatti. A sampling and discarding approach to chance-constrained optimization: Feasibility and optimality. *Journal of Optimization Theory and Applications*, 148:257–280, 2011.

- [11] L. Chen and T. Homem de Mello. Resolving stochastic programming models for airline revenue management. *Annals of Operations Research*, 177(1):91–114, 2010.
- [12] A. Damodaran. *Strategic Risk Taking, A Framework for Risk Management*. Prentice Hall, Upper Saddle River (NJ), 2008.
- [13] C.I. Fábián. Handling cvar objectives and constraints in two-stage stochastic models. *European Journal of Operational Research*, 191:888–911, 2008.
- [14] P. Jorion. *Value At Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, New York, 3rd edition, 2007.
- [15] P. Kall and J. Mayer. *Stochastic Linear Programming*. Springer, New York et al., 2nd edition, 2011.
- [16] G. Laporte and F.V. Louveaux. The integer l-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters*, 13:133–142, 1993.
- [17] G. Nemhauser and L. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, New York et al., 1999.
- [18] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.
- [19] N. Noyan. Risk-averse two-stage stochastic programming with an application to disaster management. *Computers & Operations Research*, 39:541–559, 2012.
- [20] C.G. Rawls and M.A. Turnquist. Pre-positioning of emergency supplies for disaster response. *Transportation Research Part B*, 44:521–534, 2010.
- [21] T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro. A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, 167:96–115, 2005.
- [22] G. Schildbach, L. Fagiano, and M. Morari. Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4):2479–2501, 2013.
- [23] R. Schultz, L. Stougie, and M.H. van der Vlerk. Two-stage stochastic integer programming: A survey. *Statistica Neerlandica*, 50(3):404–416, 1996.
- [24] R. Schultz and S. Tiedemann. Conditional value-at-risk in stochastic programs with mixed-integer recourse. *Mathematical Programming, Series B*, 105:365–386, 2006.
- [25] S. Sen, R.D. Doverspike, and S. Cosares. Network planning with random demand. *Telecommunication Systems*, 3:11–30, 1994.
- [26] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Stochastic Programming, Modeling and Theory*. SIAM, Philadelphia, 2009.
- [27] R.M. van Slyke and R. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17(4):638–663, 1969.

- [28] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (NJ), 1944.
- [29] R.J.B. Wets. Stochastic programs with fixed recourse: The equivalent deterministic program. *SIAM Review*, 16(3):309–339, 1974.

