Data-driven algorithms for multi-agent optimization and games

Georgios Pantazis

Lady Margaret Hall
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2022

Abstract

Multi-agent optimization and game theory constitute fundamental frameworks for modelling a plethora of engineering problems. In some systems, agents cooperate to optimize a collective cost, thus giving rise to multi-agent optimization programs. In other cases, agents act as selfish entities wishing to minimize their own individual cost subject to local or global constraints. To model such behaviour, non-cooperative game theory is required. Finally, agents though selfish can opt towards the formation of coalitions such that they receive a higher payoff. Such systems are studied under the lenses of cooperative game theory.

One of the most crucial challenges in such problems is finding solutions that are robust against the effects of uncertainty. Uncertainty represents our lack of knowledge about how the underlying system behaves and usually originates from unmodelled dynamics or external factors. Its significance is apparent in modern applications such as the electric vehicle charging problem, the optimal power flow problem with power generation from renewable energy sources, transportation and social networks.

Following a data-driven approach, this thesis builds a theoretical framework for the provision of robustness certificates for solutions to multi-agent optimization problems and non-cooperative games with an arbitrary cost function and uncertainty in the constraints. The theoretical tools developed allow the provision of collective certificates for sets (and subsets) of solutions that do not necessarily need to be optimal. Focusing on the class of aggregative games, agents' deviations from an equilibrium solution and uncertain constraints are considered. A data-driven equilibrium seeking algorithm is then proposed such that tunable collective guarantees can be provided for all possible feasible deviations from the equilibrium.

Shifting our attention towards a particular class of optimization programs with uncertain cost, prevalent in practical applications, we provide probabilistic guarantees for the randomized optimizer that do not depend on the number of agents but only on the size of the decision vector of each agent. As such, the probabilistic guarantees provided are scalable as the number of agents increases.

Finally, a data-driven theoretical framework is established for multi-agent cooperative games with uncertain value functions. Collective stability guarantees for the entire core set, a fundamental concept in cooperative game theory, are provided based on data from the value functions in a probably approximately correct learning fashion. The case where the core can be empty is considered and a methodology to accompany allocations in a relaxed core with stability certificates is proposed.

Data-driven algorithms for multi-agent optimization and games

Georgios Pantazis

Lady Margaret Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2022

This thesis is dedicated to
my parents, Nektarios and Maria
for their unconditional love and support throughout my life.

Acknowledgements

Many people have accompanied me to this rewarding and stimulating experience that the PhD was for me. This section is dedicated to everyone who, in their own way, contributed to the completion of my PhD journey.

I would like to express my deepest gratitude to my supervisor Prof. Kostas Margellos for his invaluable guidance and support, both professional and personal, in the past years. Prof. Margellos thank you for being such an amazing mentor throughout my PhD. Not only did you guide me on how to approach a research question, but you also taught me the significance of posing my own research questions in a rigorous and well-defined way. Giving me the freedom to explore different research areas and supporting my ideas has been very important for my development as a researcher. I am also thankful to Prof. Mark Cannon and Prof. Paul Goulart for the constructive feedback on my work and their advice throughout my studies and to Prof. Teodoro Alamo and Prof. Antonis Papachristodoulou for their very useful comments during the viva.

I am incredibly grateful to have worked with two amazing postdocs, Dr. Filiberto Fele and Dr. Filippo Fabiani, during my time in Oxford. Whenever I needed their advice with a problem they were more than happy to help and played a significant role in developing my research skills. I would also like to thank Yana Lishkova, Han Wang, Sophie Lekkas, Yuwen Chen, Xu Zhou, Pierre Lambert, PanZhou Dan and Liqun Zhao for making our working environment feel more vibrant. Many thanks also to Dr. Ross Drummond, Licio Romao, Niko Vertovec, Idris Kempf, Denis Lebedev, Michael Gartska, Xiaonan Lu and all other members of the control group. Special thanks to Emmanouil (Manos) Alexis, Martin Doff-Sotta and to Nikitas Rontsis for their friendship and support in the past years.

Outside of the control group, many thanks to Ada, Michele, Maria, Dimitris, Nikolas, Giorgos and Antonia for all the wonderful moments. My experience of Oxford would not have been so fulfilling without you. I would also like to thank Momo, Kumeren, Pablo, Maike, Casey, Verena and the rest of the Balliol group for all the great times we had. From walks to the university parks to college events and finally our trip to Cotswolds, these were some of the best memories I have from Oxford. Konstantina, Pavlina, Anthi, Giorgo, Katerina, Antoni and the rest of my friends in Greece, thank you for the time we spent together and your support in the past years.

Special thanks to my brother Kiriakos and my dear cousin Dimitra for making my time in Greece happier. Finally, I want to thank my parents Nektarios and Maria for their unconditional love and support throughout my life.

Abstract

Multi-agent optimization and game theory constitute fundamental frameworks for modelling a plethora of engineering problems. In some systems, agents cooperate to optimize a collective cost, thus giving rise to multi-agent optimization programs. In other cases, agents act as selfish entities wishing to minimize their own individual cost subject to local or global constraints. To model such behaviour, non-cooperative game theory is required. Finally, agents though selfish can opt towards the formation of coalitions such that they receive a higher payoff. Such systems are studied under the lenses of cooperative game theory.

One of the most crucial challenges in such problems is finding solutions that are robust against the effects of uncertainty. Uncertainty represents our lack of knowledge about how the underlying system behaves and usually originates from unmodelled dynamics or external factors. Its significance is apparent in modern applications such as the electric vehicle charging problem, the optimal power flow problem with power generation from renewable energy sources, transportation and social networks.

Following a data-driven approach, this thesis builds a theoretical framework for the provision of robustness certificates for solutions to multi-agent optimization problems and non-cooperative games with an arbitrary cost function and uncertainty in the constraints. The theoretical tools developed allow the provision of collective certificates for sets (and subsets) of solutions that do not necessarily need to be optimal. Focusing on the class of aggregative games, agents' deviations from an equilibrium solution and uncertain constraints are considered. A data-driven equilibrium seeking algorithm is then proposed such that tunable collective guarantees can be provided for all possible feasible deviations from the equilibrium.

Shifting our attention towards a particular class of optimization programs with uncertain cost, prevalent in practical applications, we provide probabilistic guarantees for the randomized optimizer that do not depend on the number of agents but only on the size of the decision vector of each agent. As such, the probabilistic guarantees provided are scalable as the number of agents increases.

Finally, a data-driven theoretical framework is established for multi-agent cooperative games with uncertain value functions. Collective stability guarantees for the entire core set, a fundamental concept in cooperative game theory, are provided based on data from the value functions in a probably approximately correct learning fashion. The case where the core can be empty is considered and a methodology to accompany allocations in a relaxed core with stability certificates is proposed.

Contents

1 Introduction

1.1 Motivation

Multi-agent systems characterize a large number of applications ranging from energy systems [1], [2] to traffic networks [3], economics [4] and the social sciences [5], [6]. Such systems comprise different entities/agents that interact with each other and make decisions, based on individual or collective criteria. Efficient coordination and control of multi-agent systems requires one to find strategies compatible with the multi-agent nature itself, taking into account computational issues as the number of agents increases, privacy concerns, personal interests and potential competition among agents. Hence, such systems intrinsically generate several challenges.

To address computational complexity and privacy concerns related to finding optimal decisions for the agents participating in a multi-agent system, several decentralised or distributed coordination schemes have been proposed [7], [8]. In the decentralised case, agents make decisions according to a certain optimization criterion and then communicate their strategies to a central authority, which is responsible for their coordination. In the distributed case, a central authority is absent and agents communicate with each other over a network, exchanging information with agents considered as neighbours given the underlying communication protocol.

Depending on their incentives, agents can either behave selfishly, being interested in minimizing only their own individual profit, or collectively, working together to minimize a common cost. The former case gives rise to a *non-cooperative game* [9], while the latter can be studied using standard optimization theory [10], [11], [7]. In other cases, coalitions among selfish agents can be allowed such that each of the members of the coalitions receives an individual payoff that is higher than what they would have received by working on their own. Such games are studied under the lenses of *cooperative game theory* [12].

Particularly in engineering, optimization and game theory have been extensively used for the analysis of communication networks [13], traffic networks [3, 14] and smart grids [15–19]. Be it a cooperative or a non-cooperative setting, uncertainty in such problems is ubiquitous and how to deal with it constitutes a fundamental challenge. Not taking uncertainty into account when proposing an optimal solution or a game equilibrium, could be agnostic, leading to erroneous behaviour of the system. As such, it is of major importance to accompany the solutions of such problems with some form of robustness certificates. At the same time, the typical assumptions made on the properties of the uncertainty, e.g., the geometry of its support set or the class of probability distributions that it follows might not be realistic in practice. As such, data-driven methodologies have been proposed to address these issues in the recent years.

A powerful data-driven technique with provable robustness certificates that waives the need for such conservative assumptions is the so called scenario approach [20–26], originally introduced to provide *a priori*[1] probabilistic guarantees for solutions of uncertain convex optimization programs. The theory has now been extended to non-convex decision making problems [24] where the probabilistic feasibility guarantees are obtained in an *a posteriori* fashion. The main advantage of the scenario approach is its applicability under very general conditions, since it belongs to the class of distribution-free techniques, i.e., knowledge of the uncertainty set and/ or the underlying probability distribution is not required, unlike other stochastic [27] or robust [28] methodologies.

[1] In this thesis the term *a priori* means that the probabilistic statements hold without requiring the observation of the data or its resulting decision. The term *a posteriori* implies that the statements made hold only after the uncertainty realizations and the resulting decision have been observed. The *a posteriori* results are usually more powerful, however they can only be provided after the decision has been made.

Even though the literature for providing feasibility certificates for yet unseen uncertainty realizations for a single optimizer is quite vast within the scenario approach literature (see [22],[29],[21],[30] to name a few), an attempt to the extension of the scenario approach to account for sets has only been seen in [31]. However, in a plethora of problems of practical interest, agents' decisions are feasible, though not necessarily optimal. This can be due to the specific numerical implementation of the solution-seeking algorithm that hinders the exact quantification of the optimal solution. Even under the assumption that the computed optimizer is exact, many optimization programs can have multiple optimal solutions, or several local optima. As such, it is important to establish a methodology where we treat all points of interest collectively as a set and aim to provide guarantees for such a set in an *a posteriori* fashion. Such a methodology has to be general enough to also address multi-agent games by providing guarantees collectively for the entire set of equilibria.

Despite recent advancements in data-driven techniques for uncertain multi-agent games, game theory still needs to overcome fundamental assumptions related to how agents make their decisions. Game-theoretic models impose on the agents a strict behavioural assumption, the so called rationality assumption. Specifically, it is assumed that agents wish to maximize their profit (or minimize their cost depending on the application) always acting as selfish entities. However, studies have shown that this is unrealistic in practice both in socioeconomics [32], [33], [34] and in engineering [35], [36] and that agents usually exhibit a bounded rationality behaviour [37], since their decision making procedure might vary depending on predispositions, computational capabilities, knowledge of the environment or time restrictions for making a decision. These factors can alter the incentives of the agents thus leading to decisions that are satisfactory, i.e., close to an equilibrium, but not optimal. A methodology for tuning in an *a priori* fashion the feasibility levels of the entire set of deviations by means of a data-driven equilibrium-seeking algorithm has not been established in the cognate literature.

Another significant challenge in the attempt to apply the scenario approach to uncertain multi-agent programs is the scalability of the robustness certificates of the optimizer as the number of agents increases. In general, the guarantees provided by

direct application of the scenario approach are not scalable as stated by the *a priori* result in [21], where it is shown that the number of support samples (see Definition 2.16) of a convex optimization problem, can in turn be bounded by the dimension of the optimization variable. This fact is then used to provide probabilistic feasibility guarantees for the optimal solution in an *a priori* fashion. Investigating specific classes of multi-agent programs common in practical applications, where the provided certificates can be theoretically proven to be independent of the number of agents has to the best of our knowledge not been explored in the literature.

In a coalitional gaming setting (see [12] and references therein), agents, though selfish, have the incentives to form coalitions to optimize their own profits. However, the total profits of agents participating in coalitions, represented in the game formulation by the so called value functions, are usually affected by uncertainty. A fundamental notion in cooperative game theory is that of *stability*, roughly defined as the lack of incentives from the agents to deviate from the already formed coalitions to originate other subcoalitions that might give them a higher payoff. To ensure that coalitions are stable even under the effect of uncertainty is a fundamental challenge. Encoding uncertainty in the value functions of a cooperative game has been explored in several works [38–45]. A data-driven methodology for the provision of provable *distribution-free* probabilistic guarantees on the stability of agents' allocations using available data from the coalitional values has not yet been established.

1.2 Thesis outline and contributions

In this section the main contributions of each chapter of the thesis are highlighted. These are summarized in Table 1.1 according to the nature of the result (*a posteriori* and/or *a priori*), the scalability of the provided guarantees and the class of problems to which they can be applied. A detailed description of each contribution is provided below.

- In Chapter 2 we define some important mathematical concepts used throughout the thesis. Notions from optimization such as feasibility and optimality are introduced. Non-cooperative games are then defined and the concept of equilibrium for standard

Class of problems	Type of solution	Nature of certificate	Scalability	Result
Feasibility programs with uncertain convex constraints	**Entire feasible set/ feasible subsets**	***A posteriori***	**Depends on the number of agents**	**Chapter 3**
Aggregative games with deviations from equilibrium and uncertain constraints	**Set of deviations around equilibrium**	***A priori***	**Depends on the number of agents**	**Chapter 4**
Optimization programs with uncertain aggregative term in the cost	**Unique optimizer**	***A priori***	**Agent independent**	**Chapter 5**
Cooperative games with uncertain value functions and nonempty core	**Core set**	***A posteriori* & *a priori***	**Depends on the number of agents**	**Chapter 6**
Cooperative games with uncertain value functions and empty core	**Solution in a relaxed core set**	***A posteriori***	**Depends on the number of agents**	**Chapter 6**

Table 1.1: Classification of main results of the thesis according to their main features

and generalized games is presented. The important class of aggregative games is also introduced. The more general class of variational inequality problems is presented as a means to design algorithms that converge to a game equilibrium and an overview of equilibrium seeking algorithms is then provided. Furthermore, cooperative games are defined in a deterministic setting and the key stability solution concepts, namely the core and the relaxed core are introduced. Finally, the chapter is concluded by an introduction to *a priori* and *a posteriori* scenario optimization.

- Chapter 3 of the thesis focuses on feasibility programs with uncertain convex constraints. It establishes a mathematical methodology for the provision of *a posteriori* probabilistic feasibility guarantees that hold for the entire feasibility region or subsets of it. As a byproduct of the proposed methodology one can provide certificates for solutions to multi-agent optimization programs and games with an arbitrary cost and uncertain constraints. It also constitutes a powerful tool that can be used for cases where the solution-seeking algorithms converge to a suboptimal or locally optimal feasible solution, or when agents deviate from the proposed optimizer/ equilibrium within the feasible domain. Numerical studies on electric vehicle charging control in both a multi-agent optimization and a game-theoretic setting corroborate the theoretical results. This chapter is based on

  - *"A posteriori probabilistic feasibility guarantees for Nash equilibria in uncertain multi-agent games", G. Pantazis, F. Fele, and K. Margellos, 21st IFAC World Congress, vol. 53, no. 2, pp. 3403–3408, 2020.*
  - *"On the probabilistic feasibility of solutions in multi-agent optimization problems under uncertainty", G. Pantazis, F. Fele, and K. Margellos, European Journal of Control, vol. 63, pp. 186–195, 2022.*

- Chapter 4 revolves around aggregative games with uncertainty in the constraints. We consider a setting where bounded deviations of agents' decisions from the equilibrium are allowed. Following a data-driven paradigm where the feasible region of the problem can be approximated through sampling, we propose a data-driven primal-dual algorithm that converges to a randomized equilibrium such that we can control in an *a priori* fashion the probabilistic feasibility levels collectively for all these possible deviations by adapting the position of the nominal equilibrium. Our algorithmic scheme is based on a novel methodology of identifying online the coupling constraints that play the most significant role to the systems' performance. Tuning of the provided guarantees is accomplished by controlling how many of these coupling constraints are allowed to intersect the region under study. Finally, numerical studies of a game with uncertain constraints on the aggregate decision support the theoretical findings. This chapter is based on

  - *"A priori control of probabilistic feasibility levels of solution sets around equilibria", G. Pantazis, F. Fele, and K. Margellos, Submitted to Automatica, 2022.*

- Chapter 5 focuses on a specific class of convex multi-agent programs, prevalent in practical applications, where agents cooperate to minimize a common cost, expressed as a function of the aggregate decision and affected by uncertainty. By exploiting the structure of the program under study and leveraging on existing results in the scenario approach literature, and in particular the so called support

rank notion [46], we provide for the optimal solution of the program distribution-free agent-independent certificates on the probability that the cost will deteriorate. This means that the constructed bound on the probability of constraint violation does not depend on the number of agents, but only on the dimension of the agents' decision. This leads to a significant improvement as it substantially reduces the number of samples required to achieve a certain level of probabilistic robustness as the number of agents increases. Our certificates can be used alongside any convex optimization algorithm centralised, decentralised or distributed, to obtain an optimal solution of the underlying problem. Our theoretical results are accompanied by a numerical example that investigates the electric vehicle charging problem and validates that the obtained robustness certificate is independent of the number of vehicles in the fleet. This chapter is based on:

  - *Agent independent probabilistic robustness certificates for robust optimization programs with uncertain quadratic cost", G. Pantazis, F. Fele, K. Margellos Proceedings of the 59th IEEE Conference on Decision and Control, 2020*
  - *"On the probabilistic feasibility of solutions in multi-agent optimization problems under uncertainty", G. Pantazis, F. Fele, and K. Margellos, European Journal of Control, vol. 63, pp. 186–195, 2022.*

- Chapter 6 addresses multi-agent cooperative games with uncertain value functions for which we establish distribution-free guarantees on the probability of allocation stability. In case the set of stable allocations, the so called core of the game, is empty, we propose a randomized relaxation of the core. We then show that those allocations that belong to this relaxed set can be accompanied by stability guarantees in a probably approximately correct fashion. Numerical experiments corroborate our theoretical findings. This chapter is based on

  - *"Probabilistically robust stabilizing allocations in uncertain coalitional games", G. Pantazis, F. Fabiani, F. Fele, K. Margellos, IEEE Control Systems Letters, vol. 6, pp. 3128-3133, 2022.*

- Finally, Chapter 7 concludes the thesis by summarizing the main contributions and proposing directions for future research.

2
Preliminaries

In this chapter we introduce key concepts from optimization, game theory and the scenario approach used throughout the thesis. We start by defining common notation and basic mathematical concepts for functions and sets. Then, we introduce constrained optimization programs and define notions such as feasibility and optimality, also distinguishing between convex and non-convex programs. To model the behaviour of selfish agents, non-cooperative games are introduced and the standard equilibrium problem, along with its generalized version, is formulated. We show that both optimization programs and equilibrium problems can be studied, under certain assumptions, using the more general mathematical framework of variational inequalities. An overview of available equilibrium seeking algorithms is then provided. We finally introduce cooperative games and two stability concepts, widely used in the game theoretic literature, namely the core and the relaxed core. This chapter concludes with an introduction to the basic concepts and most recent ideas of the scenario approach.

2.1 Notation

We denote by $\mathbb{N}$, $\mathbb{R}$, and $\mathbb{R}_+$ the set of natural, real and positive real numbers, respectively. When a matrix A is positive definite (positive semi-definite), we write $A \succ 0$ (respectively, $A \succeq 0$). We denote by $0_{q\times r}$ a $q \times r$ matrix of zeros, I_r the $r \times r$ identity matrix and by

$e_q \in \mathbb{R}^n$ the q-th vector of the canonical basis spanning $\mathbb{R}^n$. For a set S, $|S|$ denotes its cardinality, $\text{bdry}(S)$ its boundary, $\text{conv}(S)$ its convex hull, while 2^S denotes its power set, that is the collection of all subsets of S. Finally, let $\Pi_{K,D}(x) = \arg\min_{y \in K} (y-x)^T D(y-x)$ be the skewed projection operator on x.

2.2 Preliminaries on calculus and optimization

Consider a vector space $X \subseteq \mathbb{R}^n$. We introduce the notion of the norm defined as follows.

Definition 2.1 *A norm is a function $||\cdot|| : X \to \mathbb{R}$ that satisfies the following properties*

1. *For all $x \in X$, $||x|| = 0$ implies that $x = 0$.*
2. *For all $a \in \mathbb{R}$ and for all $x \in X$, we have $||ax|| = |a| \cdot ||x||$.*
3. *For all $x, y \in X$, we have $||x+y|| \leq ||x|| + ||y||$.*

The most common class of norms are the so-called *p-norms*. Their general form is

$$||x||_p = \Big(\sum_{i=1}^{n} |x_i|^p \Big)^{\frac{1}{p}},$$

where p is a real number for which it holds that $p \geq 1$. Specifically, for

1. $p = 1$, we have $||x||_1 = |x_1| + |x_2| + \cdots + |x_n|$. This norm is called the 1-norm.
2. For $p = 2$, we have $||x||_2 = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}$. This norm is called the 2-norm or Euclidean norm.
3. For $p = \infty$, we have $||x||_\infty = \max\{|x_1|, |x_2|, \ldots, |x_n|\}$. This norm is called the ∞-norm or max-norm.

We denote the closed d-dimensional p-normed ball centred at x with radius r as $\mathbb{B}_p(x, r) = \{y \in \mathbb{R}^d : ||y - x||_p \leq r\}$. Furthermore, we equip a vector space X with a norm $||\cdot||$ given by Definition 2.1, thus constructing a normed vector space and define a distance function $d : X \times X \to \mathbb{R}$ given by the relation $d(x, y) = ||x - y||$. Based on that distance function we can then define the continuity of any function as follows

Definition 2.2 *A function $f : X \to Y$, where Y is a vector space equipped with the norm $||\cdot||$, is continuous at a point $\bar{x}$ in its domain X if given any $\varepsilon > 0$, there exists $\delta > 0$ such that if $x \in X$ and $||x - \bar{x}|| < \delta$ then $||f(x) - f(\bar{x})|| < \varepsilon$. A function is continuous in its domain X if it is continuous at every point of its domain.*

If the gradient of a function f exists and is itself a continuous function, we say that f is *continuously differentiable*. A special case of continuity is the so called Lipschitz continuity defined as follows

Definition 2.3 *A function $f : X \to Y$ is Lipschitz continuous if there exists a real constant $L \geq 0$ such that, for all x and y in X, $||f(x) - f(y)|| \leq L||x - y||$.*

Another important notion in optimization and game theory is that of the convexity of a function. Since the domain of the convex function is a convex set, the definition of a convex set in Definition 2.4 precedes that of a convex function in Definition 2.5.

Definition 2.4 *A set C is convex if the line segment between any two points in C lies in C, i.e., $\forall\, x_1, x_2 \in C$, $\forall\, \theta \in [0,1]$ we have that $\theta x_1 + (1-\theta)x_2 \in C$.*

In words, a set is convex if for two points that belong to it, the line segment with boundaries these two points is also contained inside the set. We are now ready to introduce the following definition.

Definition 2.5 *A function $f : X \to \mathbb{R}$ with domain a convex set X is*

1. *Convex if $\forall\, x, y \in X$, $\forall\, \theta \in [0,1]$ we have: $f(\theta x + (1-\theta)y) \leq \theta f(x) + (1-\theta) f(y)$,*

2. *Strictly convex if $\forall\, x, y \in X$ with $x \neq y$ such that $\forall\, \theta \in [0,1]$ we have: $f(\theta x + (1-\theta)y) < \theta f(x) + (1-\theta) f(y)$,*

3. *Strongly convex if there exists $a > 0$ such that $\forall\, x, y \in X$, $\forall\, \theta \in [0,1]$ we have: $f(\theta x + (1-\theta)y) \leq \theta f(x) + (1-\theta) f(y) - a||x-y||^2$.*

The convexity of a function is important for a plethora of applications [47], [13], [48], [49] that involve the minimization of a given cost function. This is because convex functions enjoy the property that any of their local minima is also a global minimum. Note that strong convexity of a function implies strict convexity and strict convexity in turn implies convexity. Furthermore, an optimization program with a convex cost function can have possibly infinite solutions. If the cost function is strictly convex, then such program admits at most one solution. This means that if such minimum exists then it is the unique global minimum. Strong convexity of the cost implies uniqueness of the solution.

The field of optimization plays a crucial role in modern engineering, economics and social sciences. Optimization usually involves the minimization of a function on a subset of its domain. In mathematical terms, we wish to solve the following problem:

$$\begin{aligned} P: \min_{x \in X} & f(x) \\ & \text{subject to } g(x) \leq 0, \end{aligned} \tag{2.1}$$

where $f : X \subseteq \mathbb{R}^d \to \mathbb{R}$ and $g : X \subseteq \mathbb{R}^d \to \mathbb{R}^m$ (with $d, m \in \mathbb{N}$) are real-valued functions with domain X. All points $x \in X$ such that $g(x) \leq 0$ are called feasible points. The set of optimal solutions of (2.1) consists of all points x^* such that $f(x^*) \leq f(x)$ for all x in the feasible set. There are a plethora of iterative algorithms that can compute the optimal solution of (2.1) (we refer the reader to [11], [7], [10], [50], [51], [52] for a thorough overview of such algorithms). If f and g are convex functions and X is a convex set, we say that (2.1) is a convex optimization program; Otherwise we refer to (2.1) as a non-convex optimization program. Non-convex optimization programs can be encountered in applications [53], [54] and devising algorithms that converge to a global optimum is in this case very challenging, as local optima of non-convex optimization programs are not necessarily global optima.

2.3 Non-cooperative game theory

Game theory [55] is a branch of mathematics that studies decision problems in multi-agent environments. Each agent's objective is to individually optimize her own cost function that depends not only on her own strategy but also on the strategies of all the

other agents. Games can be divided in two main categories, namely, cooperative games and non-cooperative games. A game is cooperative if selfish agents are allowed to form coalitions in order to improve their payoffs and non-cooperative if agents work on their own. In this section we will focus on non-cooperative games. A discussion on cooperative game theory and its basic concepts can be found in Section 2.5.

2.3.1 Nash equilibrium problems

A game is composed of N agents, each of them choosing a decision $x_i \in \mathbb{R}^n$, aiming at minimizing their individual cost function $J_i : \mathbb{R}^{nN} \to \mathbb{R}$, which is affected by the agents' decision x_i and the decisions of all other agents $(x_j)_{j=1,j\neq i}^N$ also denoted as x_{-i}. For each $i \in \{1,\dots,N\} = \mathcal{N}$, the decision vector x_i is confined to a given set. In mathematical terms we have that given the decisions of all other agents x_{-i}, each agent solves the following optimization program

$$\min_{x_i \in X_i} J_i(x_i, x_{-i}).$$

Note that the symbols x and (x_i, x_{-i}) are used interchangeably in the thesis, depending on the context. The collection of these coupled optimization programs forms the following game, typically defined as a tuple $G := (\mathcal{N}, (J_i)_{i\in\mathcal{N}}, (X_i)_{i\in\mathcal{N}})$. In the seminal works [56], [57] a key solution concept, the well-known Nash equilibrium (NE), was introduced and its existence properties were studied. Its definition is as follows (see [9]):

Definition 2.6 *(Nash equilibrium)*
A vector $x^* = (x_i^*)_{i=1}^N$ is a NE of the game G if and only if $J_i(x_i^*, x_{-i}^*) \leq J_i(x_i, x_{-i}^*)$ for all $x_i \in X_i$ and for all $i \in \{1,\dots,N\}$.

In words, a NE is a set of agents' decisions for which each agent has no incentive to unilaterally deviate from her decision, given the decisions of all other agents. Seeking a NE of G gives rise to a so called NE problem.

From the definition of NE, it is straightforward that NE do not necessarily correspond to minimizers of the collective cost of the entire population given by the

optimization program

$$\min_{x \in X} \sum_{i=1}^{N} J_i(x_i, x_{-i}), \tag{2.2}$$

where $x = (x_i, x_{-i})$ and $X = \prod_{i=1}^{N} X_i$ is the cartesian product of the agents' decision spaces. This fact differentiates the NE from an optimal solution to (2.2), also known as the *social welfare optimum*. However, if we are able to acquire a relation between the social optimum and the NE we can use results related to the computation of the optimal solution to derive conditions for the (sub)optimality of the agents' decisions in games. This is the case for the class of so-called *potential games*. Potential games are characterized by the existence of a *potential function* whose minimization leads to a NE of the original game. Games that admit a potential function will be studied in Chapter 4. It is apparent that the smallest the distance between NE and social welfare optimum, the closer the entire population is to its optimal behaviour. Several works have focused on investigating the connection between game equilibria and social welfare optima [58],[59], [60], based on a well-known concept in algorithmic game theory termed price of anarchy [61]. In particular, the aforementioned references show that, under certain assumptions, a NE converges to the social optimum of the game as the number of agents increases and they coincide in the limiting case where the number of agents tends to infinity.

2.3.2 Generalized NE problems

Many times in applications the constraints of the problem couple the decisions of the agents via a shared decision set, thus giving rise to a more general class of games, the so-called Generalized Nash Equilibrium Problems (GNEPs). Each of the coupling constraints has the form $C = \{x \in \mathbb{R}^{nN} : g(x) \leq 0\}$, where the function $g : \mathbb{R}^{nN} \to \mathbb{R}^m$ usually satisfies certain properties such as continuity and convexity and m denotes the total number of coupling constraints. From each agent's perspective, the decisions of other agents x_{-i} are considered to be fixed, i.e., for agent i the coupling constraint C takes the form $C_i(x_{-i}) = \{x_i \in \mathbb{R}^n : g(x_i, x_{-i}) \leq 0\}$. Given the decisions of all other agents x_{-i}, each agent solves the following optimization program

$$\min_{x_i \in X_i} J_i(x_i, x_{-i}) \text{ subject to } x_i \in C_i(x_{-i})$$

The collection of these coupled optimization programs forms the following game, defined as the tuple $G' := (\mathcal{N}, (J_i)_{i\in\mathcal{N}}, (X_i)_{i\in\mathcal{N}} \cap C)$.

In the seminal work of [62] the existence of Generalized Nash Equilibria for a certain class of generalized games, that of concave N-person games, is proved. For (G)NEPs computing a (Generalized) Nash Equilibrium using the standard results in non-cooperative game theory [9] becomes a computationally demanding task as the number of players grows to a large size. The reason behind this is that each agent needs to obtain information from all the other players in order to compute and minimize her own cost function, thus forming a complex network of interactions.

The so called generalized Nash equilibrium for GNEPs is defined as follows.

Definition 2.7 *(Generalized Nash Equilibrium)*
A vector $x^* = (x_i^*)_{i=1}^N$ is a generalized Nash Equilibrium (GNE) of the game G' if and only if $J_i(x_i^*, x_{-i}^*) \leq J_i(x_i, x_{-i}^*)$ for all $x_i \in X_i \cap C_i(x_{-i}^*)$ for all $i \in \mathcal{N}$.

The class of GNEPs has been studied extensively in the operations research literature [63], [64] and has in the last decades caught the attention of the control community [65], [66], [67], [68], [69].

2.3.3 Aggregative games

Specifying the structure of the mutual interactions among the agents' strategies allows us to distinguish among several classes of games. In addition, the particular structure can also be exploited in designing tailored algorithms to compute GNE, either exact or approximate. An important class of games considered in this thesis is that of *aggregative games*.

In aggregative games, each agents' objective function is affected by the aggregate of the decisions of all the agents in the population, including its own local one see, e.g., [70]. This gives rise to the following game

$$G_a : \forall\, i \in \mathcal{N} : \min_{x_i \in X_i} J_i(x_i, \sigma(x)),$$

where $\sigma : \mathbb{R}^{nN} \to \mathbb{R}^n$ is an average aggregate function of the form $\sigma(x) = \frac{1}{N}\sum_{i\in\mathcal{N}} x_i$. Note that, sometimes, aggregative games can be formulated with σ being just the sum of

the decisions and not the average. The notion of NE and GNE as described in Definitions 2.6 and 2.7 also holds for the special case of aggregative games. Adapting Definition 2.6 we obtain the following NE definition for the case of aggregative games.

Definition 2.8 *A vector $x^* = (x_i^*)_{i=1}^N$ is a NE of the aggregative game G_a if and only if*

$$J_i(x_i^*, \sigma(x^*)) \leq J_i\left(x_i, \frac{1}{N}x_i + \frac{1}{N}\sum_{j\in\mathcal{N}\setminus\{i\}} x_j^*\right) \text{ for all } x_i \in X_i \text{ for all } i \in \mathcal{N}.$$

For the class of aggregative games another type of equilibrium is commonly used [3]. This equilibrium is called the *Wardrop* equilibrium (WE) [71], [14] defined as follows:

Definition 2.9 *(Wardrop equilibrium)*
A vector $x^* = (x_i^*)_{i=1}^N$ is a WE of the aggregative game G_a if and only if $J_i(x_i^*, \sigma(x^*)) \leq J_i(x_i, \sigma(x^*))$ for any $x_i \in X_i$ and for any $i \in \{1,\ldots,N\}$.

Note that, unlike Definition 2.8, for the case of WE in Definition 2.9, the decision of agent $i \in \mathcal{N}$ is considered to be negligible with respect to aggregate decision of the entire population. The generalized WE is a straightforward extension of the standard WE, i.e., obtained by including coupling constraints in the feasible set of each agent.

Aggregative games capture a wide range of problems in different domains, like transportation and energy problems. For example, aggregative games can be used in smart grid applications to model the charging scheduling of large fleets of electric vehicles or in transportation for coordinating a large fleet of autonomous vehicles. [3, 69, 72, 73]. As a result, several papers have exploited this structure to construct iterative algorithms that allow NE computation in aggregative games.

2.4 Variational inequalities

The content of this section is based on [63]. In its general form, a variational inequality (VI) problem is defined as follows

Definition 2.10 *Given a subset X on the Euclidean n-dimensional space $\mathbb{R}^n$ and a mapping $F : X \to \mathbb{R}^n$, the variational inequality problem $VI(X,F)$ is defined as the problem of finding a vector $x \in X$ such that*

$$(y-x)^T F(x) \geq 0, \ \forall\, y \in X. \tag{2.3}$$

The set of solutions to $\mathrm{VI}(X,F)$ is denoted as X^*. Note that from Definition 2.10 we can obtain the relation

$$y^T F(x) \geq x^T F(x),\ \forall\, y \in X. \tag{2.4}$$

This means that $x \in X^*$ if and only if x is the solution to the optimization program

$$\min_{y \in X} y^T F(x). \tag{2.5}$$

Consider the optimization program P from Section 2.1, where we additionally set $\bar{X} = \{x \in X : g(x) \leq 0\}$ for brevity. Furthermore, assume that f is defined on an open superset of $\bar{X}$ and is continuously differentiable. By the minimum principle in nonlinear programming, any local minimizer x of the optimization problem at hand satisfies

$$(y-x)^T \nabla f(x) \geq 0, \text{ for all } y \in \bar{X}.$$

The latter relation can be equivalently written as $\mathrm{VI}(\bar{X}, \nabla f)$, which is called the stationary point problem associated with the optimization problem. Its solution is called a *stationary point*. Furthermore, we know that if f is a convex function, then every stationary point is a global minimum of this optimization problem. As such, the solution set of a convex program P, as in relation (2.1), under the aforementioned assumptions is equal to the solution set X^* of $\mathrm{VI}(X, \nabla f)$. We now provide a connection between VIs and NEPs. To this end, we impose the following assumption.

Assumption 2.1 *We assume that:*

1. *for each $i \in \mathcal{N}$, X_i is a nonempty closed and convex subset of $\mathbb{R}^n$,*
2. *for each fixed tuple x_{-i} the function $J_i(x_i, x_{-i})$ is convex and continuously differentiable on X_i and that*
3. *for each $i \in \mathcal{N}$, the function $J_i(x_i, x_{-i})$ is continuously differentiable on $X = \prod_{i \in \mathcal{N}} X_i$.*

Furthermore, we define the so called *pseudo-gradient mapping* $(\nabla_{x_i} J_i(x))_{i=1}^N$ as the vector that consists of all the partial gradient vectors taken with respect to each agent's decision vector. A NE of a game can then be obtained by solving a VI problem under the proposition that follows.

Proposition 2.1 *Consider the setting of game G. Under Assumption 2.1, the tuple* $x^* = (x_i^*)_{i=1}^N$ *is a NE of G if and only if* x^* *is a solution to VI(X,* $(\nabla_{x_i} J_i(x))_{i=1}^N$*).*

Below we define properties of a mapping that play a significant role in VI analysis.

Definition 2.11 *A mapping* $F : \mathbb{R}^n \to \mathbb{R}^n$ *defined on a set X is*

1. *Monotone if* $(x-y)^T(F(x)-F(y)) \geq 0$ *for all* $x, y \in X$*;*
2. *Strictly monotone* $(x-y)^T(F(x)-F(y)) > 0$ *for all* $x, y \in X$ *with* $x \neq y$*;*
3. *Strongly monotone if there exists* $a > 0$ *such that* $(x-y)^T(F(x)-F(y)) \geq a||x-y||^2$ *for all* $x, y \in X$*.*

The following theorem adapted from [13] establishes how the monotonicity of the pseudo-gradient mapping affects the solution set of the corresponding NE problem.

Theorem 2.1 *Given an NE problem G, suppose that G satisfies Assumption 2.1 and let* $F(x) = (\nabla_{x_i} J_i(x))_{i=1}^N$*. Then we have that*

1. *If F is a monotone function on X, then G has a convex (possibly empty) solution set.*
2. *If F is strictly monotone on X, then G has at most one solution.*
3. *If F is strongly monotone on X, then G has a unique solution.*

Notice that there is a parallel between convexity in optimization/ NE problems and monotonicity in variational inequalities with respect to how many solutions their respective problems can have. Strict monotonicity implies yet again that if a solution for the VI exists, then it is unique. Strong monotonicity, as a stronger property guarantees the uniqueness of the VI solution. The parallel with optimization programs is straightforward

given that an optimization program with a (strongly or strictly) convex cost function has a (resp. strongly or strictly) monotone gradient. Similarly, for the case of NE problems, the number of possible solutions depends on the monotonicity properties of the pseudogradient as stated in Theorem 2.1. Most equilibrium seeking algorithms are based on particular assumptions with regards to the monotonicity properties of the corresponding mapping or the convexity of the agents' local costs. Furthermore, assumptions on the convexity or the linearity of both local and coupling constraints are often imposed.

Many NE seeking algorithms are centralized in the sense that a central authority/coordinator has access to the entire information of the game (agents' cost functions, local and coupling constraints etc.). However, recent work has been focused on the development of decentralised and distributed algorithms as they have several merits. First, solving the problem in a centralized manner may raise privacy concerns regarding the information being shared among agents. Secondly, the centralized problem becomes computationally prohibitive to solve as the number of agents increases.

To find a solution of a multi-agent game, some algorithms require additional information regarding the game model and are based on the assumption that each player can observe the strategies of all other players. Methods that pertain to this category are the so-called best response strategies or fictitious plays. A vast amount of literature has been dedicated to the study of best response (BR) dynamics [74–86]. BR dynamics can be viewed as the natural dynamics of the interactions among the agents that lead each agent to the computation of her optimal strategy given the strategies played by all other agents [87]. Using BR dynamics as an update scheme and under mild assumptions on the structure of the agents' network graph, [73] proposes two distributed algorithms that steer the agents' strategies to a slightly modified version of the NE, known as network aggregative NE. In [88] and [89] a special class of algorithms, the so-called relaxation algorithms are proposed for the computation of NE.

In [90] one of the first attempts towards the connection of games with variational inequalities was made by proving the existence of a NE for bimatrix games (2-player non-zero sum games) leveraging on the equivalence of the problem at hand with a linear complementarity problem, a problem equivalent to a VI problem. As evidenced by

their equivalence with (G)NEPs under certain assumptions, VIs constitute a powerful tool for the computation of NE. In [64] the notion of a special case of NE, that of the variational equilibrium is introduced which is used as a solution concept in [3] and [58]. In [65] a primal-dual distributed algorithm for the computation of variational equilibria in games with a monotone VI mapping is developed. In [72] a deterministic quadratic aggregative game is studied subject to affine coupling constraints and the existence properties of the associated variational equilibria are investigated. For the computation of the GNE, an Asymmetric Projection Algorithm (APA) [63] is used. In [3] a two-level BR algorithm of similar form to the algorithm developed in [91] is proposed, enhanced with an outer loop that updates the dual variable thus guaranteeing convergence to a Wardrop Equilibrium. However, this algorithm cannot be used for convergence to NE. To this end, an one-level APA algorithm based on gradient step is used. A modified version of the APA algorithm is proposed in Chapter 4, in a randomized setting, for controlling the probabilistic feasibility of sets around equilibria.

2.5 Cooperative games and stability of allocations

The concepts analysed in this section are detailed in [12]. This section serves as a short introduction to cooperative games. Two notions are of crucial importance for such problems, namely fairness and stability. The results of this thesis focus entirely on the latter, specifically on finding solutions with provable stability guarantees in cooperative games affected by uncertainty. Before introducing such games it is important to establish some basic notions in a deterministic setting. Note that throughout the entire dissertation the words cooperative and coalitional are used interchangeably.

Unlike non-cooperative games, in cooperative games agents are allowed to form coalitions to satisfy their own selfish interests. If $\mathcal{N} = \{1, \ldots, N\}$ is the index set of all agents, we refer to a subset S of this set as a *coalition*. The entire index set $\mathcal{N}$ is called the *grand coalition*. In this thesis we focus on the so called characteristic function games defined as follows:

Definition 2.12 *A characteristic function game G is given by a pair $(\mathcal{N}, u)$, where $\mathcal{N}$ is a finite, non-empty set of agents and $u : 2^{\mathcal{N}} \to \mathbb{R}$ is a characteristic function, which maps each coalition $S \subseteq \mathcal{N}$ to a real number $u(S)$. The output of $u(S)$ is usually referred to as the value of the coalition.*

Note that in characteristic function games a value is assigned to the coalition and not to the agents that belong to it. This implies that the characteristic function game as defined above does not dictate how the value of the coalition should be divided among its members. In this thesis we will assume that the values of coalitions in characteristic function games can always be distributed among the members of the coalition. These games are called *transferable utility (TU) games*. We then impose the following assumption

Assumption 2.2 *Consider a TU game G as in Definition 2.12. Then, we assume that*

- *the value of the empty coalition is 0, i.e., $u(\emptyset) = 0$ and*
- *the value $u(S)$ is non-negative for all coalitions $S \subseteq \mathcal{N}$, i.e., agents form coalitions to make a profit.*

This set-up can easily be adapted to account for cases where the value $u(S)$ is nonpositive for all coalitions $S \subseteq \mathcal{N}$, i.e., when agents form coalitions to share costs.

In a cooperative game, each agent $i \in \mathcal{N}$ receives a payoff allocation $x_i \in \mathbb{R}$ depending in which coalition they participate and how the value of each coalition is divided among its members. Note that each agent can be a member of only one coalition. The vector $x = (x_i)_{i \in \mathcal{N}}$ is the payoff vector of the entire population. In order to study the stability of agents' allocations, i.e., the incentives of each agent to stay in the grand coalition, we introduce the notion of the core, one of the most widely used stability concepts in cooperative game theory defined as follows

Definition 2.13 *The core $C(G)$ of a game G, is given by*

$$C(G) = \{x \in \mathbb{R}^N \mid \textstyle\sum_{i \in \mathcal{N}} x_i = u(\mathcal{N}), \sum_{i \in S} x_i \geq u(S) \text{ for all } S \subset \mathcal{N}\}.$$

In other words, if the payoff vector x lies in the core set, then any agent is better off staying in the grand coalition as the total payoff of the agents in this case is better than the total payoff that they would have obtained having chosen to participate to any other subcoalition. Finding a solution to the core of a game G amounts to solving the following optimization program

$$\begin{cases} \min\limits_{x\geq 0} & 0 \\ \text{subject to} & \sum\limits_{i\in\mathcal{N}} x_i = u(\mathcal{N}), \\ & \sum\limits_{i\in S} x_i \geq u(S) \text{ for all } S \subset \mathcal{N}. \end{cases}$$

Note that the optimization program above is equivalent to the problem of finding a point within the feasible region. For this purpose, any constant (not necessarily zero) could have been used as a cost funtion. In several cases, the core of a game G might be empty. Furthermore, the formation of subcoalitions might come at a certain cost that should then be subtracted from the profit of this subcoalition. Other times, we might wish to incentivize all agents to work together, thus imposing certain penalties to the formation of subcoalitions. Such notions cannot be captured by the notion of the core. To this end, a relaxation of the concept of the core, the so called ε-core is introduced defined as follows:

Definition 2.14 *Given some $\varepsilon \geq 0$, the ε-core of a game G coincides with the set*

$$C_\varepsilon(G) = \{x \in \mathbb{R}^N : \textstyle\sum_{i\in\mathcal{N}} x_i = u(\mathcal{N}), \sum_{i\in S} x_i \geq u(S) - \varepsilon \text{ for any } S \subseteq \mathcal{N}\}.$$

□

Finding a solution to the ε-core of a game G amounts to solving the following optimization program

$$\begin{cases} \min\limits_{x\geq 0} & 0 \\ \text{subject to} & \sum\limits_{i\in\mathcal{N}} x_i = u(\mathcal{N}), \\ & \sum\limits_{i\in S} x_i \geq u(S) - \varepsilon \text{ for all } S \subset \mathcal{N}. \end{cases}$$

In practice, engineering problems modelled as multi-agent optimization programs or non-cooperative/ cooperative games are affected by uncertainty. This can be due

to external factors affecting the system such as disturbances or uncertainty inherent in the dynamics of the system under study. In the next section we introduce a well-known data-driven methodology to deal with uncertainty in such problems, the so called scenario optimization.

2.6 Scenario optimization

There are two approaches in scenario optimization, namely the *a priori* and the *a posteriori* approach. Both are detailed in the following subsections.

2.6.1 *A priori* scenario approach

Consider the following optimization program

$$P_\delta : \min_{x \in \mathbb{R}^d} \; c^T x \qquad (2.6)$$
$$\text{subject to } x \in \bigcap_{\delta \in \Delta} X_\delta,$$

where $x \in \mathbb{R}^d$ is the decision vector, δ is the uncertainty parameter taking values in the (possibly) unknown probability space $(\Delta, \mathscr{F}, \mathbb{P})$, where Δ is the support set, $\mathscr{F}$ a σ-algebra and $\mathbb{P}$ the probability distribution of δ. X_δ denotes a convex constraint set parameterized by δ. Vector $c \in \mathbb{R}^d$ corresponds to a given cost vector; it is taken to be linear without loss of generality, as in the alternative case an epigraphic reformulation could be performed to recast the program to the form of P_δ. Note that the uncertainty set can be of infinite cardinality, thus rendering P_δ a semi-infinite optimization program, which is very challenging to solve in general.

A method to solve this problem is the well-known worst-case approach. This approach, often formulated as a robust optimization program, requires the satisfaction of the constraints that correspond to every admissible realization of the uncertainty parameter $\delta \in \Delta$. However, apart from being computationally demanding in many cases [92], this method is also very conservative as in practice the worst-case uncertainty realization rarely occurs.

An alternative approach to reduce the conservatism imposed by the worst case scenario method can be accomplished by allowing one small portion of constraints to be

violated with a certain probability, thus giving rise to the well-known chance-constrained techniques [93]. These methods fall under the realm of stochastic programming [94]. The introduction of probability-based methods to analyse the robustness of systems has been investigated in several results in the literature. Specifically, [95] provides a methodology for the use of the uniform probability distribution in the case of no *a priori* provided statistical information of the uncertain parameters, and under mild geometric assumptions imposed on the density functions under study. In [96], uncertain linear time invariant systems are studied. Randomized polynomial-time algorithms are then developed to deal with the uncertainty and provide robust solutions in the probabilistic sense. Randomized algorithms are studied extensively in the book [97]. In [98], [99] the connections between stochastic robustness and control system design is investigated. Provided that a probabilistic approach is adopted, the computational complexity of many classes of robust control problems is significantly lower compared to that of the worst case approach [100]. In addition, the application of those methods is more straightforward as there is no need to make adaptations to the real engineering problem in order to make it fit a worst case paradigm. However, chance-constrained methods usually are tractable only under assumptions on the underlying probability distribution. Furthermore, we often do not know the underlying distribution of the uncertainty, nor its support. As such, following a worst-case paradigm and/or imposing certain assumptions on the geometry of the uncertainty set or the underlying probability distribution might lead to erroneous performance. This highlights the need for distribution-free results.

Many results in statistical learning theory rely on the concept of Vapnik-Chervonenkis (VC) dimension to provide probabilistic certificates for the region of the feasible space that satisfies a given set of constraints obtained through i.i.d. samples. The general theory is developed in [101], [102], while improved complexity bounds are provided in [103] with applications in control [104], [105]. An application of the theory in stochastic Model Predictive control (MPC) can also be found in [106].

A methodology that allows to compute approximations of the region of the feasibility space satisfying chance constrained constraints is the so called probabilistic scaling technique (see [107], [108] and references therein). This methodology has found

application in stochastic MPC in the recent paper [109]. The seminal works [110], [111] characterize probabilistically the region of the feasibility space that satisfies a set of sampled linear constraints. There are several works in the literature in which the samples are drawn in a sequential manner and the point of the feasibility space that satisfies the prescribed probabilistic constraints is also obtained sequentially by solving a finite number of optimization problems. Examples of these algorithms can be found in [97], [112], [113], [114].

In this thesis we provide distribution-free probabilistic feasibility guarantees by following a non-sequential approach, i.e., our theoretical results are based on the solution of a single problem using all available i.i.d. samples. We extend the theoretical framework of a rapidly evolving methodology in the field of statistical learning, the so called scenario approach. The scenario approach first introduced in [29] is a well-established method used for providing probabilistic guarantees for convex optimization programs under uncertainty. By taking only a finite number of samples $\delta_K = (\delta^{(1)}, \delta^{(2)}, \ldots, \delta^{(K)}) \in \Delta^K$ of the uncertain parameter, we can construct the so called scenario program, including only the constraints that correspond to these samples. In mathematical terms, the scenario program is given by:

$$\begin{aligned} \mathrm{P}_K : \min_{x \in \mathbb{R}^d} \; & c^T x \\ \text{subject to } & x \in \bigcap_{k=1}^{K} X_{\delta^{(k)}}. \end{aligned} \tag{2.7}$$

Apart from its intuitive appeal, the important feature of this program is its generalization properties, i.e., there is a high probability that for a sufficiently high number of samples a constraint that corresponds to a new sample will be satisfied by the solution of P_K. To study this property of the solution, a key concept known as the probability of constraint violation is introduced.

Definition 2.15 *The probability of violation of a point $x \in \mathbb{R}^n$ is a function $V : \mathbb{R}^d \to \mathbb{R}$ of the form*

$$V(x) = \mathbb{P}\{\delta \in \Delta : x \notin X_\delta\}.$$

□

The probability of violation is a measure of how robust the system under study is against yet unknown uncertainties. Assume that the drawn samples of the uncertainty concatenated in the vector δ_K are independent and identically distributed (i.i.d.), and that P_K admits a unique[1] solution denoted by x_K^*. The fundamental result in [21] states that:

$$\mathbb{P}^K\{\delta_K \in \Delta^K : V(x_K^*) > \varepsilon\} \leq \sum_{s=0}^{d-1} \binom{K}{s} \varepsilon^s (1-\varepsilon)^{N-s}. \tag{2.8}$$

where $\mathbb{P}^K = \prod_{k=1}^K \mathbb{P}$ is the product probability measure and $\varepsilon \in (0,1)$ is referred to as the violation level. In words, the aforementioned bound implies that with confidence at least $1 - \sum_{s=1}^{d-1} \binom{K}{s} \varepsilon^s (1-\varepsilon)^{N-s}$ the probability that the optimal solution of P_K violates the constraint corresponding to a new realization of the uncertainty δ, is at most ε. Moreover, it is shown that this bound is tight, i.e., it holds with equality for the class of fully-supported problems, i.e., for convex scenario programs where the number of support constraints is exactly equal to d. As stated in [115] , an alternative interpretation of (2.8)is the following: Fix $\beta, \varepsilon \in (0,1)$ and choose the number of samples K such that:

$$K \geq \frac{2}{\varepsilon}(d - 1 + \ln\frac{1}{\beta}). \tag{2.9}$$

Then, with confidence at least $1-\beta$, the probability of constraint violation is at most equal to ε[2]. Due to the logarithmic dependence of K on β, β can be chosen to be a very small quantity without a substantial change to the number of samples required. As opposed to one-level probability results, in this thesis we focus on two-level probability results of form similar to (2.8), i.e., on results where the probability of violation is bounded by a quantity ε with high confidence $1-\beta$ [116]. The bound in (2.8) builds on the notion of the so called support samples defined as follows.

Definition 2.16 *A sample $\delta^{(j)}$ is of support for the scenario program P_K if its removal changes the solution of the program.*

[1] In case multiple solutions exist, a deterministic tie-break rule could be adopted to single-out one minimizer.

[2] Note that $K \geq \frac{2}{\varepsilon}(d-1+\ln\frac{1}{\beta})$ can be replaced by the less conservative bound $K \geq \frac{1.59}{\varepsilon}(d-1+\ln\frac{1}{\beta})$ [113].

In other words, if we consider the scenario program P_{K-1}, where we do not take into account the constraint(s) that corresponds to sample $\delta^{(j)}$ and the solution of such program is different than the solution when all samples are in place, then $\delta^{(j)}$ is a support sample. It is shown in [20] that the number of support constraints for convex optimization programs is at most equal to the dimension d of the problem at hand. I

In [30], [115] a sampling and discarding procedure is suggested, where a certain number of samples is extracted, but an *a priori* fixed number of them (in fact the ones that if removed would lead to an improved performance) is discarded prior to solving the scenario program. The number of discarded samples offers an additional tuning freedom to trade performance for constraint feasibility.

2.6.2 *A posteriori* scenario approach

Very often the optimization programs encountered in practical applications are not fully-supported and it is likely that there are significantly fewer support constraints than optimization variables (see [117] and [46]). In the previous subsection the violation level ε was a deterministic fixed quantity set to satisfy the desired robustness standards and did not depend on the multi-sample that was drawn. In [23] an alternative approach, the so-called wait-and-judge scenario approach, is introduced, providing guarantees in an *a posteriori* fashion. In this new set-up the violation level ε can be viewed as a function of the number of support constraints s_K^* that is empirically computed once the optimal solution of P_K is calculated, thus implicitly depending on the samples δ_K. As such, both $V(x_K^*)$ and $\varepsilon(s_K^*)$, are random quantities. In [24] a general framework for providing *a posteriori* probabilistic bounds even for non-convex optimization programs is developed.

Specifically, for any $\beta \in (0,1)$, the violation level is defined as a function $\varepsilon(s) : \{0,...,K\} \to [0,1]$ such that:

$$\varepsilon(K) = 1 \text{ and } \sum_{k=0}^{K-1} \binom{K}{s} (1-\varepsilon(s))^{K-s} = \beta. \tag{2.10}$$

Under very mild assumptions the probability of constraint violation can then be bounded in an *a posteriori* fashion, i.e.,

$$\mathbb{P}^K\{\delta_K \in \Delta^K : V(x_K^*) > \varepsilon(s_K^*)\} \leq \beta. \tag{2.11}$$

The function $\varepsilon(\cdot)$ and hence the quality of the bound in (2.10) is refined in [23], however, requiring that the underlying scenario program is non-degenerate. Nondegeneracy implies that for any size m of the multi-sample $\delta_k, k = 1, ..., m$, the solution to the scenario program with all constraints in place coincides with the solution to the program where only the support constraints are present with probability 1;

Furthermore, we introduce the notion of the compression set and the compression function given as follows. These notions are key for the subsequent developments.

Definition 2.17 *1. For any K, an algorithm is a mapping A_K that is fed with the multi-sample δ_K and outputs a unique solution S_K.*

2. *Given a multi-sample $\delta_K \in \Delta^K$, a compression set $\mathcal{C}$ is a subset of the elements of the entire multi-sample δ_K with cardinality $s \leq K$ such that the solution returned by A_K when fed with the vector that comprises only the elements of $\mathcal{C}$ is the same with the one obtained when the entire multi-sample is used.*

3. *A compression function B_K is a function that takes as input all the samples and returns as output only these samples that constitute an irreducible[3] compression set.* □

From Definition 2.17 note that the unique solution is not restricted to a unique singleton but can in fact be a unique set. A compression set of minimal cardinality among all irreducible compression sets is called minimal compression set. Note that the notions of compression set and compression function in Definitions 2.17.2, 2.17.3 are respectively referred to as support subsample and support subsample function in [24]. In this more general set-up, a sample δ_j is said to be of support if its removal from δ_K changes the

[3] A compression set $\mathcal{C}$ is said to be irreducible if no element can be further removed from $\mathcal{C}$ leaving the solution unchanged.

output of algorithm A_K. To establish an equivalence between the notion of the minimal compression set and the set of support samples we usually assume that the problems considered are non-degenerate.

In certain cases we might be interested in investigating the robustness properties collectively for all the points of a solution set S_K produced by Algorithm A_K to yet unseen samples, in other words in quantifying the probability that a new sample $\delta \in \Delta$ is drawn such that the constraint X_δ defined by this sample is not satisfied by some given point $x \in S_K$.

To this end, we extend the notion of the probability of violation as presented in Definition 2.15 to account for sets of feasible solutions.

Definition 2.18 *For any set $S_K \subseteq \mathbb{R}^d$ we define the probability of violation of the set S_K as a mapping $\mathbb{V} : 2^{\mathbb{R}^d} \to [0,1]$ given by the following relation:*

$$\mathbb{V}(S_K) = \sup_{x \in S_K} V(x).$$

□

Another way to obtain probabilistic certificates for a region is to make the simplification that the region is just a collection of points (e.g. a grid of points of the feasibility space), i.e., providing probabilistic certificates for a set of finite cardinality [114], [118]. In some cases, the samples of the feasibility space are obtained by means of random sampling [119], [104]. An application of this theory for probabilistic guarantees for a family of controllers can be found in [120]. Our approach is different as we focus on the feasible set obtained from the intersection of random sampled constraints. Note that standard statitical learning approaches [101] and probabilistic scaling [108] can also be used to analyze the probability of violation of a set. A connection of our work with these approaches is left for future work.

3

Probabilistic certificates for sets of feasible solutions

3.1 Introduction

Many times in problems of practical interest multiple (possibly infinite) optimal solutions can exist. The non-uniqueness of the optimal solution introduces additional challenges, especially when the problem under study is affected by uncertainty. Specifically, within the realm of the scenario approach, a general framework for the provision of collective probabilistic feasibility guarantees for the entire set of solutions has not been developed, as standard results in the literature usually focus on a single optimizer obtained after the application of a convex tie-break rule. Additional challenges in uncertain multi-agent problems are often imposed by the numerical implementation of the solution-seeking algorithm, which hinder the exact quantification of the optimal solution. Convergence to a suboptimal solution can be the result of the asymptotic nature of most existing iterative schemes and time restrictions imposed on the optimization procedure. As such, a stopping criterion is commonly used for the implementation of the algorithm. This in turn introduces sensitivity issues to the computation of the solution. Even under the restrictive assumption that the optimal decision for the problem at hand is unique, the trajectory of such algorithms, represented by the values of their iterations, depends on the initial conditions of the system, which many times cannot be controlled beforehand. As such,

depending on the initial conditions the algorithm can converge to different points within a region parameterized by the error bound of the stopping criterion. Finally, as observed across disciplines ranging from engineering to socioeconomics, agents participating in a multi-agent system might not behave as fully rational entities but rather exhibit bounded rationality behaviour. This in turn leads to deviations from the given solution.

In order to alleviate the aforementioned challenges, we establish in this chapter a more general theoretical framework that can be adapted to provide collective guarantees for solution sets, sets of suboptimal solutions or sets of deviations around a given solution. A requirement for our theoretical developments is that the set under study is feasible with respect to the observed constraint realizations of the uncertain parameter. This assumption is justified in the sense that usually we wish the decision made to satisfy the constraints available through data.

We start by considering a feasibility program with uncertain convex constraints. We, then, leverage the results of [24] in order to provide robustness certificates for the entire feasibility region. This result is interesting per se, as it complements the results in [24] allowing to provide guarantees collectively for a set of points, thus departing from the existing stream of literature in the scenario approach which typically refers to a single solution (e.g., see [20–24], [121]). Apart from the study of uncertain multi-agent programs, this general set-up can find applications in multi-agent games with uncertain convex constraints, providing guarantees collectively for the entire set of equilibria or feasible solutions close to it. As such, the considered framework offers several advantages over the standard results in the scenario approach that make it appealing also from a practical perspective. In contrast with the *a priori* result in [31], the provided guarantees are *a posteriori* and hold for the entire feasible set or a subset of interest as detailed in Remark 3.2.

This chapter is organized as follows: Section 3.2 introduces feasibility programs with uncertain convex constraints. The main theory is developed for providing probabilistic feasibility guarantees for the entire feasibility region and subsets of it. Section 3.3 focuses on the application of our main results to multi-agent optimization programs. A 2-dimensional feasibility problem and a multi-agent model of the electric vehicle charging

control problem with arbitrary cost and uncertainty in the constraints is studied under the lenses of the developed theory. In Section 3.4 our theoretical results are applied to non-cooperative games affected by uncertainty. Numerical studies involving the electric vehicle charging game under the presence of uncertainty corroborate our theoretical results. Finally, Section 3.5 concludes the chapter summarizing our main results.

3.2 Feasibility programs with uncertain convex constraints

Let $\mathcal{N} = \{1, \ldots, N\}$ be the index set of all agents, where N denotes their total number and x_i the strategy of agent i taking values in the set $X_i \subseteq \mathbb{R}^n$. We denote $x = (x_i)_{i\in\mathcal{N}} \in X = \prod_{i\in\mathcal{N}} X_i \subseteq \mathbb{R}^{nN}$ the collection of all agents' strategies[1]. Similarly, the vector $x_{-i} = (x_j)_{j\in\mathcal{N}, j\neq i} \in \prod_{j\in\mathcal{N}, j\neq i} X_j \subseteq \mathbb{R}^{n(N-1)}$ denotes the collection of the decision vectors of all agents' strategies except for that of agent i. Let δ be an uncertain parameter defined on the probability space $(\Delta, \mathscr{F}, \mathbb{P})$, where Δ is the sample space, equipped with a σ-algebra $\mathscr{F}$ and a probability measure $\mathbb{P}$. Furthermore, let $\delta_K \in \Delta^K$ be a finite collection of K i.i.d. scenarios/realisations of the uncertain vector δ, where Δ^K is the cartesian product of multiple copies of the sample space Δ; finally, $\mathbb{P}^K$ is the associated product probability measure. Consider the following feasibility program

$$\mathrm{P}_{\Delta}^{feas}: \text{ find } x \in X \text{ subject to } x \in \bigcap_{\delta\in\Delta} X_\delta, \tag{3.1}$$

where X_δ is an uncertain constraint parameterized by the uncertainty realization $\delta \in \Delta$. We seek to provide probabilistic guarantees for all feasible solutions of $\mathrm{P}_{\Delta}^{feas}$. Solving instances of $\mathrm{P}_{\Delta}^{feas}$ is generally hard without imposing any further assumptions on the uncertainty support Δ and/or the underlying probability distribution $\mathbb{P}$ [21]. We address this by considering the following approximation of $\mathrm{P}_{\Delta}^{feas}$, the so called scenario program, formulated upon a finite set δ_K of i.i.d. samples from Δ[2]. As such, we obtain the

[1] Note that the notion of strategy here does not necessarily refer to a control law but rather to a decision represented by a vector in $\mathbb{R}^n$.

[2] The results presented assume no knowledge of the probability distribution or the support set of the uncertainty, i.e., they are distribution-free. However, we stress that these developments hold only for problems for which the data satisfies the i.i.d. assumption. This is of course an approximation of reality, however, many models based on the i.i.d. assumption have successfully been used in applications. Still, this assumption can be quite restrictive. Extending the theory for non-i.i.d cases is a non-trivial task and is thus left for future work.

following scenario program

$$\mathrm{P}_K^{feas}: \text{ find } x \in X \text{ subject to } x \in \bigcap_{k=1}^{K} X_{\delta^{(k)}}. \tag{3.2}$$

Our results leverage the following assumption:

Assumption 3.1 *1. The deterministic constraint set X is non-empty, compact*[3] *and convex.*

2. *For any sample $\delta \in \Delta$, we have that $X_\delta = \{x \in \mathbb{R}^d : g(x,\delta) \leq 0\}$, where $g : \mathbb{R}^d \times \Delta \to \mathbb{R}^q$ is a vector-valued convex function.*

3. *For any fixed multi-sample $\delta_K \in \Delta^K$ the convex set $C_K = \{\bigcap_{k=1}^{K} X_{\delta^{(k)}}\} \bigcap X = \{x \in X : g(x,\delta^{(k)}) \leq 0,$ for any $k = 1,\ldots,K\}$ is non-empty.*

Note that if the support set Δ is not bounded, this can in certain cases lead to feasibility issues as the number of samples K increases (see Section 7 in [113]). In such a case the sampling and discarding approach can be employed to alleviate this issue. One could also analyse feasibility through application of Helly's theorem (see [122]). Assumption 3.1 guarantees that P_K^{feas} admits at least one solution for any chosen multi-sample $\delta_K \in \Delta^K$. The set C_K denotes the feasibility domain of problem P_K^{feas}. It follows that P_K^{feas} can be equivalently written as

$$\begin{aligned} &\text{find } x \\ &\text{subject to } x \in C_K. \end{aligned} \tag{3.3}$$

We are interested in investigating the robustness properties collectively for all the points of this domain to yet unseen uncertainty realizations. As discussed in Chapter 2, the probability of violation for a point is given by

$$V(x) = \mathbb{P}\Big\{\delta \in \Delta : \; x \notin X_\delta\Big\}. \tag{3.4}$$

In other words, $V(x)$ in (3.4) quantifies the probability of occurrence of a sample $\delta \in \Delta$ such that the corresponding constraint X_δ is not satisfied by $x \in C_K$. By Assumption 3.1,

[3] Compactness of a set means that it is closed and bounded.

the probability of violation can be equivalently written as $V(x) = \mathbb{P}\Big\{\delta \in \Delta :\ g(x,\delta) > 0\Big\}$. We can now define the probability of violation for the entire convex set C_K. Let $\mathscr{C} \subseteq 2^X$ be the set of all non-empty, compact and convex sets contained in X. For any $C_K \in \mathscr{C}$ the probability of violation of the set C_K is defined following Definition 2.18 in Chapter 2 as $\mathbb{V}(C_K) = \sup_{x \in C_K} V(x)$. Furthermore, consider an algorithm A_K as in Definition 2.17.1 that returns the feasible set C_K under study and a compression function B_K for such an algorithm. In the following, we assume that B_K always returns the minimal compression set. Such a requirement is not restrictive, and only rules out degenerate cases where samples lead to constraint accumulation [23]. If the underlying probability distribution admits a continuous density, this case happens with 0-probability.

We note that, in our set-up, a minimal compression set consists of those samples that are of support for the entire feasibility region. The following definition formalizes this notion

Definition 3.1 *(Support sample for the feasibility region) A sample $\delta_j \in \delta_K$ is said to be of support for the feasible set C_K of P_K if it belongs to the minimal compression set returned by B_K.*

As such, if $\delta^{(j)}$ belongs to the minimal compression set, then its removal may lead to an enlargement of the feasible region, i.e., $\{\bigcap_{k=1}^{K} X_{\delta^{(k)}}\} \setminus X_{\delta^{(j)}} \supset C_K$.

The number of support samples of the feasible region or, equivalently, the cardinality of the minimal compression set is denoted as F_K. The constraints that correspond to indices from the minimal compression set can be alternatively viewed as an adaptation of the notion of the facets of a polytope (see Definition 2.1 in [123]) to the more general case of compact and convex sets. Note that a single constraint may give rise to multiple "facets". Figure 3.1 illustrates the concept of the minimal compression set by showing the feasible region formed by random convex constraints. Note that only the indices of the samples $\delta^{(3)}, \delta^{(5)}, \delta^{(6)}, \delta^{(7)}$ belong to the minimal compression set, since if we feed only these samples as input into the algorithm A_K the feasible set C_K is returned.

Another important notion used in our work is the notion of an extreme point. An extreme point can be viewed as an extension of the vertex of a polytope for arbitrary

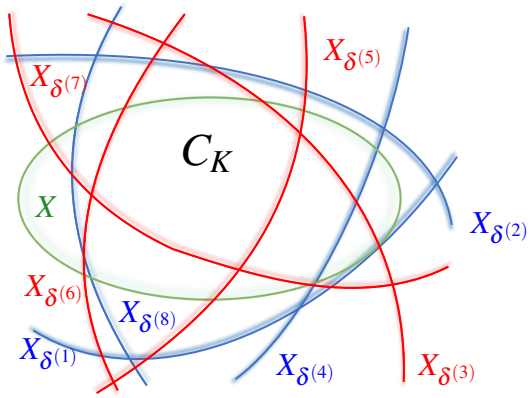

Figure 3.1: The feasibility region C_K and its connection with random convex constraints produced by eight i.i.d. samples $\{\delta^{(1)}, \delta^{(2)}, \ldots, \delta^{(8)}\}$. The constraint in green corresponds to the deterministic constraint X. Note that only the indices of the samples $\delta^{(3)}, \delta^{(5)}, \delta^{(6)}, \delta^{(7)}$ belong to the minimal compression set since their corresponding constraints (in red) form, along with the deterministic constraint X, the boundary of C_K.

compact and convex sets and is defined as a point which is not in the interior of any line segment lying entirely in the set. This property is formally presented in the following definition.

Definition 3.2 *(Extreme points) [124]. An extreme point of a convex set C is a point $x \in C$ for which the following property holds: If x can be written as a convex combination of the form $x = \lambda x_1 + (1-\lambda)x_2$ with $x_1, x_2 \in C$ and $\lambda \in [0,1]$, then $x_1 = x$ or $x_2 = x$.*

Note that the number of extreme points of a convex set depends on the geometry of the set under study and can also be infinite, e.g., in the case of a d-dimensional sphere.

Our work focuses on compact and convex sets defined over a finite-dimensional space, where the following theorem can be applied.

Theorem 3.1 *(Minkowski – Caratheodory Theorem) [124]. Let C be a compact convex subset of $\mathbb{R}^d$ of dimension d. Then any point in C is a convex combination of at most $d+1$ extreme points.*

We equip the set of extreme points of the convex set C_K with indices and denote this set of indices $E(C_K)$. It is important to emphasize that the dependence of the convex set C_K on the multi-sample δ_K implies that $|B_K|$ (and $|E(C_K)|$) are random variables that depend on $\delta_K \in \Delta^K$.

Next we define the set

$$\begin{aligned}\mathscr{C}_\delta &= \{C \in \mathscr{C} : g(x_j, \delta) \leq 0,\ \forall\ j \in E(C)\} \\ &= \{C \in \mathscr{C} : C \subseteq X_\delta\},\end{aligned} \tag{3.5}$$

of all the non-empty, compact and convex sets where elements C satisfy the constraint associated with the sample $\delta \in \Delta$. Note that if all the extreme points of the set satisfy the inequality $g(\cdot, \delta) \leq 0$, then every point $x \in C$ of the set satisfies it as well. To see this, note that x can always be expressed as a convex combination of the set's extreme points.

Our aim is to provide probabilistic guarantees for a non-empty, convex and compact set C_K constructed by the intersection of K random realizations of the uncertain convex constraint $X_\delta = \{x \in \mathbb{R}^d : g(x, \delta) \leq 0\}$, where $g : \mathbb{R}^d \times \Delta \to \mathbb{R}^q$ is a convex function with respect to the decision variable x.

The following lemma shows how the probability of violation of a convex set is related to the probability of violation of its extreme points.

Lemma 3.1 *Consider a fixed multi-sample $\delta_K \in \Delta^K$ and let $E(C_K)$ be the set of extreme points of C_K. Then,*

$$\mathbb{V}(C_K) \leq \mathbb{P}\left\{ \bigcup_{j \in E(C_K)} \left\{ \delta \in \Delta :\ g(x_j, \delta) > 0 \right\} \right\}. \tag{3.6}$$

Proof: For a fixed multi-sample $\delta_K \in \Delta^K$ consider an arbitrary point $x \in C_K$. Then,

the following inequalities are satisfied

$$
\begin{aligned}
V(x) &= \mathbb{P}\Big\{\delta \in \Delta:\ x \notin X_\delta\Big\} = \mathbb{P}\Big\{\delta \in \Delta:\ g(x,\delta) > 0\Big\} \\
&\overset{(i)}{=} \mathbb{P}\Big\{\delta \in \Delta:\ g(\sum_{j\in I_{d+1}} \lambda_j x_j, \delta) > 0\Big\} \\
&\overset{(ii)}{\leq} \mathbb{P}\Big\{\delta \in \Delta:\ \sum_{j\in I_{d+1}} \lambda_j g(x_j,\delta) > 0\Big\} \\
&\leq \mathbb{P}\Big\{\delta \in \Delta:\ \sum_{j\in I_{d+1}} \lambda_j \max_{j\in I_{d+1}} g(x_j,\delta) > 0\Big\} \\
&= \mathbb{P}\Big\{\delta \in \Delta:\ \max_{j\in I_{d+1}} g(x_j,\delta) > 0\Big\} \\
&= \mathbb{P}\Big\{\bigcup_{j\in I_{d+1}} \Big\{\delta \in \Delta:\ g(x_j,\delta) > 0\Big\}\Big\} \\
&\overset{(iii)}{\leq} \mathbb{P}\Big\{\bigcup_{j\in E(C_K)} \Big\{\delta \in \Delta:\ g(x_j,\delta) > 0\Big\}\Big\},
\end{aligned}
\tag{3.7}
$$

Equality (i) is derived from Theorem 3.1, where the set under study is the convex set C_K. In our case, the Minkowski-Caratheodory theorem states that any arbitrary point of the set $x \in C_K$ can be represented as a convex combination of at most $d+1$ extreme points of C_K, which means that there exists a subset of extreme points $\{x_j\}_{j\in I_{d+1}} \subseteq \{x_j\}_{j\in E(C_K)}$ such that $x = \sum_{j\in I_{d+1}} \lambda_j x_j$, where $\sum_{j\in I_{d+1}} \lambda_j = 1$ and $\lambda_j \geq 0,\ \forall\ j \in I_{d+1}$. Equality (ii) stems from the fact that g is a convex function of x for any given $\delta \in \Delta$. The last inequality follows from the fact that I_{d+1} is a set of indices corresponding to extreme points and as such is a subset of $E(C_K)$. Since (3.7) holds for all $x \in C_K$, it can equivalently be written as

$$
\mathbb{V}(C_K) = \sup_{x\in C_K} V(x) \leq \mathbb{P}\Big\{\bigcup_{j\in E(C_K)} \Big\{\delta \in \Delta:\ g(x_j,\delta) > 0\Big\}\Big\},
$$

which concludes the proof. ■

Lemma 3.1 shows that the probability of constraint violation of C_K is bounded by the probability that at least one of the extreme points violates the constraint and is key towards providing probabilistic feasibility guarantees for the feasible region, a result established in the following theorem.

Theorem 3.2 *Consider Assumption 3.1 and any A_K that is fed with δ_K and returns the region C_K. Furthermore, consider the compression function B_K of algorithm A_K. Fix*

$\beta \in (0,1)$ and define the violation level function $\varepsilon : \{0,...,K\} \to [0,1]$ such that

$$\varepsilon(K) = 1 \text{ and } \sum_{s=0}^{K-1} \binom{K}{s} (1-\varepsilon(s))^{K-s} = \beta. \tag{3.8}$$

We have that

$$\mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{V}(C_K) > \varepsilon(s^*) \Big\} \leq \beta,$$

where $s^* = F_K$ is the number of support samples according to Definition 3.1.

Proof: From Lemma 3.1, for any multi-sample $\delta_K \in \Delta^K$ and for any cardinality (not necessarily minimal) of the compression set $s \in \{1,\dots,K\}$, the following inequalities are satisfied:

$$\begin{aligned}
&\mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{V}(C_K) > \varepsilon(s) \Big\} \\
&\leq \mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{P}\Big\{ \bigcup_{j \in E(C_K)} \Big\{ \delta \in \Delta : g(x_j, \delta) > 0 \Big\} > \varepsilon(s) \Big\} \Big\} \\
&= \mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{P}\Big\{ \delta \in \Delta : \ \exists j \in E(C_K), g(x_j, \delta) > 0 \Big\} > \varepsilon(s) \Big\} \\
&= \mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{P}\Big\{ \delta \in \Delta : \ C_K \not\subseteq X_\delta \Big\} > \varepsilon(s) \Big\},
\end{aligned} \tag{3.9}$$

where the last equality is due to (3.5). Define now an algorithm A_K as in Definition 2.17, that returns the convex set confined by the feasibility region of C_K. By construction, A_K satisfies Assumption 1 of [24], since for any multi-sample δ_K it holds that $A_K(\delta_K) \in \mathscr{C}_{\delta^{(k)}}$, for all $k = 1,\dots,K$. The satisfaction of Assumption 3.1 paves the way for the use of Theorem 1 of [24]. In particular, for $k \geq k^* = F_K$, Theorem 1 of [24] implies that the right-hand side of (3.9) can be upper bounded by β.

As such, we have that

$$\begin{aligned}
&\mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \mathbb{P}\Big\{ \delta \in \Delta : C_K \not\subseteq X_\delta \Big\} > \varepsilon(s^*) \Big\} = \\
&\mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \mathbb{P}\Big\{ \delta \in \Delta : C_K \notin \mathscr{C}_\delta \Big\} > \varepsilon(s^*) \Big\} \leq \beta.
\end{aligned} \tag{3.10}$$

From (3.9) and (3.10) we obtain that:

$$\mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{V}(C_K) > \varepsilon(s^*) \Big\} \leq \beta, \tag{3.11}$$

thus concluding the proof. ■

As mentioned in [24] one can choose ε by splitting β evenly among the terms in the sum of (3.8). In this case, we have that

$$\varepsilon(s) = \begin{cases} 1, & s = K \\ 1 - \left(\frac{\beta}{K\binom{K}{s}}\right)^{\frac{1}{K-s}}, & \text{otherwise} \end{cases}.$$

The result of Theorem 3.2 implies that with confidence at least $1-\beta$, the probability that there exists at least one feasible point of C_K that violates the constraints for a new realization $\delta \in \Delta$, is at most equal to $\varepsilon(s^*)$. Note that our guarantees trivially hold for any subregion of the feasible set. However, even though we have extended the theoretical framework of the scenario approach to account for convex feasible regions, the compression sets cannot be easily computed in the general convex case. To circumvent this challenge we consider the presence of affine constraints.

3.2.1 The polytopic case

In this section we assume the presence of affine constraints only, i.e., that the randomized constraint set is polytopic. A polytope $\Pi \subset \mathbb{R}^d$ can be expressed by its H-representation, i.e., the intersection of a finite number of halfspaces, and also as the convex hull of its vertex set $v(\Pi) = \{x_1, ..., x_Q\}$ i.e, $\Pi = \text{conv}(v(\Pi)) = \{\sum_{j=1}^{Q} x_j \lambda_j : \sum_{j=1}^{Q} \lambda_j = 1, \lambda_j \geq 0,\ j = 1, ..., Q\}$, where $v(\cdot)$ denotes the set of vertices of the polytope. This representation is generally known as V-representation [123], [125]. We replace Assumption 3.1 with the following one:

Assumption 3.2 *Consider Assumption 3.1 and further assume that X is polytopic and $g(x,\theta) = a^T x - b \leq 0$, where $a \in \mathbb{R}^d$, $b \in \mathbb{R}$ and $\theta = (a^T\ b) \in \mathbb{R}^{d+1}$.*

We denote polytopic feasibility sets with Π and Π_K, rather than C and C_K, respectively. Under Assumption 3.2, the cardinality of the minimal compression set returned by B_K coincides by definition with the number of random facets (see Definition 2.1 in [123]) of the polytope. Then, Theorem 3.2 gives rise to the following corollary.

Corollary 3.1 *Consider Assumption 3.2 and the set-up of Theorem 3.2. Fix* $\beta \in (0,1)$ *and define the violation level* $\varepsilon : \{0,...,K\} \to [0,1]$ *as a function such that*

$$\varepsilon(K) = 1 \text{ and } \sum_{s=0}^{K-1} \binom{K}{s} (1-\varepsilon(s))^{K-s} = \beta. \tag{3.12}$$

We have that

$$\mathbb{P}^K \Big\{ \delta_K \in \Delta^K : \ \mathbb{V}(\Pi_K) > \varepsilon(s^*) \Big\} \leq \beta,$$

where $s^* = F_K$ *is the number of facets of* Π_K.

Note that, even though in the proof of our theorem we also use the extreme points (vertices) of the polytope, only the number of facets is needed to provide probabilistic guarantees for the entire feasibility region. This feature is appealing from a computational point of view as, in many practical cases, the constructed polytope has a significantly smaller number of facets than extreme points. To illustrate this, consider a finite horizon multi-agent control problem with N agents, where each agent's decision is subject to upper and lower bounds at each time instance $t \in \{1,\dots,n\}$. Hence, for a multi-sample $\delta_K \in \Delta^K$, the feasibility domain of the problem is a hyperrectangle whose number of facets $F = 2Nn$ grows linearly with respect to the number of decision variables, while the number of vertices is given by $V = 2^{Nn}$, which grows at an exponential rate with respect to Nn. Such constraints arise in several engineering applications [16], [48], [49] including the electric vehicle charging control problem presented in this chapter.

Note that, as the dimension of the decision vector increases, evaluating the minimal compression set becomes computationally challenging. Furthermore, in some cases, removing redundant constraints could be challenging from a numerical point of view as the procedure is highly dependent on the precision used to solve the involved linear optimization problems. However, several efficient algorithms have been proposed for detecting redundant constraints out of the initial set of affine constraints. The currently most efficient algorithm for redundancy detection is Clarkson's algorithm. [126]. Reducing the computational complexity of Clarkson's algorithm is still an active research area in computational geometry and combinatorics. One recent noteworthy attempt can be found in [127].

Remark 3.1 *Corollary 3.1 allows retrieving the support samples of the feasibility region by enumerating its facets. Calculating these is straightforward for a wide class of problems like those that have the structure of the example in Section 3.3.2 (see also [3], [16], [60]). The relevance of this result can be seen by considering, for example, the case where an iterative algorithm is obtained to return an optimizer/feasible solution of the underlying problem. The quantification of the support samples using the methodology suggested in [24] can be challenging due to the numerical sensitivity of the latter and the fact that it would require calling the iterative (typically with asymptotic convergence guarantees) algorithm multiple times. This difference becomes more pronounced if the objective function of the underlying problem is non-convex, thus possibly exhibiting multiple local optima.*

In case we wish to provide guarantees only for a specific subset of the feasible region, the following remark clarifies how this can be achieved by leveraging our results.

Remark 3.2 *Let A_K be an algorithm that returns only a subset of the feasible region. Following the same steps with the proof of Theorem 3.2, we can provide probabilistic feasibility guarantees for the subset of interest with B_K returning the support samples that correspond to this subset and not to the entire feasibility region. This might be useful in cases where we have prior knowledge that a certain subregion of the feasible domain is of importance, e.g., a neighbourhood around the optimizer and we wish to investigate the probabilistic feasibility of this region only. In case the cardinality of the compression set for the region is less than that of the entire feasible domain, this leads to less conservative feasibility certificates.*

3.3 Numerical studies

3.3.1 Feasibility of a random 2-dimensional polytope

We initially apply our results to a 2-dimensional example of a polytope constructed by the intersection of random halfspaces of the form $a_1x_1 + a_2x_2 - b \leq 0$, where a_1, a_2 and b are scalars following uniform distributions with support $[-4,4]$, $[-4,4]$ and $[10,15]$, respectively. Each sample δ is defined as a vector $\delta = (a_1, a_2, b) \in \mathbb{R}^3$.

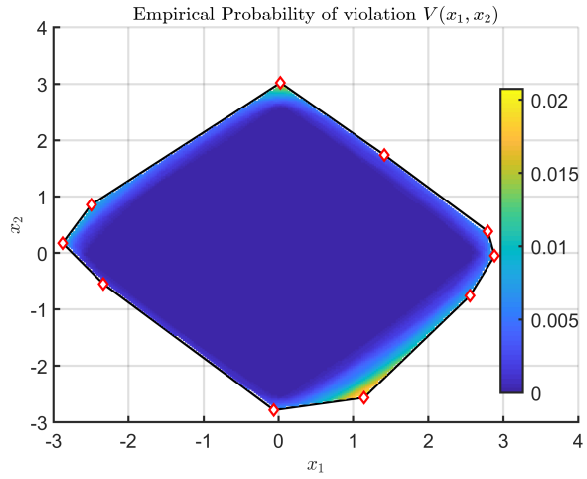

Figure 3.2: The empirical probability of violation of each point of the polytope under study constructed by 100 random realizations of the affine constraints $a_1x_1 + a_2x_2 - b \leq 0$. The number of grid points is $R = 48633$, while 10000 test samples are used. Note that the higher probability of violation occurs at one of the vertices (red diamonds).

To test the validity of our theoretical bounds in practice we need to compute the probability of violation $\mathbb{V}(\Pi_K)$ (Π_K in this instance is the randomized polytope produced by the intersection of K random constraints of the form $a_1x_1 + a_2x_2 - b \leq 0$) and compare it with the guarantees provided by Theorem 3.2. Let $\hat{\mathbb{V}}(\Pi_K)$ denote an empirical estimate of $\mathbb{V}(\Pi_K)$ and $\hat{V}(x)$ an empirical estimate of $V(x)$, where x is a point in Π_K. In particular, the former is obtained by gridding the polytope using a large enough number of points $x_r, r = 1, \ldots, R$ that cover the entire polytope and is given by $\hat{\mathbb{V}}(\Pi_K) = \max_{r=1,\ldots,R} \hat{V}(x_r)$. A sufficiently large number of test samples is used, different from those used to construct the polytope.

To this end, we generate a total number of $K = 100$ samples for the construction of the random polytope shown in Figure 3.2 and use 10000 test samples to compute the empirical probability of violation for each point of the grid. We note that the highest probability of violation occurs at a polytope vertex; the following lemma provides some theoretical support to this numerical evidence, by showing that the worst-case probability of violation is proportional to the worst-case vertex violation.

Lemma 3.2 *Consider the sets X and $X_{\delta^{(k)}}, k = 1, \ldots, K$, that satisfy Assumption 3.1. Then*

for any given multisample $\delta_K \in \Delta^K$, there exists a vertex $\hat{x} \in v(\Pi_K)$ such that:

$$\mathbb{V}(\Pi_K) \leq (d+1)\mathbb{P}\Big\{\delta \in \Delta : \ \hat{x} \notin X_\delta\Big\} = (d+1)V(\hat{x}).$$

Proof: Consider a fixed multisample and any arbitrary point $x \in \Pi_K$. Then, the following inequalities hold

$$\begin{aligned}
V(x) &\leq \mathbb{P}\Big\{ \bigcup_{j\in I_{d+1}} \Big\{\delta \in \Delta : \ g(x_j,\delta) > 0\Big\}\Big\} \\
&\leq \sum_{j\in I_{d+1}} \mathbb{P}\Big\{\delta \in \Delta : \ g(x_j,\delta) > 0\Big\} \\
&\leq \sum_{j\in I_{d+1}} \max_{j\in I_{d+1}} \mathbb{P}\Big\{\delta \in \Delta : \ g(x_j,\delta) > 0\Big\} \\
&= (d+1) \max_{j\in I_{d+1}} \mathbb{P}\Big\{\delta \in \Delta : \ g(x_j,\delta) > 0\Big\} \\
&\leq (d+1) \max_{j=1,\ldots,Q_K} \mathbb{P}\Big\{\delta \in \Delta : \ g(x_j,\delta) > 0\Big\} = (d+1)\, V(\hat{x}),
\end{aligned}$$

where $V(\hat{x}) = \max_{j=1,\ldots,Q_K} \mathbb{P}\Big\{\delta \in \Delta : \ g(x_j,\delta) > 0\Big\}$ is the maximum probability of violation among the vertices $x \in v(\Pi_K)$. This concludes our proof. ■

The inequality $V(x) \leq (d+1)\, V(\hat{x}),\ \forall\, x \in \Pi_K$ can be equivalently stated as: There exists $\hat{x} \in v(\Pi_K)$ such that $\mathbb{V}(\Pi_K) = \sup_{x\in\Pi_K} V(x) \leq (d+1)V(\hat{x})$, which means that to bound the probability of violation of the entire polytope, we only need to know the number of decision variables and the vertex with the highest probability of violation. This relation forms a bridge between the probabilities of violation of a vertex and the associated polytope.

Finally, we validate the derived theoretical bounds against the empirical probability of violation of 50 independent realizations of polytopes. This means that each polytope is constructed using a different multi-sample δ_K of size $K = 2000$. By keeping the same value for β, as in the previous case, we count the number of facets F_K of each polytope and then compute the theoretical bound of the violation level that corresponds to it. Subsequently, using 20000 test samples we compute an empirical estimate of the probability of violation for each polytopic realization, as outlined above. If there is more than one polytope among the 50 that has the same number of facets, we choose the one

with the maximum empirical probability of violation. As illustrated in Figure 3.3, $\varepsilon(s)$ constitutes an upper bound for any of the computed empirical probabilities of violation.

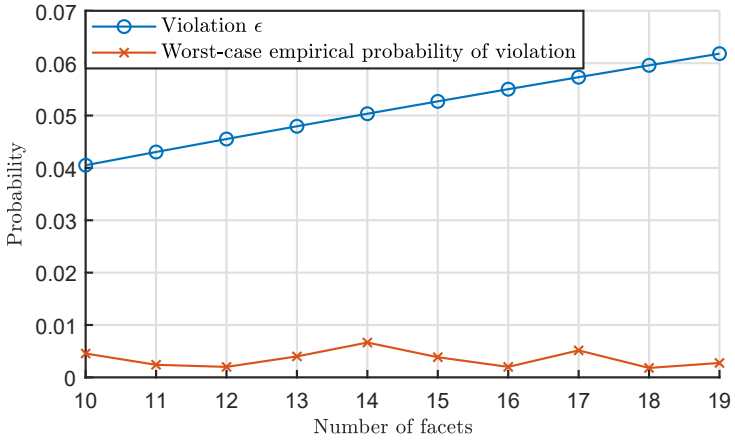

Figure 3.3: Comparison of the theoretical $\varepsilon(s)$ derived from Theorem 3.2 (blue) and the worst-case empirical probability of violation of the entire polytope that corresponds to each s (red). We generated 50 different polytopes, each of them using 2000 samples, while for each one the worst case probability corresponds to the grid point with the highest probability of constraint violation. The confidence level was fixed to $\beta = 10^{-6}$. Note that the non-monotonic behaviour for the empirical probability of violation is due to the fact that for any number of facets s the number of polytopes among which the worst-case probability of violation is calculated is not the same.

3.3.2 Application to uncertain multi-agent programs

In this subsection we consider a multi-agent optimization set-up, where the choice of the cost function can be arbitrary and the uncertainty affects only the constraints. We, then, leverage the theoretical results of this chapter in order to provide *a posteriori* robustness certificates for the entire feasible region. As a result, we account for problems with multiple optimizers or for feasible deviations from a given nominal solution.

3.3.3 Problem formulation

Consider the following optimization program

$$\mathrm{P}_{\Delta} : \min_{x \in X} J(x) \text{ subject to } x \in \bigcap_{\delta \in \Delta} X_{\delta}, \tag{3.13}$$

where the cost function $J(x)$ can be chosen arbitrarily, x is a decision vector taking values in the deterministic set $X \subset \mathbb{R}^d$, while X_{δ} is dictated by the uncertain parameter δ. The

corresponding scenario program relies on available i.i.d. samples of the uncertainty δ and it is then defined as follows.

$$\mathrm{P}_K : \min_{x\in X} J(x) \text{ subject to } x \in \bigcap_{k=1}^{K} X_{\delta^{(k)}}. \tag{3.14}$$

A direct consequence of Theorem 3.2 is that we can provide distribution-free guarantees for any optimizer x_K^* of problem P_K. To solidify this statement we introduce the following corollary

Corollary 3.2 *Consider Assumption 3.2 and the setting of Theorem 3.2. We have that*

$$\mathbb{P}^K\Big\{\delta_K \in \Delta^K :\ V(x_K^*) > \varepsilon(F_K), \textit{ for any optimizer } x_K^* \textit{ of } \mathrm{P}_K\Big\} \leq \beta.$$

In the subsection that follows, we show how Corollary 3.2 can be applied to the electric vehicle charging control problem when we additionally have uncertainty in the constraints. Our mathematical framework allows to provide probabilistic feasibility guarantees collectively for the entire set of optimal charging schedules. Furthermore, it can be leveraged to provide certificates in a probabilistic fashion for cases where vehicle users' deviations from their predicted incentives leads to changes in the systems' operating point. A cooperative scheme is considered, where agents-vehicles minimize a common electricity cost, while their charging schedules are subject to uncertain constraints.

3.3.4 Electric vehicle charging control as a multi-agent program

In this subsection we extend the framework of [16], [3], [60], where purely deterministic constraints were assumed for the electric vehicle charging control problem, by accounting for uncertainty in the constraints. Notable works towards this direction are [128] and [129], where randomized algorithms are proposed in a more centralised control scheme. In our work, we consider the following multi-agent program

$$\begin{aligned} &\min_{x\in\mathbb{R}^{nN}} J(x) \text{ subject to} \\ &x \in \bigcap_{\delta\in\Delta} \prod_{i\in\mathcal{N}} \{x_i \in [\underline{x}_i(\delta), \overline{x}_i(\delta)] : \sum_{t=1}^{n} x_i^{(t)} \geq E_i(\delta)\}. \end{aligned} \tag{3.15}$$

The two main requirements for the operation of the system under study, namely, the lower and upper bounds imposed on the power rate of each vehicle $i \in \mathcal{N}$ that is $\underline{x}_i$ and $\overline{x}_i$, respectively, and the total energy level E_i to be achieved at the end of charging, can be modelled as constraints of affine form. The variables $x_i = (x_i^{(t)})_{t=1}^n$ denote the charging schedule for all time instances $t \in \{1,\dots,n\}$.

The corresponding scenario program of (3.15) is given by

$$\begin{aligned} &\min_{x\in\mathbb{R}^{nN}} J(x) \text{ subject to} \\ &x \in \bigcap_{k=1}^{K} \prod_{i\in\mathcal{N}} \{x_i \in [\underline{x}_i(\delta^{(k)}), \overline{x}_i(\delta^{(k)})] : \sum_{t=1}^{n} x_i^{(t)} \geq E_i(\delta^{(k)})\}. \end{aligned} \tag{3.16}$$

The cost function J is allowed to have any arbitrary form. The numerical studies of this chapter and the rest of the thesis are simulations based on synthetic data. Their purpose is to illustrate the developed theoretical results. The methodology presented in this work is more general with the potential of being used in more complex data-driven applications. In our set-up we assume that the upper and lower bounds of the charging rate, $(\overline{x}_i(\delta_u))_{i\in\mathcal{N}}$, $(\underline{x}_i(\delta_l))_{i\in\mathcal{N}} \in \mathbb{R}^{nN}$ are affected by the uncertain parameters $\delta_u, \delta_l \in \mathbb{R}^{nN}$, respectively, with uncertainty representing volatile grid power restrictions. Each of the parameters' elements is extracted according to the same probability distribution $\mathcal{N}(0, 0.5)$, where $\mathcal{N}(\mu,\sigma)$ is a Gaussian distribution with mean μ and standard deviation σ. The distribution is truncated by a prespecified quantity to avoid infeasibility issues. We further assume that the uncertainty is additive, i.e., $\overline{x}(\delta_u) = \overline{x}^{nom} + \delta_u$ and $\underline{x}(\delta_l) = \underline{x}^{nom} + \delta_l$, where each element of $\overline{x}^{nom}$ is drawn from a uniform probability distribution with support $[10, 20]$ kW and $\underline{x}^{nom}$ is set to 2 kW. Finally, the energy capacity of the battery can be affected by a variety of factors such as battery aging and lithium plating for Li-ion batteries. These phenomena can have an effect on the amount of energy required by each vehicle to fully charge, thus imposing uncertainty on the final energy level to be achieved by each of them by the end of the charging cycle. In our set-up, the uncertainty in the total energy $E = (E_i)_{i\in\mathcal{N}}$ (in kWh) of each vehicle at the end of charging is yet again assumed to be additive, i.e., $E = E^{nom} + \delta_e$, where the elements of $\delta_e \in \mathbb{R}^N$ are extracted according to the probability $\mathcal{N}(0,1)$ and $E_i^{nom} \in \mathbb{R}$ is the nominal final energy demand of each agent $i \in \mathcal{N}$. The uncertainty vector is given by $\delta = [\delta_u, \delta_l, \delta_e] \in \mathbb{R}^{N(2n+1)}$.

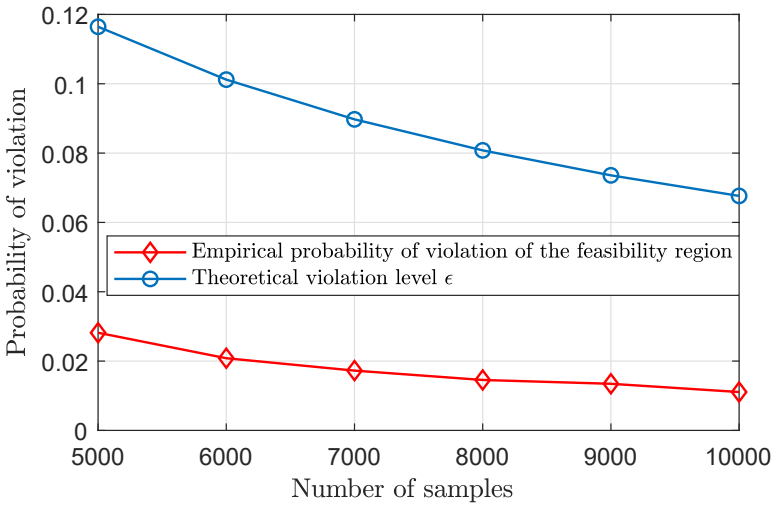

Figure 3.4: Empirical probability of violation of the feasibility region (red diamond line) versus the the theoretical violation level for a different number of samples $M = \{5000, 6000, \ldots, 10000\}$. To calculate the probability of violation a total of $M_{test} = 40000$ samples is used.

Considering $N = 5$ vehicles and $n = 12$ time slots we construct the feasibility region of the corresponding scenario program (3.16). Using $M_{test} = 40000$ test samples we empirically compute the probability that a new yet unseen constraint will be violated by at least one element of the feasible region and compare it with the theoretical violation level $\varepsilon(k)$ from Theorem 3.2. The results are shown in Figure 3.4, where the red diamond line corresponds to the empirical probability and the blue circle line corresponds to the theoretical bound. Note that an upper bound for the theoretical violation level $\varepsilon(k)$ can be obtained by counting the number of facets $k = F_K$ of the feasible set or by leveraging the geometry of our numerical example to provide an upper bound for F_K, that is $F_K \leq 2nN + N$. This bound can be derived noticing that the feasible region is in fact a cartesian product of N rectangles intersected by a halfspace that corresponds to the energy constraint. As such the worst case number of facets for the entire polytope in our example is $N(2n+1)$.

In general, for samples that give rise to more than one affine constraints, the number of facets constitutes only an upper bound for the cardinality of the minimal compression set. This bound is tight only in case there is a one to one correspondence between a sample realization and a scalar-valued constraint. This implies that the guarantees for the feasible subset of the problem under study can be significantly tighter. The reason

behind the use of the looser bound $k = 2nN + N$ in our example lies in the fact that its quantification is straightforward and the use of a compression set function is thus not required. In other cases, however, where an upper bound for k is unavailable, the methodologies for the detection of redundant affine constraints provided by [127] and references therein can be used as minimal compression set functions.

3.4 Application to uncertain non-cooperative games

The concept of Nash equilibrium (NE) typically assumes complete knowledge of the agents' cost functions and constraints. For games affected by uncertainty in the cost or constraints, the notion of NE is insufficient, as the resulting agents' strategies do not necessarily exhibit any robustness properties against yet unseen uncertainty realizations. One of the first attempts to deal with uncertainty in multi-agent games was noticed in [130]. Motivated by these developments, two main research directions are encountered: 1) NE analysis based on particular models for the probability distribution of the uncertainty [131], [132], [133] and/or the geometry of its support [134], [135] and, 2) Distribution-free NE analysis, where no assumption on the probability distribution of the uncertainty is imposed.

Only a few data-driven works for distribution-free NE seeking have appeared in the literature with [136] and [49], [48] being the most closely related to our work. These papers attempt to bridge uncertain multi-agent games with the scenario approach, focusing on the provision of guarantees for a unique NE. The latter is quite restrictive, as many games of practical interest have multiple (possibly infinite) NE. Motivated by the lack of distribution-free results for solution sets, we leverage the theoretical results developed in this chapter for the provision of probabilistic feasibility guarantees collectively for all NE of a non-cooperative game with uncertainty in the constraints.

3.4.1 Problem formulation

Under Assumption 3.2, we consider a multi-agent game, whose constraints are affected by uncertainty. Each agent $i \in \mathcal{N}$ seeks to minimize her own cost function, given the

strategies x_{-i} of all other players, by solving the following program

$$\min_{x_i \in X_i} J_i(x_i, x_{-i}) \text{ subject to } x_i \in \bigcap_{\delta \in \Delta} X_i^{\delta}. \tag{3.17}$$

For the aforementioned problem, we introduce the solution concept of *robust Nash equilibrium (NE)* as presented in Definition 3.3.

Definition 3.3 *[9] A vector $x^* = (x_i^*)_{i \in \mathcal{N}}$ is a robust NE of the associated game if and only if $J_i(x_i^*, x_{-i}^*) \leq J_i(x_i, x_{-i}^*)$ for any $x_i \in \bigcap_{\delta \in \Delta} X_i^{\delta}$ and for any $i \in \mathcal{N}$.*

A decision is a robust NE if no agent has an incentive to unilaterally deviate from their decision, given other agents' decisions, for all possible uncertainty realizations $\delta \in \Delta$. Due to the presence of uncertainty and the (possibly) infinite cardinality of Δ, problem (3.17) is very difficult to solve without imposing any assumption on the geometry of the sample set Δ or the underlying probability distribution $\mathbb{P}$. To circumvent these issues, we approximate problem (3.17) by drawing multiple i.i.d. samples $\delta_K = (\delta^{(1)}, \delta^{(2)}, \ldots, \delta^{(K)}) \in \Delta^K$ and then considering the following scenario-based NE seeking problem, where each agent $i \in \mathcal{N}$ solves the following optimization program

$$\min_{x_i \in X_i} J_i(x_i, x_{-i}) \text{ subject to } x_i \in \bigcap_{k=1,\ldots,K} X_i^{\delta^{(k)}}. \tag{3.18}$$

Given this data-driven setting, the following definition of the so called *randomized* NE is proposed:

Definition 3.4 *For a drawn multi-sample $\delta_K \in \Delta^K$, a vector $x_K^* = (x_{i,K}^*)_{i \in \mathcal{N}}$ is a randomized NE of the associated game if $J_i(x_{i,K}^*, x_{-i,K}^*) \leq J_i(x_i, x_{-i,K}^*)$ for any $x_i \in \bigcap_{k=1,\ldots,K} X_i^{\delta^{(k)}}$ and for any $i \in \mathcal{N}$.*

Our aim is to provide probabilistic feasibility guarantees for the entire set of NE of (3.18) returned by an arbitrary NE seeking algorithm, i.e., to quantify the probability that any randomized NE decision $x_{i,K}^*$ of (3.18) belongs to the constraint set X_i^{δ}, $\forall\, i \in \mathcal{N}$ for a new unseen sample $\delta \in \Delta$. To this end, our analysis is primarily focused on feasibility problems affected by uncertainty. A direct consequence of Theorem 3.2 is that we can provide distribution-free guarantees for any NE of problem (3.18). To solidify this statement we introduce the following corollary

Corollary 3.3 *Consider Assumption 3.2 and the setting of Theorem 3.2. We have that*

$$\mathbb{P}^K\Big\{\delta_K \in \Delta^K :\ V(x_K^*) > \varepsilon(F_K),\ \textit{for any NE } x_K^* \textit{ of (3.18)}\Big\} \leq \beta.$$

Note that any decentralized/distributed algorithm can be used to calculate the randomized NE. In the next subsection, we revisit the electric vehicle charging control problem but instead of assuming that users working together to optimize a common total electricity cost as in Section 3.3.4, users act as selfish agents trying to minimize their own individual cost. This gives rise to an uncertain NE problem. We then attempt to provide robustness certificates collectively for all NE of the uncertain electric vehicle charging game.

3.4.2 Electric vehicle charging games

We model the electric vehicle (EV) charging control problem as a game that comprises self-interested agents-vehicles each of them aiming at minimizing their own electricity cost, while their charging schedules are subject to certain specifications. As such, each agent $i \in \mathcal{N}$ solves the following problem

$$\begin{aligned}&\min_{x_i \in \mathbb{R}^n} J_i(x_i, x_{-i}) \text{ subject to}\\ &x_i \in \bigcap_{\delta \in \Delta} \{[\underline{x}_i(\delta), \overline{x}_i(\delta)] \bigcap \{x_i \in \mathbb{R}^n : \sum_{t=1}^{n} x_i^{(t)} \geq E_i(\delta)\}\}. \end{aligned} \tag{3.19}$$

The variables $x_i = (x_i^{(t)})_{t=1}^n$ and J_i denote, respectively, the charging schedule for all time instances $t \in \{1, \ldots, n\}$ and the electricity cost to be minimized for each vehicle $i \in \mathcal{N}$. As in Section 3.3.4, the uncertain constraint for each vehicle i comprises of uncertain lower and upper bounds $\underline{x}_i(\delta)$ and $\overline{x}_i(\delta)$ and total energy levels $E_i(\delta)$. Due to the presence of a variety of unpredictable internal and external influences contributing to the uncertainty of the system, it is very difficult to address the problem using traditional probabilistic approaches. We thus again adopt a data-based approach, a more viable alternative.

Our aim is to provide guarantees on the probability that a NE satisfies the constraints of (3.19). We assume that the upper constraint $(\overline{x}_i(\delta_u))_{i \in \mathcal{N}} \in \mathbb{R}^{nN}$ is affected by an additive uncertainty in the form of $\delta_u \in \mathbb{R}^{nN}$, whose elements are extracted according to the random parameter $l = l_1 \cdot l_2$, where l_1 follows $U(0, 0.3)$ and l_2 follows $\mathcal{N}(1, 3)$, with

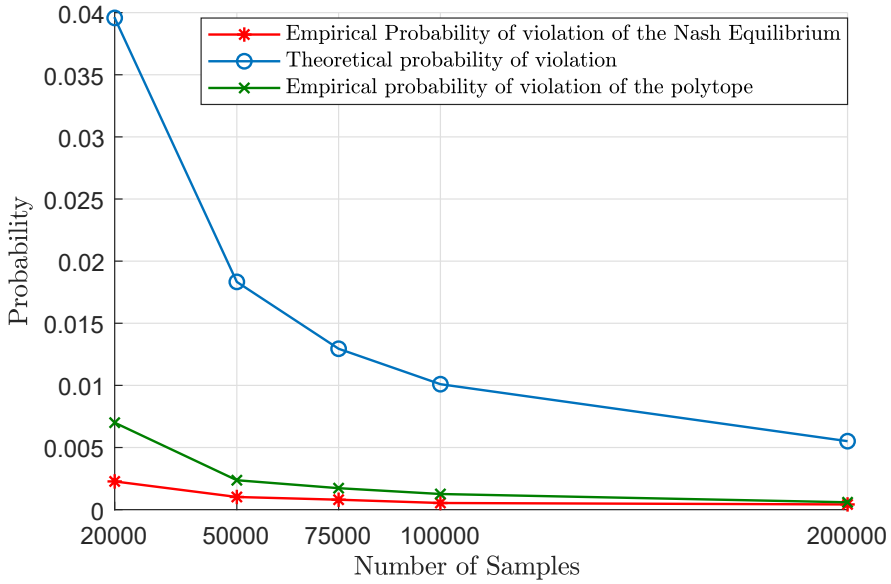

Figure 3.5: The empirical probability of violation of the entire polytope $\hat{\mathbb{V}}(\Pi_K)$ (green) and the empirical probability of violation of the computed NE returned by the algorithm of [49],[48], namely, $\hat{V}(x_K^*)$ (red) versus the theoretical violation level of Theorem 3.2 (blue) with respect to five different values of the number of samples $K = 20000, 50000, 75000, 100000, 200000$. The empirical probability of violation for both the polytope and the NE is computed using 2000000 test samples. Note that the blue line corresponds to the theoretical counterpart of the green one.

$U(0,1)$ denoting a uniform distribution with support $[0,1]$, while $\mathcal{N}(1,3)$ a gaussian distribution with mean 1 and standard deviation 3. As such, $\bar{x}(\delta_u) = \bar{x}^{nom} + \delta_u$, where $\bar{x}^{nom}$ is a given deterministic component. The total energy $E = (E_i)_{i\in\mathcal{N}}$ of each agent at the end of charging is also affected by uncertainty i.e., $E^i = (1-\delta_e)E^i_{nom}$, where $\delta_e \in \mathbb{R}^N$ and its elements are extracted according to the random parameter $0.05l_3$, with l_3 following $\mathcal{N}(0,1)$, and $E^i_{nom} \in \mathbb{R}$ is the nominal final energy demand of each agent $i \in \mathcal{N}$ drawn from $U(10,17)$. The uncertainty vector is given by $\delta = [\delta_u, \delta_e] \in \mathbb{R}^{N(n+1)}$. The lower bound is assumed to be deterministic and, particularly, $\underline{x}_i = 0$ for any $i \in \mathcal{N}$. Finally, the cost function of each vehicle $i \in \mathcal{N}$ is given by $J_i(x_i, x_{-i}) = (x_i)^T(A_0\sigma(x_i, x_{-i}) + b_0)$, where the matrix $A_0 \in \mathbb{R}^{n\times n}$ is diagonal and $\sigma(x_i, x_{-i}) = \sum_{i=1}^N x_i$. Following the work of [49], [48] the entries $\{a_t\}_{t=1}^n$ of $A_0 = \text{diag}(\{a_t\}_{t=1}^n)$ are evaluated by rescaling a winter weekday demand profile in the UK [137]. The vector $b_0 \in \mathbb{R}^n$ is set to zero.

To compute the NE we employed the algorithm of [48] (see Algorithm 1 therein). Setting the number of agents to $N = 10$ and the number of timesteps to $n = 12$, we

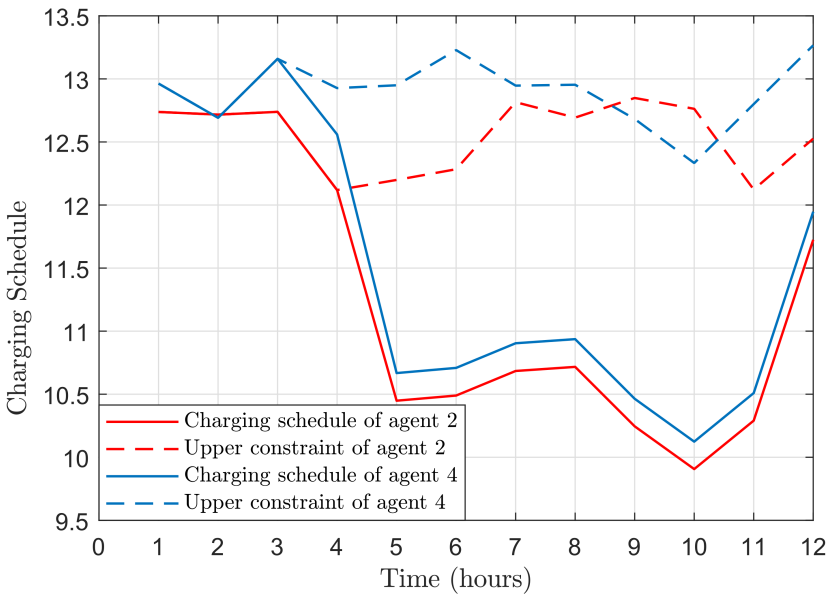

Figure 3.6: Charging schedules of agents 2 (red solid line) and 4 (blue solid line) and their respective uncertain upper bounds in dashed lines. The agents' schedules and their upper constraints coincide during the beginning and the end of the horizon, which implies that they are charging at their boundaries in their attempt to benefit from the price drop during these time instances.

run the algorithm for multi-samples of different size, namely, 20000, 50000, 75000, 100000 and 200000 and we compare the behaviour of two different probabilities with respect to the cardinality of the multisample, as illustrated in Figure 3.5. With green we illustrate the worst case empirical probability of violation for the polytope $\hat{\mathbb{V}}(\Pi_K)$ (calculated by taking the maximum violation among all grid points on the polytope), while with red we show the empirical probability $\hat{V}(x_K^*)$ of violation for the x_K^* returned by the algorithm of [49],[48]. Note that the blue line is the theoretical counterpart of the green line and corresponds to the theoretical violation level $\varepsilon(s)$, as defined in Theorem 3.2. The empirical calculation was performed using 2000000 test samples, different from those used in the NE seeking process. As expected, both empirical values are less than the theoretical bound derived by Theorem 3.2. It was also anticipated that for any multisample we would have that $\hat{\mathbb{V}}(\Pi_K) \geq \hat{V}(x_K^*)$ as the former corresponds to the collective violation of all feasible points, including x_K^*. Finally, Figure 3.6 illustrates the charging schedules of two of the agents and their respective uncertain upper bounds. Due to the high required total energy at the end of charging, we observe that the agents'

schedules and their upper constraints coincide during the beginning and the end of the horizon, which implies that they are charging at their boundaries in their attempt to exploit the low price during these time instances and minimize their charging cost.

3.5 Concluding remarks

In this chapter we addressed the challenge of providing probabilistic feasibility guarantees collectively for a set of feasible solutions. This is of particular interest in case the uncertain system at hand, whether modelled as a multi-agent program or a game admits multiple solutions, hence it is important to provide guarantees collectively for all of them. Our theory can also be applied to cases where agents decisions deviate from the solution due to changes in the agents incentives or due to the numerical implementation of the solution seeking algorithm. We first considered feasibility programs with an arbitrary cost function and uncertain convex constraints and provided *a posteriori* bounds for the probability of violation of all feasible solutions. Our results are then applied to multi-agent optimization programs and NE problems using the electric vehicle charging problem as a motivating example.

In the next chapter we show, as stated in Remark 3.2, that focusing on a certain subset of the feasible region can substantially improve the probabilistic feasibility guarantees in an *a posteriori fashion*. Furthermore, we propose a data-driven equilibrium seeking algorithm that converges to the social welfare of the problem, while simultaneously allowing *a priori* control over the probabilistic feasibility of a region around that point.

4

Control of probabilistic feasibility of sets around equilibria

In this chapter we focus on aggregative games with uncertain constraints coupling the agents' decisions. We consider a setting, where bounded deviations of agents' decisions from the equilibrium are allowed. Following a randomized approach, where uncertain constraints are only known from data, we propose an algorithm that converges to an equilibrium such that we have control, in an *a priori* fashion, over the level of probabilistic feasibility collectively for the entire set of considered deviations. We show that tuning of the provided guarantees can be accomplished by controlling the number of coupling constraints allowed to intersect the region of deviations. Since the choice of constraints that intersect can affect the efficiency of the system in terms of cost, our algorithmic scheme is based on a novel methodology that identifies online the coupling constraints that play the most significant role to the system's performance. Numerical studies corroborate our theoretical results.

4.1 Introduction

In this chapter we consider a setting where agents act as self-interested entities that interact with each other and make decisions based on their own criteria, while being subject to uncertain global or local constraints. Despite recent advancements in the

scenario approach to deal with uncertainty in games [48], [49], [136] robust game theory still needs to overcome fundamental assumptions related to how agents make their decisions. Originally established for applications in economics, standard game-theoretic models impose on the agents a strict behavioural assumption, the so called rationality assumption. Specifically, it is assumed that agents wish to maximize their profit always acting as selfish entities. However, studies have shown that this is unrealistic in practice [33], [32], [34], [35], [36] and that agents usually exhibit a bounded rationality behaviour [37], since their decision making procedure might vary depending on predispositions, lack of computational power and a restricted amount of time available for decision making. These factors can alter the incentives of the agents thus leading to decisions that are satisfactory, i.e., close to an equilibrium. Apart from economics and the social sciences, deviations from the nominal equilibrium are common in engineering applications, as well. Consider, again as an example, the electric vehicle charging control problem, where the owner of each vehicle is modelled as a selfish agent trying to optimize their own cost subject to constraints of the vehicle itself and the electricity grid. In practice, each users' cost function includes unknown charging preferences. The grid operator taking into account predictions of the users' behaviour, wishes to plan ahead by incentivizing the vehicles to charge during certain time intervals of the next day so as to minimize the overall cost of the fleet. However, due to variations from the predicted behaviour, deviations from the predicted equilibrium solution can occur.

Aiming at building a theoretical framework to address this more general setting, this chapter considers games with uncertain constraints, where agents are allowed to deviate from the equilibrium. We follow a randomized approach approximating the coupling constraints through data. In this more general setting, where deviations are considered, providing guarantees for a single solution is devoid of any meaning, as even for the same constraint data, repetition of the game might lead to a different solution in a neighbourhood around the equilibrium. This renders the computation of the support samples a very challenging, if not impossible, task. As such, a methodology for the provision of guarantees collectively for the entire set of possible deviations around the equilibrium is required. Focusing on the class of *aggregative* games [4], [69], [3] (see

Subsection 2.3.3), where the cost function of each agent depends on their own decision and the average aggregate decision of the population, our contributions are as follows.

1. In the spirit of Remark 3.2 of Chapter 3, we adopt a scenario-theoretic paradigm establishing a methodology for the provision of *a posteriori* probabilistic feasibility guarantees for the feasibility region around a randomized equilibrium of the game under study.

2. We design a data-driven equilibrium-seeking algorithm that converges to a solution such that a predefined level of probabilistic feasibility is achieved for all admissible deviations around it. The proposed algorithm allows control in an *a priori* fashion over the probabilistic feasibility of the set under study by selecting how many coupling constraint realizations intersect the set of deviations. Since the choice of constraints that intersect this set can affect the efficiency of the system in terms of cost, our algorithmic scheme is based on a novel methodology that identifies online the coupling constraints that play the most significant role to the overall systems' performance. Depending on how many facets we allow the set to intersect, this can lead to a trade-off between system performance and feasibility.

This chapter is organized as follows. In Sections 4.1 and 4.2 we provide background from game theory and the scenario approach that is key for the subsequent developments. In Section 4.3 we propose a new data-driven algorithm and prove its convergence. Furthermore, we accompany the region under study with tunable *a priori* probabilistic feasibility guarantees. In Section 4.4 we provide a detailed numerical study that corroborates our results. Finally, Section 4.5 concludes the chapter and presents future research directions. For the ease of the reader we repeat some notational and conceptual aspects introduced in Chapter 2. Some concepts are necessarily re-adapted to fit the more general framework of this chapter.

4.1.1 Aggregative games with uncertain constraints

We consider a population of agents with index set $\mathcal{N} = \{1, \ldots, N\}$. The decision vector x_i of each agent $i \in \mathcal{N}$ takes values in the set $X_i \subseteq \mathbb{R}^n$, while $x = (x_i)_{i=1}^N \in X = \prod_{i=1}^N X_i \subseteq \mathbb{R}^{nN}$ is the global decision vector that is formed by concatenating the decisions of the entire population. The vector $x_{-i} \in \mathbb{R}^{n(N-1)}$ consists of all agents' decision vectors except for that of agent i. In our set-up, we assume that the cost function of each agent $i \in \mathcal{N}$ is a real-valued function $J_i(x_i, \sigma(x))$ that depends on the decision vector of agent $i \in \mathcal{N}$ and the average aggregate decision $\sigma : \mathbb{R}^{nN} \to \mathbb{R}^n$ defined as the mapping $x \mapsto \frac{1}{N}\sum_{i=1}^N x_i$. Furthermore, we consider uncertain constraints coupling the agents' decisions[1]. Each of the coupling constraints has the form $C_\delta = \{x \in X : g(x, \delta) \leq 0\}$, where the function $g : \mathbb{R}^{nN} \times \Delta \to \mathbb{R}^2$ depends on the uncertain parameter δ that takes values in a support set Δ according to a probability distribution $\mathbb{P}$. Both Δ and $\mathbb{P}$ could possibly be unknown. In this set-up, given the decisions of all other agents x_{-i}, each agent $i \in \mathcal{N}$ solves the following optimization program

$$\begin{aligned} \bar{x}_i \in \underset{x_i \in X_i}{\arg\min}\, & J_i(x_i, \sigma(x)) \\ \text{subject to } & x_i \in \bigcap_{\delta \in \Delta} C_\delta^i(x_{-i}), \end{aligned} \tag{4.1}$$

where $C_\delta^i(x_{-i}) = \{x_i \in X_i : g(x_i, x_{-i}, \delta) \leq 0\}$ denotes the coupling constraint for an uncertain realization $\delta \in \Delta$ and fixed x_{-i}. The collection of the coupled optimization programs in (4.1) comprises, as seen in Section 3.4.1, an *uncertain game*. For convenience, we denote this game as G. Since it is impossible to solve G without any additional assumptions on the support set Δ and/or the probability distribution $\mathbb{P}$, we approximate the set of its constraints, taken over the entire support Δ of δ, with a subset obtained by drawing a finite number of independent and identically distributed (i.i.d.) samples $\delta_K = (\delta^{(1)}, \ldots, \delta^{(K)})$ from Δ, similar to the approach followed in Chapter 3..

[1]Note that assuming deterministic local constraints is without loss of generality, as uncertain local constraints can be considered as a special case of coupling constraints, that depend only on the decision vector of a single agent.

[2]Note that considering the scalar function g instead of a vector-valued function $\bar{g} : \mathbb{R}^{nN} \times \Delta \to \mathbb{R}^s$ is without loss of generality as one could always define $g(x, \delta)$ as the largest component of $\bar{g}$, i.e., $g(x, \delta) = \max_{1,\ldots,s} \bar{g}_i(x, \delta)$. Furthermore, our setting can be easily adapted for the provision of probabilistic feasibility guarantees separately for each constraint of the form $\bar{C}_\delta = \{x \in X : \bar{g}(x, \delta) \leq 0\}$, instead of the satisfaction of a joint chance-constraint.

As such, for a given collection of samples $\delta_K \in \Delta^K$ and given the decisions of all other agents x_{-i}, each agent $i \in \mathcal{N}$ solves the following scenario optimization program

$$\begin{aligned} \bar{x}_{i,K} \in & \arg\min_{x_i \in X_i} J_i(x_i, \sigma(x)) \\ & \text{subject to } \; x_i \in \bigcap_{k=1}^{K} C^i_{\delta^{(k)}}(x_{-i}). \end{aligned} \tag{4.2}$$

We refer to the collection of coupled optimization programs in (4.2) as a *scenario game* denoted as G_K, where the subscript K implies dependence on the drawn multi-sample δ_K. Under certain assumptions, a solution to the scenario game G_K exists and the problem is tractable using available solution-seeking algorithms. In this chapter, our interest lies in two solution concepts, namely the *Wardrop* and the *Nash equilibrium*. These have been defined in Chapter 2 for the deterministic case and are extended in the definition that follows to their scenario counterparts.

Definition 4.1 *For a fixed multi-sample $\delta_K \in \Delta^K$, consider the game G_K. A point $\bar{x}_K$ is called*

1. *A Nash equilibrium (NE) of G_K if*

$$\begin{aligned} & J_i(\bar{x}_{i,K}, \sigma(\bar{x}_K)) \leq J_i(x_i, \frac{1}{N} x_i + \frac{1}{N} \sum_{j \in \mathcal{N} \setminus \{i\}} \bar{x}_{j,K}) \\ & \textit{for all } x_i \in X_i \cap \bigcap_{k=1}^{K} C^i_{\delta^{(k)}}(\bar{x}_{-i,K}) \textit{ for all } i \in \mathcal{N}. \end{aligned}$$

2. *A Wardrop equilibrium (WE) of G_K if*

$$\begin{aligned} & J_i(\bar{x}_{i,K}, \sigma(\bar{x}_K)) \leq J_i(x_i, \sigma(\bar{x}_K)) \\ & \textit{for all } x_i \in X_i \cap \bigcap_{k=1}^{K} C^i_{\delta^{(k)}}(\bar{x}_{-i,K}) \textit{ for all } i \in \mathcal{N}. \end{aligned}$$

In other words, at the equilibrium each agent cannot improve her cost function by unilaterally changing her decision in the space constructed by random constraints that are obtained through sampling. We remind the reader that in the case of the WE each agents' decision can be considered negligible with respect to the average aggregate decision.

From Section 2.4 for a given multi-sample δ_K, equilibria of the scenario game G_K can be obtained as solutions to a *variational inequality* (VI) problem defined as follows

$$\begin{aligned} &\text{VI}_K : \text{Find } \bar{x}_K \in X_K \text{ such that} \\ &\qquad F^T(\bar{x}_K)(x - \bar{x}_K) \geq 0 \text{ for any } x \in X_K, \end{aligned}$$

where $X_K = \bigcap_{k=1}^{K} X_{\delta^{(k)}}$ and $X_{\delta^{(k)}} = X \cap C_{\delta^{(k)}}$. In case we are interested in finding the NE of G_K, the mapping F takes the form $F(x) = F_{NE}(x) = [\nabla_{x_i} J_i(x_i, \sigma(x))]_{i \in \mathcal{N}}$. Similarly, to obtain the WE of G_K the mapping of the VI takes the form $F(x) = F_{WE}(x) = [\nabla_{x_i} J_i(x_i, z)_{|z=\sigma(x)}]_{i \in \mathcal{N}}$. Notice that in the latter case, the second argument of the gradient is fixed and once the gradient is computed, it is set to $\sigma(x)$. To facilitate the connection between the equilibria of G_K and the solutions of VI_K, we impose the following assumptions. Note that Assumption 4.1 is similar to Assumption 2.1 but adapted for the case of aggregative games, while Assumption 4.2 is adapted from Assumptions 3.1 and 3.2.

Assumption 4.1 *For all $i \in \mathcal{N}$, the following hold:*

1. *$J_i(x_i, \sigma(x))$ is convex in x_i for any fixed $\{x_j \in X_j\}_{j \neq i}$.*
2. *$J_i(x_i, z)$ is convex in x_i for any fixed $z \in \frac{1}{N}\sum_{j=1}^{N} X_j$.*
3. *$J_i(z_1, z_2)$ is continuously differentiable with respect to both arguments over $X_i \times \frac{1}{N}\sum_{j=1}^{N} X_j$.*

Assumption 4.2 1. *For any multi-sample $\delta_K \in \Delta^K$ the feasible region $\Pi_K = X \cap \bigcap_{k=1}^{K} C_{\delta^{(k)}}$ is non-empty.*
2. *The set $X = \prod_{i=1}^{N} X_i$ is compact, polytopic and convex.*
3. *For any $\delta \in \Delta$, g is an affine function of the form $g(x, \delta) = a(\delta)^T x - b(\delta)$, where $a : \Delta \to \mathbb{R}^{nN}$ and $b : \Delta \to \mathbb{R}$.*

Assumption 4.3 *The mapping F is*

1. *a-strongly monotone, i.e.,*

$$(x-y)^T(F(x)-F(y)) \geq a||x-y||^2 \text{ for any } x,y \in X. \tag{4.3}$$

2. *L_F-Lipschitz continuous., i.e.,*

$$||F(x)-F(y)|| \leq L_F||x-y|| \text{ for any } x,y \in X. \tag{4.4}$$

Assumptions 4.1 and 4.3 are quite standard in the game-theoretic literature [63], [3]. Assumption 4.2 is relatively mild; The main advantage it imposes is the fact that the constraints are assumed to be affine, a property leveraged by the proposed data-driven algorithm (see Section 4.3), in order to converge to the desired equilibrium. The following lemma is an adaptation of Theorem 2.3.3 in [63], where convex constraints were assumed, for the special case of affine constraints as the ones used in our set-up.

Lemma 4.1 *Under Assumptions 4.1, 4.2 and 4.3, VI_K has a unique solution that is also an equilibrium of G_K.*

Note that this equilibrium could be a NE or a WE depending on the choice of F. However, the imposed assumptions capture both cases. This unique solution is called the *variational equilibrium* of G_K. For the considered class of VIs several algorithms can be employed from the literature to obtain the unique variational equilibrium of G_K. Probabilistic feasibility guarantees for the unique solution of the VI can then be provided both in an *a priori* and *a posteriori* fashion by resorting to the results in [49], [48], [136].

4.2 Tuning the probabilistic feasibility of a set around an equilibrium

4.2.1 A first a posteriori result

Following a scenario approach paradigm, we draw samples of the uncertainty parameter that affects the coupling constraints that in turn gives rise to a randomized approximation of the feasibility region $\Pi_K = X \cap \bigcap_{k=1}^{K} C_{\delta^{(k)}}$. We consider a more general set-up, where agents are allowed to deviate from a randomized equilibrium $\bar{x}_K$ of G_K due to changes

in their cost functions. These deviations are thus feasible with respect to both the local constraints and the drawn coupling constraints. The region in which agents are allowed to deviate is assumed to lie within a predefined ball $\mathbb{B}_\infty(\bar{x}_K, \rho)$, where $\bar{x}_K$ is assumed to be unique as per Lemma 4.1 and $\rho > 0$ is a fixed radius that denotes the maximum possible distance of agents' deviations from $\bar{x}_K$. As such, the region of interest is $S_K = \Pi_K \cap \mathbb{B}_\infty(\bar{x}_K, \rho)$. It is important to stress that the choice of the equilibrium seeking algorithm is not instrumental for the subsequent results. Furthermore, note that for representing the set of deviations the infinity norm is considered. Using any other norm in general requires adaptation of the proposed algorithm. However, for cases like the one in our numerical study, the proposed algorithm holds for any p-norm.

For a given multi-sample δ_K, consider an algorithm A that takes as input δ_K and returns the region $S_K = \Pi_K \cap \mathbb{B}_\infty(\bar{x}_K, \rho)$. For pictorial purposes, Figure 4.1 illustrates the construction of such a set in a two-dimensional game with scalar decisions. The equilibrium point $\bar{x}_K$ denotes the equilibrium and Π_K is formed exclusively by coupling constraints. Note that only two constraints (in colour blue) are of support for the set as removing any of the red constraints will not change the set. However, depending on the multi-sample δ_K and the location of the equilibrium, more constraints could intersect region S_K, thus increasing the number of support samples. One of the main results of this chapter (Theorem 4.3) holds under nondegeneracy, an assumption introduced in Chapter 2, repeated here for the ease of the reader.

Assumption 4.4 *For $A(\delta_K) = S_K$, if I is the set of support samples of S_K, we have that $A(\delta_K) = A(I)$.*

We can quantify the number of samples that support S_K in an *a posteriori* fashion as established in the following theorem.

Theorem 4.1 *Consider algorithm A outlined above and Assumptions 4.1, 4.2, 4.3. Then, for a fixed confidence parameter $\beta \in (0,1)$ we have that*

$$\mathbb{P}^K\Big\{\delta_K \in \Delta^K : \mathbb{V}(S_K) > \varepsilon(s^* + M)\Big\} \leq \beta,$$

$$\text{where } \varepsilon(K) = 1 \text{ and } \sum_{i=0}^{K-1} \binom{K}{i} (1-\varepsilon(i))^{K-i} = \beta,$$

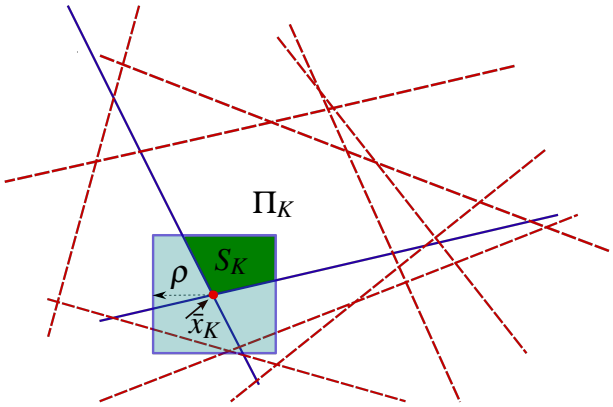

Figure 4.1: Illustration of region S_K (in green) as the intersection of the set of deviations $\mathbb{B}_\infty(\bar{x}_K, \rho)$ around the unique equilibrium $\bar{x}_K$ illustrated as a red dot.

and s^* is the number of samples that support the equilibrium $\bar{x}_K$ and M the number of facets of Π_K that intersect S_K.

Proof: The samples whose removal can lead to a change of region S_K either give rise to facets of Π_K that intersect S_K, whose number is M, or to constraints that support the equilibrium $\bar{x}_K$ (whose number is s^*). To see this note that if we remove one of these facets, then, since Assumption 4.4 holds, the region S_K is enlarged. In addition, if any of the s^* samples that support the equilibrium are removed, then the centre of the region changes position. As such, the number of samples that form a compression set for δ_K is bounded by $s^* + M$. Note that this may just be an upper bound as some sample might correspond to a constraint that intersects the region surrounding the equilibrium $\bar{x}_K$ and, at the same time, support that equilibrium. Existence of a compression set with an *a posteriori* bound on its cardinality, is sufficient for the application of Theorem 2 in [138], which then leads to the statement of the theorem. ■

We note again that one can choose ε by splitting β evenly among the terms in the sum of (3.8) as in Section 3.2 on page 40. To use the result of Theorem 4.1, one needs to count the number of samples s^* that support the randomized equilibrium $\bar{x}_K$ and the number of coupling constraints M that correspond to facets of S_K. However, this can only be done *a posteriori*, i.e., when δ_K is provided. An important question naturally arises: When we are studying a set of deviations of agents decisions around the equilibrium (instead of solely

the equilibrium point itself), is it possible to provide tunable and *a priori* guarantees for the entire set of deviations? We aim at answering this question in the subsection that follows.

4.2.2 A priori probabilistic feasibility

In this subsection we show that tunable *a priori* probabilistic feasibility guarantees can be obtained by imposing the desired guarantees as a chance-constraint of the considered scenario game, thus leading to a modified version of the problem, where the scenario game G_K is subject to

$$\mathbb{P}\{\delta \in \Delta : S_K \not\subseteq C_\delta\} \leq \varepsilon, \tag{4.5}$$

Function $\varepsilon \in [0,1]$ is the violation level which is now *a priori* fixed. Note that the modified game is very challenging, if not impossible, to solve. However, as we will see in the subsequent developments, (4.5) can be satisfied with high confidence by tuning the number M of facets of the feasibility region Π_K that the set of deviations is allowed to intersect. As such, we can instead solve G_K, while ensuring that at most M facets intersect the set of deviations. The resulting solution, denoted as x_K^* to distinguish it from the *a posteriori* solution $\bar{x}_K$ of the previous subsection, will be such that the entire set of deviations $S_K^* = \Pi_K \cap \mathbb{B}_\infty(x_K^*, \rho)$ satisfies (4.5) with high confidence. Tuning the number of facets M that can intersect the considered region of deviations can be achieved by designing a novel data-driven algorithmic scheme. In the following subsection we present the desired behaviour of such an algorithm by means of a numerical example.

4.2.3 Illustrative example and desired behaviour of the algorithm

Consider an aggregative game where agents' decisions are subject to deterministic local constraints and uncertain coupling constraints on the average aggregate decision. In this setting, each agent $i \in \mathcal{N}$ solves the following optimization problem

$$\begin{aligned} &\min_{x_i \in X_i} \; p(\sigma(x))^T x_i \\ &\qquad \text{subject to } \sigma(x) \leq b(\delta), \text{ for all } \delta \in \Delta, \end{aligned}$$

where function $p : \mathbb{R}^n \to \mathbb{R}^n$ is an affine function of the aggregate decision given by the relation $p(\sigma(x)) = C\sigma(x) + c$, where $C \in \mathbb{R}^{n\times n}$ and $c \in \mathbb{R}^n$. The set of local constraints of each agent $i \in \mathcal{N}$ is considered to be a hyperrectangle, i.e., $X_i = \{x_i \in \mathbb{R}^n : x_i \in [\underline{x}_i, \overline{x}_i]\}$, where $\underline{x}_i = (\underline{x}_{i,t})_{t=1}^n, \overline{x}_i = (\overline{x}_{i,t})_{t=1}^n \in \mathbb{R}^n$ are the lower and upper bounds for each $t \in \{1,\dots,n\}$. We denote the collection of all N coupled optimization programs as the game G'. Since we do not know the sample set Δ and/or the probability distribution $\mathbb{P}$, we draw a finite number of K samples, thus obtaining the following scenario approximation, where each agent $i \in \mathcal{N}$ solves

$$\begin{aligned} &\min_{x_i \in X_i} \; p(\sigma(x))^T x_i \\ &\quad \text{subject to } \sigma(x) \leq b(\delta^{(k)}), k = 1,\dots,K. \end{aligned}$$

We call the collection of all N coupled optimization programs G'_K. Furthermore, we assume that there exists a positive constant μ such that $C \succ \mu I_n$. Then, function $F(x) = [p(\sigma(x))]_{i\in\mathcal{N}}$ is strongly monotone when considering σ as the decision variable. Since F is additionally Lipschitz continuous, hence continuous, game G_K admits a unique aggregate equilibrium $\sigma^*_K = \sigma(x^*_K)$.

Note that the Wardrop equilibrium σ^*_K can be obtained by solving the following scenario optimization program

$$\begin{aligned} P_K : &\min_{x\in X} E(x) \\ &\quad \text{subject to } \sigma(x) \leq b(\delta^{(k)}), k = 1,\dots,K, \end{aligned}$$

where $E(x) = p(\sigma(x))^T\sigma(x)$. This game is in fact potential [139] and $E(x)$ is its potential function, meaning that minimizing $E(x)$ yields the Wardrop equilibrium.

Now suppose that bounded deviations of a maximum distance ρ from the solution are allowed. We wish to converge to a point $\sigma^*_K = \sigma(x^*_K)$ such that a certain number of randomized facets of the feasibility region Π_K intersect $\mathbb{B}(\sigma(x^*_K), \rho)$, where $\Pi_K = \{\sigma \in \sum_{i=1}^N X_i : \; \sigma \leq b(\delta^{(k)}), \; k = 1,\dots,K\}$ with $\sum_{i=1}^N X_i = \{\sigma \in \mathbb{R}^n : \sigma = \sigma((x_i)_{i\in\mathcal{N}}), x_i \in X_i \text{ for any } i \in \mathcal{N}\}$. Allowing some coupling constraints to intersect $\mathbb{B}(\sigma(x^*_K), \rho)$ can lead to improvements of the overall system performance, described by the game's potential function E. As such, a data-driven algorithm that achieves this trade-off between

feasibility and performance is required. In other words, we wish to design an algorithm that intersects only with that subset of coupling constraints, of cardinality at most M, that leads to the minimization of the potential function E that resembles the social welfare optimum. As the choice of M will affect the level of probabilistic guarantees as outlined in the next section, the lower M, the better the guarantees in terms of the probability of constraint violation. To better understand the desired behaviour of such an algorithm we simplify our example, assuming that σ is a 2-dimensional vector, i.e., $\sigma(x) = [\sigma_1(x)\ \sigma_2(x)]^T$. We can then study the problem at hand in the coordinates of the aggregate (σ_1, σ_2). Note that for this particular example, if the desired algorithm achieves intersection with at most M coupling constraints, then any other p-norm necessarily intersects with at most M coupling constraints, as well. Since our results hold for any p-norm, to better illustrate how the sets change by tuning M, we use the 1-norm.

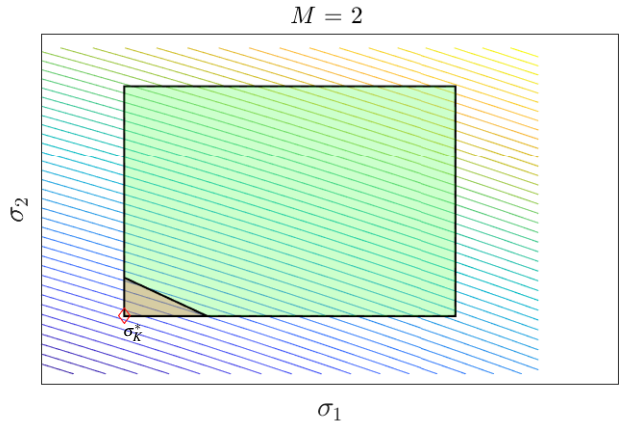

Figure 4.2: The box constraints that form the feasibility region Π_K are due to the drawn coupling constraints. The contour plot of the potential function is also shown with light yellow denoting the largest values of the potential function and dark blue the lowest ones. For $M = 2$, we allow at most two coupling constraints to intersect $\mathbb{B}_1(\sigma_K^*, \rho)$, which leads to the best possible value of the potential function for the considered multi-sample. The green region denotes the tightened feasible region of the algorithm which in this case coincides with the feasibility region Π_K.

The desired behaviour of the algorithm we propose can be summarized in Figures 4.2, 4.3, 4.4. Note that the feasibility region Π_K in our example is formed entirely by the drawn coupling constraints. This is without loss of generality as deterministic constraints could

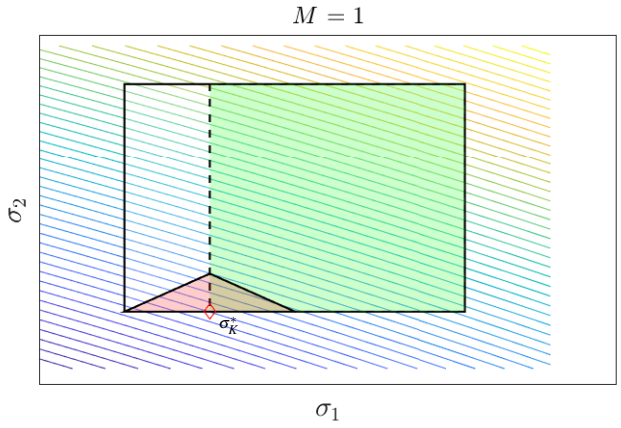

Figure 4.3: For $M = 1$, we trade some of the efficiency of the equilibrium to improve the quality of the *a priori* probabilistic certificates by restricting the number of coupling constraints allowed to intersect $\mathbb{B}_1(\sigma(x_K^*), \rho)$. The value of the potential function for the new equilibrium will be less efficient than the one in Figure 4.2, however now only the lower bound of σ_2 intersects with $\mathbb{B}_1(\sigma(x_K^*), \rho)$ leading to better theoretical certificates. Note that the algorithm should automatically choose to intersect that coupling constraint which leads to the smallest possible value of the potential function. The set in green denotes the desired virtual domain of such an algorithm.

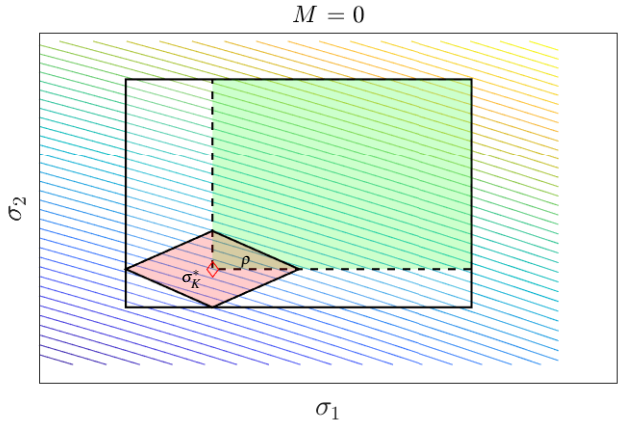

Figure 4.4: For $M = 0$ we expect the algorithm not to allow $\mathbb{B}_1(\sigma_K^*, \rho)$ to intersect any coupling constraints of the feasibility region. This is achieved by tightening both lower bounds by the radius ρ of the region. Then, the tightened feasible region of the algorithm is the region in green.

still be considered. Our algorithm iteratively tightens certain constraints of the feasible region. At convergence our solution would belong to the tightened feasibility domain. The number of constraints to be tightened depends on the choice of M. The contour plot of the $E(x)$ is also shown, with light yellow depicting its largest values and with dark blue its lowest ones. The region S_K^* (in pink) for which we provide tunable *a priori* guarantees changes depending on the value of M. The desired behaviour of the algorithm for $M = 2$

is illustrated in Figure 4.2. In that case, we allow at most two coupling constraints to intersect $\mathbb{B}(\sigma_K^*, \rho)$, which leads to the best possible value of the potential function for the considered multi-sample. As we will see, this comes at the cost of more conservative theoretical guarantees on the probabilistic feasibility of the region S_K^* under study.

In case we wish to trade efficiency of the equilibrium for improvement in the quality of the *a priori* probabilistic certificates, we can restrict the number of coupling constraints allowed to intersect $\mathbb{B}_1(\sigma(x_K^*), \rho)$ to $M = 1$, as illustrated in Figure 4.3 by tightening one constraint. As we can see in Figure 4.3, the value of the potential function for the new equilibrium will be less efficient compared to that of Figure 4.2, however now only the lower bound of σ_2 intersects with $\mathbb{B}_1(\sigma(x_K^*), \rho)$. Note that choosing to intersect the lower bound of σ_1 instead would lead to an even higher value of the potential function. As such, a desired property of the algorithm is to be able to select online which subset of at most M coupling constraints should intersect the region of deviations such that the best possible value of the potential function is returned among all possible choices. The set in green denotes the desired virtual domain of such an algorithm.

In case we wish to provide the best possible *a priori* theoretical guarantees, we expect the algorithm not to allow $\mathbb{B}_1(\sigma_K^*, \rho)$ to intersect any coupling constraints ($M = 0$) of the feasibility region. This is achieved by tightening both lower bounds by the radius ρ of the region as shown in Figure 4.4. Through this example, we can see there is a trade-off between the level of probabilistic feasibility guarantees we can offer (as detailed in the next section) that depends on M, and the cost achieved by the returned solution. Fixing M, the algorithm presented in the next section progressively tightens the feasibility region to return the solution that achieves the desired trade-off. Note that enforcing the constraints on the number of intersections M might have implications on the feasibility of the problem. In this thesis we only deal with cases where the resulting problem is feasible, hence the proposed algorithm can indeed converge to the desired solution. The following section formally introduces the proposed algorithmic scheme and proves its convergence under this condition.

4.2.4 A data-driven algorithm for tuning the probabilistic feasibility

Let $y = [x, \mu]^T$ be the vector that comprises the concatenation of the total decision vector $x = (x_i)_{i \in \mathcal{N}}$ and the Lagrange multipliers $\mu \in \mathbb{R}^m$ corresponding to the coupling constraints. Consider that the domain of y is Y, where $Y = X \times \mathcal{M}$ with $\mathcal{M} = \{\mu \in \mathbb{R}^m : \mu^{(k+1)} < \mu^{(k)} - \varphi$, for any $k = 1, \ldots, m-1\}$, where φ is a positive real. The essence of restricting the Lagrange multipliers to $\mathcal{M} \subset \mathbb{R}^m$ stems from the fact that convergence of the proposed algorithm can only be guaranteed if the mapping is continuous. Assuming that the ordered elements of the multiplier admit positive differences alleviates any discontinuity issues. In practice, restricting the set of multipliers to $\mathcal{M}$ might not be required to converge to the desired solution. In our set-up, m is the number of constraints that belong to the irredundant H-representation of the polytopic feasibility region Π_K, i.e., $\Pi_K = \{x \in \bar{X} : a(\delta_k)^T x \leq b(\delta_k),$ for all $k \in I_m\}$, where $\bar{X}$ is the subset of X that is returned after a redundancy detection algorithm is applied to the polytope Π_K and I_m is the index set (of cardinality m) of the samples that support the entire feasibility region Π_K. Let $A = (a(\delta_k))_{k \in I_m}$ and $b = (b(\delta_k))_{k \in I_m}$. The following operator is key for the proposed algorithm

$$T(y) = \begin{bmatrix} F(x) + A^T \mu \\ -(Ax - b + Q(\mu)) \end{bmatrix}.$$

The operator $Q(\mu) = P^T(\mu) R(M) P(\mu) \rho$ comprises two different mappings namely P and R.

1. Mapping $P : \mathbb{R}^m \to \mathbb{R}^{m \times m}$ is responsible for ordering elements of a multiplier vector in descending order. The reason for that is to identify which multipliers have a larger value. This is of crucial importance as in the proposed algorithm, we tune M *a priori* and compute a solution such that we allow the region of deviations $\mathbb{B}(x_K^*, \rho)$ to intersect with at most M coupling constraints [3]. However, those M constraints

[3] Note that we use a common tightening for each constraint, which based on the 2-norm can be defined as $\rho = \max_{j=1,\ldots,m} \frac{\bar{\rho}}{\|A_j\|}$, where $\bar{\rho}$ is a radius parameter that we set and j is the row index of matrix A. This worst-case tightening implies that sometimes the problem might be infeasible. For the case of box-constraints on the aggregate decision we have that $\rho = \bar{\rho}$. Even though for box-bounded constraints this approach is not conservative, for general affine constraints it can be. Extending our algorithm to account for a different radius per constraint, thus reducing the conservatism of the algorithm with respect to constraint tightening in the general case, is left for future research.

are not selected randomly among all m possible ones. In fact, they correspond to the elements of the multipliers with the highest value as the intersection of $\mathbb{B}(x_K^*, \rho)$ with these constraints leads, as numerical evidence suggests, to a smaller total cost for the population compared with relaxing other constraints.

2. The selection of how many constraints will be relaxed is the main task of mapping $R : \mathbb{N} \to \mathbb{R}^{m \times m}$. R takes as input the number of coupling constraints M we allow $\mathbb{B}(x_K^*, \rho)$ to intersect with and returns as output a matrix that coincides with the identity matrix except for the first M rows which are set to zero, i.e., $R(M) = [0_{M \times m}; [e_{M+1}^T; \ldots; e_m^T]]$. Setting some of the rows to zero allows the aforementioned intersections of $\mathbb{B}(x_K^*, \rho)$ with M coupling constraints. It is then left to the mapping P to perform a descent ordering and determine which constraints are the most important and should thus be intersected.

Algorithm 1 builds upon the class of asymmetric projection algorithms [63]. Note that without the presence of mapping $Q : \mathbb{R}_{\geq}^m \times \mathbb{N} \times \mathbb{R}_{\geq} \to \mathbb{R}_{\geq}^m$ the standard VI mapping as studied in [3] is retrieved. Unlike standard asymmetric projection algorithms though, due to Q, Algorithm 1 converges to a different solution for a different radius ρ and number of maximum intersections M. For fixed ρ and M, the returned solution x_K^* is such that $\mathbb{B}(x_K^*, \rho)$ intersects at most M facets, which tunes the theoretical violation level ε for the admissible region surrounding the equilibrium. Both ρ and M are fixed throughout the iterative process.

The proposed iterative scheme that achieves tunable probabilistic feasibility guarantees for the entire region of deviations is given in Algorithm 1 below, where $D = \begin{bmatrix} \frac{1}{\tau} I_{Mn} & 0 \\ -2A & \frac{1}{\tau} I_m \end{bmatrix}$. Note that the mapping T is a function of M, ρ and the multi-sample δ_K, which are considered fixed for the entire iterative process. As an intermediate

Algorithm 1

Initialization: $y_{(0)}, \rho, \delta_K, M$
$\kappa \leftarrow 0$
Repeat until convergence
$y_{(\kappa+1)} \leftarrow \Pi_{X \times \mathcal{M}, D} \left[y_\kappa - D^{-1} T(y_{(\kappa)}) \right]$
$\kappa \leftarrow \kappa + 1$

step towards proving convergence of the proposed data-driven algorithm, we introduce the following lemma

Lemma 4.2 *1. The mapping T is continuous on $X \times \mathcal{M}$.*

2. For any $\mu, \mu' \in \mathcal{M}$ we have that $-(\mu - \mu')^T (Q(\mu) - Q(\mu')) > h\varphi\rho$, where h is a positive integer.

□

Proof: 1) We wish to prove that the mapping T is continuous on $X \times \mathcal{M}$. Notice that T is by construction continuous on $X \times \mathcal{M}$ when the operator Q is continuous on $\mathcal{M}$. As such, it is sufficient to show that for any $\mu, z \in \mathcal{M}$ and any $\bar{\varepsilon} > 0$, there exists a $\bar{\delta} > 0$ such that if $||\mu - z|| \leq \bar{\delta}$, then $||Q(\mu) - Q(z)|| \leq \bar{\varepsilon}$. To this end, consider any $\mu \in \mathbb{R}^m$ and $z \in \mathcal{M}$, where we repeat for the ease of the reader that $\mathcal{M} = \{\mu \in \mathbb{R}^m : \mu^{(k+1)} < \mu^{(k)} - \varphi$, for any $k = 1, \ldots, m\}$ and $\varphi > 0$. Suppose that z lies within a neighbourhood around μ with radius $\frac{\varphi}{2}$, i.e., $||\mu - z|| < \frac{\varphi}{2}$. We denote $\mu_i^{(k)}$ and $z_i^{(k)}$ the k-th largest elements of vectors μ and z, respectively. Let k^* be the smallest index among $k = 1, \ldots, m$ for which the k-th elements of μ and z appear in different rows. Furthermore, let I denote the set of indices for which the ordering of the elements of μ and z agrees up to $k < k^*$, i.e., for all $i \in I$, there exists $k < k^*$ such that $\mu_i^{(k)}$ and $z_i^{(k)}$ appear in the same row i. Assume for the sake of contradiction that μ and z differ in k^*, i.e., that there exists $i \neq j$ with $i, j \notin I$ such that we have $z_j^{(k^*)}$, the k^*-th largest element of z in row j, and $\mu_i^{(k^*)}$ the k^*-th largest element of μ in row i. As such, we will have another element $z_i^{(k_1)}$ in row i, where $k_1 < k$ and another element $\mu_j^{(k_2)}$, where $k_2 > k^*$. Since $k > k^*$ and $z \in \mathcal{M}$, we have that $z_i^{(k_1)} < z_j^{(k^*)} - \varphi$. Since, $||\mu - z|| < \frac{\varphi}{2}$, we have that

$$z_i^{(k_1)} < \mu_j^{(k_2)} + \frac{\varphi}{2} - \varphi. \tag{4.6}$$

Since $z_i^{(k_1)} > \mu_i^{(k^*)} - \frac{\varphi}{2}$, $||z_i^{(k_1)} - \mu_i^{(k^*)}|| < \frac{\varphi}{2}$ and $\mu_i^{(k^*)} > \mu_j^{(k_2)}$, we have that

$$z_i^{(k_1)} > \mu_j^{(k_2)}. \tag{4.7}$$

Combining (4.6) and (4.7) we have that $\mu_j^{(k_2)} < \mu_j^{(k_2)}$, which is a contradiction. As such, all k-th largest elements of μ and z are in the same row, which implies that both of them

have the same permutation matrix, i.e, $P(\mu) = P(z)$. As such, $||Q(\mu) - Q(z)|| = 0 < \bar{\varepsilon}$ for any $\mu \in \mathbb{R}^K_{>0}$ and any $z \in \mathcal{M}$. Since $\mathcal{M} \subset \mathbb{R}^m$, this establishes the continuity of $Q(\mu)$ on $\mathcal{M}$, thus concluding the proof of the first part.

2) Let μ, μ' be any arbitrary vectors in $\mathcal{M}$. Furthermore, assume that the elements of μ_1, μ_2 can be arranged in descending order, i.e., $\mu_{i_1} > \mu_{i_2} > \cdots > \mu_{i_m}$ and $\mu'_{i'_1} > \mu'_{i'_2} > \cdots > \mu'_{i'_m}$, where $i_1, i_2, \ldots, i_m$ and $i'_1, i'_2, \ldots, i'_m$ are elements of $\{1, \ldots, m\}$. To ease notation, we set $I = \{i_1, i_2, \ldots, i_m\}$ to be the ordered indices of μ and $I' = \{i'_1, i'_2, \ldots, i'_m\}$ to be the ordered indices of μ'. Then, the first M elements of I denoted as R_I and the first M elements of I' denoted as $R_{I'}$ correspond to the elements of μ and μ', respectively, that will not be tightened by ρ as they correspond to the M largest multipliers. Equivalently, the corresponding complements R^c_I, $R^c_{I'}$ correspond to the elements of the multipliers μ and μ', respectively, that are chosen to be tightened by a preset radius ρ. Denote $(\bar{\rho}_t)_{t=1}^m = Q(\mu) - Q(\mu')$. We distinguish among the following cases

1. $t \in R^c_I \cap R_{I'}$: In that case, as $t \in R^c_I$, μ_t will be tightened by ρ through mapping Q. However, since $t \in R_{I'}$, μ'_t will not be tightened by ρ. As such, we have that $\bar{\rho}_t = \rho$.

2. $t \in R_I \cap R^c_{I'}$: In that case, as $t \in R^c_{I'}$, μ'_t will be tightened by ρ through mapping Q. However, since $t \in R_I$, μ_t will not be tightened by ρ. As such, we have that $\bar{\rho}_t = -\rho$.

3. $t \in (R_I \cap R_{I'}) \cup (R^c_I \cap R^c_{I'})$. If $t \in R_I \cap R_{I'}$, then μ_t and μ'_t will not be tightened by ρ. As such, $\bar{\rho}_t = 0$. Similarly, if $t \in R^c_I \cap R^c_{I'}$, then both μ_t and μ'_t will be tightened by ρ, which leads again to $\bar{\rho}_t = 0$.

Since the sets $R^c_I \cap R_{I'}$, $R_I \cap R^c_{I'}$, $(R_I \cap R_{I'}) \cup (R^c_I \cap R^c_{I'})$ are pairwise disjoint and

exhaust the set $\{1,\dots,m\}$, we have that

$$\begin{aligned} U &= (\mu-\mu')^T(Q(\mu)-Q(\mu')) = \sum_{t=1}^{m}(\mu_t-\mu'_t)\bar{\rho}_t \\ &= \sum_{t_1\in R_I^c\cap R_{I'}}(\mu_{t_1}-\mu'_{t_1})\rho + \sum_{t_2\in R_I\cap R_{I'}^c}(\mu_{t_2}-\mu'_{t_2})\cdot(-\rho) \\ &= (\sum_{t_1\in R_I^c\cap R_{I'}}\mu_{t_1} - \sum_{t_2\in R_I\cap R_{I'}^c}\mu_{t_2} + \sum_{t_2\in R_I\cap R_{I'}^c}\mu'_{t_2} - \sum_{t_1\in R_I^c\cap R_{I'}}\mu'_{t_1})\rho \\ &= (U_1+U_2)\rho, \end{aligned} \tag{4.8}$$

where

$$U_1 = \sum_{t_1\in R_I^c\cap R_{I'}}\mu_{t_1} - \sum_{t_2\in R_I\cap R_{I'}^c}\mu_{t_2}$$

and

$$U_2 = \sum_{t_2\in R_I\cap R_{I'}^c}\mu'_{t_2} - \sum_{t_1\in R_I^c\cap R_{I'}}\mu'_{t_1}.$$

Note that for any $t_1\in R_I^c\cap R_{I'}$ and $t_2\in R_I\cap R_{I'}^c$ it holds that $\mu_{t_1}<\mu_{t_2}$ and $\mu'_{t_2}<\mu'_{t_1}$. This is because $R_I\cap R_{I'}^c$ contains the M largest elements of μ, while $R_I^c\cap R_{I'}$ contains the rest $m-M$ smallest ones. Similarly, $R_I^c\cap R_{I'}$ contains the M largest elements of μ', while $R_I\cap R_{I'}^c$ contains the rest $m-M$ smallest ones. Furthermore, denote $R_I^c\cap R_{I'}=\{t_1^{(1)},t_1^{(2)},\dots,t_1^{(h_1)}\}$, $R_I\cap R_J^c=\{t_2^{(1)},t_2^{(2)},\dots,t_2^{(h_2)}\}$, where h_1,h_2 are positive integers. Then, $h_1=|R_I^c\cap R_{I'}|=|R_{I'}\setminus R_I|=|R_{I'}|-|R_I\cap R_{I'}|=|R_{I'}|-|R_{I'}\cap R_I|=|R_I\setminus R_{I'}|=|R_I\cap R_{I'}^c|=h_2$. This directly implies that $U_1<0$ and $U_2<0$. To obtain a lower bound we have

$$\begin{aligned} U_1 &= \sum_{t_1\in R_I^c\cap R_{I'}}\mu_{t_1} - \sum_{t_2\in R_I\cap R_{I'}^c}\mu_{t_2} \\ &= \sum_{\ell=1}^{h_1}(\mu_{t_1^{(\ell)}}-\mu_{t_2^{(\ell)}}) < -h_1\delta, \end{aligned} \tag{4.9}$$

where the last equality comes from the fact that $\mu\in\mathcal{M}$. Following the same process for U_2 we have

$$\begin{aligned} U_2 &= \sum_{t_2\in R_I\cap R_{I'}^c}\mu'_{t_2} - \sum_{t_1\in R_I^c\cap R_{I'}}\mu'_{t_1} \\ &= \sum_{\ell=1}^{h_2}(\mu'_{t_2^{(\ell)}}-\mu'_{t_1^{(\ell)}}) < -h_2\delta, \end{aligned} \tag{4.10}$$

where the last equality is again due to $\mu \in \mathcal{M}$. Setting $h = h_1 + h_2$ and combining (4.9), (4.10) and (4.8), we obtain that $U < -h\varphi\rho$, thus concluding the proof of the second part. ■

The second part of Lemma 4.2 is useful for determining an appropriate bound on the step-size τ such that we guarantee convergence of the proposed algorithm to a solution of VI(T,Y). The range of values for τ such that covergence of the iterative procedure is guaranteed is established in the theorem that follows.

Theorem 4.2 *For any fixed multi-sample δ_K and given a positive integer M, consider Assumptions 4.1, 4.2 and 4.3. Then, if*

$$\tau < \min\left\{ \frac{-\rho^2 + \sqrt{\rho^2 + 4\delta^2||A||^2}}{2||A||^2\delta}, \frac{-L_F^2 + \sqrt{L_F^4 + 4a^2||A||^2}}{2a||A||^2} \right\},$$

Algorithm 1 converges to the solution of VI(T,Y). □

Proof: We will show that Algorithm 1 converges to the solution of VI(T,Y). Define $D_s = \frac{D+D^T}{2}$ and $L = D_s^{-1/2}$. Furthermore, consider the transformation $Ly = [v,z]^T$. Since $D = \begin{bmatrix} \frac{1}{\tau}I_{Mn} & 0 \\ -2A & \frac{1}{\tau}I_m \end{bmatrix}$, we have that $D - D_s = \begin{bmatrix} 0 & A^T \\ -A & 0 \end{bmatrix}$. Note that the mapping T can equivalently be written as

$$T(y) = \begin{bmatrix} F(x) \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & A^T \\ -A & 0 \end{bmatrix} y + \begin{bmatrix} 0 \\ b - Q(\mu) \end{bmatrix}.$$

Since T is continuous over $\mathcal{M}$ due to Lemma 4.2, to establish convergence it is sufficient to find a condition for τ such that the mapping $G(y) = LT(Ly) - L(D - D_s)Ly$ is co-coercive, i.e.,

$$(y_1 - y_2)^T(G(y_1) - G(y_2)) \geq ||G(y_1) - G(y_2)||^2,$$

for any $y_1, y_2 \in Y$ (see [63], Section 12.5.1). Let $H = (y_1 - y_2)^T(G(y_1) - G(y_2)) - ||G(y_1) - G(y_2)||^2$. Then it suffices to find $\tau > 0$ such that $H \geq 0$. Since $Ly = [v,z]^T$, we have that

$$G(y) = LT(Ly) - L(D - D_s)Ly = L \begin{bmatrix} F(v) \\ 0 \\ b - Q(z) \end{bmatrix}. \tag{4.11}$$

From (4.11) and setting $Ly_1 = [v_1, z_1]^T$, $Ly_2 = [v_2, z_2]^T$ and $L^T L = W$ for any $(y_1, y_2) \in Y \times Y$ we obtain the following.

$$\begin{aligned} H &= (y_1 - y_2)^T (G(y_1) - G(y_2)) - ||G(y_1) - G(y_2)||^2 \\ &= (Ly_1 - Ly_2)^T \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix} - ||L \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix} ||^2 \\ &= \begin{bmatrix} v_1 - v_2 \\ z_1 - z_2 \end{bmatrix}^T \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix} - \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix}^T L^T L \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix} \\ &= \begin{bmatrix} v_1 - v_2 \\ z_1 - z_2 \end{bmatrix}^T \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix} - \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix}^T W \begin{bmatrix} F(v_1) - F(v_2) \\ Q(z_2) - Q(z_1) \end{bmatrix}, \end{aligned} \tag{4.12}$$

where $||\cdot||$ denotes the Euclidean norm. Matrix W can be written as $W = \begin{bmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{bmatrix}$, where $W_{11} \in \mathbb{R}^{nN \times nN}, W_{12} \in \mathbb{R}^{nN \times m}, W_{21} \in \mathbb{R}^{nN \times m}, W_{33} \in \mathbb{R}^{m \times m}$. The values for each block matrix of W are $W_{11} = \tau(I_n - \tau^2 A^T A)^{-1}$, $W_{12} = W_{21}^T = \tau^2 (I_n - \tau^2 A^T A)^{-1} A^T$, $W_{22} = \tau I_m + \tau^3 A (I_n - \tau^2 A^T A)^{-1} A^T$. Expanding the inner product in (4.12) with respect to the matrix blocks $W_{11}, W_{12}, W_{21}, W_{33}$ we have that

$$\begin{aligned} H &= (F(v_1) - F(v_2))^T [(v_1 - v_2) - W_{11}(F(v_1) - F(v_2)) - 2W_{12}(Q(z_2) - Q(z_1))] + \\ &\quad + (Q(z_2) - Q(z_1))^T [(z_1 - z_2) - W_{22}(Q(z_2) - Q(z_1))] \\ &= (F(v_1) - F(v_2))^T (v_1 - v_2) - (F(v_1) - F(v_2))^T W_{11} (F(v_1) - F(v_2)) - \\ &\quad - 2(F(v_1) - F(v_2))^T W_{12} (Q(z_2) - Q(z_1)) + (Q(z_2) - Q(z_1))^T (z_1 - z_2) - \\ &\quad - (Q(z_2) - Q(z_1))^T W_{22} (Q(z_2) - Q(z_1)). \end{aligned} \tag{4.13}$$

Setting

$$\begin{aligned} p_\tau &= (I - \tau^2 A^T A)^{-1/2} (F(v_1) - F(v_2)) \text{ and} \\ q_\tau &= \tau (I - \tau^2 A^T A)^{-1/2} A^T (Q(z_2) - Q(z_1)) \end{aligned}$$

and using the inequality $(p_\tau + q_\tau)^T (p_\tau + q_\tau) \leq 2(p_\tau^T p_\tau + q_\tau^T q_\tau)$ to relation (4.13) we obtain

$$\begin{aligned} H =& (F(v_1) - F(v_2))^T (v_1 - v_2) + (Q(z_2) - Q(z_1))^T (z_1 - z_2) \\ & - \tau (Q(z_2) - Q(z_1))^T (Q(z_2) - Q(z_1)) - \tau (p_\tau + q_\tau)^T (p_\tau + q_\tau) \\ \geq & (F(v_1) - F(v_2))^T (v_1 - v_2) + (Q(z_2) - Q(z_1))^T (z_1 - z_2) \\ & - \tau ||Q(z_2) - Q(z_1)||^2 - 2\tau (p_\tau^T p_\tau + q_\tau^T q_\tau). \end{aligned}$$

Note that F satisfies Assumption 4.3. As such, F satisfies relations (4.3) and (4.4). Furthermore, from Lemma 4.2 we have that

$$(Q(z_2)-Q(z_1))^T(z_1-z_2) \geq h\delta\rho. \tag{4.14}$$

Finally the following relation

$$||Q(z_2)-Q(z_1)||^2 \leq h\rho^2, \tag{4.15}$$

provides an upper bound for the normed difference of the mapping Q. Replacing p_τ and q_τ and taking (4.3), (4.4), (4.14) and (4.15) into account we have that

$$\begin{aligned}
&(F(v_1)-F(v_2))^T(v_1-v_2)+(Q(z_2)-Q(z_1))^T(z_1-z_2)\\
&-\tau||Q(z_2)-Q(z_1)||^2-2\tau(p_\tau^Tp_\tau+q_\tau^Tq_\tau)\\
&\geq a||v_1-v_2||^2+h\delta\rho-\tau h\rho^2\\
&-2\tau(F(v_1)-F(v_2))^T(I-\tau^2A^TA)^{-1}(F(v_1)-F(v_2))\\
&-2\tau^3(Q(z_2)-Q(z_1))^TA(I-\tau^2A^TA)^{-1}A^T(Q(z_2)-Q(z_1))\\
&\geq a||v_1-v_2||^2+h\delta\rho-\tau h\rho^2\\
&-2\tau||F(v_1)-F(v_2)||^2||(I-\tau^2A^TA)^{-1}||\\
&-2\tau^3||Q(z_2)-Q(z_1))||^2||(I-\tau^2A^TA)^{-1}||||A||^2\\
&\geq (a-2\tau L_F^2||(I-\tau^2A^TA)^{-1}||)||v_1-v_2||^2+h\delta\rho-\tau h\rho^2\\
&-2\tau^3h\rho^2\frac{1}{1-\tau^2||A||^2}||A||^2\\
&\geq (a-2\tau L_F^2||(I-\tau^2A^TA)^{-1}||)||v_1-v_2||^2\\
&+h\delta\rho-\frac{\tau h\rho^2}{1-\tau^2||A||^2}.
\end{aligned}$$

For the last inequality we assume that $\tau<1$. As such, to satisfy $H\geq 0$, we can choose a $\tau\in(0,\max\{\frac{1}{||A||},1\})$ such that

$$a-2\tau L_F^2||(I-\tau^2A^TA)^{-1}||>0 \tag{4.16}$$

and

$$h\delta\rho-\frac{\tau h\rho^2}{1-\tau^2||A||^2}>0. \tag{4.17}$$

Solving (4.16) and (4.17) with respect to τ shows that Algorithm 1 converges to the solution of VI(T,Y) if additionally

$$\tau < \min\left\{ \frac{-\rho^2+\sqrt{\rho^2+4\delta^2||A||^2}}{2||A||^2\delta}, \frac{-L_F^2+\sqrt{L_F^4+4a^2||A||^2}}{2a||A||^2} \right\}$$

is satisfied. This concludes then the proof. ■

We now wish to accompany region S_K^* around x_K^* returned by Algorithm 1 with *a priori* probabilistic feasibility guarantees tuned by M. This result is established in Theorem 4.3.

Theorem 4.3 *For a given multi-sample δ_K consider Algorithm 1 and Assumptions 4.1, 4.2, 4.3, 4.4. Fix M and the violation level $\varepsilon \in (0,1)$. Then, we have that*

$$\mathbb{P}^K\left\{ \delta_K \in \Delta^K : \mathbb{V}(S_K^*) > \varepsilon \right\} \leq \sum_{i=0}^{nN+M-1} \binom{K}{i} \varepsilon^i (1-\varepsilon)^{K-i}, \tag{4.18}$$

where S_K^* is the set returned by Algorithm 1. □

Proof: The elements of the minimal compression set I of Algorithm 1 can belong to one or both of the following subsets:

1. The subset I_1 of samples that support the solution x_K^*. Note that since Algorithm 1 converges to the point (x_K^*, μ_K^*) for a fixed choice of M, $Q(\mu_K^*)$ will be a fixed quantity. As such, the solution x_K^* of Algorithm 1 can be obtained as the unique solution to

$$\begin{aligned} &\text{Find } x_K^* \in \bar{\Pi}_K \text{ such that} \\ &F^T(x_K^*)(x - x_K^*) \geq 0 \text{ for any } x \in \bar{\Pi}_K, \end{aligned} \tag{4.19}$$

   where $\bar{\Pi}_K$ is the resulting polytope after at most M coupling constraints of the minimal H-representation of Π_K have been selected by $Q(\mu_K^*)$ to be tightened by ρ. The constraints in (4.19) are equivalent to $F^T(x_K^*)x \geq F^T(x_K^*)x_K^*$ for all $x \in \bar{\Pi}_K$. As such, x_K^* achieves the minimum value of the following problem

$$\begin{aligned} &\min_{x \in \mathbb{R}^{nN}} F^T(x_K^*)x \\ &\text{subject to } x \in \bar{\Pi}_K, \end{aligned} \tag{4.20}$$

and is also unique due to Lemma 4.1. The cost function is linear in x and $\bar{\Pi}_K$ is convex by Assumption 4.2, hence we obtain the standard scenario program considered in [21]. Applying Proposition 1 in [21] to problem (4.20), we have that $|I_1| \leq nN$, i.e., the number of support samples of its optimizer x_K^* is bounded by the dimension of the decision vector nN.

2. The subset I_2 of samples whose corresponding coupling constraints intersect $\mathbb{B}(x_K^*, \rho)$. By construction of Algorithm 1 we have that $|I_2| \leq M$.

As such, we have that $I = I_1 \cup I_2$ is a compression set with cardinality $|I| = |I_1 \cup I_2| \leq |I_1| + |I_2| \leq nN + M$. This implies by [121] that

$$\mathbb{P}^K\Big\{\delta_K \in \Delta^K : \mathbb{V}(S_K^*) > \varepsilon\Big\} \leq \sum_{i=0}^{nN+M-1} \binom{K}{i} \varepsilon^i (1-\varepsilon)^{K-i}, \tag{4.21}$$

which concludes then the proof. ■

As such, Theorem 4.3 guarantees that, for an adequate number of samples, the chance-constraint (4.5) can be satisfied with high confidence $1 - \sum_{i=0}^{nN+M-1} \binom{K}{i} \varepsilon^i (1-\varepsilon)^{K-i}$. Furthermore, the fact that the right-hand side in (4.18) depends on M provides the ability to provide *a priori* feasibility guarantees by tuning the number M of intersections through Algorithm 1. It has been observed through extensive numerical simulations that the cardinality of the minimal compression set can in fact be bounded by the tighter bound $\max\{nN, M\}$. However, this claim remains as a conjecture.

4.3 Motivating example revisited- A numerical study

We now revisit the numerical example of Subsection 4.2.3 to show how the Algorithm 1 converges to the desired solution for which we can now claim tunable *a priori* guarantees. We set the number of agents to $N = 50$ and keep the dimension of the decision vector at $n = 2$. For each agent $i \in \mathcal{N}$ and each $t \in \{1, \ldots, n\}$, $\underline{x}_i = 0$ and $\bar{x}_i = 3.5$ and $b = \max_{1,\ldots,K} b(\delta_k) = \begin{pmatrix} 0.2 & 0.4 & 1 & 1.6 \end{pmatrix}^T$. We assume that each agent can deviate no more than $\rho = 0.2$ from the aggregate equilibrium and use Algorithm 1 to find $\sigma^* = \sigma(x_K^*)$ for

different values of M. Note that we can easily modify Algorithm 1 to return the aggregate equilibrium σ_K^* by considering the mapping

$$T(\sigma,\mu)=\begin{bmatrix}\nabla_\sigma E(\sigma)+A^T\mu\\ -(\sigma-b+Q(\mu))\end{bmatrix}.$$

The following corollary provides then a tighter *a priori* result for this particular case study that does not depend on the number of agents.

Corollary 4.1 *For a given multi-sample $\delta_K\in\Delta^K$, consider game G'_K and Assumption 4.4. Fix M and the violation level $\varepsilon\in(0,1)$. Then,we have that*

$$\mathbb{P}^K\Big\{\delta_K\in\Delta^K:\ \mathbb{V}(S_K^*)>\varepsilon\Big\}\leq\sum_{i=0}^{n+M-1}\binom{K}{i}\varepsilon^i(1-\varepsilon)^{K-i} \tag{4.22}$$

where $S_K^*=\Pi_K\cap\mathbb{B}(\sigma_K^*,\rho)$ with σ_K^* being the output of Algorithm 1. □

Proof: In the spirit of the proof of Theorem 3 in [138], we have for the aggregate constraint $\sigma(x)\leq b(\delta)$ of problem P_K that $s^*\leq n$, where s^* is the number of samples that support σ_K^*. Since the upper limit in the summation in the right hand side of (4.18) is the number of support samples minus 1, the right hand side of relation (4.18) in Theorem 3 can be replaced by the tighter bound $\sum_{i=0}^{n+M-1}\binom{K}{i}\varepsilon^i(1-\varepsilon)^{K-i}$. This concludes then the proof. ■

Note that both probabilistic statements in (4.18) and (4.22) are bounded by a binomial function. One could apply the results of [113] to obtain the sample complexity K that makes this binomial function smaller than a prescribed value. These sample complexities show then how the number of required samples grow with ε.

Figures 4.5, 4.6 and 4.7 illustrate the behaviour of the algorithm for different choices of M in the aggregate space. In each of the plots, blue diamonds represent the iterates of the aggregate $\sigma_{(\kappa)}=\sigma(x_{(\kappa)})$, while the red diamond denotes the reached equilibrium for each choice of M. In Figure 4.5 we allow at most $M=2$ of the coupling constraints that are facets of Π_K to intersect $\mathbb{B}(\sigma_K^*,\rho)$ and Algorithm 1 converges to a point such that this is achieved. In this case, the tightened feasible region (in green) of Algorithm 1 coincides with the entire feasibility region.

In Figure 4.6, we allow $M=1$ coupling constraints of the feasible region Π_K to intersect $\mathbb{B}(\sigma_K^*,\rho)$. In this case, since we allow at most one constraint to intersect the

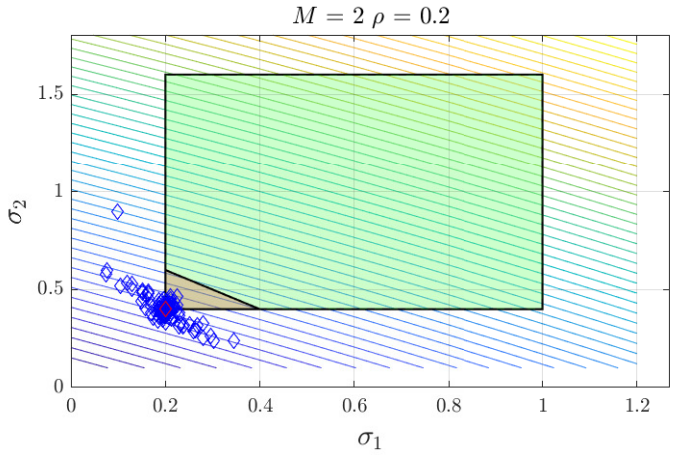

Figure 4.5: Convergence of Algorithm 1 for $M = 2$. At most two of the coupling constraints that constitute facets of Π_K intersect $\mathbb{B}(\sigma_K^*, \rho)$. Algorithm 1 follows a trajectory (blue diamonds) that converges to an equilibrium (red diamond) such that the constraint on the number of possible intersections is satisfied. Note that in this case, the tightened feasible region (in green) of Algorithm 1 coincides with the entire feasibility region.

region of deviations, Algorithm 1 converges to a configuration that does not tighten the lower bound of σ_1 such that an equilibrium is reached with the best value of potential function and at most one intersections. From the contour of the potential function it is apparent that it is a better choice to intersect only with the lower bound on σ_2 instead of the lower bound of σ_1, which would lead to an equilibrium with a higher potential value. As such, the tightened feasible region of the algorithm is the green region of Figure 4.6. Finally, in Figure 4.7 we allow $M = 0$ sampled coupling constraints to intersect the region of deviations. In this case, the tightened feasible region of Algorithm 1 consists of the feasible region where both lower bounds are tightened such that they do not intersect $\mathbb{B}(\sigma_K^*, \rho)$.

Figure 4.8 shows the values of the potential function E evaluated at each iteration $x_{(k)}$ of Algorithm 1. As expected, Algorithm 1 converges to the point with the lowest cost subject to the restrictions on the number of maximum intersections with coupling constraints imposed by the choice of the parameter M. As M increases the probabilistic statement of Corollary 4.1 becomes more conservative as the right-hand of relation (4.22) increases. At the same time, the overall cost decreases, thus achieving a trade-off between feasibility and system level efficiency.

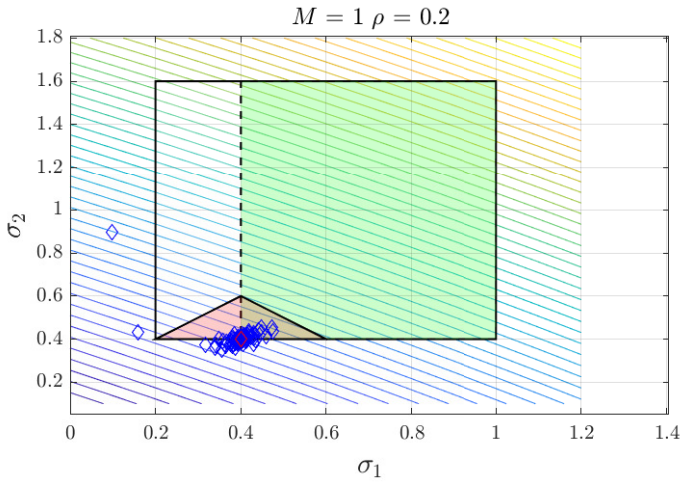

Figure 4.6: Convergence of Algorithm 1 for $M = 1$. Note that Algorithm 1 automatically chooses not to tighten the coupling constraint such that an equilibrium is reached with the smallest potential value and at most one intersections. The green region characterises the tightened feasible region of Algorithm 1.

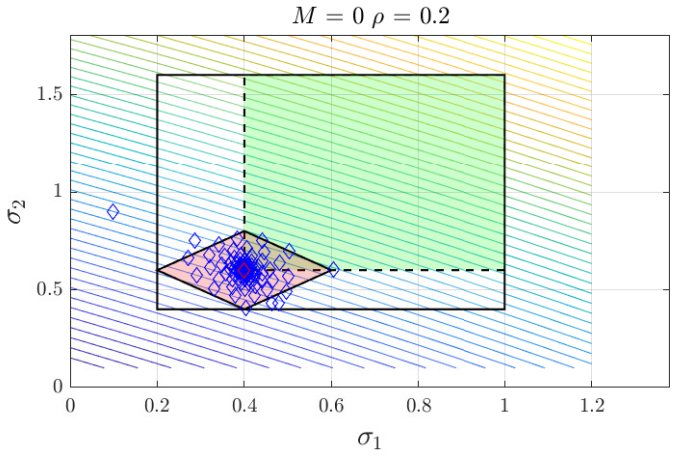

Figure 4.7: Convergence of Algorithm 1 for $M = 0$. Note that both lower bounds are tightened such that they do not intersect $\mathbb{B}(\sigma_K^*, \rho)$. The green region characterises then the tightened feasible region of Algorithm 1 and the algorithm converges to the minimum of this region.

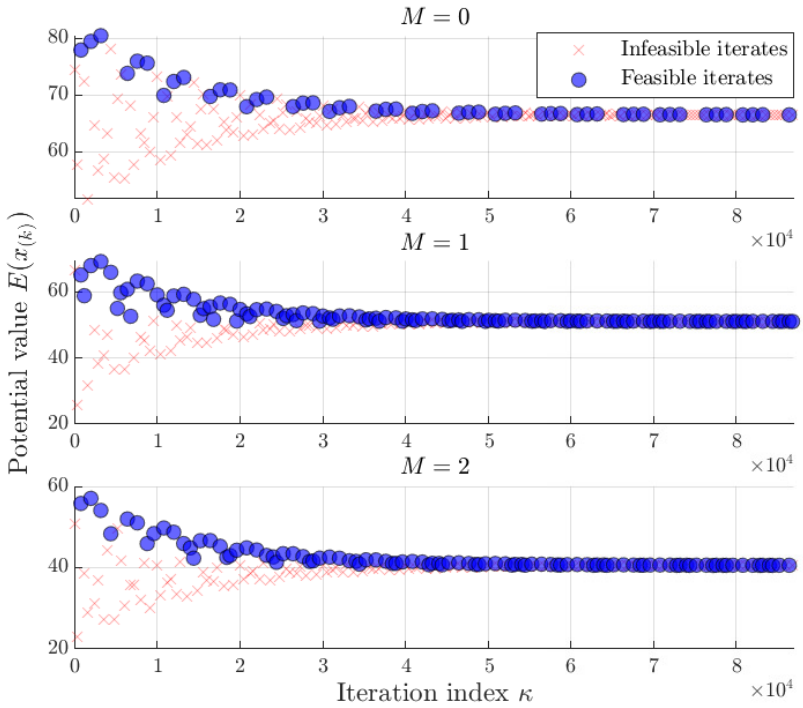

Figure 4.8: The potential function $E(x_{(\kappa)})$ evaluated at different iterations of Algorithm 1. Note that as M increases the probabilistic certificates provided for the region under study become more conservative as the right-hand of relation (4.22) increases, while the overall cost decreases. This highlights the trade-off between the probabilistic feasibility of the set of deviations around an equilibrium and the efficiency of that equilibrium.

4.4 Concluding remarks

This chapter proposes a data-driven equilibrium-seeking algorithm such that *a priori* tunable probabilistic feasibility guarantees are provided for a set around a game solution. The guarantees are tuned by selecting the number of facets that intersect this set. The scheme explored in Algorithm 1 converges to the minimizer of a potential function, such that the preselected guarantees are achieved. Formally proving that the point of convergence of Algorithm 1 exhibits this optimality property is left for future work, however, we note that our probabilistic analysis holds also for games that do not necessarily admit a potential function. Future research will additionally focus on the provision of guarantees on the probability that a new yet unseen constraint will render infeasible a certain portion of the set under study. Specifically, we wish to define a measure of how much a new constraint violates the set under study and extend our theoretical results by providing different guarantees for cases where the violation involves only a small subset of the region and cases where this subset is considerably larger.

5

Agent independent certificates for uncertain multi-agent programs

5.1 Introduction

A major challenge in a plethora of uncertain multi-agent systems is the dependence of the provided certificates on the number of agents. Given that we wish to obtain identical probabilistic guarantees as the population increases, a larger number of samples is required. However, increasing the sample size is, in general, a major issue, which can also lead to increased computational complexity. This fact hinders the provision of non-conservative guarantees for large scale applications. As such, it is of utmost importance to show that for certain classes of problems, common in practical applications, the obtained probabilistic guarantees can be agent independent.

In this chapter, we focus on optimization programs with deterministic constraints, where the cost function depends on uncertainty and admits an aggregate representation of the agents' decisions. By exploiting the structure of the program under study and leveraging the so called support rank notion, we provide agent-independent robustness certificates for the optimal solution, i.e., the constructed bound on the probability of constraint violation does not depend on the number of agents, but only on the dimension of each agent's decision space. This substantially reduces the amount of samples required to achieve a certain level of probabilistic robustness for a larger number of agents. All

robustness certificates provided in this chapter are distribution-free and can be used alongside any optimization algorithm. Our theoretical results are accompanied by a numerical case study of a charging control problem for a fleet of electric vehicles.

A problem formulation similar to ours was considered in [140] and is extended to our set-up to account for the presence of uncertainty in the cost function. Other problems whose structure shares similarities with our work can be found in [91][141], though under a purely deterministic and game-theoretic set-up. Following the recent developments in [142] we show, based on the notion of the support rank [46], that the obtained probabilistic feasibility certificates do not depend on the number of agents. This result directly outperforms probabilistic feasibility statements obtained by a direct application of the scenario approach theory [21], and shows superior scalability properties in multi-agent environments. Note that bounding the maximum number of samples can also be studied under the lenses of order statistics [143] or from the perspective of randomized algorithms [144] , [116], [120]. The problem studied in this chapter is richer as it considers a minmax problem subject to convex constraints. Furthermore, a different route is followed to tackle this problem based on scenario theoretic tools.

The rest of the chapter is organized as follows: Section 5.2 introduces the class of optimization programs under study. Section 5.3 contains the theoretical analysis and proof of the main results, namely, providing agent independent distribution-free probabilistic guarantees for this class of problems. In Section 5.4, numerical studies on the problem of charging a fleet of electric vehicles corroborate our theoretical findings. Finally, Section 5.5 concludes the chapter and provides some potential future research directions.

5.2 Optimization programs with uncertainty in the cost

We consider the following program

$$P : \min_{x \in X} J(x), \tag{5.1}$$

where $J(x) = f(x) + \max_{\delta \in \Delta} g(x, \delta)$ and $f : X \to \mathbb{R}$, $g : X \times \Delta \to \mathbb{R}$ is the deterministic and the uncertain part of the cost function, respectively, where δ is an uncertain parameter that takes values in the (possibly) unknown support Δ. In our set-up, we make no

assumptions on the probability distribution $\mathbb{P}$ that δ follows. In addition, the cost under study satisfies the following assumption

Assumption 5.1 *1. f is jointly convex with respect to all agents' decision vectors, and the set X is non-empty, compact and convex.*

2. g takes the aggregative form

$$g(x,\delta)=\sum_{i\in\mathcal{N}} g_i(x_i,x_{-i},\delta) \text{ and } g_i(x_i,x_{-i},\delta)=x_i^T(A(\delta)\sigma(x)+b(\delta)),$$

where $\sigma : X \to \mathbb{R}^n$ is a mapping $(x_i)_{i\in\mathcal{N}} \mapsto \sum_{i\in\mathcal{N}} x_i$ and $A : \Delta \to \mathbb{R}^{n\times n}$, $b : \Delta \to \mathbb{R}^n$ are uncertain mappings with $A(\delta)$ being a symmetric positive semi-definite matrix for all $\delta \in \Delta$.

Under Assumption 5.1 the function J is convex, as the pointwise maximum of an arbitrary number of convex functions is itself a convex function [50]. From Assumption 5.1.2 the uncertain counterpart of the cost function under study takes the form

$$g(x,\delta)=\sigma(x)^T(A(\delta)\sigma(x)+b(\delta)).$$

The proposed structure captures a wide class of engineering problems, including the electric vehicle charging problem detailed in Section 5.4. Since g is convex, using an epigraphic reformulation we recast P to the equivalent semi-infinite program

$$\begin{aligned} P' : \min_{x\in X,\gamma\in\mathbb{R}} \; & f(x)+\gamma \\ \text{subject to } & h(x,\gamma,\delta)\le 0,\ \forall\ \delta\in\Delta, \end{aligned} \tag{5.2}$$

where $h(x,\gamma,\delta)=g(x,\delta)-\gamma$. In addition, if (x^*,γ^*) is the optimal solution of problem P', then x^* is the optimal solution of the original problem P. Due to the presence of uncertainty and the possibly infinite cardinality of Δ, problem P' is very difficult to solve, without imposing any further assumptions on the geometry of the sample set Δ and/or the underlying probability distribution $\mathbb{P}$. To overcome this issue, we adopt again

a scenario-based scheme [145]. The corresponding scenario program of the uncertain semi-infinite program $P^{'}$ is thus given by

$$\begin{aligned} P^{'}_{SC}: & \min_{x\in X,\gamma\in\mathbb{R}} f(x)+\gamma \\ & \text{subject to } h(x,\gamma,\delta^{(k)})\leq 0,\ \forall\, i\in\mathcal{N}, \end{aligned} \tag{5.3}$$

where $\delta_K \in \Delta^K$ is an i.i.d. multi-sample of cardinality K. For the scenario program under study, we introduce the following assumption:

Assumption 5.2 *1. For any multi-sample δ_K, the scenario program $P^{'}_{SC}$ admits a feasible solution.*

2. The optimal solution (x^*,γ^*) of the scenario program $P^{'}_{SC}$ is unique.

In case multiple optimal points exist, one can use a convex tie-break rule to select a unique solution.

5.3 Agent independent probabilistic guarantees on the cost deterioration

In many practical applications there are cases where a random constraint may leave a linear subspace unconstrained for any possible sample $\delta \in \Delta$. This observation motivated the concept of the support rank as introduced in [46], which allows us to provide tighter probabilistic guarantees for the problem under study. Let $y \in \mathbb{Y} \subseteq \mathbb{R}^d$ and consider the following semi-infinite optimization program

$$\begin{aligned} & \min_{y\in\mathbb{Y}} c^T y \\ & \text{subject to } l(y,\delta)\leq 0,\ \forall\, \delta\in\Delta. \end{aligned} \tag{5.4}$$

Notice that the objective function is linear without loss of generality and in the opposite case an epigraphic reformulation could be introduced. Denoting the collection of all the linear subspaces of $\mathbb{R}^d$ as $\mathscr{L}$, we consider all the linear subspaces $L \in \mathscr{L}$ that, under

the presence of the random constraint (5.4), remain unconstrained for any uncertainty realization $\delta \in \Delta$ and any point $y \in \mathbb{Y}$, i.e., the set

$$\mathscr{U} = \bigcap_{\delta \in \Delta} \bigcap_{y \in \mathbb{Y}} \{L \in \mathscr{L} : L \subset F(y,\delta)\}, \quad \text{where } F(y,\delta) = \{\xi \in \mathbb{R}^d : l(y+\xi,\delta) = l(y,\delta)\}.$$

Definition 5.1 *(Support rank [46])*
The support rank $\rho \in \{0,\ldots,d\}$ of a random constraint equals to the dimension of the problem d minus the dimension of the maximal unconstrained linear subspace L_{max}, i.e, $\rho = d - dim(L_{max})$. By maximal unconstrained subspace we mean the unique maximal element $L_{max} \in \mathscr{L}$ for which $L \subseteq L_{max}$, for all $L \in \mathscr{L}$. □

From the support rank definition (see Lemma 3.8 in [46]) we have that Helly's dimension can be upper bounded by the support rank instead of the dimension d of the problem, which is a more conservative upper bound [21]. Keeping this relation in mind, our main goal is to obtain a bound for the support rank of the random constraint of problem P', thus improving the robustness certificates of its optimal solution. The following proposition aims at finding an upper bound for the support rank ρ of the random constraint (5.2), that is independent from the number of agents involved in the optimization program.

Proposition 5.1 *Under Assumptions 5.1 and 5.2, the support rank ρ of the random constraint (5.2) in P', has an agent independent upper bound, and in particular, $\rho \leq n+1$.*
□

Proof: The dimension of the problem under study is $d = nN+1$, due to the presence of the epigraphic variable. Let $\mathscr{L}$ be the collection of all linear subspaces in $\mathbb{R}^{nN+1}$. We aim at finding the dimension of the subspace that remains unconstrained for a scenario program subject to the random constraint $h(x,\gamma,\delta) \leq 0$ for any uncertain realisation $\delta \in \Delta$ and any decision vector $(x,\gamma) \in X \times \mathbb{R}$. We first define the collection of linear subspaces that are contained in all the sets $F(x,\gamma,\delta)$:

$$\mathscr{U} = \bigcap_{\delta \in \Delta} \bigcap_{(x,\gamma) \in \mathbb{R}^{nN+1}} \{L \in \mathscr{L} : L \subset F(x,\gamma,\delta)\}, \text{where}$$

$$F(x,\gamma,\delta) = \{(\xi,\xi') \in \mathbb{R}^{nN+1} : h(x+\xi,\gamma+\xi',\delta) = h(x,\gamma,\delta)\}.$$

In our case, $h(x+\xi,\gamma+\xi',\delta) = h(x,\gamma,\delta)$ yields:

$$\begin{aligned}
&\sigma(x+\xi)^T(A(\delta)\sigma(x+\xi)+b(\delta)) - (\gamma+\xi') = \sigma(x)^T(A(\delta)\sigma(x)+b(\delta)) - \gamma,\\
\iff\ &\sigma^T(x)A(\delta)\sigma(\xi) + \sigma^T(\xi)A(\delta)\sigma(x) + \sigma^T(\xi)A(\delta)\sigma(\xi) + \sigma^T(\xi)b(\delta) - \xi' = 0,\\
\iff\ &\sigma^T(\xi)A^T(\delta)\sigma(x) + \sigma^T(\xi)(A(\delta)\sigma(x)+b(\delta)) + \sigma^T(\xi)A(\delta)\sigma(\xi) - \xi' = 0,
\end{aligned}$$

where the first equivalence stems from the fact that $\sigma(x+\xi)$ is linear with respect to its arguments, and the second one after some algebraic rearrangement. Note that each of the terms above is scalar, which means that it is equal to its transpose for all $x \in X$ and $\delta \in \Delta$, while by Assumption 5.1, $A^T(\delta) = A(\delta)$ for any $\delta \in \Delta$. As such,

$$\sigma^T(\xi)(2A(\delta)\sigma(x)+b(\delta)) + \sigma^T(\xi)A(\delta)\sigma(\xi) - \xi' = 0. \tag{5.5}$$

Using the equalities $\sigma^T(\xi)(2A(\delta)\sigma(x)+b(\delta)) = (\mathbb{1}_{1\times N} \otimes (2A(\delta)\sigma(x)+b(\delta))^T)\xi$ and $\sigma^T(\xi)A(\delta)\sigma(\xi) = \xi^T(\mathbb{1}_{N\times N} \otimes A(\delta))\xi$, where $\mathbb{1}_{1\times N}$ denotes a row vector with all elements being equal to one and $\otimes$ denotes the Kronecker product, (5.5) can be written in the following form:

$$(\mathbb{1}_{1\times N} \otimes (2\sigma^T(x)A(\delta) + b^T(\delta)))\xi + \xi^T(\mathbb{1}_{N\times N} \otimes A(\delta))\xi - \xi' = 0, \tag{5.6}$$

Let $\tilde{C} : X \times \Delta \to \mathbb{R}^{nN}$, $\tilde{A} : \Delta \to \mathbb{R}^{nN\times nN}$, where $\tilde{C}(x,\delta) = \mathbb{1}_{1\times N} \otimes (2\sigma^T(x)A(\delta) + b^T(\delta))$ and $\tilde{A}(\delta) = \mathbb{1}_{N\times N} \otimes A(\delta)$, respectively. Then, equation (5.6) can be written as:

$$\begin{aligned}
&\tilde{C}(x,\delta)\xi + \xi^T\tilde{A}(\delta)\xi - \xi' = 0,\\
\iff\ &\begin{pmatrix}\tilde{C}(x,\delta) & -1\end{pmatrix}\begin{pmatrix}\xi\\ \xi'\end{pmatrix} + \begin{pmatrix}\xi & \xi'\end{pmatrix}\begin{pmatrix}\tilde{A}(\delta) & 0_{nN\times 1}\\ 0_{1\times nN} & 0\end{pmatrix}\begin{pmatrix}\xi\\ \xi'\end{pmatrix} = 0,\\
\iff\ &V(x,\delta)w + w^TP(\delta)w = 0,
\end{aligned}$$

where, $V(x,\delta) = \begin{pmatrix}\tilde{C}(x,\delta) & -1\end{pmatrix}$, $P(\delta) = \begin{pmatrix}\tilde{A}(\delta) & 0_{nN\times 1}\\ 0_{1\times nN} & 0\end{pmatrix}$ and $w = \begin{pmatrix}\xi\\ \xi'\end{pmatrix}$. We need to find an unconstrained linear subspace that is a subset of $F(x,\gamma,\delta)$. We define

$$L(x,\gamma,\delta) = \{w \in \mathbb{R}^{nN+1} : \begin{pmatrix}P(\delta)\\ V(x,\delta)\end{pmatrix} w = 0\}.$$

We can easily see that $L(x,\gamma,\delta) \subset F(x,\gamma,\delta)$. As such, the random constraint $h(x,\gamma,\delta) \leq 0$ cannot constrain any of the dimensions of $L(x,\gamma,\delta)$ (also denoted as L for simplicity). Let $Q(x,\gamma,\delta) = \begin{pmatrix} P(\delta) \\ V(x,\delta) \end{pmatrix}$. Then $L(x,\gamma,\delta) = \text{nullspace}(Q(x,\gamma,\delta))$ and from nullity-rank theorem [146] we have that $\dim(L(x,\gamma,\delta)) = nN + 1 - \text{rank}(Q(x,\gamma,\delta))$. Since $\text{rank}(P(\delta)) = n$ and $\text{rank}(V(x,\delta)) = 1$, this means that $\text{rank}(Q(x,\gamma,\delta)) = n+1$, which implies that $\dim(L(x,\gamma,\delta)) = nN + 1 - (n+1)$. Notice that the unconstrained subspace that we chose may not be the maximal one. This means that the support dimension is $\rho = nN + 1 - \dim(L_{max}) \leq nN + 1 - \dim(L) = n+1$, thus concluding the proof. ■

An immediate consequence of Proposition 5.1 when combined with Theorem 4.1 of [46] is the following theorem.

Theorem 5.1 *Let (x^*,γ^*) denote the optimal solution of the scenario program $P^{'}_{SC}$. Under Assumptions 5.1 and 5.2 we have that*

$$\mathbb{P}^K\{\delta_K \in \Delta^K : \mathbb{P}\{\delta \in \Delta : h(x^*,\gamma^*,\delta) > 0\} > \varepsilon\} \leq \beta, \tag{5.7}$$

$$\textit{where } \beta = \sum_{j=0}^{n} \binom{K}{j} \varepsilon^j (1-\varepsilon)^{K-j}. \tag{5.8}$$

The bound obtained from Theorem 5.1 constitutes a major improvement for this class of problems, since, irrespective of the number of agents N, the same number of samples M is required to provide identical robustness certificates given a local decision vector of size n. The proof of this theorem, is a direct application of the scenario approach theory (see Theorem 2.4 in [21], and Theorem 4.1 in [46], where the number of support constraints is replaced by the obtained bound for the support rank, namely, $n+1$). Note that in the absence of Proposition 5.1, a direct application of the scenario approach theory [21] to the problem under consideration would still result in (5.7), however, (5.8) would be replaced by

$$\beta = \sum_{j=0}^{nN} \binom{K}{j} \varepsilon^j (1-\varepsilon)^{K-j}, \tag{5.9}$$

where the dependence of the guarantees on the number of agents N is apparent.

Corollary 5.1 establishes a link between Theorem 5.1 and the initial program under study P by providing probabilistic performance guarantees for the optimal solution.

Specifically, it quantifies, in an *a priori* fashion, the probability that the cost that corresponds to the optimal value x^* of P'_{SC} will deteriorate, when a new sample $\delta \in \Delta$ is encountered. To formalise this, with a slight abuse of notation let $J(x) = J(x, \delta_K)$ be the cost function of the corresponding scenario program of program P and $J^+(x) = J(x, \delta_K \cup \{\delta\})$ the cost defined over $K+1$ scenarios by taking into account the new sample δ.

Corollary 5.1 *Under Assumptions 5.1 and 5.2 we have that*

$$\mathbb{P}^K\{\delta_K \in \Delta^K : \mathbb{P}\{\delta \in \Delta : J^+(x^*) > J(x^*)\} > \varepsilon\} \leq \beta, \tag{5.10}$$

$$\text{where } \beta = \sum_{j=0}^{n} \binom{K}{j} \varepsilon^j (1-\varepsilon)^{K-j} \tag{5.11}$$

Proof Let (x^*, γ^*) be the optimal solution of program $P^{'}$, which implies that $\gamma^* = \max\limits_{k=1,\ldots,K} g(x^*, \delta^{(k)})$. As such,

$$\begin{aligned}
&\mathbb{P}\{\delta \in \Delta : h(x^*, \gamma^*, \delta) > 0\} = \mathbb{P}\{\delta \in \Delta : g(x^*, \delta) > \gamma^*\} = \\
&\mathbb{P}\{\delta \in \Delta : g(x^*, \delta) > \max_{k=1,\ldots,K} g(x^*, \delta^{(k)})\} = \\
&\mathbb{P}\{\delta \in \Delta : \max\{g(x^*, \delta), \max_{k=1,\ldots,K} g(x^*, \delta^{(k)})\} > \max_{k=1,\ldots,K} g(x^*, \delta^{(k)})\} = \\
&\mathbb{P}\{\delta \in \Delta : J^+(x^*) > J(x^*)\},
\end{aligned} \tag{5.12}$$

where the second equality follows from the fact that $\gamma^* = \max\limits_{k=1,\ldots,K} g(x^*, \delta^{(k)})$, and the last one from the definitions of J and J^+. Direct substitution of (5.12) in (5.7) of Theorem 5.1 concludes then the proof. □.

5.4 Numerical Study

In this set-up the cost to be minimized is influenced by the electricity price, which in turn is considered to be a random variable affected by uncertainty. Uncertainty here refers to price volatility. All electric vehicles cooperate with each other choosing their charging schedules so as to minimize the total uncertain electricity cost, while satisfying their own deterministic constraints. To this end, we consider the following uncertain

electric vehicle charging problem

$$P_{EV}: \min_{x\in\mathbb{R}^{nN}} f(x)+\max_{\delta\in\Delta} g(x,\delta),$$
$$\text{subject to } x_i\in[\underline{x}_i,\overline{x}_i], \sum_{t=1}^{n} x_i^{(t)} \geq E_i \text{ for all } i\in\mathcal{N}, \tag{5.13}$$

where $f(x)=\sum_{i\in\mathcal{N}} f_i(x_i,x_{-i})=\sigma(x)^T p_0(\sigma(x))$ is the deterministic part of the electricity cost that depends on a nominal electricity price $p_0(\sigma(x))=A_0\sigma(x)+b_0$ that is, in turn, a function of the aggregate consumption of the vehicles. $g(x,\delta)=\sum_{i\in\mathcal{N}} g_i(x_i,x_{-i},\delta)=\sigma(x)^T p(\sigma(x),\delta)$ constitutes the uncertain part of the electricity cost, where the price $p(\sigma(x),\delta)=A(\delta)\sigma(x)+b(\delta)$ is additionally affected by the uncertain parameter δ extracted from the support set Δ according to a probability distribution $\mathbb{P}$, where Δ and $\mathbb{P}$ are considered unknown. The elements of $A_0\in\mathbb{R}^{n\times n}$ and $b_0\in\mathbb{R}^n$ are deterministic with A_0 being a symmetric positive semi-definite matrix, while the uncertain mappings $A:\Delta\to\mathbb{R}^{n\times n}$ and $b:\Delta\to\mathbb{R}^n$ are defined as in Section 5.3. The values of A and b can either be extracted from historical data or from a synthetic dataset obtained by some prediction model. In this particular example the electricity price is assumed to be affine with respect to the aggregate demand. Furthermore, we assume that A, b follow certain distributions. However, we stress that one can obtain values of A and b based on real data that approximate the behaviour of the electricity price by rescaling weekday demand profiles. For example, matrix A is obtained in this fashion in [48] and Section 3.4.2 of this thesis, while b is assumed to be 0 for simplicity. The vectors $\underline{x}_i,\overline{x}_i\in\mathbb{R}^n$ constitute the lower and upper bound of the charging rate of vehicle $i\in\mathcal{N}$, respectively, while $E_i\in\mathbb{R}$ is the final energy to be achieved by each vehicle $i\in\mathcal{N}$ by the end of the charging cycle.

Following the same lines as in Section 5.3, we apply an epigraphic reformulation and use samples for Δ to obtain the following scenario program

$$P_{EV}^{sc}: \min_{(x,\gamma)\in\mathbb{R}^{nN+1}} f(x)+\gamma, \tag{5.14}$$
$$\text{subject to } x_i\in[\underline{x}_i,\overline{x}_i], \sum_{t=1}^{n} x_i^{(t)} \geq E_i, \text{ for all } i\in\mathcal{N},$$
$$g(x,\delta^{(k)})\leq\gamma, \text{ for all } i\in\mathcal{N}.$$

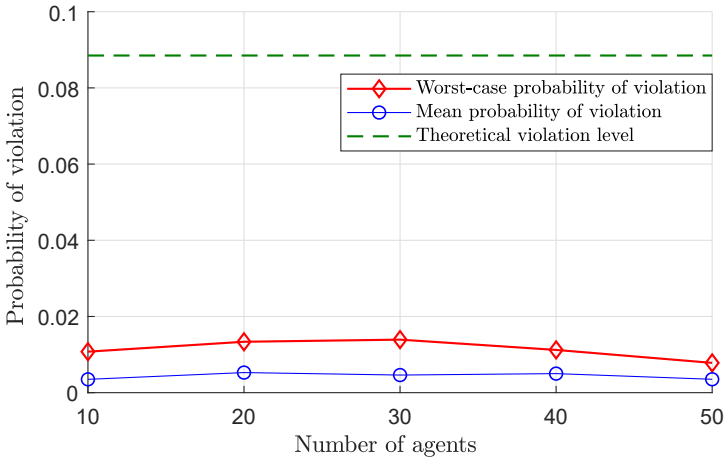

Figure 5.1: Mean and worst-case empirical probability of violation of the optimal solution with respect to the number of agents versus the theoretical violation level $\varepsilon = 0.0885$. The number of samples used is $K = 500$ and $\beta = 10^{-6}$. By drawing a different multi-sample for each choice of the number of agents $N = \{10, 20, 30, 40, 50\}$, we solve the corresponding scenario program for a fixed number of time slots $n = 12$. We then repeat this process 20 times (note that the multi-sample used for each repetition is also different) and compute the empirical probability of violation of the obtained optimal solutions, using $K_{test} = 100000$ test samples.

In our set-up $A(\delta) \in \mathbb{R}^{n\times n}$ is assumed to be a diagonal matrix with non-negative diagonal elements for any uncertain realization $\delta \in \Delta$. The diagonal elements of $A(\delta)$ and the elements of $b(\delta)$ are extracted according to uniform distributions. For each agent $i \in \mathcal{N}$ the upper bound $\bar{x}_i$ takes a random value in the set $[6, 15]$ kW, the lower bound $\underline{x}_i$ is set to 2 kW and the final energy to be achieved by the end of the charging cycle is appropriately chosen to be feasible, considering the number of timesteps n and the upper bound of the power rate of each agent. $A_0 \in \mathbb{R}^{n\times n}$ is assumed to be a diagonal matrix, whose diagonal entries are all set to 0.01 and b_0 is derived by rescaling a winter weekday demand profile in the UK.

Note that our results can be used alongside any optimization algorithm irrespective of its nature, i.e., centralised, decentralised or distributed; here we solved the problem in a centralised fashion. The number of samples we use for each problem is $K = 500$. By fixing $\beta = 10^{-6}$ and using the bound

$$\varepsilon = \frac{2}{K}(\ln\frac{1}{\beta} + n\ln 2), \tag{5.15}$$

which is a sufficient condition (see [145, (p.42)]) for satisfaction of (5.8), we obtain the theoretical violation level $\varepsilon = 0.0885$. Note that the dimension we use to provide

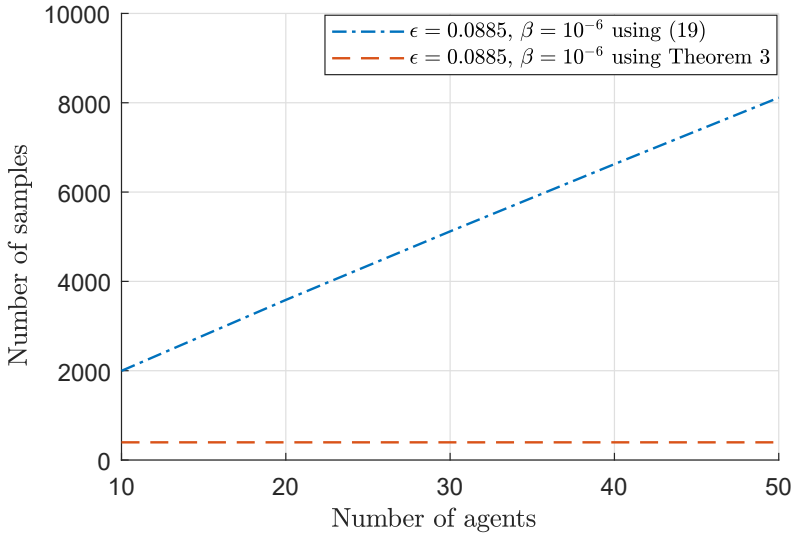

Figure 5.2: The number of samples required with respect to the number of agents $N = \{10, \ldots, 50\}$ using the results of Theorem 5.1 versus the one that would have been obtained if (5.8) is used instead. We consider a charging cycle of duration $n = 12$. The red line corresponds to Theorem 5.1, while the blue line corresponds to (5.9).

probabilistic guarantees for the optimal solution is set, in accordance to Theorem 5.1 to $n+1$ instead of $nN+1$, which circumvents the computational issues related to the rapid surge in dimension due to the multiplication of the number of agents with the number of time slots.

By drawing a different multi-sample for each choice of the number of agents $N \in \{10, 20, 30, 40, 50\}$ we solve the corresponding scenario program for a fixed number of time slots $n = 12$. We then repeat this process 20 times (note that the multi-sample used for each repetition is also different) and compute the empirical probability of violation of the obtained optimal solutions, using $K_{test} = 100000$ test samples each time. The mean and worst-case empirical probability of violation is depicted in Figure (5.1) in comparison with the theoretical violation level ε. The empirical values are always below the theoretical level of violation, which is constant as a function of the number of agents due to the agent independent nature of our Theorem 5.1. In addition, the trend in Figure 5.1 shows, as expected by Theorem 5.1, that the number of agents does not affect the empirical probability of violation.

This result highlights the fact that, for fixed number of time periods n, the number of samples K required to provide identical probabilistic guarantees, as the size of the

fleet of electric vehicles increases, remains constant. This is illustrated in Figure 5.2, where we show the number of samples required (for $\varepsilon = 0.0885$, $\beta = 10^{-6}$ and $n = 12$) using the results of Theorem 5.1 versus the number of samples needed to provide the same robustness certificates using the classic results in scenario approach for a different number of agents $N = 10, \ldots, 50$. The red dashed line corresponds to Theorem 5.1 and shows the agent independent nature of our guarantees, while the blue dash-dotted line corresponds to the conservative agent dependent result of (5.8).

5.5 Concluding remarks

In this chapter, we focused on a different class of multi-agent programs that involved an uncertain aggregative term and deterministic constraints. For such problems we provided agent independent probabilistic guarantees for the optimal solution in an *a priori* fashion. In future work we wish to extend our results to provide agent independent probabilistic guarantees for Nash equilibria of a certain class of games affected by uncertainty, as well as to investigate a more general case, where each agent participating in the game has a different uncertain function in their cost.

6

Probabilistic stability guarantees for cooperative games

In this chapter we consider multi-agent coalitional games with uncertain value functions for which we establish distribution-free guarantees on the probability of allocation stability, i.e., agents do not have incentives to defect from the grand coalition to form subcoalitions for unseen realizations of the uncertain parameter. In case the set of stable allocations, the so called core of the game, is empty, we propose a randomized relaxation of the core. We then show that those allocations that belong to this relaxed set can be accompanied by stability guarantees in a probably approximately correct fashion. Finally, numerical experiments corroborate our theoretical findings.

6.1 Introduction

As seen in Chapters 3 and 4, agents participating in a non-cooperative game are modelled as selfish entities that optimize their own payoff function. In some cases, however, due to their limited ability to increase their own utility when working on their own, agents have an incentive to form coalitions aiming at receiving a higher individual payoff. This gives rise to what we call a coalitional game setup (see [12] and references therein). Since each agent is interested in maximizing their own welfare, the problem of allocating the total value of the coalition in order to guarantee that none of the agents has an incentive to defect

it, is key from both a collective and an individual point of view. This concept is known as *stability* of the agents' allocations and it is related to the agents' coalitional values.

In real-world applications, the values of each coalition are typically affected by uncertainty. This can be due to various reasons. The most predominant ones refer to the effects of exogenous factors acting on the agents' network, or uncertainty inherent in the coalition formation process. This chapter focuses on the former case, where uncertainty in the environment results in changes in the values of agents' coalitions. Encoding uncertainty in the value functions of a coalitional game was initially explored in the seminal works [38, 39]. Therein, the authors focus on the extension of well-known solution concepts to provide stability-wise robust allocations. In [147] it is shown that for a certain class of stochastic games, properties such as the non-emptiness of the core of a deterministic game continue to hold for their stochastic counterpart. References [41] and [42] tackled uncertainty in the values of the coalitions by leveraging Bayesian learning methods, while [43] investigated which stability solution concepts maximize the probability that the allocations will be stable in an *a posteriori* fashion, that is after the samples of the uncertainty have been revealed. Finally, [44] and [45] addressed uncertainty by studying the dynamics of repeated stochastic coalitional games.

Unlike the aforementioned works, we construct a data-driven methodology that allows the provision of distribution-free guarantees on the stability of allocations in a probably approximately correct (PAC) fashion, i.e., ensure that, with high confidence, the agents' probability to defect from the grand coalition and form subcoalitions is bounded by a prespecified threshold. Specifically, this bound depends on the amount of available data, the confidence parameter and which samples from the data are crucial to obtain this particular allocation or set of allocations. Connecting PAC learning with uncertain coalitional games has also been considered in [148]. Therein, the authors use a sampling-based approach to learn the value function using a randomly generated subset out of the total number of potential coalitions. The spirit of [149] is similar, using Vapnik-Chervonenkis (VC) theory [101] to learn the winning coalitions for the class of the so-called *simple* games.

Our main contributions in comparison with the existing literature are as follows. Focusing on coalitional games with uncertain value functions, we leverage recent results from the scenario approach [24–26] to provide distribution-free PAC-type stability guarantees for agents' allocations. With respect to [148], our results do not suffer from the conservatism associated with VC-theoretic results. Furthermore, [148] focuses on a complementary problem where only a randomized subset of possible coalitions is considered; in contrast, we include all of them but account in a randomized manner for uncertainty in the value functions.

Our analysis is based on mild assumptions, as we assume no prior knowledge of the sample space or the probability distribution of the uncertainty. Note that using robust versions of the core that take into account all possible uncertainty realizations (see for example [17]), as our main solution concept for stability would pose significant challenges in such a general setting. To start with, the evaluation of the robust core under our set-up constitutes a challenging task as the uncertainty support set Δ is in general unknown in the definition of the robust core (Definition 6.1); even in cases where this is known, finding the maximum coalitional value could be computationally unviable when Δ is a continuous set, unless we impose additional structural assumptions on its geometry. Furthermore, the robust versions of the core do not provide flexibility, as [17] considers the highest value realizations that can possibly occur for each coalition. This unveils the conservatism that accompanies such worst-case paradigms, as it is not possible to account for cases where allocations are allowed to be unstable with a small probability as common in a more general chance-constrained framework. These challenges are circumvented in our work by adopting a data-driven paradigm with the scenario approach as its backbone.

Combining uncertain cooperative game theory with the scenario approach requires significant adaptation of the game-theoretic setting. Furthermore, a probabilistic stability analysis of solutions to this more general class of games under the lenses of the scenario approach is not straightforward as we will see in the subsequent developments. Our first contribution hence involves the introduction of a novel stability concept, i.e., the *scenario core*. Our approach allows then to provide stability guarantees collectively for all allocations that belong to this more relaxed set. A significant impediment on

allocation stability that spans the literature of cooperative game theory is the fact that the core of a game, i.e., the set of stable allocations, can be empty. In such cases, a relaxed deterministic version of the core is usually proposed to provide stable solutions by charging the formation of alternative sub-coalitions. This set is known in the literature as the ε-*core* of the game, where ε is a parameter that determines the level of relaxation. In certain cases, the ε-core can be viewed as a set of nearly stable solutions.

Considering a more general framework, where the amount that the agents will gain through the formation of coalitions is uncertain, our approach towards the significant challenge of a possible emptiness of the core is different. Specifically, we introduce the notion of the scenario ε-core that is now not defined beforehand, but shaped from the drawn samples. We formulate the problem of finding an allocation in a randomized ε-core of a game and we provide probabilistic stability guarantees for this allocation leveraging recent results on the scenario approach with relaxation [25, 26]. Our formulation guarantees the non-emptiness of the scenario ε-core in all cases. As a byproduct of the proposed approach, the emptiness of the original core can be easily revealed along with the part of the data that is responsible for the emptiness of the scenario core. For notational reasons, we will refer to the relaxed core as the ζ-core in the remainder of the document.

The rest of the chapter is organized as follows. Section 6.2 formulates the problem under study establishing a data-driven framework for uncertain coalitional games. Section 6.3 provides then probabilistic stability guarantees for allocations that belong to the scenario core and, in case the scenario core is empty, for those allocations inside the randomized ζ-core. Section 6.4 corroborates our theoretical findings by means of a numerical example, while Section 6.5 concludes the chapter.

6.2 A data-driven approach to uncertain coalitional games

6.2.1 Problem formulation and stability of allocations

Let $\mathcal{N} = \{1, \ldots, N\}$ be the index set of N agents and consider any subset $S \subseteq \mathcal{N}$ as a *coalition*. In the proposed setting, selfish agents have the incentive to form coalitions in order to achieve a higher individual payoff, or as the means to perform a certain

task. The value of a coalition $S \subseteq \mathcal{N}$ is represented by the so-called *value function*, denoted by $u(S),\ S \subseteq \mathcal{N}$.

We consider in this chapter a model of a coalitional game where the underlying network is affected by some exogenous uncertainty δ that takes values in a set Δ according to a probability distribution $\mathbb{P}$. Note that our model considers Δ and $\mathbb{P}$ to be fixed but possibly unknown, thus managing to capture several external factors that could potentially lead to changes of the coalitions' values. In this setting, the value functions are extended to their uncertain counterpart $u : 2^{\mathcal{N}} \times \Delta \to \mathbb{R}$ which, given a coalition $S \subset \mathcal{N}$ and an uncertainty realization $\delta \in \Delta$, returns a scalar value. The value function realizations for any subcoalition could be either available from historical data or extracted from some synthetic dataset obtained by some prediction model. In our setting we consider the value function of the grand coalition to be deterministic. An *uncertain game* is then defined as the tuple $G_\Delta = (\mathcal{N}, u, \Delta, \mathbb{P})$.

A vector $x = (x_i)_{i\in\mathcal{N}} \in \mathbb{R}^N$, where x_i is the payoff received by agent i is called an *allocation*. An allocation is *strictly feasible for a coalition S* if $\sum_{i\in S} x_i < u(S)$. Feasibility of a coalition implies that agents have an incentive to form this coalition. The so-called grand coalition $\mathcal{N}$ is called *efficient* if $\sum_{i\in\mathcal{N}} x_i = u(\mathcal{N})$. In this work we are interested in finding efficient allocations (in the grand coalition sense) that are not strictly feasible by any other coalition. Such allocations are called *stable*, as there are no incentives for agents to form coalitions different from the grand one. The set of all stable allocations is called the *core* of the game. However, in our setting, the value function is considered to be uncertain, and hence, the standard definition of the core [12] is not sufficient to capture the desired stability properties. To this end, we extend the notion of the core to account for the presence of uncertainty, as in the following definition.

Definition 6.1 (Robust core) *The* robust core $C(G_\Delta)$ *of an uncertain game* G_Δ *is given by* $C(G_\Delta) = \{x \in \mathbb{R}^N : \sum_{i\in\mathcal{N}} x_i = u(\mathcal{N}),\ \sum_{i\in S} x_i \geq \max_{\delta\in\Delta} u(S,\delta)$ *for all* $S \subset \mathcal{N}\}$. □

Definition 6.1 shares a similar spirit with Definition 6 in [17]. The robust core $C(G_\Delta)$ provides the coalitional game under study with a measure of robust stability in the sense that, for any allocation $x \in C(G_\Delta)$, agents have no incentive to defect from the grand coalition to create sub-coalitions.

6.2.2 A PAC-learning approach to allocation stability

Unfortunately, computing explicitly the robust core is hard, as we assume no knowledge on the uncertainty support Δ and the underlying probability distribution $\mathbb{P}$. To circumvent this challenge, we adopt a data-driven methodology and approximate the robust core by drawing a finite number of K independent and identically distributed (i.i.d.) samples $\delta_K := (\delta^{(1)}, \ldots, \delta^{(K)}) \in \Delta^K$, where Δ^K denotes the cartesian product consisting of K copies of Δ. We refer to vectors δ_K as multi-samples. This gives rise to the scenario game $G_K = (\mathcal{N}, u, \delta_K)$. The core of G_K, referred to as the *scenario core*, is then defined as

$$C(G_K) = \{x \in \mathbb{R}^N : \textstyle\sum_{i\in\mathcal{N}} x_i = u(\mathcal{N}), \sum_{i\in S} x_i \geq \max_{k=1,\ldots,K} u(S, \delta^{(k)}) \text{ for all } S \subset \mathcal{N}\}.$$

Furthermore, we impose the following assumption:

Assumption 6.1 *For any multi-sample $\delta_K \in \Delta^K$, the scenario core $C(G_K)$ is non-empty.*
□

Assumption 6.1 implies the existence of stable allocations. We will investigate how to waive this assumption for cases where the presence of uncertainty results in an empty scenario core in Section 6.3. On the basis of available data, we then wish to provide guarantees on the probability that payoff allocations $x \in \mathbb{R}^N$ inside the scenario core will remain stable even for future, yet unseen, uncertainty realizations. Borrowing tools from the scenario approach [138, 150], we define two probabilities of instability. The first one denotes the probability that a certain allocation will become unstable for a new unseen realization of the uncertainty. The second extends this notion to the probability of instability of the entire scenario core.

Definition 6.2 (Probability of instability)

1. *Let $V : \mathbb{R}^N \to [0,1]$. For any $x \in \mathbb{R}^N$, we call*

$$V(x) = \mathbb{P}\{\delta \in \Delta : \textstyle\sum_{i\in S} x_i < u(S, \delta) \text{ for some } S \subset \mathcal{N}\}$$

probability of allocation instability.

2. *Let* $\mathbb{V} : 2^{\mathbb{R}^N} \to [0,1]$. *We call*

$$\mathbb{V}(C(G_K)) = \mathbb{P}\{\delta \in \Delta : \exists\, x \in C(G_K) : \textstyle\sum_{i\in S} x_i < u(S,\delta), \\ \text{for some } S \subset \mathcal{N}\}$$

probability of core instability. □

The probability of core instability thus denotes the probability that there exist δ and S with value function $u(S,\delta)$, parameterized by the uncertainty δ, such that at least one of the stable allocations in the scenario core will become unstable, i.e., the agents will defect from the grand coalition to form S.

By leveraging available data, we then aim at bounding with high confidence the probability of core instability, i.e., in a PAC fashion. To achieve this we use two key concepts from statistical learning theory, namely the *algorithm* and the *compression set* [121]. These two concepts are introduced in Definition 2.17 of Chapter 2. The definition of the algorithm is adapted in this chapter as follows.

Definition 6.3 (Algorithm) *A mapping* $A : \Delta^K \to 2^{\mathbb{R}^N}$ *that takes as input a multi-sample* $\delta_K \in \Delta^K$ *and returns the scenario core of game* $C(G_K)$ *is called an* algorithm. □

Note that some samples are more important than others in the decision making procedure. In fact, only a subset of δ_K may be sufficient to produce the same scenario core. This subset, known as a compression set (see Definition 2.17) dictates the quality of the probabilistic stability guarantees that we can provide and is defined as follows. The notion of compression is adapted in the set-up of coalitional games as follows.

Definition 6.4 (Compression set) *With* $\mathbb{P}^K$*-probability one with respect to the choice of* δ_K*, a subset* $I \subseteq \{\delta^{(1)}, \ldots, \delta^{(K)}\}$ *is a* compression set *of A if the scenario core A returns does not change when only the subset of samples I is fed as input to A. In mathematical terms, I forms a compression if* $A(\delta_I) = A(\delta_K)$*, where* δ_I *is a vector whose elements are the samples included in I.* □

For the convenience of the reader we repeat that a compression set of A with minimal cardinality is hence called *minimal* compression set. A procedure that enumerates such a set is called a *compression function* (see Definition 2.17).

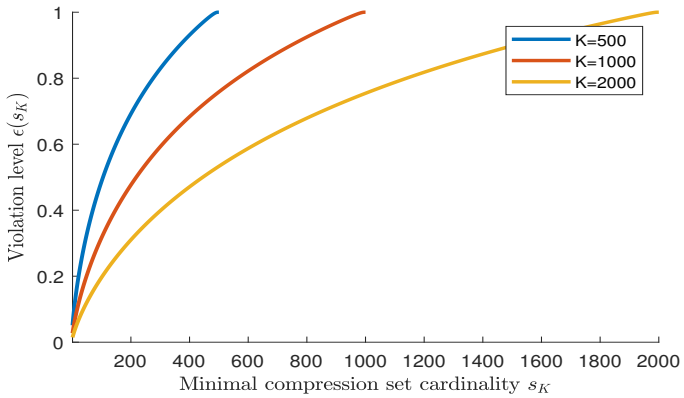

Figure 6.1: The violation level $\varepsilon(s_K)$ as a function of the cardinality of the minimal compression set for $\beta = 10^{-6}$ and three different multi-sample sizes, i.e., $K = 500, 1000, 2000$.

6.3 Probabilistic stability guarantees of allocations

6.3.1 Stability guarantees for a non-empty scenario core

The following theorem provides collective guarantees on the stability of allocations in the scenario core against yet unseen value function realizations.

Theorem 6.1 (*A posteriori* collective stability guarantees) *Consider Assumption 6.1 and algorithm A as in Definition 6.3 along with its compression function. Fix a confidence parameter* $\beta \in (0,1)$ *and define the violation level* $\varepsilon : \{0,\dots,K\} \to [0,1]$ *as a function such that*

$$\varepsilon(K) = 1 \text{ and } \sum_{s=0}^{K-1} \binom{K}{s} (1-\varepsilon(s))^{K-s} = \beta. \tag{6.1}$$

We then have that

$$\mathbb{P}^K \left\{ \delta_K \in \Delta^K : \mathbb{V}(C(G_K)) \leq \varepsilon(s_K) \right\} \geq 1 - \beta \tag{6.2}$$

where $\mathbb{P}^K = \prod_{k=1}^K \mathbb{P}$ *and* s_K *is the cardinality of the minimal compression set.* □

Proof: The proof follows by re-adapting [151, Th. 6]. ■

Roughly speaking, Theorem 6.1 states that, with confidence at least $1-\beta$, the probability that a new, yet unseen, uncertainty sample will make an allocation in the scenario core unstable is bounded by ε, a function of s_K. A simple choice of $\varepsilon(s)$ can be obtained by splitting β evenly among the K terms in the sum of (6.2), and solving with respect to

$\varepsilon(s)$. This results in [24, Eq. (7)], and is illustrated numerically in Fig. 6.1. Note that the nature of Theorem 6.1 is *a posteriori*, i.e., we can claim the probabilistic bound in (6.2) only once the samples have been revealed. In fact, this is required to quantify both the solution and the cardinality s_K of the minimal compression set.

Algorithm 2 plays the role of the compression function used to compute s_K. Note that the problem at hand has a similar structure to the problem in [152]. Algorithm 2 involves solving a linear optimization program (specifically, a feasibility problem) for each coalition $S' \subset \mathcal{N}$, where we enforce in step 3 the constraint $\sum_{i \in S'} x_i = \max_{k=1,\dots,K} u(S', \delta^{(k)})$, while all other constraints remain unchanged. If the problem is feasible, the sample that maximized the right-hand side of this equality constraint is included in the compression set I, as feasibility in this case implies that this sample belongs to the compression set of the algorithm A that forms the scenario core $C(G_K)$. Using the union in step 5 ensures that no sample is counted more than once when we compute the cardinality s_K of the compression set I in step 8.

A fundamental challenge in cooperative game theory is that the number of possible coalitions is exponential to the number of agents N. This renders Algorithm 2 not scalable with the number of agents as one has to account for a total of 2^N possible coalitions in order to find the compression. One way to address this issue is to allow the formation of only a subset of coalitions. This, however, restricts the freedom of each agent to choose beneficial coalitions. A more natural way to approach this challenge is by reformulating the original problem considering only a subset of coalitions obtained through sampling available from historical data or obtained by a synthetic model. To avoid the significant challenge that coalitions are not necessarily i.i.d, one can instead consider as a sample realization the coalitions that the population of agents forms for a given instance of the game. However, this requires a non-trivial readaptation of the problem at hand and is thus left for future work.

Instead of following the *a posteriori* methodology of Theorem 6.1, one can provide collective guarantees on the stability of the scenario core also in an *a priori* fashion, whenever the number of samples is greater than the number of possible subcoalitions. As shown in Lemma 6.1, the *a priori* bound on the cardinality of the compression

Algorithm 2 Compression function of algorithm A

1: **initialization:** $I = \emptyset$
2: **input:** δ_K
3: **for all** $S' \subset \mathcal{N}$

$$x^*_{S'} = \begin{cases} \underset{x \geq 0}{\text{argmin}} & 0 \\ \text{subject to} & \sum_{i \in \mathcal{N}} x_i = u(\mathcal{N}), \\ & \sum_{i \in S'} x_i = \max_{k=1,\ldots,K} u(S', \delta^{(k)}), \\ & \sum_{i \in S} x_i \geq \max_{k=1,\ldots,K} u(S, \delta^{(k)}) \; \forall \, S \subset \mathcal{N}. \end{cases} \tag{6.3}$$

4: **if** $x^*_{S'} \neq \emptyset$
5: $I \leftarrow I \cup \arg\max_{k=1,\ldots,K} u(S', \delta^{(k)})$
6: **endif**
7: **endfor**
8: $s_K = |I|$
9: **output:** s_K

set coincides with the number of possible coalitions that can be formed; if constraints on the latter are in place, tighter bounds could be provided. The following lemma establishes this observation.

Lemma 6.1 *Assume that $K > |2^{\mathcal{N}}|$. The cardinality of the minimal compression set is bounded by the number of possible subcoalitions.* □

Proof: By the definition of $C(G_K)$, for each $S \subset \mathcal{N}$, the largest value on the right-hand side of the inequality $\sum_{i \in S} x_i \geq \max_{k=1,\ldots,K} u(S, \delta^{(k)})$ is attained at a given sample (assuming no degenerate instances where multiple maximizers exist). Assuming $K > |2^{\mathcal{N}}|$, in a worst-case setting, for each $S \subset \mathcal{N}$ the maximizing sample may be different. Therefore, for each $S \subset \mathcal{N}$ each inequality $\sum_{i \in S} x_i \geq \max_{k=1,\ldots,K} u(S, \delta^{(k)})$ is a randomized facet of the non-redundant polytope $C(G_K)$. Each of these facets corresponds to exactly one coalition, irrespective of the number of samples. This observation concludes the proof. ■

Under Lemma 6.1, by [121, Theorem 2], for a fixed confidence level $\beta \in (0,1)$, we can claim that

$$\mathbb{P}^K \left\{ \delta_K \in \Delta^K : \mathbb{V}(C(G_K)) > \varepsilon \right\} \leq \beta,$$

where ε can be chosen *a priori* such that $\binom{K}{d}(1-\varepsilon)^{K-d} = \beta$, with d being the number of possible subcoalitions. However, note that for games with a high number of agents, the *a posteriori* result of Theorem 6.1 might be preferable as the *a priori* one tends to be conservative.

6.3.2 The case of empty core

In our framework there might exist some realization $\delta \in \Delta$ resulting in $C(G_\Delta) = \emptyset$ (and hence the same can happen to its scenario counterpart, $C(G_K)$). Establishing whether $C(G_\Delta)$ is nonempty amounts to solving the feasibility program

$$\begin{cases} \min\limits_{x \geq 0} & 0 \\ \text{subject to} & \sum\limits_{i \in \mathcal{N}} x_i = u(\mathcal{N}), \\ & \sum\limits_{i \in S} x_i \geq \max\limits_{\delta \in \Delta} u(S, \delta) \text{ for all } S \subset \mathcal{N}. \end{cases}$$

To circumvent the possibly infinite cardinality of the constraint set which must hold for all $\delta \in \Delta$, we adopt a data-driven formulation, thus obtaining the following scenario program

$$\begin{cases} \min\limits_{x \geq 0} & 0 \\ \text{subject to} & \sum\limits_{i \in \mathcal{N}} x_i = u(\mathcal{N}), \\ & \sum\limits_{i \in S} x_i \geq \max\limits_{k=1,\dots,K} u(S, \delta^{(k)}) \text{ for all } S \subset \mathcal{N}. \end{cases}$$

In this case, we lift Assumption 6.1 on non-emptiness of the scenario core and provide through Theorem 6.2 guarantees on the probabilistic stability of an allocation in a relaxed version of the scenario core, the so called *scenario ζ-core*:

Definition 6.5 (Scenario ζ-core) *Given some $\zeta \geq 0$, the scenario ζ-core of a scenario game G_K coincides with the set*

$$C_\zeta(G_K) = \{x \in \mathbb{R}^N : \textstyle\sum_{i \in \mathcal{N}} x_i = u(\mathcal{N}), \sum_{i \in S} x_i \geq \max_{k=1,\dots,K} u(S, \delta^{(k)}) - \zeta \ \forall\, S \subseteq \mathcal{N}\}.$$

□

Note that the standard notion of the scenario core is recovered as a special case of the scenario ζ-core by setting $\zeta = 0$. In other words, the scenario ζ-core is a set of

allocations based on available data, where no agent can improve its payoff by leaving the grand coalition. If it happens, then it must pay a penalty of ζ for leaving. We first aim at finding a solution inside the *least-core*, i.e., the set where allocations are stable with the minimum value of ζ. Finding such a solution in our set-up amounts to solving the following tractable optimization program, formulated in the spirit of the scenario approach with relaxation [25, 26],

$$\begin{cases} \min\limits_{x\geq 0,\ \xi\geq 0} & \sum\limits_{k=1}^{K} \xi_k \\ \text{subject to} & \sum\limits_{i\in\mathcal{N}} x_i = u(\mathcal{N}), \\ & \sum\limits_{i\in S} x_i \geq \max\limits_{k=1,\ldots,K} u(S,\delta^{(k)}) - \xi_k, \\ & \qquad\qquad \text{for all } S\subset\mathcal{N}, k=1,\ldots,K. \end{cases} \tag{6.4}$$

Note that the proposed results hold independently of the way (6.3) and (6.4) are solved. In this regard, we point out that they amount to linear programs, hence amenable to distributed computation [7], [153]. Let $\xi_K^* := \{\xi_k^*\}_{k=1}^K$ and $\zeta^* = \max_{k=1,\ldots,K} \xi_k^*$. A solution to (6.4) is a pair (x_K^*, ξ_K^*), where x_K^* is an allocation in the ζ^*-core (the least core). We make use of the following assumptions:

Assumption 6.2 (Uniqueness [25]) *For any multi-sample $\delta_K \in \Delta^K$, the solution (x_K^*, ξ_K^*) of* (6.4) *is unique.* □

In case the solution is not unique, a solution can be singled out by applying a convex tie-break rule [25].

Assumption 6.3 (Non-accumulation [25]) *For every allocation $x\in\mathbb{R}^N$, $\mathbb{P}\{\delta\in\Delta : \sum_{i\in S} x_i = u(S,\delta)\} = 0$.* □

Assumption 6.3 is related to nondegeneracy [23] and it is often satisfied when the uncertain parameter δ does not accumulate, e.g., when δ admits a (continuous) density function [25]. From [26, Thm. 1], consider the polynomial equation

$$\binom{K}{s} t^{K-s} - \frac{\beta}{2N} \sum_{i=s}^{K-1} \binom{i}{s} t^{i-s} - \frac{\beta}{6K} \sum_{i=K+1}^{4K} \binom{i}{s} t^{i-s} = 0, \tag{6.5}$$

for any $s = 0,\ldots,N-1$, while for $s = K$ consider

$$1 - \frac{\beta}{6K} \sum_{i=K+1}^{4K} \binom{i}{K} t^{i-K} = 0. \tag{6.6}$$

For any $s = 0, \dots, N-1$, equation (6.5) has exactly two solutions in $[0, +\infty)$ (see [154, Thm. 1]), which we denote as $\underline{t}(s)$ and $\bar{t}(s)$ with $\underline{t}(s) \leq \bar{t}(s)$. Equation (6.6) has only one solution in $[0, +\infty)$ which we denote as $\bar{t}(K)$, while we define $\underline{t}(K) = 0$. Let $\underline{\varepsilon}(s) = \max\{0, 1 - \bar{t}(s)\}$ and $\overline{\varepsilon}(s) = 1 - \underline{t}(s)$, $s = 0, 1, \dots, K$. We are now ready to introduce our main theorem on the provision of stability guarantees for allocations inside a scenario ζ-core.

Theorem 6.2 *Consider Assumptions 6.2 and 6.3 and fix a confidence parameter $\beta \in (0,1)$. We then have that*

$$\mathbb{P}^K\{\delta_K \in \Delta^K : \underline{\varepsilon}(s^*) \leq V(x_K^*) \leq \overline{\varepsilon}(s^*)\} \geq 1 - \beta, \tag{6.7}$$

where x_K^* is obtained from (6.4)*, s^* is the number of $\delta^{(k)}$'s for which there exists a coalition $S \subset \mathcal{N}$ such that $\sum_{i \in S} x_i^* \leq u(S, \delta^{(k)})$, and $\underline{\varepsilon}(\cdot)$, $\overline{\varepsilon}(\cdot)$ are as defined below* (6.6)*.* □

Proof: Note that for any $\delta \in \Delta$ we can equivalently rewrite $\sum_{i \in S} x_i \geq u(S, \delta)$ for all $S \subset \mathcal{N}$ as $\max_{\ell = 1, \dots, 2^N - 1} (b_\ell(\delta) - A_\ell x) \leq 0$, where the ℓ-th row A_ℓ of the matrix $A \in \mathbb{R}^{(2^N-1) \times N}$ selects those agents generating the ℓ-th sub-coalition S_ℓ among the possible $2^N - 1$, while $b(\delta) := [u(S_1, \delta) \ \cdots \ u(S_{2^N-1}, \delta)]^\top \in \mathbb{R}^{2^N-1}$. As such, the program in (6.4) takes the form

$$\begin{cases} \min\limits_{x \geq 0, \xi \geq 0} & \sum\limits_{k=1}^{K} \xi_k \\ \text{s.t.} & \sum\limits_{i \in \mathcal{N}} x_i = u(\mathcal{N}) \\ & f(x, \delta^{(k)}) \leq \xi_k, \text{ for all } k = 1, \dots, K, \end{cases} \tag{6.8}$$

with $f(x, \delta) := \max_{\ell = 1, \dots, 2^N - 1} (b_\ell(\delta) - A_\ell x)$. Since, in our setting, [26, Ass. 1] is satisfied and considering Assumptions 6.2 and 6.3, direct application of [26, Th. 1] to (6.8) allows the provision of probabilistic stability guarantees for an allocation x_K^* inside the scenario least core obtained once the values of $\{\xi_k^*\}_{k=1}^K$ have been computed. ■

The formulation in (6.4) brings several benefits. Specifically, after solving (6.4), which is always feasible, if each $\xi_k^* = 0$, then the scenario core $C(G_K)$ is nonempty, meaning that, for the collected data, the grand coalition admits at least a stable solution whose robustness

can be quantified following the discussion in Subsection 6.3.1 or the bound in Theorem 6.2. However, no conclusion can be drawn on the nonemptiness of the original robust core, $C(G_\Delta)$. In case there exists some $\xi_k^* > 0$, this means that the scenario core $C(G_K)$ is empty and the same holds for the robust core $C(G_\Delta)$ (since $C(G_\Delta) \subseteq C(G_K)$). Interestingly, Theorem 6.2 allows us to quantify the penalty to be imposed on the formation of subcoalitions in order to obtain a solution with a provable probabilistic bound on stability.

6.4 Numerical Study

We consider an uncertain coalitional game with $N = 4$ agents and the value functions for each coalition summarized in Table 6.1. Each element of the uncertainty $\delta = \{\delta_S\}_{S \subset \mathcal{N}}$ that affects the coalitional values is assumed to be drawn according to a uniform distribution with support in $[-0.5, 0.5]$. Drawing a total number of $K = \{200, 300, 400, 500, 1000, 1500, 2000\}$ samples we use Algorithm 2 to find the cardinality of the compression, which for this particular example is always upper bounded by 4. Fixing the confidence parameter to $\beta = 10^{-4}$ and considering 10^6 test samples, we calculate the empirical probability of violation $\hat{V}(C(G_K))$ by checking which of these test samples give rise to values of coalitions that render at least one allocation in the scenario core unstable. Then, we compare in Figure 6.2 this probability with the *a posteriori* theoretical violation level $\varepsilon(s_K)$ of the scenario core for each value of K. Note that the empirical probability of violation is always below the theoretical level, thus verifying the result of Theorem 6.1 numerically.

Next, we focus on instances of games where the scenario core as defined in Section 6.2 is empty due to the presence of uncertainty. We consider the same coalitional game as before where the uncertainty δ follows instead a gaussian distribution with mean 0 and standard deviation 0.3 truncated at [-1.5, 1.5]. In this case, the scenario core is empty for a large enough number of samples. Using the relaxation technique of Subsection 6.3.2 we find a solution that lies inside the randomized least core. Then, Theorem 6.2 allows to provide guarantees on the stability of an allocation x_K^* in this relaxed core. Drawing a different number of samples K each time from the set $\{200, 300, 400, 500, 1000, 1500, 2000\}$ and fixing $\beta = 10^{-5}$, we calculate s^* for each multi-sample size and compute the lower and

Table 6.1: Uncertain value functions for each coalition

Coalition	Value of coalition
$\{1\}$	$1+\delta_1$
$\{2\}$	$1.5+\delta_2$
$\{3\}$	$1+\delta_3$
$\{4\}$	$2+\delta_4$
$\{1,2\}$	$6.5+\delta_{12}$
$\{2,3\}$	$6.5+\delta_{23}$
$\{3,4\}$	$7+\delta_{34}$
$\{1,3\}$	$6+\delta_{13}$
$\{1,4\}$	$7+\delta_{14}$
$\{2,4\}$	$7.5+\delta_{24}$
$\{1,2,3\}$	$11.5+\delta_{123}$
$\{1,2,4\}$	$12.5+\delta_{124}$
$\{1,3,4\}$	$12+\delta_{134}$
$\{2,3,4\}$	$12.5+\delta_{234}$
$\{1,2,3,4\}$	17.3

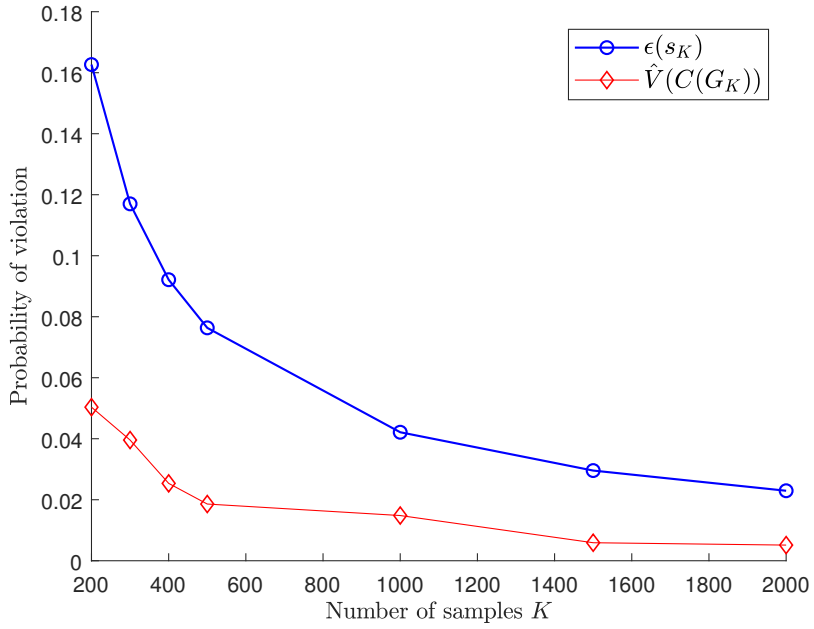

Figure 6.2: The violation level $\varepsilon(s_K)$ vs the empirical probability of violation $\hat{V}(C(G_K))$ for the scenario core of the uncertain game as defined by Table 6.1.

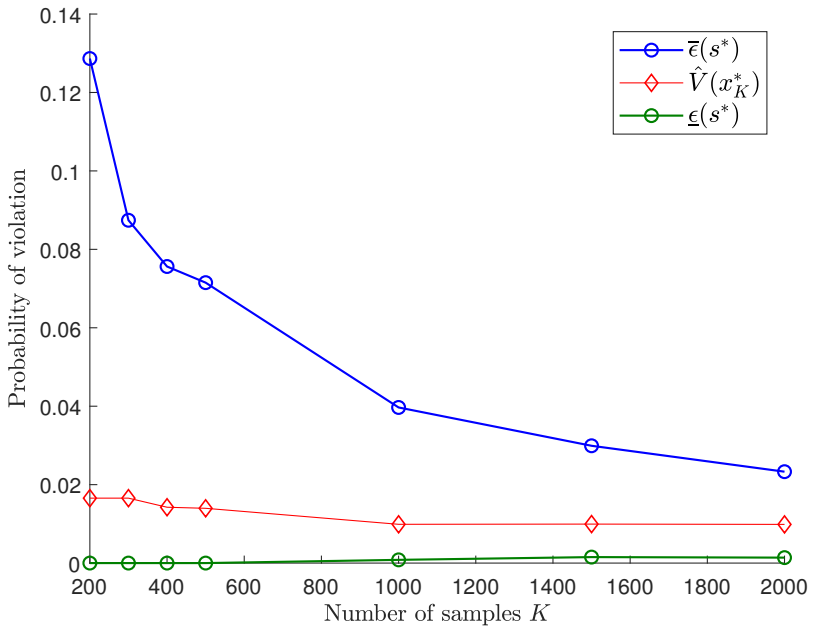

Figure 6.3: The violation levels $\overline{\varepsilon}(s^*), \underline{\varepsilon}(s^*)$ vs the empirical probability of violation $\hat{V}(x_K^*)$ for a solution inside the scenario ζ-core of the considered uncertain game.

upper violation levels $\underline{\varepsilon}(s^*), \overline{\varepsilon}(s^*)$, respectively. Using a total number of 10^6 test samples we calculate the empirical probability $\hat{V}(x_K^*)$ and we compare it with the theoretical bounds $\overline{\varepsilon}(s^*), \underline{\varepsilon}(s^*)$ (these are obtained as discussed below (6.6), using the code made available in [154]) for each multi-sample size. The results are summarized in Figure 6.3. In general, a lower probability of violation is expected by increasing the number of samples. However, note that a larger number of samples implies in our case that a larger number of relaxed constraints is generally needed to find an allocation by solving (6.4). Thus, such allocations in the relaxed core have similar or only slightly decreasing empirical probabilities of violation as the number of samples increases. Note from Figure 6.3 that the lower bound on the probability of violation returns a non-zero value and thus becomes meaningful only when a certain amount of samples is drawn.

6.5 Concluding remarks

Focusing on coalitional games with uncertain value functions, we borrow tools from the scenario theory to provide distribution-free guarantees on the probability that agents will not defect from the grand coalition to originate subcoalitions for unseen uncertainty

realizations. We provided collective stability guarantees for the entire core set in a PAC learning fashion, and accounted for the case where the scenario core can be empty, proposing a methodology to accompany allocations in a relaxed core with stability certificates. Future work will focus on establishing tighter stability guarantees on a specific core subset where desired qualities of the solution, e.g., fairness within a certain amount of tolerance, are taken into account.

7 Conclusion and Future work

This chapter concludes the thesis through a summary of our main contributions and by proposing potential research directions that build upon the results developed in the previous chapters, and aim at motivating further research to some fundamental research questions that arise directly from the topics that it explores.

7.1 Summary of contributions

In Chapter 3, we showed that guarantees on the probabilistic feasibility of the entire feasibility region or subsets of it can be provided in an *a posteriori* fashion. An immediate consequence of this result is the provision of certificates for solutions to multi-agent optimization programs and games with uncertain convex constraints. This result also alleviates any challenges related to the fact that numerical implementations of most solution-seeking algorithms converge to a suboptimal feasible solution. Furthermore, it paves the way towards the provision of collective guarantees when agents deviate from the proposed optimizer/ equilibrium within the feasible domain as shown in Chapter 4. Numerical studies on electric vehicle charging control in both a multi-agent optimization and a game-theoretic setting corroborate our theoretical results.

In Chapter 4 we focused on aggregative games with uncertainty in the coupling constraints. We consider a setting where bounded deviations of agents' decisions from

the equilibrium are allowed. Following a data-driven paradigm, we design a data-driven primal-dual algorithm, based on the class of asymmetric projection algorithms (APA), that converges to a randomized equilibrium such that we can control in an *a priori* fashion the probabilistic feasibility levels collectively for all these possible deviations by adapting the position of the nominal equilibrium. Our algorithmic scheme is based on a novel methodology of identifying online the coupling constraints that play the most significant role to the systems' performance. Tuning of the provided guarantees is accomplished by controlling how many of these coupling constraints are allowed to intersect the region under study. Finally, numerical studies of a game with uncertain constraints on the aggregate decision validate our theoretical results.

In Chapter 5 we focused on a particular class of multi-agent programs, where agents cooperate to minimize a common cost, expressed as a function of the aggregate decision and affected by uncertainty. By exploiting the structure of the program under study and leveraging on the so called support rank notion, we provide for the optimal solution of the program distribution-free certificates on the probability that the cost will deteriorate. The provided certificates do not depend on the number of agents, thus alleviating any scalability issues. The constructed bound on the probability of cost deterioration depends only on the dimension of the agents' decision. This leads to a significant improvement as it substantially reduces the number of samples required to achieve a certain level of probabilistic robustness as the number of agents increases. Our certificates can be used alongside any convex optimization algorithm centralised, decentralised or distributed, to obtain an optimal solution of the underlying problem. Finally, our theoretical results are accompanied by a numerical study on the electric vehicle charging problem that validates the agent-independence property of the probabilistic bounds.

In Chapter 6, we unified uncertain coalitional games and the scenario approach. Considering value functions parameterized by an uncertain parameter we provided guarantees on the stability of the proposed allocations, i.e., agents do not have incentives to defect from the grand coalition to form subcoalitions for unseen realizations of the uncertain parameter. For the cases when the core of the game is empty, we propose a randomized relaxation of this stability concept, the so called scenario ε-core. We

then show that those allocations that belong to this relaxed set can be accompanied by probabilistic stability guarantees leveraging results from the scenario approach with relaxation. Finally, numerical studies validate our theoretical bounds.

7.2 Future work

The contributions of Chapters 3 and 4 no longer restrict the applicability of the scenario approach to a single point of interest, which is usually an optimal solution or game equilibrium. However, there is still a fundamental challenge to be addressed when it comes to the transition from the provision of guarantees for a unique point to that of a set within the realm of the scenario approach. Our theory has not yet introduced a methodology such that we can distinguish between the case where a significant part of the set under study is violated by a new yet unseen constraint and the case where only a small portion of the set is violated. This challenge is not present when we study the probabilistic feasibility of a given point as the answer on whether a new yet unseen constraint violates or not this point is binary, i.e., it either violates or not. For the case of a solution set the notion of constraint violation is not that straightforward. As such, a methodology is required such that we can provide guarantees on the probability that a predefined maximum portion of the solution set will be violated by a new yet unseen constraint. This is a non-trivial problem for which further research is required.

In Chapter 4 the algorithm chooses to relax the coupling constraints that lead to the minimization of the potential function given the restrictions on the number of facets that are allowed to be intersected that in turn affect the guarantees on constraint violation for the region under study. The fact that the obtained equilibrium is the minimizer of the potential function has been established through extensive numerical simulations and a proof that the algorithm converges to the solution of the proposed VI problem is presented. However, a theoretical proof that the algorithmic scheme proposed converges to the optimizer of the potential function, while allowing at most M facets to intersect the region is not present in this thesis. The next step for Chapter 4 is to prove theoretically that this is actually the case by connecting the solution of the proposed VI with the optimizers of the potential function.

Simulation results in [48], indicate that agent-independent guarantees can also be provided for a class of multi-agent games with agents' cost functions with structure similar to ours, even though a detailed proof for that observation has not been provided in the research literature. In future work, we wish to find more general classes of problem for which this property holds, thus extending the work presented in Chapter 5.

With respect to Chapter 6, a fundamental challenge in uncertain coalitional games is related to how fair the proposed allocations are for each agent. An interesting research direction would involve leveraging ideas from Chapter 4, assuming that a fair solution, with respect to some measure of fairness, lies inside the scenario core in a similar way that an optimal solution or equilibrium lies within the feasible domain. Building upon this parallelism one can provide probabilistic guarantees on the stability of a set of allocations in the scenario core for the more general case, where agent allocations deviate from the fair solution.

In what follows we propose some more general ideas for further research. Taking a large number of samples can be risky as the resulting feasible region can be empty. In this case, instead of the scenario approach with relaxation, it would be interesting to connect the sampling and discarding approach [30], [115] with our methodology. Another interesting research direction revolves around the connection of our results with the recent developments in probabilistic scaling [109]. This direction might prove to be useful for the provision of guarantees that do not depend on the dimension of the decision vector. It is important to note that this methodology, though different, can also be used to provide guarantees for a region of the decision space.

Finally, the methodologies developed in this thesis are based on the assumption that the drawn samples are i.i.d. The i.i.d. assumption can be realistic in many applications of interest and eases our probabilistic model significantly. However, respecting this assumption is also related with the data collection process and there are cases where one cannot obtain easily uncorrelated data. One straightforward approach to solve this problem is to treat the parts of the data that correlate with each other as one sample. Then one can employ randomized algorithms that converge to a solution with certain probabilistic certificates. In case these samples cannot be treated as one large sample though, one

has to lift this assumption. Exploring how the scenario approach can be combined with techniques from statistical learning that lift this assumption is an interesting, though, at the same time challenging task that we would like to explore in future work.

References

[1] D. Ioli, L.Deori, A.Falsone, and M.Prandini. “A two-layer decentralized approach to the optimal energy management of a building district with a shared thermal storage”. In: *IFAC-PapersOnLine* 50 (July 2017), pp. 8844–8849.

[2] B. Gharesifard, T. Başar, and A. Domínguez-García. “Price-based coordinated aggregation of networked distributed energy resources”. English (US). In: *IEEE Transactions on Automatic Control* 61.10 (Jan. 2016), pp. 2936–2946.

[3] D. Paccagnan, B. Gentile, F. Parise, M. Kamgarpour, and J. Lygeros. “Nash and Wardrop Equilibria in Aggregative Games With Coupling Constraints”. In: *IEEE Transactions on Automatic Control* 64.4 (2019), pp. 1373–1388.

[4] D. Acemoglu and M. K. Jensen. “Aggregate comparative statics”. In: *Games and Economic Behavior* 81 (2013), pp. 27 –49.

[5] J. Ghaderi and R. Srikant. “Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate”. In: *Automatica* 50.12 (2014), pp. 3209 –3215.

[6] D. Acemoglu and A. Ozdaglar. “Opinion Dynamics and Learning in Social Networks”. In: *Dynamic Games and Applications* 1.1 (2011), pp. 3–49.

[7] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.

[8] N. Parikh and S. Boyd. “Proximal Algorithms”. In: *Foundations and Trends in Optimization* 1.3 (2014), pp. 127–239.

[9] T. Başar and G. J. Olsder. “Dynamic Non-Cooperative Game Theory”. In: *Society for Industrial and Applied Mathematics* (1999).

[10] D. Bertsekas, A. Nedic, and A. E. Ozdaglar. “Convex Analysis and Optimization”. In: *Athena Scientific* (Jan. 2003).

[11] D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.

[12] G. Chalkiadakis, E. Elkind, and M. Wooldridge. “Computational aspects of cooperative game theory”. In: *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5.6 (2011), pp. 1–168.

[13] G. Scutari, F. Facchinei, J. Pang, and D. P. Palomar. “Real and Complex Monotone Communication Games”. In: *IEEE Transactions on Information Theory* 60 (2014), pp. 4197–4231.

[14] M. J. Smith. “The existence, uniqueness and stability of traffic equilibria”. In: *Transportation Research Part B: Methodological* 13.4 (1979), pp. 295–304.

[15] W. Saad, Z. Han, H. V. Poor, and T. Basar. “Game-Theoretic Methods for the Smart Grid: An Overview of Microgrid Systems, Demand-Side Management, and Smart Grid Communications”. In: *IEEE Signal Processing Magazine* 29.5 (2012), pp. 86–105.

[16] Z. Ma, D. S. Callaway, and I. Hiskens. “Decentralized Charging Control of Large Populations of Plug-in Electric Vehicles”. In: *IEEE Transactions on Control Systems Technology* 21 (Jan. 2013), pp. 67–78.

[17] A.A. Raja and S. Grammatico. “Payoff distribution in robust coalitional games on time-varying networks”. In: *IEEE Transactions on Control of Network Systems* 1 (2021), pp. 511–520.

[18] J. Dixon, T. Morstyn, L. Han, and M. McCulloch. “Flexible cooperative game theory tool for peer-to-peer energy trading analysis”. In: *2018 IEEE Power and Energy Society General Meeting (PESGM)* (2019), pp. 1–5.

[19] H. Chen, Y. Li, R. H. Y. Louie, and B. Vucetic. “Autonomous Demand Side Management Based on Energy Consumption Scheduling and Instantaneous Load Billing: An Aggregative Game Approach”. In: *IEEE Transactions on Smart Grid* 5.4 (2014), pp. 1744–1754.

[20] G. Calafiore and M. Campi. “The Scenario Approach to Robust Control Design”. In: *IEEE Transactions on Automatic Control* 51 (June 2006), pp. 742 –753.

[21] M. Campi and S. Garatti. “The exact feasibility of randomized solutions of uncertain convex programs”. In: *SIAM Journal on Optimization* 19.3 (2008), pp. 1211–1230.

[22] M. Campi, S. Garatti, and M. Prandini. “The scenario approach for systems and control design”. In: *Annual Reviews in Control* 33 (2008), pp. 149–157.

[23] M. Campi and S. Garatti. “Wait-and-judge scenario optimization”. In: *Mathematical Programming* 167.1 (2018), pp. 155–189.

[24] M. Campi, S. Garatti, and F. A. Ramponi. “A General Scenario Theory for Nonconvex Optimization and Decision Making”. In: *IEEE Transactions on Automatic Control* 63.12 (2018), pp. 4067 –4078.

[25] M. Campi and S. Garatti. “Scenario optimization with relaxation: A new tool for design and application to machine learning problems”. In: *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE. 2020, pp. 2463–2468.

[26] M. Campi and S. Garatti. “A Theory of the Risk for Optimization with Relaxation and its Application to Support Vector Machines”. In: *Journal of Machine Learning Research* 22.288 (2021), pp. 1–38.

[27] J. R. Birge and F. Louveaux. “Introduction to Stochastic Programming”. In: *Springer Series in Operations Research and Financial Engineering* (1997).

[28] D. Bai, T. Carpenter, and J. Mulvey. “Making a Case for Robust Optimization Models”. In: *Management Science* 43 (July 1997), pp. 895–907.

[29] G. Calafiore and M. Campi. “Uncertain convex programs: randomized solutions and confidence levels”. In: *Mathematical Programming* 102.1 (2005), pp. 25–46.

[30] M. Campi and S. Garatti. “A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality”. In: *Journal of Optimization Theory and Applications* 148.2 (2011), pp. 257–280.

[31] S. Grammatico, X. Zhang, K. Margellos, P. Goulart, and J. Lygeros. “A Scenario Approach for Non-Convex Control Design”. In: *IEEE Transactions on Automatic Control* 61 (2016), pp. 334–345.

[32] C. F. Camerer. “Behavioral Game Theory: Experiments in Strategic Interaction”. In: *Princeton University Press* (2003).

[33] C. Plott and V. Smith. “Handbook of Experimental Economics Results”. In: *North Holland, Elsevier* (2008).

[34] D. Kahneman and A. Tversky. “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47.2 (1979), pp. 263–291.

[35] Z. Zhang, C. Chen, Q. Xing, Z. Ji, and X. Huang. “Comprehensive Optimal Scheduling Strategy of Multi-Element Charging Station for Bounded Rational Users”. In: *IEEE Access* 9 (2021), pp. 9442–9452.

[36] Y. Wang, W. Saad, N. B. Mandayam, and V. H. Poor. “Integrating energy storage into the smart grid: A prospect theoretic approach”. In: *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2014), pp. 7779–7783.

[37] A. Rubinstein. *Modeling Bounded Rationality*. MIT Press Books. The MIT Press, 1997.

[38] A. Charnes and D. Granot. “Coalitional and Chance-Constrained Solutions to n-Person Games I: The Prior Satisficing Nucleolus”. In: *SIAM Journal on Applied Mathematics* 31.2 (1976), pp. 358–367.

[39] A. Charnes and D. Granot. “Coalitional and Chance-Constrained Solutions to n-Person Games II: Two-Stage Solutions”. In: *Operations Research* 25.6 (1977), pp. 1013–1019.

[40] A. Charnes and D. Granot. “Prior Solutions: Extensions of Convex Nucleus Solutions to Chance-Constrained Games.” In: *Proceedings of the Computer Science and Statistics Seventh Symposium at Iowa State University* (1973), pp. 323–332.

[41] G. Chalkiadakis and C. Boutilier. “Bayesian reinforcement learning for coalition formation under uncertainty”. In: *AAMAS* (2004), pp. 1090–1097.

[42] S. Ieong and Y. Shoham. “Bayesian Coalitional Games.” In: *AAAI* (Jan. 2008), pp. 95–100.

[43] Y. Li and V. Conitzer. “Cooperative Game Solution Concepts that Maximize Stability under Noise”. In: *AAAI* (2015), pp. 979–985.

[44] G. Chalkiadakis and C. Boutilier. “Sequential decision making in repeated coalition formation under uncertainty”. In: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems* (2008), pp. 347–354.

[45] J. Guéneron and G. Bonnet. “Are Exploration-Based Strategies of Interest for Repeated Stochastic Coalitional Games?” In: *Advances in Practical Applications of Agents, Multi-Agent Systems, and Social Good. The PAAMS Collection* (2021), pp. 89–100.

[46] G. Schildbach, L. Fagiano, and M. Morari. “Randomized solutions to convex programs with multiple chance constraints”. In: *SIAM Journal on Optimization* 23 (May 2012), pp. 2479–2501.

[47] S. Grammatico, F. Parise, M. Colombino, and J. Lygeros. “Decentralized Convergence to Nash Equilibria in Constrained Deterministic Mean Field Control”. In: *IEEE Transactions on Automatic Control* 61.11 (2016), pp. 3315–3329.

[48] F. Fele and K. Margellos. “Probably Approximately Correct Nash Equilibrium Learning”. In: *IEEE Transactions on Automatic Control* 66.9 (2021), pp. 4238–4245.

[49] F. Fele and K. Margellos. “Probabilistic sensitivity of Nash equilibria in multi-agent games: a wait-and-judge approach”. In: *2019 IEEE 58th Conference on Decision and Control (CDC)*. 2019, pp. 5026–5031.

[50] S. Boyd and L. Vandenberghe. *Convex Optimization*. USA: Cambridge University Press, 2004.

[51] Y. Nesterov. *Lectures on Convex Optimization*. 2nd. Springer Publishing Company, Incorporated, 2018.

[52] A. Beck. *First-Order Methods in Optimization*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2017.

[53] P. You, S. H. Low, W. Tushar, G. Geng, C. Yuen, Z. Yang, and Y. Sun. “Scheduling of EV Battery Swapping—Part I: Centralized Solution”. In: *IEEE Transactions on Control of Network Systems* 5.4 (2018), pp. 1887–1897.

[54] P. You, S. H. Low, L. Zhang, R. Deng, G. B. Giannakis, Y. Sun, and Z. Yang. “Scheduling of EV Battery Swapping–Part II: Distributed Solutions”. In: *IEEE Transactions on Control of Network Systems* 5.4 (2018), pp. 1920–1930.

[55] J. L. von Neumann and O. Morgenstern. “The Theory of Games and Economic Behavior”. In: *Princeton University Press* (1972).

[56] J. F. Nash. “Equilibrium Points in *N*-Person Games”. In: *Proceedings of the National Academy of Sciences of the United States of America* 36 (1950), pp. 48–49.

[57] J.F. Nash. “Non-cooperative Games”. In: *Annals of Mathematics* 54 (1951), pp. 286–295.

[58] D. Paccagnan, F. Parise, and J. Lygeros. “On the Efficiency of Nash Equilibria in Aggregative Charging Games”. In: *IEEE Control Systems Letters* 2.4 (2018), pp. 629–634.

[59] L. Deori, K. Margellos, and M. Prandini. “On the connection between Nash equilibria and social optima in electric vehicle charging control games”. In: *IFAC-PapersOnLine* 50.1 (2017). 20th IFAC World Congress, pp. 14320–14325.

[60] L. Deori, K. Margellos, and M. Prandini. “Price of anarchy in electric vehicle charging control games: When Nash equilibria achieve social welfare”. In: *Automatica* 96 (2018), pp. 150 –158.

[61] E. Koutsoupias and C. Papadimitriou. “Worst-case equilibria”. In: *Computer Science Review* 3.2 (2009), pp. 65 –69.

[62] J. B. Rosen. “Existence and Uniqueness of Equilibrium Points for Concave N-Person Games”. In: *Econometrica* 33.3 (1965), pp. 520–534.

[63] F. Facchinei and J. S. Pang. “Finite-dimensional variational inequalities and complementarity problems”. In: *Springer* (2007).

[64] F. Facchinei and C. Kanzow. “Generalized Nash Equilibrium Problems”. In: *Annals of Operations Research* 175 (2010), pp. 177–211.

[65] H. Yin, U. V. Shanbhag, and P. G. Mehta. “Nash Equilibrium Problems With Scaled Congestion Costs and Shared Constraints”. In: *IEEE Transactions on Automatic Control* 56.7 (2011), pp. 1702–1708.

[66] G. Arslan, F. Demirkol, and S. Yüksel. “On Games With Coupled Constraints”. In: *IEEE Transactions on Automatic Control* 60 (2012), pp. 358–372.

[67] L. Pavel. “A doubly-augmented operator splitting approach for distributed GNE seeking over networks”. In: *2018 IEEE Conference on Decision and Control (CDC)* (2018), pp. 3529–3534.

[68] P. Yi and L. Pavel. “A distributed primal-dual algorithm for computation of generalized Nash equilibria via operator splitting methods”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. 2017, pp. 3841–3846.

[69] S. Grammatico. “Dynamic Control of Agents Playing Aggregative Games With Coupling Constraints”. In: *IEEE Transactions on Automatic Control* 62.9 (2017), pp. 4537–4548.

[70] M. K. Jensen. “Aggregative games and best-reply potentials”. In: *Economic Theory* 43.1 (2010), pp. 45–66.

[71] J. G. Wardrop. “Some theoretical aspects of road traffic research”. In: *Proceedings of the Institution of Civil Engineers* 1.3 (1952), pp. 325–362.

[72] D. Paccagnan, B. Gentile, F. Parise, M. Kamgarpour, and J. Lygeros. “Distributed computation of generalized Nash equilibria in quadratic aggregative games with affine coupling constraints”. In: *2016 IEEE 55th Conference on Decision and Control (CDC)* (2016), pp. 6123–6128.

[73] F. Parise, S. Grammatico, B. Gentile, and J. Lygeros. “Network Aggregative Games and Distributed Mean Field Control via Consensus Theory”. In: *CoRR* abs/1506.07719 (2015). URL: http://arxiv.org/abs/1506.07719.

[74] B. Swenson, R. W. Murray, and K. Soummya. “On Best-Response Dynamics in Potential Games”. English (US). In: *SIAM Journal on Control and Optimization* 56.4 (Jan. 2018), pp. 2734–2767.

[75] I. Gilboa and A. Matsui. “Social Stability and Equilibrium”. In: *Econometrica* 59.3 (1991), pp. 859–867.

[76] C. Harris. “On the Rate of Convergence of Continuous-Time Fictitious Play”. In: *Games and Economic Behavior* 22.2 (1998), pp. 238 –259.

[77] J. Hofbauer and S. Sorin. “Best response dynamics for continuous zero–sum games”. In: *Discrete and Continuous Dynamical Systems - B* 6.1 (2006), pp. 215–224.

[78] J. Hofbauer and K. Sigmund. “Evolutionary Game Dynamics”. In: *Bulletin of the American Mathematical Society* (2003), 479–519.

[79] E. Hopkins. “A Note on Best Response Dynamics”. In: *Games and Economic Behavior* 29.1 (1999), pp. 138–150.

[80] W. Yu, G. Ginis, and J. M. Cioffi. “Distributed Multiuser Power Control for Digital Subscriber Lines”. In: *IEEE Journal on Selected Areas in Communications* 20.5 (Sept. 2006), pp. 1105–1115.

[81] Z. Q. Luo and J. S. Pang. “Analysis of Iterative Waterfilling Algorithm for Multiuser Power Control in Digital Subscriber Lines”. In: *EURASIP Journal on Advances in Signal Processing* 2006.1 (2006), p. 024012.

[82] K. W. Shum, K. K. Leung, and C. W. Sung. “Convergence of Iterative Waterfilling Algorithm for Gaussian Interference Channels”. In: *IEEE Journal on Selected Areas in Communications* 25 (2007), pp. 1091–1100.

[83] G. Scutari, D. Palomar, and S. Barbarossa. “Optimal Linear Precoding Strategies for Wideband Non-Cooperative Systems Based on Game Theory—Part II: Algorithms”. In: *IEEE Transactions on Signal Processing* 56 (July 2007).

[84] G. Scutari, D. P. Palomar, and S. Barbarossa. “Competitive Design of Multiuser MIMO Systems Based on Game Theory: A Unified View”. In: *IEEE Journal on Selected Areas in Communications* 26.7 (2008), pp. 1089–1103.

[85] G. Scutari, D. P. Palomar, and S. Barbarossa. “The MIMO Iterative Waterfilling Algorithm”. In: *IEEE Transactions on Signal Processing* 57.5 (2009), pp. 1917–1935.

[86] G. Scutari, D. Palomar, and S. Barbarossa. “Cognitive MIMO Radio: A Competitive Optimality Design Based on Subspace Projections”. In: *CoRR* abs/0808.0978 (Aug. 2008).

[87] A. Cortés and S. Martínez. “Self-Triggered Best-Response Dynamics for Continuous Games”. In: *IEEE Transactions on Automatic Control* 60.4 (2015), pp. 1115–1120.

[88] T. Basar. “Relaxation techniques and asynchronous algorithms for on-line computation of noncooperative equilibria”. In: *26th IEEE Conference on Decision and Control* 26 (1987), pp. 275–280.

[89] L. Shu and T. Basar. “Distributed algorithms for the computation of noncooperative equilibria”. In: *Automatica* 23.4 (Jan. 1987), pp. 523–533.

[90] C. E. Lemke and J. T. Howson. “Equilibrium Points of Bimatrix Games”. In: *Journal of the Society for Industrial and Applied Mathematics* 12.2 (1964), pp. 413–423.

[91] S. Grammatico. “Aggregative control of competitive agents with coupled quadratic costs and shared constraints”. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. 2016, pp. 3597–3602.

[92] A. Ben-Tal and A. Nemirovski. “Robust Convex Optimization”. In: *Mathematics of Operations Research* 23 (Nov. 1998).

[93] B. L. Miller and H. M. Wagner. “Chance Constrained Programming with Joint Constraints”. In: *Operations Research* 13.6 (1965), pp. 930–945.

[94] A. Prékopa. “Stochastic Programming”. In: *Elsevier Science Publishers B. V.* (1986).

[95] B. R. Barmish and C. M. Lagoa. “The uniform distribution: rigorous justification for its use in robustness analysis”. In: *Proceedings of 35th IEEE Conference on Decision and Control* 3 (1996), pp. 3418–3423.

[96] G. Calafiore, F. Dabbene, and R. Tempo. “Randomized algorithms for probabilistic robustness with real and complex structured uncertainty”. In: *IEEE Transactions on Automatic Control* 45.12 (2000), pp. 2218–2235.

[97] G. Calafiore R. Tempo and F. Dabbenee. “Randomized Algorithms for Analysis and Control of Uncertain Systems: with Applications.” In: *Springer Science & Business Media* (2012).

[98] R. F. Stengel and L. E. Ryan. “Stochastic robustness of linear time-invariant control systems”. In: *IEEE Transactions on Automatic Control* 36.1 (1991), pp. 82–87.

[99] L. R. Ray and R. F. Stengel. “A monte carlo approach to the analysis of control system robustness”. In: *Automatica* 29.1 (1993), pp. 229 –236.

[100] P. Khargonekar and A. Tikku. "Randomized algorithms for robust control analysis and synthesis have polynomial complexity". In: *Proceedings of 35th IEEE Conference on Decision and Control* 3 (1996), pp. 3470–3475.

[101] V. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, 1998.

[102] M. Vidyasagar. *A Theory of Learning and Generalization: With Applications to Neural Networks and Control Systems*. Berlin, Heidelberg: Springer-Verlag, 1997.

[103] T. Alamo, R. Tempo, and E. Camacho. "Randomized Strategies for Probabilistic Solutions of Uncertain Feasibility and Optimization Problems". In: *IEEE Transactions on Automatic Control* 54.11 (2009), pp. 2545–2559.

[104] M. Vidyasagar. "Randomized algorithms for robust controller synthesis using statistical learning theory". In: *Automatica* 37.10 (2001), pp. 1515–1528.

[105] M. Vidyasagar and V. D. Blondel. "Probabilistic solutions to some NP-hard matrix problems". In: *Automatica* 37.9 (2001), pp. 1397–1405.

[106] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer. "Stochastic MPC with offline uncertainty sampling". In: *Automatica* 81 (2017), pp. 176–183.

[107] T. Alamo, V. Mirasierra, F. Dabbene, and M. Lorenzen. "Safe approximations of chance constrained sets by probabilistic scaling". In: *2019 18th European Control Conference (ECC)* (2019), pp. 1380–1385.

[108] M. Mammarella, V. Mirasierra, M. Lorenzen, T. Alamo, and F. Dabbene. "Chance-constrained sets approximation: A probabilistic scaling approach". In: *Automatica* 137 (2022), pp. 110–108.

[109] M. Mammarella, T. Alamo, F. Dabbene, and M. Lorenzen. "Computationally efficient stochastic MPC: a probabilistic scaling approach". In: *2020 IEEE Conference on Control Technology and Applications (CCTA)* (2020), pp. 25–30.

[110] D.P. de Farias and B. Van Roy. "On constraint sampling in the linear programming approach to approximate linear programming". In: *42nd IEEE International Conference on Decision and Control* 3 (2003), pp. 2441–2446.

[111] D. P. de Farias and B. Van Roy. "The Linear Programming Approach to Approximate Dynamic Programming". In: *Operations Research* 51 (2003), pp. 850–865.

[112] Y. Oishi. "Polynomial-time algorithms for probabilistic solutions of parameter-dependent linear matrix inequalities". In: *Automatica* 43 (2007), 538–545.

[113] T. Alamo, R. Tempo, A. Luque, and D. R. Ramirez. "Randomized methods for design of uncertain systems: Sample complexity and sequential algorithms". In: *Automatica* 52 (2015), pp. 160–172.

[114] G. Calafiore, F. Dabbene, and R. Tempo. "Research on probabilistic methods for control system design". In: *Automatica* 47.7 (2011), pp. 1279–1293.

[115] G. Calafiore. "Random Convex Programs". In: *SIAM Journal on Optimization* 20.6 (2010), pp. 3427–3464.

[116] T. Alamo, J.M. Manzano, and E. Camacho. "Robust Design Through Probabilistic Maximization: In Honor of Roberto Tempo". In: *Systems and Control: Foundations and Applications* (Jan. 2018), pp. 247–274.

[117] A. Carè, S. Garatti, and M. Campi. "FAST—Fast Algorithm for the Scenario Technique". In: *Operations Research* 62.3 (2014), pp. 662–671.

[118] R. Tempo T. Alamo and A. Luque. “On the sample complexity of randomized approaches to the analysis and design under uncertainty”. In: *Proceedings of the American Control Conference, Baltimore, USA, 2010* (2010).

[119] Y. Fujisaki and Y. Kozawa. “Probabilistic robust controller design: probable near minimax value and randomized algorithms”. In: *42nd IEEE International Conference on Decision and Control* 2 (2003), pp. 1938–1943.

[120] T. Alamo B. Karg and S. Lucia. “Probabilistic performance validation of deep learning-based robust NMPC controllers”. In: *arXiv preprint 1910.13906* (2019).

[121] K. Margellos, M. Prandini, and J. Lygeros. “On the connection between compression learning and scenario based single-stage and cascading optimization problems”. In: *IEEE Transactions on Automatic Control* 60.10 (2015), pp. 2716–2721.

[122] R. Rockafellar. *Convex analysis*. Princeton University Press, 1970.

[123] G. M. Ziegler. “Lectures on polytopes”. In: *Springer-Verlag New York, Inc* (1995).

[124] S. Barry. “Convexity: An Analytic Viewpoint”. In: *Cambridge University Press*. Cambridge Tracts in Mathematics (2011).

[125] Arne. Brøndsted. *An introduction to convex polytopes.* Graduate texts in mathematics. New York: Springer-Verlag, 1982.

[126] K. L. Clarkson. “More Output-Sensitive Geometric Algorithms”. In: *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*. SFCS ’94. USA: IEEE Computer Society, 1994, 695–702.

[127] K. Fukuda, B. Gärtner, and M. Szedlak. “Combinatorial Redundancy Detection”. In: *Annals of Operations Research* 265 (June 2018).

[128] P. Pflaum, M. Alamir, and M. Y. Lamoudi. “Probabilistic Energy Management Strategy for EV Charging Stations Using Randomized Algorithms”. In: *IEEE Transactions on Control Systems Technology* 26.3 (2018), pp. 1099–1106.

[129] P. Pflaum, M. Alamir, and M.Y. Lamoudi. “Battery sizing for PV power plants under regulations using randomized algorithms”. In: *Renewable Energy* 113 (2017), pp. 596–607.

[130] J. C. Harsanyi. “Bargaining in ignorance of the opponent’s utility function”. In: *Journal of Conflict Resolution* 6.1 (1962), pp. 29–38.

[131] J. C. Harsanyi. “Games with Incomplete Information Played by "Bayesian" Players, I-III. Part II. Bayesian Equilibrium Points”. In: *Management Science* 14.5 (1968), pp. 320–334.

[132] P. Couchman, B. Kouvaritakis, M. Cannon, and F. Prashad. “Gaming strategy for electric power with random demand”. In: *IEEE Transactions on Power Systems* 20.3 (2005), pp. 1283–1292.

[133] V. V. Singh, O. Jouini, and A. Lisser. “Existence of Nash equilibrium for chance-constrained games”. In: *Operations Research Letters* 44.5 (2016), pp. 640 –644.

[134] M. Aghassi and D. Bertsimas. “Robust Game Theory”. In: *Mathematical Programming* 107.1-2 (2006), pp. 231–273.

[135] S. Hayashi, N. Yamashita, and M. Fukushima. “Robust Nash Equilibria and Second-Order Cone Complementarity Problems”. In: *Journal of Nonlinear and Convex Analysis* 6 (2005).

[136] D. Paccagnan and M. Campi. “The Scenario Approach Meets Uncertain Game Theory and Variational Inequalities”. In: *2019 IEEE 58th Conference on Decision and Control (CDC)* (2019), pp. 6124–6129.

[137] “National Grid. Historical Demand Data”. In: *https://www.nationalgrideso.com/balancing-data* (2019).

[138] G. Pantazis, F. Fele, and K. Margellos. “On the probabilistic feasibility of solutions in multi-agent optimization problems under uncertainty”. In: *European Journal of Control* 63 (2022), pp. 186–195.

[139] M. K. Jensen. “Aggregative games and best-reply potentials”. In: *Economic Theory* 43 (2010), pp. 45–66.

[140] L. Deori, K. Margellos, and M. Prandini. “Regularized Jacobi Iteration for Decentralized Convex Quadratic Optimization With Separable Constraints”. In: *IEEE Transactions on Control Systems Technology* 27 (2016), pp. 1636–1644.

[141] S. Grammatico, F. Parise, M. Colombino, and J. Lygeros. “Decentralized convergence to Nash equilibria in constrained deterministic mean field control”. In: *IEEE Transactions on Automatic Control* 61.11 (2016), pp. 3315–3329.

[142] G. Pantazis, F. Fele, and K. Margellos. “Agent independent probabilistic robustness certificates for robust optimization programs with uncertain quadratic cost”. In: *Proceedings of the 59th IEEE Conference on Decision and Control* (2020), pp. 554–559.

[143] H. A. David and H. N. Nagaraja. “Order statistics”. In: *John Wiley & Sons,* (2004).

[144] R. Tempo, E.W. Bai, and F. Dabbene. “Probabilistic robustness analysis: Explicit bounds for the minimum number of samples”. In: *Systems & Control Letters* 30.5 (1997), pp. 237–242.

[145] M. Campi and S. Garatti. “Introduction to the Scenario Approach”. In: *Society of industrial and applied mathematics (SIAM)* (2018).

[146] S. Axler. “Linear Algebra Done Right”. In: *Springer*. Undergraduate Texts in Mathematics (1997).

[147] J.P.M. Suijs, P.E.M. Borm, A.M.B. De Waegenaere, and S.H. Tijs. “Cooperative games with stochastic payoffs”. In: *European Journal of Operational Research* 113.1 (1999), pp. 193–205.

[148] M. F. Balcan, A. D. Procaccia, and Y. Zick. “Learning Cooperative Games”. In: *IJCAI* (2015), 475–481.

[149] A.D. Procaccia and J. S. Rosenschein. “Learning to Identify Winning Coalitions in the PAC Model”. In: *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems* (2006), 673–675.

[150] F. Fabiani, K. Margellos, and P. J. Goulart. “Probabilistic feasibility guarantees for solution sets to uncertain variational inequalities”. In: *Automatica* 137 (2022), p. 110120.

[151] G. Pantazis, F. Fele, and K. Margellos. “A posteriori probabilistic feasibility guarantees for Nash equilibria in uncertain multi-agent games”. In: *IFAC-PapersOnLine* 53.2 (2020). 21st IFAC World Congress, pp. 3403–3408.

[152] F. Fabiani, K. Margellos, and P. J. Goulart. “On the robustness of equilibria in generalized aggregative games”. In: *2020 59th IEEE Conference on Decision and Control (CDC)* (2020), pp. 3725–3730.

[153] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers”. In: *Foundations and Trends in Machine Learning* 3 (Jan. 2011), pp. 1–122.

[154] S. Garatti and M. C. Campi. “Risk and Complexity in Scenario Optimization”. In: *Mathematical Programming* 191.1 (2022), 243–279.