

A posteriori probabilistic feasibility guarantees for Nash equilibria in uncertain multi-agent games^{*}

George Pantazis, Filiberto Fele, Kostas Margellos

Department of Engineering Science, University of Oxford, OX1 3PJ, UK,
(e-mail: georgios.pantazis@lmh.ox.ac.uk, filiberto.fele@eng.ox.ac.uk,
kostas.margellos@eng.ox.ac.uk)

Abstract: In this paper a distribution-free methodology is presented for providing robustness guarantees for Nash equilibria (NE) of multi-agent games. Leveraging recent *a posteriori* developments of the scenario approach, we provide probabilistic guarantees for feasibility problems with polytopic constraints. This result is then used in the context of multi-agent games, allowing to provide robustness certificates for constraint violation of any NE of a given game. Our guarantees can be used alongside any NE seeking algorithm. Finally, by exploiting the structure of our problem, we circumvent the need of employing computationally prohibitive algorithms to find an irreducible support subsample, a concept at the core of the scenario approach. Our theoretical results are accompanied by simulation studies that investigate the robustness of the solutions of two different problems, namely, a 2-dimensional feasibility problem and an electric vehicle (EV) charging control problem.

Keywords: Scenario approach, Multi-agent games, Nash equilibria, Feasibility guarantees, Electric vehicles

1. INTRODUCTION

Decentralized optimization and control of large scale systems is of significant interest for a variety of fields ranging from engineering and biology to economics and social sciences. In many cases, systems of this kind can be modelled as a network comprising self-interested entities/agents that interact/compete with each other in order to meet their own individual goals, thus giving rise to a non-cooperative set-up. Aiming at the resolution of the inherent conflict among the agents, a vast amount of research work has been focused on game-theoretic approaches. Especially in engineering, this framework has been extensively used for the analysis of communication (Scutari et al., 2014) and traffic (Smith, 1979) networks, smart grids (Saad et al., 2012), and electricity markets (Chen et al., 2014).

In real-world applications, however, the system is affected by uncertainty that can be due to several factors. This gives rise to games of incomplete information (Harsanyi, 1968). The concept of Nash equilibrium (NE) (see Definition 2) typically assumes complete information. For incomplete information games the notion of NE appears insufficient to deal with the presence of uncertainty, as the resulting agents' strategies do not necessarily exhibit any robustness properties against uncertainty. Motivated by this observation, two main research directions are encountered: 1) NE analysis based on particular models for the probability distribution of the uncertainty, (Couchman et al., 2005), (Singh et al., 2016) and/or the geometry of its support (Aghassi and Bertsimas, 2006), (Hayashi et al., 2005) and, 2) Distribution-free NE analysis, where no assumption on the probability distribution of the uncertainty is imposed.

In this paper we focus on distribution-free NE seeking, where algorithmic developments have been quite restricted. Motivated by the lack of distribution-free results, we leverage the recent developments of the so called scenario approach (Campi et al., 2018) and represent the uncertainty by means of scenarios that could be either available as historical data, or extracted via some prediction model. We attempt to transfer concepts from the scenario approach, well-understood in the context of optimization, to multi-agent games aiming to accompany game equilibria with guarantees on the probability of constraint satisfaction.

The scenario approach is a well-established mathematical technique (Campi and Calafiore, 2006), (Campi and Garatti, 2008), originally introduced to provide *a priori* probabilistic guarantees for solutions of uncertain convex optimization programs. Very recently, the theory has been extended to non-convex decision making problems (Campi et al., 2018), where the probabilistic guarantees are obtained in an *a posteriori* fashion. The main advantage of the scenario approach is its applicability under very general conditions, since it does not require the knowledge of the uncertainty set or the probability distribution, key assumptions in robust (Bai et al., 1997) and stochastic optimization (Birge and Louveaux, 1997), respectively. According to the scenario approach, the original problem can be approximated by solving a computationally tractable approximate problem, the so called scenario program consisting of a finite number of constraints, each of them corresponding to a different realization of the uncertain parameter.

Only a few data-driven works for distribution-free NE seeking have appeared in the literature with (Paccagnan and Campi, 2019) and (Fele and Margellos, 2019a), (Fele and Margellos, 2019b) being the most closely related to our work. Both papers attempt to bridge multi-agent games with the scenario approach, thus providing a distribution-free way to determine

^{*} Research was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant agreement EP/P03277X/1.

NE with quantifiable robustness properties. Specifically, in (Paccagnan and Campi, 2019), the authors focus on variational inequalities affected by uncertainty represented by means of scenarios. Due to the connection between variational inequalities and NE, their developments naturally find applications to multi-agent games. However, the theoretical analysis requires strong monotonicity of the operator associated with the variational inequalities, which in a gaming setting implies uniqueness of the NE. The latter is quite restrictive, as many games of practical interest have multiple (possibly infinite) NE. In (Fele and Margellos, 2019a) and (Fele and Margellos, 2019b), multi-agent games with uncertainty affecting agents' cost functions are considered, and once again uncertainty is tackled by means of the scenario approach. However, agents' constraint sets are assumed to be deterministic. Uncertain constraints could still be considered, under the restrictive assumption that they are decoupled. The approach proposed here is instead applicable to general uncertain constraints. The contributions of our paper with respect to the aforementioned works are the following:

- (1) Leveraging the recent results of (Campi et al., 2018), we provide *a posteriori* robustness certificates for the entire feasibility region of feasibility programs with polytopic constraints.
- (2) Focusing on multi-agent games, we provide distribution-free probabilistic guarantees for the entire set of NE in an *a posteriori* fashion. We extend the results of (Fele and Margellos, 2019a) to account for uncertainty in agents' (possibly) coupling constraints, while we do not require uniqueness of the associated NE as in (Paccagnan and Campi, 2019).
- (3) The probabilistic results of recent works in an *a posteriori* context rely on certain algorithms for quantifying a so called irreducible support subsample (see Definition 2 in (Campi et al., 2018)). Restricting our attention to multi-agent games, the use of such algorithms alongside NE seeking iterative algorithms, apart from being computationally prohibitive, leads to erratic behaviour due to numerical issues. In our case, the cardinality of the support subsample coincides with the number of facets of the polytopic constraint set, which is directly available.

The rest of the paper is organized as follows: Section 2 introduces the problem under study and offers a motivating application. Section 3 provides the theoretical analysis and proof of the main results, namely, providing *a posteriori* guarantees for feasibility problems. Section 4 provides numerical examples and revisits the electric vehicle application of Section 2. Finally, Section 5 concludes the paper and provides some directions for future work.

2. SCENARIO-BASED MULTI-AGENT GAMES WITH UNCERTAIN CONSTRAINTS

Let M be the total number of agents and $x^m = (x_t^m)_{t=1}^d$ the decision vector of dimension $d \in \mathbb{N}$ of agent $m \in \mathcal{M} = \{1, \dots, M\}$ taking values in the uncertain set $X_\delta^m \subseteq \mathbb{R}^d$ parameterized by the uncertain parameter δ . The uncertainty parameter is defined on the (possibly unknown) probability space $(\Delta, \mathcal{F}, \mathbb{P})$, where Δ is the sample space, equipped with a σ -algebra \mathcal{F} and a probability measure \mathbb{P} . Similarly, we define $x^{-m} = (x^1, \dots, x^{m-1}, x^{m+1}, \dots, x^M) \in X_\delta^{-m}$, where $X_\delta^{-m} = \prod_{j \in \mathcal{M}, j \neq m} X_\delta^j \subseteq \mathbb{R}^{d(M-1)}$, as the vector comprising the decisions of all other agents except for that of agent m . Finally, let

$\{\delta_i\}_{i=1}^N \in \Delta^N$ be a finite collection of independent and identically distributed (i.i.d.) scenarios/realisations of the uncertain vector δ , where Δ^N is the cartesian product of multiple copies of the sample space Δ . In our set-up, agents are considered as self-interested entities, i.e., they are interested in minimizing their own deterministic cost function $J_m : X \rightarrow \mathbb{R}$, where $X = \prod_{m=1}^M X^m \subseteq \mathbb{R}^{dM}$. Furthermore, we impose the following assumption:

Assumption 1. Each agent's decision set is formed by the intersection of a deterministic decision space X^m and uncertain affine constraints affected by any realization $\delta \in \Delta$, i.e., $\bigcap_{\delta \in \Delta} X_\delta^m = \{x^m \in X^m : g(x^m, x^{-m}, \delta) \leq 0\}$, $\forall m \in \mathcal{M}^1$, where $g : X \times \Delta \rightarrow \mathbb{R}$ is an affine function with respect to its first two arguments.

Under Assumption 1, we consider a multi-agent game, whose constraints are affected by uncertainty. Each agent $m \in \mathcal{M}$ seeks to minimize her own cost function, given the strategies x^{-m} of all other players, by solving the following program

$$\min_{x^m \in X^m} J_m(x^m, x^{-m}) \text{ subject to } x^m \in \bigcap_{\delta \in \Delta} X_\delta^m. \quad (1)$$

For the aforementioned problem, we consider the solution concept of NE as presented in Definition 2.

Definition 2. (Adapted from (Başar and Olsder, 1999)) A vector $x_{NE} = (x_{NE}^m)_{m \in \mathcal{M}}$ is a NE of the associated game if and only if $J_m(x_{NE}^m, x_{NE}^{-m}) \leq J_m(x^m, x_{NE}^{-m})$ for any $x^m \in \bigcap_{\delta \in \Delta} X_\delta^m$ and for any $m \in \mathcal{M}$.

Due to the presence of uncertainty and the (possibly) infinite cardinality of Δ , problem (1) is very difficult to solve, without imposing any assumptions on the geometry of the sample set Δ or the underlying probability distribution \mathbb{P} . To circumvent those issues, we approximate problem (1) by drawing multiple i.i.d. samples $\{\delta_i\}_{i=1}^N \in \Delta^N$ and then considering the following scenario-based NE seeking problem, where each agent $m \in \mathcal{M}$ solves the following optimization program

$$\min_{x^m \in X^m} J_m(x^m, x^{-m}) \text{ subject to } x^m \in \bigcap_{i=1, \dots, N} X_{\delta_i}^m. \quad (2)$$

Our aim is to provide probabilistic feasibility guarantees for the entire set of NE of (2) returned by an arbitrary NE seeking algorithm, i.e., to quantify the probability that any NE strategy x_{NE}^m of (2) belongs to the constraint set X_δ^m , $\forall m \in \mathcal{M}$ for a new unseen sample $\delta \in \Delta$. To this end, our analysis is primarily focused on feasibility problems affected by uncertainty. Based on the derived results we revisit problem (1) and attempt to provide robustness certificates for the obtained NE for a motivating application which fits in this class of games, i.e, the electric vehicle (EV) charging control problem

The EV-charging control problem can be treated as a non-cooperative game comprised of self-interested agents-vehicles each of them aiming at minimizing their own electricity cost, while their charging schedules are subject to certain specifications. The two main requirements for the operation of the system, namely, the lower and upper bounds imposed on the charging schedule and the total energy level to be achieved at the end of charging, can be modelled as constraints of affine form. However, most of the work up to this point assumed

¹ Formally, each agent's decision set possibly depends on the strategies of all other agents, thus allowing the use of our results in generalised NE problems. For simplicity we drop this dependence in the subsequent analysis, as our main focus is the treatment of uncertainty.

that these constraints are purely deterministic (Ma et al., 2013), (Paccagnan et al., 2019), (Deori et al., 2018). We extend this framework by imposing uncertainty on the constraints. In this case, each agent $m \in \mathcal{M}$ solves the following problem

$$\min_{x^m \in \mathbb{R}^d} J_m(x^m, x^{-m}) \text{ subject to}$$

$$x^m \in \bigcap_{\delta \in \Delta} \{[\underline{x}^m(\delta), \bar{x}^m(\delta)] \cap \{x^m \in \mathbb{R}^d : \sum_{t=1}^d x_t^m \geq E^m(\delta)\}\}. \quad (3)$$

The variables $x^m = (x_t^m)_{t=1}^d$ and J_m denote, respectively, the charging schedule for all time instances $t \in \{1, \dots, d\}$ and the electricity cost to be minimized for each vehicle $m \in \mathcal{M}$. The uncertain constraint for each vehicle m comprises of uncertain lower and upper bounds $\underline{x}^m(\delta)$ and $\bar{x}^m(\delta)$ and total energy levels $E^m(\delta)$. 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 adopt a data-based approach, a more viable alternative.

3. PROBABILISTIC GUARANTEES OF FEASIBILITY PROGRAMS

Motivated by the EV-charging control problem discussed in the previous section, we leverage the recent results of the scenario approach (Campi et al., 2018) to provide probabilistic guarantees for the (possibly) multiple NE of (3) in a computationally efficient manner. As NE constitute feasible solutions as far as constraint satisfaction is concerned, we focus on obtaining robustness certificates for a more general class of problems, that of feasibility problems under uncertain constraints, where the NE problem (3) emerges as a special case. As such, we consider the following feasibility problem:

$$P_\Delta : \text{find } x \in X, \text{ subject to } x \in \bigcap_{\delta \in \Delta} X_\delta,$$

where x is a decision vector belonging to the set $X \subset \mathbb{R}^d$ and δ is a random variable, defined as in Section 2, that encodes the uncertainty parameterizing agents' constraint sets. Note that for the gaming set-up, as introduced in Section 2, $X = \mathbb{X} = \prod_{m \in \mathcal{M}} X^m$ and $X_\delta = \prod_{m \in \mathcal{M}} X_\delta^m$. For the feasibility problem under study, we define the following scenario program:

$$P_N : \text{find } x \in X \text{ subject to } x \in \bigcap_{i=1, \dots, N} X_{\delta_i},$$

where N denotes the number of samples $\delta_i, i = 1, \dots, N$, drawn in an i.i.d. fashion from Δ according to the probability distribution \mathbb{P} and X_{δ_i} their respective constraints. Our results depend on an affine constraint structure, thus we impose the following standing assumption:

Assumption 3.

- (1) The deterministic constraint set X is a non-empty, compact and convex polytope.
- (2) For each $\delta \in \Delta$, $X_\delta = \{x \in \mathbb{R}^d : g(x, \delta) = a^T x - b \leq 0\}$, where g is an affine function given by the mapping $g : X \times \Delta \rightarrow \mathbb{R}$, where $a \in \mathbb{R}^d$, $b \in \mathbb{R}$ and $\delta = (a^T \ b) \in \mathbb{R}^{d+1}$.
- (3) For each multi-sample $\{\delta_i\}_{i=1}^N$ the polytope $\Pi_N = \{\bigcap_{i=1}^N X_{\delta_i}\} \cap X = \{x \in X : g(x, \delta_i) \leq 0, i = 1, \dots, N\}$ has a non-empty interior.

Note that vector valued affine functions are also captured by our framework; see example of Section 4. Assumption 3 guarantees that the polytope Π_N is compact and P_N admits at least one so-

lution for any chosen multisample $\{\delta_i\}_{i=1}^N$. Under Assumption 3 the feasibility problem P_N can be equivalently written as

$$P_N : \text{find } x \text{ subject to } x \in \Pi_N.$$

Upon finding the feasibility domain Π_N of the problem P_N , we are interested in investigating the robustness properties collectively for all the points of this domain 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 $x \in \Pi_N$. This concept, which is of crucial importance for our work, is known in the literature as the probability of violation and is adapted in our context to represent the probability of violation of a set. By Definition 1 in (Campi and Calafiore, 2006) the probability of violation of a given point $x \in \Pi_N$ is defined as

$$V(x) = \mathbb{P}\{\delta \in \Delta : x \notin X_\delta\}. \quad (4)$$

By Assumption 3, the probability of violation can be equivalently written as $V(x) = \mathbb{P}\{\delta \in \Delta : g(x, \delta) > 0\}$. We can now define the probability of violation of the polytope Π_N .

Definition 4. Let $\mathcal{P} \subseteq 2^X$ be the set of all non-empty, compact and convex polytopes that are subsets of X . For any $\Pi_N \in \mathcal{P}$ we define the probability of violation of the set Π_N as a mapping $\mathbb{V} : \mathcal{P} \rightarrow [0, 1]$ given by the following relation:

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

Definition 4 can be considered a special case of Definition 2 in (Grammatico et al., 2016). In Definition 5 three concepts of crucial importance are introduced.

- Definition 5.* (1) For any N , an algorithm is a mapping $A_N : \Delta^N \rightarrow \mathcal{P}$ that associates the multisample $\{\delta_i\}_{i=1}^N$ to a unique polytope $\Pi_N \in \mathcal{P}$.
- (2) Given a multisample $\{\delta_i\}_{i=1}^N \in \Delta^N$, a support subsample $S \subseteq \{\delta_i\}_{i=1}^N$ is a subset of the entire multisample with cardinality $k \leq N$ so that for $i_1 < i_2 < \dots < i_k$ and $I_k = \{i_1, i_2, \dots, i_k\}$, $S = \{\delta_i\}_{i \in I_k}$ is such that $A_k(\{\delta_i\}_{i \in I_k}) = A_N(\{\delta_i\}_{i=1}^N)$, i.e., the solution returned by an algorithm when fed with the subsample is the same with the one obtained when the entire multisample is used.
- (3) A support subsample function is a function of the form $B_N : \{\delta_i\}_{i=1}^N \rightarrow \{i_1, \dots, i_k\}$ that takes as input all the samples and returns as output the indices of only these samples that constitute an irreducible² support subsample.

Note that the notions of support subsample and support subsample function in Definitions 5.2, 5.3 are respectively referred to as compression set and compression function in (Margellos et al., 2015).

Let $K_N = |v(\Pi_N)|$ and F_N be the number of vertices and the number of facets of Π_N , respectively. It is important to emphasize that the dependence of the polytope Π_N on the multi-sample $\{\delta_i\}_{i=1}^N$ implies that both F_N and K_N are random variables that depend on $\{\delta_i\}_{i=1}^N$. Next we define the set

$$\mathcal{P}_\delta = \{\Pi \in \mathcal{P} : g(x, \delta) \leq 0, \forall x \in v(\Pi)\}$$

$$= \{\Pi \in \mathcal{P} : \Pi \subseteq X_\delta\}, \quad (5)$$

of all the non-empty, compact and convex polytopes Π that satisfy the constraint associated with the sample $\delta \in \Delta$. Note

² A support subsample $S = \{\delta_i\}_{i \in I_k} \subseteq \{\delta_i\}_{i=1}^N$ is said to be irreducible if no element can be further removed from S leaving the solution unchanged.

that if all the vertices of the polytope satisfy the inequality $g(\cdot, \delta) \leq 0$, then every point $x \in \Pi$ of the polytope satisfies it as well, since x can always be expressed as a convex combination of the polytope's vertices.

We now introduce the following theorem, which is the main result of our paper.

Theorem 6. Consider Assumption 3 and any A_N, B_N as in Definition 5. Fix $\beta \in (0, 1)$ and define the violation level $\varepsilon : \{0, \dots, N\} \rightarrow [0, 1]$ as a function such that

$$\varepsilon(N) = 1 \text{ and } \sum_{k=0}^{N-1} \binom{N}{k} (1 - \varepsilon(k))^{N-k} = \beta. \quad (6)$$

$$\text{We have that } \mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{V}(\Pi_N) > \varepsilon(F_N) \right\} \leq \beta, \quad (7)$$

where $\mathbb{P}^N = \prod_{i=1}^N \mathbb{P}$ is the product probability measure, and F_N is the number of facets of Π_N .

Proof 1. The first part of the proof closely follows that of Theorem 1 in (Grammatico et al., 2016). For a fixed multisample $\{\delta_i\}_{i=1}^N \in \Delta^N$ consider an arbitrary point $x \in \Pi_N$. Then, the following inequalities are satisfied

$$\begin{aligned} V(x) &= \mathbb{P}\{\delta \in \Delta : x \notin X_\delta\} = \mathbb{P}\{\delta \in \Delta : g(x, \delta) > 0\} \\ &\stackrel{(i)}{=} \mathbb{P}\{\delta \in \Delta : g(\sum_{j \in I_{d+1}} \lambda_j x_j, \delta) > 0\} \stackrel{(ii)}{=} \mathbb{P}\{\delta \in \Delta : \sum_{j \in I_{d+1}} \lambda_j g(x_j, \delta) > 0\} \\ &\leq \mathbb{P}\{\delta \in \Delta : \max_{j \in I_{d+1}} g(x_j, \delta) > 0\} = \mathbb{P}\left\{\bigcup_{j \in I_{d+1}} \{\delta \in \Delta : g(x_j, \delta) > 0\}\right\} \\ &\stackrel{(iii)}{\leq} \mathbb{P}\left\{\bigcup_{j=1}^{K_N} \{\delta \in \Delta : g(x_j, \delta) > 0\}\right\}. \end{aligned} \quad (8)$$

Equality (i) is derived from Caratheodory's Theorem, where the set under study is the polytope Π_N . In our case, Caratheodory's Theorem states that any arbitrary point of the polytope $x \in \Pi_N$ can be represented as a convex combination of at most $d + 1$ vertices from the set $v(\Pi_N)$, which means that there exists a subset of indices $I_{d+1} \subseteq \{1, \dots, K_N\}$ 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 an affine function of x for any given $\delta \in \Delta$ due to Assumption 3. The last inequality follows from the fact that $I_{d+1} \subseteq \{1, \dots, K_N\}$, due to the fact that $K_N \geq d + 1$ as the polytope has a non-empty interior by Assumption 3 (3). Since (8) holds for all $x \in \Pi_N$, we have that

$$\mathbb{V}(\Pi_N) = \sup_{x \in \Pi_N} V(x) \leq \mathbb{P}\left\{\bigcup_{j=1}^{K_N} \{\delta \in \Delta : g(x_j, \delta) > 0\}\right\}.$$

Therefore, for any multisample $\{\delta_i\}_{i=1}^N$ and for any cardinality (not necessarily irreducible) of the support subsample $k \in \{1, \dots, N\}$ the following inequalities are satisfied:

$$\begin{aligned} &\mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{V}(\Pi_N) > \varepsilon(k) \right\} \\ &\leq \mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{P}\left\{\bigcup_{j=1}^{K_N} \{\delta \in \Delta : g(x_j, \delta) > 0\}\right\} > \varepsilon(k) \right\} \\ &= \mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{P}\{\delta \in \Delta : \exists x \in v(\Pi_N), g(x, \delta) > 0\} > \varepsilon(k) \right\} \\ &= \mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{P}\{\delta \in \Delta : \Pi_N \not\subseteq X_\delta\} > \varepsilon(k) \right\}, \end{aligned} \quad (9)$$

where the last inequality is due to (5). Define now an algorithm A_N as in Definition 5.1, that returns the polytope confined by the feasibility region of P_N . By construction, A_N satisfies Assumption 1 of (Campi et al., 2018), since for any multisample $\{\delta_i\}_{i=1}^N$ it holds that $A(\{\delta_i\}_{i=1}^N) \in \mathcal{P}_{\delta_i}$, for all $i = 1, \dots, N$.

The satisfaction of Assumption 1 paves the way for the use of Theorem 1 of (Campi et al., 2018). In particular, Theorem 1 of (Campi et al., 2018) implies that the right-hand side of (9) can be upper bounded by β , for k being the cardinality of an irreducible support subsample (see Definition 5.1). For our case, an irreducible (in fact minimal) subsample coincides with the minimum number of facets that construct the polytope Π_N , i.e., $k = F_N = \text{rank}([H \ L])$, where matrices H, L are of appropriate dimension and constitute the H-representation of Π_N , i.e., $\Pi_N = \{x \in \mathbb{R}^d : Hx \leq L\}$. As such, we have that

$$\begin{aligned} &\mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{P}\{\delta \in \Delta : \Pi_N \not\subseteq X_\delta\} > \varepsilon(F_N) \right\} = \\ &\mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{P}\{\delta \in \Delta : \Pi_N \notin \mathcal{P}_\delta\} > \varepsilon(F_N) \right\} \leq \beta. \end{aligned} \quad (10)$$

From (9) and (10) we obtain that:

$$\mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : \mathbb{V}(\Pi_N) > \varepsilon(F_N) \right\} \leq \beta, \quad (11)$$

thus concluding the proof. \square

The cardinality of the minimal support subsample in our case coincides with the number of facets, thus circumventing the need of employing the greedy algorithm in (Campi et al., 2018). A direct consequence of Theorem 6 is that we can provide distribution-free guarantees for any NE of problem (1). To solidify this statement we introduce the following corollary

Corollary 7. Consider Assumption 3 and the setting of Theorem 6. We have that

$$\mathbb{P}^N \left\{ \{\delta_i\}_{i=1}^N \in \Delta^N : V(x_{NE}) > \varepsilon(F_N), \text{ for any } x_{NE} \text{ of (2)} \right\} \leq \beta.$$

Note that the choice of algorithm to determine a NE of the problem is arbitrary. Corollary 7 can be easily extended to provide guarantees on the probability that the set of Nash equilibria remains the same after the occurrence of a new yet unseen sample.

4. NUMERICAL EXAMPLES

We initially apply our results to a 2-dimensional example of a polytope constructed by the intersection of random halfspaces of the form $a_1 x_1 + a_2 x_2 - b \leq 0$, 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 $\delta = (a_1, a_2, b) \in \mathbb{R}^3$. To test the validity of our theoretical bounds in practice we need to compute the probability of violation $\mathbb{V}(\Pi_N)$ and compare it with the guarantees provided by Theorem 6. Let $\hat{\mathbb{V}}(\Pi_N)$ denote an empirical estimate of $\mathbb{V}(\Pi_N)$ and $\hat{V}(x)$ an empirical estimate of a point $x \in \Pi_N$. By gridding the polytope using a large enough number of points $x_r, r = 1, \dots, R$ that cover the entire polytope we have that $\hat{\mathbb{V}}(\Pi_N) = \max_{r=1, \dots, R} V(x_r)$. 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 $N = 100$ samples for the construction of the random polytope shown in Figure 1 and use 10000 test samples to compute the empirical probability of violation for each point of the grid. The highest probability of violation occurs at a polytope vertex; Lemma 1 in the extended version of the current paper (Pantazis et al., 2020) 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. A tighter relation for certain class of programs is currently under investigation.

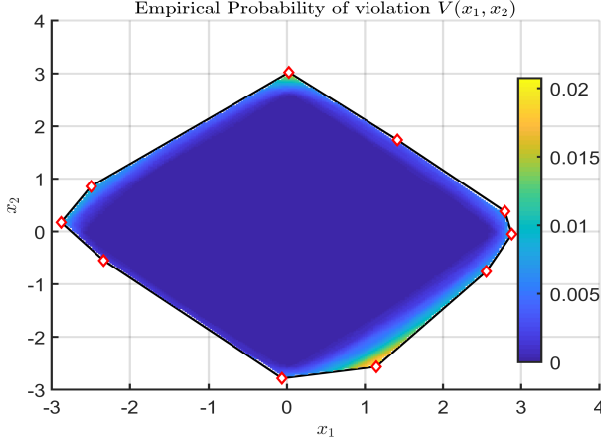


Fig. 1. The empirical probability of violation of each point of the polytope under study constructed by 100 random realizations of the affine constraints. The number of grid points is $R = 48633$, while 10000 test samples are used. Note that the higher probability of violation occurs at one of the vertices (red diamonds).




Fig. 2. Comparison of the theoretical $\varepsilon(k)$ derived from Theorem 6 (blue) and the worst-case empirical probability of violation of the entire polytope that corresponds to each k (red). Note that the non-monotonic behaviour for the empirical probability of violation is due to the fact that for any number of facets k the number of polytopes among which the worst-case probability of violation is calculated is not the same.

Finally, we validate the derived theoretical bounds against the empirical probability of violation of 50 independent realizations of polytopes. Each polytope is constructed using a different multi-sample $\{\delta_i\}_{i=1}^{2000}$. By keeping the same value for β , as in the previous case, we count the number of facets F_N 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 polytope realization, as outlined above. If there are more than one polytope among the 50 that share the same number of facets, we choose the one with the maximum empirical probability of violation. As anticipated, $\varepsilon(k)$ constitutes an upper bound for any of the computed empirical probabilities of violation.

We now revisit the EV charging control game of Section 2. Our aim is to provide guarantees on the probability that a NE satisfies the constraints of (3). We assume that the upper constraint $(\bar{x}^m(\delta_u))_{m \in \mathcal{M}} \in \mathbb{R}^{Md}$ is affected by an additive uncertainty in the form of $\delta_u \in \mathbb{R}^{Md}$, whose elements are extracted according to the probability distribution $0.3U(0,1)\mathcal{N}(1,3)$,

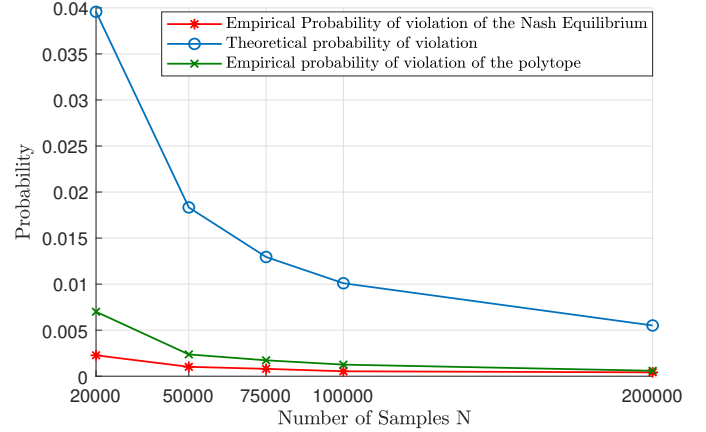


Fig. 3. The empirical probability of violation of the entire polytope $\hat{V}(\Pi_N)$ (green) and the empirical probability of violation of the computed NE returned by the NE seeking algorithm $\hat{V}(x_{NE})$ (red) versus the theoretical violation level of Theorem 6 (blue) with respect to the number of samples. Note that the blue line corresponds to the theoretical counterpart of the green one.

where $U(0,1)$ is a uniform distribution with support $[0,1]$ and $\mathcal{N}(1,3)$ is a gaussian distribution with mean 1 and standard deviation 3. As such, $\bar{x}^m(\delta_u) = \bar{x}^{nom} + \delta_u$, where \bar{x}^{nom} is a given deterministic component. The total energy $E = (E^m)_{m \in \mathcal{M}}$ of each agent at the end of charging is also affected by uncertainty, i.e., $E = (1 - \delta_e)E_{nom}$, where $\delta_e \in \mathbb{R}^M$ and its elements are extracted according to the probability $0.05\mathcal{N}(0,1)$ and $E_{nom}^m \in \mathbb{R}$ is the nominal final energy demand of each agent $m \in \mathcal{M}$ drawn from $U(10,17)$. The uncertainty vector is given by $\delta = [\delta_u, \delta_e] \in \mathbb{R}^{M(d+1)}$. The lower bound is assumed to be deterministic and, particularly, $x^i = 0$ for any $i \in \mathcal{M}$. Finally, the cost function of each vehicle $m \in \mathcal{M}$ is given by $J_m(x^m, x^{-m}) = (x^m)^T(A_0\sigma(x^m, x^{-m}) + b_0)$, where the matrix $A_0 \in \mathbb{R}^{d \times d}$ is diagonal and $\sigma(x^m, x^{-m}) = \sum_{m=1}^d x^m$. Following the work of (Fele and Margellos, 2019b) the entries $\{a_t\}_{t=1}^d$ of $A_0 = \text{diag}(\{a_t\}_{t=1}^d)$ are evaluated by rescaling a winter weekday demand profile in the UK (NationalGrid, 2019). The vector $b_0 \in \mathbb{R}^d$ is set to zero.




Fig. 4. Charging schedules of agents 2 (red solid line) and 4 (blue solid line) and their respective uncertain upper bounds. 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 die during these time instances.

To compute the NE of (3) we employed the algorithm of (Fele and Margellos, 2019b) (see Algorithm 1 therein). Setting the number of agents to $M = 10$, the number of timesteps to $d = 12$ and $\beta = 10^{-5}$, we 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. With green we illustrate the worst case empirical probability of violation for the polytope $\hat{V}(\Pi_N)$, while with red we show the empirical probability $\hat{V}(x_{NE})$ of violation for the x_{NE} returned by the algorithm of (Fele and Margellos, 2019a),(Fele and Margellos, 2019b). Note that the blue line is the theoretical counterpart of the green line and corresponds to the theoretical violation level $\varepsilon(k)$, as defined in Theorem 6. 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 6. It was also anticipated that for any multi-sample we would have that $\hat{V}(\Pi_N) \geq \hat{V}(x_{NE})$ as the former corresponds to the collective violation of all feasible points, including x_{NE} . Finally, Figure 4 illustrates the charging schedules of two of the agents and their respective uncertain upper bounds.

5. CONCLUDING REMARKS

Considering a feasibility problem under uncertain polytopic constraints, we provided probabilistic guarantees for the entire feasibility set in an *a posteriori* fashion. The importance of this result is better shown in the context of the EV-charging control problem where computationally efficient robustness certificates are obtained for the NE returned by any algorithm. Effort is being made towards extending our results to include feasibility problems subject to uncertain convex constraints and in investigating the case of uncertain games where each agent's samples are drawn from her own private uncertainty set. Finally, simulation results indicate that the highest probability of violation among all the points of the polytope occurs at (at least) one of its vertices for some problems; we aim to investigate theoretically the class of programs for which this indication is valid.

REFERENCES

- Aghassi, M. and Bertsimas, D. (2006). Robust game theory. *Math. Program.*, 107(1-2), 231–273.
- Başar, T. and Olsder, G. (1999). Dynamic non-cooperative game theory.
- Bai, D., Carpenter, T., and Mulvey, J. (1997). Making a case for robust optimization models. *Management Science*, 43, 895–907. doi:10.1287/mnsc.43.7.895.
- Birge, J.R. and Louveaux, F. (1997). *Introduction to Stochastic Programming*. Springer-Verlag, New York, NY, USA.
- Campi, M.C. and Garatti, S. (2008). The exact feasibility of randomized solutions of uncertain convex programs. *SIAM J. on Optimization*, 19(3), 1211–1230.
- Campi, M. and Calafiore, G.C. (2006). The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5), 742–753.
- Campi, M.C., Garatti, S., and Ramponi, F.A. (2018). A general scenario theory for nonconvex optimization and decision making. *IEEE Transactions on Automatic Control*, 63(12), 4067 – 4078.
- Chen, H., Li, Y., Louie, R.H.Y., and Vucetic, B. (2014). Autonomous demand side management based on energy consumption scheduling and instantaneous load billing: An aggregative game approach. *IEEE Transactions on Smart Grid*, 5(4), 1744–1754.
- Couchman, P., Kouvaritakis, B., Cannon, M., and Prashad, F. (2005). Gaming strategy for electric power with random demand. *IEEE Transactions on Power Systems*, 20(3), 1283–1292.
- Deori, L., Margellos, K., and Prandini, M. (2018). Price of anarchy in electric vehicle charging control games: When Nash equilibria achieve social welfare. *Automatica*, 96, 150 – 158.
- Fele, F. and Margellos, K. (2019a). Probabilistic sensitivity of Nash equilibria in multi-agent games: a wait-and-judge approach. *2019 IEEE Conference on Decision and Control (CDC) (to appear)*.
- Fele, F. and Margellos, K. (2019b). Probably approximately correct Nash equilibrium learning. *IEEE Transactions on Automatic Control*, 2019, submitted.
- Grammatico, S., Zhang, X., Margellos, K., Goulart, P., and Lygeros, J. (2016). A scenario approach for non-convex control design. *IEEE Transactions on Automatic Control*, 61(2), 334–345.
- Harsanyi, J.C. (1968). Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science*, 14(5), 320–334.
- Hayashi, S., Yamashita, N., and Fukushima, M. (2005). Robust Nash equilibria and second-order cone complementarity problems. *Journal of Nonlinear and Convex Analysis*, 6.
- Ma, Z., S. Callaway, D., and Hiskens, I. (2013). Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Transactions on Control Systems Technology*, 21, 67–78.
- Margellos, K., Prandini, M., and Lygeros, J. (2015). On the connection between compression learning and scenario based single-stage and cascading optimization problems. *IEEE Transactions on Automatic Control*, 60(10), 2716–2721.
- NationalGrid (2019). Mar. 2019, historical demand data. <https://www.nationalgrideso.com/balancing-data>.
- Paccagnan, D., Gentile, B., Parise, F., Kamgarpour, M., and Lygeros, J. (2019). Nash and Wardrop equilibria in aggregative games with coupling constraints. *IEEE Transactions on Automatic Control*, 64(4), 1373–1388.
- Paccagnan, D. and Campi, M. (2019). The scenario approach meets uncertain variational inequalities and game theory. *2019 IEEE Conference on Decision and Control (CDC) (to appear)*.
- Pantazis, G., Fele, F., and Margellos, K. (2020). A posteriori probabilistic feasibility guarantees for nash equilibria in uncertain multi-agent games. <https://arxiv.org/abs/2003.11307>.
- Saad, W., Han, Z., Poor, H.V., and Basar, T. (2012). Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications. *IEEE Signal Processing Magazine*, 29(5), 86–105.
- Scutari, G., Facchinei, F., Pang, J., and Palomar, D.P. (2014). Real and complex monotone communication games. *IEEE Transactions on Information Theory*, 60(7), 4197–4231.
- Singh, V.V., Jouini, O., and Lisser, A. (2016). Existence of Nash equilibrium for chance-constrained games. *Operations Research Letters*, 44(5), 640 – 644.
- Smith, M.J. (1979). The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological*, 13(4), 295–304.