

A priori data-driven robustness guarantees on strategic deviations from generalised Nash equilibria

Pantazis, Georgios; Fele, Filiberto; Margellos, Kostas

DOI

[10.1016/j.automatica.2024.111746](https://doi.org/10.1016/j.automatica.2024.111746)

Publication date

2024

Document Version

Final published version

Published in

Automatica

Citation (APA)

Pantazis, G., Fele, F., & Margellos, K. (2024). A priori data-driven robustness guarantees on strategic deviations from generalised Nash equilibria. *Automatica*, 167, Article 111746.
<https://doi.org/10.1016/j.automatica.2024.111746>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



A priori data-driven robustness guarantees on strategic deviations from generalised Nash equilibria[☆]

Georgios Pantazis^a, Filiberto Fele^{b,*}, Kostas Margellos^c

^a Delft Center for Systems and Control, TU Delft, 2628 CD Delft, Netherlands

^b Department of Systems Engineering and Automation, University of Seville, Spain

^c Department of Engineering Science, University of Oxford, OX1 3PJ, UK

ARTICLE INFO

Article history:

Received 11 April 2023

Received in revised form 24 November 2023

Accepted 15 April 2024

Available online xxxx

Keywords:

Generalised equilibrium problem

Randomised methods

Stochastic game theory

Multi-agent systems

ABSTRACT

In this paper we focus on noncooperative games with uncertain constraints coupling the agents' decisions. We consider a setting where bounded deviations of agents' decisions from the equilibrium are possible, and uncertain constraints are inferred from data. Building upon recent advances in the so called scenario approach, we propose a randomised algorithm that returns a nominal equilibrium such that a *pre-specified* bound on the probability of violation for yet unseen constraints is satisfied for an entire region of admissible deviations surrounding it—thus supporting neighbourhoods of equilibria with probabilistic feasibility certificates. For the case in which the game admits a potential function, whose minimum coincides with the social welfare optimum of the population, the proposed algorithmic scheme opens the road to achieve a trade-off between the guaranteed feasibility levels of the region surrounding the nominal equilibrium, and its system-level efficiency. Detailed numerical simulations corroborate our theoretical results.

© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The study of noncooperative games plays a significant role in a panoply of applications ranging from smart-grids (Saad, Han, Poor, & Basar, 2012) to communication (Scutari, Facchinei, Pang, & Palomar, 2014) and social networks (Acemoglu & Ozdaglar, 2011). In these setups, agents can be modelled as self-interested entities that interact with each other and make decisions based on possibly misaligned individual criteria, while being subject to constraints (local or global) that restrict the domain of their choices. Even though a variety of systems can be analysed by means of deterministic game-theoretic tools (Grammatico, Parise, Colombino and Lygeros, 2016; Paccagnan, Gentile, Parise, Kamgarpour, & Lygeros, 2019; Scutari et al., 2014), in many applications the decision making procedure is affected by uncertainty. A number of results in the literature have explicitly addressed uncertainty in a noncooperative setting. Specifically, Bopardikar, Borri, Hespanha, Prandini, and Di Benedetto (2013) follow a randomised approach for the special case of stochastic zero-sum

games. Most results rely on specific assumptions on the probability distribution (Couchman, Kouvaritakis, Cannon, & Prashad, 2005; Singh, Jouini, & Lissner, 2016) and/or the geometry of the uncertainty set (Aghassi & Bertsimas, 2006; Hayashi, Yamashita, & Fukushima, 2005; Nishimura, Hayashi, & Fukushima, 2009).

To circumvent these limitations, recent developments adopt a data-driven perspective, focusing on the connection of game theory with the so called scenario approach (Calafiore & Campi, 2006). This is based on the idea that an optimisation problem with constraints parametrised by an uncertain parameter – with fixed but possibly unknown support set and probability distribution – can be approximated by drawing samples of that parameter and solving the problem subject to the constraints generated by those samples only; this approximation is known as the *scenario program*. Standard results in the scenario approach (Calafiore, 2010; Campi & Garatti, 2008; Campi, Garatti, & Prandini, 2009) provide certificates on the probability that a new yet unseen constraint will violate the randomised solution obtained by the scenario program.

While the aforementioned results apply to uncertain *convex* optimisation problems, the works Campi and Garatti (2018) and Campi, Garatti, and Ramponi (2018) paved the way towards the provision of data-driven robustness guarantees to solutions of more general nonconvex problems. In Paccagnan and Campi (2019) and Fele and Margellos (2019, 2021), these theoretical advancements were leveraged for the first time in a game-theoretic

[☆] The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Gurdal Arslan under the direction of Editor Florian Dorfler.

* Corresponding author.

E-mail addresses: G.Pantazis@tudelft.nl (G. Pantazis), ffele@us.es (F. Fele), kostas.margellos@eng.ox.ac.uk (K. Margellos).

context, for the formulation of distribution-free probabilistic feasibility guarantees for randomised Nash equilibria. These works provide guarantees for one specific equilibrium point (often assumed to be unique); this was extended in [Pantazis, Fele, and Margellos \(2020, 2022\)](#), by providing *a posteriori* feasibility guarantees for the entire domain. Besides the game-theoretic context, alternative methodologies for set-oriented probabilistic feasibility guarantees have been proposed in the seminal works ([Alamo, Tempo, & Camacho, 2009](#); [de Farias & Roy, 2003](#)), which a priori characterise probabilistic feasibility regions constructed out of sampled constraints using statistical learning theoretic results. More recently, the so called probabilistic scaling ([Alamo, Mirasiera, Dabbene, & Lorenzen, 2019](#); [Mammarella, Mirasiera, Lorenzen, Alamo, & Dabbene, 2022](#)) has been proposed to obtain a posteriori guarantees on the probability that a polytope generated out of samples is a subset of some chance-constrained feasibility region. Following an approach similar to [Pantazis et al. \(2020\)](#), the works ([Fabiani, Margellos, & Goulart, 2020, 2022](#)) deliver tighter guarantees by focusing on variational-inequality (VI) solution sets.

The results above follow a standard approach in the game-theoretic literature, where a strict behavioural assumption – the so called *rationality* – is imposed on the players' decision making. Namely, the players are viewed as rational agents wishing to maximise their profit (expressed by some given cost function). However, studies have shown that this is unrealistic in practice ([Camerer, 2003](#); [Kahneman & Tversky, 1979](#); [Plott & Smith, 2008](#); [Wang, Saad, Mandayam, & Poor, 2016](#)) and that agents usually exhibit a *boundedly rational* behaviour ([Rubinstein, 1998](#)), i.e., their decisions can deviate from rationality due to individual biases, behavioural inertia, restricted computational power/time, etc. The consequences of this become relevant in engineering applications, as the human role in technical systems evolves beyond mere users and consumers to active agents, operators, decision-makers and enablers of efficient, resilient and sustainable infrastructures ([Lamnabhi-Lagarigue et al., 2017](#)).

To bridge this gap between real-world applications and the cognate literature, here we study games with uncertain constraints, where deviations from a *nominal* equilibrium are explicitly considered. We follow a randomised approach to approximate the coupling constraints by means of data. In this more general setting, where deviations are considered, providing guarantees for a single solution is devoid of any meaning: indeed, repetition of the game might lead to a different solution in a neighbourhood around the nominal equilibrium, irrespective of the employed dataset. Technically speaking, this renders the identification of the data samples that support the solution (cf. sample compression ([Margellos, Prandini, & Lygeros, 2015](#))) a challenging task. Focusing on the class of generalised Nash/Wardrop equilibrium seeking problems (GNEs/GWEs) ([Facchinei & Kanzow, 2010](#)), we contribute to the provision of data-driven robustness guarantees for the collection of possible deviations from the equilibrium as follows:

- (1) Adopting a scenario-theoretic paradigm, we establish a methodology for the provision of *a posteriori* probabilistic feasibility guarantees for a region around the randomised equilibrium of the game under study. This result ([Theorem 1](#)), complements ([Fabiani et al., 2020, 2022](#); [Pantazis et al., 2020, 2022](#)), that instead focus on the entire feasibility region. Focusing on a circumscribed region around a GNE/GWE allows offering tighter probabilistic bounds, while the results of [Fabiani et al. \(2020, 2022\)](#) and [Pantazis et al. \(2020, 2022\)](#), can be obtained as a limiting case of [Theorem 1](#).

- (2) We design a data-driven algorithm that returns a GNE/GWE and show that all points in a predefined admissible region surrounding it enjoy *a priori* probabilistic feasibility guarantees. This result ([Theorems 2 and 3](#)), unlike [Theorem 1](#), offers an *a priori* statement valid for a region that is tunable by the user, modelling possible deviations from a nominal equilibrium that a designer wishes to take into account when incentivising a certain operation profile.

A distinctive feature of this result is that it provides *a priori* guarantees for a set rather than single points ([Fele & Margellos, 2019, 2021](#); [Paccagnan & Campi, 2019](#)). These guarantees depend on a prespecified quantity, which in turn can affect the location of the nominal equilibrium and the size of the region for which these probabilistic guarantees hold. As such this region is tunable, unlike [Grammatico, Zhang, Margellos, Goulart and Lygeros \(2016\)](#) where *a priori* guarantees for a set of solutions are provided, but this set is not controlled by the user and could be arbitrarily narrow. Moreover, the results of [Grammatico, Zhang et al. \(2016\)](#) do not focus on games and follow a fundamentally different approach.

Furthermore, when the game under study admits a potential function – whose minimum coincides with some social welfare optimum – our methodology provides a new perspective for trading off the probabilistic feasibility of the region surrounding the nominal equilibrium and its system-level efficiency.

- (3) We propose an equilibrium seeking algorithm as the machinery to obtain a region surrounding a GNE/GWE over which the aforementioned feasibility guarantees hold. The algorithm relies on a primal-dual scheme and is inspired by seminal developments in [Facchinei and Pang \(2003\)](#). However, the mapping that characterises the algorithm updates differs from those typically encountered in the literature (e.g., see [Facchinei and Pang \(2003, Ch. 12\)](#)). This requires showing that the ad-hoc mapping enjoys certain continuity and co-coercivity properties, thus extending the proof-line of [Facchinei and Pang \(2003\)](#) (see [Lemmas 2 & 3](#), and proof of [Theorem 2](#)), a task which is interesting per se.

Our contributions compared to the cognate literature are summarised in [Table 1](#). The rest of the paper is organised as follows. In [Section 2](#) we provide fundamentals of game theory and the scenario approach. In [Section 3.1](#) we show how the feasibility guarantees for a region around the game solution can be a posteriori quantified. In [Section 3.2](#) we propose a data driven algorithm and prove its convergence to an equilibrium such that the considered neighbourhood of strategic deviations can satisfy prespecified probabilistic feasibility requirements. An illustrative example in [Section 4](#) corroborates our theoretical analysis. [Section 5](#) concludes the paper and presents future research directions. To streamline the presentation of our results, some proofs are deferred to the [Appendix](#).

2. Preliminaries

Notation: All vectors are column unless otherwise indicated. \mathbb{R}_+^n is the nonnegative orthant in \mathbb{R}^n . For an $n \times n$ matrix A , we write $A \succ 0$ ($A \succeq 0$) when $x^T A x > 0$ ($x^T A x \geq 0$), for any $x \in \mathbb{R}^n$. We denote by $\mathbf{0}_{q \times r}$ the $q \times r$ null matrix, by I_r the $r \times r$ identity matrix, and by $\mathbf{1}_r$ the vector of r ones; dimensions can be omitted when clear from the context. e_q is the unit vector whose q th element is 1 and all other elements are 0, $\|\cdot\|_p$ the p -norm operator, and $(\cdot)_r$ denotes the r th component of its vector argument. $\mathbb{B}_p(x, \rho) = \{y \in \mathbb{R}^d : \|y - x\|_p < \rho\}$ is the open p -normed ball centred at x with radius ρ ; when p is omitted, any

Table 1
Classification of results and comparison with cognate literature.

Problem class	Solution sets	Type of feasibility guarantees	Tuning	References
Affine feasibility problems	Entire feasible set	A posteriori	✗	Pantazis et al. (2020, Thm. 6)
Convex feasibility problems	Entire feasible set	A posteriori	✗	Pantazis et al. (2022, Thm. 2)
Uncertain GNEs	GNE solution set	A posteriori	✗	
Uncertain GNEs/GWEs	Subset of feasible deviations around GNE/GWE	A posteriori	✗	Theorem 1
Uncertain VIs	Unique solution	A priori/A posteriori	✗	Fele and Margellos (2019, Thm. 5), Fele and Margellos (2021, Thm. 8) and Paccagnan and Campi (2019, Cor. 1)
Convex feasibility problems	(Arbitrary) inner approx. of feasible set	A priori	✗	Grammatico, Zhang et al. (2016, Thm.2)
Uncertain GNEs/GWEs	Tunable subset of feasible deviations around GNE/GWE	A priori	✓	Theorem 3

choice of norm is valid. For a set S , $|S|$ denotes its cardinality, while 2^S denotes its power set, i.e., the collection of all subsets of S . Finally, given $D \succ 0$, $\text{proj}_{K,D}[x] := \arg \min_{y \in K} (y - x)^T D (y - x)$ is the skewed projection of x onto the set K .

2.1. Games with uncertain constraints

We consider a population of agents with index set $\mathcal{N} = \{1, \dots, N\}$. The decision vector x_i of each agent $i \in \mathcal{N}$ takes value 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 $x_{-i} \in \mathbb{R}^{n(N-1)}$ comprises all agents' decisions except for those of agent i . In our setup, the cost incurred by agent $i \in \mathcal{N}$ is expressed by a real-valued function $J_i(x_i, x_{-i})$ that depends on local decisions as well as on the decisions from other agents $j \in \mathcal{N} \setminus \{i\}$. In the following, with a slight abuse of notation, we can exchange x for (x_i, x_{-i}) to single out agent i 's decision from the global decision vector. Furthermore, we consider *uncertain* constraints coupling the agents' decisions. These can be expressed in the form¹

$$C_\delta = \{x \in X : g(x, \delta) \leq 0\}, \quad \delta \in \Delta, \quad (1)$$

where $g : \mathbb{R}^{nN} \times \Delta \rightarrow \mathbb{R}$ depends on some uncertain parameter δ taking values in a support set Δ according to a probability measure \mathbb{P} .

Feasible collective decisions under this setup can be found by letting every agent $i \in \mathcal{N}$ solve the following optimisation program, where the decisions x_{-i} of all other agents are given,

$$G : \min_{x_i \in X_i} J_i(x_i, x_{-i}) \quad \left\{ \begin{array}{l} \text{subject to } x_i \in \bigcap_{\delta \in \Delta} C_{\delta}^i(x_{-i}) \end{array} \right\} \quad \forall i \in \mathcal{N}; \quad (2)$$

here, $C_{\delta}^i(x_{-i}) = \{x_i \in X_i : g(x_i, x_{-i}, \delta) \leq 0\}$ is the projection of the coupling constraint on X_i for fixed x_{-i} and uncertain realisation $\delta \in \Delta$. The collection of coupled optimisation programs in (2) for all $i \in \mathcal{N}$ constitutes an *uncertain noncooperative game*; we denote it as G .

Note that (2) follows a worst-case paradigm, taking into account all possible coupling constraints that can be realised by variations of the uncertain parameter $\delta \in \Delta$. This can render the solutions of G rather conservative. Furthermore, it is in general not possible to compute a solution for G without an accurate knowledge of, and/or additional assumptions on, the support set Δ and the probability distribution \mathbb{P} . To circumvent these limitations, we follow a data-driven paradigm and approximate G by means of a finite number of samples drawn from Δ , namely

the K -multisample $\delta_K = (\delta_1, \dots, \delta_K) \in \Delta^K$. In the sequel, we hold on to the standing assumption that these samples are independent and identically distributed (i.i.d.). Apart from this, no other knowledge on the support set Δ and the probability distribution \mathbb{P} of the uncertain parameter is required. Then, for a given multi-sample δ_K , (2) can be rewritten as

$$G_K : \min_{x_i \in X_i} J_i(x_i, x_{-i}) \quad \left\{ \begin{array}{l} \text{subject to } x_i \in \bigcap_{k=1}^K C_{\delta_k}^i(x_{-i}) \end{array} \right\} \quad \forall i \in \mathcal{N}. \quad (3)$$

Instead of considering all possible uncertainty realisations $\delta \in \Delta$ as in (2), we let the data encoded in δ_K lead agents to their decision by solving (3). We refer to the collection of coupled optimisation programs in (3) as the *scenario game* G_K (the subscript K implies dependence on the drawn multi-sample δ_K). Under standard assumptions, a solution to the scenario game G_K exists and the problem is, in contrast to G , tractable using state-of-the-art equilibrium seeking algorithms.

2.2. Variational inequalities and game equilibria

Notably – under certain assumptions detailed next – solutions to the game G_K can be retrieved as solutions to a variational inequality (VI), for specific choices of the mapping $F : X \rightarrow \mathbb{R}^{nN}$ (Facchinei & Kanzow, 2010, Thm 3.9):

$$VI_K : \text{Find } x^* \in \Pi_K \text{ such that} \quad (x - x^*)^T F(x^*) \geq 0 \text{ for any } x \in \Pi_K, \quad (4)$$

where $\Pi_K := X \cap \bigcap_{k=1}^K C_{\delta_k}$ denotes the problem domain. A classic game solution concept, which encounters wide application in the literature, is the Nash equilibrium (NE) (Nash, 1950). At a NE, no agent can decrease their cost by unilaterally changing their decision. Formally, this can be stated as follows.

Definition 1. A point $x^* = (x_i^*, x_{-i}^*) \in \Pi_K$ is called a generalised Nash equilibrium (GNE) of G_K if, for all $i \in \mathcal{N}$,

$$J_i(x_i^*, x_{-i}^*) \leq J_i(y_i, x_{-i}^*), \quad \forall y_i \in X_i \cap \bigcap_{k=1}^K C_{\delta_k}^i(x_{-i}^*).$$

For our analysis, we rely on the following conditions:

Assumption 1. For all $i \in \mathcal{N}$, $J_i(x_i, x_{-i})$ is convex and continuously differentiable in x_i for any fixed x_{-i} .

Assumption 2.

- (1) For any multi-sample $\delta_K \in \Delta^K$, the domain Π_K is non-empty.

¹ This formulation can describe deterministic and/or local constraints as special cases.

- (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 $a : \Delta \rightarrow \mathbb{R}^N$ and $b : \Delta \rightarrow \mathbb{R}$.

Note that convexity of the cost function with respect to the agent's local decision is crucial for the design of tailored algorithms with theoretical convergence guarantees for Nash equilibrium seeking. Under these assumptions, we can determine a GNE as in [Definition 1](#) by solving (4) with

$$F(x) = F_{NE}(x) := \begin{bmatrix} \nabla_{x_1} J_1(x_1, x_{-1}) \\ \vdots \\ \nabla_{x_N} J_N(x_N, x_{-N}) \end{bmatrix}. \quad (5)$$

A class of problems of common interest can be modelled by the so called *aggregative games* ([Acemoglu & Jensen, 2013](#); [Jensen, 2010](#); [Kukushkin, 2004](#)), where the cost incurred by agents depends on some aggregate measure – typically the average – of the decision of the entire population. Such a cost can be expressed in (3) by the function $J_i(x_i, \sigma(x))$, where the aggregate $\sigma : \mathbb{R}^{nN} \rightarrow \mathbb{R}^n$ is defined as the mapping $x \mapsto \frac{1}{N} \sum_{i=1}^N x_i$. A solution frequently linked to this class of games is the Wardrop equilibrium (WE), a concept akin to the NE but specifically defined in the context of transportation networks ([Beckmann, McGuire, & Winsten, 1956](#)). The variational WEs of G_K can be expressed by using $F(x) = F_{WE}(x) := [\nabla_{x_i} J_i(x_i, z)]_{i \in \mathcal{N}}|_{z=\sigma(x)}$; notice that in this case the second argument of J_i is fixed and set to $\sigma(x)$, consistently with the notion of WE where agents neglect the impact of their decision on others.

We restrict the considered class of variational mappings as follows:

Assumption 3. The mapping F is

- (1) α -strongly monotone, i.e., $(x - y)^T (F(x) - F(y)) \geq \alpha \|x - y\|^2$ for any $x, y \in X$,
- (2) L_F -Lipschitz continuous, i.e., $\|F(x) - F(y)\| \leq L_F \|x - y\|$ for any $x, y \in X$.

[Assumptions 1](#) and [3](#) are standard in the game-theoretic literature ([Facchinei & Pang, 2003](#); [Scutari et al., 2014](#)). [Assumption 2](#) is relatively mild; the affine form of the constraints is exploited in the proposed algorithm (see [Section 3](#)) for the convergence to an equilibrium bearing the desired robustness properties.

We point out that in general only a subset of solutions to G_K can be retrieved through (4): these are referred to as *variational equilibria*, and enjoy favourable properties over nonvariational ones, as with the former the coupling constraints' burden is equally split among agents ([Harker, 1991](#); [Kulkarni & Shanbhag, 2012](#)). The following lemma, adapted from [Facchinei and Pang \(2003, Thm. 2.3.3\)](#), formalises the connection between the solutions to VI_K and the GNEs (or GWEs) of G_K .

Lemma 1. Under [Assumptions 1–3](#), VI_K has a unique solution that is also an equilibrium of G_K .

For the considered class of VIs, several algorithms from the literature can be employed to obtain a variational equilibrium of G_K ; see, e.g., [Facchinei and Kanzow \(2010\)](#) and [Paccagnan et al. \(2019\)](#). We remark that, even if not explicitly shown for ease of notation, any solution x^* to G_K is itself a function of the drawn multisample $\delta_K \in \Delta^K$. Probabilistic feasibility guarantees for the unique solution of VI_K can then be provided both in an *a priori* and *a posteriori* fashion by resorting to the results in [Fele and Margellos \(2019, 2021\)](#) and [Paccagnan and Campi \(2019\)](#). However, these results are tailored to the provision of probabilistic feasibility guarantees for a single point (namely the solution of a VI): any strategic deviation from the equilibrium is not covered by such guarantees. We cover this issue in [Section 3](#). First, we provide some background on the scenario approach.

2.3. Basic concepts in the scenario approach

A fundamental notion in the scenario approach is the *probability of violation* of an uncertain constraint.

Definition 2.

- (1) The probability of violation $\mathcal{V} : \mathbb{R}^{nN} \rightarrow [0, 1]$ of a point $x \in \Pi_K$ is defined as the probability that a new yet unseen sample $\delta \in \Delta$ will give rise to a constraint C_δ (as defined in (1)) such that $x \notin C_\delta$, i.e., $\mathcal{V}(x) := \mathbb{P}\{\delta \in \Delta : x \notin C_\delta\}$.
- (2) The probability of violation $\mathbb{V} : 2^{\mathbb{R}^{nN}} \rightarrow [0, 1]$ of a set $S \subseteq \Pi_K$ is defined as the worst-case \mathcal{V} among all the points in S , i.e., $\mathbb{V}(S) = \sup_{x \in S} \mathcal{V}(x) = \sup_{x \in S} \mathbb{P}\{\delta \in \Delta : x \notin C_\delta\}$.

A data-driven decision-making process can be formally characterised by a mapping – the *algorithm* – that takes as input the data encoded by the samples and returns a set of decisions.

Definition 3. An algorithm is a function $\mathcal{A} : \Delta^l \rightarrow 2^{\mathbb{R}^{nN}} \times \mathbb{R}^{nN}$ that takes as input an l -multisample and returns the pair $(x^*, S_l^*(x^*))$, namely, a point x^* and a solution set S_l^* parametrised by x^* .

In the setting we consider, we have $x^* \in S_l^*(x^*)$; however, this ought not to be the case in general. In the following, we interpret the above definition as context-dependent, in that the size l of the input multisample is admitted to vary—all else remaining fixed for a given algorithm \mathcal{A} .

A key notion, strongly linked to that of algorithm, is the *minimal compression set* ([Margellos et al., 2015](#)). This concept springs from the observation that typically only a subset of the sampled data is relevant to a decision or set of decisions, and all other samples are redundant.

Definition 4 (Compression Set). Consider an algorithm \mathcal{A} as in [Definition 3](#). A subset of samples $I \subseteq \delta_K$ is called a *compression* for $\mathcal{A}(\delta_K)$ if $\mathcal{A}(I) = \mathcal{A}(\delta_K)$.² As multiple subset of samples can exist that fulfil this property, the ones with the minimal cardinality are called *minimal compression sets*.

If we feed the algorithm with the set of samples corresponding to a compression, then the same decision will be returned as when we feed the algorithm with the entire multi-sample. As established in [Margellos et al. \(2015\)](#), the compression set is related to the notion of support samples ([Campi & Garatti, 2008](#)) and that of essential constraints ([Calafiore, 2010](#)). Under certain non-degeneracy assumptions these concepts coincide.

3. Probabilistic feasibility of sets around equilibria

3.1. A first a posteriori result

Returning to the scenario game G_K in (3), we now consider a more general setup where agents are allowed to deviate from x^* following, e.g., unmodelled changes in their cost functions; while we suppose that these deviations are feasible with respect to the local constraints, we want to study the feasibility as regards the coupling constraints obtained through sampling. Specifically, the region in which agents' strategies can deviate from the nominal equilibrium is assumed to lie within a predefined open ball $\mathbb{B}(x^*, \rho)$, where $\rho > 0$ is a fixed radius that denotes the maximum possible distance of agents' deviations from x^* ; the latter

² With some abuse of notation, in the remainder the symbol δ_K is interpreted as either the i.i.d. sample vector $\delta_K \in \Delta^K$, or the set comprising its components, i.e., $\delta_K = \{\delta_1, \dots, \delta_K\} \subseteq \Delta$, depending on the context.




Fig. 1. Region S_K^* (in green) obtained as the intersection of the set of deviations $\mathbb{B}_\infty(x^*, \rho)$ around the equilibrium x^* (red dot) with the domain Π_K . The samples producing the constraints in blue are in the compression set of x^* , while those associated with the red constraint are in the compression set of S_K^* ; discarding these does not change x^* . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

is assumed to be unique as per [Lemma 1](#). As such, the region of interest is $S_K^* = \Pi_K \cap \mathbb{B}(x^*, \rho)$.

This is depicted in [Fig. 1](#) using the ∞ -norm (any other norm could have been used instead): an algorithm \mathcal{A} (see [Section 2.3](#)) takes as input a multi-sample δ_K and returns the region S_K^* around the solution $x^* \in \mathbb{R}^2$ of a game with two players whose decisions are defined as scalar quantities. For this pictorial example, Π_K is shaped exclusively by sampled coupling constraints. Any compression set as per [Definition 4](#) for \mathcal{A} must be associated with the solid blue constraints (these form a compression for x^*), and with the dashed red constraint that intersects $\mathbb{B}(x^*, \rho)$ —as its removal would change S_K^* .

We can quantify the number of samples that form a compression set for the algorithm that returns S_K^* in an *a posteriori* fashion as established in [Theorem 1](#). To this end, for a fixed confidence $\beta \in (0, 1)$, let the violation level be defined as a function $\epsilon : \{0, \dots, K\} \rightarrow [0, 1]$ satisfying ([Campi et al., 2018, Eq. \(7\)](#))

$$\epsilon(K) = 1 \quad \text{and} \quad \sum_{i=0}^{K-1} \binom{K}{i} (1 - \epsilon)^{K-i} = \beta. \quad (6)$$

Theorem 1. Under [Assumptions 1–3](#), let algorithm \mathcal{A} return a pair $(x^*, S_K^*(x^*))$. Fix a confidence parameter $\beta \in (0, 1)$ and a violation level $\epsilon : \{0, \dots, K\} \rightarrow [0, 1]$ that satisfies (6). We have that

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

where s^* is the number of samples that form a compression set for the equilibrium mapping x^* and M the number of uncertain constraints of Π_K that intersect S_K^* .

Proof. Let (x^*, S_K^*) be the solution returned by \mathcal{A} for some given δ_K , according to [Definition 3](#). We aim at determining a compression set for $\mathcal{A}(\delta_K)$, and use its cardinality to reach the theorem's conclusion by means of [Theorem 2](#) in [Pantazis et al. \(2022\)](#). This would be the union of: (i) the samples that form a compression set for x^* — i.e., solving the problem using only these would result in the same equilibrium obtained by using all samples —, and (ii) any other sample (not in the compression set of x^*) whose removal can still lead to a change of the region S_K^* .

Case (i): Determining a (possibly non-minimal) compression set for x^* can be achieved, as suggested in [Campi et al. \(2018\)](#), by progressively removing samples till a subset that leaves the solution unchanged is determined. We denote its cardinality by s^* . With reference to [Fig. 1](#), this set would be associated with the blue constraints active at x^* . Case (ii): We need to count the samples whose removal does not change x^* but yields a larger region S_K^* (red constraint in [Fig. 1](#)). Their number can be upper bounded by the M facets of Π_K that intersect S_K^* .

Hence, the number of samples that form a compression set for $\mathcal{A}(\delta_K)$ is bounded by $s^* + M$. Existence of a compression set I

with a bound on its cardinality is sufficient for the application of [Theorem 2](#) in [Pantazis et al. \(2022\)](#). The fact that for the minimal compression set $|I^*| \leq |I| \leq s^* + M$ always holds leads then to the statement of this theorem. ■

It is important to stress that the application of [Theorem 1](#) is agnostic on the choice of the equilibrium seeking algorithm. To use the result of [Theorem 1](#), one needs to quantify (an upper bound of) the number of samples s^* that form a compression set for x^* and (an upper bound of) the number M of coupling constraints that correspond to facets of S_K^* . While $s^* \leq nN$ under [Assumptions 1](#) and [3](#),³ an upper bound for M can in general only be achieved *a posteriori*, i.e., once δ_K is sampled. In the next section we show how we could obtain *a priori* bounds for the same quantity.

3.2. A priori probabilistic certificates

Consider the scenario game G_K and suppose that bounded deviations from the solution are allowed. We model such deviations as a ball of radius ρ around the equilibrium, as in [Section 3.1](#). In contrast to the *a posteriori* nature of the result therein, our goal here is to achieve an *a priori* bound. Namely, we aim at establishing the main statement of [Theorem 1](#) with a *prespecified* violation level, which does not depend on the given multi-sample δ_K . In other words, we seek a statement — holding with known confidence — of the form $\mathbb{V}(\Pi_K \cap \mathbb{B}(x^*, \rho)) < \bar{\epsilon}$, with $\bar{\epsilon} \in (0, 1)$ *a priori* fixed.

To achieve this we build upon the previous conclusions, which expose a link between the probability of constraint violation and the number M of facets of Π_K (each originated from some uncertainty sample) that $\mathbb{B}(x^*, \rho)$ intersects. In particular, a monotonic relationship follows from (6): the smaller M the better, i.e., less conservative, the theoretical feasibility guarantees on constraint violation for the strategies belonging to the feasible region S_K^* surrounding the equilibrium. Also, a smaller value of M can result in a larger region for which the guarantees of [Theorem 1](#) hold—due to a smaller portion of $\mathbb{B}(x^*, \rho)$ being cut off by intersection with Π_K . This motivates us to study the role of M as a modulating parameter for the robustness of the feasibility certificates offered for the region S_K^* , as well as the extent of deviation from the nominal equilibrium covered by such certificates.

3.2.1. GNE-seeking algorithm with a priori robustness guarantees

We consider an iterative scheme to determine a solution of VI_K in (4). In particular, since the problem involves coupling constraints, we build our [Algorithm 1](#) upon a primal–dual scheme, where constraint satisfaction is achieved by the use of Lagrange multipliers; similar developments hold for both GNE and GWE problems. To this end, we define the augmented vector $y := (x, \mu) \in \mathbb{R}^{nN+m}$ by stacking the global decision vector x and the Lagrange multipliers $\mu = (\mu_\ell)_{\ell=1}^m \in \mathcal{M} \subseteq \mathbb{R}_+^m$. The set \mathcal{M} denotes the domain of μ ; in the sequel we impose some structure on \mathcal{M} once some necessary theoretical ingredients are introduced. As deterministic constraints do not play a role in the evaluation of the robustness guarantees, suppose for ease of exposition that Π_K only comprises uncertain coupling constraints. Let $A \in \mathbb{R}^{m \times nN}$ and $b \in \mathbb{R}^m$ such that

$$\Pi_K = \{x \in X : Ax \leq b\}, \quad (7a)$$

$$\|a_\ell\|_2 = 1, \quad \ell = 1, \dots, m, \quad (7b)$$

³ By arguments similar to those in [Pantazis et al. \(2022\)](#) and [Schildbach, Fagiano, and Morari \(2012\)](#) it can be shown that a tighter bound $s^* \leq n$ holds for the game G_K in case coupling constraints only concern the aggregate variable.

Algorithm 1 A priori robust GNE seeking algorithm**Require:** $y^{(0)}, \rho \in \mathbb{R}_+, \delta_K \in \Delta^K, M \leq m, \xi \geq 0$

```

1:  $\kappa \leftarrow 0$ 
2: repeat
3:    $y^{(\kappa+1)} \leftarrow \text{proj}_{X \times \mathcal{M}, D} [y^{(\kappa)} - D^{-1}T(y^{(\kappa)}, \rho, M)]$ 
4:    $\kappa \leftarrow \kappa + 1$ 
5: until  $\|y^{(\kappa+1)} - y^{(\kappa)}\| \leq \xi$ 
6:  $y^* \leftarrow y^{(\kappa+1)}$ 
7: return  $y^*$  and  $\Pi_K \cap \mathbb{B}(x^*, \rho)$ 

```

where a_ℓ^\top denotes the ℓ th row of A . Eq. (7) is the irredundant H -representation of the polytopic feasibility region Π_K defined in (4), where the rows of matrix A are unit vectors. Property (7b) is key to the second statement in Lemma 2. It entails no loss of generality, since for any A, b forming an equivalent H -representation of Π_K , (7) can be obtained by normalising each row of A and the corresponding component of b by the row-vector norm. Thus, the pair (A, b) encodes the set of randomised coupling constraints that constitute facets of Π_K .⁴

The main step of Algorithm 1 (step 3) is a projected gradient descent (ascent) update for $x(\mu)$ through the mapping $T : \mathbb{R}^{nN+m+1} \times \mathbb{N} \rightarrow \mathbb{R}^{nN+m}$ given by

$$T(y, \rho, M) := \begin{bmatrix} F(x) + A^\top \mu \\ -(Ax - b + Q(\mu, M)\rho) \end{bmatrix}. \quad (8)$$

T follows from the primal-dual conditions of the game solution; see Facchinei and Kanzow (2010, Sec. 4.2), Facchinei and Pang (2003, Sec. 1.4.1). F is the pseudo-gradient mapping defined as in Section 2.2, A, b are as in (7), and $\rho := c\rho \mathbf{1}_m$, where c is a constant scaling factor (see Section 3.2.2) and M a nonnegative integer. In the second block-row of (8), the $m - M$ least relevant (based on the multipliers value) coupling constraints are tightened by an amount $c\rho$ through the mapping $Q : \mathbb{R}_+^m \times \mathbb{N} \rightarrow \{0, 1\}^{m \times m}$. Finally, the asymmetric projection matrix $D \succ 0$ includes the step-size parameter $\tau > 0$ and is defined as

$$D := \begin{bmatrix} \frac{1}{\tau} I_{nN} & 0 \\ -2A & \frac{1}{\tau} I_m \end{bmatrix}. \quad (9)$$

Note that the constraint tightening performed in the second block-row of T is equivalent to preventing $\mathbb{B}(x^*, \rho)$ from intersecting these constraints. In other words, Q ensures that the number of facets of Π_K intersecting $\mathbb{B}(x^*, \rho)$ is at most M , which in turn enables to obtain an *a priori* estimate of the number of samples that form a compression for S_K^* and hence on $\mathbb{V}(\Pi_K \cap \mathbb{B}(x^*, \rho))$; this is formalised by Theorems 2 and 3. Since $m - M$ coupling constraints are tightened, smaller values for M can result in a more robust and possibly larger region S_K^* ; however, they can also move the location of the nominal equilibrium x^* to a somewhat less efficient point towards the interior of Π_K . As we will demonstrate numerically in the sequel, this is the case with potential games (Facchinei, Piccialli, & Sciandrone, 2011).

3.2.2. Constraint tightening via mapping Q

We define the mapping Q as

$$Q(\mu, M) := P^\top(\mu)R(M), \quad (10)$$

where

⁴ Formally, $A : \Delta^K \rightarrow \mathbb{R}^{m \times nN}$ and $b : \Delta^K \rightarrow \mathbb{R}^m$ are mappings from the K -multisample to the space of real $m \times nN$ matrices and m -dimensional vectors, respectively.

- $P : \mathbb{R}^m \rightarrow \{0, 1\}^{m \times m}$ returns a permutation matrix such that $P(\mu)\mu$ is the vector composed by the elements of μ arranged in decreasing order.
- $R : \mathbb{N} \rightarrow \{0, 1\}^{m \times m}$ takes as input the number of coupling constraints $M \leq m$ we allow $\mathbb{B}(x^*, \rho)$ to intersect with and returns as output the matrix

$$R(M) = \begin{bmatrix} \mathbf{0}_{m \times M} & \mathbf{0}_{M \times m-M} \\ & I_{m-M} \end{bmatrix}. \quad (11)$$

Compatibly with the definition of $P(\cdot)$, $R(M)P(\mu)\rho = (\mathbf{0}_M^\top \ c\rho \mathbf{1}_{m-M}^\top)^\top = R(M)\rho$, where the last equality holds since all components of ρ are equal.

As discussed in Section 3.2.1, $Q(\cdot, M)$ allows to tighten the constraints corresponding to the smallest $m - M$ multipliers. For this, we use the radius of the sphere that circumscribes $\mathbb{B}(x^*, \rho)$. This is $\rho_\ell = c\rho \|a_\ell\|_2 = c\rho$, where the last equality is due to (7b); depending on the choice of norm, $c = 1$ if $\mathbb{B}(\cdot, \rho)$ is expressed by a p -norm with $p \leq 2$, $c = \sqrt{n}$ otherwise. Conversely, at most M constraints can intersect $\mathbb{B}(x^*, \rho)$ upon convergence of the algorithm. Let $\mathcal{L}(M) \subseteq \{1, \dots, m\}$ contain the indices of the M largest multipliers. Then, $\ell \in \mathcal{L}(M) \Leftrightarrow (Q(\mu, M)\rho)_\ell = 0$, and the second block row of T in (8) expresses

$$\begin{cases} a_\ell^\top x \leq b_\ell & \text{if } (Q(\mu, M)\rho)_\ell = 0, \\ a_\ell^\top x \leq b_\ell - c\rho & \text{if } (Q(\mu, M)\rho)_\ell = c\rho. \end{cases} \quad (12)$$

Illustrative example: Let Π_K result from the intersection of 3 hyperplanes and allow $\mathbb{B}(\cdot, \rho)$ to intersect at most $M = 1$ of them. From (11), $R(M) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. At iteration κ of Algorithm 1, let the multiplier vector $\mu^{(\kappa)} = (\mu_\ell^{(\kappa)})_{\ell=1}^3$ be such that $\mu_2^{(\kappa)} > \mu_1^{(\kappa)} > \mu_3^{(\kappa)}$. Then, $P(\mu^{(\kappa)}) = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ is the permutation matrix such that $P(\mu^{(\kappa)})\mu^{(\kappa)} = (\mu_2^{(\kappa)} \ \mu_1^{(\kappa)} \ \mu_3^{(\kappa)})^\top$. So $Q(\mu^{(\kappa)}, M)\rho = P^\top(\mu^{(\kappa)})R(M)\rho = (c\rho \ 0 \ c\rho)^\top$, where $P^\top(\cdot)$ applies the correct ordering to the vector $R(M)\rho$. Suppose $\mu_2^{(j)} > \mu_1^{(j)} > \mu_3^{(j)}$ holds for all $j \geq \kappa$. Then, at convergence, it follows from (12) that $\mathbb{B}(x^*, \rho)$ will not be intersecting the constraints associated to μ_1 and μ_3 , whereas it could be intersecting the hyperplane associated to μ_2 .

3.3 Convergence analysis and main result

Due to Q , mapping T is discontinuous on $X \times \mathbb{R}^m$. To circumvent this, we restrict the multipliers to the set \mathcal{M} on which we impose some structure granting continuity of T on $X \times \mathcal{M}$. To this end, let $\mathcal{Z} := [\zeta, +\infty) \cup \{0\}$, for some small $\zeta > 0$, i.e., $\mathcal{Z} \subset \mathbb{R}$ contains all nonnegative scalars which take value greater than ζ when nonzero.

Assumption 4. Let Λ be an arbitrarily large compact set. Λ admits the form

$$\mathcal{M} := \{\mu \in \Lambda : (P(\mu)\mu)_{\ell+1} \leq (P(\mu)\mu)_\ell - \zeta, \forall \ell = 1, \dots, m-1\} \cap \mathcal{Z}^m. \quad (13)$$

Recalling that $P(\mu)\mu$ rearranges the multipliers in descending order, the set \mathcal{M} contains all vectors where the difference between every pair of strictly positive components—and the distance of the smallest of these from zero—is no less than ζ . We note that (13) is the union of $q = m! + m + 1$ disjoint convex subsets of \mathbb{R}_+^m , each of which we denote as \mathcal{M}_j , i.e., $\mathcal{M} = \bigcup_{j=1}^q \mathcal{M}_j$; Fig. 2 illustrates this set for $m = 3$. It is therefore possible to compute the projection in line 3 of Algorithm 1 by, e.g., projecting on \mathcal{M}_j , for $j = 1, \dots, q$, and then setting $y^{(\kappa+1)}$ to be the solution among these that results in the minimum distance




Fig. 2. Domain \mathcal{M} of the Lagrange multipliers associated to the coupling constraints, for $\zeta = 0.2$ and $m = 3$. This results in $q = 10$ convex subsets, including the origin and a portion of the axes.

from $y^{(k)} - D^{-1}T(y^{(k)}, \rho, M)$. Still, the projection on \mathcal{M} can be computationally intensive if q is large.

Imposing on \mathcal{M} the structure of (13) endows T with the desired nonexpansiveness properties that are exploited in the proof of Lemma 3. In the numerical implementation of the algorithm, ensuring $\mu \in \mathcal{M}$ can possibly introduce small perturbations in the multipliers – compared to standard formulations where $\mu \in \mathbb{R}_+^m$ – which in turn could produce a slight violation of the constraints (this can be controlled through the magnitude of ζ). We note that \mathcal{M} is compact by construction due to the intersection with the compact set Λ in (13) which can, however, be arbitrarily large thus not impacting the result numerically. Compactness is used in the proof of Theorem 2; Remark 1 discusses cases where this requirement can be lifted.

Lemma 2. Define T as in (8)–(11), where A, b satisfy (7). Then, for any $\mu, \mu' \in \mathcal{M}$, $\mu \neq \mu'$, there exists an integer $0 \leq h \leq M$ such that

$$(\mu - \mu')^\top (Q(\mu, M) - Q(\mu', M))\rho \leq -h\zeta c\rho. \quad (14)$$

Lemma 3. Consider T as in (8)–(11), where A, b satisfy (7) and $\mathcal{M} = \bigcup_{j=1}^q \mathcal{M}_j$ as in (13). For each $j = 1, \dots, q$, let $VI(X \times \mathcal{M}_j, T)$ denote the VI problem defined by the map T restricted to the subdomain $X \times \mathcal{M}_j$. Under Assumptions 1–3 the following holds:

- (1) T is continuous on $X \times \mathcal{M}$.
- (2) Let D as in (9) and set $\tau > 0$ such that

$$\tau < \min \left\{ \frac{-L_F^2 + \sqrt{L_F^4 + 4\alpha^2 \|A\|^2}}{2\alpha \|A\|^2}, \frac{-\rho(1 + \|A\|^2) + \sqrt{\rho^2(1 + \|A\|^2)^2 + 16\zeta^2 \|A\|^2}}{4\zeta \|A\|^2} \right\}. \quad (15)$$

Then, for any $j = 1, \dots, q$, Algorithm 1 converges to a solution of $VI(X \times \mathcal{M}_j, T)$, when the gradient step in line 3 is projected on the corresponding subdomain, for any $y^{(0)} \in X \times \mathcal{M}_j$.

Continuity of the mapping is essential for the theoretical convergence of Algorithm 1. The second part of Lemma 3 provides an admissible range of values for τ such that Algorithm 1 converges to a solution of $VI(X \times \mathcal{M}_j, T)$ if at each iteration the projection in line 3 is performed on the (convex) subdomain $\mathcal{M}_j \subset \mathcal{M}$, $j \in \{1, \dots, q\}$. The stepsize τ is chosen such that conditions

standard in NE seeking are satisfied and oscillations among multiple equilibria are avoided. Still, we are interested in establishing convergence on the entire domain \mathcal{M} , so at each iteration the projected solution might belong to a different subdomain. This does not trivially follow from the second part of Lemma 3; therefore, by Lemmas 2 and 3 we establish an additional condition on τ such that Algorithm 1 retrieves a solution of $VI(X \times \mathcal{M}, T)$.

Theorem 2. Consider Assumptions 1–4. Fixed $0 \leq M \leq m$, assume the domain Π_K is nonempty for any of the $\binom{m}{M}$ combinations of constraints tightened as in (12). Let $D > 0$ be defined as in (9), where τ satisfies (15) and

$$\tau < \frac{-(\bar{C} + \bar{R}) + \sqrt{(\bar{C} + \bar{R})^2 + 2\zeta \bar{R}}}{2\bar{R}}, \quad (16)$$

where $\bar{R} = \max\{\sup_{x \in X} \sup_{\mu \in \mathcal{M}} \|2A(F(x) + A^\top \mu)\|, \sup_{x \in X} \|Ax - b\|\}$ and $\bar{C} = c\rho\sqrt{m - M}$.

Then Algorithm 1 converges to a solution of $VI(X \times \mathcal{M})$ for any initial condition $y^{(0)} \in X \times \mathcal{M}$.

Note that as $\mu^{(k)} \rightarrow \mu^*$, we have $Q(\mu^{(k)}) \rightarrow Q(\mu^*) =: Q^*$. Then, the solution returned by Algorithm 1 is the equilibrium of a variant of G_K with $m - M$ tightened constraints (follows from (12) with $Q(\mu)$ replaced by Q^*).

Remark 1 (Relaxing Compactness). Theorem 2 still holds when $\Lambda = \mathbb{R}^m$ in the definition of \mathcal{M} in (13) if for all multi-samples, (i) A is full row-rank, or (ii) all elements of A are positive.

(i) To show this, consider mapping T and matrix D in (9). The multipliers' update involves projecting (weighted according to D) on \mathcal{M} , the term

$$\mu^{(k)} - 2\tau^2 A(F(x^{(k)}) + A^\top \mu^{(k)}) + \tau(Ax^{(k)} - b + Q(\mu^{(k)}, M)\rho). \quad (17)$$

Since X is compact, there exists a subsequence $\{\kappa_i\}_{i \in \mathbb{N}}$ such that $\lim_{i \rightarrow \infty} \kappa_i = \infty$, $\lim_{i \rightarrow \infty} x^{(\kappa_i)} = \bar{x}$, for some $\bar{x} \in X$. It suffices to show that the sequence of multipliers $\{\mu^{(\kappa_i)}\}_{i=1}^\infty$ remains bounded (all arguments in the proof of Theorem 2 from (C.1) onwards remain unaltered).

For the sake of contradiction, assume that there exists at least one element of $\mu^{(\kappa_i)}$ that tends to infinity across the considered subsequence. Let then $\mu^{(\kappa_i)} = (\mu_{\infty}^{(\kappa_i)}, \mu_F^{(\kappa_i)})$, where based on our contradiction hypothesis $\lim_{i \rightarrow \infty} \|\mu_{\infty}^{(\kappa_i)}\| = \infty$ while $\lim_{i \rightarrow \infty} \|\mu_F^{(\kappa_i)}\| < \infty$. (Taking the first elements of $\mu^{(\kappa_i)}$ to be the ones that tend to infinity is only to simplify notation and is without loss of generality.) Let then $A = [A_{\infty}^\top A_F^\top]^\top$, $b = (b_{\infty}^\top b_F^\top)^\top$ be the corresponding partition of A and b , respectively, where A_{∞}, b_{∞} are non-empty by hypothesis. To have $\|\mu_{\infty}^{(\kappa_i)}\| \rightarrow \infty$, we need the terms that are integrated in the multipliers' update, i.e., last two terms in (17), to be positive for all i (in fact across a subsequence), which since $\tau > 0$ is equivalent to

$$A_{\infty} A^\top \mu^{(\kappa_i)} < -A_{\infty} F(x^{(\kappa_i)}) + \frac{1}{2\tau} (A_{\infty} x^{(\kappa_i)} - b_{\infty} + (Q(\mu^{(\kappa_i)}, M)\rho)_{\infty}), \quad (18)$$

where $(\cdot)_{\infty}$ denotes the elements of its argument corresponding to $\mu_{\infty}^{(\kappa_i)}$. Notice that $A_{\infty} A^\top \mu^{(\kappa_i)} = A_{\infty} A_{\infty}^\top \mu_{\infty}^{(\kappa_i)} + A_{\infty} A_F^\top \mu_F^{(\kappa_i)}$. As such, we have

$$A_{\infty} A_{\infty}^\top \mu_{\infty}^{(\kappa_i)} < -A_{\infty} A_F^\top \mu_F^{(\kappa_i)} - A_{\infty} F(x^{(\kappa_i)}) + \frac{1}{2\tau} (A_{\infty} x^{(\kappa_i)} - b_{\infty} + (Q(\mu^{(\kappa_i)}, M)\rho)_{\infty}). \quad (19)$$

However, $\lim_{i \rightarrow \infty} x^{(\kappa_i)} = \bar{x} \in X$, and $(Q(\mu^{(\kappa_i)}, M)\rho)_{\infty} \leq c\rho$ for all i , while by Lemma 3, F is continuous over the domain of multipliers

satisfying (13). Moreover, $\mu_F^{(\kappa_i)}$ contains the components of $\mu^{(\kappa_i)}$ that remain finite. Therefore, the limit as $i \rightarrow \infty$ of the right-hand side of (19) is finite. Due to the assumed full row-rank structure of A , matrix $A_\infty A_\infty^\top$ is invertible, hence (19) implies $\lim_{i \rightarrow \infty} \|\mu^{(\kappa_i)}\| < \infty$, establishing a contradiction showing that the subsequence $\{\mu^{(\kappa_i)}\}$ remains bounded.

(ii) If all elements of A are positive, and since $a_\ell^\top a_\ell = 1$, for all $\ell = 1, \dots, m$, all arguments of case (i) remain the same with the only difference that we directly have that $\|A_\infty A_\infty^\top \mu_\infty^{(\kappa_i)}\| \geq \|\mu_\infty^{(\kappa_i)}\|$.

The next result accompanies the region $S_K^* = \Pi_K \cap \mathbb{B}(x^*, \rho)$ of strategic deviations from the equilibrium x^* with *a priori* probabilistic feasibility guarantees that can be tuned by means of M . It should be noted that Theorem 2 establishes that there exists a choice of τ to guarantee convergence of Algorithm 1. The admissible range of values for τ is explicit via (15), (16), but difficult to quantify due to \bar{R} . Numerical evidence suggests that selecting a small enough value is sufficient for convergence.

Theorem 3. Consider Assumptions 1–4. Let x^* and $S_K^* = \Pi_K \cap \mathbb{B}(x^*, \rho)$ be returned by Algorithm 1; fix $\bar{\epsilon} \in (0, 1)$ and M . We then have that

$$\begin{aligned} \mathbb{P}^K \left\{ \delta_K \in \Delta^K : \mathbb{V}(S_K^*) \leq \bar{\epsilon} \right\} \\ \geq 1 - \sum_{i=0}^{nN+M-1} \binom{K}{i} \bar{\epsilon}^i (1 - \bar{\epsilon})^{K-i}. \end{aligned} \quad (20)$$

By Definition 2, Theorem 3 guarantees that for any point in S_K^* , the probability of constraint violation is bounded by $\bar{\epsilon}$, with confidence at least $1 - \sum_{i=0}^{nN+M-1} \binom{K}{i} \bar{\epsilon}^i (1 - \bar{\epsilon})^{K-i}$. The dependence of this term on M gives us an additional degree of freedom in trading the robustness of the solution for its associated probabilistic confidence. The choice of M can also have an effect on the size of S_K^* , as well as on the location of x^* , thus resulting in a trade-off between performance and robustness.

For the case in which the coupling constraints concern exclusively the aggregate variable, it can be shown that the upper limit of the summation in the right-hand side of (20) can be replaced by $n + M - 1$, as n is the dimension of the aggregate vector. This allows to state (20) with a much higher confidence of $1 - \sum_{i=0}^{n+M-1} \binom{K}{i} \bar{\epsilon}^i (1 - \bar{\epsilon})^{K-i}$; for details, we refer the reader to Pantazis et al. (2022), where the notion of support rank is exploited (Schildbach et al., 2012).

4. Numerical example

Consider a game with N agents whose decisions are subject to deterministic local constraints and uncertain coupling constraints on the aggregate decision:

$$\left. \begin{aligned} \min_{x_i \in X_i} \quad & x_i^\top (C\sigma(x) + d) \\ \text{subject to} \quad & \underline{b}_{\delta_k} \leq \sigma(x) \leq \bar{b}_{\delta_k}, \\ & k = 1, \dots, K \end{aligned} \right\} \forall i \in \mathcal{N}, \quad (21)$$

where $C \succ \alpha I_n$, for some $\alpha > 0$, and $d \in \mathbb{R}^n$. Note that a structure similar to our numerical example has been considered in applications of aggregative games such as electric vehicle charging and traffic management under uncertainty (Fele & Margellos, 2019, 2021; Paccagnan et al., 2019). We impose no knowledge of Δ and \mathbb{P} ; we rely instead on a scenario-based approximation of the game, whereby each sample $\delta_k \in \delta_K$ gives rise to $\underline{b}_{\delta_k}, \bar{b}_{\delta_k}$. Eq. (21) is an aggregative game in the form of (3). In this instance, we assume each agent's action has negligible effect on the aggregate, and accordingly consider a GWE-seeking problem.

Following the definition of F_{WE} (Section 2.2), we get $F(x) = F_{\text{WE}} = [C\sigma(x) + c]_{i \in \mathcal{N}}$. It can be verified that F is Lipschitz continuous and strongly monotone with respect to σ : by Facchinei and Pang (2003, Thm. 2.3.3), (21) admits a unique aggregate equilibrium $\sigma^* = \sigma(x^*)$.⁵

We employ Algorithm 1 to seek a WE x^* such that, by fixing M , a prespecified theoretical violation level is guaranteed for the set $\Pi_K \cap \mathbb{B}(x^*, \rho)$. Due to uniqueness of σ^* , all sets $\mathbb{B}(\cdot, \rho)$ – parametrised by any x^* solving (21) – are projected on the unique ball $\mathbb{B}(\sigma^*, \rho/N)$ in the aggregate space. Also note that by definition of σ , at most n non-redundant samples will contribute to define the domain $\Pi_K^\sigma := \frac{1}{N}(\mathbf{1}_N^\top \otimes I_n)X \cap \left(\bigcap_{k=1}^K C_{\delta_k}\right)$ in (21). For the derivation of the robustness guarantees, we can thus restrict our attention to $S_K^* = \Pi_K^\sigma \cap \mathbb{B}(\sigma^*, \rho/N) \subseteq \mathbb{R}^n$. As remarked at the end of Section 3.3, we can apply (20) with the upper limit in the summation involved replaced by $n + M - 1$. For the case $n = 2$, $N = 50$, and different choices of M , Fig. 3 depicts the projected iterations $\{\sigma(x^{(\kappa)})\}$, $\kappa = 1, 2, \dots$ generated by Algorithm 1 on the space Π_K^σ . It can be observed how the region S_K^* changes as the value of M is modified.

It is worth noting that in this case $F(x)$ is integrable—this can be inferred by Facchinei and Pang (2003, Thm. 1.3.1) since the Jacobian of the game is symmetric, i.e., $\nabla_x F(x) = \nabla_x F(x)^\top$. Therefore, a GWE x^* can also be obtained by solving

$$\begin{aligned} \min_{x \in X} \quad & \sigma(x)^\top C\sigma(x) + d^\top \sigma(x) \\ \text{subject to} \quad & \underline{b}_{\delta_k} \leq \sigma(x) \leq \bar{b}_{\delta_k}, \quad k = 1, \dots, K. \end{aligned} \quad (22)$$

In other words, this game admits a *potential function* $E(x) := \sigma(x)^\top C\sigma(x) + d^\top \sigma(x)$, whose minimisers correspond to GWEs. E can be interpreted as the total cost incurred by the population of agents, and its minimisation leads to the optimum social welfare. The contour lines of E are depicted in Fig. 3: since x^* minimises $E(\cdot)$, σ^* lies on the contour associated to the minimum value of E within the feasible domain. Lower values of M result in larger regions for which guarantees are provided. Fig. 4 shows how the sequence $\{E(x^{(\kappa)})\}_{\kappa=1,2,\dots}$ converges to the minimum *potential* within the possibly tightened feasibility region. It can be observed how in this case the efficiency of the equilibrium decreases as smaller values of M are chosen. The three panels in Fig. 4 show the trade-off between system level efficiency and the guaranteed robustness levels. The lower the value of M , the lower the empirical constraints violation—corresponding to a better confidence bound in the right-hand side of (20).

5. Concluding remarks

This work proposes a data-driven equilibrium-seeking algorithm such that probabilistic feasibility guarantees are provided for a region surrounding a game equilibrium. These guarantees are *a priori* and the region that is accompanied with such a probabilistic certificate is tunable. For games that admit a potential function, the proposed scheme is shown to achieve a trade-off between cost and the level of probabilistic feasibility guarantees. In fact, our scheme returns the most efficient equilibrium such that the predefined guarantees are achieved. Proving this conjecture is left for future work. Moreover, current work investigates a distributed implementation of the proposed equilibrium seeking algorithm.

⁵ We note that this case slightly transcends the conditions in Theorem 2, as F does not comply with Assumption 3–(1). Convergence of Algorithm 1 (following from the nonexpansiveness of T on each subdomain \mathcal{M}_i) can still be ensured here due to the affine structure of F ; cf. Facchinei and Pang (2003, Sec. 12.5.1).




Fig. 3. Iterates generated by Algorithm 1 (blue diamonds) for different choices of M . In this numerical instance, $N = 50$, $\rho = 10$, and $X_i := \{x_i \in \mathbb{R}^n : x_i \in [\underline{x}_i, \bar{x}_i]\}$, with $\underline{x}_i = (0, 0)$, $\bar{x}_i = (3.5, 3.5)$. The randomly generated coupling constraints form the rectangular feasibility region Π_K^c (delineated by the solid black line). The red-shaded region represents the intersection between the latter and the ball $\mathbb{B}_1(\sigma^*, \rho/N)$ around the aggregate equilibrium σ^* (red diamond marker). In this instance, its volume increases as larger values for M are chosen. The value associated to the contour lines of the potential function E decreases from top-right to bottom-left, showing that σ^* is the unique minimiser in the admissible region (shaded in green) after constraint tightening is performed by the algorithm (see Section 3.2.2). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

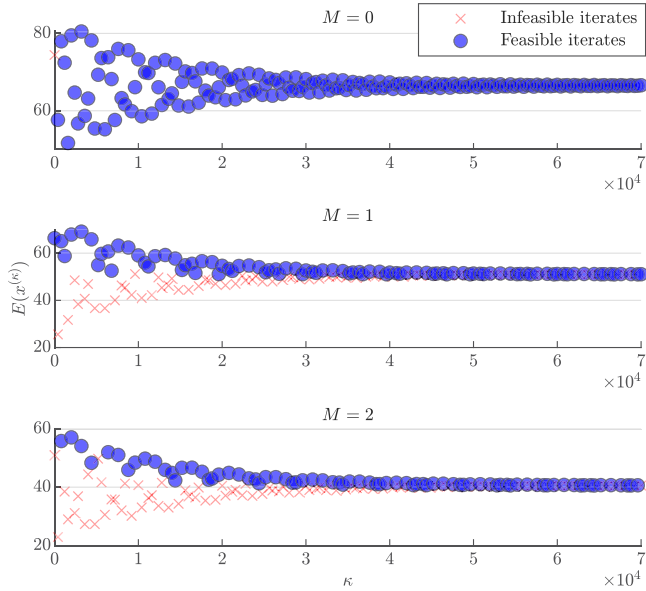


Fig. 4. Potential function $E(x^{(k)})$ evaluated along the iterations of Algorithm 1. Lower values of M yield better confidence on the theoretical robustness certificates for the considered region (see Theorem 3), which results in a lower empirical probability of constraint violation. On the other hand, the system-level efficiency of the equilibrium increases for higher values of M .

Acknowledgements

F. Fele gratefully acknowledges support from grant RYC2021-033960-I funded by MICIU/AEI/10.13039/501100011033 and European Union NextGenerationEU/PRTR, as well as from grant PID2022-142946NA-I00 funded by MICIU/AEI/10.13039/501100011033 and by ERDF/EU. The work of K. Margellos was partially supported by EPSRC under Grant EPP03277X/1.

Appendix A. Proof of Lemma 2

Let μ, z be arbitrary vectors in \mathcal{M} and, as in the proof of Lemma 3, define $\bar{\mu}, \bar{z}$ as the vectors composed by rearranging the elements of μ, z in decreasing order. According to this arrangement, let $\mathcal{I}_\mu = \{i_1, i_2, \dots, i_m\}$ be the ordered set of indices of μ , i.e., $i_k : \mu_{i_k} = \bar{\mu}_k, k = 1, \dots, m$; as a result, i_1 and i_m will be the indices of the largest and smallest components of μ , respectively. Applying a similar definition to z , we denote the corresponding

set $\mathcal{I}_z := \{j_1, j_2, \dots, j_m\}$. Then, the first M indices in \mathcal{I}_μ and \mathcal{I}_z , denoted as \mathcal{L}_μ and \mathcal{L}_z , respectively, are relative to the constraints not tightened by the application of $Q(\cdot, M)$. In other words, for all $\ell \in \mathcal{L}_\mu$, $(Q(\mu, M)\rho)_\ell = 0$ – and similarly for z . Vice versa, the complementary sets $\mathcal{L}_\mu^c = \mathcal{I}_\mu \setminus \mathcal{L}_\mu$ and $\mathcal{L}_z^c = \mathcal{I}_z \setminus \mathcal{L}_z$ are such that for all $\ell \in \mathcal{L}_\mu^c$, $(Q(\mu, M)\rho)_\ell = c\rho$, and for all $\ell \in \mathcal{L}_z^c$, $(Q(z, M)\rho)_\ell = c\rho$. Let $q = [Q(\mu, M) - Q(z, M)]\rho$. We distinguish between the following cases:

- (1) $\ell \in \mathcal{L}_\mu^c \cap \mathcal{L}_z$: we have $(Q(\mu, M)\rho)_\ell = c\rho$ since $\ell \in \mathcal{L}_\mu^c$, while $(Q(z, M)\rho)_\ell = 0$ as $\ell \in \mathcal{L}_z$. Then, $q_\ell = c\rho$.
- (2) $\ell \in \mathcal{L}_\mu \cap \mathcal{L}_z^c$: from $\ell \in \mathcal{L}_z^c$ we have $(Q(z, M)\rho)_\ell = c\rho$. On the other hand, since $\ell \in \mathcal{L}_\mu$, $(Q(\mu, M)\rho)_\ell = 0$. This results in $q_\ell = -c\rho$.
- (3) $\ell \in (\mathcal{L}_\mu \cap \mathcal{L}_z) \cup (\mathcal{L}_\mu^c \cap \mathcal{L}_z^c)$. If $\ell \in \mathcal{L}_\mu \cap \mathcal{L}_z$ then $(Q(\cdot, M)\rho)_\ell = 0$ for both μ and z . Therefore, $q_\ell = 0$. Conversely, if $\ell \in \mathcal{L}_\mu^c \cap \mathcal{L}_z^c$, then $(Q(\cdot, M)\rho)_\ell = c\rho$ for both μ and z , which results again in $q_\ell = 0$.

The sets $\mathcal{L}_\mu^c \cap \mathcal{L}_z$, $\mathcal{L}_\mu \cap \mathcal{L}_z^c$, $(\mathcal{L}_\mu \cap \mathcal{L}_z) \cup (\mathcal{L}_\mu^c \cap \mathcal{L}_z^c)$ are pairwise disjoint and exhaust the set $\{1, \dots, m\}$. Hence we can write

$$\begin{aligned} U &= (\mu - z)^\top (Q(\mu, M) - Q(z, M))\rho = \sum_{\ell=1}^m (\mu_\ell - z_\ell) q_\ell \\ &= \sum_{i \in \mathcal{L}_\mu^c \cap \mathcal{L}_z} (\mu_i - z_i) c\rho + \sum_{j \in \mathcal{L}_\mu \cap \mathcal{L}_z^c} (\mu_j - z_j) \cdot (-c\rho) \\ &= \underbrace{\left(\sum_{i \in \mathcal{L}_\mu^c \cap \mathcal{L}_z} \mu_i - \sum_{j \in \mathcal{L}_\mu \cap \mathcal{L}_z^c} \mu_j \right)}_{=: U_1} + \underbrace{\left(\sum_{j \in \mathcal{L}_\mu \cap \mathcal{L}_z^c} z_j - \sum_{i \in \mathcal{L}_\mu^c \cap \mathcal{L}_z} z_i \right)}_{=: U_2} c\rho, \end{aligned} \quad (\text{A.1})$$

Now, notice that for any $i \in \mathcal{L}_\mu^c \cap \mathcal{L}_z \subseteq \mathcal{L}_\mu^c$ and $j \in \mathcal{L}_\mu \cap \mathcal{L}_z^c \subseteq \mathcal{L}_\mu$, we have by definition of \mathcal{L}_μ and \mathcal{L}_μ^c that $\mu_i \leq \mu_j$ (which by (13) only holds with equality if $\mu_i = \mu_j = 0$). With analogous reasoning, we have $z_i \geq z_j$ for any $i \in \mathcal{L}_\mu^c \cap \mathcal{L}_z \subseteq \mathcal{L}_z$ and $j \in \mathcal{L}_\mu \cap \mathcal{L}_z^c \subseteq \mathcal{L}_z^c$. Let h_1 be the cardinality of the set $\mathcal{L}_\mu^c \cap \mathcal{L}_z$, and h_2 that of $\mathcal{L}_\mu \cap \mathcal{L}_z^c$. Then,

$$\begin{aligned} h_1 &= |\mathcal{L}_\mu^c \cap \mathcal{L}_z| \stackrel{(a)}{=} |\mathcal{L}_z \setminus \mathcal{L}_\mu| = |\mathcal{L}_z| - |\mathcal{L}_\mu \cap \mathcal{L}_z| \\ &\stackrel{(b)}{=} |\mathcal{L}_\mu| - |\mathcal{L}_z \cap \mathcal{L}_\mu| = |\mathcal{L}_\mu \setminus \mathcal{L}_z| = |\mathcal{L}_\mu \cap \mathcal{L}_z^c| = h_2, \end{aligned}$$

where (a) holds since $\mathcal{L}_\mu, \mathcal{L}_z \subseteq \{1, \dots, M\}$, and (b) follows from $|\mathcal{L}_\mu| = |\mathcal{L}_z| = M$. Therefore $h_1 = h_2 =: h$ and $0 \leq h \leq M$, which implies $U_1 \leq 0$ and $U_2 \leq 0$ in (A.1). We can observe that $U_1 < 0$ and $U_2 < 0$ if $\mathcal{L}_\mu \cap \mathcal{L}_z^c$ and $\mathcal{L}_\mu^c \cap \mathcal{L}_z$ are nonempty and the corresponding components of μ and z are nonzero. In such a

case $h \geq 1$ and we can write

$$U_1 = \sum_{i \in \mathcal{L}_\mu \cap \mathcal{L}_z} \mu_i - \sum_{j \in \mathcal{L}_\mu \cap \mathcal{L}_z^c} \mu_j \leq -h\zeta, \quad (\text{A.2})$$

where the inequality follows from (13) and the above discussion. A similar reasoning holds for U_2 . Lastly, note that if $\mu \neq z$ and $h \geq 1$, then at least one of $U_1 \leq -h\zeta$ and $U_2 \leq -h\zeta$ will hold. By (A.1), we can thus conclude $U \leq -h\zeta c\rho$ for any $\mu, z \in \mathcal{M}$, $\mu \neq z$. ■

Appendix B. Proof of Lemma 3

Part (1): To prove that the mapping T is continuous on its domain, we first notice that T is by construction continuous on $X \times \mathcal{M}$ when the operator $Q(\cdot, M)$ is continuous on \mathcal{M} (as the parameter M is fixed). Therefore, it is sufficient to show that for any $\mu, z \in \mathcal{M}$ and any $\eta > 0$, there exists $\delta > 0$ such that

$$\|\mu - z\| < \delta \Rightarrow \|Q(\mu, M) - Q(z, M)\| \|\rho\| < \eta, \quad (\text{B.1})$$

where $\rho = c\rho \mathbf{1}_m \neq \mathbf{0}$. To this end, consider any $\mu, z \in \mathcal{M}$ such that $\|\mu - z\| < \frac{\zeta}{2}$, with ζ as defined in (13).⁶ Let $\tilde{\mu}$ and \tilde{z} denote the vectors μ and z sorted in decreasing order; thus, $\tilde{\mu}_\ell$ is the ℓ th largest element of μ (and similarly for z). For any given ℓ , let $i : \mu_i = \tilde{\mu}_\ell, j : z_j = \tilde{z}_\ell$, and $\bar{\ell} := \min_{\{1, \dots, m\}} \ell : i \neq j$. In words, $\bar{\ell}$ is the smallest index for which the ℓ th largest elements of μ and z do not appear at the same row of their respective vectors. We then let \mathcal{I} be the set of indices for which the ordering of the elements of μ and z agrees, i.e., for all $k \in \mathcal{I}$, there exists $\ell < \bar{\ell}$ such that $i = j = k$, with $i : \mu_i = \tilde{\mu}_\ell$ and $j : z_j = \tilde{z}_\ell$.

We prove our statement by contradiction. Suppose there exists $i, j \notin \mathcal{I}$ such that $i : \mu_i = \tilde{\mu}_\ell$ and $j : z_j = \tilde{z}_\ell$ for some $\ell > \bar{\ell}$, where $\mu_i < \mu_j$ and $z_i > z_j$. First, we note that such an instance exists by hypothesis, as otherwise the only possible case is where $i = j$, which contradicts $i, j \notin \mathcal{I}$ and implies $Q(\mu, M) = Q(z, M)$. Since $z \in \mathcal{M}$, it further holds $z_j < z_i - \zeta$, which by $\|\mu_i - z_i\| \leq \|\mu - z\| < \frac{\zeta}{2}$ implies

$$z_j < \mu_i + \frac{\zeta}{2} - \zeta. \quad (\text{B.2})$$

We bound (B.2) from below by noting $z_j > \mu_j - \frac{\zeta}{2}$, which holds since $\|\mu_j - z_j\| < \frac{\zeta}{2}$, obtaining $\mu_j - \frac{\zeta}{2} < \mu_i + \frac{\zeta}{2} - \zeta$, or equivalently $\mu_j < \mu_i$, which contradicts our hypothesis. Hence the elements of any pair of vectors $\mu, z \in \mathcal{M}$ such that $\|\mu - z\| < \frac{\zeta}{2}$ must follow the same ordering. By definition of $P(\cdot)$, this implies $P(\mu) = P(z)$ and, in turn, $\|Q(\mu, M) - Q(z, M)\| = 0$. This validates (B.1) with $\delta = \frac{\zeta}{2}$ and any $\eta > 0$, establishing the continuity of $Q(\cdot, M)$ on \mathcal{M} and concluding the proof of the first part.

Part (2): We show that the mapping T fulfils certain nonexpansiveness properties required for the convergence of Algorithm 1, for compatible choices of τ . In particular, we provide a sufficient condition for which the iteration

$$y^{(k+1)} = \text{proj}_{X \times \mathcal{M}_j, D} [y^{(k)} - D^{-1}T(y^{(k)}, \rho, M)], \quad (\text{B.3})$$

converges to a solution of $\text{VI}(X \times \mathcal{M}_j, T)$, where $j \in \{1, \dots, q\}$ is fixed, for any $y^{(0)} \in X \times \mathcal{M}_j$. Notice that in (B.3) the skew projection is performed on the convex subdomain $X \times \mathcal{M}_j$. (B.3) is the solution of the $\text{VI}(X \times \mathcal{M}_j, T_D^{(k)})$ (see Facchinei and Pang (2003, Sec. 12.5.1)), where $T_D^{(k)}(y) := T(y^{(k)}, \rho, M) + D(y - y^{(k)})$ is strongly monotone due to $D > 0$ and $(T(y, \rho, M) - T(y', \rho, M))^T(y - y') \geq 0$, for all $y, y' \in X \times \mathcal{M}$, which in turn follows from Assumption 3 and Lemma 2. The fixed-point iteration (B.3) is an instance of the forward-backward splitting method: we thus resort to standard

results in the literature to prove its convergence. Following the notation in Facchinei and Pang (2003, Sec. 12.5.1), we let $D := D_s^{-1/2}(D - D_s)D_s^{-1/2}$, where $D_s := \frac{D+D^T}{2}$. Also, $\mathcal{U}_j := \{D_s^{1/2}y : y \in X \times \mathcal{M}_j\}$, $\mathcal{U} = \bigcup_{j=1}^q \mathcal{U}_j$, and $\tilde{T}(w) := D_s^{-1/2}T(D_s^{-1/2}w, \rho, M)$, for all $w \in \mathcal{U}$. To ease notation, we drop the dependence of \tilde{T} and \tilde{T}_D on ρ, M , as they remain fixed throughout the proof. According to Facchinei and Pang (2003, Thm. 12.5.2) (see also Zhu and Marcotte (1996, Sec. 4.3)), to ensure convergence of (B.3) to a solution of the $\text{VI}(X \times \mathcal{M}_j, T)$ it is sufficient to show that $\tilde{T}_D = \tilde{T} - \tilde{D}$ is β -cocoercive on \mathcal{U}_j , i.e.,

$$(\tilde{T}_D(v) - \tilde{T}_D(w))^T(v - w) \geq \beta \|\tilde{T}_D(v) - \tilde{T}_D(w)\|^2, \quad (\text{B.4})$$

for some $\beta > \frac{1}{2}$ and all $v, w \in \mathcal{U}_j, j \in \{1, \dots, q\}$. In fact, we will go a step forward and demonstrate here that \tilde{T}_D is co-coercive on \mathcal{U} with $\beta > \frac{1}{2}$. Due to the saddle problem structure of the mapping in (8), we adopt the procedure in Facchinei and Pang (2003, Prop. 12.5.4) and define D as in (9) (see also Paccagnan et al. (2019)). It then follows from the above definitions that $\tilde{T}_D(w)$, for any $w \in \mathcal{U}$, reduces to

$$\tilde{T}_D(D_s^{1/2}y) = D_s^{-1/2} \begin{bmatrix} F(x) \\ b - Q(\mu, M)\rho \end{bmatrix}, \quad \forall y \in X \times \mathcal{M}, \quad (\text{B.5})$$

which can be easily seen by rewriting (8) as

$$T(y, \rho, M) = \begin{bmatrix} F(x) \\ 0 \end{bmatrix} + \underbrace{\begin{bmatrix} 0 & A^T \\ -A & 0 \end{bmatrix}}_{D-D_s} y + \begin{bmatrix} 0 \\ b - Q(\mu, M)\rho \end{bmatrix}.$$

Define $W := (D_s^{-1/2})^T D_s^{-1/2} = D_s^{-1}$, and let $\tilde{Q}(\cdot)$ be a shorthand for $Q(\cdot, M)\rho$ (as M is a fixed parameter). Then, for any $w_a, w_b \in \mathcal{U}$, we can expand (B.4) by using (B.5), obtaining

$$\begin{aligned} (w_a - w_b)^T(\tilde{T}_D(w_a) - \tilde{T}_D(w_b)) - \beta \|\tilde{T}_D(w_a) - \tilde{T}_D(w_b)\|^2 \\ = (D_s^{-1/2}w_a - D_s^{-1/2}w_b)^T \begin{bmatrix} F(x_a) - F(x_b) \\ \tilde{Q}(\mu_b) - \tilde{Q}(\mu_a) \end{bmatrix} \\ - \beta \left\| D_s^{-1/2} \begin{bmatrix} F(x_a) - F(x_b) \\ \tilde{Q}(\mu_b) - \tilde{Q}(\mu_a) \end{bmatrix} \right\|^2 \\ = \begin{bmatrix} x_a - x_b \\ \mu_a - \mu_b \end{bmatrix}^T \begin{bmatrix} F(x_a) - F(x_b) \\ \tilde{Q}(\mu_b) - \tilde{Q}(\mu_a) \end{bmatrix} \\ - \beta \begin{bmatrix} F(x_a) - F(x_b) \\ \tilde{Q}(\mu_b) - \tilde{Q}(\mu_a) \end{bmatrix}^T W \begin{bmatrix} F(x_a) - F(x_b) \\ \tilde{Q}(\mu_b) - \tilde{Q}(\mu_a) \end{bmatrix}, \end{aligned} \quad (\text{B.6})$$

for all $y_a, y_b \in X \times \mathcal{M}$, where the last equality follows from the definition of \mathcal{U}_j and by expanding the norm. Matrix W can be written as $W = \begin{bmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{bmatrix}$, where $W_{11} \in \mathbb{R}^{nN \times nN}$, $W_{12} \in \mathbb{R}^{nN \times m}$, $W_{33} \in \mathbb{R}^{m \times m}$ are:

$$\begin{aligned} W_{11} &= \tau(I_n - \tau^2 A^T A)^{-1}, \\ W_{12} &= W_{21}^T = \tau^2(I_n - \tau^2 A^T A)^{-1} A^T, \\ W_{22} &= \tau I_m + \tau^3 A(I_n - \tau^2 A^T A)^{-1} A^T. \end{aligned}$$

Expanding the inner product in (B.6) with respect to the matrix blocks $W_{11}, W_{12}, W_{21}, W_{33}$ we obtain

$$\begin{aligned} \beta(F(x_a) - F(x_b))^T \left[\frac{1}{\beta}(x_a - x_b) \right. \\ - W_{11}(F(x_a) - F(x_b)) - 2W_{12}(\tilde{Q}(\mu_b) - \tilde{Q}(\mu_1)) \\ + \beta(\tilde{Q}(\mu_b) - \tilde{Q}(\mu_a))^T \left[\frac{1}{\beta}(\mu_a - \mu_b) \right. \\ \left. \left. - W_{22}(\tilde{Q}(\mu_b) - \tilde{Q}(\mu_a)) \right] \right] \end{aligned}$$

⁶ The proof of this part also holds for $\mu \in \mathbb{R}_+^m \supset \mathcal{M}$.

$$\begin{aligned}
&= (F(x_a) - F(x_b))^T(x_a - x_b) \\
&\quad - \beta(F(x_a) - F(x_b))^T W_{11}(F(x_a) - F(x_b)) \\
&\quad - 2\beta(F(x_a) - F(x_b))^T W_{12}(\bar{Q}(\mu_b) - \bar{Q}(\mu_a)) \\
&\quad + (\bar{Q}(\mu_b) - \bar{Q}(\mu_a))^T(\mu_a - \mu_b) \\
&\quad - \beta(\bar{Q}(\mu_b) - \bar{Q}(\mu_a))^T W_{22}(\bar{Q}(\mu_b) - \bar{Q}(\mu_a)).
\end{aligned}$$

Setting $p_\tau := (I - \tau^2 A^T A)^{-1/2}(F(x_a) - F(x_b))$ and $q_\tau := \tau(I - \tau^2 A^T A)^{-1/2} A^T(\bar{Q}(\mu_b) - \bar{Q}(\mu_a))$ above we obtain

$$\begin{aligned}
&(F(x_a) - F(x_b))^T(x_a - x_b) \\
&\quad + (\bar{Q}(\mu_b) - \bar{Q}(\mu_a))^T(\mu_a - \mu_b) \\
&\quad - \beta\tau(\bar{Q}(\mu_b) - \bar{Q}(\mu_a))^T(\bar{Q}(\mu_b) - \bar{Q}(\mu_a)) \\
&\quad - \beta\tau(p_\tau + q_\tau)^T(p_\tau + q_\tau) \\
&\geq \alpha\|x_a - x_b\|^2 + 2h\zeta c\rho \\
&\quad - 2\beta\tau h(c\rho)^2 - 2\beta\tau(p_\tau^T p_\tau + q_\tau^T q_\tau),
\end{aligned} \tag{B.7}$$

where for the last inequality we used, in order, (i) strong monotonicity of F (cf. [Assumption 3](#)), (ii) [Lemma 2](#), (iii) $\|\bar{Q}(\mu_b) - \bar{Q}(\mu_a)\|^2 \leq 2h(c\rho)^2$ – which follows from the same arguments used in the proof of [Lemma 2](#) – and (iv) $(p_\tau + q_\tau)^T(p_\tau + q_\tau) \leq 2(p_\tau^T p_\tau + q_\tau^T q_\tau)$. Expanding the term containing p_τ, q_τ in (B.7) we get

$$\begin{aligned}
&\alpha\|x_a - x_b\|^2 + 2h\zeta c\rho - 2\beta\tau h(c\rho)^2 \\
&\quad - 2\beta\tau(F(x_a) - F(x_b))^T(I_n - \tau^2 A^T A)^{-1}(F(x_a) - F(x_b)) \\
&\quad - 2\beta\tau^3(\bar{Q}(\mu_b) - \bar{Q}(\mu_a))^T \\
&\quad \cdot A(I_n - \tau^2 A^T A)^{-1} A^T(\bar{Q}(\mu_b) - \bar{Q}(\mu_a)) \\
&\stackrel{(a)}{\geq} \alpha\|x_a - x_b\|^2 + 2h\zeta c\rho - 2\beta\tau h(c\rho)^2 \\
&\quad - 2\beta\tau\|F(x_a) - F(x_b)\|^2 \cdot \|(I_n - \tau^2 A^T A)^{-1}\| \\
&\quad - 2\beta\tau^3\|\bar{Q}(\mu_b) - \bar{Q}(\mu_a)\|^2 \cdot \|(I_n - \tau^2 A^T A)^{-1}\| \cdot \|A\|^2 \\
&\stackrel{(b)}{\geq} (\alpha - 2\beta\tau L_F^2 \|(I_n - \tau^2 A^T A)^{-1}\|)\|x_a - x_b\|^2 \\
&\quad + 2h\zeta c\rho - 2\beta\tau h(c\rho)^2 \left(1 + \frac{2\tau^2}{1 - \tau^2\|A\|^2} \|A\|^2\right),
\end{aligned} \tag{B.8}$$

where (a) is obtained by applying the Cauchy–Schwarz inequality, and in (b) we use the Lipschitz continuity of F (cf. [Assumption 3](#)), $\|\bar{Q}(\mu_b) - \bar{Q}(\mu_a)\|^2 \leq 2h(c\rho)^2$, and the triangle inequality. Notice that for the last term in (B.8),

$$\begin{aligned}
&2\beta\tau h(c\rho)^2 \left(1 + \frac{2\tau^2}{1 - \tau^2\|A\|^2} \|A\|^2\right) \\
&= 2\beta\tau h(c\rho)^2 \frac{1 + \tau^2\|A\|^2}{1 - \tau^2\|A\|^2} \leq 2\beta\tau h(c\rho)^2 \frac{1 + \|A\|^2}{1 - \tau^2\|A\|^2}
\end{aligned} \tag{B.9}$$

holds for any choice of $\tau \in (0, \max\{\frac{1}{\|A\|}, 1\})$. Recall that by invoking ([Facchinei & Pang, 2003](#), Thm. 12.5.2), our objective is to show that (B.4) holds for some $\tau > 0$ and $\beta > \frac{1}{2}$. Then, by inspecting (B.8) and using (B.9), to achieve this it is sufficient to guarantee

$$\begin{aligned}
&\alpha - \tau L_F^2 \|(I_n - \tau^2 A^T A)^{-1}\| > 0, \\
&2h\zeta c\rho - \tau h(c\rho)^2 \frac{1 + \|A\|^2}{1 - \tau^2\|A\|^2} > 0, \text{ if } 1 \leq h \leq M.
\end{aligned}$$

Solving the quadratic expressions above with respect to τ results in the admissible range of values in (15) (these are also satisfying $\tau \in (0, \max\{\frac{1}{\|A\|}, 1\})$, required for (B.9) to hold). Therefore, for any τ satisfying this condition, \tilde{T}_D is co-coercive with $\beta > \frac{1}{2}$ on the entire domain \mathcal{U} , which in turn implies that co-coercivity of \tilde{T}_D holds on each subdomain $\mathcal{U}_j, j = 1, \dots, q$, with the same modulus.




Fig. C.1. Domain \mathcal{M} of the Lagrange multipliers associated to the coupling constraints, for the case $m = 2$. Notice the minimum distance ζ between any two subdomains of \mathcal{M} involves the origin as one of these subdomains.

By [Facchinei and Pang \(2003, Thm. 12.5.2\)](#), this is sufficient to guarantee the convergence of (B.3) to a solution of the VI($X \times \mathcal{M}_j, T$), thus concluding the proof. ■

Appendix C. Proof of Theorem 2

Fix any τ satisfying the conditions of [Lemma 3](#) and (16). The sequence $\{y^{(k)}\}_{k=1,2,\dots}$ (where $y^{(k)} = (x^{(k)}, \mu^{(k)})$) generated by Algorithm 1 lives in a compact set since X and \mathcal{M} are compact (see [Assumption 4](#)). As such, by the Bolzano–Weierstrass theorem ([Rudin, 1976](#), Thm. 3.6), there exist convergent subsequences, i.e., the set

$$\begin{aligned}
\Omega &:= \{\bar{y} = (\bar{x}, \bar{\mu}) : \exists \text{ subsequence } \{\kappa_i\}_{i \in \mathbb{N}} \\
&\text{such that } \lim_{i \rightarrow \infty} \kappa_i = \infty, \lim_{i \rightarrow \infty} y^{(\kappa_i)} = \bar{y}\},
\end{aligned} \tag{C.1}$$

containing the limit points of $\{y^{(k)}\}$ is non-empty; see, e.g., [Rudin \(1976, p. 48\)](#). We will show that Ω is a singleton for any τ satisfying (15)–(16), which implies that the iterates generated by Algorithm 1 have a unique limit point, hence they converge. To achieve this, we assume for the sake of contradiction that there exist two cluster points $\bar{y}_1, \bar{y}_2 \in \Omega$, where $\bar{y}_1 = (\bar{x}_1, \bar{\mu}_1)$ and $\bar{y}_2 = (\bar{x}_2, \bar{\mu}_2)$. Moreover, we assume that $\bar{\mu}_1 \in \mathcal{M}_i$, and $\bar{\mu}_2 \in \mathcal{M}_j$, with $i \neq j$. Note that if this were not the case, then we would be in a trivial case where $\bar{y}_1 = \bar{y}_2$, due to co-coercivity of T (see [Lemma 3](#))—by which Algorithm 1 converges to a unique solution when restricted to any convex subdomain $X \times \mathcal{M}_j, j = 1, \dots, q$. To ease the notation in the remainder of the proof, we assume without loss of generality that $\bar{\mu}_1 \in \mathcal{M}_1, \bar{\mu}_2 \in \mathcal{M}_2$ (see [Fig. C.1](#)). By (C.1) there exist an infinite subsequence $\{\kappa_i\}_{i \in \mathbb{N}}$ of the iterates generated by Algorithm 1 whose elements get arbitrarily close to $\bar{\mu}_1$ while staying in \mathcal{M}_1 where this cluster point belongs (similarly for $\bar{\mu}_2$). We then have that for any $\delta > 0$, there exists $\tilde{\kappa}$ such that for all $\kappa_i \geq \tilde{\kappa}$, $\|y^{(\kappa_i)} - \bar{y}_1\| \leq \delta$; this implies $\|x^{(\kappa_i)} - \bar{x}_1\| \leq \delta$ and $\|\mu^{(\kappa_i)} - \bar{\mu}_1\| \leq \delta$.

Due to our contradiction hypothesis (recall that $\{\kappa_i\}_{i \in \mathbb{N}}$ is a subsequence), the sequence of iterates generated by Algorithm 1 would be leaving \mathcal{M}_1 towards \mathcal{M}_2 infinitely often. Denote then by $\bar{\kappa} > \tilde{\kappa}$ the smallest index of the subsequence such that $\mu^{(\bar{\kappa})} \in \mathcal{M}_1$, but $\mu^{(\bar{\kappa}+1)} \in \mathcal{M}_2$, i.e., after the $\bar{\kappa}$ th iterate the original sequence would jump to \mathcal{M}_2 (for the first time after $\bar{\kappa}$). For this jump to occur, the unprojected solution for the Lagrange multipliers must be “closer” to \mathcal{M}_2 than to any other sub-domain of \mathcal{M} . To see this more formally, let D_μ^{-1} denote the lower block-row of $D^{-1} = \begin{bmatrix} \tau I_{nN} & 0 \\ 2A\tau^2 & \tau I_m \end{bmatrix}$, corresponding to the Lagrange multiplier

update in line 3 of Algorithm 1. By definition of \mathcal{M} , such a jump requires the Euclidean distance between the unprojected gradient step at $\bar{\kappa} + 1$ and $\mu^{(\bar{\kappa})}$ to satisfy

$$\|\mu^{(\bar{\kappa})} - D_{\mu}^{-1}T(y^{(\bar{\kappa})}, \rho, M) - \mu^{(\bar{\kappa})}\| > \zeta/2. \quad (\text{C.2})$$

Fig. C.1 illustrates this construction: (C.2) describes the minimum distance for a jump to occur. This is when the ellipsoidal contour levels according to which the projection is performed (skew projection defined by matrix D) have their major axis aligned between subdomains as in Fig. C.1 (solid red ellipses). For this two-dimensional example this distance would then be half the width of the white stripe, i.e., $\zeta/\sqrt{2}$. We rather impose $\zeta/2$ (which is smaller) in (C.2), to account for the case where one of the subdomains is the origin (\mathcal{M}_3). However,

$$\begin{aligned} & \|\mu^{(\bar{\kappa})} - D_{\mu}^{-1}T(y^{(\bar{\kappa})}, \rho, M) - \mu^{(\bar{\kappa})}\| \\ &= \tau \| -2\tau A(F(x^{(\bar{\kappa})}) + A^T \mu^{(\bar{\kappa})}) \\ & \quad + Ax^{(\bar{\kappa})} - b + Q(\mu^{(\bar{\kappa})}, M)\rho \| \\ &= \tau \| -2\tau A(F(x^{(\bar{\kappa})}) - F(\bar{x}_1) + A^T(\mu^{(\bar{\kappa})} - \bar{\mu}_1)) \\ & \quad + A(x^{(\bar{\kappa})} - \bar{x}_1) + Q(\mu^{(\bar{\kappa})}, M)\rho \\ & \quad - 2\tau A(F(\bar{x}_1) + A^T \bar{\mu}_1) + (A\bar{x}_1 - b) \| \\ &\leq \tau^2 \|2A(F(\bar{x}_1) + A^T \bar{\mu}_1)\| + \tau \|A\bar{x}_1 - b\| \\ & \quad + \tau \|Q(\mu^{(\bar{\kappa})}, M)\rho\| + \tau \|A\| \|x^{(\bar{\kappa})} - \bar{x}_1\| \\ & \quad + 2\tau^2 (\|A(F(x^{(\bar{\kappa})}) - F(\bar{x}_1))\| + \|AA^T(\mu^{(\bar{\kappa})} - \bar{\mu}_1)\|) \\ &\leq (\tau^2 + \tau)\bar{R} + \tau c\rho\sqrt{m-M} \\ & \quad + \tau\delta(2\tau(L_F\|A\| + \|AA^T\|) + \|A\|), \end{aligned} \quad (\text{C.3})$$

where the first equality follows from the definition of D_{μ}^{-1} and T , and the second one by adding and subtracting $F(\bar{x}_1)$, $A^T \bar{\mu}_1$ and $A\bar{x}_1$. The first inequality is due to the triangle inequality, while the last one follows from the previous one by upper-bounding (i) the first two terms using the definition of \bar{R} ; (ii) $\|Q(\mu^{(\bar{\kappa})}, M)\rho\|$ by $c\rho\sqrt{m-M}$ based on its definition; and (iii) the last three terms using $\|F(x^{(\bar{\kappa})}) - F(\bar{x}_1)\| \leq L_F \|x^{(\bar{\kappa})} - \bar{x}_1\|$ by Assumption 3, and $\|x^{(\bar{\kappa})} - \bar{x}_1\| \leq \delta$, $\|\mu^{(\bar{\kappa})} - \bar{\mu}_1\| \leq \delta$. By (C.3), and choosing τ as in (16), we have that

$$\|\mu^{(\bar{\kappa})} - D_{\mu}^{-1}T(y^{(\bar{\kappa})}, \rho, M) - \mu^{(\bar{\kappa})}\| < \frac{\zeta}{2} + \bar{K}\delta, \quad (\text{C.4})$$

where \bar{K} is a constant, emanating from the coefficient of δ in (C.3) when substituting for τ the upper-bound in (16). Note that $\bar{\kappa}$ is a function of δ , as it depends on $\bar{\kappa}$, which in turn depends on δ . Since δ is arbitrary, taking $\limsup_{\delta \rightarrow 0}$ in (C.4) and $\liminf_{\delta \rightarrow 0}$ in (C.2) leads to

$$\limsup_{\delta \rightarrow 0} \|\mu^{(\bar{\kappa})} - D_{\mu}^{-1}T(y^{(\bar{\kappa})}, \rho, M) - \mu^{(\bar{\kappa})}\| < \frac{\zeta}{2}, \quad (\text{C.5})$$

$$\liminf_{\delta \rightarrow 0} \|\mu^{(\bar{\kappa})} - D_{\mu}^{-1}T(y^{(\bar{\kappa})}, \rho, M) - \mu^{(\bar{\kappa})}\| > \frac{\zeta}{2}, \quad (\text{C.6})$$

establishing a contradiction. Then $\bar{\mu}_2, \bar{\mu}_1 \in \mathcal{M}_1$, i.e., all cluster points should be in the same subdomain of \mathcal{M} . As Lemma 3 establishes co-coercivity of T on each subdomain $X \times \mathcal{M}_j$, $j = 1, \dots, q$, it must be $\bar{\mu}_2 = \bar{\mu}_1$, i.e., Ω is a singleton, implying that Algorithm 1 converges. ■

Appendix D. Proof of Theorem 3

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

- (1) The subset I_1 of samples that form a minimal compression for x^* . Note that since Algorithm 1 converges to the point

(x^*, μ^*) for a fixed choice of M , $Q(\mu^*, M)$ will be a fixed quantity. Then Algorithm 1 will converge to a solution of

Find $x^* \in \hat{\Pi}_K$ such that

$$F(x^*)^T(x - x^*) \geq 0 \text{ for any } x \in \hat{\Pi}_K, \quad (\text{D.1})$$

where $\hat{\Pi}_K$ denotes the polytope obtained from Π_K by tightening at most M coupling constraints, as dictated by (12) with $Q(\mu^*, M)$. The constraints in (D.1) are equivalent to $F(x^*)^T x \geq F(x^*)^T x^*$ for all $x \in \hat{\Pi}_K$. Then, x^* is the minimiser of

$$\begin{aligned} & \min_{x \in \mathbb{R}^n} F(x^*)^T x \\ & \text{subject to } x \in \hat{\Pi}_K, \end{aligned} \quad (\text{D.2})$$

which is unique due to Lemma 1. Since the cost function is linear in x and $\hat{\Pi}_K$ is convex by Assumption 2, we obtain a scenario program as in Campi and Garatti (2008). Applying (Campi & Garatti, 2008, Prop. 1) or Margellos et al. (2015, Section III-B) to (D.2), we have that $|I_1| \leq nN$, i.e., the cardinality of a minimal compression set for x^* is bounded by the dimension of the decision vector nN .

- (2) The subset I_2 of samples whose corresponding coupling constraints intersect $\mathbb{B}(x^*, \rho)$. 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 $|I| = |I_1 \cup I_2| \leq |I_1| + |I_2| \leq nN + M$. Then, by Corollary 2 in Margellos et al. (2015), (20) follows. ■

References

- Acemoglu, D., & Jensen, M. K. (2013). Aggregate comparative statics. *Games and Economic Behavior*, 81, 27–49.
- Acemoglu, D., & Ozdaglar, A. (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1), 3–49.
- Aghassi, M., & Bertsimas, D. (2006). Robust game theory. *Mathematical Programming*, 107(1–2), 231–273.
- Alamo, T., Mirasierra, V., Dabbene, F., & Lorenzen, M. (2019). Safe approximations of chance constrained sets by probabilistic scaling. In *2019 18th European control conference ECC*, (pp. 1380–1385).
- Alamo, T., Tempo, R., & Camacho, E. (2009). Randomized strategies for probabilistic solutions of uncertain feasibility and optimization problems. *IEEE Transactions on Automatic Control*, 54(11), 2545–2559.
- Beckmann, M. J., McGuire, C. B., & Winsten, C. B. (1956). *Studies in the economics of transportation*. Yale University Press.
- Bopardikar, S. D., Borri, A., Hespanha, J. P., Prandini, M., & Di Benedetto, M. D. (2013). Randomized sampling for large zero-sum games. *Automatica*, 49(5), 1184–1194.
- Calafiore, G. C. (2010). Random convex programs. *SIAM Journal on Optimization*, 20(6), 3427–3464.
- Calafiore, G. C., & Campi, M. C. (2006). The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5), 742–753.
- Camerer, C. F. (2003). *Behavioral game theory: experiments in strategic interaction*. Princeton University Press.
- Campi, M. C., & Garatti, S. (2008). The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3), 1211–1230.
- Campi, M. C., & Garatti, S. (2018). Wait-and-judge scenario optimization. *Mathematical Programming*, 167(1), 155–189.
- Campi, M. C., Garatti, S., & Prandini, M. (2009). The scenario approach for systems and control design. *Annual Reviews in Control*, 33(2), 149–157.
- Campi, M. C., Garatti, S., & Ramponi, F. A. (2018). A general scenario theory for nonconvex optimization and decision making. *IEEE Transactions on Automatic Control*, 63(12), 4067–4078.
- Couchman, P., Kouvaritakis, B., Cannon, M., & Prashad, F. (2005). Gaming strategy for electric power with random demand. *IEEE Transactions on Power Systems*, 20(3), 1283–1292.
- de Farias, D. P., & Roy, B. V. (2003). The linear programming approach to approximate dynamic programming. *Operations Research*, 51, 850–865.
- Fabiani, F., Margellos, K., & Goulart, P. J. (2020). On the robustness of equilibria in generalized aggregative games. In *2020 59th IEEE conference on decision and control CDC*, (pp. 3725–3730).
- Fabiani, F., Margellos, K., & Goulart, P. J. (2022). Probabilistic feasibility guarantees for solution sets to uncertain variational inequalities. *Automatica*, 137, Article 110120.

- Facchinei, F., & Kanzow, C. (2010). Generalized Nash equilibrium problems. *Annals of Operations Research*, 175(1), 177–211.
- Facchinei, F., & Pang, J.-S. (2003). *Finite-dimensional variational inequalities and complementarity problems*. Springer-Verlag New York.
- Facchinei, F., Piccialli, V., & Sciandrone, M. (2011). Decomposition algorithms for generalized potential games. *Computational Optimization and Applications*, 50(2), 237–262.
- Fele, F., & Margellos, K. (2019). Probabilistic sensitivity of Nash equilibria in multi-agent games: A wait-and-judge approach. In *2019 IEEE 58th conference on decision and control CDC*, (pp. 5026–5031).
- Fele, F., & Margellos, K. (2021). Probably approximately correct Nash equilibrium learning. *IEEE Transactions on Automatic Control*, 66(9), 4238–4245.
- Grammatico, S., Parise, F., Colombino, M., & Lygeros, J. (2016). Decentralized convergence to Nash equilibria in constrained deterministic mean field control. *IEEE Transactions on Automatic Control*, 61(11), 3315–3329.
- Grammatico, S., Zhang, X., Margellos, K., Goulart, P., & Lygeros, J. (2016). A scenario approach for non-convex control design. *IEEE Transactions on Automatic Control*, 61(2), 334–345.
- Harker, P. T. (1991). Generalized Nash games and quasi-variational inequalities. *European Journal of Operational Research*, 54(1), 81–94.
- Hayashi, S., Yamashita, N., & Fukushima, M. (2005). Robust Nash equilibria and second-order cone complementarity problems. *Journal of Nonlinear and Convex Analysis*, 6.
- Jensen, M. K. (2010). Aggregative games and best-reply potentials. *Economic Theory*, 43(1), 45–66.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Kukushkin, N. S. (2004). Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48(1), 94–110.
- Kulkarni, A. A., & Shanbhag, U. V. (2012). On the variational equilibrium as a refinement of the generalized Nash equilibrium. *Automatica*, 48(1), 45–55.
- Lamnabhi-Lagarrigue, F., Annaswamy, A., Engell, S., Isaksson, A., Khargonekar, P., Murray, R. M., et al. (2017). Systems & control for the future of humanity, research agenda: Current and future roles, impact and grand challenges. *Annual Reviews in Control*, 43, 1–64.
- Mammarella, M., Mirasiera, V., Lorenzen, M., Alamo, T., & Dabbene, F. (2022). Chance-constrained sets approximation: A probabilistic scaling approach. *Automatica*, 137, 110108.
- Margellos, K., Prandini, M., & 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.
- Nash, J. F. (1950). Equilibrium points in N -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(48–49).
- Nishimura, R., Hayashi, S., & Fukushima, M. (2009). Robust Nash equilibria in n -person non-cooperative games: Uniqueness and reformulation. *Pacific Journal of Optimization*, 5(2), 237–259.
- Paccagnan, D., & Campi, M. C. (2019). The scenario approach meets uncertain game theory and variational inequalities. In *2019 IEEE 58th conference on decision and control CDC*, (pp. 6124–6129).
- Paccagnan, D., Gentile, B., Parise, F., Kamgarpour, M., & Lygeros, J. (2019). Nash and Wardrop equilibria in aggregative games with coupling constraints. *IEEE Transactions on Automatic Control*, 64(4), 1373–1388.
- Pantazis, G., Fele, F., & Margellos, K. (2020). A posteriori probabilistic feasibility guarantees for Nash equilibria in uncertain multi-agent games. *IFAC-PapersOnLine*, 53(2), 3403–3408, 21st IFAC World Congress.
- Pantazis, G., Fele, F., & Margellos, K. (2022). On the probabilistic feasibility of solutions in multi-agent optimization problems under uncertainty. *European Journal of Control*, 63, 186–195.
- Plott, C. R., & Smith, V. L. (Eds.). (2008). *Vol. 1, Handbook of experimental economics results*. Elsevier.
- Rubinstein, A. (1998). *Modeling bounded rationality*. MIT Press.
- Rudin, W. (1976). *Principles of mathematical analysis* (3rd ed.). McGraw-Hill, Inc.
- Saad, W., Han, Z., Poor, H. V., & 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.
- Schildbach, G., Fagiano, L., & Morari, M. (2012). Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23.
- Scutari, G., Facchinei, F., Pang, J., & Palomar, D. P. (2014). Real and complex monotone communication games. *Institute of Electrical and Electronics Engineers. Transactions on Information Theory*, 60(7), 4197–4231.
- Singh, V. V., Jouini, O., & Lisser, A. (2016). Existence of Nash equilibrium for chance-constrained games. *Operations Research Letters*, 44(5), 640–644.
- Wang, Y., Saad, W., Mandayam, N. B., & Poor, H. V. (2016). Load shifting in the smart grid: To participate or not?. *IEEE Transactions on Smart Grid*, 7(6), 2604–2614.
- Zhu, D. L., & Marcotte, P. (1996). Co-coercivity and its role in the convergence of iterative schemes for solving variational inequalities. *SIAM Journal on Optimization*, 6(3), 714–726.



Georgios Pantazis received the Diploma in electrical and computer engineering from the University of Patras, Greece, in 2017, and the DPhil in Engineering Science from the University of Oxford, United Kingdom, in 2022. He is currently a postdoctoral researcher at the Delft Center for Systems and Control (DCSC), TU Delft, the Netherlands. His research interests include optimisation and control of uncertain multi-agent systems with applications to smart grids and autonomous driving.



Filiberto Fele is Ramón y Cajal Research Fellow at the Department of Systems Engineering and Automation, University of Seville, Spain. He received the M.Sc. in Mechatronics Engineering from Politecnico di Torino, Italy, in 2008, and the Ph.D. in Control Engineering from the University of Seville, Spain, in 2017. He held postdoctoral appointments at the Electrical and Electronic Engineering Department, Imperial College London, in 2017, and the Department of Engineering Science, University of Oxford, 2018–2022, where he was Research Associate with the Energy and Power Group and the Control Group. From 2019 to 2022 he served as Lecturer at Christ Church and St Edmund Hall Colleges, Oxford. He is part of the EUCA Conference Editorial Board. His interests include data-driven robust optimisation and control of multi-agent systems, with a focus on the smart grid and the transition to net zero.



Kostas Margellos received the Diploma in electrical engineering from the University of Patras, Greece, in 2008, and the Ph.D. in control engineering from ETH Zurich, Switzerland, in 2012. He spent 2013, 2014 and 2015 as a postdoctoral researcher at ETH Zurich, UC Berkeley and Politecnico di Milano, respectively. In 2016 he joined the Control Group, Department of Engineering Science, University of Oxford, where he is currently an Associate Professor. He is also a Fellow of Reuben College and a Lecturer at Worcester College. He is currently serving as Associate Editor in *Automatica* and in the *IEEE Control Systems Letters*, and is part of the Conference Editorial Board of the IEEE Control Systems Society and EUCA. His research interests include optimisation and control of complex uncertain systems, with applications to energy and transportation networks.