Proceedings of Machine Learning Research vol 242:1702–1714, 2024 6th Annual Conference on Learning for Dynamics and Control

Probably approximately correct stability of allocations in uncertain coalitional games with private sampling

George Pantazis G.PANTAZIS@TUDELFT.NL
Delft Center for Systems and Control, TU Delft, The Netherlands

Filiberto Fele FFELE@US.ES
Department of Systems Engineering and Automation, University of Seville, Spain

Filippo Fabiani FILIPPO.FABIANI@IMTLUCCA.IT
IMT School for Advanced Studies Lucca, Italy

Sergio Grammatico S.GRAMMATICO@TUDELFT.NL
Delft Center for Systems and Control, TU Delft, The Netherlands

Kostas Margellos KOSTAS.MARGELLOS@ENG.OX.AC.UK
Department of Engineering Science, University of Oxford, UK

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

We study coalitional games with exogenous uncertainty in the coalition value, in which each agent is allowed to have private samples of the uncertainty. As a consequence, the agents may have a different perception of stability of the grand coalition. In this context, we propose a novel methodology to study the out-of-sample coalitional rationality of allocations in the set of stable allocations (i.e., the *core*). Our analysis builds on the framework of probably approximately correct learning. Initially, we state *a priori* and *a posteriori* guarantees for the entire core. Furthermore, we provide a distributed algorithm to compute a compression set that determines the generalization properties of the a posteriori statements. We then refine our probabilistic robustness bounds by specialising the analysis to a single payoff allocation, taking, also in this case, both *a priori* and *a posteriori* approaches. Finally, we consider a relaxed ζ-core to include nearby allocations and also address the case of empty core. For this case, probabilistic statements are given on the eventual stability of allocations in the ζ-core.

Keywords: Uncertain coalitional games; Statistical learning; Data privacy

1. Introduction

Multi-agent systems are pervasive across various fields, including engineering (Raja and Grammatico, 2021; Karaca and Kamgarpour, 2020; Han et al., 2019; Fele et al., 2017, 2018), economics, and social sciences (McCain, 2008). Even though agents often behave as self-interested entities, their limited ability to improve their utility can, in certain scenarios, motivate them to form coalitions to achieve a higher individual payoff. This situation can be modelled through coalitional games (Chalkiadakis et al., 2011). As each agent's participation in a coalition is subject to their own payoff being maximised, a question emerges about how to allocate the total value of the coalition such that no agent deviates from it. This problem is known as *stability* of agents' allocations.

In real-world scenarios, the value of a coalition is typically affected by uncertainty. The consideration of uncertainty in the coalitional values of a game finds its roots in seminal works such as Charnes and Granot (1976, 1977, 1973), which provided solution concepts to yield stable allocations against uncertainty. Suijs et al. (1999) discussed non-emptiness of the core for a particular class of stochastic games. Chalkiadakis and Boutilier (2004) and Ieong and Shoham (2008) addressed uncertainty through Bayesian learning methods, while Li and Conitzer (2015) investigated different solution concepts that maximize the probability of obtaining stable allocations. Chalkiadakis and Boutilier (2008) and Guéneron and Bonnet (2021) explored the dynamics of repeated stochastic coalitional games.

The connection between probably approximately correct (PAC) learning and uncertain coalitional games has been explored in Balcan et al. (2015). The spirit of Procaccia and Rosenschein (2006) is similar, using Vapnik-Chervonenkis (VC) theory to learn the winning coalitions for the class of the so-called *simple* games. Balcan et al. (2015) focused on a complementary problem where only a randomized subset of coalitions is considered. Both these works evaluate the sample complexity from the VC theory perspective, and hence their results suffer from the associated conservativeness. Pantazis et al. (2022a) leverage the scenario approach (Campi and Garatti, 2021, 2018a,b; Garatti and Campi, 2022) to provide distribution-free guarantees on the stability of allocations in a PAC manner, based on samples from the exogenous uncertainty affecting the value functions. On a parallel line of research, Pantazis et al. (2023) propose a data-driven Wasserstein-based distributionally robust approach for allocations' stability.

Unlike the aforementioned works, we consider a more general setting where the uncertainty data is privately drawn by each agent. As such, information is heterogeneous across agents, and the samples are regarded as a private resource. When data samples are commonly shared among agents, the work of Pantazis et al. (2022a) provide guarantees on stability of allocations. In particular, the developments in Pantazis et al. (2022a) rely on the notion of scenario core – a data-driven approximation of the robust core based on the hypothesis that all uncertainty samples are shared among the agents (Raja and Grammatico, 2021). Given this common set of samples, the perception of a stable allocation based on the available data is identical across agents. This is no longer the case for private sampling, as different samples of the exogenous uncertainty can result in a different perception of stability by each agent. We adopt a PAC learning approach and leverage the concept of compression (Margellos et al., 2015), i.e., the set of samples essential for the reconstruction of the scenario core and whose cardinality affects the generalization properties of the provided guarantees. We give *a priori* and *a posteriori* certificates for the entire scenario core obtained by private sampling, and propose a distributed algorithm to calculate a compression set.

We then focus on the specific allocation returned by some algorithm (akin, e.g., to the distributed payoff algorithm by Raja and Grammatico (2021)). We then provide *a priori* coalition stability guarantees for this allocation (Theorem 3). We show how the dimension of the problem – in our case the number of agents – plays a key role. Less conservative probabilistic bounds can often be obtained by taking an *a posteriori* approach (Campi et al., 2018), i.e., based on observation of the realized uncertainty. In this case, the probabilistic bounds can be significantly improved and even a tighter *a priori* bound can be obtained (see Theorem 4). Finally, we consider a relaxed ζ-core to include nearby allocations and also address the case of empty core. We leverage recents results by Campi and Garatti (2021) and study the probabilistic stability of allocations in the ζ-core.

2. Scenario approach for stochastic coalitional games

We consider a coalitional game with N agents identified by the set $\mathcal{N} = \{1,\dots,N\}$. We denote the number of possible subcoalitions, excluding the grand coalition, by M, i.e., $M = 2^N - 1$. The worth of a coalition $S \subseteq \mathcal{N}$ is given by the so called *value function*; here we consider transferable utility problems, i.e., the value of coalitions is expressed by a real value which can be split among participating agents. We posit that the value of each coalition is subject to uncertainty. Then, the value of a coalition $S \subset \mathcal{N}$ is a function $u_S : \Xi \to \mathbb{R}$ that given an uncertainty realization $\xi \in \Xi$ returns the total payoff for the agents in S. The value of the grand coalition is the deterministic quantity $u_{\mathcal{N}} \in \mathbb{R}$.

The uncertain coalitional game is then defined as the tuple $G_\Xi = \langle \mathcal{N}, \{u_S\}_{S\subseteq\mathcal{N}}, \Xi, \mathbb{P}\rangle$, where $\mathbb{P}$ is some unknown probability measure over Ξ. A vector $x := (x_i)_{i\in\mathcal{N}} \in \mathbb{R}^N$, where x_i is the payoff received by agent i, is called an *allocation*. For a given uncertainty realization $\xi \in \Xi$, an allocation is *strictly rational* for the members of S if $\sum_{i\in S} x_i > u_S(\xi)$; the latter implies that agents have an incentive to form this coalition for this particular uncertainty realization. An allocation is *efficient* if $\sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}}$. Efficient allocations such that there are no incentives for agents to deviate the grand coalition are called *stable*. The set of all such allocations is the *core* of the game.

The notion of robust core is proposed to account for the presence of uncertainty; see, e.g., Pantazis et al. (2022a); Raja and Grammatico (2021); Nedić and Bauso (2013). We formally define it as $C(G_\Xi) := \{x \in \mathbb{R}^N \colon \sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}}, \sum_{i\in S} x_i \geq \max_{\xi\in\Xi} u_S(\xi) \text{ for all } S \subset \mathcal{N}\}$. For any possible realization of the uncertainty $\xi \in \Xi$, any allocation $x \in C(G_\Xi)$ gives the agents no incentive to defect from the grand coalition and form sub-coalitions. Unfortunately, computing explicitly the robust core is hard, as we assume no knowledge on the uncertainty support Ξ (nor on the underlying probability distribution $\mathbb{P}$). To circumvent this challenge, we adopt a data-driven methodology and approximate the robust core by drawing a finite number K of independent and identically distributed (i.i.d.) samples $\boldsymbol{\xi} := (\xi^{(1)},\dots,\xi^{(K)}) \in \Xi^K$, where Ξ^K denotes the K-fold cartesian product of Ξ; we refer to vectors $\boldsymbol{\xi}$ as multi-samples. This constitutes the scenario game $G_K = \langle \mathcal{N}, \{u_S\}_{S\subseteq\mathcal{N}}, \boldsymbol{\xi}\rangle$, whose core is the set $\{x \in \mathbb{R}^N : \sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}} \text{ and } \sum_{i\in S} x_i \geq \max_{k=1,\dots K} u_S(\xi^{(k)}), \forall S \subset \mathcal{N}\}$, referred to as the *scenario core*.

We now take the notion of allocation stability a step further, considering the more general setting where every agent i has only access to a private set of samples $\boldsymbol{\xi}_i$ from Ξ. Let $\boldsymbol{\xi}_i \in \Xi^{K_i}$ be the multi-sample privately drawn by agent i. We say that an allocation x of G_K is *stable with respect to* $\boldsymbol{\xi}_i$ if $\sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}}$ and $\sum_{i\in S} x_i \geq \max_{k=1,\dots K_i} u_S(\xi_i^{(k)}), \forall S \subset \mathcal{N}$ s.t. $S \supseteq \{i\}$.

This immediately leads to the following extension of the scenario core:

Definition 1 *Let $\boldsymbol{\xi}_i = (\xi_i^{(1)},\dots,\xi_i^{(K_i)}) \in \Xi^{K_i}$ be some multi-sample drawn by agent i. The scenario core with private sampling is given by*

$$C(G_K) = \left\{ x \in \mathbb{R}^N : \sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}} \text{ and } \sum_{i\in S} x_i \geq \max_{i\in S} \max_{k=1,\dots K_i} u_S(\xi^{(k)}), \forall S \subset \mathcal{N} \right\}. \tag{1}$$

Unless differently specified, we assume the following conditions hold throughout the paper:

Assumption 1

(i) Each agent $i \in \mathcal{N}$ draws K_i independent samples from the probability distribution $\mathbb{P}$. The samples drawn by any agent are independent from those drawn by other agents.

(ii) $C(G_K)$ *is non-empty for any multi-sample* $(\xi_i)_{i\in\mathcal{N}} \in \Xi^K$.

In the following, we let $K = \sum_{i\in\mathcal{N}} K_i$, and $\xi = (\xi_i)_{i=1}^N$. Also, for brevity, we will use $S \supseteq \{i\}$ to denote all $S \subset \mathcal{N}$ which allow agent i as a member.

On the basis of privately available data, we wish to provide guarantees on the probability that allocations $x \in C(G_K)$ will remain stable (i.e., within $C(G_K)$) for any future, yet unseen, uncertainty realization. Capitalizing on Pantazis et al. (2022b); Fabiani et al. (2022); Pantazis et al. (2022a), we define two probabilistic notions of instability. In particular, the first refers to a particular allocation in the core, whereas the second involves the entire set $C(G_K)$.

Definition 2

(i) *Let* $V : \mathbb{R}^N \to [0,1]$. *For any* $x \in \mathbb{R}^N$, $V(x) \coloneqq \mathbb{P}\{\xi \in \Xi : \exists S \subset \mathcal{N}, \sum_{i\in S} x_i < u_S(\xi)\}$ *is the* probability of allocation instability.

(ii) *Let* $\mathbb{V} : 2^{\mathbb{R}^N} \to [0,1]$. *We call* $\mathbb{V}(C(G_K)) \coloneqq \mathbb{P}\{\xi \in \Xi : \exists x \in C(G_K), S \subset \mathcal{N} \colon \sum_{i\in S} x_i < u_S(\xi)\}$ probability of core instability.

$\mathbb{V}(C(G_K))$ thus denotes the probability with respect to the realizations of ξ that, for some S with value function $u_S(\xi)$, at least one of the allocations in the scenario core will become unstable, i.e., will be dominated by the option of defecting from the grand coalition to form S.

To bound the probability of core instability in a PAC fashion, we introduce two key concepts from statistical learning theory, namely the *algorithm* and the *compression set* (Margellos et al., 2015). The latter refers to the fact that only a subset of data from ξ may be sufficient to produce the same scenario core. As we will see later this underpins the quality of the provided probabilistic stability guarantees.

Definition 3 *A mapping* $\mathcal{A} : \Xi^K \to 2^{\mathbb{R}^N}$ *that takes as input a multi-sample* $\xi \in \Xi^K$ *and returns the scenario core of game* $C(G_K)$ *is called an* algorithm. *With* $\mathbb{P}^K$*-probability one w.r.t. the choice of* ξ*, a subset* $I \subseteq \{\xi^{(1)}, \dots, \xi^{(K)}\}$ *is a* compression set *for* ξ *if* $\mathcal{A}(I) = \mathcal{A}(\xi)$.

Any compression set of least cardinality is called *minimal* compression set. Another important notion used in our derivations is the *support rank* (Schildbach et al., 2012).

Definition 4 *Consider the maximal unconstrained subspace* $L \in \mathcal{L}$ *of a constraint in the form* $f(x,\xi) \le 0$, *i.e.,* $L' \subseteq L$ *for all* $L' \in \mathcal{L}$, *where* $\mathcal{L} = \bigcap_{\xi\in\Xi} \bigcap_{x\in\mathbb{R}^N} \{\bar{L}$ *is a linear subspace in* $\mathbb{R}^N$ *and* $\bar{L} \subset F(x,\xi)\}$, *with* $F(x,\xi) = \{\xi \in \mathbb{R} : f(x+\xi,\delta) = f(x,\xi)\}$. *Then, the support rank* ρ *of this constraint is given by* $\rho = N - dim(L)$.

In words, the support rank of an uncertain constraint is equal to the dimension of the problem at hand minus the dimension of the maximal unconstrained space of the constraint.

3. PAC stability guarantees for the scenario core with private sampling

3.1. A posteriori collective stability guarantees

In the following theorem we bound with high confidence the probability that some allocation $x \in C(G_K)$ will become unstable, i.e., that the scenario core – computed on the basis of $K = \sum_{i\in\mathcal{N}} K_i$ samples – will be reduced after a new uncertainty realization.

Theorem 1 *Suppose that each agent independently draws a multi-sample $\xi_i \in \Xi^{K_i}$ and let $\xi = (\xi_i)_{i=1}^N$. Fix a confidence parameter $\beta \in (0,1)$ and choose $\beta_i > 0$ such that $\sum_{i\in\mathcal{N}} \beta_i = \beta$. Consider an algorithm that takes as input ξ and returns the scenario core $C(G_K)$. It holds:*

$$\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{V}(C(G_K)) \le \sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K})\Big\} \ge 1-\beta, \tag{2}$$

where $s_{i,K}$ is the cardinality of the subset of a compression relative to agent i, quantified a posteriori, and ε_i satisfies

$$\varepsilon_i(K_i) = 1, \quad \sum_{k=1}^{K_i-1} \binom{K_i}{k} (1-\varepsilon_i(k))^{K_i-k} = \beta_i. \tag{3}$$

Proof: First, note that $\xi \in \Xi^K$ due to Assumption 1-(i). Then, the following inequalities hold.

$$\begin{aligned}
&\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{V}(C(G_K)) \le \sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K})\Big\} \\
=&\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists (i, S \supseteq \{i\}, x) : \sum_{i\in S} x_i < u_S(\xi)\} \le \sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K})\Big\} \\
=&\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\bigcup_{i\in\mathcal{N}} \{\xi \in \Xi : \exists (i, S \supseteq \{i\}, x) : \sum_{i\in S} x_i < u_S(\xi)\} \le \sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K})\Big\} \\
\ge&\mathbb{P}^K\Big\{\xi \in \Xi^K : \sum_{i\in\mathcal{N}} \mathbb{P}\{\xi \in \Xi : \exists (i, S \supseteq \{i\}, x) : \sum_{i\in S} x_i < u_S(\xi)\} \le \sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K})\Big\} \\
\ge&\mathbb{P}^K\Big\{\bigcap_{i\in\mathcal{N}} \{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists (i, S \supseteq \{i\}, x) : \sum_{i\in S} x_i < u_S(\xi)\} \le \varepsilon_i(s_{i,K})\}\Big\} \\
\ge&1-\sum_{i\in\mathcal{N}} \mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists (i, S \supseteq \{i\}, x) : \sum_{i\in S} x_i < u_S(\xi)\} > \varepsilon_i(s_{i,K})\Big\}
\end{aligned}$$

Given some new uncertainty realization ξ, let $C(G_{K+\xi})$ designate the scenario core built from the multisample $(\zeta^{(1)}, \dots, \zeta^{(K)}, \xi)$. The first equality stems from Definition 2, and expresses the fact that for $C(G_{K+\xi})$ to be a *strict* subset of $C(G_K)$, there must exist some allocation in $C(G_K)$ which, due to ξ, violates the rationality condition for some subcoalition S, i.e., $\sum_{i\in\mathcal{N}} x_i < u_S(\xi)$, causing the departure of some agent $i \in S$. The third inequality is obtained by applying the subadditivity property, while the second to last from applying Bonferroni's inequality (Margellos et al., 2018; Falsone et al., 2020). Note now that $C(G_K)$ can be found as the solution of the feasibility problem

$$\begin{aligned}
&\text{Find all } x \in \mathbb{R}^N \\
&\text{s.t. } \sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}}, \\
&\text{and } \sum_{i\in S} x_i \ge \max_{k=1,\dots K_i} u_S(\xi^{(k)}),\ \forall\, S \supseteq \{i\}, \forall i \in \mathcal{N},
\end{aligned} \tag{4}$$

Let $\mathscr{C}_\xi^i = \{C_i \subset 2^{R^N} : \sum_{i\in S} x_i \ge u_S(\xi), \forall S \supseteq \{i\}, \forall x \in C_i\}$ be the collection of (sub)sets of allocations satisfying the coalitional constraints obtained from data corresponding to agent $i \in \mathcal{N}$. Considering

Algorithm 1 Distributed compression algorithm

Input: Multi-sample ξ_i, coalition values $\{u_S(\cdot)\}_{S\supseteq\{i\}}$;
Output: Compression set I_i;
Initialization: $I_i=\varnothing$;
Each agent $i\in\mathcal{N}$ **performs**
  For all $S'\supseteq\{i\}$:

$$
\begin{aligned}
& x^*_{S'}\in\arg\min_{x\geq 0}\ 0\\
\text{s.t. } & \sum_{i\in S'} x_i=\max_{k=1,\dots,K_i} u_{S'}(\xi_i^{(k)}),\\
& \sum_{i\in S} x_i\geq \max_{k=1,\dots,K_i} u_S(\xi_i^{(k)}), \forall S\supseteq\{i\} \text{ and } S\neq S'.
\end{aligned}
$$

  If $x^*_{S'}\neq\varnothing$
    $I_i\leftarrow I_i\cup\arg\max_{k=1,\dots,K_i} u_{S'}(\xi_i^{(k)})$;
  End If
  End For

the unique set $\bar{C}_i=\{x\in\mathbb{R}^N:\sum_{i\in S}x_i\geq \max_{k=1,\dots K_i}u_S(\xi_i^{(k)}), \forall S\supseteq\{i\}\}$, note that $\bar{C}_i\in\bigcap_{k=1}^{K_i}\mathscr{C}^i_{\xi_i^{(k)}}$, i.e., $\bar{C}_i$ satisfies the consistency assumption required by Campi et al. (2018, Th. 1). From this,

$$
\begin{aligned}
&\mathbb{P}^K\Big\{\xi\in\Xi^K:\mathbb{P}\{\xi\in\Xi:\exists(S\supseteq\{i\},x):\sum_{i\in\mathcal{N}}x_i<u_S(\xi)\}>\varepsilon_i(s_{i,K})\Big\}\\
&\qquad\qquad\qquad\qquad=\mathbb{P}^K\Big\{\xi\in\Xi^K:\mathbb{P}\{\xi\in\Xi:\bar{C}_i\notin\mathscr{C}^i_{\xi}\}>\varepsilon_i(s_{i,K})\Big\}\leq\beta_i,
\end{aligned}
$$

from which (2) follows. ∎

A compression (sub)set I_i originated from agent's i uncertainty samples can be obtained by means of Algorithm 1; its cardinality can then be obtained as $s_{i,K}=|I_i|$. Note that differently from the compression algorithm in Pantazis et al. (2022a), Algorithm 1 is distributed among agents. At each iteration of Algorithm 1, a feasibility program is solved where coalitional rationality is enforced with equality for each coalition S' in which agent i is allowed to participate. This allows to identify whether the sample that maximizes $u_S(\cdot)$ is critical for the construction of a stable allocation set for agent $i\in\mathcal{N}$. The latter is verified if the problem is feasible, and the sample collected as part of the compression set I_i. Note that unless coordination is imposed among agents, the compression set obtained through Algorithm 1 is non-minimal.

3.2. A priori collective stability guarantees

As a byproduct of our *a posteriori* analysis, an *a priori* bound can be obtained. Specifically one can obtain an a priori bound by considering the worst case value of $\sum_{i\in\mathcal{N}}\varepsilon_i(s_{i,K})$ over all possible combinations for which $\sum_{i\in\mathcal{N}}s_{i,K}\leq M$ (where $M=2^N-1$). The following theorem provides then an *a priori* bound for the entire core with private samples.

Theorem 2 *Fix $\beta \in (0,1)$ and choose $\beta_i \in (0,1)$ such that $\sum_{i\in\mathcal{N}} \beta_i = \beta$. It holds that*

$$\mathbb{P}^K\{\xi \in \Xi^K : \mathbb{V}(C(G_K)) \leq \varepsilon^*\} \geq 1-\beta, \tag{5}$$

where $\varepsilon^* = \max\{\sum_{i\in\mathcal{N}} \varepsilon_i(s_i) \colon \sum_{i\in\mathcal{N}} s_i \leq M,\, s_i \in \mathbb{N}\}$, with ε_i defined as in (3), is a quantity independent from the given multi-sample.

Proof: We apply the a posteriori result in Pantazis et al. (2022b, Th. 1) – which admits the number of facets of a randomized feasibility domain as an upper bound to the cardinality of the minimal compression set – to $C(G_K)$, where each facet is associated to some subcoalition in $2^{\mathcal{N}}$. Now, let $s_{i,K}$ denote the cardinality of the minimal (sub)compression set relative to agent i, quantified *a posteriori*: by Pantazis et al. (2022b, Th.1) it holds that $\sum_{i\in\mathcal{N}} s_{i,K} \leq M$. Then, $\sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K}) \leq \max_{\{s_i\}_{i\in\mathcal{N}}} \sum_{i\in\mathcal{N}} \varepsilon_i(s_i)$, for any $\{s_i\}_{i\in\mathcal{N}}$ such that $\sum_{i\in\mathcal{N}} s_i \leq M$. By definition of ε^* it follows

$$\mathbb{P}^K\{\xi \in \Xi^K : \mathbb{V}(C(G_K)) \leq \varepsilon^*\} \geq \mathbb{P}^K\left\{\xi \in \Xi^K : \mathbb{V}(C(G_K)) \leq \sum_{i\in\mathcal{N}} \varepsilon_i(s_{i,K})\right\} \geq 1-\beta,$$

where the last inequality holds because of Theorem 1. ■

The results above can be applied to the entire set of allocations that are stable w.r.t. the observed data; because of their generality, these theoretical guarantees tend to be conservative. In what follows we specialise our analysis to a single allocation.

4. PAC stability of a single allocation with private sampling

Let x^* be the unique allocation in $C(G_K)$ which maximises some convex utility function $f(\cdot)$.[1] Recalling the notion of support rank as per Definition 4, the following result holds for x^*.

Theorem 3 *Suppose that each agent draws a multi-sample $\xi_i \in \Xi^{K_i}$ and let $\xi = (\xi_i)_{i\in\mathcal{N}}$. Fix $\varepsilon \in (0,1)$, and choose $\varepsilon_i > 0$ such that $\sum_{i\in\mathcal{N}} \varepsilon_i = \varepsilon$. Then,*

$$\mathbb{P}^K\big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : x^* \notin C(G_K)\} \leq \varepsilon\big\} \geq 1 - \sum_{i\in\mathcal{N}} \beta_i, \tag{6}$$

where $\beta_i = \sum_{j=1}^{\rho_i} \binom{K_i}{j} \varepsilon_i^j (1-\varepsilon_i)^{K_i - j}$, with $\rho_i \leq N$ being the support rank of the coalitional rationality constraints corresponding to agent i, i.e., relative to all $S \subset \mathcal{N}$ such that $S \supseteq \{i\}$.

1. Uniqueness of the maximiser is without loss of generality. If multiple maximisers are possible, a convex tie-break rule can be applied to single out a unique solution.

Proof: Similarly to the proofline of Theorem 1, we have that

$$\begin{aligned}
&\mathbb{P}^K\Big\{\xi \in \Xi^K : V(x^*) \leq \varepsilon\Big\} \\
=&\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists (i \in \mathcal{N}, S \supset \{i\}) : \sum_{i\in S} x_i^* < u_S(\xi)\} \leq \sum_{i\in\mathcal{N}} \varepsilon_i\Big\} \\
=&\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\bigcup_{i\in\mathcal{N}} \{\xi \in \Xi : \exists S \supseteq \{i\} : \sum_{i\in S} x_i^* < u_S(\xi)\}\} \leq \sum_{i\in\mathcal{N}} \varepsilon_i\Big\} \\
\geq&\mathbb{P}^K\Big\{\xi \in \Xi^K : \sum_{i\in\mathcal{N}} \mathbb{P}\{\xi \in \Xi : \exists S \supseteq \{i\} : \sum_{i\in S} x_i^* < u_S(\xi)\} \leq \sum_{i\in\mathcal{N}} \varepsilon_i\Big\} \\
\geq&\mathbb{P}^K\Big\{\bigcap_{i\in\mathcal{N}} \{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists S \supseteq \{i\} : \sum_{i\in S} x_i^* < u_S(\xi)\} \leq \varepsilon_i\}\Big\} \\
\geq&1 - \sum_{i\in\mathcal{N}} \mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists S \supseteq \{i\} : \sum_{i\in S} x_i^* < u_S(\xi)\} > \varepsilon_i\Big\} \geq 1 - \sum_{i\in\mathcal{N}} \beta_i.
\end{aligned}$$

To obtain the last inequality consider the following optimization program

$$\begin{aligned}
& x^* = \operatorname*{arg\,min}_{x\in\mathbb{R}^N} f(x) \\
& \text{s.t. } \sum_{i\in\mathcal{N}} x_i = u_{\mathcal{N}}, \\
& \text{and } \sum_{i\in S} x_i \geq \max_{k=1,\ldots K_i} u_S(\xi_i^{(k)}),\ \forall S \supseteq \{i\},\ \forall i \in \mathcal{N}.
\end{aligned} \tag{7}$$

We then group the constraints depending on which agent they correspond to. Let $X_i(\xi) \coloneqq \{x \in \mathbb{R}^N\colon \sum_{i\in S} x_i \geq u_S(\xi), \forall S \supseteq \{i\}\}$. With β_i defined as above, from Schildbach et al. (2012) we have

$$\begin{aligned}
\beta_i &\geq \mathbb{P}^K\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : x^* \notin X_i(\xi)\} > \varepsilon_i\} \\
&= \mathbb{P}^K\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists (i \in \mathcal{N}, S \supseteq \{i\}) : \sum_{i\in S} x_i^* < u_S(\xi)\} > \varepsilon_i\},
\end{aligned}$$

which concludes the proof. ■

The confidence parameter β_i in Theorem 3 depends on the number of samples K_i of agent $i \in \mathcal{N}$ and the violation level ε_i set individually by the agent to determine the probability of satisfaction of the rationality constraints. Most importantly, β_i depends on the so called support rank ρ_i of the coalitional constraints corresponding to $i \in \mathcal{N}$, which is in any instance upper bounded by the dimension of the decision variable in (7). Theorem 3 is an *a priori* result and in certain cases can be conservative. The following result provides an improved probabilistic bound, based on the observation of the uncertainty realization. Let

$$\varepsilon_i(s) = 1 - \left(\frac{\beta_i}{(N+1)\binom{K_i}{s}}\right)^{\frac{1}{K_i - s}}. \tag{8}$$

Note that (8) is compatible with (3), and allows to explicitly compute ε_i as a function of β_i.

Theorem 4 *Fix $\beta \in (0,1)$ and choose $\beta_i > 0$ such that $\sum_{i\in\mathcal{N}} \beta_i = \beta$. Set ε_i as in (8). Then, for the unique allocation $x^* \in C(G_K)$ it holds*

$$\mathbb{P}^K\Big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : x^* \notin C(G_K)\} \leq \sum_{i\in S} \varepsilon_i(s_{i,K})\Big\} \geq 1 - \beta. \tag{9}$$

Recall that, as in Theorem 4, $s_{i,K}$ is the cardinality of the (sub)compression set associated with the rationality constraints relative to all $S \supseteq \{i\}$, quantified a posteriori, e.g., by means of the procedure illustrated in Campi et al. (2018, §II).

Proof: For each agent $i \in \mathcal{N}$, taking the conditional probability over all other agents' multi-samples, Theorem 1 in Campi et al. (2018) allows to state

$$\mathbb{P}^K\big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \sum_{i\in S} x_i^* > u_S(\xi)\} \leq \varepsilon_i(s_{i,K}) | (\xi_j)_{j\neq i} \in \Xi^{K-K_i}\big\} \geq 1-\beta_i,$$

which, by integrating w.r.t. the probability of realization of all other agents' scenarios, becomes

$$\mathbb{P}^K\big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \sum_{i\in S} x_i^* > u_S(\xi)\} \leq \varepsilon_i(s_{i,K})\big\} \geq 1-\beta_i. \tag{10}$$

The proof is concluded by applying relation (10) in the proofline of Theorem 3. ■

Finally, an a priori bound can be derived from Theorem 4 by considering the worst case among all a posteriori observable compressions $\{s_{i,K}\}$; this bound is nontrivial by noticing that $\sum_{i\in\mathcal{N}} s_{i,K} \leq N$, which follows from Calafiore and Campi (2006, Th. 3).

Corollary 1 *Let $\varepsilon^* = \max\{\sum_{i\in\mathcal{N}} \varepsilon_i(s_i) :\ \sum_{i\in\mathcal{N}} s_i \leq N,\ s_i \in \mathbb{N}\}$, with ε_i defined as in (8). Then*

$$\mathbb{P}^K\big\{\xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : x^* \notin C(G_K)\} \leq \varepsilon^*\big\} \geq 1-\beta. \tag{11}$$

4.1. The case of empty core

Lifting Assumption 1-(ii) on non-emptiness of the scenario core, we define a relaxed version of the latter, the so called ζ-core for the case of private uncertainty sampling as follows:

Definition 5 *For any $(\xi_i)_{i\in\mathcal{N}} \in \Xi^K$, fixed $\bar{\zeta}_i \geq 0$ for all $i \in \mathcal{N}$, the scenario ζ-core is given by*

$$C_\zeta(G_K) = \Big\{x \in \mathbb{R}^N : \sum_{i\in\mathcal{N}} x_i = u_\mathcal{N} \text{ and } \sum_{i\in S} x_i \geq \max_{i\in S} \max_{k=1,\dots,K_i} u_S(\xi_i^{(k)}) - \bar{\zeta}_i,\ \forall S \subset \mathcal{N}\Big\}. \tag{12}$$

It is worth pointing out that the ζ-core allows to extend the analysis to allocations "closest" to the core, and the interest in it may not be necessarily restricted to cases where the standard core is empty (this is recovered by setting $\bar{\zeta}_i = 0$ for all $i \in \mathcal{N}$). Also, Definition 5 contemplates a different relaxation parameter $\bar{\zeta}_i$ for every agent, allowing to specialise the definition according to individual preferences, or possibly different information on, e.g., each agent's sampling.

On these grounds, here we address the problem of i) providing an allocation with (relaxed) coalitional stability certificates and ii) measuring how agents' multi-samples contribute to lack of coalitional stability. These questions can be answered at once by solving

$$\min_{x,\{\zeta_i\geq 0\}_{i\in\mathcal{N}}} \sum_{i\in\mathcal{N}} \sum_{k=1}^{K_i} \zeta_i^{(k)}, \tag{13a}$$

$$\text{s.t.} \sum_{i\in\mathcal{N}} x_i = u_\mathcal{N} \tag{13b}$$

$$\text{and} \sum_{i\in S} x_i \geq \max_{k=1,\dots,K_i} u_S(\xi^{(k)}) - \zeta_i^{(k)},\ \forall S \supseteq \{i\},\ \forall i \in \mathcal{N}. \tag{13c}$$

In what follows, we consider without loss of generality that for any multi-sample $(\xi_i)_{i\in\mathcal{N}} \in \Xi^K$, (13) returns a unique pair $(x^*, \zeta^*) \in \mathbb{R}^N \times \mathbb{R}^K$ (this can be ensured by using a convex tie-break rule).

Assumption 2 *For every allocation $x \in \mathbb{R}^N$, $\mathbb{P}\{\xi \in \Xi : \sum_{i \in S} x_i = u_S(\xi)\} = 0$ for any $S \subset \mathcal{N}$.*

Assumption 2 is related to non-degeneracy and is often satisfied when ξ does not accumulate as is the case for continuous probability distributions. Now, for $i \in \mathcal{N}$, consider the polynomial equation (Campi and Garatti, 2021, Th. 1)

$$\binom{K_i}{s} t^{K_i - s} - \frac{\beta_i}{2N} \sum_{j=s}^{K_i - 1} \binom{i}{s} t^{j-s} - \frac{\beta_i}{6K_i} \sum_{j=K_i+1}^{4K_i} \binom{j}{s} t^{j-s} = 0 \tag{14}$$

For $s \in \{0, \ldots, K_i - 1\}$, (14) has two solutions in $[0, +\infty)$. We denote the smallest as $\underline{t}_i(s)$, and further let $\underline{t}_i(s) = 0$ for $s = K_i$. We can now propose the following a posteriori statement on x^*:

Theorem 5 *Fix $\beta \in (0,1)$ and choose $\beta_i > 0$ such that $\sum_{i \in \mathcal{N}} \beta_i = \beta$. Let (x^*, ζ^*) be the solution of (13), and denote by s_i^* the cardinality of $\{k\colon \zeta_i^{*(k)} > 0\}$. Under Assumption 2,*

$$\mathbb{P}^K \{ \xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : \exists (i, S \supseteq \{i\}) : \sum_{i \in S} x_i^* < u_S(\xi)\} \leq \sum_{i \in \mathcal{N}} \bar{\varepsilon}_i(s_i^*)\} \geq 1 - \beta, \tag{15}$$

where $\bar{\varepsilon}_i(s_i^*) = 1 - \underline{t}_i(s_i^*)$, for all $i \in \mathcal{N}$.

Proof: Let $f_i(x, \xi) := \max_{j=1,\ldots,M}(b_j(\xi) - a_j^\top x)$, where the i-th component of a_j is 1 if $i \in S$ (and 0 otherwise), and $b_j(\xi) = (u_S(\xi))_{S \supseteq \{i\}}$. Then (13c) can be rewritten as

$$f_i(x, \xi_i^{(k)}) \leq \zeta_i, \ \forall k = 1, \ldots, K_i, \ \forall i \in \mathcal{N}.$$

We note that in our setting Campi and Garatti (2021, Assum. 1) is satisfied; under Assumption 2, it is then possible to apply Campi and Garatti (2021, Th. 1), which yields

$$\mathbb{P}^K \left\{ \xi \in \Xi^K : \mathbb{P}\{\xi \in \Xi : f_i(x^*, \xi) > 0\} > \bar{\varepsilon}_i(s_i^*) \right\} \leq \beta_i, \tag{16}$$

where $\bar{\varepsilon}_i(\cdot)$ is derived from (14) as described above. By applying arguments similar to those used in proving Theorems 3 and 4, we obtain (15). ■

The statement in Theorem 5 can be interpreted as follows: as a result of (13), x^* is an allocation in the ζ-core defined by $\{\bar{\zeta}_i\}_{i \in \mathcal{N}}$, with $\bar{\zeta}_i = \max_{k=1,\ldots,K_i} \zeta_i^{*(k)}$. It holds with confidence $1 - \beta$ that x^* is a coalitionally stable allocation with probability $1 - \sum_{i \in \mathcal{N}} \bar{\varepsilon}_i$, under all possible realizations of the uncertain parameter $\xi \in \Xi$.

5. Conclusion

In this work we considered uncertain coalitional games and proposed a data-driven methodology to study the stability of allocations for the general setting where uncertainty is privately sampled by the agents. Future work will investigate stability of allocations through a distributionally robust framework, as well as analyse the case where the value of the grand coalition can also be affected by uncertainty.

Acknowledgments

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 S. Grammatico and G. Pantazis has been partially supported by the ERC under project COSMOS (802348).

References

Maria-Florina Balcan, Ariel D. Procaccia, and Yair Zick. Learning cooperative games. In *IJCAI15: Proceedings of the 24th International Conference on Artificial Intelligence*, 2015.

Giuseppe C. Calafiore and Marco C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006. doi: 10.1109/TAC.2006.875041.

Marco C. Campi and Simone Garatti. Introduction to the scenario approach. *Society for Industrial and Applied Mathematics*, 2018a.

Marco C. Campi and Simone Garatti. Wait-and-judge scenario optimization. *Mathematical Programming*, 167:155–189, 2018b.

Marco C. Campi and Simone Garatti. A theory of the risk for optimization with relaxation and its application to support vector machines. *Journal of Machine Learning Research*, 22(288):1–38, 2021.

Marco C. Campi, Simone Garatti, and Federico Alessandro Ramponi. A general scenario theory for nonconvex optimization and decision making. *IEEE Transactions on Automatic Control*, 63(12): 4067–4078, 2018. doi: 10.1109/TAC.2018.2808446.

Georgios Chalkiadakis and Craig Boutilier. Bayesian reinforcement learning for coalition formation under uncertainty. In *AAMAS 2004: Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems.*, pages 1090–1097, 2004.

Georgios Chalkiadakis and Craig Boutilier. Sequential decision making in repeated coalition formation under uncertainty. In *AAMAS08: Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, 2008.

Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(6): 1–168, 2011.

Abraham Charnes and Daniel Granot. Prior solutions: Extensions of convex nucleus solutions to chance-constrained games. In *Proceedings of the Computer Science and Statistics Seventh Symposium at Iowa State University*, pages 323–332, 1973.

Abraham Charnes and Daniel Granot. Coalitional and chance-constrained solutions to n-person games I: The prior satisficing nucleolus. *SIAM Journal on Applied Mathematics*, 31(2):358–367, 1976.

Abraham Charnes and Daniel Granot. Coalitional and chance-constrained solutions to n-person games II: Two-stage solutions. *Operations Research*, 25(6):1013–1019, 1977.

Filippo Fabiani, Kostas Margellos, and Paul J. Goulart. Probabilistic feasibility guarantees for solution sets to uncertain variational inequalities. *Automatica*, 137:110120, 2022. ISSN 0005-1098.

Alessandro Falsone, Kostas Margellos, Maria Prandini, and Simone Garatti. A scenario-based approach to multi-agent optimization with distributed information. *IFAC-PapersOnLine*, 53(2): 20–25, 2020. ISSN 2405-8963. doi: 10.1016/j.ifacol.2020.12.034. 21st IFAC World Congress.

Filiberto Fele, José M. Maestre, and Eduardo F. Camacho. Coalitional control: Cooperative game theory and control. *IEEE Control Systems Magazine*, 37(1):53–69, 2017.

Filiberto Fele, Ezequiel Debada, José M. Maestre, and Eduardo F. Camacho. Coalitional control for self-organizing agents. *IEEE Transactions on Automatic Control*, 63(9):2883–2897, 2018. doi: 10.1109/TAC.2018.2792301.

Simone Garatti and Marco C. Campi. Risk and complexity in scenario optimization. *Mathematical Programming*, 191:243–279, 2022.

Josselin Guéneron and Grégory Bonnet. Are exploration-based strategies of interest for repeated stochastic coalitional games? *Advances in Practical Applications of Agents, Multi-Agent Systems, and Social Good. The PAAMS Collection*, pages 89–100, 2021.

Liyang Han, Thomas Morstyn, and Malcolm McCulloch. Incentivizing prosumer coalitions with energy management using cooperative game theory. *IEEE Transactions on Power Systems*, 34 (1):303–313, 2019.

Samuel Ieong and Yoav Shoham. Bayesian coalitional games. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, pages 95–100, 2008.

Orcun Karaca and Maryam Kamgarpour. Core-selecting mechanisms in electricity markets. *IEEE Transactions on Smart Grid*, 11(3):2604–2614, 2020. doi: 10.1109/TSG.2019.2958710.

Yuqian Li and Vincent Conitzer. Cooperative game solution concepts that maximize stability under noise. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, page 979–985, 2015.

Kostas Margellos, Maria Prandini, and John Lygeros. On the connection between compression learning and scenario-based single-stage and cascading optimization problems. *IEEE Transactions on Automatic Control*, 60(10):2716–2721, 2015.

Kostas Margellos, Alessandro Falsone, Simone Garatti, and Maria Prandini. Distributed constrained optimization and consensus in uncertain networks via proximal minimization. *IEEE Transactions on Automatic Control*, 63(5):1372–1387, 2018. doi: 10.1109/TAC.2017.2747505.

Roger A. McCain. Cooperative games and cooperative organizations. *The Journal of Socio-Economics*, 37(6):2155–2167, 2008. ISSN 1053-5357.

Angelia Nedić and Dario Bauso. Dynamic coalitional tu games: Distributed bargaining among players' neighbors. *IEEE Transactions on Automatic Control*, 58(6):1363–1376, 2013. doi: 10.1109/TAC.2012.2236716.

George Pantazis, Filippo Fabiani, Filiberto Fele, and Kostas Margellos. Probabilistically robust stabilizing allocations in uncertain coalitional games. *IEEE Control Systems Letters*, 6:3128–3133, 2022a. doi: 10.1109/LCSYS.2022.3182152.

George Pantazis, Filiberto Fele, and Kostas Margellos. On the probabilistic feasibility of solutions in multi-agent optimization problems under uncertainty. *European Journal of Control*, 63:186–195, 2022b. ISSN 0947-3580.

George Pantazis, Barbara Franci, Sergio Grammatico, and Kostas Margellos. Distributionally robust stability of payoff allocations in stochastic coalitional games. *Accepted for publication at the IEEE Conference on Decision and Control 2023*, 2023.

Ariel D. Procaccia and Jeffrey S. Rosenschein. Learning to identify winning coalitions in the PAC model. In *AAMAS06 - 5th International Joint Conference on Autonomous Agents and Multi-agent Systems*, page 673–675. Association for Computing Machinery, NY, US, 2006. doi: 10.1145/1160633.1160751.

Aitazaz Ali Raja and Sergio Grammatico. Payoff distribution in robust coalitional games on time-varying networks. *IEEE Transactions on Control of Network Systems*, 2021.

Georg Schildbach, Lorenzo Fagiano, and Manfred Morari. Randomized solutions to convex programs with multiple chance constraints. *SIAM J. Optim.*, 23:2479–2501, 2012.

Jeroen Suijs, Peter Borm, Anja De Waegenaere, and Stef Tijs. Cooperative games with stochastic payoffs. *European Journal of Operational Research*, 113(1):193–205, 1999. ISSN 0377-2217. doi: 10.1016/S0377-2217(97)00421-9.