

Distributionally robust stability of payoff allocations in stochastic coalitional games

George Pantazis, Barbara Franci, Sergio Grammatico and Kostas Margellos

Abstract—We consider multi-agent coalitional games with uncertainty in the coalitional values. We provide a novel methodology to study the stability of the grand coalition in the case where each coalition constructs ambiguity sets for the (possibly) unknown probability distribution of the uncertainty. As a less conservative solution concept compared to worst-case approaches for coalitional stability, we consider a stochastic version of the so-called core set, i.e., the expected value core. Unfortunately, without exact knowledge of the probability distribution, the evaluation of the expected value core is an extremely challenging task. Hence, we propose the concept of distributionally robust (DR) core. Leveraging tools from data-driven DR optimization under the Wasserstein distance, we provide finite-sample guarantees that any allocation which lies in the DR core is also stable with respect to the true probability distribution. Furthermore, we show that as the number of samples grows unbounded, the DR core converges almost surely to the true expected value core. We dedicate the last section to the computational tractability of finding an allocation in the DR core.

I. INTRODUCTION

Coalitional games [1] are prevalent in applications ranging from engineering [2]–[4] to economics and social sciences [5]. Even though agents in such systems typically act as selfish entities, they are incentivized to form coalitions aiming at receiving higher individual gains or reducing their own costs. A challenging task, due to the agents’ individual interests, is to distribute their payoffs in such a way that none of them has an incentive to deviate from the so-called *grand coalition*, i.e., the coalition where all agents work together. In the literature of coalitional game theory this problem is known as *stability of the grand coalition* and the set of payoffs for which stability is achieved is known as the *core* of the game. Due to its conceptual simplicity, the core has been widely used as a stability concept in coalitional games [1] and in turn intense research has been dedicated to finding allocations that lie within the core.

Stability of the grand coalition is fundamentally connected to the values of each coalition. However, coalitional values are typically subject to uncertainty. As such, the mathematical framework of deterministic coalitional games needs to be revisited and extended. The seminal works [6]–[8] are the first on stochastic coalitional games. The work in

[9] also studies uncertain coalitional games and shows that for a particular class, certain properties of the game, such as the non-emptiness of the core continue to hold when uncertainty is introduced. Uncertain coalitional games were studied under the lenses of Bayesian learning in [10], [11], while the work in [12] investigates which stability solution concepts maximize the probabilistic stability of allocations after the samples of the uncertainty have been revealed. Moreover, [13] and [14] focus on the dynamic evolution of repeated stochastic games.

In [2], the concept of the so called *robust core* is introduced as a generalization to the traditional deterministic core. In this setting, the range of possible coalitional values is assumed to be known. The work in [15] extends the notion of the robust core to that of the *scenario core* accounting for the more general case where both the support set and the probability distribution of the uncertainty affecting the coalitional value are unknown. As an alternative to the robust core [2] and its data-driven counterpart [15], in this paper we consider instead the significantly milder concept of stability in the mean sense that in turn gives rise to the so-called *expected value core*. Studying allocation stability in the mean sense might circumvent the possible emptiness of the core set, a fundamental technical challenge in coalitional game theory. Apart from very mild assumptions on the probability distribution of the uncertainty, here we consider both the support set and the probability distribution to be unknown. In other cases, the uncertain parameter affecting the coalitional game might not even admit a single distribution, but a range of possible distributions, quantified through data-driven approaches. As such, evaluating the expected core in this setting is extremely challenging.

To address this, we follow an approach based on distributionally robust (DR) optimization [16]–[22] thus considering ambiguity sets that represent empirical sets in which the true probability distribution (in case the uncertainty admits one) is likely to be contained. The consideration of ambiguity sets leads to allocations that are *distributionally stable*. We call the set of all distributionally stable allocations the *distributionally robust (DR) core* of the game.

Leveraging results from data-driven DR optimization under the Wasserstein distance [23], [24], we provide finite sample guarantees on the probability that any allocation in the DR core of the DR game approximation is also in the expected value core of the original game with a given confidence (Section III). Moreover, we prove almost-sure asymptotic convergence of the Wasserstein DR core to the expected core of the original game (Section IV.A).

The authors are with Delft Center for Systems & Control, TU Delft ({g.pantazis, s.grammatico}@tudelft.nl), the Department of Advanced Computing Sciences in Maastricht University (b.franci@maastrichtuniversity.nl) and the Department of Engineering Science, University of Oxford, (kostas.margellos@eng.ox.ac.uk). This work has been partially funded by the ERC under project COSMOS (802348).

Finally, we provide the means to calculate an allocation in the Wasserstein DR core (Section IV.B). Specifically, we show that under certain conditions, the problem of finding such an allocation can be recast as a convex optimization problem, whose complexity both in the number of decision variables and constraints is inherently connected to the number of possible subcoalitions. Numerical simulations corroborate our theoretical findings (Section V).

II. STOCHASTIC COALITIONAL GAMES

A. Allocation mean stability and expected value core

We consider a coalitional game with N agents parameterized by the index set $\mathcal{N} = \{1, \dots, N\}$. We denote the number of possible subcoalitions except for the grand coalition by M , i.e., $M = 2^N - 1$. In this setting, the agents, though selfish, wish to form coalitions if that implies an increase in their individual payoffs. The total gain for each coalition is given by the so called *value function*, which, depending on the coalition $S \subseteq \mathcal{N}$, takes a real value representing the total payoff that agents participating in it would obtain from its formation. However, the values of each coalition are subject to uncertainty thus rendering the value function of each coalition stochastic.

Definition 1: (Stochastic value function) The value function of a coalition $S \subset \mathcal{N}$ is a function $u_S : 2^{\mathcal{N}} \times \Xi \rightarrow \mathbb{R}$ that, given the value of the uncertainty realization $\xi \in \Xi \subseteq \mathbb{R}^p$, returns the total payoff for the agents forming a coalition S . The value function of the grand coalition is deterministic, i.e., $u_{\mathcal{N}} : 2^{\mathcal{N}} \rightarrow \mathbb{R}$. \square

An uncertain coalitional game is then defined as the tuple $G_{\mathbb{P}} = \{\mathcal{N}, \{u_S\}_{S \subseteq \mathcal{N}}, \Xi, \mathbb{P}\}$, where \mathbb{P} denotes the probability distribution that the uncertainty $\xi \in \Xi$ follows. To circumvent the fundamental issue of the emptiness of the robust core as defined in [2], let us consider the concept of stability of allocations in the mean sense, defined as follows.

Definition 2: (Stability in the mean sense) An allocation $x = (x_i)_{i \in \mathcal{N}}$ of the game $G_{\mathbb{P}} = \{\mathcal{N}, \{u_S\}_{S \subseteq \mathcal{N}}, \Xi, \mathbb{P}\}$ is *stable in the mean sense* if i) $\sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}}$ and ii) $\sum_{i \in S} x_i \geq \mathbb{E}_{\mathbb{P}}[u_S(\xi)], \forall S \subset \mathcal{N}$. \square

The first condition is called the efficiency condition. Due to our assumption that the grand coalition is deterministic, it means that the total increase in gains when all agents work together is known with certainty. This is the case when agents might know how efficient a fully-cooperative scheme is but have some level of uncertainty/ambiguity with respect to the potential outcomes of the subcoalitions. The second condition implies that the allocation x is not strictly feasible, hence agents do not have an incentive to form S . Otherwise, if $\sum_{i \in S} x_i < \mathbb{E}_{\mathbb{P}}[u_S(\xi)]$ agents would have the incentive to leave the grand coalition and form S , thus receiving $\mathbb{E}_{\mathbb{P}}[u_S(\xi)]$ as opposed to $\sum_{i \in S} x_i$. In this setting, we wish to study the stability of the grand coalition, where no agent has an incentive to deviate and create other subcoalitions. To this end, let us introduce an extension to the classic notion of the core, the so-called *expected value core*, which is the set of all stable allocations in the mean sense as defined next.

Definition 3: (Expected value core) The expected value core $C_E(G_{\mathbb{P}})$ of the game $G_{\mathbb{P}}$ is defined as the set

$$C_E(G_{\mathbb{P}}) = \{x \in \mathbb{R}^N : \sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}}, \sum_{i \in S} x_i \geq \mathbb{E}_{\mathbb{P}}[u_S(\xi)], \forall S \subset \mathcal{N}\},$$

where

$$\mathbb{E}_{\mathbb{P}}[u_S(\xi)] = \int_{\Xi} u_S(\xi) \mathbb{P}(d\xi). \quad (1)$$

\square

We now impose the following mild technical assumptions.

Assumption 1: (Light-tailed probability distribution) For the true probability measure \mathbb{P} , there exists $a > 1$ such that

$$A := \mathbb{E}_{\mathbb{P}}[\exp(\|\xi\|^a)] = \int_{\Xi} \exp(\|\xi\|^a) \mathbb{P}(d\xi) < \infty.$$

Assumption 2: For the probability measure \mathbb{P} it holds that $\mathbb{E}_{\mathbb{P}}[\|\xi\|] = \int_{\Xi} \|\xi\| \mathbb{P}(d\xi) < \infty$.

Assumption 1 requires the tail of the true probability distribution \mathbb{P} to decay at an exponential rate. In case Ξ is a compact set this assumption trivially holds. Assumption 2 requires that \mathbb{P} admits a finite first-order moment. In such a general set-up it is challenging, if not impossible, to compute the expected value integral given in (1), i.e., one cannot evaluate the expected-valued core $C_E(G_{\mathbb{P}})$. To circumvent this challenge, we propose a methodology based on distributionally robust optimization.

B. Distributionally robust stability of allocations

In our setting, agent coalitions $S \subset \mathcal{N}$ can construct ambiguity sets of the probability distribution \mathbb{P} of the uncertainty $\xi \in \Xi$ that affect their coalitional values $u_S(\xi)$. This is due to lack of knowledge of \mathbb{P} . In other words, we do not only have uncertainty affecting the coalitional game, but uncertainty about the distribution of the uncertain parameter.

We postulate that each coalition $S \subset \mathcal{N}$ is allowed to construct their own ambiguity sets. The heterogeneity of the coalitional ambiguity sets provides the necessary modelling freedom for our theory to be flexible for application purposes. To this end, we assume that each coalition $S \subset \mathcal{N}$ has access to their own i.i.d. samples $\xi_{K_S} = (\xi^{(1)}, \dots, \xi^{(K_S)}) \in \Xi^{K_S}$ and consider the distributionally robust version $G_{\hat{\mathcal{P}}_K}$ of the original game $G_{\mathbb{P}}$ defined as the tuple $G_{\hat{\mathcal{P}}_K} = \{\mathcal{N}, \{u_S\}_{S \subseteq \mathcal{N}}, \Xi, \hat{\mathcal{P}}_K\}$, where $K = \{K_S\}_{S \subseteq \mathcal{N}}$, while $\hat{\mathcal{P}}_K = \{\hat{\mathcal{P}}_{K_S}\}_{S \subseteq \mathcal{N}}$ is the collection of ambiguity sets constructed based on the available data ξ_{K_S} of each subcoalition $S \subset \mathcal{N}$. We now proceed to defining the notion of distributional stability of an allocation.

Definition 4: (Distributionally robust stability of allocations) For a given number of i.i.d. drawn samples $\xi_{K_S} = (\xi^{(1)}, \dots, \xi^{(K_S)}) \in \Xi^{K_S}$ per coalition $S \subset \mathcal{N}$, an allocation $x = (x_i)_{i \in \mathcal{N}}$ is distributionally stable with respect to the coalitional ambiguity sets $\hat{\mathcal{P}}_{K_S}, S \subset \mathcal{N}$ if

- 1) $\sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}}$ and
- 2) $\sum_{i \in S} x_i \geq \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_Q[u_S(\xi)], \forall S \subset \mathcal{N}$. \square

III. DISTRIBUTIONALLY ROBUST COALITIONAL GAMES BASED ON THE WASSERSTEIN DISTANCE

A. Background on distributional robustness

In this section, we introduce some basic concepts from distributionally robust optimization under the Wasserstein metric [17], [23], [24]. We show how one can leverage this framework in order to provide certificates of stability with respect to the true unknown coalitional game along with a tractable approximation of its expected value core. We start by imposing a mild assumption on the probability distributions serving as candidates of the true distribution \mathbb{P} . Specifically, we consider all distributions with bounded first-order moments, i.e., $\mathbb{Q} \in \mathcal{M}(\Xi)$, where $\mathcal{M}(\Xi)$ is the set of probability distributions with support Ξ that satisfy Assumption 2. We then need a measure of distance between two probability distributions to quantify how close a candidate probability distribution is to the true probability distribution \mathbb{P} ; Let us thus use the Wasserstein distance defined as follows.

Definition 5: (Wasserstein distance, [25]) The Wasserstein distance $d_W : \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \rightarrow \mathbb{R}_{\geq 0}$ between two probability distributions $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\Xi)$ is defined as

$$d_W(\mathbb{Q}_1, \mathbb{Q}_2) = \inf_{\Pi} \left\{ \int_{\Xi^2} \|\xi_1 - \xi_2\| \Pi(d\xi_1, d\xi_2) : \right.$$

Π is a joint distribution of ξ_1 and ξ_2

with marginals \mathbb{Q}_1 and \mathbb{Q}_2 , respectively $\left. \right\}$,

where $\|\cdot\|$ can be any norm on \mathbb{R}^p .

An alternative dual interpretation can be derived by the so-called Kantorovich-Rubinstein theorem:

Theorem 1: (Kantorovich-Rubinstein theorem, [25]) Given two probability distributions $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\Xi)$ we have that

$$d_W(\mathbb{Q}_1, \mathbb{Q}_2) = \sup_{f \in \mathcal{F}} \left\{ \int_{\Xi} f(\xi) \mathbb{Q}_1(d\xi) - \int_{\Xi} f(\xi) \mathbb{Q}_2(d\xi) \right\},$$

where \mathcal{F} is the space of all Lipschitz continuous functions for which $|f(\xi) - f(\xi')| \leq \|\xi - \xi'\|$ for all $\xi, \xi' \in \Xi$. \square

B. Finite-sample guarantees for the distributionally robust core

In the subsequent developments, we consider that each coalition $S \subset \mathcal{N}$ has their own independent samples from \mathbb{P} . Any given coalition $S \subset \mathcal{N}$ constructs their respective ambiguity set based on their collected data $\xi_{K_S} = (\xi^{(k_S)})_{k_S=1}^{K_S} \in \Xi^{K_S}$. For each coalition $S \subset \mathcal{N}$ each ambiguity set is given by:

$$\mathbb{B}_{\varepsilon_S}(\hat{\mathbb{P}}_{K_S}) = \{\mathbb{Q}_S \in \mathcal{M} : d_W(\hat{\mathbb{P}}_{K_S}, \mathbb{Q}_S) \leq \varepsilon_S\}, \quad (2)$$

where $\hat{\mathbb{P}}_{K_S} = \sum_{k_S=1}^{K_S} \delta(\xi - \xi^{(k_S)})$ is the empirical probability distribution of each coalition S on the basis of K_S i.i.d samples from the support set Ξ by coalition S with $\delta(\xi - \xi^{(k_S)}) = \frac{1}{K_S}$ if $\xi = \xi^{(k_S)}$, $k_S = \{1, \dots, K_S\}$ and 0 otherwise.

The following result relies on Assumption 1 to provide guarantees on the probability that a multi-sample will be drawn from coalition $S \subset \mathcal{N}$ such that true probability measure lies within the constructed Wasserstein ball with a given confidence.

Lemma 1: Let Assumption 1 hold and for any coalition $S \subset \mathcal{N}$ fix $\varepsilon_S > 0$. We have that

$$\mathbb{P}^{K_S} \left\{ \xi_{K_S} \in \Xi^{K_S} : d_W(\mathbb{P}, \hat{\mathbb{P}}_{K_S}) \leq \varepsilon_S \right\} \geq 1 - \beta_S,$$

where

$$\beta_S = \begin{cases} c \exp(-q K_S \varepsilon_S^{\max\{p, 2\}}), & \text{if } \varepsilon_S \leq 1 \\ c \exp(-q K_S \varepsilon_S^a), & \text{if } \varepsilon_S > 1, \end{cases} \quad (3)$$

for all $K_S \geq 1$ and $p \neq 2$, where c, q are positive constants that only depend on the parameters a, A in Assumption 1 and the dimension of the support set p .

Proof: The proof is an adaptation of Theorem 2 in [17] applied to each Wasserstein ball $\hat{\mathcal{P}}_{K_S} = \mathbb{B}_{\varepsilon_S}(\hat{\mathbb{P}}_{K_S})$ constructed around the empirical probability distribution $\hat{\mathbb{P}}_{K_S}$ and a radius ε_S chosen by each coalition $S \subset \mathcal{N}$. \blacksquare

Lemma 1 paves the way towards establishing finite sample guarantees for the following Wasserstein-based version of the distributionally robust core.

Definition 6: (Distributionally robust core) The distributionally robust core $C_{DR}(G_{\hat{\mathcal{P}}_K})$ of the game $G_{\mathbb{P}}$ based on the Wasserstein distance is defined as the set

$$C_{DR}(G_{\hat{\mathcal{P}}_K}) = \left\{ x \in \mathbb{R}^N : \sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}}, \right. \\ \left. \sum_{i \in S} x_i \geq \sup_{\mathbb{Q}_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{\mathbb{Q}_S}[u_S(\xi_S)], \forall S \subset \mathcal{N} \right\},$$

where $\hat{\mathcal{P}}_{K_S} = \mathbb{B}_{\varepsilon_S}(\hat{\mathbb{P}}_{K_S})$. \square

Throughout we assume that for all multi-samples, the ambiguity sets are such that a non-empty DR core is returned.

Theorem 2: For each $S \subset \mathcal{N}$ fix a Wasserstein radius ε_S and consider a multi-sample size K_S . It holds that

$$\mathbb{P}^K \left\{ \xi_K \in \Xi^K : C_E(G_{\mathbb{P}}) \supseteq C_{DR}(G_{\hat{\mathcal{P}}_K}) \right\} \geq \beta, \quad (4)$$

where $\beta = \prod_{S \subset \mathcal{N}} (1 - \beta_S)$ and each β_S is given by (3).

Proof: We have that

$$\begin{aligned} & \mathbb{P}^K \left\{ \xi_K \in \Xi^K : C_E(G_{\{\Xi, \mathbb{P}\}}) \supseteq C_{DR}(G_{\{\Xi, \mathbb{P}\}}) \right\} \\ & \geq \mathbb{P}^K \left\{ \xi_K \in \Xi^K : \mathbb{E}_{\mathbb{P}}[u_S(\xi)] \leq \sup_{\mathbb{Q}_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{\mathbb{Q}_S}[u_S(\xi)], \forall S \subset \mathcal{N} \right\} \\ & = \mathbb{P}^K \left\{ \bigcap_{S \subset \mathcal{N}} \left\{ \xi_K \in \Xi^K : \mathbb{E}_{\mathbb{P}}[u_S(\xi)] \leq \sup_{\mathbb{Q}_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{\mathbb{Q}_S}[u_S(\xi)] \right\} \right\} \\ & = \prod_{S \subset \mathcal{N}} \mathbb{P}^{K_S} \left\{ \xi_{K_S} \in \Xi^{K_S} : \mathbb{E}_{\mathbb{P}}[u_S(\xi)] \leq \sup_{\mathbb{Q}_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{\mathbb{Q}_S}[u_S(\xi)] \right\} \end{aligned} \quad (5)$$

The last equality is due to the fact that each coalition constructs its ambiguity sets based on its own (independent) samples. From Lemma 1 for each coalition $S \subset \mathcal{N}$ we have that

$$\mathbb{P}^{K_S} \left\{ \xi_{K_S} \in \Xi^{K_S} : d_W(\mathbb{P}, \hat{\mathbb{P}}_{K_S}) \leq \varepsilon_S \right\} \geq 1 - \beta_S.$$

Therefore,

$$\begin{aligned} & \prod_{S \subset \mathcal{N}} \mathbb{P}^{K_S} \left\{ \xi_{K_S} \in \Xi^{K_S} : \mathbb{E}_{\mathbb{P}}[u_S(\xi)] \leq \sup_{\mathbb{Q}_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{\mathbb{Q}_S}[u_S(\xi_S)] \right\} \\ & \geq \prod_{S \subset \mathcal{N}} \mathbb{P}^{K_S} \left\{ \xi_{K_S} \in \Xi^{K_S} : d_W(\mathbb{P}, \hat{\mathbb{P}}_{K_S}) \leq \varepsilon_S \right\} \\ & \geq \prod_{S \subset \mathcal{N}} (1 - \beta_S), \end{aligned} \quad (6)$$

where the first inequality follows from Theorem 3.5 in [23]. \blacksquare

The following result provides guarantees when all coalitions use the same parameters.

Corollary 1: Consider Assumption 1. For each coalition $S \subset \mathcal{N}$ fix a common Wasserstein radius ε and assume that the same multi-sample ξ_K is used among coalitions. Then, it holds that

$$\mathbb{P}^K \{ \xi_K \in \Xi^K : C_E(G_{\mathbb{P}}) \supseteq C_{DR}(G_{\hat{\mathbb{P}}_K}) \} \geq (1 - \beta)^M,$$

where β is given by (3) by setting $\varepsilon_S = \varepsilon$ and $K_S = K$. \square

Proof: Since $\varepsilon_S = \varepsilon$ for all $S \subset \mathcal{N}$ and the same number of samples $K \in \mathbb{N}$ is drawn for all $S \subset \mathcal{N}$, then by (3), $\beta_S = \beta$ for all $S \subset \mathcal{N}$. By Lemma 1 we then have that for each coalition $S \subset \mathcal{N}$

$$\mathbb{P}^K \{ \xi_K \in \Xi^K : d_W(\mathbb{P}, \hat{\mathbb{P}}_K) \leq \varepsilon \} \geq 1 - \beta,$$

where β is given by (3). Following the same proofline as in Theorem 2 and setting $\beta_S = \beta$ to the right hand side of (6) concludes the proof. \blacksquare

Remark 1: The confidence parameter in Theorem 2 and Corollary 1 depends on the number of possible subcoalitions $S \subset \mathcal{N}$. Corollary 1 shows that the dependence on the number of possible coalitions of the original problem is also inherited by the provided finite-sample guarantees. Though under a different approach, this is also observed in the *a priori* results in [15], where the authors apply the results of [26], [27] to construct a data-driven version of the robust core. \square

C. Agent-level sampling

Assume now that the ambiguity sets are constructed on the basis of samples drawn by the individual agents, i.e., with a slight abuse of notation each agent draws a multi-sample ξ_{K_i} , which is then used by any coalition $S \subset \mathcal{N}$ for which $i \in S$. Fix a confidence parameter $\bar{\beta}_S \in (0, 1)$ for each $S \subset \mathcal{N}$ and let $\xi_{\bar{K}} = (\xi_{\bar{K}_S})_{S \subset \mathcal{N}}$, where $\xi_{\bar{K}_S}$ are the samples of each coalition $S \subset \mathcal{N}$. Then, it is easy to show that

$$\mathbb{P}^{\bar{K}} \{ \xi_{\bar{K}} \in \Xi^{\bar{K}} : C_E(G_{\mathbb{P}}) \supseteq C_{DR}(G_{\hat{\mathbb{P}}_{\bar{K}}}) \} \geq \bar{\beta},$$

where $\bar{\beta} = \sum_{S \subset \mathcal{N}} (1 - \bar{\beta}_S) - M + 1$ and each $\bar{\beta}_S$ is given by (3) setting $\varepsilon_S = \bar{\varepsilon}_S$ and $K_S = \bar{K}_S = \sum_{i \in S} K_i$. This result is obtained by applying Bonferroni's inequality to the third step of the proof of Theorem 2 and thus constitutes a rather conservative bound. Further work is required to leverage the data on the agents' level and translate it to guarantees on a coalitional level taking into account sharing of data among coalitions and thus improving those theoretical guarantees. This is a challenging task and thus left for future work.

IV. ASYMPTOTIC CONSISTENCY AND COMPUTATIONAL TRACTABILITY OF THE DR CORE

A. Asymptotic consistency of the DR core

In this subsection we show that under Lipschitz continuity of the value functions and careful choice of the radius of the Wasserstein ball ε_S and of the confidence parameter β_S for each coalition $S \subset \mathcal{N}$, the DR core based on the Wasserstein distance converges almost surely to the true expected value core of the original problem. Let us impose the following assumption:

Assumption 3: For each coalition $S \subset \mathcal{N}$ the value function u_S is L_S -Lipschitz continuous in ξ with $L_S \geq 0$, i.e., $\|u_S(\xi) - u_S(\xi')\| \leq L_S \|\xi - \xi'\|$ for all $\xi, \xi' \in \Xi$. \square

Such assumptions are common in stability analysis for stochastic programming [24], [28] and also provide interesting insights in our setting. In fact, the increase in the error between a DR coalitional value and the corresponding expected coalitional value is proportional to the estimation error between the true and the empirical distribution of that coalition amplified by at most the Lipschitz constant of the corresponding value function.

As opposed to the developments of the previous section, where we fix the radius ε_S and the number of samples K_S for any $S \subset \mathcal{N}$ and calculate β_S based on (3), we now solve (3) with respect to ε_S , thus obtaining the Wasserstein radius as a function of the confidence parameter β_S and the number of samples K_S . In particular, we have that

$$\varepsilon_S(\beta_S, K_S) = \begin{cases} \left(\frac{\ln(\frac{c}{\beta_S})}{qK_S} \right)^{\frac{1}{\max\{p, 2\}}} & \text{if } K_S \geq \frac{\ln(\frac{c}{\beta_S})}{q} \\ \left(\frac{\ln(\frac{c}{\beta_S})}{qK_S} \right)^{\frac{1}{p}} & \text{if } K_S < \frac{\ln(\frac{c}{\beta_S})}{q}. \end{cases}$$

The following theorem establishes almost-sure convergence of the DR core to the expected value-core as the number of samples increases.

Theorem 3: Let Assumptions 1 and 3 hold. Suppose that for each $S \subset \mathcal{N}$, $\beta_S^{K_S} \in (0, 1)$, $K_S \in \mathbb{N}$ satisfies $\sum_{K_S=1}^{\infty} \beta_S^{K_S} < \infty$ and $\lim_{K_S \rightarrow \infty} \varepsilon_S(\beta_S, K_S) = 0$. Any sequence of the Wasserstein-based DR cores $\{C_{DR}(G_{\hat{\mathbb{P}}_K})\}_{K \in \mathbb{N}^M}$, where $K = (K_S)_{S \subset \mathcal{N}}$, converges \mathbb{P}^∞ -almost surely to the true expected core $C_E(G_{\mathbb{P}})$ as $K_S \rightarrow \infty$ for all $S \subset \mathcal{N}$. \square

Proof: For each coalition $S \subset \mathcal{N}$ we have that

$$\mathbb{P}^{K_S} \{ \xi_{K_S} \in \Xi^{K_S} : \mathbb{E}_{\mathbb{P}}[u_S(\xi)] \leq \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)] \}$$

$$\geq \mathbb{P}^{K_S} \{ \xi_{K_S} \in \Xi^{K_S} : d_W(\mathbb{P}, \hat{\mathbb{P}}_{K_S}) \leq \varepsilon_S(\beta_S^{K_S}, K_S) \} \geq 1 - \beta_S^{K_S},$$

where the last inequality is due to Lemma 1. Letting $K_S \rightarrow \infty$ and since $\lim_{K_S \rightarrow \infty} \beta_S^{K_S} = 0$ we have that

$$\mathbb{E}_{\mathbb{P}}[u_S(\xi)] \leq \lim_{K_S \rightarrow \infty} \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)] \quad (7)$$

\mathbb{P}^∞ -almost surely. Following a methodology similar in spirit to [23], for each coalition $S \subset \mathcal{N}$ and for every $K_S \in \mathbb{N}$ by the definition of supremum for any $\delta_S > 0$ there exists \hat{Q}_{K_S} such that

$$\sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)] \leq \mathbb{E}_{\hat{Q}_{K_S}}[u_S(\xi)] + \delta_S$$

By the Kantorovich-Rubinstein theorem (Theorem 1), we have that

$$L_S d_W(\hat{Q}_{K_S}, \mathbb{P}) \geq L_S \left(\mathbb{E}_{\hat{Q}_{K_S}} \left[\frac{1}{L_S} u_S(\xi) \right] - \mathbb{E}_{\mathbb{P}} \left[\frac{1}{L_S} u_S(\xi) \right] \right),$$

since, under Assumption 3, $\frac{u_S(\xi)}{L_S}$ is a Lipschitz continuous function with Lipschitz constant less than or equal to 1. The relation above can then be written as

$$\mathbb{E}_{\hat{Q}_{K_S}}[u_S(\xi)] \leq \mathbb{E}_{\mathbb{P}}[u_S(\xi)] + L_S d_W(\hat{Q}_{K_S}, \mathbb{P}).$$

We then have that:

$$\begin{aligned} \lim_{K_S \rightarrow \infty} \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)] &\leq \lim_{K_S \rightarrow \infty} \mathbb{E}_{\hat{Q}_{K_S}}[u_S(\xi)] + \delta_S \\ &\leq \lim_{K_S \rightarrow \infty} \{ \mathbb{E}_{\mathbb{P}}[u_S(\xi)] + L_S d_W(\hat{Q}_{K_S}, \mathbb{P}) \} + \delta_S = \mathbb{E}_{\mathbb{P}}[u_S(\xi)] + \delta_S \end{aligned}$$

\mathbb{P}^∞ -almost surely, since by adapting [23, Lemma 3.7] in our setting, we have that for each $S \subset \mathcal{N}$

$$\lim_{K_S \rightarrow \infty} d_W(\mathbb{P}, \hat{Q}_{K_S}) = 0, \text{ } \mathbb{P}^\infty\text{-almost surely.}$$

Letting $\delta_S \downarrow 0$ we have that

$$\lim_{K_S \rightarrow \infty} \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)] \leq \mathbb{E}_{\mathbb{P}}[u_S(\xi)]. \quad (8)$$

From relations (7) and (8) we have that for any $S \subset \mathcal{N}$

$$\lim_{K_S \rightarrow \infty} \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)] = \mathbb{E}_{\mathbb{P}}[u_S(\xi)], \text{ } \mathbb{P}^\infty\text{-almost surely.}$$

As such, by Definitions 3 and 6, any sequence of DR cores $\{C_{\text{DR}}(G_{\hat{\mathcal{P}}_K})\}_{K \in \mathbb{N}^M}$, for which $S \subset \mathcal{N}$, $\beta_S^{K_S} \in (0, 1)$, $K_S \in \mathbb{N}$ satisfies $\sum_{K_S=1}^{\infty} \beta_S^{K_S} < \infty$ and $\lim_{K_S \rightarrow \infty} \varepsilon_S(\beta_S) = 0$ for all $S \subset \mathcal{N}$ with $K = (K_S)_{S \subset \mathcal{N}}$, converges \mathbb{P}^∞ -almost surely to the true expected core $C_E(G_{\mathbb{P}})$ as $K_S \rightarrow \infty$ for all $S \subset \mathcal{N}$. ■

B. Finding allocations inside the DR core

The results of this subsection hold irrespective of what type of sampling we perform. Leveraging results from [23], we show that an allocation inside the DR core can be computed by solving a finite-dimensional convex optimization problem. Here we impose the following assumption.

- Assumption 4:* 1) For any $S \subset \mathcal{N}$ the value function $u_S(\xi)$ can be written as $u_S(\xi) = \max_{m_S=1, \dots, M_S} u_{m_S}(\xi)$, where $-u_{m_S}(\xi)$ is proper, convex and lower semi-continuous for all $m_S \in \{1, \dots, M_S\}$ and any $S \subset \mathcal{N}$.
2) For any $S \subset \mathcal{N}$ u_S does not take the value $-\infty$ on Ξ .
3) The support set Ξ is closed and convex. □

Under these assumptions we have the following result:

Lemma 2: Let Assumption 4 hold. By drawing K_S samples and considering the dual variables $\lambda_S, \ell_{k_S}, z_{k_S m_S}, v_{k_S m_S}$ that correspond to the Wasserstein ball constraint of each coalition $S \subset \mathcal{N}$, an allocation inside the DR core is found by solving the optimization problem

$$P : \begin{cases} \min_{x, \{\lambda_S, \ell_{k_S}, z_{k_S m_S}, v_{k_S m_S}\}_{S \subset \mathcal{N}}} \|x\|_2^2 \\ \text{s.t. } \sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}} \\ \lambda_S \varepsilon_S + \frac{1}{K_S} \sum_{k_S=1}^{K_S} \ell_{k_S} \leq \sum_{i \in S} x_i, \forall S \subset \mathcal{N} \\ [-u_{m_S}]^*(z_{k_S m_S} - v_{k_S m_S}) + \sigma_{\Xi}(v_{k_S m_S}) - z_{k_S}^\top \xi^{(k_S)} \leq \ell_{k_S}, \\ \quad \forall k_S, \forall m_S, \forall S \subset \mathcal{N} \\ \|z_{k_S m_S}\|_* \leq \lambda_S \forall k_S, \forall m_S, \forall S \subset \mathcal{N} \end{cases}$$

where $[f]^*$ denotes the conjugate function of a function f , i.e., $[f]^*(y) = \sup_{x \in \text{dom}(f)} (y^\top x - f(x))$. and $\|\cdot\|_*$ is the dual norm, while σ_X denotes the conjugate of the characteristic function. □

Proof: We wish to solve the following feasibility problem

$$\begin{cases} \min_{x \in \mathbb{R}^N} \|x\|_2^2 \\ \text{s.t. } x \in C_{\text{DR}}(G_{\hat{\mathcal{P}}_K}), \end{cases}$$

which, by Definition 6 and considering the data-driven Wasserstein ball as the ambiguity set of each coalition, is equivalent to

$$\begin{cases} \min_{x \in \mathbb{R}^N} \|x\|_2^2 \\ \text{s.t. } \sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}}, \\ \sum_{i \in S} x_i \geq \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)], \forall S \subset \mathcal{N}. \end{cases}$$

Note that calculating the DR core based on the drawn K_S samples from each coalition S boils down to solving the worst-case expected value problem:

$$P'_S : \sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)], \forall S \subset \mathcal{N}.$$

This is an infinite-dimensional optimization problem over probability measures. Under Assumption 4 each of these programs parameterized by S can be rewritten as a finite-dimensional convex program under Theorem 4.2 in [23], thus taking the dual form:

$$\begin{cases} \min_x \|x\|_2^2 \\ \text{s.t. } \sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}} \text{ and } \forall S \subset \mathcal{N}, \\ \sum_{i \in S} x_i \geq \begin{cases} \min_{\lambda_S, \ell_{k_S}, z_{k_S m_S}, v_{k_S m_S}} \lambda_S \varepsilon_S + \frac{1}{K_S} \sum_{k_S=1}^{K_S} \ell_{k_S} \\ \text{s.t. } [-u_{m_S}]^*(z_{k_S m_S} - v_{k_S m_S}) + \\ \sigma_{\Xi}(v_{k_S m_S}) - z_{k_S}^\top \xi^{(k_S)} \leq \ell_{k_S} \forall k_S, \forall m_S \\ \|z_{k_S m_S}\|_* \leq \lambda_S \forall k_S, \forall m_S. \end{cases} \end{cases}$$

The problem above can then be rewritten as

$$\begin{cases} \min_x \|x\|_2^2 \\ \text{s.t. } \sum_{i \in \mathcal{N}} x_i = u_{\mathcal{N}} \text{ and } \forall S \subset \mathcal{N}, \\ \begin{cases} \exists \lambda_S, \ell_{k_S}, z_{k_S m_S}, v_{k_S m_S} : \lambda_S \varepsilon_S + \frac{1}{K_S} \sum_{k_S=1}^{K_S} \ell_{k_S} \leq \sum_{i \in S} x_i \\ \text{s.t. } [-u_{m_S}]^*(z_{k_S m_S} - v_{k_S m_S}) + \\ \sigma_{\Xi}(v_{k_S m_S}) - z_{k_S}^\top \xi^{(k_S)} \leq \ell_{k_S} \forall k_S, \forall m_S \\ \|z_{k_S m_S}\|_* \leq \lambda_S \forall k_S, \forall m_S. \end{cases} \end{cases}$$

Due to the existential operator we can minimize with respect to $\lambda_S, \ell_{k_S}, z_{k_S m_S}, v_{k_S m_S}$ for all $S \subset \mathcal{N}$. This concludes then the proof. ■

Note that the additional decision variables of P correspond to Wasserstein distance constraints in the primal problems P'_S for each $S \subset \mathcal{N}$ [23, Theorem 4.2].

V. NUMERICAL EXAMPLE

We consider a stochastic coalitional game with $N = 3$ agents and coalitional values $u_1 = 2 + \xi$, $u_2 = 1.5 + \xi$, $u_3 = 2.5 + \xi$ and $u_{12} = 6 + \xi, u_{23} = 6.5 + \xi, u_{13} = 7 + \xi$, where ξ is an uncertain parameter following a Gaussian distribution with mean value $\mu = 1$ and variance $\sigma^2 = 1$, truncated at the interval $\Xi = [0, 1]$. We assume that each coalition has a certain level of ambiguity with respect to the distribution of the uncertain parameter. Though our results hold under the consideration of different Wasserstein balls




Fig. 1. DR coalitional values for the expectation of $u_S(\xi)$ evaluated over 500 simulations (red shaded area) vs corresponding expected coalitional value (blue solid line) for $K_S \in \{5, 10, 30, 50, 100, 200, 500\}$ and $\varepsilon = 0.3$. For $K_S \geq 10$ the DR core is contained with empirical confidence 1 in the expected valued core.




Fig. 2. DR coalitional values for the expectation of $u_S(\xi)$ evaluated over 500 simulations (red shaded area) vs corresponding expected coalitional value (blue solid line) for $K_S \in \{5, 10, 30, 50, 100, 200, 500\}$ and $\varepsilon = 0.03$. It is observed that decrease in ε affects the confidence with which the DR core is contained within the expected value core. At $K_S = 500$, however, only a small portion of DR values is below the expected value.

and multi-sample sizes per coalition $S \subset \mathcal{N}$, we assume here the same radius and the same multi-sample size for illustration purposes (see Corollary 1). Initially, we consider for each coalition a Wasserstein ball of radius $\varepsilon = 0.3$.

Figure 1 focuses on coalition $S = \{1\}$ and shows the range of normalized DR values for the expectation of $u_S(\xi)$, denoted by $\sup_{Q_S \in \hat{\mathcal{P}}_{K_S}} \mathbb{E}_{Q_S}[u_S(\xi)]$ over 500 simulations per multi-sample size K_S . We note that a similar behaviour is exhibited in all other coalitions, with the graphs being centered at different values. As the observed pattern is similar across coalitions, it is not shown to avoid repetition.

For the same number of samples per coalition varying in $[5, 500]$, we observe that as the number of samples increases, all the DR coalitional values, illustrated by the red shaded area in Figure 1, are above the corresponding expected value (blue solid line). Since the same pattern is observed across all coalitions, this implies (see Equation (6) in the proof of Theorem 2) that the DR core is contained within the expected value core as K_S increases for each $S \subset \mathcal{N}$. This observation is in line with Theorem 1, since for an




Fig. 3. Coalitional values (red shaded area), among 500 simulations with increasing Wasserstein radius ε for $K_S = 100$. As the radius increases, the DR coalitional values of the expectation of $u_S(\xi)$ are above their corresponding expected value (blue solid line) and thus, since the same pattern is repeated across coalitions, the DR core lies within the expected valued core.




Fig. 4. Coalitional values (red shaded area), among 500 simulations with increasing Wasserstein radius ε for $K_S = 250$. When samples are drawn from distributions that admit a density with certain concentration properties, we conjecture that due to the more accurate empirical probability across simulations, the DR coalitional values variability is smaller compared to Figure 3.

increasing multi-sample size per coalition, the confidence β in the provided theoretical guarantees tends to 1. Figure 2 shows that following the same approach as Figure 1 for a significantly smaller Wasserstein radius $\varepsilon = 0.03$ leads to a smaller empirical confidence that is improved the more samples we obtain. At $K_S = 500$, however, only a small portion of DR values is below the expected value, which implies stability of DR stable allocations in the mean sense with high confidence.

Figure 3 illustrates the DR coalitional values for the empirical expectation of $u_S(\xi)$ (red shaded areas), over 500 simulations, compared to the corresponding expected value (blue solid line) as the Wasserstein radius ε increases (for a fixed number of samples $K_S = 100$). We note again that for a radius ε larger than a certain threshold, all DR coalitional values are above their corresponding expected value and therefore, since this holds for all coalitions when the same number of samples is used, the DR core lies within the expected value core. This behaviour is expected because for a fixed number of samples the larger the Wasserstein ball the

more likely it is to include the true probability distribution. As such, obtaining allocations stable in the mean sense can also be achieved, even for a small number of samples, by tuning the Wasserstein radius of each coalition accordingly. We note that the comparison of Figures 1 and 2 is consistent with Figure 3. The percentage of the width of the shaded area acts as an empirical estimate of the confidence. The higher ε or K_S are, the lower β_S and as a result the higher the confidence.

Compared to Figure 3, in Figure 4 for each simulation a larger number $K_S = 250$ of samples is generated. Drawing conclusions for the general case is not straightforward, however, when samples are drawn from distributions that admit a density and with certain concentration properties we conjecture that the more samples are used, the more likely it is that the resulting empirical distributions \hat{P}_{K_S} are closer to each other across simulations (i.e., for different multi-samples). As such the centres of the Wasserstein balls would be closer, which in turn implies that the DR value for the expectation of $u_S(\xi)$ would be closer to each other as well. In other words, the higher the number of samples the smaller the variability of the resulting DR value for the expectation of $u_S(\xi)$ across simulations. In line with this intuition, the width of the shaded area in Figure 4 (an empirical estimate of variability) is smaller compared to that of Figure 3.

VI. CONCLUSION

We have introduced the concept of distributionally robust core for coalitional games subject to distributional uncertainty, namely a set of payoff allocations that is robust in the expected value sense. We showed both theoretically and numerically that the concept of distributionally robust stability implies stability in the mean sense as more data becomes available given a certain radius. Furthermore, one can obtain payoff allocations in the expected core by tuning the Wasserstein radius even for a small sample size. This paper takes a first step towards studying the class of distributionally robust chance-constrained coalitional games. Future work will focus on improving the probabilistic guarantees and on designing distributed payoff allocation algorithms.

REFERENCES

- [1] G. Chalkiadakis, E. Elkind, and M. Wooldridge, "Computational aspects of cooperative game theory," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 5, no. 6, pp. 1–168, 2011.
- [2] A. A. Raja and S. Grammatico, "Payoff distribution in robust coalitional games on time-varying networks," *IEEE Transactions on Control of Network Systems*, pp. 1–1, 2021.
- [3] L. Han, T. Morstyn, and M. McCulloch, "Incentivizing prosumer coalitions with energy management using cooperative game theory," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 303–313, 2019.
- [4] F. Fele, J. M. Maestre, and E. F. Camacho, "Coalitional control: Cooperative game theory and control," *IEEE Control Systems Magazine*, vol. 37, no. 1, pp. 53–69, 2017.
- [5] R. A. McCain, "Cooperative games and cooperative organizations," *The Journal of Socio-Economics*, vol. 37, no. 6, pp. 2155–2167, 2008, special Issue: Cooperatives and the Economy.
- [6] A. Charnes and D. Granot, "Coalitional and chance-constrained solutions to n-person games I: The prior satisficing nucleolus," *SIAM Journal on Applied Mathematics*, vol. 31, no. 2, pp. 358–367, 1976.
- [7] —, "Coalitional and chance-constrained solutions to n-person games II: Two-stage solutions," *Operations Research*, vol. 25, no. 6, pp. 1013–1019, 1977.
- [8] —, "Prior solutions: Extensions of convex nucleus solutions to chance-constrained games," 1973.
- [9] J. Suijs, P. Borm, A. De Waegenaere, and S. Tijs, "Cooperative games with stochastic payoffs," *Journal of Games and Economic Behaviour*, vol. 1995-88, 1995.
- [10] G. Chalkiadakis and C. Boutilier, "Bayesian reinforcement learning for coalition formation under uncertainty," *AAMAS 2004*, pp. 1090–1097, 2004.
- [11] S. Jeong and Y. Shoham, "Bayesian coalitional games," *AAAI*, pp. 95–100, 01 2008.
- [12] Y. Li and V. Conitzer, "Cooperative game solution concepts that maximize stability under noise," *AAAI*, 2015.
- [13] G. Chalkiadakis and C. Boutilier, "Sequential decision making in repeated coalition formation under uncertainty," in *AAMAS*, 2008.
- [14] J. Guéron and G. 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*, pp. 89–100, 2021.
- [15] G. Pantazis, F. Fabiani, F. Fele, and K. Margellos, "Probabilistically robust stabilizing allocations in uncertain coalitional games," *IEEE Control Systems Letters*, vol. 6, pp. 3128–3133, 2022.
- [16] A. Shapiro and A. Kleywegt, "Minimax analysis of stochastic problems," *Optimization Methods and Software*, vol. 17, no. 3, pp. 523–542, Jan. 2002.
- [17] N. Fournier and A. Guillin, "On the rate of convergence in Wasserstein distance of the empirical measure," *Probability Theory and Related Fields*, vol. 162, no. 3-4, pp. 707–738, Aug. 2015.
- [18] G. C. Calafiore and L. E. Ghaoui, "On Distributionally Robust Chance-Constrained Linear Programs," *Journal of Optimization Theory and Applications*, vol. 130, no. 1, pp. 1–22, Dec. 2006.
- [19] J. Goh and M. Sim, "Distributionally Robust Optimization and Its Tractable Approximations," *Operations Research*, vol. 58, no. 4-part-1, pp. 902–917, Aug. 2010.
- [20] E. Delage and Y. Ye, "Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems," *Operations Research*, vol. 58, no. 3, pp. 595–612, Jun. 2010.
- [21] A. Cherukuri and J. Cortes, "Cooperative Data-Driven Distributionally Robust Optimization," *IEEE Transactions on Automatic Control*, vol. 65, no. 10, pp. 4400–4407, Aug. 2020.
- [22] R. Gao and A. Kleywegt, "Distributionally Robust Stochastic Optimization with Wasserstein Distance," *Mathematics of Operations Research*, Aug. 2022.
- [23] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, Sep. 2018.
- [24] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, "Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning," *Operations Research & Management Science in the Age of Analytics*, pp. 130–166, Oct. 2019.
- [25] L. Kantorovich and G. S. Rubinstein, "On a space of totally additive functions," *Vestnik Leningrad. Univ.*, vol. 13, pp. 52–59, 1958.
- [26] G. Pantazis, F. Fele, and K. Margellos, "A posteriori probabilistic feasibility guarantees for Nash equilibria in uncertain multi-agent games," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 3403–3408, 2020.
- [27] —, "On the probabilistic feasibility of solutions in multi-agent optimization problems under uncertainty," *European Journal of Control*, vol. 63, pp. 186–195, 2022.
- [28] W. Römisch and R. Schultz, "Stability analysis for stochastic programs," *Annals of Operations Research*, vol. 30, no. 1, pp. 241–266, Dec. 1991.