



Learning with Adaptive Conservativeness for Distributionally Robust Optimization: Incentive Design for Voltage Regulation

Preprint

Zhirui Liang,¹ Qi Li,¹ Joshua Comden,² Andrey Bernstein,² and Yury Dvorkin¹

1 Johns Hopkins University

2 National Renewable Energy Laboratory

Presented at the 63rd IEEE Conference on Decision and Control (CDC 2024)

Milan, Italy

December 16 - 19, 2024

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5D00-90853
December 2024



Learning with Adaptive Conservativeness for Distributionally Robust Optimization: Incentive Design for Voltage Regulation

Preprint

Zhirui Liang,¹ Qi Li,¹ Joshua Comden,² Andrey Bernstein,² and Yury Dvorkin¹

1 Johns Hopkins University

2 National Renewable Energy Laboratory

Suggested Citation

Liang, Zhirui, Qi Li, Joshua Comden, Andrey Bernstein, and Yury Dvorkin. 2024. *Learning with Adaptive Conservativeness for Distributionally Robust Optimization: Incentive Design for Voltage Regulation: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-5D00-90853. <https://www.nrel.gov/docs/fy25osti/90853.pdf>.

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5D00-90853
December 2024

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Learning with Adaptive Conservativeness for Distributionally Robust Optimization: Incentive Design for Voltage Regulation

Zhirui Liang, Qi Li, Joshua Comden, Andrey Bernstein, Yury Dvorkin

Abstract—Information asymmetry between the Distribution System Operator (DSO) and Distributed Energy Resource Aggregators (DERAs) obstructs designing effective incentives for voltage regulation. To capture this effect, we employ a Stackelberg game-theoretic framework, where the DSO seeks to overcome the information asymmetry and refine its incentive strategies by learning from DERA behavior over multiple iterations. We introduce a model-based online learning algorithm for the DSO, aimed at inferring the relationship between incentives and DERA responses. Given the uncertain nature of these responses, we also propose a distributionally robust incentive design model to control the probability of voltage regulation failure and then reformulate it into a convex problem. This model allows the DSO to periodically revise distribution assumptions on uncertain parameters in the decision model of the DERA. Finally, we present a gradient-based method that permits the DSO to adaptively modify its conservativeness level, measured by the size of a Wasserstein metric-based ambiguity set, according to historical voltage regulation performance. The effectiveness of our proposed method is demonstrated through numerical experiments.

I. INTRODUCTION

Grid-edge Distributed Energy Resources (DERs), encompassing electricity sources, storage, and demand response programs linked to distribution networks, pose both challenges and opportunities for the Distribution System Operator (DSO) [1]. To tackle issues like forecast uncertainty, monitoring difficulties, and constrained control capabilities, the integration of DER aggregators (DERAs) into a hierarchical control structure has been advocated [2]. This approach allows DSOs to more effectively harness DER flexibility for services like voltage regulation, frequency response, and reserve, though it necessitates well-crafted incentives to boost DER participation. This paper studies the incentive design

Zhirui Liang, Qi Li, and Yury Dvorkin are with the Johns Hopkins University, Baltimore MD, U.S. ({zliang31, qli112, ydvorki1}@jhu.edu). Joshua Comden and Andrey Bernstein are with the National Renewable Energy Laboratory, Golden CO, U.S. ({joshua.comden, andrey.bernstein}@nrel.gov).

This work was supported in part by the Risk Aware Power System Control, Dispatch and Market Incentives through ARPA-E PERFORM under Award DE-AR0001300; in part by NSF under Award 2330450, Award 2245507, and Award 2245506.

This work was authored in part by the National Renewable Energy Laboratory (NREL), operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

problem for the DSO, a topic that has gained significant attention [3]–[5]. Specifically, this paper highlights an often-overlooked aspect: DERAs might withhold some information from the DSO, such as the cost and flexibility region of individual DERs, to leverage information asymmetry for profit gains [6]. This necessitates a strategic approach for the DSO to learn the decision-making patterns of DERAs.

Previous research offers various strategies for learning the decision-making models of DERAs, navigating the challenge that the DSO can only observe the responses of DERAs to incentives without direct access to the internal parameters of the DERA decision model. Some studies opt for data-driven and model-free approaches, like the use of neural networks in [7], while others apply reverse engineering techniques, such as inverse optimization, to deduce the unseen parameters from observed optimal responses [8]. However, these approaches tend to capture a static relationship between incentives and DER responses, overlooking the time-varying nature of parameters within the demand response model.

Several studies have recognized the uncertainty in demand response models and have developed distributionally robust strategies for downstream control tasks. For instance, [9], [10] introduced online learning methods to estimate the price sensitivity of demand response participants, incorporating this insight into a distributionally robust optimal power flow model. Similarly, [11] formulated distributionally robust control policies using Markov Decision Processes for demand response dispatch. However, these studies typically assume a fixed level of decision conservativeness, which should evolve as parameter learning continues to mirror the increasing confidence in the acquired knowledge. In distributionally robust optimization, this conservativeness level can be quantified by the size of the ambiguity set, as outlined in our previous research [12].

In this paper, we present a model-based online learning algorithm for the DSO to infer the relationship between incentives and DERA responses. Leveraging the learned information, we develop a distributionally robust incentive design model aimed at enhancing voltage regulation in electric power distribution systems. Additionally, we introduce a gradient-based strategy enabling the DSO to adaptively adjust its conservativeness level based on historical voltage regulation performance. This distributionally robust decision-making model with adaptive conservativeness level differs from the *adaptive distributionally robust optimization* proposed in [13], which aims to make the best decision for multi-stage problems. It also differs from the decision-aware learning models, e.g., uncertainty set learning as in [14], [15],

which minimize the distance between the expected cost and the best-case cost rather than the worst-case cost as in this paper.

II. OVERVIEW OF INCENTIVE DESIGN PROBLEM

A. Incentive-based Voltage Regulation

The process of incentive-based voltage regulation is characterized by a series of strategic decisions. Initially, the DSO assesses the voltage at all nodes in the distribution system, and then computes incentives for a desired DER response in areas with voltage violations. Then, the DERAs adjust DER settings, leveraging comprehensive control and monitoring over the DER output. Often, this is an iterative process since DER responses will influence voltages and thus the optimal incentives in the future.

This incentive-based voltage regulation problem can be modeled with a Stackelberg game-theoretic framework, which captures the interactions between two players in a leader-follower dynamic. Fig. 1 shows these interactions for a distribution system with N_b nodes, where the DSO, acting as the single leader, initiates by offering incentives, while the DERAs, as multiple followers in the game, strategically respond by re-dispatching the DER operations. In this paper, we assume a single DERA per node for clarity, though real-world scenarios might feature multiple DERAs per node or a single DERA managing DERs across different nodes.

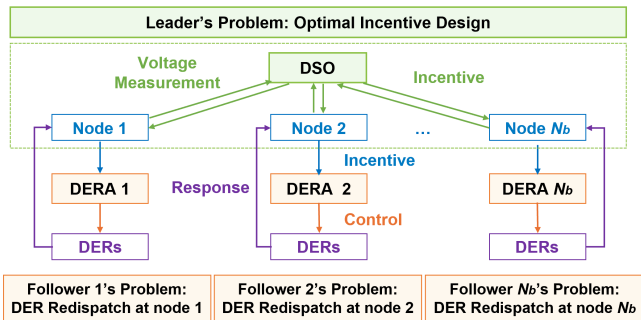


Fig. 1. Incentive-based voltage regulation modeled as a Stackelberg game between the DSO (leader) and DERAs (followers).

B. Leader's Problem: Optimal Incentive Design

The decision-making problem for the DSO is designing incentives to achieve voltage regulation in the most efficient way, which can be formulated as:

$$\min_{k_n^P, k_n^Q, \Delta V_n} f = \sum_{n=1}^{N_b} (k_n^P \Delta P_n + k_n^Q \Delta Q_n) \quad (1a)$$

$$\text{s.t. } n = 1, \dots, N_b$$

$$\Delta V_n = \sum_{m=1}^{N_b} \alpha_{n,m} \Delta P_m + \sum_{m=1}^{N_b} \beta_{n,m} \Delta Q_m \quad (1b)$$

$$\Delta V_n^{\min} \leq \Delta V_n \leq \Delta V_n^{\max} \quad (1c)$$

$$\Delta P_n, \Delta Q_n \in \text{argmin } g(\Delta P_n, \Delta Q_n | k_n^P, k_n^Q), \quad (1d)$$

where the decision variables k_n^P and k_n^Q are the incentives for active and reactive power response at node n . Objective (1a) aims to minimize the cost associated with voltage regulation.

This cost is calculated by summing the products of the incentives and the DERA responses at each node, with ΔP_n and ΔQ_n representing the total active and reactive DER responses at node n , as determined by (1d). We assume that the DERs are capable of providing bidirectional responses, meaning the values of k_n^P , k_n^Q , ΔP_n , and ΔQ_n can be either positive or negative. The auxiliary variable ΔV_n denotes the voltage change at node n , which can be derived from the linearized power flow equation in (1b). In this equation, $\alpha_{n,m}$ and $\beta_{n,m}$ are coefficients linking the voltage change at node n to the DERA response at node m [16]. Furthermore, (1c) specifies a permissible range for voltage changes, which is time-varying based on voltage measurements in the system. The response of each DERA is determined by (1d), where function g will be introduced in the following subsection.

C. Follower's Problem: DER Redispatch

Considering the incentives provided by the DSO, along with the operational flexibility and associated costs of each DER, the DERA can identify which set of DERs to re-dispatch, aiming at maximizing its overall profit. The decision-making model for the DERA at node n is:

$$\min_{\Delta P_r, \Delta Q_r} g = \sum_{r=1}^{R_n} s_r \Delta P_r - k_n^P \Delta P_n - k_n^Q \Delta Q_n \quad (2a)$$

$$\text{s.t. } r = 1, \dots, R_n$$

$$\Delta P_r^{\min} \leq \Delta P_r \leq \Delta P_r^{\max} \quad (2b)$$

$$\phi_r^{\min} \Delta P_r \leq \Delta Q_r \leq \phi_r^{\max} \Delta P_r \quad (2c)$$

$$\Delta P_n = \sum_{r=1}^{R_n} \Delta P_r \quad (2d)$$

$$\Delta Q_n = \sum_{r=1}^{R_n} \Delta Q_r \quad (2e)$$

$$k_n^P, k_n^Q \in \text{argmin } f(k_n^P, k_n^Q | \Delta P_n, \Delta Q_n), \quad (2f)$$

where R_n is the total number of DERs at node n , ΔP_r and ΔQ_r are the active and reactive power response of the r -th DERs at node n , respectively. For simplicity, it is assumed that only ΔP_r incurs a specific unit cost s_r , whereas ΔQ_r is a by-product of ΔP_r and does not entail any direct cost. The linear feasible regions of ΔP_r and ΔQ_r are enforced by (2b) and (2c), respectively, while the aggregated DER responses at node n are calculated by (2d) and (2e).

III. LEARNING FOR INCENTIVE DESIGN

If the DSO has complete knowledge about the decision-making models used by DERAs, including the cost and flexibility of each DER, the optimal incentive design problem could be framed as a bi-level optimization problem, where the upper level model is (1) and the lower level model is (2). However, the limited knowledge available to the DSO complicates the direct implementation of this bi-level optimization strategy and necessitates a learning approach. In this section, we introduce an online learning algorithm, where the DSO can infer the underlying decision-making models of DERAs by continuously observing DERA responses to incentives within the Stackelberg game context.

A. Learning with Surrogate Model

The DSO seeks to identify the mapping from the incentives k_n^P, k_n^Q to the responses $\Delta P_n, \Delta Q_n$ through a learning process and to use this mapping in (1d). The relationship between incentives and responses is essentially determined by the decision-making model of the DERA detailed in (2). However, the DSO cannot directly infer the model in (2) due to two main factors. First, the DSO has access only to the aggregated demand response at each node (ΔP_n and ΔQ_n) without visibility into individual DER responses (ΔP_r and ΔQ_r). Second, several parameters in (2), including the costs s_r and the feasible ranges for ΔP_r and ΔQ_r determined by (2b) and (2c), are time-varying due to the variability of DER resources. Therefore, the DSO has to learn this mapping relationship based on a surrogate model. Without losing generality (as discussed in Section III-B below), we assume the DSO employs the following model to approximate (2):

$$\begin{aligned} \min_{\Delta P_n, \Delta Q_n} \quad & c_n \Delta P_n^2 + d_n \Delta P_n - k_n^P \Delta P_n - k_n^Q \Delta Q_n \quad (3a) \\ \text{s.t.} \quad & \Delta Q_n = \phi_n \Delta P_n. \quad (3b) \end{aligned}$$

where c_n and d_n are the coefficients of the quadratic and linear terms, respectively, in the aggregated cost function at node n , while ϕ_n is the ratio of the total reactive to active power responses at node n . Given the time-varying nature of values in the set of $\mathcal{S} = \{c_{n,t}, d_{n,t}, \phi_{n,t}\}_{n=1}^{N_b}$, they can be modeled as uncertain parameters, with their distribution characterized by samples. We will then demonstrate the process for deriving a set of samples for these uncertain parameters using observations at time t , denoted by $\{\Delta P_{n,t}, \Delta Q_{n,t}, k_{n,t}^P, k_{n,t}^Q\}_{n=1}^{N_b}$.

According to (3b), $\phi_{n,t}$ can be calculated directly by:

$$\phi_{n,t} = \Delta Q_{n,t} / \Delta P_{n,t} \quad (4)$$

By replacing $\Delta Q_{n,t}$ with $\phi_{n,t} \Delta P_{n,t}$ in (3), we transform it into an unconstrained optimization problem. For optimality, it requires the following first-order condition to hold:

$$2c_{n,t} \Delta P_{n,t} + d_{n,t} - k_{n,t}^P - k_{n,t}^Q \phi_{n,t} = 0 \quad (5)$$

However, (5) alone does not suffice to uniquely determine the specific values of $c_{n,t}$ and $d_{n,t}$. We need another time instance t' to derive an additional equation involving $c_{n,t'}$ and $d_{n,t'}$. By assuming that $c_{n,t} = c_{n,t'}$ and $d_{n,t} = d_{n,t'}$, we can calculate their values using these two equations. With this new set of estimated parameter values, the DSO can refine the distributions of c_n , d_n , and ϕ_n by incorporating these new samples into their historical sample sets.

Consider a sample set of $\{\mathcal{S}\}_{N_s}$ with N_s samples for each uncertain parameter. Based on each sample, we can derive the optimal response of the DERA at node n as a linear function of the incentives at node n based on the Karush-Kuhn-Tucker (KKT) conditions of (3):

$$\Delta P_{n,i} = \frac{1}{2c_{n,i}} k_n^P + \frac{\phi_{n,i}}{2c_{n,i}} k_n^Q - \frac{d_n}{2c_{n,i}}, \quad i = 1, \dots, N_s \quad (6)$$

$$\Delta Q_{n,i} = \frac{\phi_{n,i}}{2c_{n,i}} k_n^P + \frac{\phi_{n,i}^2}{2c_{n,i}} k_n^Q - \frac{\phi_{n,i} d_{n,i}}{2c_{n,i}}, \quad i=1, \dots, N_s \quad (7)$$

B. Properties of Surrogate Model

The selection of the surrogate model in (3) is justified by three key reasons:

- 1) The surrogate model (3) preserves certain characteristics of the original DERA decision model in (2). For instance, the quadratic function in (3a) is a good approximation of the piece-wise linear function with an incrementally increasing slope in (2a).
- 2) The surrogate model (3) features unknown parameters that can be accurately inferred using a relatively limited set of observations, i.e., through (4) and (5).
- 3) The surrogate model (3) facilitates the derivation of an explicit mapping relationship from k_n^P, k_n^Q to $\Delta P_n, \Delta Q_n$, as shown in (6) and (7).

We recognize that (3) is not the only option for a surrogate model. Any model meeting the aforementioned three criteria could serve as an effective surrogate of (2) and be used in the learning process.

The surrogate model (3), while sharing characteristics with the actual DERA decision model in (2), diverges in terms of modeling granularity. Specifically, (3) operates at the node level, whereas (2) focuses on individual DERs. This difference raises questions about the effectiveness of the proposed learning algorithm in capturing the relationship between incentives and DERA responses.

We claim that, despite the difference in granularity, (3) can effectively capture the essential relationship between k_n^P, k_n^Q and $\Delta P_n, \Delta Q_n$ determined by (2). This effectiveness is primarily because the DSO is concerned with the aggregated response at each node, rather than individual DER responses. For instance, if the constraint of $\Delta P_r \leq \Delta P_r^{\max}$ is not binding, variations in ΔP_r^{\max} will not impact the aggregate responses ΔP_n and ΔQ_n . Define the set of uncertain parameters in (2) as $\mathcal{S}^{\text{full}} = \{c_r, \Delta P_r^{\min}, \Delta P_r^{\max}, \phi_r^{\min}, \phi_r^{\max}, r = 1, \dots, R_n, n = 1, \dots, N_b\}$. Essentially, the distribution of \mathcal{S} originates from a specific subset of the distribution of $\mathcal{S}^{\text{full}}$, but only this subset is vital for the mapping from k_n^P, k_n^Q to $\Delta P_n, \Delta Q_n$. Consequently, the distribution of \mathcal{S} captures all the variations in $\mathcal{S}^{\text{full}}$ that matter to the DSO.

Additionally, the effectiveness of the proposed learning strategy can be further improved. For instance, instead of directly learning distributions $P(c_n)$, $P(d_n)$, and $P(\phi_n)$, it may be more beneficial to learn conditional distributions $P(c_n | \Psi_n)$, $P(d_n | \Psi_n)$, and $P(\phi_n | \Psi_n)$, where Ψ_n includes all pertinent external factors like time and temperature that could affect the availability of flexible DERs.

IV. DISTRIBUTIONALLY ROBUST INCENTIVE DESIGN

By replacing (1d) with (6) and (7), we convert (1) into a single-level robust optimal incentive design model that is more computationally tractable. However, employing robust optimization strategies can lead to overly conservative decisions. Additionally, future instances of uncertain parameters may deviate from their historical patterns, potentially leading to unforeseen adverse scenarios. In response, this section introduces a distributionally robust incentive design approach

that accounts for the empirical distribution of uncertainties based on all available samples and accommodates deviations within a predetermined range from the empirical distribution. We model voltage limits as a chance constraint, which effectively manages the risk of voltage level violations at each node. We also provide a convex reformulation of this distributionally robust incentive design model, facilitating its solution with standard optimization solvers.

A. Distributionally Robust Incentive Design Model

To streamline our notation, we introduce a new vector ξ_n for node n , encapsulating five uncertainties derived from the three initial uncertain parameters c_n, d_n, ϕ_n in (3):

$$\xi_n = [\xi_{n,1}, \dots, \xi_{n,5}] = \left[\frac{1}{2c_n}, \frac{\phi_n}{2c_n}, \frac{-d_n}{2c_n}, \frac{\phi_n^2}{2c_n}, \frac{-\phi_n d_n}{2c_n} \right]$$

The distributionally robust incentive design model is:

$$\inf_{k_n^P, k_n^Q} \sup_{\mathbb{P} \in \mathcal{A}} \mathbb{E}^{\mathbb{P}} \left[\sum_{n=1}^{N_b} (k_n^P \Delta P_n + k_n^Q \Delta Q_n) \right] \quad (8a)$$

$$\text{s.t. } n = 1, \dots, N_b$$

$$\Delta P_n = \xi_{n,1} k_n^P + \xi_{n,2} k_n^Q + \xi_{n,4} \quad (8b)$$

$$\Delta Q_n = \xi_{n,2} k_n^P + \xi_{n,3} k_n^Q + \xi_{n,5} \quad (8c)$$

$$\Delta V_n = \sum_{m=1}^{N_b} \alpha_{n,m} \Delta P_m + \sum_{m=1}^{N_b} \beta_{n,m} \Delta Q_m \quad (8d)$$

$$\inf_{\mathbb{P} \in \mathcal{A}} \mathbb{P} \left\{ \begin{array}{l} \Delta V_n^{\min} \leq \Delta V_n \\ \Delta V_n \leq \Delta V_n^{\max}, n=1, \dots, N_b \end{array} \right\} \geq 1-\gamma \quad (8e)$$

This model assumes that the DSO adopts a risk-averse stance regarding the voltage regulation problem. The objective in (8a) aims to minimize the voltage regulation cost in the worst-case scenario. The joint chance constraint in (8e) guarantees that the uncertain value of ΔV_n remains within the prescribed limits, ΔV_n^{\min} and ΔV_n^{\max} , with a minimum probability of $1 - \gamma$. The distribution of ξ_n is permitted to vary within an ambiguity set \mathcal{A} . In this paper, we construct this ambiguity set through the application of the Wasserstein metric due to its superior out-of-sample performance [17].

B. Wasserstein Metric-based Ambiguity Set

We first give the definition of the Wasserstein metric, a measure of distance between two probability distributions:

Definition 1: (Wasserstein Distance) The Wasserstein metric $d_W(\mathbb{P}_1, \mathbb{P}_2) : \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \rightarrow \mathbb{R}$ is defined by:

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

where Ξ is the domain of a random variable ξ (i.e., $\xi \in \Xi$), $\mathcal{M}(\Xi)$ contains all distributions defined on Ξ , $\Pi(\mathbb{P}_1, \mathbb{P}_2)$ is the set of joint probability distributions of ξ_1, ξ_2 with marginal distributions of \mathbb{P}_1 and \mathbb{P}_2 [17].

Given N_s historical samples $\{\hat{\xi}_i\}_{i=1}^{N_s}$, the empirical distribution $\hat{\mathbb{P}}_{N_s}$ can be derived as:

$$\hat{\mathbb{P}}_{N_s} = \frac{1}{N_s} \sum_{i=1}^{N_s} \delta_{\hat{\xi}_i}, \quad (9)$$

where $\delta_{\hat{\xi}_i}$ is a Dirac distribution centered at $\hat{\xi}_i$. The ambiguity set based on these historical samples is defined as a Wasserstein ball centered at $\hat{\mathbb{P}}_{N_s}$:

$$\mathcal{B}_\epsilon(\hat{\mathbb{P}}_{N_s}) = \left\{ \mathbb{P} \in \mathcal{M}(\Xi) \mid d_W(\mathbb{P}, \hat{\mathbb{P}}_{N_s}) \leq \epsilon \right\}, \quad (10)$$

where ϵ is the radius of the Wasserstein ball. Here we treat ϵ as a predefined value. In Section V, we will study the method for adaptively adjusting the value of ϵ to optimize performance in voltage regulation.

C. Convex Formulation of Distributionally Robust Model

The distributionally robust incentive design model in (8) presents challenges for direct solution until it is converted into a convex optimization problem. This conversion leverages the conditional value-at-risk (CVaR) and the Wasserstein metric-based ambiguity set. To maintain clarity and conciseness, we omit the detailed steps of this reformulation. Readers seeking a deeper understanding of these steps are directed to [17] for further information.

For leaner notation in the following derivations, we collect all decision variables in vector $x \in \mathbb{R}^{6N_b}$ and all uncertainties in vector $\xi \in \mathbb{R}^{6N_b}$, i.e.,

$$\begin{aligned} x &= [(k_1^P)^2, \dots, (k_{N_b}^P)^2, 2k_1^P k_1^Q, \dots, 2k_{N_b}^P k_{N_b}^Q, \\ &\quad (k_1^Q)^2, \dots, (k_{N_b}^Q)^2, k_1^P, \dots, k_{N_b}^P, k_1^Q, \dots, k_{N_b}^Q, \mathbf{1}_{N_b}]^\top \\ \xi &= [\xi_{1,1}, \dots, \xi_{N_b,1}, \dots, \xi_{1,5}, \dots, \xi_{N_b,5}, \mathbf{0}_{N_b}]^\top, \end{aligned}$$

where $\mathbf{1}_{N_b}$ and $\mathbf{0}_{N_b}$ denote all-one and all-zero vectors of length N_b . Then, we can re-write (8) as:

$$\inf_{x \in \mathcal{F}} \sup_{\mathbb{P} \in \mathcal{B}_\epsilon} \mathbb{E}^{\mathbb{P}} [\langle x, \xi \rangle] \quad (11a)$$

$$\text{s.t. } \inf_{\mathbb{P} \in \mathcal{B}_\epsilon} \mathbb{P} \left\{ \max_{k=1, \dots, 2N_b} [a'_k \xi + b'_k] \leq 0 \right\} \geq 1-\gamma, \quad (11b)$$

where \mathcal{F} is the feasible space of x , and \mathbb{P} is the possible distribution of ξ constrained by a Wasserstein ball \mathcal{B}_ϵ . In (11b), a'_k is the k -th row of this $2N_b \times 6N_b$ matrix:

$$\begin{bmatrix} (-W_1 x)^\top \\ \vdots \\ (-W_{N_b} x)^\top \\ (W_1 x)^\top \\ \vdots \\ (W_{N_b} x)^\top \end{bmatrix}, \quad W_n = \begin{bmatrix} & A_n & \mathcal{O}_{1 \times 1} & \mathcal{O}_{1 \times 1} \\ & B_n & A_n & \mathcal{O}_{1 \times 1} \\ \mathcal{O}_{6 \times 3} & \mathcal{O}_{1 \times 1} & B_n & \mathcal{O}_{1 \times 1} \\ & \mathcal{O}_{1 \times 1} & \mathcal{O}_{1 \times 1} & A_n \\ & \mathcal{O}_{1 \times 1} & \mathcal{O}_{1 \times 1} & B_n \\ & \mathcal{O}_{1 \times 1} & \mathcal{O}_{1 \times 1} & \mathcal{O}_{1 \times 1} \end{bmatrix}$$

$$A_n = \mathcal{D}(\alpha_{n,1}, \dots, \alpha_{n,N_b}), \quad B_n = \mathcal{D}(\beta_{n,1}, \dots, \beta_{n,N_b}),$$

where $\mathcal{D}(\cdot)$ creates a diagonal matrix from a vector, $\mathcal{O}_{m \times n}$ is a all-zero matrix with a dimension of $mN_b \times nN_b$. Similarly, b'_k is the k -th entry of the following $2N_b \times 1$ vector:

$$[\Delta V_1^{\min}, \dots, \Delta V_{N_b}^{\min}, -\Delta V_1^{\max}, \dots, -\Delta V_{N_b}^{\max}]^\top$$

Following [17], we can transform (11) into the following

convex problem:

$$\min_{x, \tau, \lambda_j^{co}, \lambda_j^{cc}, s_{ji}^{co}, s_{ji}^{cc}} \sum_{j=1}^{6N_b} \left(\lambda_j^{co} \epsilon + \frac{1}{N_s} \sum_{i=1}^{N_s} s_{ji}^{co} \right) \quad (12a)$$

$$\text{s.t. } i = 1, \dots, N_s, j = 1, \dots, 6N_b, k = 0, \dots, 2N_b$$

$$(\rho_{ij}^1) s_{ji}^{co} \geq x_j \hat{\xi}_{ji} \quad (12b)$$

$$(\rho_{ij}^2) s_{ji}^{co} \geq x_j \bar{\xi}_j - \lambda_j^{co} |\bar{\xi}_j - \hat{\xi}_{ji}| \quad (12c)$$

$$(\rho_{ij}^3) s_{ji}^{co} \geq x_j \underline{\xi}_j - \lambda_j^{co} |\underline{\xi}_j - \hat{\xi}_{ji}| \quad (12d)$$

$$(\mu) \gamma \tau + \sum_{j=1}^{6N_b} \left(\lambda_j^{cc} \epsilon + \frac{1}{N_s} \sum_{i=1}^{N_s} s_{ji}^{cc} \right) \leq 0 \quad (12e)$$

$$(\phi_{ikj}^1) s_{ji}^{cc} \geq a_{kj} \hat{\xi}_{ji} + b_k \quad (12f)$$

$$(\phi_{ikj}^2) s_{ji}^{cc} \geq a_{kj} \bar{\xi}_j + b_k - \lambda_j^{cc} |\bar{\xi}_j - \hat{\xi}_{ji}| \quad (12g)$$

$$(\phi_{ikj}^3) s_{ji}^{cc} \geq a_{kj} \underline{\xi}_j + b_k - \lambda_j^{cc} |\underline{\xi}_j - \hat{\xi}_{ji}| \quad (12h)$$

$$\lambda_j^{co}, \lambda_j^{cc} \geq 0, \quad (12i)$$

Eq. (12a)-(12d) are the *exact* reformulation of the objective in (11a), while (12e)-(12h) are the *inner* approximation of the chance constraint in (11b). This reformulation is derived by introducing $\bar{a}_0 = \mathbf{0}_{6N_b}$, $\bar{b}_0 = 0$, $\bar{a}_k = a'_k$, and $\bar{b}_k = b'_k - \tau$, $k = 1, \dots, 2N_b$. Then, a_{kj} is the j -th coordinate of \bar{a}_k , $b_k = \bar{b}_k / 6N_b$, $\underline{\xi}_j$ and $\bar{\xi}_j$ are the lower and upper bounds, respectively, of the j -th coordinate of the domain of ξ . Additionally, τ , λ_j^{co} , λ_j^{cc} , s_{ji}^{co} and s_{ji}^{cc} are auxiliary variables. The superscript in λ_j^{co} and s_{ji}^{co} highlights its connection to the cost function reformulation. Similarly, the superscript in λ_j^{cc} and s_{ji}^{cc} highlight its connection to the chance-constraint reformulation. Greek letters listed on the left side of (12e)-(12d) are dual variables of the corresponding constraints.

Model (12) constitutes a finite-dimensional convex problem that can be efficiently solved using commercial optimization solvers. The number of constraints in (12) correlates directly with N_s , which is the number of samples factored into the empirical distribution $\hat{\mathbb{P}}_{N_s}$.

V. CONSERVATIVENESS UPDATE IN INCENTIVE DESIGN

In the distributionally robust incentive design model (12), the size of the Wasserstein metric-based ambiguity set, indicated by the value of ϵ , determines the conservativeness of the decision-makers, mirroring their confidence in the empirical distribution of ξ . A larger ϵ incorporates a broader range of possible distributions in the decision making model, leading to decisions that are more conservative. Intuitively, as learning in Section III progresses, the DSO will obtain an empirical distribution of ξ that increasingly better approximates its true distribution. Consequently, the DSO should become less conservative in incentive design.

In this section, we propose a method for assessing the suitability of the conservativeness level *post hoc* and a conservativeness updating algorithm to optimize voltage regulation performance. Fig.2 illustrates how the conservativeness level can be updated within the proposed framework. The inner loop depicted in Fig.2 combines the distributionally robust incentive design model introduced in Section IV (i.e., the green block in Fig. 2) with the DERA decision-making

model outlined in Section II (i.e., the orange block in Fig. 2). After executing the inner loop several times, such as over a day or multiple days, the DSO can use the learning approach detailed in Section III to refine its understanding of the empirical distribution (i.e., the red block in Fig. 2). Meanwhile, the DSO can adjust its level of conservativeness using the algorithm described in this section (i.e., the yellow block in Fig. 2). This belief update process, or outer loop, is repeated until the empirical distribution of uncertainty approximates the true distribution acceptably well, and the conservativeness level ϵ stabilizes at an optimal value.

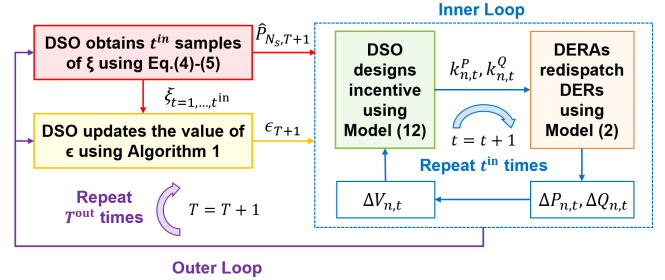


Fig. 2. Distributionally robust incentive design for voltage regulation with learning and adaptive conservativeness level.

A. Adequacy of Current Conservativeness Level

To determine the adequacy of the current level of conservativeness, we need to understand the rationale behind employing a distributionally robust approach for incentive design. The DSO has two objectives when solving the distributionally robust incentive design model in (8):

- 1) Objective 1: Minimize the discrepancy between the expected and actual worst-case costs associated with voltage regulation, as outlined in (8a).
- 2) Objective 2: Ensure the expected worst-case voltage violation rate, denoted by γ in (8e), aligns closely with the actual occurrence rate of such violations.

In this paper, the conditional value-at-risk (CVaR) is used to measure the violation rate of a chance-constraint. As demonstrated in [18], enforcing a $\text{CVaR} \leq 0$ ensures that the corresponding chance constraint is satisfied.

Throughout the time steps in the inner loop, the expected worst-case cost at each specific time t , labeled as $\text{Cost}_t^{\text{exp}}$, is equal to the objective value of (12a) at time t . The expected CVaR, denoted as CVaR^{exp} , is a statistical value derived from results over all time steps. Specifically, we have $\text{CVaR}^{\text{exp}} = 0$ if the chance constraint in (8e) is consistently met in the incentive design process.

On the other hand, the actual cost and CVaR, denoted by $\text{Cost}_t^{\text{act}}$ and CVaR^{act} respectively, are determined by the actual DERA responses over the past t^{in} steps:

$$\text{Cost}_t^{\text{act}} = \sum_{n=1}^{N_b} \left(k_{n,t}^P \Delta P_{n,t} + k_{n,t}^Q \Delta Q_{n,t} \right), \quad t = 1, \dots, t^{\text{in}} \quad (13)$$

$$\text{CVaR}^{\text{act}} = \inf_{\gamma} \gamma \tau + \frac{1}{t^{\text{in}}} \sum_{t=1}^{t^{\text{in}}} \left[\max_{k=0,1,\dots,2N_b} \langle \bar{a}_k, \xi_t \rangle + \bar{b}_k \right]. \quad (14)$$

Calculating CVaR^{act} requires solving the optimization problem in (14) using ξ_t , which are the realizations of uncertainty ξ over the past t^{in} steps. These new realizations of ξ are estimated by the DSO using the learning method in Section III.

We establish the following loss function to evaluate the performance of ϵ over the preceding t^{in} time steps:

$$\mathcal{L}(\epsilon) = (\text{Cost}_{\underline{t}}^{exp} - \text{Cost}_{\underline{t}}^{act})^2 + (\text{CVaR}^{exp} - \text{CVaR}^{act})^2, \quad (15)$$

where \underline{t} denotes the time when the actual worst-case scenario occurs, i.e., when the discrepancy between the expected worst-case cost and the actual cost is minimized:

$$\underline{t} = \arg \min_{t=1, \dots, t^{in}} |\text{Cost}_t^{exp} - \text{Cost}_t^{act}| \quad (16)$$

It is sufficient to focus on \underline{t} in the loss function (15) because the actual costs in scenarios other than the worst-case do not reflect the performance of risk-averse incentives and, therefore, should not be employed to update ϵ .

B. Conservativeness Update Algorithm

Based on the loss function in (15), we can use the gradient descent method to find the optimal value of ϵ following the update rule of:

$$\epsilon' = \epsilon - \chi \frac{\partial \mathcal{L}(\epsilon)}{\partial \epsilon}, \quad (17)$$

where χ is a predefined learning rate. The gradient term in (17) can be further expressed as:

$$\begin{aligned} \frac{\partial \mathcal{L}(\epsilon)}{\partial \epsilon} &= 2(\text{Cost}_{\underline{t}}^{exp} - \text{Cost}_{\underline{t}}^{act}) \frac{\partial \text{Cost}(\epsilon)}{\partial \epsilon} \\ &\quad + 2(\text{CVaR}^{exp} - \text{CVaR}^{act}) \frac{\partial \text{CVaR}(\epsilon)}{\partial \epsilon}, \end{aligned} \quad (18)$$

Following the convex formulation in (12), the worst-case cost and CVaR can be written as functions of ϵ :

$$\text{Cost}(\epsilon) = \sum_{j=1}^{6N_b} \left(\lambda_j^{co} \epsilon + \frac{1}{N_s} \sum_{i=1}^{N_s} s_{ji}^{co} \right) \quad (19)$$

$$\text{CVaR}(\epsilon) = \mu \left[\tau \gamma + \sum_{j=1}^{6N_b} \left(\lambda_j^{cc} \epsilon + \frac{1}{N_s} \sum_{i=1}^{N_s} s_{ji}^{cc} \right) \right], \quad (20)$$

where (19) is the same as the objective function in (12a), and (20) is the complementary slackness condition based on (12e). Therefore, the two partial derivatives in (18) can be calculated as:

$$\frac{\partial \text{Cost}(\epsilon)}{\partial \epsilon} = \sum_{j=1}^{6N_b} \lambda_{j,\underline{t}}^{co}, \quad (21)$$

$$\frac{\partial \text{CVaR}(\epsilon)}{\partial \epsilon} = \frac{1}{t^{in}} \sum_{t=1}^{t^{in}} \mu_t \sum_{j=1}^{6N_b} \lambda_{j,t}^{cc}, \quad (22)$$

where $\lambda_{j,t}^{co}$ and $\lambda_{j,t}^{cc}$ are auxiliary variables in (12) and can be obtained directly from solving (12), while μ_t is a dual variable in (12) and can be returned directly by some optimization solvers or can be derived by the KKT conditions of (12). The two partial derivatives in (21) and (22) are always non-negative since (12) requires $\lambda_{j,t}^{co}, \lambda_{j,t}^{cc}, \mu_t \geq 0$ for any j and t . Algorithm 1 summarizes this gradient-based update of ϵ .

Algorithm 1: Adaptive Conservativeness Level

input : $\{k_{n,t}^P, k_{n,t}^Q, \Delta P_{n,t}, \Delta Q_{n,t}\}_{n=1, \dots, N_b, t=1, \dots, t^{in}}$;
 $\{\lambda_{j,t}^{co}, \lambda_{j,t}^{cc}, \mu_t\}_{j=1, \dots, 6N_b, t=1, \dots, t^{in}}$;
Initial values $\{\text{cost}_t^{exp}\}_{t=1, \dots, t^{in}}$;
Current conservativeness level ϵ_T ;
Maximum iterations allowed count^{\max} ;
Convergence threshold $\Delta \epsilon^{\min}$;

output: Updated conservativeness level ϵ_{T+1} ;

begin

Calculated cost^{act} and CVaR^{act} using (13) and (14);
Calculated two partial derivatives using (21) and (22);
 $\Delta \epsilon \leftarrow \epsilon_T$; $\epsilon \leftarrow \epsilon_T$; $\text{count} \leftarrow 0$

while $\text{count} < \text{count}^{\max}$ and $\Delta \epsilon > \Delta \epsilon^{\min}$ **do**

Calculate gradient following (18);

Obtain ϵ' following (17);

Re-run (12) for t^{in} times with ϵ' to obtain updated values of $\{\text{cost}_t^{exp}\}_{t=1, \dots, t^{in}}$;

$\Delta \epsilon \leftarrow |\epsilon - \epsilon'|$; $\epsilon \leftarrow \epsilon'$; $\text{count} \leftarrow \text{count} + 1$

end

$\epsilon_{T+1} \leftarrow \epsilon$

return ϵ_{T+1}

end

C. Convergence Analysis of Proposed Algorithm

We first intuitively analyze the expected outcomes of Algorithm (1) at each iteration T . Given that the partial derivative terms in (18) are always non-negative, the direction in which ϵ is updated, indicated by the sign of $\frac{\partial \mathcal{L}(\epsilon)}{\partial \epsilon}$, depends on relative magnitudes of the expected versus actual cost and CVaR values. If $\text{Cost}_{\underline{t}}^{exp} > \text{Cost}_{\underline{t}}^{act}$, indicating that the DSO is overly conservative by planning for adverse scenarios that never materialize, then the value of ϵ should be decreased. Similarly, if $\text{CVaR}^{exp} > \text{CVaR}^{act}$, suggesting that the DSO has prepared excessively for constraint violations that occur less frequently in practice, then ϵ should also be reduced.

It is also possible that the two terms in (18) may indicate differing directions for updating ϵ . This is because (8) considers two worst-case distributions that may not coincide: one associated with the maximum cost in (8a) and the other corresponding to the maximum violation rate of the chance constraint in (8e). In this case, ϵ will be updated using a weighted sum of the two directions suggested by both the cost and CVaR, and the weights are the partial derivatives in (21) and (22). As highlighted in [19], the values of λ_j^{co} and λ_j^{cc} can be interpreted as the marginal values of ϵ in the optimal solution of (12). Therefore, the weights of the cost and the CVaR terms in (18) indicate their relative importance within the context of risk-averse decision-making.

Formally, we propose and prove the following proposition for the convergence of ϵ_T based on Algorithm 1:

Proposition 1: (Convergence of ϵ_T) For each iteration $T \leq T^{\text{out}}$, Algorithm 1 will return a near optimal value of ϵ_T , under the condition that count^{\max} is adequately large, and both the learning rate χ and the convergence threshold $\Delta \epsilon^{\min}$ are set to sufficiently small values.

Proof: The essence of proving this proposition lies in showing the convexity of loss function $\mathcal{L}(\epsilon)$ with respect to ϵ ,

in accordance with the well-established convergence theorem of the gradient descent method for convex functions [20]. This requires an analysis of the second-order derivative:

$$\frac{\partial^2 \mathcal{L}(\epsilon)}{\partial \epsilon^2} = 2 \left(\frac{\partial \text{Cost}(\epsilon)}{\partial \epsilon} \right)^2 + 2 \left(\frac{\partial \text{CVaR}(\epsilon)}{\partial \epsilon} \right)^2$$

which is always non-negative. Therefore, the $\mathcal{L}(\epsilon)$ is convex in ϵ , implying the convergence of the gradient-based method in Algorithm 1. ■

The optimal value of ϵ conveys a tangible significance as it represents the weighted average of the two Wasserstein distances. The first is the distance between the empirical distribution $\hat{\mathbb{P}}_{N_s}$ and the distribution that results in $\text{Cost}_t^{\text{act}}$, while the second is the distance between $\hat{\mathbb{P}}_{N_s}$ and the distribution leading to CVaR^{act} . Essentially, the optimal value of ϵ reflects the distance between the empirical and true distributions of ξ .

D. Conservativeness Change in Long Term

Regrading the change of ϵ_T when T increases, we propose and prove the following proposition:

Proposition 2: (Convergence of ϵ_∞) Assuming the true distribution of ξ is stationary (i.e., it does not change over time) and the iterations within the inner loop $t^{\text{in}} \rightarrow \infty$, then the optimal value of $\epsilon_T \rightarrow 0$ when $T \rightarrow \infty$.

Proof: Proving this proposition entails the proofs of three key aspects: i) the optimal value of ϵ_T decreases monotonically as T increases; ii) the minimal possible value for ϵ_T is 0, and iii) the optimal value of ϵ_T is stable when $T \rightarrow \infty$. First, with the assumption that $t^{\text{in}} \rightarrow \infty$ and the stationary nature of the true distribution, the sample distribution characterized by $\{\xi_t\}_{t=1}^{t^{\text{in}}}$ effectively mirrors the true distribution of ξ for each iteration T . As T progresses, the inclusion of additional samples into the historical dataset implies that the Wasserstein distance between the empirical distribution $\hat{\mathbb{P}}_{N_s}$ and the true distribution should decrease monotonically. Second, $\epsilon = 0$ implies a perfect match between the empirical and true distributions of ξ , which is achievable through the learning algorithm outlined in Section III given sufficient learning iterations T . In such scenarios, the expected and actual costs, as well as the expected and actual CVaR, converge and thus yield a zero gradient in (18). This outcome indicates the stability of ϵ_∞ , thereby proving the third point. ■

However, the two assumptions made in Proposition 2 might not always apply in real-world scenarios. First, the inner loop is limited to a finite number of iterations, which means the sample distribution might not accurately represent the true distribution of ξ . Consequently, gradient estimates could be imprecise or noisy due to the reliance on a limited number of samples, causing fluctuations in the convergence trajectory. This scenario bears resemblance to the behavior of stochastic gradient descent (SGD) as opposed to gradient descent (GD), where the former is known for its variability in the path towards convergence due to its sample-based gradient updates.

On the other hand, the true distribution of ξ may not be stationary due to factors like the installation of new DERs, which may change nodal flexibility regions, or fluctuations in residential electricity prices which may affect DER participation costs. In this case, the empirical distribution learned by the DSO may not perfectly match the true distribution, indicating that the optimal value of ϵ will not converge to 0. Regarding this, we have the following remark:

Remark 1: If changes in the true distribution of ξ are slow and smooth, the optimal value of ϵ_T is expected to stabilize at a small positive value when $T \rightarrow \infty$. This reflects the rate of change in the true distribution.

Proposition 2 also informs what steps should decision-makers take once ϵ_T reaches 0. With the optimal level of conservativeness being 0, a transition from a risk-averse to a risk-neutral stance is recommended. This shift involves minimizing the expected cost rather than the worst-case costs in the objective and considering the expected rate of constraint violations instead of the worst-case violation rate.

If T keeps increasing, more samples will be added into the empirical distribution without providing new information, it may be appropriate to stop the learning process. Furthermore, given that the computational complexity of a sample-based decision-making model depends on the number of samples, reducing the sample count through sample selection techniques or diminishing the size of the uncertainty set could be considered to improve the efficiency.

VI. NUMERICAL EXPERIMENTS

This section considers interactions between one DSO and one DERA to show the effectiveness of the proposed method. We set the initial value of $\epsilon = 0.01$ and the learning rate $\chi = 0.001$. The iteration number in the inner loop is $t^{\text{in}} = 5$, and the true distribution of ξ is assumed to follow a uniform distribution between -1 and 2. Furthermore, We assume that initially the DSO has access to 10 samples of ξ . This initial count, set lower than what might be found in real-world situations, is chosen to more effectively highlight the convergence trend of ϵ as the number of available samples increases. All the numerical experiments are conducted using the Gurobi solver [21] in Python and are performed on a PC with an Intel Core i9, 2.20 GHz with 16 GB RAM.

We first examine the convergence of the optimal conservativeness level ϵ when ϵ_T is obtained using Algorithm 1 at each iteration T . Fig. 3 (a) shows the value of ϵ_T as T increases, where the black line is a representative convergence trajectory, and the shaded area represents the confidence interval derived from conducting the experiment for 50 times. It is clear that ϵ_T converges to 0, with some fluctuations along the convergence trajectory. These fluctuations are attributed to using only five samples of ξ in the inner loop to approximate its true distribution.

The computational challenge presented by Algorithm 1 stems from iteratively solving the distributionally robust incentive design problem in (12). A practical simplification approach involves updating ϵ just once per iteration T using (17), rather than continuously updating until the optimal

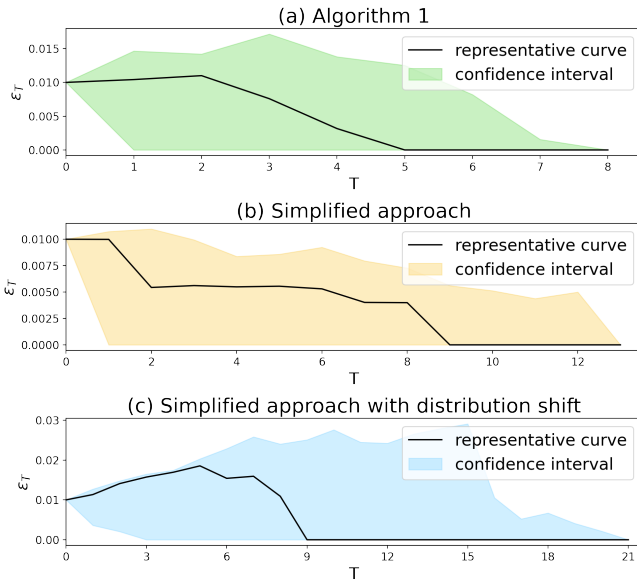


Fig. 3. Convergence trajectory of optimal conservativeness level.

value of ϵ_T is reached. Fig.3 (b) illustrates the convergence trajectory of ϵ_T using this simplified approach. While the convergence is slower compared to the trajectory shown in Fig.3 (a), ϵ_T still approaches 0 after several iterations. This highlights a viable pathway for enhancing the computational efficiency of the proposed algorithm.

The robustness of this simplified method is further evaluated in a more realistic scenario with shifts in the true distribution of ξ . At each iteration T , the DERA responses in the inner loop are generated based on a varying distribution of ξ . More specifically, the bounds of the uniform distribution of ξ , initially set at -1 and 2, are expanded at a constant rate of 0.1 in every iteration. Fig. 3 (c) shows that ϵ converges to 0 despite the distribution shift.

VII. CONCLUSION AND FUTURE WORK

This paper proposes to learn parameters of the distributionally robust optimization applied to incentive design for voltage regulation with adjustable conservativeness. The incentive design problem is reformulated as a convex model, facilitating efficient resolution with standard solvers. We also prove the convergence of the gradient-based method for updating the conservativeness level. Additionally, we analyze the optimal conservativeness level, establishing a connection between risk-averse decision-making and the learning process. Numerical experiments on a single DSO and a single DERA scenario shows the effectiveness of our proposed approach.

One limitation of the proposed method is its assumption of a fully rational DERA, meaning it presumes the DER responses are dictated by an optimization model. In future research, we plan to employ a more realistic decision-making model for DERAs, for instance, by incorporating the concept of bounded rationality [22].

REFERENCES

- [1] F. Ding, W. Liu, U. Kumar, and Y. Yao, "Unleash values from grid-edge flexibility: An overview, experience, and vision for leveraging grid-edge distributed energy resources to improve grid operations," *IEEE Electrification Magazine*, vol. 10, no. 4, pp. 29–37, 2022.
- [2] Y. Y. Zhang, C. Feng, P. Shaffery, and R. Yang, "Grid edge visibility: Gaps and a road map," *IEEE Electrification Magazine*, vol. 10, no. 4, pp. 10–19, 2022.
- [3] X. Zhou, E. Dall'Anese, L. Chen, and A. Simonetto, "An incentive-based online optimization framework for distribution grids," *IEEE Trans. Autom. Control*, vol. 63, no. 7, pp. 2019–2031, 2017.
- [4] M. Yu, S. H. Hong, Y. Ding, and X. Ye, "An incentive-based demand response (dr) model considering composited dr resources," *IEEE Trans. Ind. Electron. Control Instrum.*, vol. 66, no. 2, pp. 1488–1498, 2018.
- [5] J. Yang, M. R. Alam, W. Tushar, and T. K. Saha, "Incentivizing prosumer voltage regulation for unbalanced radial distribution networks," *IEEE Trans. Sustain. Energy*, 2023.
- [6] S. Bialek, Y. Dvorkin, J. Kim, and B. Unel, "Who knows what: Information barriers to efficient der roll-out," *USAEE Working Paper*, 2021.
- [7] D. Zhang, S. Li, M. Sun, and Z. O'Neill, "An optimal and learning-based demand response and home energy management system," *IEEE Trans. Smart Grid*, vol. 7, no. 4, pp. 1790–1801, 2016.
- [8] Y. Bian, N. Zheng, Y. Zheng, B. Xu, and Y. Shi, "Demand response model identification and behavior forecast with optnet: A gradient-based approach," in *Proc. of the 13th ACM Int. Conf. on Future Energy Systems*, 2022, pp. 418–429.
- [9] R. Mieth and Y. Dvorkin, "Online learning for network constrained demand response pricing in distribution systems," *IEEE Trans. Smart Grid*, vol. 11, no. 3, pp. 2563–2575, 2019.
- [10] N. Tucker, A. Moradipari, and M. Alizadeh, "Constrained thompson sampling for real-time electricity pricing with grid reliability constraints," *IEEE Transactions on Smart Grid*, vol. 11, no. 6, pp. 4971–4983, 2020.
- [11] A. Hassan, R. Mieth, D. Deka, and Y. Dvorkin, "Stochastic and distributionally robust load ensemble control," *IEEE Trans. Power Syst.*, vol. 35, no. 6, pp. 4678–4688, 2020.
- [12] Q. Li, Z. Liang, A. Bernstein, and Y. Dvorkin, "Revealing decision conservativeness through inverse distributionally robust optimization," *IEEE Control Systems Letters*, 2024.
- [13] D. Bertsimas, M. Sim, and M. Zhang, "Adaptive distributionally robust optimization," *Management Science*, vol. 65, no. 2, pp. 604–618, 2019.
- [14] I. Wang, C. Becker, B. V. Parys, and B. Stellato, "Learning for robust optimization," *arXiv preprint arXiv:2305.19225*, 2023.
- [15] R. Mieth and H. V. Poor, "Prescribed robustness in optimal power flow," *Electric Power Systems Research*, vol. 235, p. 110704, 2024.
- [16] A. Bernstein and E. Dall'Anese, "Linear power-flow models in multiphase distribution networks," in *2017 IEEE PES Innovative Smart Grid Technologies Conference Europe*. IEEE, 2017, pp. 1–6.
- [17] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations," *Math. Program.*, vol. 171, no. 1, pp. 115–166, 2018.
- [18] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–42, 2000.
- [19] R. Mieth, J. M. Morales, and H. V. Poor, "Data valuation from data-driven optimization," *IEEE Transactions on Control of Network Systems*, pp. 1–12, 2024.
- [20] S. Boyd and L. Vandenberghe, *Convex optimization*. New York, USA: Cambridge university press, 2004.
- [21] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: <https://www.gurobi.com>
- [22] A. Marín Radoszynski and P. Pinson, "Electric demand response and bounded rationality: mean-field control for large populations of heterogeneous bounded-rational agents," *Philosophical Transactions of the Royal Society A*, vol. 379, no. 2202, p. 20190429, 2021.