

# Constrained Thompson Sampling for Real-Time Electricity Pricing with Grid Reliability Constraints

Nathaniel Tucker<sup>†</sup> Ahmadreza Moradipari<sup>†</sup> Mahnoosh Alizadeh

<sup>†</sup>Authors have equal contribution

**Abstract**—We consider the problem of an aggregator attempting to learn customers’ load flexibility models while implementing a load shaping program by means of broadcasting daily dispatch signals. We adopt a multi-armed bandit formulation to account for the stochastic and unknown nature of customers’ responses to dispatch signals. We propose a constrained Thompson sampling heuristic, Con-TS-RTP, that accounts for various possible aggregator objectives (e.g., to reduce demand at peak hours, integrate more intermittent renewable generation, track a desired daily load profile, etc) and takes into account the operational constraints of a distribution system to avoid potential grid failures as a result of uncertainty in the customers’ response. We provide a discussion on the regret bounds for our algorithm as well as a discussion on the operational reliability of the distribution system’s constraints being upheld throughout the learning process.

**Index Terms**—Constrained optimization, distribution network, multi-armed bandit, real-time pricing, demand response, Thompson sampling.

## I. INTRODUCTION

In order to integrate the increasing volume of intermittent renewables in modern power grids, aggregators are exploring various methods to manipulate both residential and commercial loads in real-time. As a result, various demand response (DR) frameworks are gaining popularity because of their ability to shape electricity demand by broadcasting time-varying signals to customers; however, most aggregators have not implemented complex DR programs beyond peak shaving and emergency load reduction initiatives. One reason for this is the customers’ unknown and time-varying responses to dispatch signals, which can lead to economic uncertainty for the aggregator and reliability concerns for the grid operator relying on DR performance [1]. The aggregator could explicitly request response information from its customers; however, this process would have a large communication overhead. More importantly, most customers cannot readily characterize their response, and even if they could, they might not be willing to share this private information. With this in mind, it is evident that future load shaping initiatives for renewable integration (i.e., more complex objectives than peak shaving) need to *passively* learn customers’ response only from historical data of past interactions [2].

Recently, much work has been done for aggregators attempting to learn customers’ price responses whilst implementing peak shaving DR programs. The authors of [3] present a data-driven strategy to estimate customers’ demands and develop prices for DR. In [4], the authors use linear regression models to derive estimations of customers’ responses to DR signals. Similarly, [5] develops a joint online learning and pricing algorithm based on linear regression. In [6], the authors present a contract-based DR strategy to learn customer behavior while broadcasting DR signals. The authors of [7] present an online learning approach based on piecewise linear stochastic approximation for an aggregator to sequentially adjust its DR prices based on the behavior of the customers in the past. In [8], the authors develop a risk-averse learning approach for aggregators operating DR programs. In [9], a learning algorithm for customers’ utility functions is developed and it is assumed that the aggregator acts within a two-stage (day-ahead and real-time) electricity market. Furthermore, a multi-armed bandit (MAB) formulation is used in [10], [11] to determine which customers to target with load reduction signals for DR programs.

In addition to learning how customers respond to DR signals, an aggregator must also consider power system constraints to ensure reliable operation (e.g., nodal voltage, transformer capacities, and line flow limits). In real distribution systems, it is critical that these constraints are satisfied at every time step to ensure customers receive adequate service and to avoid potential grid failures even without sufficient knowledge about how customers respond to price signals (i.e., in early learning stages) [12], [13]. One paper that considers these realistic constraints, [14], presents a least-square estimator approach to learn customer sensitivities and implements DR in a distribution network. However, the proposed learning approach does not have a regret guarantee compared to the clairvoyant solution that has full information of customer sensitivities.

Similar to the aforementioned papers, the work presented in this manuscript considers the problem of an aggregator *passively* learning the customers’ price sensitivities while running a load shaping program. However, our approach permits more complex load shaping objectives (e.g., tracking a daily target load profile) and varies in terms of both load modeling and learning approach from all the above papers. Specifically, we present a modified multi-armed bandit (MAB) heuristic akin to Thompson sampling (TS) to tackle

This work was supported by NSF grants #1847096 and #1737565 and UCOP Grant LFR-18-548175. N. Tucker, A. Moradipari, and M. Alizadeh are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106 USA (email: nathaniel\_tucker@ucsb.edu)

the trade-off between exploration of untested price signals and exploitation of well-performing price signals while ensuring grid reliability (A preliminary version of this work was published in [15]; however, it did not account for distribution system constraints). It is important to note that standard TS cannot guarantee that grid reliability constraints are upheld during the learning process. As such, we present a *modified* version of TS while retaining the fundamental principles TS is based on. Furthermore, we provide discussion on how the constraints are upheld (i.e., operational reliability), discussion on the performance of the heuristics compared to a clairvoyant solution, and simulation results highlighting the strengths of the method.

The remainder of the paper is organized as follows: Section II presents the aggregator’s daily objective as well as the customers’ load model. Section III describes the multi-armed bandit formulation for the electricity pricing problem, presents the modified TS heuristic, and discusses its performance and reliability. Section IV presents simulation results that showcase the efficacy of the approach.

## II. PROBLEM FORMULATION

### A. The Aggregator’s Objective

The aggregator’s main objective is to select dispatch signals to manipulate customer demand according to a given optimization objective that varies daily. Specifically, we consider the case where the aggregator broadcasts a dispatch signal  $\mathbf{p}_\tau = [p(t)]_{t=1,\dots,T}$  to the population of customers each day (we use  $t = 1, \dots, T$  to index time of day and  $\tau = 1, \dots, \mathcal{T}$  to index days). The set of dispatch signals available for use by the aggregator is denoted as  $\mathcal{P}$ . In this paper, without a loss of generality, we will assume that the dispatch signal sent to customers for load shaping purposes is a real-time pricing (RTP) signal\*.

On each day,  $\tau$ , the aggregator’s cost function is a fixed and known nonlinear function  $f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau)$  that depends on the load profile  $\mathbf{D}_\tau(\mathbf{p}_\tau)$  of the population in response to the daily price  $\mathbf{p}_\tau$  and a random exogenous parameter  $\mathbf{V}_\tau$ . The exogenous and given vector  $\mathbf{V}_\tau$  varies daily and can correspond to a daily target profile reflecting renewable generation forecasts, weather predictions, and grid conditions.

Moreover, the aggregator must ensure that the broadcasted price signals do not result in load profiles that violate distribution system reliability constraints (e.g., nodal voltage, transformer capacities, or line flow limits). As such, if the aggregator had full information about how the population responds to price signals (i.e., full knowledge of  $\mathbf{D}_\tau(\mathbf{p}_\tau)$ ),

the aggregator can solve the following optimization problem on day  $\tau$  to select the optimal price  $\mathbf{p}_\tau^*$ :

$$\mathbf{p}_\tau^* = \arg \min_{\mathbf{p}_\tau \in \mathcal{P}} f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau) \quad (1)$$

$$\text{s.t.} \quad g_j(\mathbf{D}_\tau(\mathbf{p}_\tau)) \leq 0, \quad \forall j = 1, \dots, J \quad (2)$$

where  $g_j(\cdot)_{j=1,\dots,J}$  represents the reliability constraints for the distribution system as a function of load injections. Specifically, these general functions represent distribution system parameters (i.e., the nodal voltage  $u_\tau(t)$  and power flow through distribution lines  $f_\tau(t)$ ) that should obey the following constraints:

$$u_\tau(t) \geq u^{min}, \quad \forall t, \tau, \quad (3)$$

$$u_\tau(t) \leq u^{max}, \quad \forall t, \tau, \quad (4)$$

$$f_\tau(t) \leq S^{max}, \quad \forall t, \tau, \quad (5)$$

where  $u^{min}$ ,  $u^{max}$ , and  $S^{max}$  correspond to the lower voltage limit, upper voltage limit, and power flow limit, respectively, for the population’s connection to the distribution grid. The power flow model we use to derive  $u_\tau(t)$  and  $f_\tau(t)$  from the load profile  $\mathbf{D}_\tau(\mathbf{p}_\tau)$  is given in Section IV-B.

However, the aggregator cannot simply solve (1). As explained in the introduction, knowledge of customers’ price response is unavailable to the aggregator. Recall, 1) the aggregator does not want to directly query customers for their response function, 2) most customers cannot readily characterize their response, and 3) customers might not be willing to share this private information. Accordingly, the aggregator needs a method to sequentially choose daily prices to simultaneously 1) control the daily incurred cost; 2) learn the customers’ price response models; and 3) ensure the distribution system constraints are not violated at any time.

### B. Load Flexibility Model

It is hard to approach the problem of learning the response of a population of customers to complex dispatch signals such as RTP as a complete “black box problem”, i.e., by just observing the broadcasted price and the load response. There are many reasons for this, including 1) the existence of random or exogenous parameters which lead to variability in the temporal and geographical behavior of electricity demand; 2) the variability of the control objective on a daily basis (e.g., due to randomness in renewable generation outputs, market conditions, or baseload); and 3) the small size of the set of observations that one can gather compared to the high dimensional structure of the load (there are only 365 days in a year, so only 365 set of prices can be posted). Hence, in this paper, we will be *exploiting the known physical structure* of the problem and making use of our statistical prior knowledge of how the load behaves to lower the problem dimensionality.

Specifically, to lower the dimensionality for the learning problem, we explore the fact that flexible loads only show limited number of “load signatures” (potentially due to the automated nature of load response through home energy management systems and the limited types of flexible appliances).

\*The reader should not that this choice is not fundamental to the development of the modified learning heuristics we present in this paper. It only allows us to provide a concrete characterization of the response to dispatch signals by mathematically modeling the customers as cost-minimizing agents equipped with home energy management systems in Section II-C.

Let us assume that electric appliances can belong to a finite number of clusters  $c \in \mathcal{C}$ . For each cluster  $c$ , we denote the set of feasible daily power consumption schedules that satisfy the energy requirements of the corresponding appliances by  $\mathcal{D}_c$ . Any power consumption schedule,  $[d_c(t)]_{t=1,\dots,T} = \mathbf{D}_c \in \mathcal{D}_c$ , would satisfy the daily power needs of an appliance in cluster  $c$ . To give an example, consider a cluster that represents plug-in electric vehicles (EVs) that require  $E_c$  kWh in the time interval  $[t_1, t_2]$  with a maximum charging rate of  $\rho$  kW. Accordingly, the set  $\mathcal{D}_c$  of daily feasible power consumption schedules is given by:

$$\mathcal{D}_c = \{\mathbf{D}_c \mid \sum_{t=t_1}^{t_2} d_c(t) = E_c; 0 \leq d_c(t) \leq \rho\}. \quad (6)$$

For discussion on characterizing the sets  $\mathcal{D}_c$  for other flexible appliances, including interruptible, non-interruptible, and thermostatically controlled loads, we refer the reader to [16].

### C. Price Response Model

In this section, we discuss how the total population's load responds to prices given the fact that flexible appliances belong to a finite number of clusters  $c \in \mathcal{C}$ . The price signals affects the power consumption in two ways:

1) *Automated per cluster response*: Within each load cluster  $c$  (i.e., given prespecified preferences such as EV charging deadlines or AC temperature set points), we assume that the customer chooses the power consumption profile  $\mathbf{D}_c \in \mathcal{D}_c$  that minimizes their electricity cost dependent on the daily broadcasted price  $\mathbf{p}_\tau$ . For appliances in cluster  $c$  on day  $\tau$ , we assume all will choose the same minimum cost power consumption profile:

$$\tilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau) = \arg \min_{\mathbf{D}_c \in \mathcal{D}_c} \sum_{t=1}^T p(t) d_c(t). \quad (7)$$

Due to the automated nature of home energy management systems, each cluster selecting its cost minimizing profile is a reasonable assumption once the customers have defined their flexibility preferences, e.g., the desired charge amounts and deadlines for EVs [17], [18].

2) *Preference Adjustment*: We also consider the fact that customers may respond to price signals by adjusting their preferences. For example, some customers are willing to pay more to keep their AC temperature set points lower, or charge their EV less. This means that the number of appliances in each cluster, denoted by  $a_c$ , also depends on the daily posted price vector  $\mathbf{p}_\tau$ .

Combining the *automated per cluster response* and *preference adjustment*, we can define the population's load on day  $\tau$  in response to the posted price  $\mathbf{p}_\tau$  as follows:

$$\mathbf{D}_\tau^*(\mathbf{p}_\tau) = \sum_{c \in \mathcal{C}} a_c(\mathbf{p}_\tau) \tilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau). \quad (8)$$

As stated before, if the aggregator has full knowledge of the customers' price responses, which reduces to having full knowledge of the preference adjustments  $a_c(\mathbf{p}_\tau)$ , then the aggregator can pick the daily price vector  $\mathbf{p}_\tau^*$  in order to shape the population's power consumption according to (1).

However, the functions  $a_c(\mathbf{p}_\tau)$  are unknown to the aggregator and also exhibit inherent stochasticities due to variations of daily customer needs. As such, we will model the  $a_c(\mathbf{p}_\tau)$ 's as random variables with parameterized distributions,  $\phi_c$ , based on the posted price signal  $\mathbf{p}_\tau$  and an unknown but constant parameter vector  $\theta^*$ . Here,  $\theta^*$  represents the *true model* for the customers' sensitivity to the price signals. This allows for the complex response of the customer population to be represented through a single unknown vector, thus reducing the dimensionality of the learning problem. With this in mind, we would like to highlight three important properties of the price response model we adopt:

- 1) The preference adjustment models  $a_c(\mathbf{p}_\tau)$  are stochastic and their distributions  $\phi_c$  are parameterized by  $\mathbf{p}_\tau$  and  $\theta^*$ . This is due to exogenous factors outside of the aggregator's scope that influence customers' power consumption profiles resulting in a level of stochasticity in the responses to prices (i.e., customers will not respond to prices in the same fashion each day).
- 2) The probability distributions of  $a_c(\mathbf{p}_\tau)$  (i.e.,  $\phi_c$ ) are unknown to the aggregator, i.e., the aggregator does not know the true parameter  $\theta^*$  of the stochastic model.
- 3) The realizations of  $a_c(\mathbf{p}_\tau)$  are not directly observable by the aggregator. The aggregator can only monitor the population's total consumption profile  $\mathbf{D}_\tau$  and cannot observe the decomposed response of each cluster  $a_c(\mathbf{p}_\tau) \tilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau)$  independently.

Because we have introduced stochasticity to customers' price response models, we appropriately alter the aggregator's optimization problem for selecting the price signal on day  $\tau$  to account for the distributions  $\phi_c$ :

$$\mathbf{p}_\tau^* = \arg \min_{\mathbf{p}_\tau \in \mathcal{P}} \mathbb{E}_{\{\phi_c\}_{c \in \mathcal{C}}} [f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau)] \quad (9)$$

$$\text{s.t. } \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [g_j(\mathbf{D}_\tau(\mathbf{p}_\tau)) \leq 0] \geq 1 - \mu, \quad \forall j \quad (10)$$

where the distribution system constraints have to be satisfied with a given probability  $1 - \mu$ . In (9), the aggregator now considers minimizing an expected cost and is subject to probabilistic reliability constraints in (10) that depend on the distributions  $\phi_c$  of the preference adjustment models  $a_c(\mathbf{p}_\tau)$ .

Clearly, the aggregator needs to learn the underlying parameters of the stochastic models  $\phi_c$  of how customers respond to price signals in order to select price signals for load shaping initiatives (i.e., the aggregator needs to learn  $\theta^*$ ). Our proposed learning approach and pricing strategy for an electricity aggregator is detailed in the next section.

## III. REAL-TIME PRICING VIA MULTI-ARMED BANDIT

### A. Multi-Armed Bandit Overview

We utilize the multi-armed bandit (MAB) framework to model the iterative decision making procedure of an aggregator implementing a daily load shaping program [19]–[21]. Moreover, the MAB framework exemplifies the *exploration-exploitation* trade-off dilemma faced by an aggregator each day in the electricity pricing problem. Namely, should the

aggregator choose to broadcast untested prices (i.e., *explore*) to learn more information about the customers? Or should the aggregator choose to broadcast well-performing prices (i.e., *exploit*) to manipulate the daily electricity demand?

To evaluate the performance of an algorithm that aims to tackle the exploration-exploitation trade-off, one commonly examines the algorithm's *regret*. Formally, regret is the cumulative difference in cost incurred over  $\mathcal{T}$  days between a clairvoyant algorithm (i.e., the optimal strategy that is aware of the customers' price responses) and any proposed algorithm that does not know the customers' price responses:

$$R_{\mathcal{T}} = \sum_{\tau=1}^{\mathcal{T}} \left\{ \mathbb{E}_{\{\phi_c\}_{c \in \mathcal{C}}} [f(\mathbf{D}_{\tau}(\mathbf{p}_{\tau}), \mathbf{V}_{\tau})] - \mathbb{E}_{\{\phi_c\}_{c \in \mathcal{C}}} [f(\mathbf{D}_{\tau}(\mathbf{p}^*), \mathbf{V}_{\tau})] \right\}. \quad (11)$$

Instead of the above more standard definition, an alternative metric for regret that is easier to bound for more complex bandits is to count the number of times that suboptimal price signals are selected over the  $\mathcal{T}$  days. For this, we introduce the following notation: let  $\mathbf{p}^{\mathbf{V}_{\tau}, \boldsymbol{\theta}^*}$  denote the optimal price signal for the true model of the population's price response  $\boldsymbol{\theta}^*$  when the daily exogenous parameter  $\mathbf{V}_{\tau}$  is observed on day  $\tau$ . Any price signal  $\mathbf{p}_{\tau} \neq \mathbf{p}^{\mathbf{V}_{\tau}, \boldsymbol{\theta}^*}$  is considered a suboptimal price. Moreover, we denote  $N_{\tau}(\mathbf{p}, \mathbf{V})$  as the number of times up to day  $\tau$  that the algorithm simultaneously observes the exogenous parameter  $\mathbf{V}$  and selects the price signal  $\mathbf{p}$ . As such, the total number of times that suboptimal price signals are selected over  $\mathcal{T}$  days is:

$$\sum_{\mathbf{V} \in \mathcal{V}} \sum_{\mathbf{p} \in \{\mathcal{P} \setminus \mathbf{p}^{\mathbf{V}, \boldsymbol{\theta}^*}\}} N_{\mathcal{T}}(\mathbf{p}, \mathbf{V}) = \sum_{\tau=1}^{\mathcal{T}} \mathbb{1}_{\{\mathbf{p}_{\tau} \neq \mathbf{p}^{\mathbf{V}_{\tau}, \boldsymbol{\theta}^*}\}}, \quad (12)$$

where  $\mathbb{1}_{\{\cdot\}}$  is the indicator function that is set equal to one if the criteria is met and zero otherwise. Subsequently, in an iterative decision making problem such as this, the question arises: *how can an aggregator learn to price electricity with bounded regret, and what are the regret bounds we can provide for a proposed algorithm given dynamically changing grid conditions and reliability constraints?* In the following sections, we present a modified Thompson sampling heuristic for the electricity pricing problem to simultaneously learn the true model  $\boldsymbol{\theta}^*$  for the population, select the daily price signals, ensure grid reliability, and provide a regret guarantee.

### B. Thompson Sampling

Thompson sampling is a well-known MAB heuristic for choosing actions in an iterative decision making problem with the *exploration-exploitation* dilemma [22]–[24]. In summary, the integral characteristic of Thompson sampling is that the algorithm's knowledge on day  $\tau$  of the unknown parameter  $\boldsymbol{\theta}^*$  is represented by the prior distribution  $\pi_{\tau-1}$ . Each day the algorithm samples  $\tilde{\boldsymbol{\theta}}_{\tau}$  from the prior distribution, and selects a price assuming that the sampled parameter is the true parameter. The algorithm then makes an observation dependent on the chosen price and the hidden parameter and performs a Bayesian update on the parameter's distribution

$\pi_{\tau}$  based on the new observation. Because TS samples parameters from the prior distribution, the algorithm has a chance to explore (i.e., draw new parameters) and can exploit (i.e., draw parameters that are likely to be the true parameter) throughout the run of the algorithm.

### C. Constrained Thompson Sampling

In this section, we present the MAB heuristic titled Con-TS-RTP adopted to the electricity pricing problem. Con-TS-RTP is a modified Thompson sampling algorithm where the daily optimization problem is subject to constraints (standard TS algorithms do not have constraints in the daily optimization) [25]. Each day, the algorithm observes the daily target profile  $\mathbf{V}_{\tau}$ , draws a parameter  $\tilde{\boldsymbol{\theta}}_{\tau}$  from the prior distribution, broadcasts a price signal to the customers, observes the load profile of the population in response to the broadcasted price, and then performs a Bayesian update on the parameter's distribution  $\pi_{\tau}$  based on the new observation.

The observation on day  $\tau$  is denoted as  $\mathbf{Y}_{\tau} = \mathbf{D}_{\tau}^*(\mathbf{p}_{\tau})$  and we assume that each  $\mathbf{Y}_{\tau}$  comes from the observation space  $\mathcal{Y}$  that is known a priori. When performing the Bayesian update, the algorithm makes use of the following *likelihood function*:  $\ell(\mathbf{Y}_{\tau}; \mathbf{p}, \boldsymbol{\theta}) = \mathbb{P}_{\boldsymbol{\theta}}(\mathbf{D}_{\tau}^*(\mathbf{p}_{\tau}) = \mathbf{Y}_{\tau} | \mathbf{p}_{\tau} = \mathbf{p})$ . This function calculates the likelihood of observing a specific load profile when broadcasting price  $\mathbf{p}$  and the true parameter is  $\boldsymbol{\theta}$ . The pseudocode for Con-TS-RTP applied to the constrained electricity pricing problem is presented in Algorithm 1. The reader will notice that the optimization in line 4 is presented under two different sets of constraints: *Constraint Set A*: The operational constraints of the distribution system are formulated with respect to the drawn parameter  $\tilde{\boldsymbol{\theta}}_{\tau}$  (i.e., the constraints are not necessarily enforced for  $\boldsymbol{\theta}^*$ ); *Constraint Set B*: The operational constraints of the distribution system are formulated with respect to the current prior distribution  $\pi_{\tau-1}$  on the unknown parameter  $\boldsymbol{\theta}^*$  (as chance constraints). We will see that under certain assumptions, this means that the constraints are upheld for the true unknown parameter  $\boldsymbol{\theta}^*$  at each round (with high probability).

### D. Discussion on Regret Performance of Con-TS-RTP

The regret analysis of Con-TS-RTP is inspired by the results in [26] for TS with nonlinear cost functions. The analysis in the aforementioned paper provides bounds on the total number of times that suboptimal price signals selected by the algorithm over  $\mathcal{T}$  days as specified in equation (12). The regret guarantee we provide in this work extends the result further, allowing for constraints in the daily optimization that are dependent on the sampled  $\tilde{\boldsymbol{\theta}}_{\tau}$  and on the exogenous target profiles  $\mathbf{V}_{\tau}$ . As such, our regret guarantee applies to the Con-TS-RTP algorithm with constraints as formulated in

---

**Algorithm 1** CON-TS-RTP

---

**Input:** Parameter set  $\Theta$ ; Price set  $\mathcal{P}$ ; Observation set  $\mathcal{Y}$ ; Voltage constraints  $u^{min}, u^{max}$ ; Power flow constraint  $S^{max}$ , Reliability metrics  $\mu, \nu$

**Initialize**  $\pi_0$ .

- 1: **for** Day index  $\tau = 1 \dots \mathcal{T}$  **do**
- 2:   Sample  $\tilde{\theta}_\tau$  from distribution  $\pi_{\tau-1}$ .
- 3:   Observe the daily exogenous parameter  $\mathbf{V}_\tau$ .
- 4:   Broadcast the daily price signal:

$$\hat{\mathbf{p}}_\tau = \arg \min_{\mathbf{p}} \mathbb{E}_{\{\phi_c\}_{c \in \mathcal{C}}} [f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau) | \boldsymbol{\theta} = \tilde{\theta}_\tau]$$

Subject to:

**Constraint Set A:**

$$\begin{cases} \mathbf{A.1:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \geq u^{min} | \boldsymbol{\theta} = \tilde{\theta}_\tau] \geq 1 - \mu, \quad \forall t \\ \mathbf{A.2:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \leq u^{max} | \boldsymbol{\theta} = \tilde{\theta}_\tau] \geq 1 - \mu, \quad \forall t \\ \mathbf{A.3:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [f_\tau(t) \leq S^{max} | \boldsymbol{\theta} = \tilde{\theta}_\tau] \geq 1 - \mu, \quad \forall t \end{cases}$$

**Constraint Set B:**

$$\begin{cases} \mathbf{B.1:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \geq u^{min} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \\ \mathbf{B.2:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \leq u^{max} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \\ \mathbf{B.3:} \mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [f_\tau(t) \leq S^{max} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \end{cases}$$

- 5:   Observe  $\mathbf{Y}_\tau = \mathbf{D}_\tau^*(\hat{\mathbf{p}}_\tau)$ .
- 6:   Posterior update:

$$\forall S \subseteq \Theta : \pi_\tau(S) = \frac{\int_S \ell(\mathbf{Y}_\tau; \hat{\mathbf{p}}_\tau, \boldsymbol{\theta}) \pi_{\tau-1}(d\boldsymbol{\theta})}{\int_\Theta \ell(\mathbf{Y}_\tau; \hat{\mathbf{p}}_\tau, \boldsymbol{\theta}) \pi_{\tau-1}(d\boldsymbol{\theta})}$$

- 7: **end for**
- 

Constraint Set A<sup>†</sup> in Algorithm 1.

**Assumption 1.** (Finitely many price signals, observations, exogenous vectors).  $|\mathcal{P}|, |\mathcal{Y}|, |\mathcal{V}| < \infty$ .

**Assumption 2.** (Finite Prior, “Grain of truth”) The prior distribution  $\pi$  is supported over finitely many particles:  $|\Theta| < \infty$ . The true parameter exists within the parameter space:  $\boldsymbol{\theta}^* \in \Theta$ . The initial distribution  $\pi_0$  has non-zero mass on the true parameter  $\boldsymbol{\theta}^*$  (i.e.,  $\mathbb{P}_{\pi_0}[\boldsymbol{\theta}^*] > 0$ ).

**Assumption 3.** The exogenous vectors  $\mathbf{V}$  are i.i.d. drawn from a distribution defined on a finite sample space  $\mathcal{V}$ , with each outcome drawn with a nonzero probability.

**Assumption 4.** (Unique optimal price signal). There is a unique optimal price signal  $\mathbf{p}^{\mathbf{V},*}$  for each exogenous parameter  $\mathbf{V} \in \mathcal{V}$ .

<sup>†</sup>We note that the regret result in Theorem 1 only applies to the Con-TS-RTP algorithm with the daily optimization (line 4) subject to *Constraint Set A*. This is due to *Constraint Set B* (which is dependent on  $\pi_{\tau-1}$ ) potentially prohibiting the aggregator from selecting the optimal price signals throughout the learning process due to constraint violations from parameters  $\boldsymbol{\theta} \neq \boldsymbol{\theta}^*$ . However, *Constraint Set B* provides reliability guarantees that *Constraint Set A* cannot provide (see Section III-E). A thorough analysis on the effect of stage-wise reliability-constraints such as *Constraint Set B* on the growth of regret can be found in [27] but only for the case of stochastic MABs with linear costs and linear constraints.

**Theorem 1.** Under assumptions 1-4 and Constraint Set A in Algorithm 1, for  $\delta, \epsilon \in (0, 1)$ , there exists  $\mathcal{T}^* \geq 0$  such that for all  $\mathcal{T} \geq \mathcal{T}^*$ , with probability  $1 - \delta$ :

$$\sum_{\mathbf{V} \in \mathcal{V}} \sum_{\mathbf{p} \in \{\mathcal{P} \setminus \mathbf{p}^{\mathbf{V},*}\}} N_{\mathcal{T}}(\mathbf{p}, \mathbf{V}) \leq B + C(\log \mathcal{T}), \quad (13)$$

where  $B \equiv B(\delta, \epsilon, \mathcal{P}, \mathcal{Y}, \Theta)$  is a problem-dependent constant that does not depend on  $\mathcal{T}$ , and  $C(\log \mathcal{T})$  depends on  $\mathcal{T}$ , the sequence of selected price signals, and the Kullback-Leibler divergence properties of the bandit problem (i.e., the marginal Kullback-Leibler divergences of the observation distributions  $KL[\ell(\mathbf{Y}; \mathbf{p}, \boldsymbol{\theta}^*), \ell(\mathbf{Y}; \mathbf{p}, \boldsymbol{\theta})]$  (The complete description of the  $C(\log \mathcal{T})$  term is left to the appendix).

*Proof.* The proof is in the appendix.

In the next section, we discuss the distribution system reliability issues that could arise from *Constraint Set A* and a modification to the Con-TS-RTP algorithm to ensure the constraints are enforced on all days (i.e., *Constraint Set B*).

### E. Con-TS-RTP with Improved Reliability Constraints

In order for the aggregator to ensure safe operation of the distribution grid while running the Con-TS-RTP algorithm, the reliability constraints need to hold for the true price response model  $\boldsymbol{\theta}^*$  each day. However, with the constraints formulated as in Algorithm 1’s *Constraint Set A*, the distribution system constraints are only enforced for the sampled  $\tilde{\theta}_\tau$  and not necessarily the true parameter  $\boldsymbol{\theta}^*$ . This entails that the distributions  $\{\phi_c\}_{c \in \mathcal{C}}$  are parameterized by the sampled  $\tilde{\theta}_\tau$ ; therefore, they are inaccurate if any parameter  $\tilde{\theta}_\tau \neq \boldsymbol{\theta}^*$  is sampled. This could potentially lead to many constraint violations throughout the run of the algorithm resulting in inadequate service for the customers and grid failures.

Due to the importance of reliable operation of the distribution system, we present a modification to the Con-TS-RTP algorithm (i.e., replacing *Constraint Set A* with *Constraint Set B* in Algorithm 1) to increase the reliability of the selected prices and resulting load profiles with respect to the grid constraints. Specifically, we propose alternate constraints that depend on the algorithm’s current knowledge of the true parameter, instead of the sampled parameter. In other words, instead of depending on  $\tilde{\theta}_\tau$ , the proposed alternate constraints depend on the prior distributions  $\pi_{\tau-1}$  as follows:

$$\mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \geq u^{min} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \quad (14)$$

$$\mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [u_\tau(t) \leq u^{max} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \quad (15)$$

$$\mathbb{P}_{\{\phi_c\}_{c \in \mathcal{C}}} [f_\tau(t) \leq S^{max} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1 - \nu, \quad \forall t \quad (16)$$

where  $\nu$  is a small constant (detailed in Proposition 1). When considering constraints (14)-(16) in Con-TS-RTP, the algorithm will select more conservative price signals each day that can guarantee the distribution system’s constraints are met with high probability by using the information in the updated prior distributions. Before analyzing the modified algorithm’s reliability, we make the following assumption:




Fig. 1. Radial distribution system.

**Assumption 5.** *There exists  $\xi^* > 0$ ,  $\lambda \geq 0$ ,  $\delta \in (0, 1)$  such that for all  $\theta \neq \theta^*$ ,  $KL[\ell(\mathbf{Y}; \mathbf{p}, \theta^*), \ell(\mathbf{Y}; \mathbf{p}, \theta)] \geq \xi^*$ , where*

$$\xi_{\theta, \mathbf{p}}^* = \max_{x \in \mathbb{Z}^{>0}} \left\{ \frac{-\lambda}{x} + \frac{4}{\sqrt{x}} \sqrt{\frac{\log |\mathcal{Y}| |\mathcal{P}|}{\delta} + \frac{\log x}{2}} \right. \\ \left. \times \sum_{\mathbf{Y} \in \mathcal{Y}} \left| \log \frac{\ell(\mathbf{Y}; \mathbf{p}, \theta^*)}{\ell(\mathbf{Y}; \mathbf{p}, \theta)} \right| \right\}$$

and

$$\xi^* = \max_{\theta \in \Theta, \mathbf{p} \in \mathcal{P}} \xi_{\theta, \mathbf{p}}^*.$$

Assumption 5 ensures that if the aggregator observes  $\mathbf{Y}_\tau$  on day  $\tau$ , the algorithm's Bayesian updates of the prior distribution  $\pi_\tau$  will never decrease the mass of the true parameter  $\theta^*$  below a certain threshold. Specifically, with Assumption 5, it can be shown (as in [26]) that with probability  $1 - \delta\sqrt{2}$  the following holds for all  $\tau \geq 1$ :

$$\pi_\tau(\theta^*) \geq \pi_0(\theta^*) e^{-\lambda|\mathcal{P}|}, \quad (17)$$

where  $\lambda \geq 0$  is a chosen parameter (from Assumption 5) that dictates the minimum reachable mass of the true parameter via Bayesian updates. With the modified constraints (14)-(16) and the minimum mass of the true parameter (17), the reliability of Con-TS-RTP can be characterized as follows:

**Proposition 1.** *Under assumptions 1-5, with  $\nu$  in equations (14)-(16) chosen such that  $\nu = \mu\pi_0(\theta^*)e^{-\lambda|\mathcal{P}|}$ , with probability  $1 - \delta\sqrt{2}$ , the Con-TS-RTP algorithm with Constraint Set B will uphold the probabilistic distribution system constraints as formulated in (10) for each day  $\tau = 1, \dots, \mathcal{T}$ .*

*Proof.* The proof is in the appendix.

## IV. EXPERIMENTAL EVALUATION

### A. Test Setup: Radial Distribution System

In this section we describe the power distribution system and the corresponding network parameters for the test case. We consider an actual radial distribution system from the ComEd service territory in Illinois, USA (adopted from [28] and shown in Fig. 1) represented by the undirected graph  $\mathcal{G}$  which includes a set of nodes (vertices)  $\mathcal{N}$  and a set of power lines (edges)  $\mathcal{L}$ . In this work, we consider each node as one population with its own daily load profile; however, each node could be an aggregation of smaller entities downstream

Line	R ( $10^{-3}\Omega$ )	X ( $10^{-3}\Omega$ )	$S^{max}$ (KVA)	Line	R ( $10^{-3}\Omega$ )	X ( $10^{-3}\Omega$ )	$S^{max}$ (KVA)
1	24.2	48.2	54	20	129.5	30.9	10.8
2	227.3	743.5	84	21	15.1	5.4	14.4
3	76.3	18.2	10.8	22	50.8	12.1	10.8
4	43.6	142.7	84	23	69.1	16.5	10.8
5	25.8	84.4	84	24	31.6	11.2	14.4
6	10.5	10.7	40.2	25	96.3	23	10.8
7	23.2	23.6	40.2	26	110.7	112.6	40.2
8	75.1	26.7	14.4	27	2.1	0.7	14.4
9	114.4	27.3	10.8	28	242.1	86.2	14.4
10	110.8.3	67.7	14.4	29	27.3	27.8	40.2
11	63.7	22.7	14.4	30	174.6	62.1	16.2
12	278.7	99.2	14.4	31	43	15.3	10.8
13	254.2	10.8.5	14.4	32	207.8	74	10.8
14	21.8	5.2	10.8	33	109.4	38.9	14.4
15	57.3	20.4	14.4	34	50.5	18	14.4
16	126.7	45.1	14.4	35	165.2	58.8	14.4
17	48.6	11.6	10.8	36	49.5	17.6	14.4
18	95.1	22.7	10.8	37	5.8	2.1	14.4
19	137.3	32.8	10.8				

TABLE I  
DISTRIBUTION SYSTEM PARAMETERS.

of the local distribution connection point. The undirected graph is organized as a tree, with the root node representing the distribution system's substation where it is connected to the regional transmission system. We denote  $N$  as the total number of nodes in the network excluding the root node. The nodes are indexed as  $i = 0, \dots, N$ , and the node corresponding to  $i = 0$  (i.e., the root node) is the substation. The power lines are indexed by  $i = 1, \dots, N$  where the  $i$ -th line is directly upstream of node  $i$  (i.e., line  $i$  feeds directly to node  $i$ ). In the following, we denote the parent vertex of node  $i$  as  $\mathcal{A}_i$  and the set of children vertices of node  $i$  as  $\mathcal{K}_i$ .

Furthermore, we assume the aggregator has access to measurement data at each node's local connection point. Specifically, the aggregator measures the active and reactive power demands at each node  $i$  at time  $t$  on day  $\tau$  denoted as  $d_{i,\tau}^P(t)$  and  $d_{i,\tau}^Q(t)$ , respectively. In order to ensure the delivered power is suitable for the electricity customers, the aggregator also monitors node  $i$ 's local voltage at time  $t$  on day  $\tau$  denoted as  $v_{i,\tau}(t)$ . In the following, we denote the active power daily load profile of node  $i$  on day  $\tau$  as  $\mathbf{D}_{i,\tau}^P = [d_{i,\tau}^P(t)]_{t=1, \dots, T}$ . Additionally, the aggregator records the active and reactive power flows,  $f_{i,\tau}^P(t)$  and  $f_{i,\tau}^Q(t)$ , respectively, on each line  $i \in \mathcal{L}$ . Each line in the distribution system has its own internal resistance denoted as  $R_i$ , reactance denoted as  $X_i$ , and apparent power limit denoted as  $S_i^{max}$ . The parameters for the distribution system are listed in Table I.

### B. Power Flow Model

In order to solve for the power flow and nodal voltages of the power distribution system, we make use of the *LinDistFlow* model [29], which is a linear approximation for the AC power flow model. The *LinDistFlow* model has been




Fig. 2. Evolution of prior distribution  $\pi_\tau$  for node 10. From left to right: Day 1 (initialized to uniform distribution, i.e., no knowledge of the true parameter), Day 15 prior, Day 30 prior, Day 90 prior, and Day 180 prior.

extensively studied and verified to be competitive to the nonlinear AC flow model on many realistic feeder topologies including radial [30]–[33]. The *LinDistFlow* model reduces computational complexity by making use of the following linear power flow and voltage equations:

$$d_{i,\tau}^P(t) - g_{i,\tau}^P(t) + \sum_{j \in \mathcal{K}_i} f_{j,\tau}^P(t) = f_{\mathcal{A}_i,\tau}^P(t); \quad \forall t, \tau, i, \quad (18)$$

$$d_{i,\tau}^Q(t) - g_{i,\tau}^Q(t) + \sum_{j \in \mathcal{K}_i} f_{j,\tau}^Q(t) = f_{\mathcal{A}_i,\tau}^Q(t); \quad \forall t, \tau, i, \quad (19)$$

$$u_{\mathcal{A}_i,\tau}(t) - 2(f_{i,\tau}^P(t)R_i + f_{i,\tau}^Q(t)X_i) = u_{i,\tau}(t); \quad \forall t, \tau, i. \quad (20)$$

In (20) we make use of the operator  $u_{i,\tau}(t) = (v_{i,\tau}(t))^2$  to provide a linear voltage drop relationship across the distribution system. For the scope of this work, we assume that the substation connection to the regional transmission system (node  $i = 0$ ) is regulated and has a fixed voltage  $v_{0,\tau}(t) = 120\text{V}, \forall t, \tau$ .

### C. Distribution System Operational Constraints

The nodal voltages and line flows calculated in (18)–(20) should obey the following constraints for reliable operation:

$$u_{i,\tau}(t) \geq u_i^{\min}, \quad \forall t, \tau, i \in \mathcal{N}, \quad (21)$$

$$u_{i,\tau}(t) \leq u_i^{\max}, \quad \forall t, \tau, i \in \mathcal{N}, \quad (22)$$

$$(f_{i,\tau}^P(t))^2 + (f_{i,\tau}^Q(t))^2 \leq (S_i^{\max})^2, \quad \forall t, \tau, i \in \mathcal{L}, \quad (23)$$

where (21)–(22) are the nodal voltage constraints and (23) is the apparent power constraint for each distribution line.

### D. Load Model and Multi-armed Bandit Formulation

In this test case, the goal of the aggregator is to integrate varying levels of intermittent solar generation into the distribution system (i.e., the aggregator wants the customers to take advantage of the available renewable energy and consume all the solar generation each day). To model this, we consider 10 unique target load profile vectors, with the daily target profile  $\mathbf{V}_{i,\tau}$  for node  $i$  for day  $\tau$  drawn from a uniform distribution each morning. Each of the 10 target load profile vectors corresponds to the forecasted solar generation in each time slot at node  $i$ . In this setup, we consider 6 time slots each day, each 4 hours long and the aggregator transmits daily price signals  $\mathbf{p}_{i,\tau}$  to each node within the system. The

aggregator has a high and low price for each of the 6 time slots resulting in  $2^6$  possible daily price signals. Each node has a cost function that is dependent on the node’s demand as well as the target profile. In this test case, we assume the cost function is the squared deviation of the node’s demand from the target profile, thus equally penalizing over-usage and under-usage.

We consider 20 unique load flexibility clusters in this test case. Each node in the distribution system is comprised of these 20 load clusters with its own unique sensitivities  $a_{i,c}(\mathbf{p}_\tau)$  for each cluster. Each sensitivity parameter is selected as  $a_{i,c}(\mathbf{p}_\tau) \sim \mathcal{N}(\frac{\beta_c}{\theta_i^* \mathbf{p}_{i,\tau}}, \sigma^2)$  each day where  $\beta_c$  is a cluster specific constant known by the aggregator. Each node’s price sensitivity, i.e., parameter to be learned,  $\theta_i^*$ , is a vector of length 6 and the set of possible parameters,  $\Theta$ , contains 10 unique vectors.

### E. Results

We simulated the Con-TS-RTP algorithm for 365 days for an aggregator attempting to learn the sensitivities of the nodes in the system and shape their demands. In the following, we highlight the results of the simulation at node 10 of the radial distribution system. Figure 2 presents the evolution of the prior distribution for node 10’s hidden parameter.

Figure 3 presents the regret performance (both the cumulative regret and number of suboptimal price signals selected) of Con-TS-RTP at node 10. As seen in Figure 3, the regret curve flattens after day 130 as the algorithm never chooses a suboptimal price signal after this day.

Figure 4 presents node 10’s deviation from a specific daily target profile. On days 2, 3, 4, 53, and 365 the same target profile (i.e.,  $\mathbf{V}_2 = \mathbf{V}_3 = \mathbf{V}_4 = \mathbf{V}_{53} = \mathbf{V}_{365}$ ) was drawn and the aggregator selected different price signals to shape the node’s demand. As seen in Fig. 4, the deviation from the target profile on day 365 is less than the deviation on the other days as the algorithm has learned the true parameter and selects the optimal price signal to shape the load.

In Figure 5, we present the distribution system constraint violations that were avoided by using Con-TS-RTP instead of an unconstrained TS algorithm. Clearly, in the early learning stages, the unconstrained TS algorithm does not have

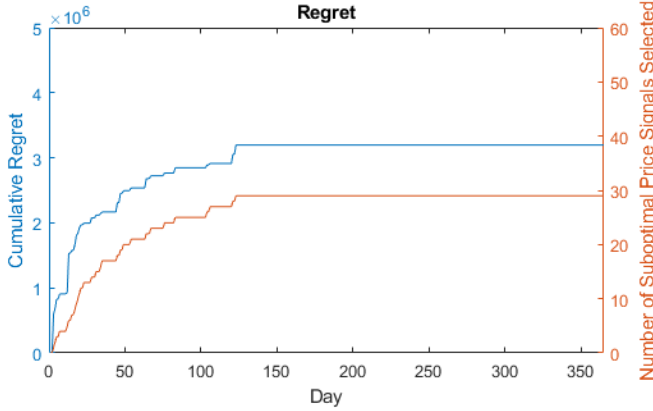


Fig. 3. Regret performance of Con-TS-RTP at node 10.




Fig. 4. Deviation of node 10's demand from a specific daily target profile.

accurate knowledge of the hidden parameters and violates the distribution system constraints often. Con-TS-RTP is more conservative with its exploration of untested price signals and avoids the constraint violations made by the unconstrained TS algorithm. The simulation was implemented with Matlab/CVX, an i7 processor, and 16gb RAM in  $< 5$  minutes.

## V. CONCLUSION

We presented a multi-armed bandit problem formulation for an electricity aggregator implementing a real-time pricing program for load shaping (e.g., reduce demand at peak hours, integrate more intermittent renewables, track a desired daily load profile, etc). We made use of a constrained Thompson sampling heuristic, Con-TS-RTP, as a solution to the *exploration/exploitation* problem of an aggregator *passively* learning customers' price sensitivities while broadcasting price signals that influence customers to alter their demand. The Con-TS-RTP algorithm permits day-varying target load profiles and takes into account the actual operational constraints of a distribution system to ensure that the customers receive adequate service and to avoid potential grid failures.

We discussed a regret guarantee for the proposed Con-TS-RTP algorithm which bounds the total number of suboptimal price signals broadcasted by the aggregator. Furthermore, we




Fig. 5. Top: Distribution system constraint violations at node 10 avoided by using Con-TS-RTP instead of an unconstrained TS. Bottom: Distribution system constraint violations across the entire system avoided by using Con-TS-RTP.

discussed an operational reliability guarantee that ensures the power distribution system constraints are upheld with high probability throughout the run of the Con-TS-RTP algorithm.

## REFERENCES

- [1] C. Eid, E. Koliou, M. Valles, J. Reneses, and R. Hakvoort, "Time-based pricing and electricity demand response: Existing barriers and next steps," *Utilities Policy*, vol. 40, pp. 15 – 25, 2016.
- [2] V. Gomez, M. Chertkov, S. Backhaus, and H. J. Kappen, "Learning price-elasticity of smart consumers in power distribution systems," in *2012 SmartGridComm*. IEEE, 2012, pp. 647–652.
- [3] Z. Xu, T. Deng, Z. Hu, Y. Song, and J. Wang, "Data-driven pricing strategy for demand-side resource aggregators," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 57–66, 2016.
- [4] P. Li and B. Zhang, "Linear estimation of treatment effects in demand response: An experimental design approach," *arXiv preprint arXiv:1706.09835*, 2017.
- [5] P. Li, H. Wang, and B. Zhang, "A distributed online pricing strategy for demand response programs," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 350–360, 2017.
- [6] K. Khezeli, W. Lin, and E. Bitar, "Learning to buy (and sell) demand response," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6761–6767, 2017.
- [7] L. Jia, L. Tong, and Q. Zhao, "An online learning approach to dynamic pricing for demand response," *arXiv preprint arXiv:1404.1325*, 2014.
- [8] K. Khezeli and E. Bitar, "Risk-sensitive learning and pricing for demand response," *IEEE Transactions on Smart Grid*, vol. 9, no. 6, pp. 6000–6007, 2017.
- [9] L. Jia, Q. Zhao, and L. Tong, "Retail pricing for stochastic demand with unknown parameters: An online machine learning approach," in *2013 51st Allerton*. IEEE, 2013, pp. 1353–1358.
- [10] Y. Li, Q. Hu, and N. Li, "Learning and selecting the right customers for reliability: A multi-armed bandit approach," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 4869–4874.
- [11] D. Kalathil and R. Rajagopal, "Online learning for demand response," in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2015, pp. 218–222.
- [12] R. Mieth and Y. Dvorkin, "Data-driven distributionally robust optimal power flow for distribution systems," *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 363–368, 2018.
- [13] E. Dall'Anese, K. Baker, and T. Summers, "Chance-constrained ac optimal power flow for distribution systems with renewables," *IEEE Transactions on Power Systems*, vol. 32, no. 5, pp. 3427–3438, 2017.
- [14] R. Mieth and Y. Dvorkin, "Online learning for network constrained demand response pricing in distribution systems," *arXiv:1811.09384*, 2018.
- [15] A. Moradipari, C. Silva, and M. Alizadeh, "Learning to dynamically price electricity demand based on multi-armed bandits," in *2018 IEEE GlobalSIP*, Nov 2018, pp. 917–921.



- [16] M. Alizadeh, A. Scaglione, A. Applebaum, G. Kesidis, and K. Levitt, "Reduced-order load models for large populations of flexible appliances," *IEEE Transactions on Power Systems*, vol. 30, no. 4, 2015.
- [17] T.-H. Chang, M. Alizadeh, and A. Scaglione, "Coordinated home energy management for real-time power balancing," in *2012 IEEE Power and Energy Society General Meeting*. IEEE, 2012, pp. 1–8.
- [18] M. Alizadeh and A. Scaglione, "Least laxity first scheduling of thermostatically controlled loads for regulation services," in *2013 IEEE GlobalSIP*. IEEE, 2013, pp. 503–506.
- [19] A. Krishnamurthy, Z. S. Wu, and V. Syrgkanis, "Semiparametric contextual bandits," in *International Conference on Machine Learning*, 2018, pp. 2781–2790.
- [20] D. Foster, A. Agarwal, M. Dudik, H. Luo, and R. Schapire, "Practical contextual bandits with regression oracles," *Proceedings of Machine Learning Research*, vol. 80, 2018.
- [21] T. Xu, Y. Yu, J. Turner, and A. Regan, "Thompson sampling in dynamic systems for contextual bandit problems," *arXiv preprint arXiv:1310.5008*, 2013.
- [22] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, Z. Wen *et al.*, "A tutorial on thompson sampling," *Foundations and Trends® in Machine Learning*, vol. 11, no. 1, pp. 1–96, 2018.
- [23] D. Russo and B. Van Roy, "Learning to optimize via posterior sampling," *Mathematics of Operations Research*, vol. 39, no. 4, pp. 1221–1243, 2014.
- [24] S. Agrawal and N. Goyal, "Analysis of thompson sampling for the multi-armed bandit problem," in *Conference on Learning Theory*, 2012, pp. 39–1.
- [25] V. Saxena, J. Jaldén, J. E. Gonzalez, I. Stoica, and H. Tullberg, "Constrained thompson sampling for wireless link optimization," *arXiv preprint arXiv:1902.11102*, 2019.
- [26] A. Gopalan, S. Mannor, and Y. Mansour, "Thompson sampling for complex online problems," in *International Conference on Machine Learning*, 2014, pp. 100–108.
- [27] S. Amani, M. Alizadeh, and C. Thrampoulidis, "Linear stochastic bandits under safety constraints," *arXiv:1908.05814*, 2019.
- [28] P. Andrianesis, M. Caramanis, R. Masiello, R. Tabors, and S. Bahramirad, "Locational marginal value of distributed energy resources as non-wires alternatives," *IEEE Transactions on Smart Grid*, 2019.
- [29] M. E. Baran and F. F. Wu, "Optimal capacitor placement on radial distribution systems," *IEEE Transactions on Power Delivery*, vol. 4, no. 1, pp. 725–734, Jan 1989.
- [30] H. J. Liu, "Decentralized optimization approach for power distribution network and microgrid controls," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2017.
- [31] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3794–3803, 2016.
- [32] P. Šulc, S. Backhaus, and M. Chertkov, "Optimal distributed control of reactive power via the alternating direction method of multipliers," *IEEE Transactions on Energy Conversion*, vol. 29, no. 4, 2014.
- [33] M. Farivar, L. Chen, and S. Low, "Equilibrium and dynamics of local voltage control in distribution systems," in *52nd IEEE Conference on Decision and Control*. IEEE, 2013, pp. 4329–4334.

## APPENDIX

The following contains the supplementary material for the manuscript entitled: *Constrained Thompson Sampling for Real-Time Electricity Pricing with Grid Reliability Constraints* by N. Tucker, A. Moradipari, and M. Alizadeh.

### A. Discussion on Regret Performance

In this section, we describe the necessary background for Theorem 1 and then present the full version of the Theorem. In the following,  $\mathbf{p}^{\mathbf{V},*}$  denotes the optimal price signal for the true model of the population's price response  $\theta^*$  when the daily exogenous parameter  $\mathbf{V}_\tau$  is observed on day  $\tau$ . Any price signal  $\mathbf{p}_\tau \neq \mathbf{p}^{\mathbf{V},*}$  is considered a suboptimal price.

We now briefly explain how the posterior updates affect the regret performance. When price  $\mathbf{p}$  is posted on day  $\tau$ , the prior density is updated as

$$\pi_\tau(d\theta) \propto \exp\left(-\log \frac{l(\mathbf{Y}_\tau; \mathbf{p}, \theta^*)}{l(\mathbf{Y}_\tau; \mathbf{p}, \theta)}\right) \pi_{\tau-1}(d\theta). \quad (24)$$

Now, denote by  $KL(\theta_p^* || \theta_p)$  the marginal Kullback-Leibler divergence between the distribution  $\{l(Y; \mathbf{p}, \theta^*) : Y \in \mathcal{Y}\}$  and  $\{l(Y; \mathbf{p}, \theta) : Y \in \mathcal{Y}\}$ . As in [26], we can approximately write (24) as:

$$\pi_\tau(d\theta) \propto \exp\left(-\sum_{\mathbf{p} \in \mathcal{P}} N_\tau(\mathbf{p}) KL(\theta_p^* || \theta_p)\right) \pi_{\tau-1}(d\theta), \quad (25)$$

where  $N_\tau(\mathbf{p}) = \sum_{\mathbf{V} \in \mathcal{V}} N_\tau(\mathbf{p}, \mathbf{V})$ , and  $N_\tau(\mathbf{p}, \mathbf{V})$  is the number of times up to day  $\tau$  that the algorithm simultaneously observes a target profile  $\mathbf{V}$  and posts a price  $\mathbf{p}$ .

Furthermore, we define  $\mathbf{N}_\tau = [N_\tau(\mathbf{p})]_{\mathbf{p} \in \mathcal{P}}$  as a vector consisting of the number of times each price is posted up to day  $\tau$ . We can consider the quantity in the exponent of (25) as a loss suffered by model  $\theta$  up to day  $\tau$ . Since the term in the exponent of (25) is equal to 0 when  $\theta = \theta^*$ , we can see that Thompson sampling samples  $\theta^*$  and hence posts the optimal price with at least a constant probability at each day, i.e.,  $N_\tau(\mathbf{p}^{\mathbf{V},*}, \mathbf{V})$  grows linearly with  $\tau$  for all  $\mathbf{V}$ .

For each price, we define  $S_p(\mathbf{V}) := \{\theta \in \Theta : \mathbf{p}_\tau = \mathbf{p} | \mathbf{V}_\tau = \mathbf{V}\}$  to be the set of parameters  $\theta \in \Theta$  whose optimal price when observing a daily target load profile  $\mathbf{V}$  is  $\mathbf{p}$ . Furthermore, define  $S'_p(\mathbf{V}) := \{\theta \in S_p(\mathbf{V}) : KL(\theta_{\mathbf{p}^{\mathbf{V},*}}^* || \theta_{\mathbf{p}^{\mathbf{V},*}}) = 0\}$  which is the set of models  $\theta$  that exactly match  $\theta^*$  in marginal distribution of  $Y$  when the true model  $\theta^*$  is selected and the optimal price  $\mathbf{p}^{\mathbf{V},*}$  is posted, and  $S''_p(\mathbf{V}) := S_p(\mathbf{V}) \setminus S'_p(\mathbf{V})$ .

For each of the models  $\theta$  in  $S''_p(\mathbf{V})$ ,  $\mathbf{p} \neq \mathbf{p}^{\mathbf{V},*}$ ,  $KL(\theta_{\mathbf{p}^{\mathbf{V},*}}^* || \theta_{\mathbf{p}^{\mathbf{V},*}}) > \epsilon > 0$ . As we have assumed that the probability of observing any target profile  $\mathbf{V} \in \mathcal{V}$  is bounded away from zero,  $N_\tau(\mathbf{p}^{\mathbf{V},*})$  grows linearly with  $\tau$  for all  $\mathbf{V} \in \mathcal{V}$ . Hence, any such model  $\theta$  is sampled with probability exponentially decaying in  $\tau$  in (25) and the regret from such  $S''_p(\mathbf{V})$ -sampling is negligible. We define the set of all such models as  $\theta \in \Theta'' = \cup_{\mathbf{V} \in \mathcal{V}} S''_p(\mathbf{V})$ .

A model  $\theta \in S'_p(\mathbf{V})$  will only face loss whenever the algorithm posted a suboptimal price  $\mathbf{p}$  for which  $KL(\theta_p^* || \theta_p) > 0$ . For  $\mathbf{V}$ , a suboptimal price  $\mathbf{p}_k^{\mathbf{V}} \neq \mathbf{p}^{\mathbf{V},*}$  may still be posted if any of the set of models in  $S'_{\mathbf{p}_k^{\mathbf{V}}}(\mathbf{V})$  may still be drawn with non-negligible probability. Hence, a price will be eliminated after the probability of drawing all  $\theta \in S'_{\mathbf{p}_k^{\mathbf{V}}}(\mathbf{V})$  is negligible. For each  $\mathbf{V}$ , suboptimal prices are eliminated one after the other at times  $\tau_k^{\mathbf{V}}, k = 1, \dots, |\mathcal{P}| - 1$ . We refer the reader to [26] for a full discussion of when a suboptimal price  $\mathbf{p}$  is considered statistically eliminated, which is used to write constraints in (26) below.

**Theorem 1.** (Expanded Version) Under assumptions 1-4 and Constraint Set A in Algorithm 1, for  $\delta, \epsilon \in (0, 1)$ , there exists  $\mathcal{T}^* \geq 0$  s.t. for all  $\mathcal{T} \geq \mathcal{T}^*$ , with probability  $1 - \delta$ :

$$\sum_{\mathbf{V} \in \mathcal{V}} \sum_{\mathbf{p} \in \{\mathcal{P} \setminus \mathbf{p}^{\mathbf{V}, *}\}} N_{\mathcal{T}}(\mathbf{p}, \mathbf{V}) \leq B + C(\log \mathcal{T}),$$

where  $B \equiv B(\delta, \epsilon, \mathcal{P}, \mathcal{Y}, \Theta)$  is a problem-dependent constant that does not depend on  $\mathcal{T}$ , and  $C(\log \mathcal{T})$  depends on  $\mathcal{T}$ , the sequence of selected price signals, and the Kullback-Leibler divergence properties of the bandit problem (i.e., the marginal Kullback-Leibler divergences of the observation distributions  $KL[\ell(\mathbf{Y}; \mathbf{p}, \boldsymbol{\theta}^*), \ell(\mathbf{Y}; \mathbf{p}, \boldsymbol{\theta})]$ ). Specifically, the  $C(\log \mathcal{T})$  term is defined as follows:

$$\begin{aligned} C(\log \mathcal{T}) &\equiv & (26) \\ &\max_{\mathbf{V} \in \mathcal{V}} \sum_{k=1}^{|\mathcal{P}|-1} N_{\tau_k^{\mathbf{V}}}(\mathbf{p}, \mathbf{V}) \\ &\text{s.t. } \forall \mathbf{V} \in \mathcal{V}, \forall j > 1, \forall 1 \leq k \leq |\mathcal{P}| - 1 : \\ &\min_{\boldsymbol{\theta} \in \left\{ \mathbf{s}'_{\mathbf{p}_k^{\mathbf{V}}}(\mathbf{V}) - \boldsymbol{\theta}'' \right\}} \langle \mathbf{N}_{\tau_k^{\mathbf{V}}}, KL_{\boldsymbol{\theta}} \rangle \geq \frac{1+\epsilon}{1-\epsilon} \log \mathcal{T}, \\ &\min_{\boldsymbol{\theta} \in \left\{ \mathbf{s}'_{\mathbf{p}_k^{\mathbf{V}}}(\mathbf{V}) - \boldsymbol{\theta}'' \right\}} \langle \mathbf{N}_{\tau_k^{\mathbf{V}}} - e^{(j)}, KL_{\boldsymbol{\theta}} \rangle < \frac{1+\epsilon}{1-\epsilon} \log \mathcal{T}, \end{aligned}$$

where  $e^{(j)}$  denotes the  $j$ -th unit vector in finite-dimensional Euclidean space. The last two constraints ensure that price  $\mathbf{p}_k^{\mathbf{V}}$  is eliminated at time  $t_k^{\mathbf{V}}$  (no earlier and no later).

*Proof.* In Con-TS-RTP with Constraint Set A, the aggregator's daily objective and constraints are dependent on the sampled parameter  $\tilde{\boldsymbol{\theta}}_{\tau}$ . The only difference between Con-TS-RTP and the daily optimization in [15] is the added constraints. Since the constraints are only enforced for the sampled parameter, each sampled parameter  $\tilde{\boldsymbol{\theta}}_{\tau}$  still has a unique optimal price signal, and more importantly, the constraints do not prohibit the algorithm from selecting the optimal price for the sampled parameter. As such, the addition of constraints that depend only on the daily sampled parameter does not alter the bandit problem, and the regret analysis follows from [26].  $\square$

## B. Discussion on Operational Reliability

**Proposition 1.** (Repeated) Under assumptions 1-5, with  $\nu$  in equations (14)-(16) chosen such that  $\nu = \mu\pi_0(\boldsymbol{\theta}^*)e^{-\lambda|\mathcal{P}|}$ , with probability  $1 - \delta\sqrt{2}$ , the Con-TS-RTP algorithm with Constraint Set B will uphold the probabilistic distribution system constraints as formulated in (10) for each day  $\tau = 1, \dots, \mathcal{T}$ .

*Proof.* In [26], it is shown that with probability  $1 - \delta\sqrt{2}$  the mass of the true parameter never decreases below  $\pi_0(\boldsymbol{\theta}^*)e^{-\lambda|\mathcal{P}|}$  in the prior distribution during the learning process. As such, the desired reliability metric on the RHS of the constraints (14)-(16), i.e.,  $1 - \nu$ , can be selected such that the constraints *must* hold for the true parameter. Let  $\pi_{min}^* = \pi_0(\boldsymbol{\theta}^*)e^{-\lambda|\mathcal{P}|}$  be the minimum reachable mass of the true parameter in the prior distribution. Furthermore, we abuse

notation and denote  $\mathbb{P}_j^{safe} = \mathbb{P}_{\{\phi_c\}_{c \in C}}[g_j(\mathbf{D}_{\tau}(\mathbf{p}_{\tau})) \leq 0]$  as the probability that constraint  $j$  is upheld. Now, assuming the aggregator only has knowledge of the true parameter given by the prior distribution  $\pi_{\tau}$  on day  $\tau$ , the aggregator can calculate the probability of satisfying the constraint as follows:

$$\sum_{\hat{\boldsymbol{\theta}} \in \Theta} \pi_{\tau}(\hat{\boldsymbol{\theta}})(\mathbb{P}_j^{safe} | \boldsymbol{\theta} = \hat{\boldsymbol{\theta}}). \quad (27)$$

This can be split into two terms for the true parameter  $\boldsymbol{\theta}^*$  and all other parameters  $\boldsymbol{\theta} \neq \boldsymbol{\theta}^*$ :

$$\pi_{\tau}(\boldsymbol{\theta}^*)(\mathbb{P}_j^{safe} | \boldsymbol{\theta} = \boldsymbol{\theta}^*) + (1 - \pi_{\tau}(\boldsymbol{\theta}^*))(\mathbb{P}_j^{safe} | \boldsymbol{\theta} \neq \boldsymbol{\theta}^*). \quad (28)$$

Now, we can rewrite the probability assuming that  $\boldsymbol{\theta}^*$  has reached the minimum mass  $\pi_{min}^*$  in the prior distribution:

$$\pi_{min}^*(\mathbb{P}_j^{safe} | \boldsymbol{\theta} = \boldsymbol{\theta}^*) + (1 - \pi_{min}^*)(\mathbb{P}_j^{safe} | \boldsymbol{\theta} \neq \boldsymbol{\theta}^*). \quad (29)$$

Recall, the aggregator wants constraint  $j$  to hold with probability at least  $1 - \mu$  for the true parameter  $\boldsymbol{\theta}^*$ , so we can replace  $(\mathbb{P}_j^{safe} | \boldsymbol{\theta} = \boldsymbol{\theta}^*)$  with  $1 - \mu$ . Furthermore,  $(\mathbb{P}_j^{safe} | \boldsymbol{\theta} \neq \boldsymbol{\theta}^*) \leq 1$  and we replace it accordingly yielding:

$$\pi_{min}^*(1 - \mu) + (1 - \pi_{min}^*). \quad (30)$$

Now, we want this probability to be the minimum allowable probability across the prior  $\pi$  for constraint  $j$  to hold so we set it equal to the reliability metric:

$$\pi_{min}^*(1 - \mu) + (1 - \pi_{min}^*) = 1 - \nu, \quad (31)$$

which yields

$$\nu = \mu\pi_{min}^*. \quad (32)$$

By selecting  $\nu = \mu\pi_{min}^*$ , the aggregator ensures that constraint  $j$  will be upheld with probability at least  $1 - \mu$  for the true parameter  $\boldsymbol{\theta}^*$ . (i.e., the total mass of the incorrect parameters  $\boldsymbol{\theta} \neq \boldsymbol{\theta}^*$  in the prior distribution  $\pi_{\tau}$  can never be large enough to satisfy the constraint's inequality without the true parameter also satisfying the constraint).  $\square$