Dynamic Estimation of the Price-Response of Deadline-Constrained Electric Loads under Threshold Policies

Mesrob I. Ohannessian, Mardavij Roozbehani, Donatello Materassi, and Munther A. Dahleh

***Abstract*—The paper presents a consistent and unbiased estimator for dynamic, one-step-ahead prediction of the aggregate response of a large number of individual loads to a common price signal, using only aggregate past response data. The price per unit of consumption is an exogenous signal which is updated at discrete time intervals. It is assumed that individual loads arrive in the system at random times with random demands and random consumption deadlines, and may defer their consumption up to the deadline in order to minimize their total cost. It is further assumed that the individual loads adopt a threshold policy in the sense that they only consume when the price is below a certain threshold. A dynamic aggregate model is constructed from models of independent individual loads. A consistent and unbiased estimator which only uses aggregate data, i.e., the price and aggregate consumption time-series is presented for estimating the aggregate consumption as a function of price.**

***Index Terms*—Load Scheduling, Estimation.**

I. Introduction

Real-time demand response from *intelligent* loads in response to the real-time wholesale electricity market conditions has been identified as one of the distinctive features of future power grids [1], and both the theory and the practical implementation aspects of this paradigm change have been the subjects of a growing number of studies, see, e.g. [2], [3], [4], [5]. Such active participation of loads in real-time operation of power systems is primarily advocated as a mechanism to help match supply and demand, absorb exogenous supply shocks, and improve the overall efficiency of the system. It is important to note that sending price signals to electric loads is not the only way to close the loop between wholesale electricity markets and retail consumers, and other viable schemes such as long-term contracts along with automation and control have been considered and studied [6], [7], [8].

Given that a significant amount of the price elasticity of consumers will come from deferring flexible loads, several papers have examined optimal scheduling policies under various models for deadline-constrained deferrable loads under real-time or day-ahead pricing [9], [10], and [11]. An interesting question of practical significance that has received less attention is how to predict the aggregate consumption in response to real-time electricity prices? What information and at what level of detail is needed in order to reach an acceptable level of prediction accuracy?

Due to load-shifting and deadline constraints, taking the dynamics of consumption into consideration becomes important, and static price-demand estimation approaches will have limited applicability when it comes to real-time system operation in which more accurate demand predictions are needed. For these reasons, developing mathematical models of price-responsiveness of consumers with flexible loads, smart appliances, or electric vehicles, as well as the emerging aggregate behaviors from distributed load-scheduling is fundamental for reliable and efficient operation of future power systems. In particular, predicting the response of these systems to a given price signal from limited aggregate information is a problem of paramount importance.

In this paper, we consider a model in which individual consumers with random demands and random deadlines arrive in the system at random times and implement a threshold policy to minimize their individual costs. We present an explicit formula for computing the aggregate demand as a function of price history and the arrival rates. We then present a consistent and unbiased estimator which only uses the history of price versus total consumption for estimating the arrival rates. These estimated arrival rates are then used to estimate the aggregate demand as a function of price. This is the main contribution of this paper and is important for at least two reasons: First, in practice, the only available data might be the price-consumption history, and the actual arrival rates of consumers may be unknown or hard to estimate. Second, given that prices induce correlation in the consumption of different consumers, it is not a-priori clear that consumer arrival rates can be estimated from price-consumption history.

For brevity and due to space limitations, the proofs of the mathematical statements are omitted but can be found in the extended version [12].

II. Model

A. Individual Consumer Model

Throughout the paper, time is discrete and denoted by $k \in \mathbb{Z}_+$. Let $N > 1$ be a positive integer. Let the *type* of an individual consumer be specified by a parameter $n \in \{1, \ldots, N\}$, which represents a deadline constraint, that is, the number of periods within which the demand must be fulfilled. The model of an individual load/consumer is then

Mesrob Ohannessian is the ERCIM Postdoctoral Fellow at the Département de Mathématiques of the Universitè Paris-Sud. mesrob@gmail.com.

Mardavij Roozbehani is a Principal Research Scientist at the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology (MIT). mardavij@mit.edu.

Donatello Materassi is a Postdoctoral Research Associate at LIDS, MIT. donnie13@mit.edu.

Munther A. Dahleh is a Professor of Electrical Engineering and Computer Science (EECS) at MIT. dahleh@mit.edu.

defined as follows:

$$x(1) = u(1) - d(1)$$
$$x(k+1) = x(k) + u(k) - d(k), \qquad k = 1, \ldots, n-1,$$

where, $x(k)$ is the backlog state, $d(k)$ can be interpreted as demand or exogenous disturbance during period k, and $u(k)$ is the consumption decision according to the consumer's policy.

A general class of policies is *threshold policies* characterized by a set of *thresholds* that depend on the time-to-go until the deadline. Such a class can be described by a vector $(\tau_t)_{t\in\{1,\ldots,N\}} \in \mathbb{R}^N$, which is an N-tuple of parameters indexed by the time-to-go t until the deadline[1]. The corresponding threshold policy is described by:

$$u(k) = \begin{cases} -x(k) + d(k), & \lambda(k) \le \hat{\tau}, \\ 0, & \lambda(k) > \hat{\tau}. \end{cases} \tag{1}$$

where $\lambda(k)$ is the price per unit of consumption paid by each consumer during period k, and $\hat{\tau}$ is the corresponding threshold at time k. We assume that the price process is a stationary stochastic process which is exogenously determined, i.e., it is independent of the decisions the consumers. In what follows, we assume that any individual consumer follows the threshold policy (1). This is motivated by the fact that cost-optimal policies[2] under various stochastic models for the price fall in this category (see, for example, [11],[13],[9]).

B. Aggregate Model and Price-Consumption Function

The main idea behind this section is to present the aggregate consumption as a function of price by keeping track of history of arrivals as well as the prices, and how the consumers would have backlogged. In the aggregate model the consumers enter the system according to a stochastic arrival process, respond to the real-time price by consuming or shifting their demands based on the threshold policy (1), and exit the system at their respective deadlines. We denote the aggregate *price-consumption function* by

$$u(k, \lambda),$$

which represents the aggregate consumption as a function of price $\lambda(k)$ at each time step k, acknowledging that it depends implicitly on the histories of price, demand, and agent arrivals.

Definition 1. *An* arrival model *is a non-negative, integer-valued, N-dimensional vector stochastic process $(A(k))_{k\in\mathbb{N}_+}$, where for each component $n \in \{1, \cdots, N\}$ and each time k, the sample path $A_n(k)$ represents the number of new consumers that arrive with deadline n at time k.*

Consider two attributes of an individual consumer in this model at any specific time k: (1) the time-to-go until her individual deadline which we denote by $t \in \{1, \cdots, N\}$, and (2) the time she has been in the system so far which we denote by $s = n - t \in \{0, \cdots, N-t\}$. For convenience in notation, let $\mathcal{G}_{s,t}(k)$ denote the group of all consumers at time k, that have time-to-go t, and have been in the system for s periods. All such consumers use the same threshold for responding to the price signal, though they may have different backlogs and there is heterogeneity in the intensity of their responses.

Let $u_{s,t}(k, \lambda)$ be the consumption of all consumers in $\mathcal{G}_{s,t}(k)$ in response to price λ. Then, we clearly have:

$$u(k, \lambda) = \sum_{t=1}^{N} \sum_{s=0}^{N-t} u_{s,t}(k, \lambda)$$

Let $x_{s,t}(k)$ denote the total backlog of consumers in $\mathcal{G}_{s,t}(k)$ and let $d_{s,t}(k)$ denote the total demand of consumers in $\mathcal{G}_{s,t}(k)$, that is:

$$d_{s,t}(k) = \sum_{j\in\mathcal{G}_{s,t}(k)} d^j(k)$$

Due to the threshold policy of the loads, all consumers in $\mathcal{G}_{s,t}(k)$ consume their backlog in addition to their current demand if and only if $\lambda(k)$ is less than or equal to τ_t, and consume nothing otherwise. Thus, $u_{s,t}(k, \lambda)$ can be expressed as:

$$u_{s,t}(k, \lambda) = \begin{cases} -x_{s,t}(k) + d_{s,t}(k), & \lambda \le \tau_t, \\ 0, & \lambda > \tau_t. \end{cases} \tag{2}$$

Therefore, we can characterize the aggregate consumption as a function of price by computing the backlog for each group in the system. Since the maximal deadline is N, the backlog accumulation duration is completely determined by the price history from time $k - N + 1$ to time $k - 1$. Let $\overleftarrow{\lambda}(k) = (\lambda(k-N+1), \cdots, \lambda(k-1))$ be a shorthand notation for this price history, and let $\overleftarrow{\Lambda} = (\Lambda_{-N+1}, \cdots, \Lambda_{-1})$ represent a vector of price history without reference to a specific time k. Define by $B_{s,t}(\overleftarrow{\Lambda})$ to be the exact duration over which any consumer in groups $\mathcal{G}_{s,t}(\cdot)$ has accumulated backlog when the price history is $\overleftarrow{\Lambda}$. This duration cannot exceed s, the time in the system, and stretches as far back as the last time the price fell below these consumers' common threshold. Formally, we have:

$$B_{s,t}(\overleftarrow{\Lambda}) = \min\left(\{s\} \cup \{b \in \mathbb{N} : \Lambda_{-b-1} \le \tau_{t+b+1}\}\right). \tag{3}$$

With a slight abuse of notation, we denote the backlog accumulation duration of $\mathcal{G}_{s,t}(k)$ by $B_{s,t}(\overleftarrow{\lambda}(k))$. It can now be seen that the total backlog of all consumers in $\mathcal{G}_{s,t}(k)$ is given by:

$$x_{s,t}(k) = -\sum_{\ell=k-B_{s,t}(\overleftarrow{\lambda}(k))}^{k-1} d_{s,t}(\ell). \tag{4}$$

Therefore, by using (4) and (2), the aggregate price-consumption function at time k can be expressed explicitly in

[1]Note that for a particular consumer with deadline n, only the first n thresholds are relevant.

[2]All admissible policies that guarantee the entire demand would be fulfilled by the deadline must satisfy $\tau_1 \ge \lambda_{\max} = \max_k \lambda(k)$.

terms of the price, demand and arrival processes as follows:

$$u(k,\lambda) = \sum_{t=1}^{N}\sum_{s=0}^{N-t} u_{s,t}(k,\lambda) = \sum_{t:\lambda\le\tau_t}\sum_{s=0}^{N-t}\left(\sum_{\ell=k-B_{s,t}(\overleftarrow{\lambda}(k))}^{k}\ \sum_{j\in\mathcal{G}_{s,t}(k)} d^j(\ell)\right). \quad (5)$$

C. Averaged Model

As a proxy for $u(k,\lambda)$, which, in the absence of the exact knowledge of all arrivals and demands is impossible to compute, we propose the *averaged price-consumption function*:

$$\overline{u}(\overleftarrow{\Lambda},\lambda) := \mathbf{E}[u(k,\lambda)\,|\,\overleftarrow{\lambda}(k)=\overleftarrow{\Lambda}],$$

where the averaging is over the unknown arrivals and demands. This is motivated by the fact that a system operator may have knowledge of the model, including the stochastic description of the demand and arrivals, but not of their actual instances.

Assumption 1 (Stationarity Condition)**.** *The arrival process is independent for different deadlines, is independent from the demand process, and is i.i.d. across time, with mean* $\overline{A}_n = \mathbf{E}[A_n(k)]$ *and variance* $\sigma^2_{A_n} = \mathrm{var}[A_n(k)] < \infty$ *for all* n *and* k*. Each consumer has an independent instance of the same i.i.d. demand process, with mean* $\overline{d} = \mathbf{E}[d^j(k)]$ *and variance* $\sigma^2_d = \mathrm{var}[d^j(k)] < \infty$ *for all users* j *and time* k*.*

With these assumptions, once we have the average rates $\overline{d}$ and $\overline{A}_n$, we can readily compute the averages of the quantities in Equations (4), (2), and (5) in the following way:

$$\mathbf{E}[x_{s,t}(k)\,|\,\overleftarrow{\lambda}(k)=\overleftarrow{\Lambda}] = -\overline{d}\;\overline{A}_{s+t}\,B_{s,t}(\overleftarrow{\Lambda}), \quad (6)$$

$$\mathbf{E}[u_{s,t}(k,\lambda)\,|\,\overleftarrow{\lambda}(k)=\overleftarrow{\Lambda}] = \begin{cases} \overline{d}\,\overline{A}_{s+t}\,(B_{s,t}(\overleftarrow{\Lambda})+1), & \lambda\le\tau_t,\\ 0, & \lambda>\tau_t. \end{cases} \quad (7)$$

Therefore,

$$\overline{u}(\overleftarrow{\Lambda},\lambda) = \sum_{t:\lambda\le\tau_t}\sum_{s=0}^{N-t}\overline{d}\,\overline{A}_{s+t}\,(B_{s,t}(\overleftarrow{\Lambda})+1). \quad (8)$$

III. Main Results

In this section we present our main results. We first describe the relative accuracy of the averaged model (Equation (8)) versus the exact aggregate model (Equation (5)). This result characterizes the regimes (in terms of model parameters) for which using the averaged price-consumption function as an analysis and forecasting tool is justified.

A. Accuracy of the Averaged Model

It is imperative to ask how does the computed expected consumption perform as a predictor of the actual consumption? In the following theorem we present conditions under which (8) provides an *accurate estimate* of the *actual* demand.

Theorem 1. *Given a price history* $\overleftarrow{\lambda}(k)=\overleftarrow{\Lambda}$*, and a price* $\lambda(k)=\lambda$*, define* $e(k,\lambda)$ *as the relative error of predicting the actual demand based on the averaged demand. More precisely, let:*

$$e(k,\lambda) = \frac{u(k,\lambda)}{\overline{u}(\overleftarrow{\Lambda},\lambda)} - 1. \quad (9)$$

Suppose that Assumption 1 is satisfied. Then, for all $\overleftarrow{\Lambda}$ *and* λ *the mean-square relative prediction error, averaged over demands and arrivals, satisfies:*

$$\mathbf{E}[e^2(k,\lambda)|\overleftarrow{\lambda}(k)=\overleftarrow{\Lambda}] \le \frac{1}{\sum_n \overline{A}_n}\sigma_d^2/\overline{d}^2 + N\frac{\max_n \overline{A}_n}{\sum_n \overline{A}_n}\ \max_n\ \sigma^2_{A_n}/\overline{A}_n^2. \quad (10)$$

Inequality (10) has an interesting and intuitive explanation. Indeed, $\sigma_d^2/\overline{d}^2$ and $\max_n\ \sigma^2_{A_n}/\overline{A}_n^2$ represent the mean-square relative errors of using averages to approximate the demands and arrivals, respectively. These terms contribute additively to the relative error of predicting consumption, uniformly across all prices. In addition, there are two critical properties of the arrivals that affect the accuracy of prediction. These are: the arrival sum-rate and the arrival rate *heterogeneity*. We obtain this additional insight by examining the factors in (10):

- The demand relative error is scaled inversely by the arrival sum-rate $\sum_n \overline{A}_n$. Therefore when the latter is large, demand randomness becomes less relevant. This can be explained by the increased averaging within each group $\mathcal{G}_{s,t}(\cdot)$.
- The arrival relative error is scaled by a measure of heterogeneity among arrival rates for various deadlines: the factor $\left(\max_n \overline{A}_n\right)/\left(\sum_n \overline{A}_n\right)$. To justify this interpretation, note that it is indeed minimal (equal to $1/N$) when all arrival rates are equal, i.e., extreme homogeneity. Conversely, it is maximal (equal to 1) when a single arrival rate dominates all others, i.e., extreme heterogeneity.

Examining the case when the arrival means and variances are of the same order (e.g. for the discretized Poisson arrivals where $\sigma^2_{A_n}/\overline{A}_n = 1$, c.f. Section IV) provides additional insight. In this case, we have $\max_n\ \sigma^2_{A_n}/\overline{A}_n^2 = 1/\min_n \overline{A}_n$. The contribution to the relative error by the arrival randomness then becomes:

$$N\ \frac{\max_n \overline{A}_n/\min_n \overline{A}_n}{\sum_n \overline{A}_n}.$$

We see that while the numerator is still a measure of heterogeneity, the arrival-sum rate scales the error inversely.

We may thus conclude that the demand-averaged and arrival-averaged perspective is particularly accurate when arrival rates are large and homogeneous.

B. Estimating the Averaged Price-Consumption Function

1) Price-consumption data and consistent estimation: The estimator (8) allows us to forecast consumption as a function of price if the average demand and arrival rates at various deadlines are known. We show that when these rates are not a-priori known, by using the structure of the model we can learn them from the price-consumption history.

Definition 2. *The* price-consumption data *of size K consists of the price history $\lambda(k),\ k = 1, \cdots, K$, and the consumption history $u(k, \lambda(k)),\ k = 1, \cdots, K$. We denote this by $(\lambda, u)_K$.*

Our goal is to design a consistent estimator of the averaged price-consumption function (8) from this limited information.

Definition 3. *An* estimator *of $\overline{u}(\cdot, \cdot)$ is a sequence of functions $\widehat{u}_K$ that map $(\lambda, u)_K$ to a real-valued function $\widehat{u}_K(\cdot, \cdot)$ on $\mathbb{R}^N$. A* consistent estimator *is one that guarantees, as $K \to \infty$, that:*

$$\mathbf{E}\left[\sup_{\overleftarrow{\Lambda}, \lambda} \left|\overline{u}(\overleftarrow{\Lambda}, \lambda) - \widehat{u}_K(\overleftarrow{\Lambda}, \lambda)\right|\right] \to 0.$$

The above definition is motivated by the fact that one would like to obtain an estimator that uniformly converges to $\overline{u}$ over all possible sets of prices. Note that $\overline{u}$ is parametrized by the average demand-weighted arrival rate vector $\overline{dA}$. Consequently, our approach is based on estimating this combined vector in order to design an estimator of $\overline{u}$. We present the following proposition as a direct result of Equation 8 and the fact that $B_{s,t}(\overleftarrow{\Lambda}) + 1$ is bounded by N.

Proposition 1. *Let $\overline{dA}$ and $\overline{dA}'$ be two average demand-weighted arrival rates, and denote by $\overline{u}(\overleftarrow{\Lambda}, \lambda)$ and $\overline{u}'(\overleftarrow{\Lambda}, \lambda)$ the corresponding averaged price-consumption functions. Then:*

$$\sup_{\overleftarrow{\Lambda}, \lambda} \left|\overline{u}(\overleftarrow{\Lambda}, \lambda) - \overline{u}'(\overleftarrow{\Lambda}, \lambda)\right| \leq N \sum_{n=1}^{N} n |\overline{dA}_n - \overline{dA}'_n|.$$

This motivates the following definitions:

Definition 4. *An* estimator *of $\overline{dA}$ is a sequence of functions $\widehat{dA}(K)$, that map $(\lambda, u)_K$ to a vector in $\mathbb{R}^N$. A* consistent estimator *is one that guarantees, for each component n as $K \to \infty$, that:*

$$\mathbf{E}\left[\left|\overline{dA_n} - \widehat{dA_n}(K)\right|\right] \to 0.$$

Definition 5. *A* π-estimator *of $\overline{u}$ corresponding to an estimator $\widehat{dA}(K)$ is an estimator $\widehat{u}(K)$ which is obtained by evaluating the averaged price-consumption function with $\widehat{dA}(K)$ instead of $\overline{dA}$ in Equation (8).*

It is then only a quick corollary of Proposition 1 to show that if $\widehat{dA}(K)$ is consistent, then the corresponding π-estimator $\widehat{u}(K)$ is consistent as well. We therefore move our attention to estimating $\overline{dA}$. It is tempting to use the linear structure of Equation (8) and perform linear regression in order to construct an estimator of $\overline{dA}$. However, since it is the actual consumption and not the averaged consumption that is observed, Equation (5) does not lend itself easily to a straightforward regression on $\overline{dA}$. This intuition, however, serves as the basis for our approach.

2) Matrix representation: For convenience, we first introduce a matrix representation for the dynamics of the system. The following definitions transform the price-consumption data into a consumption vector and a backlog duration matrix, and represent the (unobserved) arrivals as a stacked vector. To bypass the transient dynamics we can consider consumption only at times $k \geq N$.

Definition 6. *Given the price-consumption data $(\lambda, u)_K$, define the* consumption vector $\mathbf{u}(K)$ *to be a $K - N + 1$-dimensional vector indexed by $k \in \{N, \cdots, K\}$:*

$$\mathbf{u}_k(K) = u(k, \lambda(k)). \tag{11}$$

Define also the arrivals vector $\mathbf{A}(K)$ *as an NK-dimensional vector indexed by $n \in \{1, \cdots, N\}$ and $i \in \{1, \cdots, K\}$:*

$$\mathbf{A}_{n+N(i-1)}(K) = A_n(i).$$

Proposition 2. *If Assumption 1 holds, then for every $k \in \{N, \cdots, K\}$ the consumption $\mathbf{u}_k(K)$ is the sum of $(\mathbf{B}(K)\mathbf{A}(K))_k$ demands, where $\mathbf{B}(K)$ is a* backlog duration matrix *that depends only on the price history within the price-consumption data $(\lambda, u)_K$.*

These demands correspond to the $d^j(\ell)$ that appear in Equation (5). They are i.i.d. within and across sums for various k, and each is distributed as in Assumption 1. Without explicit indexing this can be expressed as:

$$\mathbf{u}_k(K) = \sum_{1}^{(\mathbf{B}(K)\mathbf{A}(K))_k} d. \tag{12}$$

This proposition makes it transparent that the consumption has the structure of a compound model, comparable to the compound Poisson process. The linear structure is in the random number (that depends on backlog and arrivals) of terms in a sum of i.i.d. random variables (the demands).

3) A consistent estimator of $\overline{dA}$: Before we state Theorem 2, we introduce an identifiability condition. Let $(\cdot)^\sharp$ denote the pseudo-inverse of a matrix. Let $\mathbf{R}(K)$ indicate K vertically stacked $N \times N$ identity matrices:

$$\mathbf{R}(K) = [\mathbf{I}_N \mid \ldots \mid \mathbf{I}_N]^\intercal.$$

One could think of $\mathbf{R}(K)$ as a block-replicating operator when acting on a matrix from the left, and a block-summing operator when multiplying from the right. To state the identifiability condition concisely, note that if we partition the space of prices based on the partition of the real numbers by the thresholds $\tau_1, \cdots, \tau_N$, then any variation within each part does not alter the dynamics of the system. Therefore such a partition defines a (finite) *equivalence class* of prices.

Assumption 2 (Identifiability Condition)**.** *There exists some $\delta > 0$ such that for K large enough, the price-consumption data $(\lambda, u)_K$ contains at least a fraction δ of each possible value of $(\overleftarrow{\lambda}(k), \lambda(k))$, up to equivalence.*

Theorem 2. *Consider the estimator $\widetilde{dA}(K)$ of the average demand-scaled arrival rates $\overline{dA}$ given by*

$$\widetilde{dA}(K) = \left(\mathbf{B}(K)\mathbf{R}(K)\right)^{\sharp} \mathbf{u}(K). \tag{13}$$

If Assumption 2 holds, then $\widetilde{dA}(K)$ is consistent. More precisely, for large enough K, $\widetilde{dA}_n(K)$ are unbiased for all n, and their variances decay like $\mathcal{O}(\frac{1}{K})$.

IV. Simulations

In this section, we assume that each $\{A_n(k)\}_{k \in \mathbb{N}_+}$ for $t = 1, \cdots, N$, is a discretized version of a Poisson process with rate α_n. The discretization is in the sense that all arrivals within each unit interval are assumed to enter the system at an endpoint of the interval, and the time axis is normalized such that each unit interval corresponds to a unit of time in our model. It then follows that $\overline{A}_n = \alpha_n$.

Example 1 - Learning, and prediction: We first examine the performance of the proposed estimator for the case of constant demand $d = 1$, leaving investigating variable demand to the next example. We simulate the case with 20 independent instances of $\widetilde{dA}$ learned from $K = 100$, 1000, and 10000 samples. In Figure 1, we plot a sample $u(k, \lambda)$,

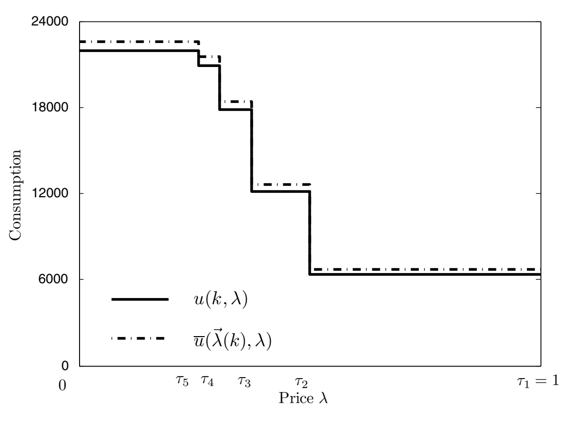

Fig. 1. A sample price-consumption function $u(k, \lambda)$ at an arbitrary time k, together with the demand-averaged and arrival-averaged approximation $\overline{u}(\vec{\lambda}(k), \lambda)$, which the system operator can compute.

at some arbitrary k. Note the piecewise constant property, with changes only at $\tau_N, \tau_{N-1}, \cdots, \tau_2$, due to the threshold policy implemented by the consumers. The sloped response, with higher consumption at lower prices, captures the price-elasticity of the aggregate demand. Figure 1 also shows the demand-averaged and arrival-averaged price-consumption function $\overline{u}(\overleftarrow{\lambda}(k), \lambda)$ at the same time k, illustrating the close approximation.

Example 2 - Accuracy of the average approximation: To quantify the goodness of the approximation, we consider the relative error (Equation (9)) between the true price-consumption function and the demand-averaged and arrival-averaged approximations. To average out this error over arrivals, demands, and price histories, we simulate $K = 10,000$ steps, and evaluate the sample root-mean-square $\sqrt{\frac{1}{K}\sum(k) e^2(k, \lambda)}$.

Figure 2 shows the dependence of the approximation error on the arrival rates, using three cases:

- The case of the previous example, i.e. $\alpha_n = 100(2n + 1)$, for $t = 1, 2, 3, 4, 5$. The error is at most of the order of $\pm 1.7\%$.
- The case when the arrival rates are doubled, i.e. $\alpha_n = 200(2n + 1)$, for $t = 1, 2, 3, 4, 5$. The error is approximately $1/\sqrt{2}$ times the previous example's error, which agrees with what (10) predicts.
- The case when the arrival rates are halved i.e. $\alpha_n = 50(2n + 1)$, for $t = 1, 2, 3, 4, 5$. The error is approximately $\sqrt{2}$ times the previous example's error, as predicted.

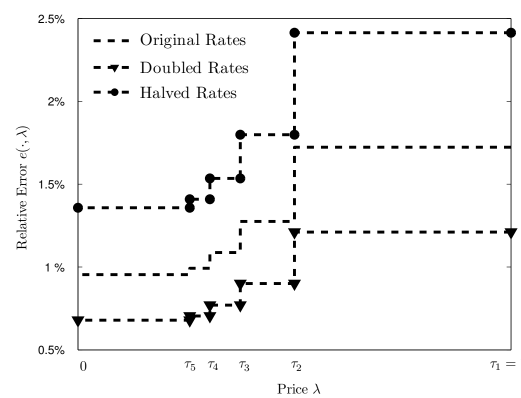

Fig. 2. The sample root-mean-square relative approximation error, computed using $10,000$ samples, for each of three arrival configurations.

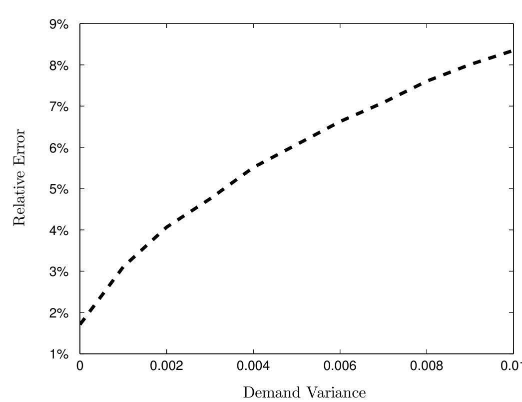

Fig. 3. The largest sample root-mean-square relative approximation error over 10,000 samples, as a function of demand variance.

So far we have assumed constant demand $d = 1$. We expect the error to be worse with demand randomness. To explore this, it is easiest to simulate all consumers using the same demand process, and therefore highly correlated. This is a departure from the simplified model presented in Section II-C. However, when various consumers have independent demands, the relative error of the approximation is necessarily smaller, by virtue of Theorem 1 and the discussion in Section III-A. The proposed analysis can therefore be conservative.

Example 3 - The effect of heterogeneity of demand: We fix the arrival rates to that of the previous examples, and consider a sequence of uniformly distributed demands, all centered at 1, but with variances ranging from 0 to 0.01. As described, these demands are independent between time steps but identical among all consumers. In Figure 3, we plot the largest (over price) root-mean-square relative error that results under each such demand distribution. Despite this more conservative analysis, these observations are consistent with the prediction of (10). These experiments suggest that time-heterogeneity of demands does impact the goodness of the price-consumption function approximation, but only moderately so, and the latter remains a viable tool for predicting aggregate consumption in response to time-varying prices.

V. Conclusions and Future Work

From the model of an individual consumer we developed a dynamic macro model which maps the price history and the arrivals of individual consumers to an aggregate consumption as a function of price. When the actual instances of consumer arrivals and their demands are not known, the expected aggregate consumption in response to a given price, conditioned on the price history, can be computed as a function of the rates of arrivals and the average demands. We showed that this expected aggregate consumption can be used as an effective predictor for actual aggregate consumption, by characterizing its mean relative error. Moreover, when the rates themselves are not known, we developed a consistent and unbiased estimator to estimate them, using only the history of consumption versus price.

Our results also highlight the effect of heterogeneity of consumers in prediction error. As one may expect, the relative error of the prediction decreases as the mean rates increase or their variances decrease. However, when the consumers are more heterogeneous, they pose an additional source of uncertainty. In particular, if consumer arrival rates are highly variable among various deadlines, then the error between predicted consumption and actual consumption can be higher. These results suggest that in practice, predicting the actual response of a large population of consumers to price signals with high accuracy can be a challenging task for system operators.

VI. Acknowledgments

Th authors thank The National Science Foundation (NSF-CPS 1135843), Draper Laboratories, the MIT-Masdar collaborative project, and the MIT-Skoltek initiative (SDP) for financial support of this research.

References

[1] K. Afridi *et al*, *The Future of the Electric Grid*. MIT Press, 2011.

[2] S. Borenstein, M. Jaske, and A. Rosenfeld, "Dynamic pricing, advanced metering and demand response in electricity markets," *UC Berkeley: Center for the Study of Energy Markets.*, 2002.

[3] W. Hogan, "Demand response pricing in organized wholesale markets," *IRC Comments, Demand Reponse Notice of Proposed Rulemaking, FERC Docket RM10-17-000*, May 2010.

[4] M. Roozbehani, M. Dahleh, and S. Mitter, "Volatility of power grids under real-time pricing," *Power Systems, IEEE Transactions on*, vol. 27, no. 4, pp. 1926–1940, 2012.

[5] ——, "Dynamic pricing and stabilization of supply and demand in modern electric power grids," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, 2010, pp. 543–548.

[6] S. Meyn, P. Barooah, A. Busic, and J. Ehren, "Ancillary service to the grid from deferrable loads: The case for intelligent pool pumps in florida," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, Dec 2013, pp. 6946–6953.

[7] E. Bitar and S. Low, "Deadline differentiated pricing of deferrable electric power service," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 2012, pp. 4991–4997.

[8] C.-W. Tan and P. Varaiya, "Interruptible electric power service contracts," *Journal of Economic Dynamics and Control*, vol. 17, no. 3, pp. 495–517, 1993.

[9] A. Papavasiliou and S. Oren, "Supplying renewable energy to deferrable loads: Algorithms and economic analysis," in *Power and Energy Society General Meeting, 2010 IEEE*. IEEE, 2010, pp. 1–8.

[10] T. Kim and H. Poor, "Scheduling power consumption with price uncertainty," *Smart Grid, IEEE Transactions on*, vol. 2, no. 3, pp. 519–527, 2011.

[11] M. Roozbehani, D. Materassi, M. Ohannessian, and M. A. Dahleh, "Robust and optimal consumption policies for deadline-constrained deferrable loads," *IEEE Transactions on Smart Grids*, vol. *In Press*, 2014.

[12] M. Ohannessian, M. Roozbehani, D. Materassi, and M. A. Dahleh, "Dynamic estimation of the price-response of deadline-constrained electric loads with threshold policies: Extended online version," Sept 2013. [Online]. Available: http://www.mit.edu/~mardavij/publications.html

[13] M. J. Neely, A. S. Tehrani, and A. G. Dimakis, "Efficient algorithms for renewable energy allocation to delay tolerant consumers," in *First IEEE International Conference on Smart Grid Communications*, 2010.