Robust and Optimal Consumption Policies for Deadline-Constrained Deferrable Loads

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

***Abstract*—The paper analyzes the optimal response of an individual smart load with deferrable demand for electricity, to exogenous and stochastic price processes. It is assumed that a smart load can delay a time-flexible energy demand up to a fixed deadline. Under mild assumptions on the regularity of the stochastic price process, it is shown that the optimal strategy is to consume only when the price is less than or equal to a certain threshold that depends only on the time left to the deadline and the price statistics. This analysis is performed under both perfect and partial information about the statistics of the price process. Robust policies with performance guarantees are derived for the partial information case. Such performance bounds are also used for deriving upper and lower bounds on the economic value of load-shifting. Numerical simulation results based on price data from wholesale electricity markets suggest that when information must be empirically extracted from data, the robust policies provide various theoretical and practical advantages over the full information policies.**

I. Introduction

It is envisioned that in future smart grids, consumers or *smart loads* on their behalf will adjust their consumption in real-time to help mitigate the effects of the intermittencies of renewable resources. New mechanisms for demand response through contracts with industrial and residential consumers, and through dynamic pricing to induce shifting of flexible loads are possible enablers of this technological change [1]. Indeed, it has been argued that at any given time, a considerable amount of the total generated power is supplied to flexible loads that are shiftable in time by a few minutes to a few hours at little or no cost to the final user [2], [3]. Examples abound and can be found in the areas of material processing, electric vehicle (EV) charging, heating, ventilation, air conditioning, refrigeration, and agricultural irrigation. Dynamic pricing mechanisms can incentivize smart loads with flexible demand to backlog their consumption when prices are high, thereby relieving the grid of congestion or other strenuous situations related to shortage of supply. By the same token, smart loads will tend to consume when prices are low and the grid is underutilized [4], [5].

Motivated by these considerations, we present a mathematical model to investigate the optimal response of smart loads to exogenous and stochastic electricity prices. Although the characterization of the associated optimal policy is closely related to some variations of the well-studied inventory control problem (see, e.g., [6], [7], [8] and the references therein), the existing literature related to energy consumption under price uncertainty and deadline constraints is relatively narrow. Several papers including [9], [10], [11], [12], [2], [13], and [14] have introduced the general concepts and examined detailed models for optimization of energy consumption (including load-shifting) under price and/or supply uncertainty.

Compared to many previous quantitative studies, our model is more abstract, since it does not consider physical parameters such as temperature, internal dynamics of heterogenous appliances, or the mechanisms by which consumers or smart appliances may react to prices. Instead, we include more general descriptive features for the characteristics of energy consumption under real-time pricing: the dynamics of backlog demand, the deadline constraints, the statistical knowledge of the price process, and the disutility of delay. Our model does not include bounds on the possible consumption during a single time interval and assumes the demand profile of the task is known. These hypotheses are not very restrictive and are met in many application domains. Indeed, technological improvements are leading to batteries that can be rapidly charged [15], and in many application scenarios, such as EV charging, the energy demand profile of a device can be assumed known (to the device) *a-priori*. As we will show in this article, our model is sufficiently tractable to allow for several insightful results about the optimal response of individual loads. Furthermore, such abstract and generic models are more likely to lead to development of tractable models of aggregate consumption in response to real-time prices. Such models are needed for feedback design in power distribution grids with demand response [16]. In order to be useful for feedback design purposes, the models need to be simple and tractable, and cannot capture all the details. Robust control theory then allows for feedback design based on a simplified nominal model with robustness against uncertainties and unmodeled dynamics.

From a methodological perspective, studies on energy consumption optimization based on stochastic dynamic programming (DP) or Lyapunov optimization, such as the recent papers by [2], [17], and [14], are closest to our work. In particular, a result similar to ours on the affine structure of the value function associated with the optimal load-shifting problem was obtained in [2] in a setup which directly couples renewable generation with flexible loads and uses stochastic dynamic programming to match intermittent generation with consumption. In [17] a problem setup analogous to ours is posed for optimal scheduling of both interruptible and non-interruptible loads and optimal threshold policies are derived. Similar threshold policies are obtained in [14] for optimal load-shifting with the objective of minimizing the time-averaged cost of buying electricity from the grid in the presence of a free renewable resource. However, none of these papers studies the case of correlated prices or robustness analysis. Recently, approaches based on a combination of Monte Carlo simulations and mixed-integer linear programming (MILP) for stochastic optimization, or robust optimization and MILP for hedging against price uncertainty have been developed for relatively detailed models of energy consumption [18]. The context and purpose

M. Roozbehani, D. Materassi, and M. Dahleh are with Laboratory for Information and Decision Systems, Massachusetts Institute of Technology. Emails: {mardavij, donnie13, dahleh}@mit.edu. M. Ohannessian is currently an ERCIM postdoctoral fellow at the Département de Mathématiques of the Université Paris-Sud.

of these detailed studies are clearly different from ours for the reasons that we already discussed above.

The contributions of this paper can be summarized as follows. We propose a mathematical model for load-shifting in response to a stochastic price process. Under the assumptions of *Markov contractivity in the mean* and *stochastic monotonicity*, we show that the optimal policy is to consume when the price falls below or is on par with a time-varying threshold that depends only on the statistics of the price process and the time left to the deadline. This analysis considers a price process that exhibits correlation in the time domain. From a practical perspective, it is a challenging task to determine an accurate statistical model for the price process from historical data. Indeed, while the first and second moments of the price process can be estimated in a relatively easy way, large quantities of data would be required to obtain a non-parametric Bayesian model. At the same time, the choice of an appropriate parametric model is not trivial. In order to avoid these difficulties, we derive policies that are robust with respect to partial information about the available price statistics. This partial information is represented in the form of inclusion within a family of distributions.

We begin with specializing our analysis to the case of a price process that is independent and identically distributed (i.i.d.) and then show how this result can be leveraged for deriving threshold policies that are robust against adversarial price processes. In other words, at each time step the price can be sampled from an adversarially selected distribution from within a family of possible distributions. As a special case, this includes the scenario where only the first and the second moments of the price process are known. For this case, we present explicit formulas for the computation of the thresholds associated with the robust policy as a function of the first two moments. We also use the structure of the optimal policy and the specific dependence of the optimal cost on the thresholds to obtain upper and lower bounds on the *economic value of load-shifting*, defined as the expected savings that result from optimal load-shifting compared with consuming energy on demand. Numerical experiments based on actual wholesale market prices are used to compare the various consumption strategies showing the beneficial effects of the robust policies.

In order to improve the clarity of the exposition and the flow of the results, all of the proofs are gathered in the Appendix.

II. A Dynamic Model of Consumption with Shiftable Demand

Each individual smart load has a predetermined demand for energy to be satisfied within a known time horizon consisting of n intervals, indexed by $k = 0, ..., n-1$. We denote the demand for the k-th interval by d_k, $k \in \{0, \cdots, n-1\}$. It is assumed that the vector $d = [d_0, \cdots, d_{n-1}]$ is known at the beginning of the optimization horizon. Each period's demand d_k is inelastic (independent of price) but shiftable and can be satisfied during any period $l \in \{k, \cdots, n-1\}$.

The backlogged demand at the k-th interval is denoted by $x_k \in (-\infty, 0]$. The state variable x_k is updated according to

$$x_{k+1} = x_k + u_k - d_k, \tag{1}$$

where, $u_k \geq 0$ is the amount of energy that is withdrawn from the grid during the k-th period. We enforce the deadline for meeting the total demand by setting $x_n = 0$ as the terminal condition.

In general, there is a disutility associated with backlogging the demand, which represents the inconvenience of deferring consumption. We characterize this disutility via a sequence of linear penalty functions defined by $p_k(x_k) = -p_k x_k$, with $p_k \in [0, \infty)$ for $k \in \{0, \cdots, n-1\}$.

The price is declared at the beginning of each interval and remains constant during that interval. We denote the price for the k-th interval by λ_k, and assume that the price process is an exogenous Markov process satisfying the two properties of *Markov contractivity in the mean* [19], [20], and *stochastic monotonicity* [21], [22] as defined below.

Definition 1 (Markov contractivity)**.** *Consider a Markov process $\{\lambda_k\}$. We say that $\{\lambda_k\}$ is* Markov-contracting in the mean, *if for some $0 \leq \alpha < 1$ and for all $\lambda^{(1)}$ and $\lambda^{(2)}$ we have*

$$\left|\mathbf{E}[\lambda_{k+1}|\lambda_k = \lambda^{(2)}] - \mathbf{E}[\lambda_{k+1}|\lambda_k = \lambda^{(1)}]\right| \leq \alpha|\lambda^{(2)} - \lambda^{(1)}|. \tag{2}$$

Definition 2 (Stochastic Monotonicity)**.** *A Markov process $\{\lambda_k\}$ is said to be* stochastically monotonic *if for all $\hat{\lambda} \in \mathbb{R}$ the following property holds*

$$\begin{aligned} &\lambda^{(1)} \leq \lambda^{(2)} \Rightarrow \\ &\qquad \mathbf{P}\{\lambda_{k+1} \leq \hat{\lambda}|\lambda_k = \lambda^{(1)}\} \geq \mathbf{P}\{\lambda_{k+1} \leq \hat{\lambda}|\lambda_k = \lambda^{(2)}\}. \end{aligned} \tag{3}$$

The interpretation of the Markov contractivity condition is that prices tend to correct (in the average sense) towards an equilibrium. The stochastic monotonicity condition describes a form of "continuity", or "stickiness", in the price dynamics: a low price at time k is more likely (than a high price) to lead to a low price at time $k+1$ and, similarly, a high price at time k is more likely (than a low price) to lead to a high price at time $k+1$. Both these conditions are relatively mild and represent simple forms of regularity in the price dynamics. Moreover, it is immediate that an independently and identically distributed (i.i.d.) stochastic process satisfies both Markov-contractivity (with $\alpha = 0$) and stochastic monotonicity (with equality holding in Equation (3)). Thus, the assumptions of Markov-contractivity and stochastic monotonicity encompass the frequently utilized i.i.d. assumption as a special case, and are, hence, less restrictive.

The Optimization-Based Model

The consumer's energy management problem can now be formulated as a finite-horizon dynamic programming problem as follows

$$\begin{aligned} \min \quad & \mathbf{E}\left[\sum_{k=0}^{n-1} \lambda_k u_k - p_k x_k\right] \qquad (4) \\ \text{s.t.} \quad & x_{k+1} = x_k + u_k - d_k, \qquad x_n = 0 \\ & x_k \leq 0, \\ & u_k \geq 0. \end{aligned}$$

where the price process $\{\lambda_k\}$ is a Markovian process satisfying (2) and (3).

III. Robust and Optimal Policies for Load-Shifting

A. Perfect Information about the Price Distribution

The first result that we present is the solution to the optimization problem (4).

Theorem 3 (Characterization of the Optimal Policy). *Consider the optimal load-shifting problem (4). Let $P_k(\lambda_{k+1} \mid \lambda_k)$ denote the conditional probability measure of λ_{k+1} given the price λ_k at the k-th interval. Define, for $k = 0, ..., n-1$, a sequence of functions $\psi_k : \mathbb{R} \mapsto \mathbb{R}$ as follows:*

$$\psi_n(\lambda) = \lambda, \quad \forall \lambda \in \mathbb{R}. \tag{5}$$

$$\psi_k(\lambda) = p_k + \int \min\{\theta, \psi_{k+1}(\theta)\} dP_k(\theta \mid \lambda), \quad \forall \lambda \in \mathbb{R}. \tag{6}$$

The following statements hold.

(i) *The optimal policy is a threshold policy characterized by:*

$$u_k^* = \begin{cases} 0 & \lambda_k > t_{k+1} \\ d_k - x_k & \lambda_k \le t_{k+1} \end{cases} \tag{7}$$

where $t_n = +\infty$, and each threshold t_i, $i \in \{0, ..., n-1\}$, is a solution to the equation

$$\psi_i(\lambda) = \lambda. \tag{8}$$

Furthermore, Equation (8) has a unique solution.

(ii) *The value function is affine in the backlog state x_k and nonlinear in the price, of the form:*

$$W_k(x_k, \lambda_k) = \eta_k(\lambda_k) - (p_k + \min\{\lambda_k, \psi_{k+1}(\lambda_k)\}) x_k, \tag{9}$$

where the sequence of functions $\eta_k(\cdot)$ is defined recursively as follows, $\forall \lambda \in \mathbb{R}$,

$$\eta_n(\lambda) = 0, \tag{10}$$

$$\eta_k(\lambda) = d_k \min\{\lambda, \psi_{k+1}(\lambda)\} + \int \eta_{k+1}(\theta) dP_k(\theta \mid \lambda). \tag{11}$$

The interpretation of Theorem 3 is as follows. The optimal policy for consumption is determined by a sequence of thresholds $(t_1, ..., t_n)$ that can be computed with the complete knowledge of price statistics. If the price λ_k falls strictly above the threshold t_{k+1}, we have $u_k^* = 0$, otherwise, $u_k^* = d_k - x_k$, meaning that the user consumes enough to meet both the standing demand for the period and the backlogged demand if any. The backward iteration step (5)–(6) has a straightforward economic interpretation once we assume that this threshold policy is optimal. Neglecting for simplicity the penalty term p_k, the function $\psi_{k+1}(\theta)$ represents the average price that will be paid if the threshold policy is applied in the time horizon $[k+1, n]$, given that the price in the k-th interval is θ. Equation (6), thus characterizes a choice that must be made between consumption at time k with price θ or load shifting at an average price $\psi_{k+1}(\theta)$, properly weighted according to the probability distribution $P_k(\theta \mid \lambda)$. The result of this computation provides the average price that it is going to be paid, i.e., $\psi_k(\lambda)$, by applying the threshold policy in the interval $[k, n]$ given that the price at time $k-1$ is λ.

The iterations (10)–(11) have a similar interpretation. The function $\eta_k(\lambda)$ represents the average cost of consumption that occurs in the interval $[k, n]$, given that the price at time k is λ. It is obtained by summing the average cost of consumption in the horizon $[k+1, n]$ given that the price at time k is λ, i.e. $\mathbf{E}[\eta_{k+1}(\theta \mid \lambda)]$, and the average cost incurred by consuming d_k at time k, i.e., $d_k \min\{\lambda, \psi_{k+1}(\lambda)\}$.

While these iterations require the function ψ to be computed for each time, it has to be stressed that they can be implemented relatively efficiently since ψ is involved only in an integral expression.

Next, we present a corollary that specializes Theorem 3 to the i.i.d. price scenario. This corollary will play a pivotal role in determining consumption strategies that are going to be adversarially robust. Before we proceed, we introduce the following definition.

Definition 4. *The modulated expectation (ME) function associated with a probability measure P is a concave function $\Gamma_P : \mathbb{R} \to \mathbb{R}_-$ defined according to*

$$\Gamma_P(x) = \int \min\{\theta - x, 0\} dP(\theta) \tag{12}$$

Corollary 5 (Optimal Policy for i.i.d. Prices). *Consider the optimal load-shifting problem (4) with an i.i.d. price process $\{\lambda_k\}$ defined by a probability measure P. Let $\Gamma_P(\cdot)$ be the associated ME as defined in (12). The following statements hold.*

(i) *The optimal policy is a threshold policy characterized by (7), with the corresponding thresholds computed recursively via the following equations:*

$$t_n = +\infty, \qquad t_k = p_k + t_{k+1} + \Gamma_P(t_{k+1}) \tag{13}$$

(ii) *The expected cost-to-go (expected value function) is affine in the backlog state, satisfying:*

$$V_k(x_k) = \int W_k(x_k, \theta) dP(\theta) = -t_k x_k + e_k, \tag{14}$$

where the constant terms e_k satisfy the following recursive equations:

$$e_n = 0, \quad e_k = e_{k+1} + d_k(t_{k+1} + \Gamma_P(t_{k+1})) \tag{15}$$

(iii) *Given $x_0 = 0$, the optimal expected cost is a demand-weighted sum of the differences between the stage thresholds and the marginal disutilities, i.e.,*

$$V_0(0) = \sum_{k=0}^{n-1} d_k(t_k - p_k) \tag{16}$$

Observe that, since $t_n = +\infty$, the computation of t_{n-1} is obtained by computing the following limit:

$$\lim_{t_{k+1} \to +\infty} p_k + t_{k+1} + \Gamma_P(t_{k+1}).$$

It is straightforward to prove that this limit always exists and is finite. Also note that from (13) we always have $t_{n-1} = \mathbf{E}[\lambda] + p_{n-1}$.

Remark 6. *It is straightforward to show that in the case of an i.i.d. price process, for non-interruptible loads with duration longer than one period, our results are readily applicable. The only decision is when to start consumption, and that will be determined based on a similar threshold policy. The only change is that the load's time to deadline must be adjusted,* *and thus, the thresholds are shifted by a number of periods that is equal to the load's duration. The same applies to the robust policies that are presented in the sequel.*

B. Partial Information about the Price Distribution

In this section we relax the *perfect information* assumption and propose threshold policies that are robust against an adversarial set of distributions. We also provide performance guarantees for the robust policies. The analysis is general and relies on lower and upper bounds of the ME function Γ_P. As a special case of this analysis we consider the case where only the mean and the variance of the price distribution is known, and present analytic expressions for the approximations.

Definition 7. *Given a sequence $\{t\} \stackrel{def}{=} (t_1, \cdots, t_n) \in \mathbb{R}^n$, we say that $\pi_t : \mathbb{R}^2 \mapsto \mathbb{R}$ is a threshold policy associated with $\{t\}$ if the control policy $u_k^* = \pi_t(x_k, \lambda_k)$ is of the form* (7) *with thresholds $t_1, \cdots, t_n$. We denote by J_{P,π_t} the expected cost under policy π_t and price distribution function P.*

Theorem 8 (Robust Performance Guarantees)**.** *Let $\mathcal{P}$ be a set of probability measures. Let functions $\overline{\Gamma}$, and $\underline{\Gamma}$ mapping $[0,\infty)$ to $(-\infty, 0]$ be such that for all $x \in [0,\infty)$, we have*

$$\underline{\Gamma}(x) \leq \Gamma_P(x) \leq \overline{\Gamma}(x), \quad \forall P \in \mathcal{P}. \tag{17}$$

Define the midmost *approximation function $\widehat{\Gamma} : [0,\infty) \mapsto (-\infty, 0]$ as follows:*

$$\widehat{\Gamma}(x) = \frac{\underline{\Gamma}(x) + \overline{\Gamma}(x)}{2}, \quad \forall x \in [0,\infty). \tag{18}$$

Let $\{\bar{t}\}$, $\{\underline{t}\}$, and $\{\hat{t}\}$, be sequences recursively generated according to (13) *with $\overline{\Gamma}$, $\underline{\Gamma}$, and $\widehat{\Gamma}$ respectively, and let $\overline{\pi}$, $\underline{\pi}$, and $\hat{\pi}$, be the associated threshold policies. Define:*

$$\overline{J} = \sum_{k=0}^{n-1} d_k(\bar{t}_k - p_k), \qquad \underline{J} = \sum_{k=0}^{n-1} d_k(\underline{t}_k - p_k),$$
$$\widehat{J} = \sum_{k=0}^{n-1} d_k(\hat{t}_k - p_k).$$

The following statements hold:

(i) *The threshold policy $\overline{\pi}$ is robust in the sense that*

$$\underline{J} \leq J_{P,\overline{\pi}} \leq \overline{J}, \quad \forall P \in \mathcal{P}. \tag{19}$$

(ii) *For any distribution $P \in \mathcal{P}$, the optimal cost is bounded between $\underline{J}$ and $\overline{J}$. That is,*

$$\underline{J} \leq J_P^* \stackrel{\text{def}}{=} \inf_{\pi} J_{P,\pi} \leq \overline{J}, \quad \forall P \in \mathcal{P}. \tag{20}$$

(iii) *The expected cost under the* midmost *threshold policy $\hat{\pi}$ is close to $\widehat{J}$ in the following sense:*

$$\left| J_{P,\hat{\pi}} - \widehat{J} \right| \leq \sum_{k=1}^{n-1} f(k) \left| \Gamma_P\left(\hat{t}_k\right) - \widehat{\Gamma}\left(\hat{t}_k\right) \right|, \quad \forall P \in \mathcal{P} \tag{21}$$

$$\leq \frac{1}{2} \sum_{k=1}^{n-1} f(k) \left| \overline{\Gamma}\left(\hat{t}_k\right) - \underline{\Gamma}\left(\hat{t}_k\right) \right|, \tag{22}$$

where

$$f(k) = \min\left\{ \frac{\max\{d_k\}}{F(\hat{t}_0)}, \ \sum_{i=0}^{k-1} d_i \right\}.$$

Remark 9. *Note that the results of Theorem 8 do not rely on the assumption that the price distribution is i.i.d. In particular, the policy $\overline{\pi}$ is robust adversarially, in the sense that at each stage, the price can be sampled from a different distribution that is chosen by an adversary. As long as $\overline{\Gamma}$ is an upper bound on the ME function of these distributions the policy is robust and the bounds are valid. In the case of correlated prices this means that the ME of all the conditional distributions must be upper bounded by $\overline{\Gamma}$.*

Theorem 10 (Bounding the ME function)**.** *Suppose that the price process has support over a bounded interval $[\lambda_{\min}, \lambda_{\max}] \subset \mathbb{R}$. Let $\lambda_{\min} = 0$ and $\lambda_{\max} = 1$. Given a mean $\mu \in [0,1]$ and an achievable variance σ^2, let $\mathcal{P}$ be the set of all distributions with support on $[0,1]$ that have mean μ and variance σ^2, and let $P \in \mathcal{P}$. Then, Γ_P can be bounded from above and below as follows:*

$$\underline{\Gamma}(x) \leq \Gamma_P(x) \leq \overline{\Gamma}(x), \quad \forall x \in [0,1], \tag{23}$$

where,

$$\overline{\Gamma}(x) = \begin{cases} 0 & , \quad x \in \left[0, \mu - \frac{\sigma^2}{1-\mu}\right] \\ (1-\mu)(\mu - x) - \sigma^2 & , \quad x \in \left[\mu - \frac{\sigma^2}{1-\mu}, \mu + \frac{\sigma^2}{\mu}\right] \\ \mu - x & , \quad x \in \left[\mu + \frac{\sigma^2}{\mu}, 1\right] \end{cases} \tag{24}$$

$$\underline{\Gamma}(x) = \begin{cases} \frac{-\sigma^2}{\sigma^2+\mu^2} x & , \quad x \in \left[0, \frac{\mu^2+\sigma^2}{2\mu}\right] \\ \frac{-\sigma^2\sqrt{(\mu-x)^2+\sigma^2}}{\sigma^2 + (\mu - x + \sqrt{(\mu-x)^2+\sigma^2})^2} & , \quad x \in \left[\frac{\mu^2+\sigma^2}{2\mu}, \frac{1-\mu^2-\sigma^2}{2(1-\mu)}\right] \\ \frac{-(1-\mu)^2(x-1)}{(1-\mu)^2+\sigma^2} + \mu - 1 & , \quad x \in \left[\frac{1-\mu^2-\sigma^2}{2(1-\mu)}, 1\right] \end{cases} \tag{25}$$

Furthermore, both of these bounds are tight point-wise, in the sense that for every $x \in [0,1]$ there exists a distribution $\overline{P} \in \mathcal{P}$ for which $\Gamma_{\overline{P}}(x) = \overline{\Gamma}(x)$ and another distribution $\underline{P} \in \mathcal{P}$ for which $\Gamma_{\underline{P}}(x) = \underline{\Gamma}(x)$. Finally, with a slight abuse of notation, the bounds for arbitrary $\lambda_{\min}$ and $\lambda_{\max}$ can be expressed via the following transformations:

$$\overline{\Gamma}(x; \lambda_{\min}, \lambda_{\max}, \mu, \sigma) = l\overline{\Gamma}\left(x_l; 0, 1, \mu_l, \sigma_l\right), \tag{26}$$
$$\underline{\Gamma}(x; \lambda_{\min}, \lambda_{\max}, \mu, \sigma) = l\underline{\Gamma}\left(x_l; 0, 1, \mu_l, \sigma_l\right), \tag{27}$$

where $l = \lambda_{\max} - \lambda_{\min}$, $x_l = (x - \lambda_{\min})/l$, $\mu_l = (\mu - \lambda_{\min})/l$, and $\sigma_l = \sigma/l$.

Consider the class of probability distributions with support over $[0,1]$, with mean $\mu = 1/2$ and variance $\sigma^2 = 1/12$. Figure 1(a), shows the ME function (Γ_P) for the special case of a uniform distribution, as well as the upper and lower bounds defined in Theorem 10, and the midmost approximation of equation (18) for all distributions within this class. Note that the midmost approximation closely approximates the uniform distribution.

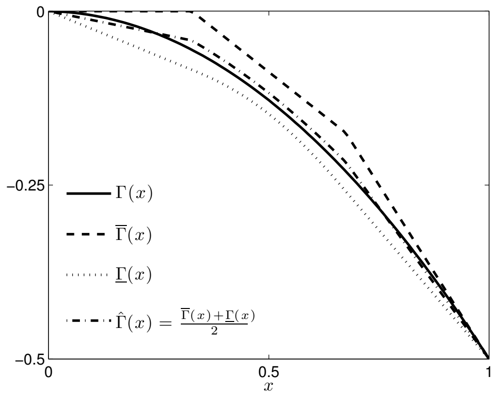

(a)

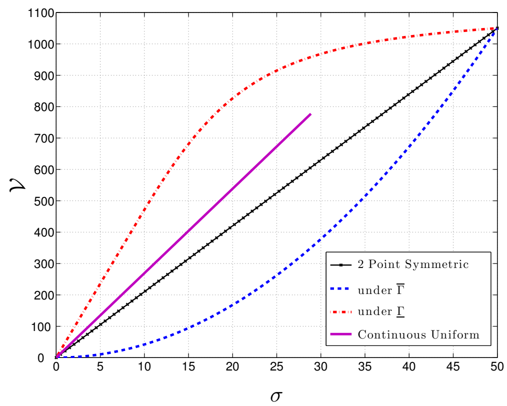

(b)

Fig. 1. (a) For distributions over $[0, 1]$, with mean $1/2$ and variance $1/12$, we plot the ME function of the uniform distribution, as well as partial information upper and lower bounds, and the midmost approximation. (b) Illustration of upper and lower bounds on the value of load-shifting as a function of the standard deviation, under all possible distributions with support over $[0, 100]$ and mean $\mu = 50$. The shifting horizon is $n = 24$, $d_k = 1$, and $p_k = 0$ for all k. Also shown are the value under 2-point symmetric distribution and continuous uniform distribution with the same mean $\mu = 50$.

In Section IV, we use the results of Theorems 8 and 10 to establish bounds of the cost of the optimal load-shifting problem (4). We apply the threshold policy (7) under various perfect and partial information structure assumptions studied in this section to actual wholesale electricity market data and discuss the simulation results.

IV. The Value of Load-Shifting

We define the *value of load-shifting* as the expected saving that results from adopting the optimal policy (7) compared with consuming electricity on demand. The formal definition is as follows:

$$\mathcal{V} = \sup_{\pi_t} \left\{ \mathbf{E}\left[\sum_{k=0}^{n-1} d_k \lambda_k \right] - J_{\pi_t} \right\}, \tag{28}$$

where J_{π_t} is the expected cost of problem (4) under a threshold policy π_t. For the case of i.i.d. price distribution with full information we can express (28) as $\mathcal{V} = \mathbf{E}\left[\sum_{k=0}^{n-1} d_k\lambda_k\right] - V_0(0)$, where $V_k(\cdot)$ is the expected value function defined in (14). Thus, we are benchmarking the performance of the optimal threshold policy (7) against the average cost in the absence of load-shifting. Intuitively, we expect that the value of load-shifting would increase with price volatility (measured by the variance). The following corollary establishes that when the average price is held constant, for all possible distributions the value of load-shifting is lower-bounded by a quadratic function of the standard deviation.

Corollary 11. *Let $\mathcal{V}$ be the value, as defined in* (28)*, of the load-shifting problem (4). Let $p = [p_0, \cdots, p_{n-1}]$ be the vector of disutilities and $d = [d_0, \cdots, d_{n-1}]$ be the vector of demand. Then, for all distributions (not necessarily i.i.d.) with fixed mean μ, support over a bounded subset $[\lambda_{\min}, \lambda_{\max}] \subset \mathbb{R}$, and variance $\sigma^2 \leq \sigma^2_{\max}$, where $\sigma^2_{\max}$ is the maximum achievable variance, the following statements hold:*

(i) *There exist functions $C(p)$ and $D(p, d)$ such that*

$$\sum_{k=0}^{n-1} d_k(\mu + p_k - \bar{t}_k) = C(p) + D(p, d)\sigma^2 \leq \mathcal{V} \leq \leq \sum_{k=0}^{n-1} d_k(\mu + p_k - \underline{t}_k) \tag{29}$$

where $\{\underline{t}_k\}$ and $\{\bar{t}_k\}$ are recursively generated as in (13) *with $\underline{\Gamma}$, and $\overline{\Gamma}$ given by* (24) *and* (25) *respectively. Furthermore, $C(0) = 0$.*

(ii) *Regardless of the the information structure about the price process, both bounds in* (29) *are tight for $n = 2$.*

Figure 1(b) shows the upper and lower bounds on $\mathcal{V}$ as the standard deviation σ varies from 0 to the maximum achievable $\overline{\sigma} = 50$ for the case where $\lambda_{\min} = 0$, $\lambda_{\max} = 100$, and the mean is $\mu = 50$.

Numerical experiments

In order to evaluate the implications of our assumptions—particularly about the stochastic price process—on performance, we have simulated implementations of the threshold policy under different assumptions, against actual wholesale electricity market data. We have considered hourly averaged real-time Locational Marginal Prices (LMP's) in the PJM (Pennsylvania-Jersey-Maryland) Interconnection for 676 weeks during the period 1/10/1999 to 12/31/2011 [23], and have adjusted them for inflation according to the monthly Consumer Price Index (CPI). Since energy consumption is not homogeneous during the span of a day, we have taken into account only data in a 16-hour lapse (08:00-24:00) that manifest a better level of uniformity. Scatter plots of the price at a certain hour versus the price at the next hour are shown in Figure 2(a) for four different times of the day.

These scatter plots highlight both non-stationarity across the hours of the day in the data and also an observable degree of correlation in the time series. In addition to this short term correlation, data exhibit also correlation at other time-scales, for example across various weeks of the year because of seasonal and long term phenomena.

Given an empirical probability distribution (obtained from a properly normalized, uniformly spaced histogram of the price data) corresponding to any given time window, it is immediate to determine the associated (empirical) ME function Γ_P using (12) (see Figure 2(b)). The same procedure has been followed to compute 13 empirical estimates of the ME function associated with sub-periods of 52 weeks each. The 13 yearly estimates have been used to determine a "robust" piecewise linear upper bound for $\Gamma_P(\cdot)$, as illustrated in Figure 2(b). The conditional probability distributions are estimated in a similar way. The price range has been divided into identically spaced intervals and a two-dimensional histogram of the realized pairs $(\lambda_{k-1}, \lambda_k)$ has been created. The normalized histogram has been taken as the estimate of the conditional probability distributions.

In this experiment, we have considered smart loads that have a single unit of energy demand at the beginning of the optimization horizon and no demand during the other time slots. For different time horizons (from 1 to 16 hours) we have implemented and compared the following strategies

- Immediate consumption on demand (no load-shifting).
- Corollary 5 (i.i.d. model) with empirical Γ_P estimated from the whole data set.
- Corollary 5 (i.i.d. model) with an empirical Γ_P estimated from a rolling time window of 4 weeks leading to consumption time.
- The (robust) threshold policy of Theorem 8 with empirical piecewise linear bound $\overline{\Gamma}$ (see Figure 2(b)).
- The (robust) threshold policy of Theorem 8 with the "tighter" bound $\overline{\Gamma}$ from Theorem 10 with empirical mean, variance, minimum and maximum prices estimated from a rolling time window of 4 weeks leading to consumption time.
- The (midmost) threshold policy of Theorem 8 with $\hat{\Gamma}$ given by the average of the two functions $\overline{\Gamma}$ and $\underline{\Gamma}$ as provided by Theorem 10 with empirical mean, variance minimum and maximum prices estimated from a rolling time window of 4 weeks leading to consumption time.
- Theorem 3 (Markov model) with conditional distributions estimated from the whole data set.
- A procedure based on Certainty Equivalent Model Predictive Control (CE-MPC) (see, e.g., [24]) where future prices have been assumed equal to their conditional expectation given the current price using the empirical conditional distribution obtained from the whole data set.

In order to test these different strategies, we have run series of Monte Carlo simulations for 16 different optimization horizons sampling random initial times in the considered period of 676 weeks. From the data we have obtained an empirical probability distribution of price sequences that has been assumed exact.

For each optimization horizon we have considered 10000 price sequences, so, considering all time horizons, each policy has been tested over more than one hundred thousand times. The initial time is sampled randomly in such a way that the optimization horizon never crosses midnight. In other words, only contiguous data in the interval $8-24$ are used. The initial times have been drawn with resampling from the empirical distribution (thus the same sequence might have been used multiple times). Resampling is a quite common approach when dealing with empirical data, especially for the evaluation of the variability of data [25]. In our case, we have just resampled data in order to obtain smoother curves (indeed the same analysis without resampling has led to identical conclusions). In all tests the disutility of delay has been set to zero.

We have considered different metrics to evaluate the strategies. In Figure 3(a), we report the mean cost of consumption

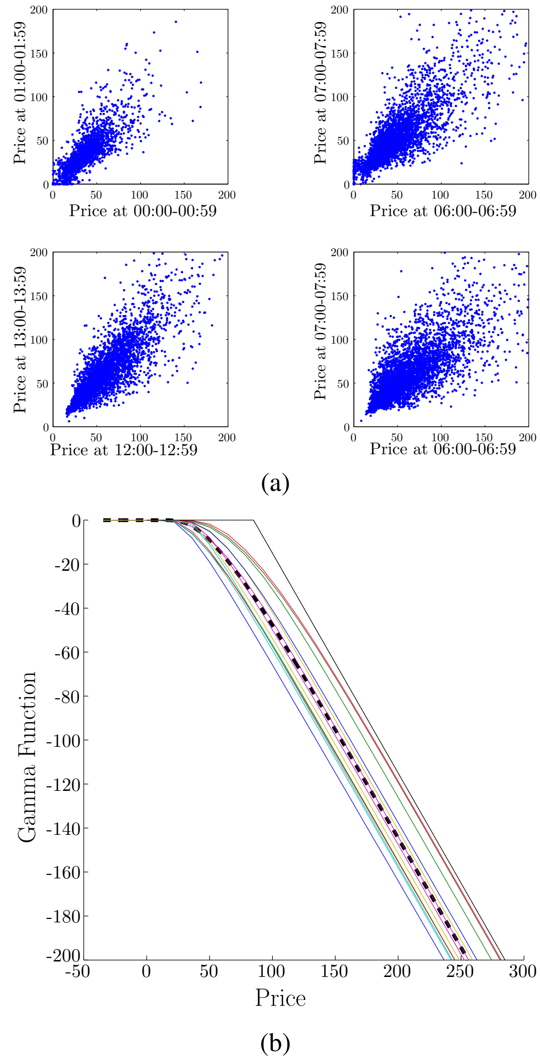

Fig. 2. (a) Scatter plot of price versus price at the next hour at four different times of the day. (b) Γ_P obtained using the entire set of data (black, thick, dashed), Γ_P obtained using sub-periods of 52 weeks each (color, thin, solid) and a piecewise upper bound on Γ_P (black, thick, solid) which can be regarded as an empirical upper bound on the $\Gamma(\cdot)$ functions. This is a piecewise linear function with one breakpoint only. The breakpoint and the slopes are chosen such that all the empirical $\Gamma(\cdot)$ functions fall below this piecewise linear approximation. Note that it is straightforward to construct tighter piecewise linear upper bounds with more breakpoints to improve accuracy.

for each strategy as a function of the time horizon.

Consumption on demand gives a constant expected cost that does not depend on the horizon. All load shifting strategies perform remarkably better. When the entire data set is used for estimating the probability distributions, the strategy based on the Markovian assumption provides only a slightly lower average cost than the strategy based on the i.i.d. assumption. In contrast, due to non-stationarity, strategies that use only a rolling time window of 4 weeks lead to significantly higher savings. The policy that provides the lowest average cost is the the robust one based on Theorem 10.

While the proposed strategies result in a lower average cost than consuming on demand, for certain price sample paths, load shifting may actually result in a higher cost. Indeed, shifting introduces a risk component in the total cost of consumption. In Figure 3(b), we report the probability of

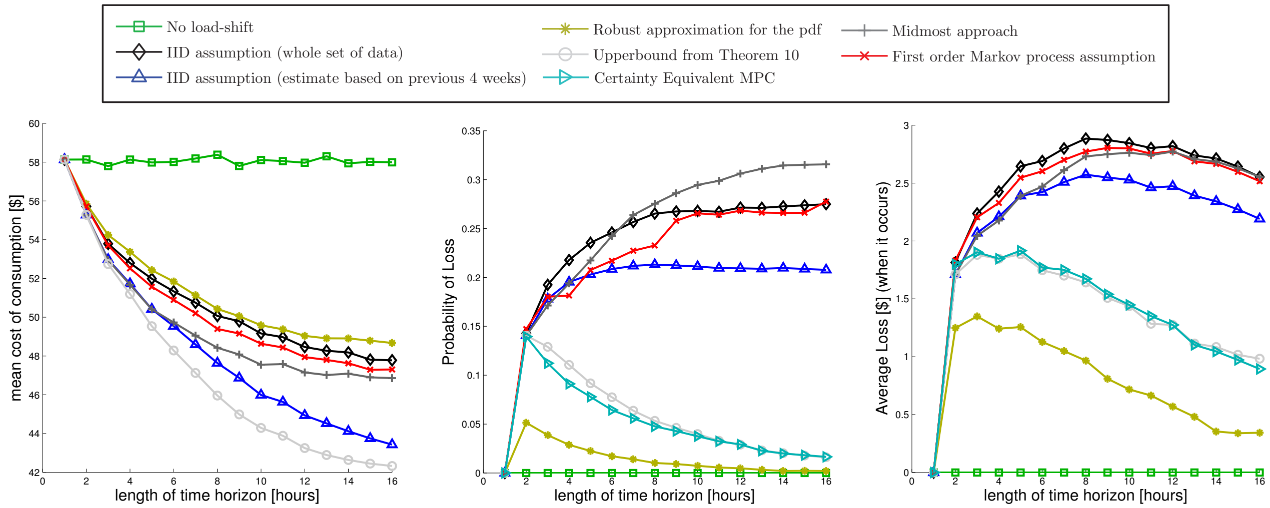

(a) (b) (c)

Fig. 3. Performance and measure of risk associated with different implementations of the threshold policy under different assumptions about the price process. The square-marked line represents consumption on demand and is reported for comparison. The diamond-marked and upper-triangle-marked lines are associated with the implementation of the optimal policy under the assumption of i.i.d. price. For the diamond-marked lines the probability distribution of the price is obtained using the whole set of data, while for the upper-triangle-marked ones the probability distribution is estimated using only the four weeks preceding the consumption horizon. The right-triangle-marked curves are the results of the implementation of the Certainty Equivalent Model Predictive Control approach. The asterisk-marked, circle-marked and plus-marked lines are associated with robust implementations of the policy, again under the assumption of an i.i.d. process. They correspond to the computations of the thresholds adopting, respectively, the piecewise upper bound of the function Γ depicted in Figure 2(b), the function $\overline{\Gamma}$ as given by Theorem 10 and the function $\hat{\Gamma}$ obtained as the mean of the two functions $\overline{\Gamma}$ and $\underline{\Gamma}$ as given by Theorem 10. The cross-marked lines are associated with the implementation of the optimal strategy estimating the conditional probability distribution of the price, given the price at the previous interval. (a) Average cost (b) Probability of loss (incurring in a higher cost than consumption on demand) (c) Average loss when it occurs.

suffering a loss when a load shifting strategy is adopted (the benchmark is consuming on demand), and in Figure 3(c), we report the average loss incurred in those cases. The graphs in Figure 3(b) and in Figure 3(c) illustrate measures of the risks associated with different load shifting strategies. We observe that the strategies based on upper bounding the ME function keep both the probability of loss and the average loss at low levels. Interestingly, the CE-MPC policy shows similar risk to the robust policy (whether it is measured as probability of loss or average loss). The reasons behind this observations may deserve further investigation in future work. Nevertheless, the CE-MPC policy has a higher average cost and is outperformed by the robust policy. Numerical approaches based on MPC for integration of heterogeneous resources and automated demand response are popular, and often applicable to relatively detailed models of energy networks, see, e.g., [26] for simulation models and a case study.

Note that it is not possible to test the Markov strategy using most recent price data leading to consumption time because significantly more data is needed for constructing empirical conditional distributions. A different approach (e.g., parametric modeling) is needed for estimating the conditional distributions.

Remark 12. *In all these simulations the disutility of delay was set to zero. We can explain how these results would change with increased disutility. As the disutility increases, the thresholds increase for all presented policies. This means that in Figure 3(a), all curves would be shifted upwards towards the green curve that represents the consumption on demand or "no load shifting" policy. The curvature or convexity of these curves would also decrease as disutility increases and in the limit, the curves approach the straight line that corresponds to consumption on demand. The qualitative effect of increased disutility on Figure 3(b) and Figure 33(c) is similar. Higher disutility means higher thresholds, less shifting, lower probability of loss and lower average loss. All curves in Figure 3(b) and Figure 33(c) approach the straight line green curve representing the no load shifting policy as disutility increases. As a result, the sensitivity of the savings, or the losses to the adopted policy decreases, as all policies approach the no shifting policy.*

Remark 13. *An attractive feature of our approach is the simplicity of the proposed threshold policies, which do not demand the online solution of an optimization problem. The basic step required for the backward iterations amounts to the computation of an integral under a pre-specified probability measure. While given the current device standards these methodologies might not be deployment-ready, there are several ways in which they could be implemented in practice. If demand response technologies are going to be widely adopted in the future, it would be natural to see the introduction of more sophisticated devices with higher computational power. Alternatively, specific hardware implementations could be used to make these computations more efficient since numerical integration would be the only basic operation that is required. Finally, a third possibility is based on centralized computations: a central entity can compute the thresholds and communicate them to smart devices.*

V. Conclusions and Future Work

We proposed a dynamic model of consumption in response to time-varying electricity prices, and derived optimal policies for load-shifting when the price process is an exogenous

Markovian process satisfying assumptions of Markov contractivity and stochastic monotonicity. As a special case, this includes the case where the price is an i.i.d. process. We have leveraged the solution of the i.i.d. case to derive robust threshold policies with performance guarantees for the case where, at each time step, the price is sampled from an adversarially selected distribution among a family of distribution functions. We have compared these different policies via numerical experiments using actual wholesale market prices. In particular we have considered the performance of these policies in terms of two metrics: expected reduced cost compared with consumption on-demand, and expected loss compared with consumption on-demand given that a loss has occurred. The first metric represents how on average a consumer will save by adopting the policy while the second one can be understood as a surrogate metric for the risk associated with these policies. We have shown that the robust policies offer a good compromise between performance and volatility without requiring detailed information about the price process. Future works include extending these results to the case with finite consumption constraints on individual loads, and/or aggregate congestion constraints on the consumption of several loads.

REFERENCES

[1] K. Afridi, A. M. Farid, J. M. Grochow, W. W. Hogan, H. D. Jacoby, J. L. Kirtley, H. G. Michaels, I. Perez-Arriaga, D. J. Perreault, N. L. Rose, and G. L. Wilson, *The Future of the Electric Grid*. MIT Press, 2011.

[2] 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.

[3] K. Vanthournout, R. D'hulst, D. Geysen, and G. Jacobs, "A smart domestic hot water buffer," *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 2121–2127, 2012.

[4] D. Caves, K. Eakin, and A. Faruqui, "Mitigating price spikes in wholesale markets through market-based pricing in retail markets," *The Electricity Journal*, vol. 13, no. 3, pp. 13–23, 2000.

[5] D. Kirschen, "Demand-side view of electricity markets," *Power Systems, IEEE Transactions on*, vol. 18, no. 2, pp. 520–527, 2003.

[6] P. H. Zipkin, *Foundations of Inventory Management*. McGraw-Hill, Boston, MA, 2000.

[7] M. Sobel and R. Zhang, "Inventory policies for systems with stochastic and deterministic demand," *Operations research*, vol. 49, no. 1, pp. 157–162, 2001.

[8] K. C. Frank, R. Q. Zhang, and I. Duenyas, "Optimal policies for inventory systems with priority demand classes," *Operations Research*, vol. 51, pp. 993–100, 2003.

[9] P. Constantopoulos, F. Schweppe, and R. Larson, "Estia: a real-time consumer control scheme for space conditioning usage under spot electricity pricing," *Computers & operations research*, vol. 18, no. 8, pp. 751–765, 1991.

[10] J. Black, "Integrating demand into the us electric power system: technical, economic, and regulatory frameworks for responsive load," Ph.D. dissertation, Massachusetts Institute of Technology, 2005.

[11] J. Kueck, B. Kirby, T. Laughner, and K. Morris, "Spinning reserve from responsive load," no. Technical report, Oak Ridge National Laboratory (ORNL), 2009.

[12] D. Livengood and R. Larson, "The energy box: Locally automated optimal control of residential electricity usage," *Service Science*, vol. 1, no. 1, pp. 1–16, 2009.

[13] A. Mohsenian-Rad and A. Leon-Garcia, "Optimal residential load control with price prediction in real-time electricity pricing environments," *IEEE Transactions on Smart Grid*, vol. 1, pp. 120–133, 2010.

[14] 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*. IEEE, 2010.

[15] H. Zhang, X. Yu, and P. V. Braun, "Three-dimensional bicontinuous ultrafast-charge and-discharge bulk battery electrodes," *Nature Nanotechnology*, vol. 6, no. 5, pp. 277–281, 2011.

[16] W. Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin, "Aggregate model for heterogeneous thermostatically controlled loads with demand response," in *Power and Energy Society General Meeting, 2012 IEEE*. IEEE, 2012, pp. 1–8.

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

[18] Z. Chen, L. Wu, and Y. Fu, "Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization," *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1822–1831, 2012.

[19] E. Altman, "Constrained markov decision processes with total cost criteria: Lagrangian approach and dual linear program," *Mathematical methods of operations research*, vol. 48, no. 3, pp. 387–417, 1998.

[20] P. Bertail, P. Doukhan, and P. Soulier, *Dependence in probability and statistics*. Springer Science+ Business Media, 2006, vol. 187.

[21] I. Herbst and L. Pitt, "Diffusion equation techniques in stochastic monotonicity and positive correlations," *Probability theory and related fields*, vol. 87, no. 3, pp. 275–312, 1991.

[22] S. Lee, O. Linton, and Y.-J. Whang, "Testing for stochastic monotonicity," *Econometrica*, vol. 77, no. 2, pp. 585–602, 2009.

[23] The PJM Independent System Operator (PJM-ISO), "http://www.pjm.com/pub/account/lmpmonthly/index.html."

[24] R. Soeterboek, *Predictive control: a unified approach*. Prentice-Hall, Inc., 1992.

[25] B. Efron, *The jackknife, the bootstrap and other resampling plans*. SIAM, 1982, vol. 38.

[26] I. Lampropoulos, N. Baghina, W. Kling, and P. Ribeiro, "A predictive control scheme for real-time demand response applications," *Smart Grid, IEEE Transactions on*, vol. 4, no. 4, pp. 2049–2060, 2013.

[27] T. Tao, *An introduction to measure theory*. AMS Bookstore, 2011, vol. 126.

APPENDIX

Proof of Theorem 3:
We need to introduce three lemmas before presenting the proof of Theorem 3. The first lemma provides a decomposition of a probability measure P as the sum of n measures $\{P_1, ..., P_n\}$ defined on n intervals. Interested readers can consult [27] for introductory concepts in measure theory.

Lemma 14. *Let P be a probability measure on $\mathbf{R}$. For any $n \in \mathbf{N} \setminus \{0\}$, it is possible to find a covering of the real line $\{I_1 = (-\infty, b_1], I_2 = [a_2, b_2], ..., I_k = [a_k, b_k], \dots I_n = [a_n, +\infty)\}$ and n measures $P_1, ..., P_n$ with the following properties*

- $a_{k+1} = b_k$ *for* $k = 1, ..., n-1$
- *for every measurable set $A \subseteq \mathbf{R}$ we have* $\int_A dP = \sum_{k=1}^n \int_A dP_k$
- $\int_{I_i} dP_j = \delta_{ij}/n$

where $\delta_{ij} = 1$ if $i = j$, and $\delta_{ij} = 0$ otherwise.

Proof: Fixed n, for any $k = 1, ..., n-1$, define

$$b_k := \sup_q \left\{ \int_{-\infty}^q dP \leq \frac{k}{n} \right\} \tag{30}$$

$$a_{k+1} := b_k. \tag{31}$$

For any set $A \subseteq \mathbf{R}$, define the measures $\{Q_1, ..., Q_n\}$ in the following manner

$$\int_A dQ_k := \int_{A \cap I_k} dP \tag{32}$$

Defining

$$dP_k := dQ_k + \left(\frac{k}{n} - \int dQ_k \right) \delta_{b_k} \tag{33}$$

where δ_z is the measure associated with the Dirac delta centered in z, the properties in the lemma statement are immediately verified. ■

The following lemma provides conditions on probability measures preserving, in some sense, the monotonicity and the Lipschitz property of a function.

Lemma 15. *Consider two probability measures $dP^{(1)}$ and $dP^{(2)}$ on the real line, such that, for every q,*

$$\int_{-\infty}^{q} dP^{(1)} \leq \int_{-\infty}^{q} dP^{(2)}. \tag{34}$$

Consider a monotonic non decreasing function $\phi : \mathbf{R} \to \mathbf{R}$ with Lipschitz constant less or equal to 1. Then, we have that

$$0 \leq \int \phi(\theta) dP^{(2)}(\theta) - \int \phi(\theta) dP^{(1)}(\theta) \leq \tag{35}$$

$$\leq \int \theta dP^{(2)}(\theta) - \int \theta dP^{(1)}(\theta) \tag{36}$$

Proof: Fix a natural number $n > 0$ and decompose $P^{(1)}$ and $P^{(2)}$ as indicated in Lemma 14

$$P^{(1)} = \sum_{k=1}^{n} P_k^{(1)} \qquad P^{(2)} = \sum_{k=1}^{n} P_k^{(2)}. \tag{37}$$

Also, let the two coverings be $\{I_1^{(1)} = (-\infty, b_1^{(1)}], I_2^{(1)} = [a_2^{(1)}, b_2^{(1)}], ..., I_k^{(1)} = [a_k^{(1)}, b_k^{(1)}], I_n^{(1)} = [a_n^{(1)}, +\infty)\}$ $\{I_1^{(2)} = (-\infty, b_1^{(2)}], I_2^{(2)} = [a_2^{(2)}, b_2^{(2)}], ..., I_k^{(2)} = [a_2^{(2)}, b_2^{(2)}], I_n^{(2)} = [a_n^{(2)}, +\infty)\}$, respectively. Since $\phi(\cdot)$ is monotonic and has Lipschitz constant $L \leq 1$, it satisfies

$$\theta_1 \leq \theta_2 \Rightarrow 0 \leq \phi(\theta_2) - \phi(\theta_1) \leq \theta_2 - \theta_1. \tag{38}$$

Observe that

$$\int \phi(\theta) dP^{(1)}(\theta) = \sum_{k=1}^{n} \int_{I_k^{(1)}} \phi(\theta) dP^{(1)}(\theta) \tag{39}$$

$$= \frac{1}{n} \sum_{k=1}^{n-1} \phi(b_k^{(1)}) + \varepsilon_n^{(1)} = \frac{1}{n} \sum_{k=2}^{n} \phi(a_k^{(1)}) + \varepsilon_n^{(1)} \tag{40}$$

and analogously

$$\int \phi(\theta) dP^{(2)}(\theta) = \frac{1}{n} \sum_{k=2}^{n} \phi(a_k^{(2)}) + \varepsilon_n^{(2)}. \tag{41}$$

where, for the monotonicity of ϕ, $\varepsilon_n^{(1)}$ and $\varepsilon_n^{(2)}$ are two quantities converging to 0 for n that goes to $+\infty$. By contradiction, assume that there exists $\alpha > 0$ such that

$$\int \phi(\theta) dP^{(2)}(\theta) - \int \phi(\theta) dP^{(1)}(\theta) < -\alpha. \tag{42}$$

Then, by observing that $a_k^{(1)} \geq a_k^{(2)}$, we have that

$$\varepsilon_n^{(2)} - \varepsilon_n^{(1)} < -\alpha \tag{43}$$

contradicting the fact the two quantities converge to 0 for $n \to +\infty$. Analogously, consider

$$\int \phi(\theta) dP^{(2)} - \int \phi(\theta) dP^{(1)} = \tag{44}$$

$$= \frac{1}{n} \sum_{k=1}^{n-1} (\phi(b_k^{(2)}) - \phi(b_k^{(1)})) + \varepsilon_n^{(2)} - \varepsilon_n^{(1)} \leq \tag{45}$$

$$\leq \varepsilon_n^{(2)} - \varepsilon_n^{(1)} + \frac{1}{n} \sum_{k=1}^{n-1} (b_k^{(2)} - b_k^{(1)}) \tag{46}$$

where the last inequality follows form the fact that the Lipschitz constant of ϕ is less of equal to 1. We obtain the assertion of the lemma by letting $n \to +\infty$. ■

Lemma 16. *Consider a sequence of classes of probability measures $\{P_k(\theta|\lambda)\}$, for $k = 1, ..., n$, parameterized by a scalar λ. Consider the iterations for $k = 0, ..., n-1$.*

$$\psi_k(\lambda) = p_k + \int \min\{\theta, \psi_{k+1}(\theta)\} dP_{k+1}(\theta|\lambda) \tag{47}$$

where p_k are fixed real positive numbers and the terminal condition is $\psi_n(\lambda) = \lambda$. Assume that the map $\lambda \to \int \theta dP_k(\theta|\lambda)$ is a contraction and that the measures $dP_k(\theta|\lambda)$ are stochastically monotonic in λ. Then we have that

- *$\psi_k(\lambda)$ is monotonic and Lipschitz with constant $L < 1$ for any $k \in \{0, ..., n-1\}$.*

Proof: Define $\phi_{k+1}(\theta) := \min\{\theta, \psi_{k+1}(\theta)\}$ and observe that if $\psi_{k+1}(\theta)$ is monotonic non decreasing with Lipschitz constant less than or equal to 1, also $\phi_{k+1}(\theta)$ is monotonic non decreasing with Lipschitz constant less than or equal to 1. From Lemma 15 we have, for $\lambda_1 < \lambda_2$,

$$0 \leq \psi_k(\lambda_2) - \psi_k(\lambda_1) = E[\phi_{k+1}(\theta)|\lambda_2] - E[\phi_{k+1}(\theta)|\lambda_1] \leq \\ \leq E[\theta|\lambda_2] - E[\theta|\lambda_1] < \alpha(\lambda_2 - \lambda_1)$$

where the last inequality is given by the contraction property. ■

We are now ready to present the proof of Theorem 3.

Proof of Theorem 3 : The Bellman equation corresponding to the underlying dynamic programming problem associated with (4) is

$$W_k(x_k, \lambda_k) = \\ \min_{u_k} \{-p_k x_k + \lambda_k u_k + E[W_{k+1}(x_{k+1}, \lambda_{k+1})|x_k, \lambda_k]\}.$$

Define the Bellman function as

$$W_k(x_k, \lambda_k) = -p_k x_k - \min\{\lambda_k, \psi_{k+1}(\lambda_k)\} x_k + \eta_k(\lambda_k). \tag{48}$$

Check by inspection that it satisfies the Bellman equation

$$\begin{aligned}
&W_k(x_k, \lambda_k) \\
&= \min_{0 \leq u_k \leq d_k - x_k} \{-p_k x_k + \lambda_k u_k + E[W_{k+1}(x_{k+1}, \lambda_{k+1})|x_k, \lambda_k]\} \\
&= \min_{0 \leq u_k \leq d_k - x_k} \{-p_k x_k + \lambda_k u_k + \\
&-\mathbf{E}\left[(\min\{\lambda_{k+1}, \psi_{k+2}(\lambda_{k+1})\} + p_{k+1}) x_{k+1} - \eta_{k+2}(\lambda_{k+1})|x_k, \lambda_k\right]\} \\
&= \min_{0 \leq u_k \leq d_k - x_k} \{-p_k x_k + \lambda_k u_k + \mathbf{E}[\eta_{k+2}(\lambda_{k+1})|\lambda_k] + \\
&-\mathbf{E}\left[(\min\{\lambda_{k+1}, \psi_{k+2}(\lambda_{k+1})\} + p_{k+1}) (x_k + u_k - d_k)|x_k, \lambda_k\right]\} \\
&= -p_k x_k + E[\eta_{k+2}(\lambda_{k+1})|\lambda_k] + \\
&\quad + \min_{0 \leq u_k \leq d_k - x_k} \{\lambda_k u_k - \psi_{k+1}(\lambda_k)(x_k + u_k - d_k)\} \\
&= -p_k x_k + E[\eta_{k+2}(\lambda_{k+1})|\lambda_k] + \min\{\lambda_k, \psi_{k+1}(\lambda_k)\}(d_k - x_k) \\
&= E[\eta_{k+2}(\lambda_{k+1})|\lambda_k] + \min\{\lambda_k, \psi_{k+1}(\lambda_k)\} d_k \\
&\quad - (p_k + \min\{\lambda_k, \psi_{k+1}(\lambda_k)\}) x_k.
\end{aligned}$$

Observe that

$$\min_{0 \leq u_k \leq d_k - x_k} \{\lambda_k u_k - \psi_{k+1}(\lambda_k)(x_k + u_k - d_k)\} = \\ = \min\{\lambda_k, \psi_{k+1}(\lambda_k)\}(d_k - x_k)$$

is obtained for

$$u_k = \begin{cases} 0 & \lambda_k > t_{n-k-1} \\ d_k - x_k & \lambda_k \leq t_{n-k-1} \end{cases} \tag{49}$$

where t_{k+1} is the solution of $\psi_{k+1}(\lambda_k) = \lambda_k$. The solution is unique because of Lemma 16. Indeed since $\psi_{k+1}(\lambda_k)$ has Lipschitz constant smaller than 1 it admits a unique intersection with a line with slope qual to 1. ■

Proof of Corollary 5: The proof follows from Theorem 3. Consider the iterations (5). Since the price is an i.i.d. process the function $\psi_k(\lambda)$, for $k = 0, ..., n-1$, is constant and equal to the threshold t_k. Furthermore we have

$$\psi_k = p_k + \int_{-\infty}^{+\infty} \min\{\theta, \psi_{k+1}\} dP(\theta) \tag{50}$$

$$= p_k + \psi_{k+1} + \int_{-\infty}^{+\infty} (\min\{\theta, \psi_{k+1}\} - \psi_{k+1}) dP(\theta) \tag{51}$$

$$= p_k + \psi_{k+1} + \Gamma(\psi_{k+1}). \tag{52}$$

Analogously, in the i.i.d case, the difference equation defining η_k becomes

$$\eta_k(\lambda) = \eta_{k+1}(\lambda) + \min\{\lambda, \psi_{k+1}\} d_k. \tag{53}$$

As a consequence, we have

$$e_k = \mathbf{E}[\eta_k(\theta)] = e_{k+1} + (t_{k+1} + \Gamma_P(t_{k+1})) d_k. \tag{54}$$

The computation of $V_0(0)$ is straightforward from the first two results of the theorem. ■

Proof of Theorem 8: Before we proceed with the proof of Theorem 8, we introduce a proposition and a lemma.

Proposition 17. *Let the function $\Upsilon : \mathbb{R}^2 \mapsto \mathbb{R}$ be defined as follows*

$$\Upsilon(\tau, \gamma) = \int_{\lambda_{\min}}^{\tau} (\theta - \gamma) \, dP(\theta) \tag{55}$$

For any given sequence $t = (t_1, \cdots, t_n) \in \mathbb{R}^n$, with $t_n = \lambda_{\max}$, the expected cost-to-go of problem (4) under the threshold policy π_t is affine, of the form

$$V_k^t(x_k) = -g_k x_k + h_k, \quad \forall k \in \{0, \cdots, n-1\}, \tag{56}$$

where g_k and h_k satisfy the following recursive equations:

$$g_n = \lambda_{\max}, \;\; h_n = 0, \tag{57}$$

$$g_k = p_k + g_{k+1} + \Upsilon(t_{k+1}, g_{k+1}) \tag{58}$$

$$h_k = h_{k+1} + d_k (g_{k+1} + \Upsilon(t_{k+1}, g_{k+1})) \tag{59}$$

Proof of Proposition 17 : The proof is by backward induction. Since $t_n = \lambda_{\max}$, we have $u_{n-1} = d_{n-1} - x_{n-1}$ and $V_{n-1}^t(x_{n-1}) = -p_{n-1}x_{n-1} - \overline{\lambda} x_{n-1} + \overline{\lambda} d_{n-1}$, where $\overline{\lambda}$ is the expected value of the price. It follows from (57) – (59) that $g_{n-1} = \overline{\lambda} + p_{n-1}$ and $h_{n-1} = \overline{\lambda} d_{n-1}$, and therefore, (56) holds for $k = n-1$. Suppose that $V_{k+1}^t(\cdot)$ is of the form (56) for some $k \leq n-2$. Then, we have:

$$\begin{aligned}
&J_k(x_k, \lambda_k) \\
&\quad \overset{\text{def}}{=} \min \mathbf{E}\left[-p_k x_k + \lambda_k u_k + \sum_{i=k+1}^{n-1} -p_i x_i + \lambda_i u_i \,\middle|\, x_k, \lambda_k \right] \\
&\quad = \min_{u_k} \left\{ -p_k x_k + \lambda_k u_k + V_{k+1}^t(x_{k+1}) \right\} \\
&\quad = \min_{u_k} \left\{ -p_k x_k + \lambda_k u_k - g_{k+1}(x_k + u_k - d_k) + h_{k+1} \right\} \\
&\quad = h_{k+1} - p_k x_k + \begin{cases} g_{k+1}(d_k - x_k), & \lambda_k > t_{k+1} \\ \lambda_k (d_k - x_k), & \lambda_k \leq t_{k+1}. \end{cases}
\end{aligned}$$

Thus,

$$\begin{aligned}
V_k^t(x_k) &\overset{\text{def}}{=} \mathbf{E}[J_k(x_k, \lambda_k)] \\
&= h_{k+1} - p_k x_k + \\
&\quad (d_k - x_k)\left[g_{k+1}(1 - F(t_{k+1})) + \mathbf{E}[\lambda_k | \lambda_k \leq t_{k+1}] F(t_{k+1}) \right] \\
&= h_{k+1} - p_k x_k + (d_k - x_k)\left[g_{k+1} + \int_{\lambda_{\min}}^{t_{k+1}} (\theta - g_{k+1}) \, dP(\theta) \right] \\
&= h_{k+1} + d_k (g_{k+1} + \Upsilon(t_{k+1}, g_{k+1})) + \\
&\qquad - x_k (p_k + g_{k+1} + \Upsilon(t_{k+1}, g_{k+1})) \\
&\overset{\text{def}}{=} h_k - g_k x_k
\end{aligned}$$

Proof is complete by induction. ■

Lemma 18. *The function $G(x) \overset{def}{=} x + \Gamma_P(x)$ where Γ_P is a modulated expectation function is non-decreasing.*

Proof of Lemma 18 : Take $x_1 < x_2$ By definition, we have

$$\begin{aligned}
G(x_1) &= \int \min\{\theta, x_1\} dP(\theta) \\
&\leq \int \min\{\theta, x_2\} dP(\theta) = G(x_2)
\end{aligned}$$

■

We are now ready to present the proof of Theorem 8.

Proof of Theorem 8 :

(i) We prove the upper bound. The lower bound follows automatically from the lower bound in (20). By definition, we have
$$J_{P,\overline{\pi}} = V_0^t(0) = \overline{h}_0$$
where $\overline{h}_0$ is computed recursively via (57) – (59). It follows from (59) that
$$\overline{h}_0 = \sum_{k=0}^{n-1} d_k \left(\overline{g}_{k+1} + \Upsilon\left(\overline{t}_{k+1}, \overline{g}_{k+1}\right) \right) = \sum_{k=0}^{n-1} d_k \left(\overline{g}_k - p_k \right)$$
Thus, to prove the upper bound it is sufficient to show that $\overline{g}_k \leq \overline{t}_k$ for all $k = 0, ...n$ This is in turn proven by induction. Note that we have $\overline{g}_n = \overline{t}_n = \lambda_{\max}$. Suppose that $\overline{g}_k \leq \overline{t}_k$. for some $0 \leq k < n$. Then we have:
$$\overline{g}_k = \overline{g}_{k+1} + \Upsilon\left(\overline{t}_{k+1}, \overline{g}_{k+1}\right) \tag{60a}$$
$$\leq \overline{t}_{k+1} + \Upsilon\left(\overline{t}_{k+1}, \overline{t}_{k+1}\right) \tag{60b}$$
$$= \overline{t}_{k+1} + \Gamma\left(\overline{t}_{k+1}\right) \tag{60c}$$
$$\leq \overline{t}_{k+1} + \overline{\Gamma}\left(\overline{t}_{k+1}\right) \tag{60d}$$
$$= \overline{t}_k \tag{60e}$$

where (60a) follows from (58), (60c) follows from the definitions of Υ and Γ in (55) and (4), (60d) follows from the right hand side inequality in (17), (60e) holds by definition, and the proof of (60b) is as follows:

$$\begin{aligned}&\left(\overline{g}_{k+1}-\overline{t}_{k+1}\right)+\Upsilon\left(\overline{t}_{k+1},\overline{g}_{k+1}\right)-\Upsilon\left(\overline{t}_{k+1},\overline{t}_{k+1}\right)\\&=\left(\overline{g}_{k+1}-\overline{t}_{k+1}\right)+\int_{\lambda_{\min}}^{\overline{t}_{k+1}}\left(\overline{t}_{k+1}-\overline{g}_{k+1}\right)dP\left(\theta\right)\\&=\left(\overline{g}_{k+1}-\overline{t}_{k+1}\right)\left(1-F\left(\overline{t}_{k+1}\right)\right)\leq 0.\end{aligned}$$

The proof is completed by induction.

(ii) The upper bound follows immediately from the upper bound in (19). We prove the lower bound. It follows from (16) that

$$J_P^*=\sum_{k=0}^{n-1}d_k(t_k-p_k)$$

where t_k are the optimal thresholds computed via (13). Therefore, it is sufficient to prove that $t_k\geq \underline{t}_k$, for all k. This is in turn proven by induction. Note that $t_n=\underline{t}_n=\lambda_{\max}$. Suppose that $t_k\geq\underline{t}_k$ for some $k\geq 0$. Then,

$$t_{k-1}=t_k+\Gamma\left(t_k\right)\geq\underline{t}_k+\Gamma\left(\underline{t}_k\right)\geq\underline{t}_k+\underline{\Gamma}\left(\underline{t}_k\right)=\underline{t}_{k-1}$$

where the leftmost inequality follows from $t_k\geq\underline{t}_k$ and monotonicity of the function $G\left(x\right)=x+\Gamma\left(x\right)$ (Lemma 18), and the rightmost inequality follows from the left hand side inequality in (17). The proof is complete by induction. Note that this also proves the lower bound in (19).

(iii) It is sufficient to prove (21). The inequality (22) is immediate from the definitions of $\overline{\Gamma}$, $\underline{\Gamma}$, and $\widehat{\Gamma}$. It can be shown that

$$\left|J_{P,\hat{\pi}}-\widehat{J}\right|=\left|\sum_{k=0}^{n-1}d_k(\hat{g}_k-\hat{t}_k)\right|\leq\sum_{k=0}^{n-1}d_k\left|\hat{g}_k-\hat{t}_k\right| \tag{61}$$

where $\hat{g}_{n-k-1}$ are computed recursively via (57) – (59). For convenience in notation let $a_k=1-F(\hat{t}_k)$. It can be verified that :

$$\hat{g}_k-\hat{t}_k=\left(\hat{g}_{k+1}-\hat{t}_{k+1}\right)\hat{a}_{k+1}+\Gamma(\hat{t}_{k+1})-\widehat{\Gamma}(\hat{t}_{k+1}).$$

In addition, the sequence a_k is monotonic: $\hat{a}_0\leq\hat{a}_1\leq\cdots\leq\hat{a}_n$. Therefore,

$$\left|\hat{g}_k-\hat{t}_k\right|\leq\left|\hat{g}_{k+1}-\hat{t}_{k+1}\right|\hat{a}_0+\left|\Gamma(\hat{t}_{k+1})-\widehat{\Gamma}(\hat{t}_{k+1})\right|. \tag{62}$$

Summing both sides of the above inequality from $k=0$ to $k=n-1$, followed by adding a term $\left|\hat{g}_0-\hat{t}_0\right|\hat{a}_0$ to the right hand side yields (note that $\hat{g}_n=\hat{t}_n$):

$$\sum_{k=0}^{n-1}\left|\hat{g}_k-\hat{t}_k\right|\leq\hat{a}_0\sum_{k=0}^{n-1}\left|\hat{g}_k-\hat{t}_k\right|+\sum_{k=0}^{n-1}\left|\Gamma(\hat{t}_{k+1})-\widehat{\Gamma}(\hat{t}_{k+1})\right|$$

If $\hat{a}_0<1$, the above inequality yields

$$\left|J_{P,\hat{\pi}}-\widehat{J}\right|\leq\frac{\max\{d_k\}}{F(\hat{t}_n)}\sum_{k=0}^{n-1}\left|\Gamma(\hat{t}_{k+1})-\widehat{\Gamma}(\hat{t}_{k+1})\right|.$$

The distribution free upper bound with $f(k)=\sum_{i=0}^{k-1}d_k$ can be obtained by recursively plugging in (62) with $\hat{a}_0=0$ in (61).

■

Proof of Theorem 10: We provide a quick sketch of the proof. For simplicity, we assume $p_k=0$ for $k=0,...,n-1$. The proof can be modified in order to take into account the case where the disutility for backlogging is not identically zero. When the mean and the variance of a bounded random variable are constrained, the cumulative distribution is consequently constrained, and so is its integral, i.e. Γ. To provide tight upper bounds and lower bounds we therefore argue first that no distribution can fall beyond those bounds while satisfying the mean and variance constraints. This is in the same spirit as Markov's inequality.

For brevity, we do not consider all six cases individually, and instead give the first case of the upper bound in detail, to illustrate the approach. As x grows away from zero, one could place no probability whatsoever before x without violating the mean and variance constraints, up to some point $\overline{x}_1$. It follows that we have, $\overline{\Gamma}$ is 0 for all $x\in[0,\overline{x}_1]$. The transition happens when no such distribution exists. Since the mean can be met even with an impulse on μ, the variance constraint will fail first. A distribution supported on and $[x,1]$, with mean μ, has the largest variance when it is a two-impulse distribution concentrated on x and 1. Say these masses are p and $1-p$ respectively. The mean constraint dictates that $p=\frac{1-\mu}{1-x}$. Therefore the second moment is $x^2\frac{1-\mu}{1-x}+\frac{\mu-x}{1-x}$. At the breakpoint, this maximal second moment will match $\sigma^2+\mu^2$ exactly, i.e. we must have:

$$(1-\mu)\overline{x}_1^2+(\sigma^2+\mu^2-1)\overline{x}_1+(\mu-\sigma^2-\mu^2)=0.$$

It follows that:

$$\overline{x}_1=\frac{\mu(1-\mu)-\sigma^2}{1-\mu}.$$

The other cases can be obtained using the same procedure. First constraint violations are identified, then resulting transition points are marked down. Across each transition the optimal distribution changes, and consequently so do the upper and lower bounds on Γ. ■

Proof of Corollary 11: For simplicity, we assume $p=0$. The proof can be easily modified to take into account nonzero disutility of delay.

(i) Both bounds follow from Theorem 8, equation (19) for partial information and equation (20) for full information. It remains to show that the lower bound is quadratic in σ. This in turn follows from the quadratic dependence of $\overline{\Gamma}(\cdot)$ on σ as defined in Theorem 10, which preserves the quadratic dependence of all the thresholds $\overline{t}_k$ on σ in backward computation.

(ii) Note that $\overline{t}_{n-1}=\underline{t}_{n-1}=t_{n-1}=\mu$ regardless of the information structure. Thus we always have the optimal threshold in the stage $k=n-1$. Moreover, as stated in Theorem 10, the bounds $\overline{\Gamma}(\cdot)$ and $\underline{\Gamma}(\cdot)$ are pointwise tight. In particular, they are tight at $x=\mu$. Thus when $n=2$, there always exists a pair of distributions that achieve the upper and lower bounds respectively.

■

Mardavij Roozbehani Received the B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 2000, the M.Sc. degree in Mechanical and Aerospace Engineering from the University of Virginia, Charlottesville, VA in 2003, and the Ph.D. degree in Aeronautics and Astronautics from the Massachusetts Institute of Technology (MIT), Cambridge, MA in 2008. Since 2012, he has held a Principal Research Scientist position at the Laboratory for Information and Decision Systems (LIDS) at MIT, where he previously held postdoctoral, course instructor, and research scientist positions between 2008 and 2011. His main research interests include distributed and networked control systems, software and finite-state systems, and dynamics and economics of power systems with an emphasis on robustness and risk. Dr. Roozbehani is a recipient of the 2007 AIAA graduate award for safety verification of real-time software systems.

Donatello Materassi received a Laurea in ?Ingegneria Informatica? and a ?Dottorato di Ricerca? in electrical engineering/nonlinear dynamics and complex systems from Università degli Studi di Firenze, Florence, Italy. He has been a Research Associate at University of Minnesota, Minneapolis. Currently, he is with the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge and also holds a position as a lecturer at Harvard University. His research interests include nonlinear dynamics, system identification, and classical control theory with applications to power grids, biophysics, statistical mechanics, and quantitative finance.

Mesrob I Ohannessian received a B.Eng. from the American University of Beirut, Lebanon, and an S.M. in Information Technology and a Ph.D. in Electrical Engineering and Computer Science, both from the Massachusetts Institute of Technology, Cambridge, MA. He was a Marie Curie Fellow at UniversitÈ Paris-Sud, France, under an ERCIM Alain Bensoussan Fellowship. He is currently a postdoctoral researcher at the Microsoft Research - Inria Joint Centre in France. His research interests are broadly in statistics, information theory, machine learning, and their applications.

Munther A Dahleh received the B.S. degree from Texas A & M university, College Station, Texas in 1983, and his Ph.D. degree from Rice University, Houston, TX, in 1987, all in Electrical Engineering. Since then, he has been with the Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, where he is now the Associate Department Head. Previously, he was the acting director of the Laboratory for Information and Decision Systems. He has been a visiting Professor at the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, for the Spring of 1993. He has held consulting positions with several companies in the US and abroad.

Dr. Dahleh is interested in problems at the interface of robust control, filtering, information theory, and computation which include control problems with communication constraints and distributed agents with local decision capabilities. In addition to methodology development, he has been interested in the application of distributed control in the future electric grid and the future transportation system with particular emphasis in the management of systemic risk. He is also interested in various problems in network science including distributed computation over noisy networks as well as information propagation over complex engineering and social networks. Dr. Dahleh is the co-author (with Ignacio Diaz-Bobillo) of the book *Control of Uncertain Systems: A Linear Programming Approach*, published by Prentice-Hall, and the co-author (with Nicola Elia) of the book *Computational Methods for Controller Design* published by Springer.