RC25408 (WAT1210-090) October 31, 2012
Mathematics

IBM Research Report

A Framework for the Analysis of Probabilistic Demand Response Schemes

Pavithra Harsha, Mayank Sharma, Ramesh Natarajan, Soumyadip Ghosh
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 208
Yorktown Heights, NY 10598
USA

A Framework for the Analysis of Probabilistic Demand Response Schemes

Pavithra Harsha, Mayank Sharma, Ramesh Natarajan, Soumyadip Ghosh

***Abstract*—We describe the class of *probabilistic demand response (PDR)* schemes, which are particularly suited for dynamic load management in the residential sector. Our main contribution is a new methodology for implementing and analyzing these schemes based on an operational objective function that balances the total cost of meeting demand, which includes the costs of supply generation, and spinning reserves, with the total revenue from the met demand and the gain from storage/deferment. We derive structural results for the design of PDR schemes in terms of sufficient conditions that yield a well-posed joint optimization problem for the two decision variables: the planned supply generation level and the real-time PDR signal magnitude. These results are used to evaluate the suitability of various proposed PDR schemes in single-period and multiple-period contexts. Finally, using simulations, we illustrate the application and effectiveness of the proposed methodology for a collection of thermostatically-controlled residential loads.**

***Index Terms*—Smart Grid, Demand Response, Load Curtailment, Spinning Reserve, Newsvendor Model**

I. Introduction

A. Background

In the traditional power grid, the supply side consists of multiple sources of generation with varying response time-scales, and the demand side is characterized by passive electric loads. The program operator is responsible for matching supply and demand, and for ensuring the reliable operation of the power grid. These objectives are achieved primarily by controlling the supply side: first, based on the forecasted demand, a day-ahead plan is derived for the generators with slow response times; then, any real-time mismatch is balanced using the reserve generators with fast response times. Demand side controls are rarely used, although load shedding may be used as a last resort.

In the emerging smart grid, *Demand Response* (DR) schemes, which involve active load management, are widely recognized to have equivalent power-balancing outcomes, in principle, to the supply-side controls [1]. Furthermore, as [2] note, DR schemes are highly flexible, and their response times are comparable to that of the reserve generators on the supply side. Therefore, the use of relatively-inexpensive DR schemes to offset some of the reserve generation capacity can be of considerable economic value for the grid operation. Additionally, DR schemes can also be used by the program operator to achieve other operational objectives, such as peak load shaving, or to provide ancillary services for power balancing and regulation.

A schematic for a prototypical DR scheme is shown in Fig. 1 (although for the specific class of PDR schemes described further in Section I-B below). Here, for the time interval of interest, a signal z is sent to each consumer, and the k'th consumer responds by modifying their base load L_k in the absence of this signal, to the responsive load value $L_k(z)$, which are then aggregated to provide the demand response $D(z)$. The signal z is always deterministic, whereas the base load L_k and responsive load $L_k(z)$ are invariably stochastic. Various DR schemes can be identified based on the realized form of $L_k(z)$, and the methodology used by the program operator to determine the signal z.

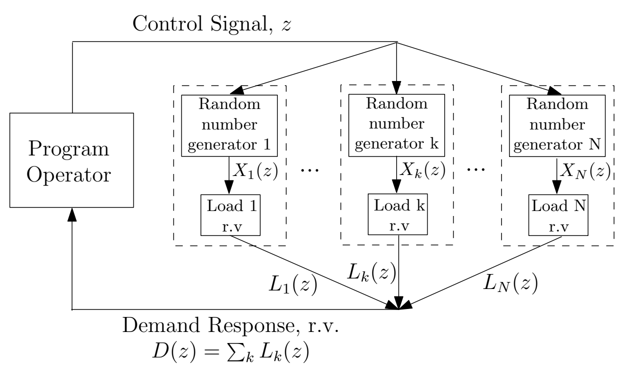

Fig. 1. Schematic of a probabilistic demand response (PDR) scheme.

The two main classes of DR schemes that have been extensively studied in the literature are dynamic pricing and direct control. In dynamic pricing schemes

All the authors are with IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.
Corresponding author: Email - pharsha@us.ibm.com

(e.g., [3], [4], [5] and the references therein), the signal z reflects the real-time cost of electricity generation, and can be interpreted as the mechanism by which the program operator passes on this cost to the consumer. The corresponding responsive load generally takes a form $L_k(z) = L_k - \zeta z + \epsilon$, where ζ reflects the price sensitivity for electricity consumption, and ϵ is a noise term that captures the idiosyncratic demand effects [6]. In direct control schemes, the signal z can be used to directly modify the on-off states [7], [8], or to reset the operating set-points [9], [10], [11], of the individual consumer loads according to certain preset, deterministic rules. The corresponding responsive load often takes the form $L_k(z) = zL_k$, which reflects a proportional modification of the base load.

The benefit of these two classes of DR schemes for the industrial or commercial sectors is readily apparent, since only a few large loads are typically involved in the planning, and their corresponding responsive load characteristics can be easily ascertained. However, in the smart grid of the future, the program operator will have to contend with significantly greater short-term variability and uncertainty on the supply side, due to the increasing adoption and integration of intermittent sources of renewable energy in the grid operations. This increased supply-side variability will invariably require DR schemes to be extended and implemented for the residential sector as well. Indeed, this is a primary impetus for the recent interest in providing the smart grid infrastructure for residential DR schemes.

Nevertheless, the extension of dynamic pricing or direct control DR schemes to the residential sector raises many issues, which include, (1) *Nature of Load:* Unlike the large industrial loads, the individual residential loads are numerous, often operate only at a discrete set of power usage levels, and comprise of legacy equipment that are expensive to retrofit with DR controls. The use of direct control schemes in this situation is computationally burdensome, since the program operator has to decide which of the numerous devices to turn on and off, or which set points to modify, in order to achieve the specified operational objectives. (2) *Estimation of Model Parameters:* The use of dynamic pricing schemes requires the estimation of key parameters in the responsive load $L_k(z)$ beyond the base load L_k, such as the sensitivity ζ and the model parameters for the noise (ϵ); these parameters are influenced by changes in the comfort-level settings and residential usage profiles, and vary considerably across individual households. (3) *Consumer Satisfaction:* Any DR scheme must meet the service-level expectations of residential consumers, who may discern that certain schemes frequently "tamper" with their base load in an undesirable way, leading to program opt-outs. For direct control schemes, the already onerous computational burden of the program operator is further increased by having to ensure that the service-level expectations are maintained over time. (4) *Fairness:* Consumers also need to perceive that the DR schemes are fair relative to their neighbors over time [1], [12]. For instance, dynamic pricing schemes that expose the real-time prices to all consumersregardless of the desirable responsiveness or their economic capability raises questions of overall fairness in the program implementation. Similarly, direct control schemes must ensure that no individual household is unfairly penalized with frequent drops in load with respect to the norm in its local neighborhood.

B. Overview and Contributions

This paper formally introduces the class of *Probabilistic Demand Response* (PDR) schemes, which are perhaps best described using a specific example from the Olympic Peninsula GridWise Project [13]. The project included field demonstrations of several smart grid technologies in a residential electric-feeder network. The specific example concerned a collection of 50 residential water heaters, whose individual base loads L_k were subject to their usual set-point based operation, and were hence uncertain quantities. In each time period, the control modules for these water heaters received a signal z from the program operator. Based on this signal magnitude, and taking into account the user-defined household comfort settings, a curtailment probability was then explicitly derived. Further, for each water heater, based on this curtailment probability, a sample Bernoulli outcome $I_k(z)$ was obtained using a random number generator (RNG). This sample outcome was either $I_k(z) = 0$ or $I_k(z) = 1$, and the corresponding load response of the water heater to the signal z then took the form $L_k(z) = I_k(z)L_k$. So, an outcome $I_k(z) = 0$ led to full curtailment of the k-th load for that time period.

The key characteristics of such PDR schemes, which are explicitly illustrated in Fig. 1, can be described as follows. The program operator communicates a deterministic signal z to all participating loads. The magnitude of z is used to set the parameters of a *suitably designed* probability distribution, from which, for each load, a sample outcome is then obtained using a RNG. This sampled outcome is used to modify the operation of the corresponding load according to some pre-defined rule. Note that even though the same signal z is sent to each load, and even if for purposes of argument, the comfort settings and base loads are identical across the entire collection, the individual responsive loads will not be identical due to the variability in the randomly sampled outcomes. Consequently, the base load L_k, whose magnitude is already uncertain as noted above, is further modified, based on the magnitude of signal z to obtain the randomized responsive load $L_k(z)$. Three specific examples of such PDR schemes are described in Section II, in order to exemplify the design and implementation of PDR schemes.

In our view, PDR schemes can overcome the impediments that were described earlier. First, these schemes are quite simple to implement for a variety of loads, including discrete or continuous loads (criterion 1). Second, they do not require any additional estimation beyond that for the base load (criterion 2), nor do they impose an extensive computation burden for obtaining the individual load settings; in fact, the distribution of the aggregate load response for a PDR scheme can be obtained directly using statistical methods, which minimizes the two-way communication requirements and is hence particularly desirable for a DR scheme [14] (e.g., the utility does not need to know the sample outcomes of the RNG for each individual load). Third, in principle, these schemes are also intrinsically fair (criterion 4) because, other things being equal, the same signal z elicits a statistically identical response from all the participating loads. A caveat here is the requirement for implementing random sampling over a large and diverse collection of participating loads; however, state-of-the-art RNGs that pass statistical independence tests in very high dimensions can be implemented with moderate effort [15], so that this fairness objective is achievable not only in theory but also in practice. Fourth, as shown below in Section IV, the issues related to the consumer satisfaction over multiple periods (criterion 3) can be addressed by including constraints on consumer service levels over multiple periods.

We note that probabilistic or randomized protocols are widely used in many key areas of computer networking and communication infrastructure (e.g., for ethernet multiple access resolution, router queue management, multicast and routing in unreliable networks, and router desynchronization), primarily due to their simplicity and scalability, and due to their ability to redistribute the load and reduce resource contention by desynchronizing network activity.

The PDR schemes described here assume that the load curtailment effected are always accepted by the customers as long as a certain level of service is maintained, and there are sufficient incentives – economic or otherwise – for program participation. For example, one driver for increased buy-in is clearly system-wide reduction in the cost of electricity due to better planning because of improved grid efficiency and reliability due to lower use of reserves. The design of participation incentives or compliance rebates (provided over longer time horizons) is an interesting problem in itself, but is not within the scope of this paper.

The main contribution of this article is a framework for implementing a given PDR scheme i.e., setting the signal value z. The three examples of the PDR schemes are analysed using this framework, and are seen to provide demand-side controls with a distinct flavor. The general framework jointly optimizes the planning and real-time decision variables, enabling the program operator to match supply and demand in a systemically beneficial manner i.e., through an operational objective function that incorporates various cost (planned generation and spinning reserve costs) and revenue streams (including the gain from storage/deferment). In addition, we consider using the framework for the special case where the planning decision variable, the supply generation level, is fixed to an exogenously specified value. This flexibility allows the program operator to use the same framework for the longer time scale planning, as well as for the shorter time scale control of the electric grid through demand response.

As [16] have noted, the control mechanisms in the emerging smart grid must consider risk-based approaches for ensuring operational reliability over short time scales. This objective is similar in spirit to the loss-of-load probability (LOLP) considerations that are used to plan generation capacity in traditional power grids, albeit over longer time scales [17]. However, while the focus in the literature has hitherto been on ensuring generation adequacy to handle supply intermittency due to increased capacity from renewables, this paper proposes that the demand-side uncertainty that is induced in the PDR schemes can also be assessed within a similar short-term, risk-based operational framework.

The objective function of the framework is similar to generalizations of the well-known *newsvendor model* described in inventory theory ([18] and the references therein). Accordingly, our analysis broadly follows this theory, but we propose modifications and extensions, and obtain new structural results, for applying this theory to this smart grid context. In particular, for the PDR schemes, we derive sufficient conditions that yield a well-posed joint optimization problem with non-trivial optimal solutions for the decision variables. These sufficient conditions impose certain requirements on the distribution function of the responsive loads in the PDR scheme, which are then used to evaluate the suitability of the proposed schemes in the single-period and multiple-period contexts respectively.

Finally, we present a simulation study for a realistic scenario consisting of a large population of thermostatically controlled loads. Our results demonstrate that PDR schemes can significantly improve the program operator's objective of meeting residential demand, while minimizing the use of expensive spinning reserves and the risk of undersupply, without compromising consumer service expectations. In this particular scenario, we are able to demonstrate an increase in operational profit over a typical day of over 57%, through a reduction of the reserve requirements by 64%, when compared to a baseline no-DR scenario. This is achieved at very minimal discomfort to the consumers: the average thermostat readings differ at worst by only 0.29^0C.

C. Organization

Section II describes three specific examples of PDR schemes that are analyzed in the paper. Section III describes our methodology and framework for implementing and evaluating these schemes based on the proposed operational objective function. Section IV considers the multi-period generalization of PDR schemes incorporating user-behavior dependencies across the individual periods. Section V describes results for a simulated case study of PDR schemes in a realistic scenario. Section VI provides the summary and conclusion.

II. Probabilistic demand response schemes

Consider the program operator in the schematic in Fig. 1 who sends out a load management signal $z \in Z$ to a group of N customers every period (in implementations, this is usually 5 minutes in duration). The domain Z is a closed and compact set.

The following are three example PDR scheme designs:

P.Br: *Bernoulli ON/OFF*

$$L_k(z) = I_k(z)L_k,$$

where I_k is 0 with probability z and 1 otherwise. Here, the signal $z \in [0,1]$ and can be mapped to the probability of curtailment of the scheme. This scheme was implemented in the Olympic Peninsula project [13].

P.DU: *Discrete Uniform scaling*

$$L_k(z) = \frac{1}{q}X_k L_k,$$

where X_k is discrete uniform r.v. in $\{q-z,...,q\}$, q is a positive integer, and $z \in \{0,...,q\}$. Here, the signal z can be mapped to a level of the load that is probabilistically curtailed, with q fixed *a priori*.

P.CU: *Continuous Uniform scaling*

$$L_k(z) = (1 - z + zX_k)L_k,$$

where X_k is $U(0,1)$. Here, the signal $z \in [0,1]$ and can be mapped to the probabilistically controlled load fraction, while $1-z$ is the guaranteed load fraction. More generally, we can have $L_k(z) = \big(1 - z + g(z)X_k\big)L_k$ where the r.v. X_k with values in the range $[0,1]$ is independent of z, and $g(z) : [0,1] \to [0,1]$ is a positive convex function of z such that $g(z) \le z$ (the reason for this special structure is clarified in Section III below).

In the above schemes, a deterministic signal z induces a probabilistic outcome via randomization, which in turn induces a probabilistic load curtailment as the demand response. The choice of the elements of the scheme, including the probability distribution that yields $L_k(z)$ and the mechanism that sets the signal value z given the program operators objectives and constraints, are part of the design elements of a demand response scheme for a load. The example schemes above are practical schemes that can be used to adapt existing household devices such as washer/dryers, dishwashers and other thermostatic load to the smart grid

The aggregate controllable load observed by the program operator is given by the r.v.

$$D(z) = \sum_{k=1}^{N} L_k(z). \tag{1}$$

Note that this aggregate includes two sources of uncertainty: 1) the uncertainty in individual loads (L_k) in the absence of any signal, which may have some dependence across the population due to shared patterns of appliance usage, and 2) the uncertainty induced by the signal z by design through the probability distribution of the random number generator. The RNG is part of the scheme's design, and thus designed to be independent and identically distributed (i.i.d.) across individual loads.

Setting the signal value z requires an in depth understanding of the properties of $D(z)$. Estimation of the aggregate load distribution is simplified for large N if the uncertainty in the individual loads, L_k, are also i.i.d., in which case $D(z)$ can be well approximated using the central limit theorem. Observe that for a practical implementation of a PDR scheme, it suffices for the program operator to statistically estimate $D(z)$, as opposed to measure individual $L_k(z)$ for each k.

This can simplify the process of learning $D(z)$, thereby reducing the amount of two-way communication needed by the scheme.

III. Demand management

The implementation and evaluation framework for the probabilistic schemes is based on the well-known *price-sensitive newsvendor model* described in the operations management literature [19]. This model is adapted to the smart grid context, and for this setting we obtain some newer results as described below.

Consider a period in which the program operator provides its customers with a signal $z \in Z$ with aggregate demand response $D(z)$. For simplicity, and without loss of generality in the analysis or results, we omit the portion of load that cannot be controlled, and assume that $D(z)$ has a form described in Eq. (1). We assume that the program operator incurs a unit cost c for any planned generation, and charges customers a unit price p for their consumption. We also assume a unit cost m to meet any supply shortfalls, from the ancillary services market, and a salvage price s from selling any supply excess. The cost from the ancillary services market is a blended, possibly pre-negotiated, cost of ancillary services and the spot market price of electricity. Our view of salvage is a contracted selling price with a bulk storage farm, or any other form of return on excess supply including unused fuel. This is just set to zero if this supply is just grounded. To avoid trivial solutions, we assume that $p > c > s$ and $m > c$. Letting $x \in \mathbb{R}^+$ denote the amount of planned generation, the profit-maximization objective of the program operator can be formulated as follows:

$$P: \quad \max_{x \in \mathbb{R}^+,\ z \in Z} \quad \pi(x, z), \tag{2}$$

where

$$\begin{aligned} \pi(x,z) \quad = \quad & pE[D(z)] - cx - mE\left[(D(z)-x)^+\right] \\ & +sE\left[(x-D(z))^+\right] \end{aligned}$$

In this formulation, we allow both x and z to be decision variables as it incorporates the flexibility of jointly optimizing for the planned generation and the demand signal. For example, this joint optimization may be possible in one-hour ahead electricity market and it allows the planned generation to match the predicted signalled-demand. However, depending on the flexibility (i.e., ramp constraints) of planned generators, x may also be predetermined and fixed, in which case, it is not treated as a decision variable in the optimization. We refer to this latter problem as the signaling problem that is solved at a finer time scale (e.g., real-time, say every 5 minutes) and the former problem as the joint planning and signaling problem that is solved at a coarser time scale (say every hour). The advantage of this unified evaluation framework is that it incorporates the sequential aspect of energy planning and signaling (first turn on conventional generators and follow it with reserve generators and demand response in real-time for balancing).

The use of formulation 2 for the evaluation of demand response has other advantages as well. First, it incorporates all the different costs and revenue streams that are faced by a program operator in order to obtain the demand response signal in a way that makes economic sense, which is crucial for incentivizing the implementation of demand response schemes. Second, it incorporates the notion of managing risk due to real-time uncertainty with signaling. This can be seen by rewriting the expected profit by substituting $(x-D)^+$ with $x - D + (D-x)^+$, to obtain:

$$\pi(x,z) = (p-s)E[D(z)] - (c-s)J_\beta(x,z), \tag{3}$$

where

$$J_\beta(x,z) = x + \frac{1}{(1-\beta)} E\left[(D(z)-x)^+\right].$$

In [20], it is shown that minimizing $J_\beta(x,z)$ w.r.t x gives the *conditional value at risk*, $\mathrm{CVaR}_\beta(z)$, of the r.v. $D(z)$:

$$\mathrm{CVaR}_\beta(z) = \min_x J_\beta(x,z).$$

The quantity CVaR is a well-known risk measure, which is also referred to in the literature as the average value at risk, mean excess loss, or mean shortfall, and is also closely related to the the widely-used reliability metric *loss of load in expectation* (LOLE) in the energy literature [17]. Third, the formulation can be easily extended to include boundary constraints on generation and piecewise linear generation costs with minor modifications to the results presented in this paper.

Note that the formulation 2 is similar to the price-sensitive newsvendor model with emergency ordering, where the excess demand is satisfied by the market at a high but constant cost, m [18]; in fact, these are equivalent where the decision variable z is the price p. In the same spirit as [21] for the price-sensitive newsvendor model with lost-sales ($m = p$), we obtain optimality conditions for general demand distributions for which the expected profit problem is concave (however, the approach in [21] does not generalize to the formulation 2 because $m \neq p$). More specifically, we derive conditions under which:

1) $\pi(x,z)$ is jointly concave in (x,z).
2) $\pi(x,z)$ is concave in z given x.

Clearly, the conditions for 2 are more general, and must hold for 1 to be satisfied. The conditions 1 and 2 are important for developing efficient demand response algorithms. They also ensure that the optimal solutions are non-trivial and useful in practice, and as shown below, provide the criteria for evaluating the suitability of various PDR schemes.

It is well known that $\pi(x, z)$, the objective function for problem P, is a concave function in x for given z, and the solution of the newsvendor problem has the form [19]:

$$x^*(z) = \inf\left\{x \geq 0 : F_{D(z)}(x) \geq \beta = \frac{m-c}{m-s}\right\}, \quad (4)$$

where $F_{D(z)}(.)$ is the distribution function for $D(z)$, and β is referred to as the critical newsvendor quantile. But extending this result to obtain the joint concavity of $\pi(x, z)$ in (x, z) is not always guaranteed, as we will see below. We first analyze the scheme **P.Br**.

Claim 1. *Under scheme* **P.Br**, $\pi(x, z)$ *is* not *jointly concave in* (x, z) *but concave in* z *given* x.

The proof of this claim is provided in the appendix, where a counter example is provided which shows that the objective is convex (rather than concave) in z after substituting the optimal $x^*(z)$ from Eq. (4) in formulation 2. The intuition for this comes from the nature of the Bernoulli r.v. that tends to increase the standard deviation for small non-zero values of z (see Fig. 2 (a)). This is equivalent to increasing risk of undersupply while decreasing mean demand and is undesirable. As a result, the **P.Br** scheme has lower flexibility from the perspective of our evaluation framework when the decision variables, i.e., the planned generation and the signal magnitude must be optimized simultaneously. However, for fixed x, we retrieve the concavity of the objective in z (see Claim 1) so that the signal magnitude can always be fine-tuned when planned generation has been fixed (this is demonstrated in some of the simulation results in Section V).

The analysis for this scheme indicates that a preferred curtailment scheme would allow for the mean demand and the risk of undersupply to be simultaneously reduced. Such a preferred design, which is illustrated in Fig. 2(b) can motivate the design of other schemes which have this desirable property, as shown below.

First we consider the class of probabilistic schemes where the induced noise is independent of the signal z, viz., $D(z) = d(z, \epsilon)$ where $d(.,.)$ is a deterministic function and ϵ is a r.v. with density function h and distribution function H respectively independent of z. Examples of such schemes include the case when the induced probability distributions for the demand response can be reduced to a standard form with location and scale parameters. The scheme **P.CU** is such an example, although the schemes **P.Br** and **P.DU** are not.

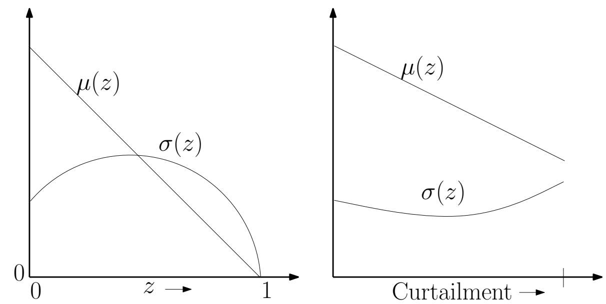

(a) Bernoulli ON/OFF scheme (b) Example of preferred scheme

Fig. 2. Mean $\mu(z)$ and standard deviation $\sigma(z)$ of the aggregate demand $D(z)$ as a function of the signal z for (a) the **P.Br** scheme, and (b) a preferred scheme in our framework. Note that this is a schematic and the y-axis is not drawn to scale.

Claim 2. *Suppose* $d(z, \epsilon)$ *is convex in* z *for all* ϵ*, and* $E[d(z, \epsilon)]$ *is linear in* z *then* $\pi(x, z)$ *is jointly concave in* (x, z).

Proof: From Theorem 2 in [20], we know that $J_\beta(x, z)$ is jointly convex in (x, z) when $d(z, \epsilon)$ is convex in z for every ϵ. Since $E[D(z)] = J_\beta(0, z)$ and the latter is also convex. Hence, the linearity in the expected demand is the minimum requirement that can be imposed to guarantee the joint concavity of $\pi(x, z)$. ■

Remark 3. *Consider the case when* $D(z, \epsilon) = A + Bd(z)$*, with* A *and* B *are positive independent r.v.'s with* $E[A] = 0$ *and* $E[B] = 1$*. It is shown in [18], that if* $d(z)$ *is continuous and decreasing in* z *then* $\pi(x, z)$ *is jointly concave in* (x, z)*. Their proof first shows concavity in the vector* $(x, \mu) = (x, d(z))$*, which is a special case of Claim 2. From the 1-to-1 mapping between the expected demand,* $d(z)$*, and* z*, they recover the result for* (x, z).

The above remark has an implication that it is suffices to impose the conditions of Claim 2 on the net signal, and transform it to the signal space, z, with a one to one correspondence.

Now consider the most general case of schemes in which the noise in the demand response also depends on the signal z. Here, general second-order optimality conditions can be derived when the signal z is continuous, however, it is often easier to analyze these schemes on a case-by-case basis based on the structural properties of individual schemes. This is because the conditions involve the derivatives of the density function and the distribution function of $D(z)$ which are convolution of the respective load responses of the individual customers, and are usually hard to obtain in closed form. However, depending on the structural properties of the scheme, it may be possible to perform these convolution operations for various sub-collections of the customer loads, in order to carry out the analysis (e.g., Claim 1 is an example of such an analysis). Here, we provide the analysis for scheme **P.DU** under two limiting conditions:

Claim 4. *Under the scheme* **P.DU***,* $\pi(x, z)$ *is jointly concave in* (x, z) *for large* N *with i.i.d loads when* $\frac{\mu_L}{\sigma_L} < 12q^2 + 12q - 1$ *or for large* q*.*

Discussion: The main observation from all the above analysis is as follows: *demand response schemes should be designed so that the control signal not only decreases the mean load but also decreases the risk of undersupply.* Otherwise the usefulness of the demand response scheme is undermined, both at the level of the individual customer who is unsure of planning specific activities in the light of control actions, and more so at level of the aggregate customer population, where the specification of the planned generation becomes more difficult. This observation is particularly crucial for PDR schemes which are specifically designed to induce uncertainty in the demand response.

Note that probabilistic schemes are not beneficial for deterministic loads as they increase the variance of the demand. In fact, no demand response scheme is beneficial because one can use planned generation for matching these loads quite accurately unless there is an unforeseen supply shortfall. On the other hand, highly variable loads (e.g., AC loads during high-temperature weather conditions) are good candidates for the implementation of PDR schemes.

A. Multiple Load Classes

All controllable loads of a customer have been assumed to get the same curtailment signal. The program operator may however choose to segment the customer base into load classes, for instance, based on similar load consumption patterns. Members of a class would then get the same signal, but the operator could choose to send different value to different classes. Suppose aggregate load of each class, i, is denoted by $D_i(z_i)$ where $i \in \{1, ..., M\}$. The aggregate load faced by the program operator is $D(\mathbf{z}) = \sum_{i=1}^{M} D_i(z_i)$ where $\mathbf{z}$ which is a vector of the signals z_i $\forall i$. This segmentation can help the program operator gets higher profits by jointly managing all load classes together than separately. The risk of undersupply is also lowered with aggregation. This is because LOLP or CVaR, our primary risk objective, is subadditive [20], i.e., for any two r.v.'s the CVAR of their sum is smaller than the sum of their CVAR's.

A relevant question here is whether each class is less curtailed on average, when the class are optimized together rather than separately. We address this question in a special case.

Claim 5. *Consider an aggregation of independent load classes each of which is normally distributed with an affine mean* $\mu_i(z_i)$ *and a convex standard deviation* $\sigma_i(z_i)$*. If the program operator jointly optimizes formulation 2 with all the load classes together, each class is curtailed by a smaller magnitude than if they are optimized separately.*

The concept of multiple load classes can be extended to multiple contract classes. Customers may chose to sign onto different contract classes depending on their willingness to undergo higher curtailment magnitudes in return for higher economic participation incentives. Extending our analysis to optimally price such portfolios of contract classes is of future interest.

IV. Multi-period Demand Management

Multi-period model frameworks capture constraints and behaviors that apply over multiple periods or vary over time including certain load characteristics or ramp constraints. There are some important multi-period interactions that need to be modeled in the context of PDR schemes. Foremost are the issues of customer satisfaction with respect to demand management across time. Sample paths where a customer encounters significant discomfort due to repeated curtailment (e.g., a long run of 0's from the **P.Br** scheme) can be realized under this scheme. This may lead to such customers perceiving themselves as being disproportionately penalized, and as a result, scaling down or opting out of the demand response program.

Two approaches can be pursued to rectify this situation. First, a customer-specific service constraint could be imposed on the program operator's signal z decision problem. For instance, one may wish to minimize the probability that the ratio of the non-curtailed load to the maximum load requested by each household over a longer duration, say a day, exceeds a specified bound, such as $P(\sum_{t=1}^{T} L_t(z_t) \geq \alpha \sum_{t=1}^{T} L_t^{\max}) \geq 1 - \epsilon$, where the service level α is offered to the consumer with probability $(1 - \epsilon)$. In another instance, one may express a similar constraint by just limiting the signaling capability without monitoring each household, e.g., for the scheme **P.CU**, satisfy $P(\sum_{t=1}^{T}(1 - z_t + z_t X_t) \geq \alpha T) \geq 1 - \epsilon$ or $E[\sum_{t=1}^{T}(1 - z_t + z_t X_t)] \geq \alpha T$. These types of service measures are (1) hard for the utility to monitor, (2) not transparent from the perspective of a consumer (i.e., unverifiable from an electricity bill), (3) easy to game by a consumer, and (4) computationally intensive on the utility as there can be a constraint for each household or load class that must be guaranteed.

The second, and preferred, approach is to impose a customer satisfaction rule that is based on the realizations of the uncertainty at each household. An implementation of this could allow smart meters to temporarily opt out of the program for one or more period(s) if a multi-period service level constraint gets violated. In scheme **P.Br**, it can be as simple as stipulating that there will be no ℓ successive periods of curtailment. For schemes **P.DU** and **P.CU**, the customer could opt-out if $\sum_{t=1}^{T}(1 - z_t + z_t X_t^{\omega}) > \alpha T$ for some chosen value for α where the ω represents the realization of the r.v.s in successive periods. These transparent service measures are more attractive because the program operator's signal-selection task is simplified: one needs to only statistically estimate the distribution of the number of people opting out of the demand response scheme at any specific time from the history of signals, and incorporate this information in the decision problem.

Consider a PDR scheme implemented for a population of N i.i.d. loads, where O customers have opted out in the current epoch. Then, for the current period, $D(z) = \sum_{k=1}^{O} L_k + \sum_{k=1}^{N-O} L_k(z_k)$. Using the same single-period optimization formulation, the program operator would thus send a *stronger* signal to the remaining $(N - O)$ users in the program in this epoch to achieve the same response as when there is no opt-out strategy. The realization of the r.v. O depends on the history of the signals. But because $O << (N-O)$, the r.v. $D(z)$ is well approximated by replacing O with $E[O]$, which is easy to estimate. For example, in scheme **P.Br** with the provision that customers with ℓ successive curtailments opt out, $E[O] = \Pi_{\tau=1}^{\ell} z_{t-\tau} N$. And for scheme **P.CU** where the fairness criterion is set up as $\sum_{t=1}^{T}(1-z_t+z_t X_t^{\omega}) > \alpha T$, $E[O] = P(\sum_{t=1}^{T}(1 - z_t + z_t X_t) \leq \alpha T)N$.

Observe that the carefully chosen customer fatigue constraint resulted in only very few people having bad sample paths, and a complex multi-period problem was reduced to an enhanced but easily solvable single period problem. This can be generalized further with a new participation map $\xi_k(z)$ of each for customer k in each epoch. The operator observes the $k-$th load response as the composite r.v. $L_k(\xi_k(z))$. The participation (random) map ξ_k is the identity map ($\xi_k(z) = z$) in most instances, but in rare instances may modify z: as a function of the complete sequence of prior curtailment signals, the service level guarantees or other contractually agreed-to conditions, the no-curtailment value ($\xi_k(z) = 1$); or, the map ξ_k may also be used to model any incentives/penalties that discourage serial opt-outs for other reasons by setting $\xi_k(z) < z$. The generalization of ξ_k therefore encapsulates a broad set of multi-period behaviors in the PDR scheme. A detailed treatment of how it may be estimated and used in the context of multi-period models is of further research interest.

V. Simulation Case Study

This section describes a simulation case study of PDR schemes being used with thermostatically controlled loads (TCLs), which are considered suitable for demand response because they provide the ability to shift electricity consumption in time due to thermal capacitance effects [1]. The simulations we present consider a setting where a utility manages a heterogeneous population of residential air-conditioning (AC) load, but the method easily extends to other TCL categories such as electric water heaters and refrigerators.

A. Simulation Setting

We study the operation of $N = 1000$ ACs over a typical warm day in Southern California: Fig. 3 plots the ambient temperature as a function of the hour-of-day. We compare two PDR schemes **P.Br** and **P.CU** against a baseline scheme with no signal-based demand response. Under each scheme, the utility operates using a two-time-scale decision making process. It first solves hourly planning problems at the beginning of each hour to decide the next hour's optimal planned power generation level x^*_{hour}kW. The second process operates every five minutes, where the utility runs demand matching operations. For the baseline, this leads to any excess demand being purchased from the spot market (ancillary services such as regulation and reserve including spinning and non-spinning sources) at m_t prices, whose values over time is plotted in Fig. 3 on the secondary axis. Under either PDR scheme, demand matching also takes advantage of the DR capability by deciding the optimal signal z^*_t to modulate the demand itself.

The joint optimization problem P from Section III is used for both the hourly planning and the five-minute signaling. The baseline scheme solves the planning problem with a no-curtailment demand $D(z)$ model. Under the PDR schemes, the planning solves a joint optimization problem to determine (x^*_{hour}, z^*_{hour}), and the x^*_{hour} is subsequently the fixed power level for the

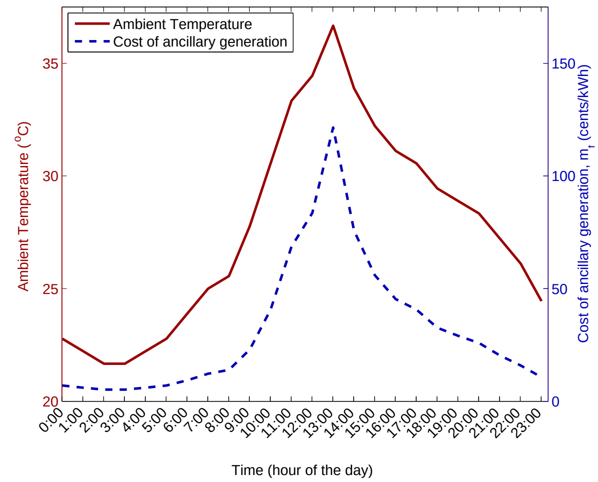

Fig. 3. Ambient temperature and spot market prices vs. time of day.

rest of the hour. Note that the planning problem sets $z^*_{hour} = 0$ for scheme **P.Br**, i.e. assumes no curtailment, which is appropriate in light of Claim 1 and Fig. 2. Demand matching operations under the PDR schemes modulates the suggested optimal signal z^*_{hour} by solving the same optimization problem every 5 min for optimal z_t with the fixed planned power generation level x^*_{hour}. With these assumptions, the planning and signaling problems are convex for all schemes, and can be solved efficiently. The input parameters of the optimization problem P are chosen to fall within the typical ranges observed in CAISO [22], [23]. The cost of planned generation is $c = 8$¢/kWh, the price of electricity is $p = 12$¢/kWh and the price of unused generation treated as the savings from unused fuel is $s = 2$¢/kWh. The hourly planning problems use a blended average spot-market cost $m_{hour} = \frac{1}{12}\sum_{t\in hour} m_t$.

B. Load Model and Forecasting

Accurate modeling and forecasting of heterogeneous TCL loads by using only aggregate easily-observable measurements has been recently proposed [11], [14] for demand response models. While similar methods are of high interest in implementing PDR schemes, our focus here is primarily in emphasizing the benefits of the proposed evaluation framework proposed. We will therefore assume that the utility has perfect knowledge of the individual TCL load at any time t to obtain the aggregate load distribution $D_t(z_t)$ and optimize for the DR signal z_t. The dynamics of each TCL is modeled with the following commonly used ([7], [9]) discrete time state-update equations (We set the time discretization h to observe the thermodynamic evolution at 5 minutes, though this can be chosen to be finer than the demand-signaling interval):

$$\theta_{k,t} = a_k\theta_{k,t-1} + (1-a_k)(\theta^{amb}_{k,t-1} - q_{k,t-1}R_kP_k(z_{t-1})) + \epsilon_{k,t-1}, \text{ and} \quad (5a)$$

$$q_{k,t} = \begin{cases} 0, & \theta_{k,t} < \theta^-_k = \theta^{set}_k - \frac{\delta_k}{2} \\ 1, & \theta_{k,t} > \theta^+_k = \theta^{set}_k + \frac{\delta_k}{2} \\ q_{k,t-1} & \text{otherwise.} \end{cases} \quad (5b)$$

Here, at time t and for the $k-$th device, $\theta_{k,t}$ denotes the actual temperature in the house, $\theta^{amb}_{k,t}$ the ambient temperature and $\theta^{set}_{k,t}$ the desired (or set) temperature. The dimensionless parameter $a_k = e^{-h/C_kR_k}$, where C_k and R_k are respectively the thermal capacitance and resistance. The energy transfer (or usage) rate P_k to operate the TCL is typically a positive constant equivalent to the device's power rating. For PDR schemes, we modify this to a r.v. $P_k(z)$ that depends on the real-time signal z. In particular, in the scheme **P.Br**, $P_k(z) = I_kP_k$, where I_k is 1 with probability $1-z$ and 0 otherwise, and in the scheme **P.CU**, $P_k(z) = (1 - z + zX_k)P_k$ where $X_k \sim U(0,1)$. These simulations do not model the opt-out possibilities discussed in Section IV. The noise process $\epsilon_{k,t}$ models the heat gain or loss that is not explicitly modeled, e.g. due to the opening and closing of doors, or the operation of other loads, or changes in the number of people within the house. We assume $\epsilon_{k,t} \sim \text{Normal}(0, h\sigma^2)$. The TCL control $q_{k,t}$ switches the $k-$th thermostat on or off at time t governed by Eq. (5b), and depends on the width of the dead band δ_k around the set point $\theta^{set}_{k,t}$. The parameters of the heterogenous TCL population are sampled uniformly from the range of parameter values in Table I. Communications delays

Parameter	Definition	Value*
θ^{set}	Temperature set point	20–28^{o}C
δ	Dead-band width	0.75–1.5^{o}C
R	Thermal resistance	1.5–2.5^{o}C/kW
P	Energy transfer rate	10–18kW
η	Load efficiency	2.5
σ	Standard deviation of ϵ	0.01-0.02^{o}Cs$^{-\frac{1}{2}}$

TABLE I
AIR CONDITIONER PARAMETERS.

*Values are for a 250m^2 house adapted from [14]

usually result in households observing the latest signal z_t in a staggered fashion over the next 5 minutes, and this indeed also helps avoid synchronization of the load responses. We simulate this by assigning a delay factor λ_k to the $k-$th device sampled uniformly over $[0,1]$. Each AC's load response at time t is then a combination of the responses to the earlier and current period's signals:

	Scheme	Total Profit ($)	Total Energy Consumed (kWh)	Total Planned Generation (kWh)	Total Reserve (kWh)
All day	**P.CU**	788.5 (1.3)	20,853 (12.8)	21,180 (10.2)	236 (5.1)
	P.Br	763.5 (1.1)	20,679 (16.6)	21,043 (11)	281.2 (6.6)
	Baseline	486.6 (6.6)	20,935 (13.7)	21,043 (11)	773.4 (9.9)
*Peak**	**P.CU**	388.9 (0.8)	10,154 (6.4)	10,365 (5.1)	6.7 (0.5)
	P.Br	371 (1.0)	9,994 (11.9)	10,280 (5.7)	17 (2)
	Baseline	187.7 (6.6)	10,145 (8.4)	10,280 (5.7)	229.2 (7.6)

TABLE II
MAIN RESULTS
Performance averaged over 20 instances (with standard error)

$L_{k,t}(z_t) = \frac{1}{\eta_k}[\lambda_k q_{k,t-1} P_k(z_{t-1}) + (1-\lambda_k) q_{k,t} P_k(z_t)]$, where η_k is the load efficiency parameter. The total power $D_t(z_t)$ for the TCL load population is the corresponding sum over the individual loads.

C. Results

The main results described in Table V-C are obtained by averaging over 20 day-long sample paths (after stabilization). The upper half of the table summarizes the results for the entire day, while the bottom half summarizes it for the duration of the peak period, from 11:00AM through 3:00PM, during which the ambient temperatures were greater than 32^{o}C. The total profit for the utility over a day, which is the objective for our evaluation framework, improves by 62% and 57% for **P.CU** and **P.Br** respectively over the baseline scheme, and by 107% and 97% respectively during the peak period. These savings are in part are due to the decrease in the reserve requirements by 69% and 64% respectively over the baseline scheme, and by 97% and 92% respectively during the peak periods. Observe also that the standard errors of the total profit and total reserve are significantly lower for the probabilistic schemes over the baseline.

The improved performance of the **P.CU** relative to the **P.Br** scheme (both in absolute numbers and in the standard errors) is because the former is able to jointly optimize the planned generation and signal magnitudes, whereas in the latter, in the light of Claim 1 and Fig. 2, the planned generation is the same as in the baseline scheme. The smaller values for the standard errors, in turn, leads to a slightly larger planned generation levels (0.7% over a day and 0.8% during peak) that is offset with a lower reserve requirement and hence a higher profitability. The **P.Br** scheme does have a marginally lower total energy consumption compared to the baseline. This not the case for the **P.CU** scheme, especially during peak hours, because peak shaving is not its primary objective, but rather is achieved by the framework as a trade-off in improving overall profitability. Nevertheless, both probabilistic schemes are seen to be highly promising in terms of their benefits: significantly improved profitability by properly managing the reserve required and minimizing its variability.

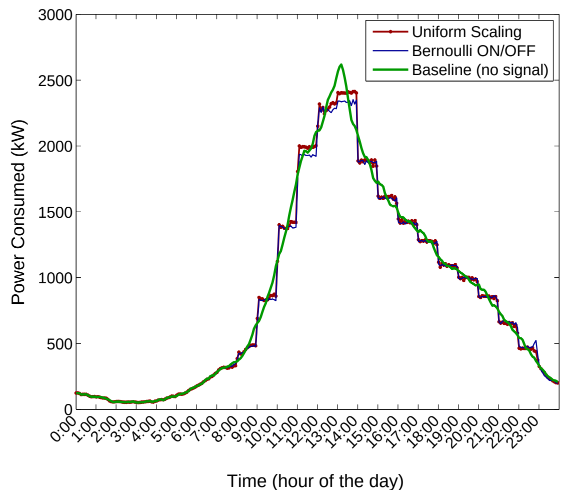

Fig. 4. Total power consumed by TCL population

Fig. 4 plots the total power consumed on a 5-minute time scale throughout the day. We observe a stepwise behavior of the total power consumption for the probabilistic schemes throughout the day, pronounced when the ambient temperatures are above the TCL set-points. This is the combined effect of the hourly planning cycle that fixes the hourly delivered power to x^*_{hour}, and the demand modulation via signals z_t. Under the PDR schemes, the hourly scheduled power is set to be in the range of power consumption expected in the hour (with signaling for **P.CU** and without for **P.Br**), and signals z_t manage to keep consumption below this level throughout the hour. An interesting side-effect is noticedwhen the ambient temperatures are rising: the power consumed is higher under PDR schemes compared to the baseline for the first part of the hour. This is because the curtailment signal z_t is low in this part of the hour because of the gap between base load (same as baseline) and planned level $x*_{hour}$, and so ACs that were curtailed in the earlier periods are recovering to cool the households in the absence of a curtailment signal for the beginning of the hour.

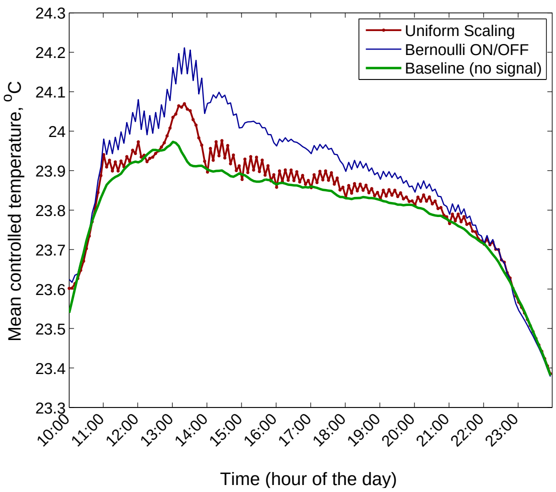

Fig. 5. Average household temperature at peak period

A higher fraction of the ACs are on (i.e. have a $q_{k,t} > 0$ in (5b)) under the PDR schemes during peak periods. This can be attributed to the load curtailments in the DR schemes that effectively result in increasing temperature in the households over the baseline. Fig. 5 shows that the (mean) temperature deviations (across the TCL population) are at most 0.14^oC and 0.29^oC for the **P.CU** and **P.Br** schemes respectively occurring around the peak temperature hour of the day. Thus, the discomfort experienced by the consumers under the PDR schemes is very minimal. (The jagged patterns in the results for the PDR schemes is due to chatter in the TCL operation at the upper boundary of the dead-band, in conjunction with the demand response. These can be smoothed by introducing constraints on TCL control operation.)

Finally, Fig. 6 shows the load curtailment effect of the signals under the PDR schemes, plotting the average fractional reduction in load vs the fraction of the day when reduction of said magnitude were encountered. There is little or no signaling over a large portion of the day (>54%), and large signals are encountered over proportionately smaller fractions of the overall day. Note that **P.CU** curtails only up to half the base load.

VI. Conclusions

In this paper, we describe a new class of PDR schemes that induce load randomization as part of the demand response, and provide three examples of such PDR schemes. We introduce a methodological framework to

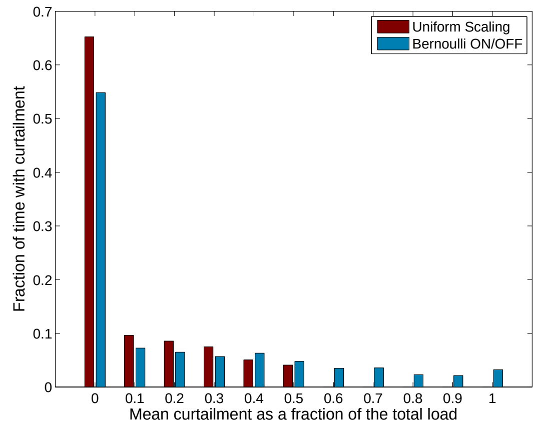

Fig. 6. Bar chart of the fraction of time against the mean depth of curtailment in the two PDR schemes, corresponding to signal z.

implement and analyse PDR schemes, and obtain structural results within this framework to help judge their suitability. In particular, using the proposed framework, we show that one of the example PDR schemes (the *Bernoulli ON/OFF scheme*, **P.Br**) that was previously also implemented in a smart grid demonstration [13], has certain limitations. The other two example PDR schemes are free of these limitations within the framework, and yet can be implemented with the same minimal infrastructure, communication and control requirements as the **P.Br** scheme. And finally, we provide a detailed simulation case-study that illustrates the advantages of PDR schemes for a realistic residential scenario.

Further theoretical analysis and simulation studies, along with practical field experiments will be of considerable interest for the development of PDR schemes for the smart grid. Some other topics of future interest concern the use of signaling to increase consumption to accommodate increased availability of power, say due to excess wind; and an in-depth analysis of how consumer incentives, economic or other, can be incorporated into the design of PDR schemes for increasing consumer acceptance and program participation.

Appendix A
Proof of Claim 1

Proof: We prove the first part of the theorem by counterexample. Suppose the loads, L_k, are i.i.d. random variables.By design, the signal-induced uncertainties are i.i.d as well. Since N the number of customers in each load class is large, from the central limit theorem, the aggregate demand $D(z)$ can be approximated as a normal r.v. with mean $\mu(z) = N\mu_L(1-z)$ and variance $\sigma^2(z) = N[(1-z)\sigma_L^2 + z(1-z)\mu_L^2]$. For a normally distributed r.v., $X \sim \mathcal{N}(\mu, \sigma^2)$,

$$\text{CVaR}_X(\beta) = \mu + h(\beta)\sigma, \quad (6)$$

where the function $h(\beta) = (\sqrt{2\pi}\exp(\text{erf}^{-1}(2\beta-1))^2(1-\beta)^{-1}$ which depends only the distribution function for a standard Normal r.v., $\mathcal{N}(0,1)$. Substituting Eq. (6), the expected profit reduces to:

$$\pi(x^*(z),z) = (p-c)\mu(z) - (c-s)h(\beta)\sigma(z). \quad (7)$$

It can be verified that $\sigma(z)$ is a concave function because $\sigma^2(z)$ is a concave function. This implies that $\pi(x^*(z),z)$ is a convex function of z (rather than concave as desired) and hence the result.

To prove the second part of the theorem we perform the second derivative test with respect to z. For this scheme since $E[D(z)]$ is linear, it suffices to show that $g(z) := E[(D(z)-x)^+]$ is convex in z.

$$g(z) = E[(D(z)-x)^+] = \int_x^\infty [1 - F_{D(z)}(y)]\, dy, \quad (8)$$

where $F_{D(z)}(.)$ is the distribution function of $D(z)$ We want to show that $g''(z) = -\int_x^\infty \frac{\partial^2 F_{D(z)}(y)}{\partial z^2} dy \geq 0$.

$$\begin{aligned} F_{D(z)}(x) &= P\left(\sum_{i=1}^N I_i L_i \leq x\right) \\ &= \sum_{k=0}^N \sum_{j_k \in J_k} z^k (1-z)^{N-k} F_{\bar{L}_{j_k}}(x). \end{aligned}$$

where J_k is the set of all possible $(N-k)$ different loads that are turned on. We represent the aggregate load that is turned on by $\bar{L}_{j_k}$ for simplicity. Observe that when the loads are i.i.d the second summation can be replaced by $\binom{N}{k}$.

$$\begin{aligned} \frac{\partial F_{D(z)}(x)}{\partial z} &= \sum_{k=0}^N \sum_{j_k \in J_k} z^{k-1}(1-z)^{N-k-1}[k - Nz]\, F_{\bar{L}_{j_k}}(x) \\ &= \frac{1}{1-z} \sum_{i=1}^N \left[P(B_i \leq x) - P\left(\sum_{k=1}^N I_k L_k \leq x\right)\right] \\ &= \sum_{i=1}^N \left[P(B_i \leq x) - P(L_i + B_i)\right], \end{aligned}$$

where $B_i = \sum_{\substack{k=1 \\ k \neq i}}^N I_k L_k$.

$$\begin{aligned} \frac{\partial^2 F_{D(z)}(x)}{\partial z^2} &= \sum_{k=0}^N \sum_{j_k \in J_k} z^{k-2}(1-z)^{N-k-2} F_{\bar{L}_{j_k}}(x) * \\ &\quad \left[n(n-1)z^2 - 2k(N-1)z + k(k-1)\right] \\ &= \frac{1}{(1-z)^2} \sum_{i \neq j} \left[P\left(\sum_{k=1}^N I_k L_k \leq x\right) - P(B_i \leq x)\right. \\ &\quad \left. - P(B_j \leq x) + P(A_{i,j} \leq x)\right] \\ &= \sum_{i \neq j} \left[P(L_i + L_j + A_{i,j} \leq x) + P(A_{i,j} \leq x)\right. \\ &\quad \left. - P(L_i + A_{i,j} \leq x) - P(L_j + A_{i,j} \leq x)\right], \end{aligned}$$

where $A_{i,j} = \sum_{\substack{k=1 \\ k \neq i,j}}^N I_k L_k$.

Therefore, in order to show that $g''(z) \geq 0$, it suffices to show the following that for any i, j

$$\begin{aligned} &E[(L_i + L_j + A_{i,j} - x)^+] - E[(L_i + A_{i,j} - x)^+] \\ &- E[(L_j + A_{i,j} - x)^+] + E[(A_{i,j} - x)^+] \geq 0. \end{aligned}$$

The transformation to the expectation from the distribution function is similar to that used in Eq. (8). We will prove the above equation for every realization of the r.v.'s and hence the result will be true also in expectation. Consider a realization of the r.v.'s for which we use the same notation for simplicity. Without loss of generality, say, $L_i \leq L_j$. We also note that all the r.v.'s are non-negative. Consider the following different cases: (a) $A_{i,j} \leq x$ (b) $L_i + A_{i,j} \leq x$ (c) $L_j + A_{i,j} \leq x$ and (d) $L_i + L_j + A_{i,j} \leq x$, we can show that

$$\begin{aligned} &(L_i + L_j + A_{i,j} - x)^+ - (L_i + A_{i,j} - x)^+ \\ &- (L_j + A_{i,j} - x)^+ + (A_{i,j} - x)^+ \geq 0. \end{aligned}$$

Since the above is true for every realization of the r.v.'s, it is also true in expectation, and hence the proof. ■

Appendix B
Proof of Claim 4

Proof: For i.i.d. loads and large N, from the central limit theorem, we have $\mu(z) = N\mu_L\left(1 - \frac{z}{2q}\right)$ and variance $\sigma^2(z) = \frac{N}{q^2}\left[\left(q - \frac{z}{2}\right)^2 \sigma_L^2 + \frac{(z+1)^2 - 1}{12}(\mu_L^2 + \sigma_L^2)\right]$. We now have to identify conditions when $\sigma(z)$ is a convex function in z. With these, using Eq. (7) for a normal distribution, we can deduce that $\pi(x,z)$ is jointly concave in (x,z).

Suppose $h(z) = \sqrt{f(z)}$ and $f(z) \geq 0\ \forall z$ then $h(z)$ is convex if and only if $2f''(z)f(z) > [f'(z)]^2$. Translating these conditions to the case when $f(z) = \sigma^2(z)$, we get

$$f'(z) = \frac{N}{q^2}\left[-\left(q - \frac{z}{2}\right)\sigma_L^2 + \frac{(z+1)}{6}(\mu_L^2 + \sigma_L^2)\right],$$

$$f''(z) = \frac{N}{q^2}\left[\frac{1}{2}\sigma_L^2 + \frac{(z+1)}{6}(\mu_L^2 + \sigma_L^2)\right],$$

and finally simplifying, we get $2f''(z)f(z) - [f'(z)]^2 > 0$ if and only if $\sigma^2(12q^2 + 12q - 1) - \mu^2 > 0$. This proves the first part of the result.

For large q, scheme 2 is a special case of scheme 3, where we show that the signal can be separated from the noise using the $U[0,1]$ distribution. This result then follows from Claim 2. ■

Appendix C
Proof of Claim 5

Proof: From Eq. (6) and Eq. (7), we know that

$$\pi(x^*(\mathbf{z}), \mathbf{z}) = (p-c)\left[\sum_{i=1}^M \mu_i(z_i)\right] - (c-s)h(\beta)\sigma(\mathbf{z}), \quad (9)$$

where $\sigma^2(\mathbf{z}) = \sum_i \sigma^2(z_i)$. The first order conditions with an affine mean simplifies to

$$\frac{\sigma_i(z_i)}{\sigma(\mathbf{z})}\frac{\partial \sigma_i(z_i)}{\partial z_i} = \frac{(p-c)b_i}{(c-s)h(\beta)}, \qquad \forall i = 1, ..., M,$$

where $\mu_i(z_i) = a_i + b_i z$. We compare the solution to the above equations to the case when $\sigma(\mathbf{z}) = \sigma(z_i)$. The difference is in the multiplier to the partial derivative which happen to be less than 1. With each $\sigma_i(z_i)$ being convex, it is easy to deduce by plotting the σ_i's and identifying the optima that each class i gets curtailed by a smaller amount when optimized together than separately. ■

REFERENCES

[1] D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.

[2] B. Kirby, "Spinning reserve from responsive loads," Oak Ridge National Laboratory, Tech. Rep. ORNL/TM-2003/19, 2003.

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

[4] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *IEEE Transactions on Power Systems*, 2011, accepted.

[5] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *IEEE Power & Energy Society General Meeting*, 2011.

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

[7] B. Ramanathan and V. Vittal, "A framework for evaluation of advanced direct load control with minimum disruption," *IEEE Transactions on Power Systems*, vol. 23, no. 4, pp. 1681–1688, 2008.

[8] K.-Y. Huang and Y.-C. Huang, "Integrating direct load control with interruptible load management to provide instantaneous reserves for ancillary services," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1626–1634, 2004.

[9] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.

[10] S. Bashash and H. K. Fathy, "Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads," in *American Control Conference*, 2011, pp. 4546–4553.

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

[12] K. Bhattacharyya and M. Crow, "A fuzzy logic based approach to direct load control," *IEEE Transactions on Power Systems*, vol. 11, no. 2, pp. 708–714, 1996.

[13] D. Hammerstrom *et al.*, "Pacific northwest gridwise testbed demonstration projects part I. olympic peninsula project." Pacific Northwest National Laboratory, Tech. Rep. PNNL No. 17167, 2007.

[14] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.

[15] P. L'Ecuyer, "Random number generators," in *Encyclopedia of Operations Research and Management Science*, 3rd ed., S. I. Gass and M. C. Fu, Eds. Springer-Verlag, 2012.

[16] P. Varaiya, F. F. Wu, and J. W. Bialek, "Smart operation of the smart grid: Risk-limiting dispatch," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 40–57, 2011.

[17] R. Billinton and R. N. Allan, *Reliability evaluation of power systems*. Plenum Press, 1996.

[18] D. Simchi-Levi, X. Chen, and J. Bramel, *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management*. Springer, 2005.

[19] N. C. Petruzzi and M. Dada, "Pricing and the newsvendor problem: A review with extensions," *Operations Research*, vol. 47, no. 2, pp. 183–194, 1999.

[20] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–41, 2000.

[21] A. Kocabiyikolu and I. Popescu, "An elasticity approach to the newsvendor with price-sensitive demand," *Operations Research*, vol. 59, no. 2, pp. 301–312, 2011.

[22] "California ISO's OASIS," 2012, energy prices. [Online]. Available: http://oasis.caiso.com/mrioasis/logon.do;jsessionid=8C17DCC679A3A38C54C17E556F2FC494

[23] "Ancillary services market," CAISO, Tech. Rep., April 2007. [Online]. Available: http://www.caiso.com/1bb7/1bb77a7a16e20.pdf

PLACE PHOTO HERE

Pavithra Harsha is a member of the Business Analytics and Mathematics (BAMS) department at the IBM T. J. Watson Research Center since 2011. She is also a visiting scientist at the Sloan School of Management, MIT. She received a Ph.D. in Operations Research from MIT in 2008 and a B.Tech. from IIT Madras in 2003. Her current research interests are in the areas of pricing, supply chain, robust and stochastic optimization with applications to the smart grid and retail.

PLACE PHOTO HERE

Mayank Sharma received a B. Tech. from the Indian Institute of Technology, Delhi, India in 1997; and a M.S. and Ph.D. from Stanford University in 2005, all in Electrical Engineering. Since 2005, he has been a Research Staff Member in the Stochastic Processes and Applications group within the Business Analytics and Mathematical Sciences department at the IBM T. J. Watson Research Center, Yorktown Heights, NY. His interests include the mathematical analysis, modeling and optimization of complex stochastic systems, drawing from areas such as stochastic processes theory, probability theory, stochastic optimization and control, algorithm design, information theory, and mathematical statistics. Dr. Sharma is a recipient of the 2010 Daniel H. Wagner Prize for Excellence in Operations Research Practice. He is a member of the ACM, DIMACS, IEEE, and INFORMS.

PLACE PHOTO HERE

Ramesh Natarajan is a member of the Business Analytics and Mathematics (BAMS) department at the IBM T. J. Watson Research Center. His research interests are in the areas of statistical modeling, data mining, database applications and parallel computing. He is the recipient of 3 IBM research division awards and the Fourth Plateau Invention award.

PLACE PHOTO HERE

Soumyadip Ghosh is a member of the Mathematics division (BAMS) at the IBM T. J. Watson Research Center. His current research interests lie in simulation based optimization techniques for stochastic optimization problems, with a focus on applications in Energy and Power systems and supply chain management.