2019 IEEE 58th Conference on Decision and Control (CDC)
Palais des Congrès et des Expositions Nice Acropolis
Nice, France, December 11-13, 2019

Customer Incentives for Gaming Demand Response Baselines

Douglas Ellman and Yuanzhang Xiao

Abstract— **We present a novel multi-stage stochastic dynamic programming model that optimizes customer actions under a baseline-based demand response program, in order to discover incentives for gaming behaviors. Customers can decide how much to increase or decrease electricity consumption at each time period in order to maximize expected rewards, with imperfect knowledge of when future demand response events will occur. Customers may decide to take action during events (desired behavior) and during corresponding hours of non-event days in order to manipulate baselines (gaming behavior). Analytical results for special cases show fundamental drivers of gaming incentives. Simulation results reveal gaming incentives and impacts to program performance for a real-world baseline-based demand response program, and how incentives vary with changing program and customer parameters.**

I. INTRODUCTION

Demand response (DR) aims to modify customers' electricity consumption in order to improve operation of the electric grid (e.g., reduce costs, improve reliability, reduce emissions). A wide variety of DR mechanisms have been proposed or implemented [1]. Most of these mechanisms fall into three categories: price-based DR, auction-based DR, and incentive-based DR. In price-based DR (e.g., real-time pricing [2][3]), the price of electricity varies over time, which encourages customers to reduce consumption when prices are high. In auction-based DR, the customers submit bids (e.g., supply functions [4], utility of consumption [5][6]) to compete for payment for demand reduction.

Incentive-based DR, where customers are paid incentives based on participation or performance in the DR program, is widely implemented [7]. An important class of incentive-based DR is baseline-based DR [8][9][10][11][12][13], where incentive payments are based on the difference between baseline consumption (e.g., average consumption during previous days) and consumption during the DR event. Baseline-based DR can be implemented as an opt-in program that does not require changing customer electricity rates, which can make public acceptance and regulatory approval easier to obtain. Baselines are also central to DR participation in U.S. wholesale energy markets, as regulated by the Federal Energy Regulatory Commission [14]. Given its prevalence in practice, baseline-based DR will be the focus of our paper.

With the growing use of baseline-based DR, it is important to understand incentives and behavior of participating customers, since customer actions are the primary determinant of overall DR program performance and cost. Baseline-based DR can create complicated customer incentives, including incentives to take action both during DR events (i.e., desired behavior) and during non-event times in order to manipulate baselines (i.e., gaming behavior). In general, gaming behavior can be difficult to predict, because optimal decisions for each time period are interrelated and can depend on the customer's current baseline consumption, customer costs, the demand response program structure, and uncertain future values (such as when DR events will occur).

It is not new that baseline-based DR is susceptible to gaming behavior. Randomized control trials [15] and legal settlements [16] provide real-world evidence that customers may game baselines. Interested readers can also see [17] for a high-level discussion of the potential for inflating baselines in baseline-based DR. In cases with incentives for gaming behaviors, ignoring these behaviors can lead to significant errors in projections of required DR quantities, DR costs, and demand forecasts. As the use of DR for critical grid services grows and as customer responses to DR programs become increasingly sophisticated, these errors could cause significant impacts to the cost and reliability of the grid.

Although we are aware of the potential for gaming behavior in baseline-based DR, we may not yet have satisfactory understandings of this behavior. Some works, such as [8], compare alternative DR baseline structures based on metrics such as accuracy of the baseline, but do not consider gaming incentives. Some works [9][11][12] recognize the potential for gaming behaviors in baseline-based DR and propose mechanisms to mitigate gaming such as self-reported baselines and profit sharing, but do not analyze customer incentives under commonly implemented baseline-based DR frameworks. The work closest to ours is [10], which models baseline-based DR as a multistage stochastic decision problem. However, it only provides results when there is one single baseline stage and one DR event, and the customer knows with certainty when the DR event will occur. Therefore, to better understand gaming behavior, we need a more comprehensive model that includes critical factors in the customer's decision making process, such as the uncertainty of the demand response event schedule and multiple baseline and event stages. As a demonstration of the importance of more comprehensive models, we will show non-obvious results (in Section IV) on how penalty (i.e., negative payment for increasing the load in the DR event) affects the system performance.

We build on the prior work and develop a novel model to identify optimal customer decisions under a baseline-based DR program. Our model solves the customer's multistage stochastic decision problem with a dynamic programming algorithm, thus providing an optimal policy for how much

The authors are with Hawaii Advanced Wireless Technologies Institute and Department of Electrical Engineering, University of Hawaii at Manoa, Honolulu, HI 96822, USA. Email:{dellman, yxiao8}@hawaii.edu.

978-1-7281-1398-2/19/$31.00 ©2019 IEEE

to reduce or increase electricity consumption based on time, current customer baseline, and whether a DR event is active. The key contribution of our approach is representation of the customer's imperfect knowledge of when DR events will occur over the course of a DR season, which can have significant impacts on the optimal policy and is a primary uncertainty for real-world DR participants. Our major contributions are as follows.

- To the best of our knowledge, we propose the first baseline-based DR model that includes a variety of commonly implemented baseline-based DR structures with *multi-day* baselines and *multiple uncertain* events over a DR season.
- For a special, yet widely adopted, class of baseline-based DR, we provide structural results on the customer's optimal policy, which sheds light on how to efficiently design the DR program.
- For a general baseline-based DR model, we develop an efficient dynamic programming algorithm by exploiting the structure of the problem, which is capable of providing solutions quickly for extended DR seasons of several months or longer.

The remainder of this paper is organized as follows: Section II describes the model, Section III provides analytical results for special cases, Section IV provides numerical results from a computer implementation of the model, and Section V summarizes our conclusions.

II. MODEL

A. System Setup

We study a demand response program that occurs over a demand response season of $T \in \mathbb{N}$ days. At the start of each day $t \in \{1,\dots,T\}$, the utility notifies the customer whether the current day is a demand response *event day*. On event days, the utility desires participating customers to reduce electricity consumption during a pre-specified window of hours (e.g., 2 p.m – 6 p.m. on all event days). We denote the probability that day t is an event day by $p_{E,t} \in [0,1]$.

In each day t, the customer has a default load level l_t, and may deviate from the default load level by a_t, resulting in an actual load level of $a_t + l_t$. The default load level l_t is a model parameter, and the deviation a_t is the customer's decision variable. We model the deviation, instead of the actual load, as the decision variable in order to better represent the customer's decision of whether to increase or decrease the load. We denote the set of available load deviation by $A_t \subset \mathbb{R}$, a compact set that includes both positive and negative values.

On an event day t, the customer receives a rebate from the utility for load reduction. We define the rebate function as

$$\begin{aligned} r_t \quad : \quad & \mathbb{R} \times A_t \to \mathbb{R} \\ & (\bar{s}_{B,t}, a_t) \mapsto r_t(\bar{s}_{B,t}, a_t), \end{aligned} \tag{1}$$

which depends on the *baseline* $\bar{s}_{B,t}$ and the actual load $l_t + a_t$. The baseline is usually determined by the utility according to the actual load over the previous N_B non-event days. We denote the load over the N_B non-event days prior to day t by a vector $\mathbf{s}_{B,t} \in \mathbb{R}^{N_B}$, and write the rule of determining the baseline as a function

$$f : \mathbf{s}_{B,t} \mapsto \bar{s}_{B,t}. \tag{2}$$

Often, the baseline $\bar{s}_{B,t}$ is defined as the average load of the preceding N_B non-event days or the average load over $N'_B < N_B$ days selected based on some rule [8]. A rebate can be calculated by

$$r_t(\bar{s}_{B,t}, a_t) = r_{DR,t} \cdot [\bar{s}_{B,t} - (l_t + a_t)], \tag{3}$$

where $r_{DR,t} \in \mathbb{R}_+$ is the rebate per unit of load reduction. In this rebate scheme, the rebate could be negative, which means that the customer could pay penalty for increasing load relative to the baseline. In some demand response programs, the rebate is capped above zero, so the rebate is calculated by

$$r_t(\bar{s}_{B,t}, a_t) = r_{DR,t} \cdot [\bar{s}_{B,t} - (l_t + a_t)]^+. \tag{4}$$

For example, [18] caps performance payments above zero for each aggregation of participants, but allows individual customers within an aggregation to make negative contributions. A reasonable rebate should be non-decreasing in the baseline and non-increasing in the actual load, in order to incentivize the customer to reduce the load.

The customer incurs a cost when deviating from the default load. We define the cost of changing the load on day t as

$$c_t : A_t \to \mathbb{R}. \tag{5}$$

A typical cost function could be a piece-wise linear function

$$c_t(a_t) = c_{R,t} \cdot (-a_t)^+ + c_{I,t} \cdot (a_t)^+, \tag{6}$$

where $c_{R,t}$ is the marginal cost of reducing the load on day t, $c_{I,t}$ is the marginal cost of increasing the load on day t, and $(\cdot)^+ = \max\{\cdot, 0\}$. The cost of reducing the load may be the discomfort cost of foregone electricity use (e.g., the discomfort of turning off the air conditioning) [1], and the cost of increasing the load may be the additional payment for electricity. Lower-cost options to increase and reduce load during DR windows include load shifting [1], which can be done via rescheduling the use of appliances or through energy storage.

The goal of the customer is to maximize its expected total profit (i.e., rebate minus cost) during the demand response season. Therefore, the customer has incentives to reduce the load during the event days if the rebate is large enough to offset the cost of reducing load. Moreover, the customer may also increase the load during non-event days in order to inflate the baseline and thus increase the future rebate, although increasing the load would only result in a higher cost in that non-event day. In other words, the customer may "game the system" during non-event days. We are interested in whether and to what extent such gaming behavior occurs.

B. Dynamic Programming Formulation

We formulate the customer's decision making problem as a stochastic dynamic program over a finite horizon of T days. Next, we describe the key components of the dynamic program in detail.

The state at each day t consists of the baseline state $\mathbf{s}_{B,t} \in \mathbb{R}^{N_B}$ and the event state $s_{E,t} \in \{0,1\}$. As discussed before, the baseline state $\mathbf{s}_{B,t}$ is a vector of the load levels on the previous N_B non-event days prior to day t. The event state $s_{E,t}$ indicates whether day t is an event day, with $s_{E,t} = 1$ indicating that day t is an event day. We write the complete state for day t as $s_t = (s_{B,t}, s_{E,t})$. The action at each day t is $a_t \in A_t$, the deviation from the default load.

The state s_t and the action a_t determine the next baseline state $s_{B,t+1}$. If day t is an event day, the action a_t will not be counted in determining the future baselines. In this case, the baseline state $\mathbf{s}_{B,t+1} = \mathbf{s}_{B,t}$ stays the same, and the event state $s_{E,t+1}$ follows Bernoulli distribution with parameter $p_{E,t+1}$.

If day t is a non-event day, the action a_t will be counted towards determining future baselines. In this case, the baseline state $\mathbf{s}_{B,t+1}$ is a concatenation of the last $N_B - 1$ elements of the previous baseline state $\mathbf{s}_{B,t}$ and the actual load $l_t + a_t$. Mathematically, we have

$$\mathbf{s}_{B,t+1} = \left([\mathbf{s}_{B,t}]_{2:N_B}, l_t + a_t\right), \tag{7}$$

where $[\mathbf{s}_{B,t}]_{2:N_B}$ is the vector of last $N_B - 1$ elements of $\mathbf{s}_{B,t}$. Similar as in the event day, the event state $s_{E,t+1}$ is drawn from the Bernoulli distribution with parameter $p_{E,t+1}$.

The reward function $u_t : \mathbb{R}^{N_B} \times \{0,1\} \times A_t \to \mathbb{R}$ is the rebate minus the cost, namely

$$u_t(s_t, a_t) = s_{E,t} \cdot r_t\left(f(\mathbf{s}_{B,t}), a_t\right) - c_t(a_t). \tag{8}$$

Without loss of optimality, we focus on Markov policies, which depend on the current state and time only. We write a Markov policy as

$$\pi = (\pi_1, \ldots, \pi_T) \text{ with } \pi_t : \mathbb{R}^{N_B} \times \{0,1\} \to A_t, \tag{9}$$

and the set of all Markov policies as Π_M.

Our goal is to find the optimal Markov policy that maximizes the expected total reward, namely solving the following dynamic program

$$\pi^* = \arg\max_{\pi \in \Pi_M} \mathbb{E}^{\pi}\left\{\sum_{t=1}^{T} u_t(s_t, a_t)\right\}, \tag{10}$$

where the expectation $\mathbb{E}^{\pi}$ depends on the policy π.

The dynamic program (10) can be solved through backward induction.

III. ANALYSIS

In this section, we focus on a special case of the dynamic program (10), which we can solve analytically. As a result, we can provide structural results on the customer's optimal policy. Our structural results will provide insights on the customer's decision making process and on the design of the demand response program.

We consider the demand response programs where the baseline is determined as the average load of the previous N_B non-event days, namely

$$\bar{s}_{B,t} = f(\mathbf{s}_{B,t}) = \frac{\sum_{i=1}^{N_B} [\mathbf{s}_{B,t}]_i}{N_B}, \tag{11}$$

where $[\mathbf{s}_{B,t}]_i$ is the i-th element of the vector $\mathbf{s}_{B,t}$. Note that this is a common way to determine the baseline in many demand response programs [8]. We also assume that the rebate function is linear as in (3). The linear rebate could apply, for example, to an individual customer in [18] that is part of an aggregation.

Before stating our analytical results, we need to define a few useful quantities. The first one is the myopically optimal action in an event day:

$$a_{E,t} = \arg\max_{a_t \in A_t} r_t\left(\bar{s}_{B,t}, a_t\right) - c_t(a_t). \tag{12}$$

For each non-event day t, an important quantity, denoted by M_t, is the expected number of event days whose baselines depend on the load a_t in day t. Given M_t, we define

$$a_{N,t} = \arg\max_{a_t \in A_t} \frac{M_t \cdot r_{DR,t}}{N_B} \cdot a_t - c_t(a_t). \tag{13}$$

Proposition 1: Assume that the rebate is linear as in (3), and that the baseline is determined according to (11). The optimal policy π^* satisfies that for all $t = 1, \ldots, T$,

$$\pi_t(\mathbf{s}_{B,t}, 1) = a_{E,t} \text{ and } \pi_t(\mathbf{s}_{B,t}, 0) = a_{N,t}. \tag{14}$$

Proof: Please see the online appendix [19]. ■

Proposition 1 characterizes the customer's optimal policy. In event days, the customer will choose an action that myopically optimizes the current reward. This is reasonable because the load in the event day will not be counted towards establishing future baselines. Therefore, the customer will focus on the current reward without worrying about the impact of its action on the future rewards. From the utility's perspective, our result also guarantees the simplicity of designing the rebate scheme: to ensure certain level of demand response, the utility needs to consider the customer's cost of changing load only, but not other factors such as the probabilities of events, the number of non-event days used to calculate the baseline, and so on.

In non-event days, the customer will choose the action $a_{N,t}$ defined in (13). Note that this is not a myopic action, because the current reward in the non-event day is $-c_t(a_t)$. The objective function in (13) includes a term $\frac{M_t \cdot r_{DR,t}}{N_B} \cdot a_t$, which represents the expected future benefit of increasing the current load by a_t. The benefit comes from the inflated future baselines, which would result in higher future rebates. By exploiting the structure of the problem, we are able to characterize this benefit analytically as $\frac{M_t \cdot r_{DR,t}}{N_B} \cdot a_t$, leading to sharp structural results.

Proposition 1 also suggests that we can solve the dynamic program without using backward induction. In particular, we can avoid the "curse of dimensionality" if the problems in (12) and (13) are easy to solve, which is often the case.

For example, when the rebate scheme is linear as in (3), the problem in (12) reduces to

$$a_{E,t} = \arg\max_{a_t \in A_t} -r_{DR,t} \cdot a_t - c_t(a_t), \tag{15}$$

which admits analytical solutions if the cost function is linear or quadratic. Similarly, we can solve the problem in (13) analytically if the cost function is linear or quadratic.

The remaining difficulty lies in how to compute M_t. Note that M_t depends only on the Bernoulli distributions of the event states only, but not on the baseline states, the actions taken, etc. In other words, we can compute M_t based on the probabilities $p_{E,1}, \ldots, p_{E,T}$. For some special cases, we are able to obtain analytical expressions for M_t.

Lemma 2: For a non-event day t, the expected number of event days whose baselines depend on the load a_t in day t can be computed analytically in the following cases.

- When $N_B = 1$, we have

$$M_t = \sum_{s=t+1}^{T} (s-t) \cdot \left(\Pi_{\tau=t+1}^{s} p_{E,\tau}\right) \cdot (1 - p_{E,s+1}). \tag{16}$$

- When the probabilities of events are the same (i.e., $p_{E,t} = p_E, \ \forall t$), we have

$$M_t = \sum_{r=1}^{T-t} \sum_{s=(r-N_B)^+}^{r-1} \binom{r-1}{s} \cdot p_E^{s+1} \cdot (1-p_E)^{r-s-1}. \tag{17}$$

Proof: Please see the online appendix [19]. ■

IV. Simulation Results

In this section, we present numerical results from a computer implementation of the model. We model several scenarios based on hypothetical customers in a DR program like Consolidated Edison's "Commercial System Relief Program - Voluntary Option," or "CSRP-V" (see Rider T in [20], [18][21][22]). CSRP-V is a baseline-based DR program with day-ahead event notice, thus matching the structure of our model. Results for the selected scenarios demonstrate how our model can reveal non-obvious relationships between program design, customer characteristics, and customer actions.

A. Scenario Descriptions

The model parameters for the scenarios are summarized in the upper section of Table I, and are described below.

We model customers participating in CSRP-V for the summer DR season (roughly 150 days). CSRP-V pays \$3 per kWh of demand reduction during a pre-assigned 4-hour window on event days. For direct participants, payments are capped above zero, but for participants represented by an aggregator, negative event performance is counted and netted against contributions from other customers. The utility advises to expect 3 events per year, on peak load days. CSRP-V offers a "5 in 10" baseline methodology, where baseline load is average load over the 5 highest load days out of the preceding 10 similar days. We also consider an alternative baseline of the average load over the prior 5 similar days, or "5 in 5." For the "5 in 10" and "5 in 5" baseline scenarios, we include 10 or 5 days, respectively, before the start of the DR season in order to establish an initial baseline.

Scenario 1 models a customer who cannot take gaming actions. The customer has three choices for actions and costs each day: (1) no load change, \$0; (2) reduce load by 1 kWh, \$0.02; or (3) reduce load by 2 kWh, \$2.02. This represents a customer who can use a battery to shift 1 kWh of energy (paying 10% of a \$0.20/kWh electricity rate for efficiency losses), and can reduce appliance use to lower consumption 1 kWh more (costing \$2/kWh for lost utility). The customer has no additional insight about when events will occur, so the probability of event each day is the total number of expected events per season divided by the number of days in the season (3/150). The customer is not represented by an aggregator, so negative payments are not allowed, and we set the default load to be the same each day.

Scenario 2 adds two additional action and cost options: (4) increase load by 1 kWh, \$0.02; or (5) increase load by 2 kWh, \$0.22. This represents a customer who can use a battery to shift 1 kWh of load (again paying for 10% losses), and can use appliances to increase load an additional 1 kWh (costing the electricity rate of \$0.20/kWh).

Scenario 3 is a modification of scenario 2, representing a customer with more information about when events are likely to occur. This represents a customer that predicts event based on weather forecasts. This scenario includes two consecutive days with 0.5 probability of event (high probability days), and lower event probability for remaining days, so the total number of expected events is still 3.

Scenarios 4 through 8 include the alternative "5 in 5" baseline. Scenario 4 is identical to scenario 2 besides the alternative baseline type. Scenarios 5 through 8 include large day-to-day variations in default load, where the deviations are greater than the available customer load modifying actions. Scenarios 5 and 6 compare the impact of variable default load with negative payments not allowed and negative payments allowed, respectively. Scenarios 7 and 8 are identical to scenarios 5 and 6, except the customer has a less efficient battery with 25% round trip losses, so the initial 1 kWh of reduction or increase costs \$0.05.

B. Model Outputs and Findings

Selected numerical results are presented in Table I. Figures 1 and 2 show daily probability of event (left axis) and mean optimal policy for event and non-event days (right axis) for each scenario. The mean optimal policy is an average over the baseline states, which enables visualizing the policy when the number of states is large. The results reveal several findings, some of which corroborate our analytical results, and some of which address more complicated cases that would be difficult to analyze manually.

Scenario 1 shows that when gaming is not possible, the true DR quantity (load reductions due to actions during events) matches the apparent DR quantity (the difference between the potentially inflated baseline load and the event load), so that the utility payment per unit of true DR matches the CSRP-V program incentive rate (\$3/kWh). Alternatively, in scenario 2, the customer has the ability and incentive to take gaming actions. This case has the same true DR quantity,

TABLE I

SIMULATION SCENARIOS AND RESULTS

Scenario Number	1	2	3	4	5	6	7	8
Gaming Allowed	no	yes	yes	yes	yes	yes	yes	yes
Probability of Event	flat	flat	spike	flat	flat	flat	flat	flat
Negative Payments	no	no	no	no	no	yes	no	yes
Default Load	flat	flat	flat	flat	variable	variable	variable	variable
Baseline Type	5 in 10	5 in 10	5 in 10	5 in 5	5 in 5	5 in 5	5 in 5	5 in 5
Customer Costs	standard	standard	standard	standard	standard	standard	higher	higher
Expected True DR (kWh)	6.0	6.0	6.0	6.0	3.7	6.0	2.9	6.0
Expected Apparent DR (kWh)	6.0	9.0	10.1	9.0	6.3	8.9	2.9	8.8
Expected DR Payments ($)	18.0	27.0	30.3	27.0	18.9	26.6	8.8	26.4
Expected Customer Costs ($)	6.1	7.6	8.6	9.0	6.4	9.0	3.3	13.3
Expected Customer Net Benefits ($)	11.9	19.4	21.8	17.9	12.5	17.6	5.4	13.1
Expected Payment per Unit True DR ($/kWh)	3.0	4.5	5.1	4.5	5.2	4.4	3.0	4.4

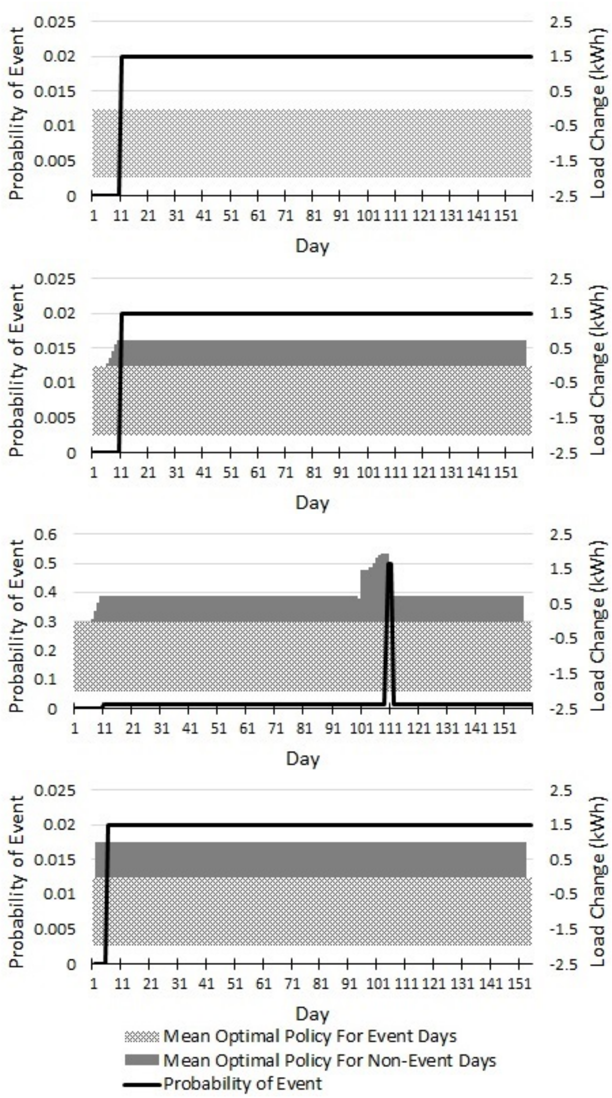

Fig. 1. Optimal policy and probability of event by day for scenario 1 (top), 2 (second), 3 (third), and 4 (bottom).

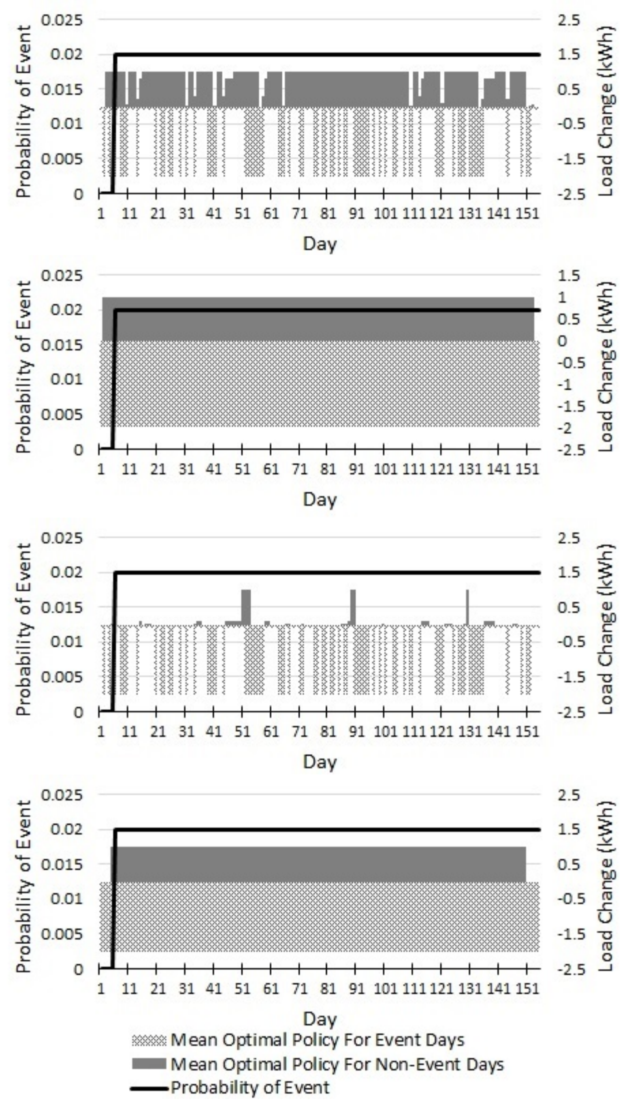

Fig. 2. Optimal policy and probability of event by day for scenario 5 (top), 6 (second), 7 (third), and 8 (bottom).

but a larger apparent DR quantity due to the inflated baseline, thus resulting in 63% greater customer net benefits and a 50% higher utility cost per unit true DR.

Scenario 2 also shows that even with poor information about when events will occur, customers can have incentive to take inexpensive gaming actions (such as load shifting with a battery) to inflate baselines. With better information about event probabilities, as in scenario 3, customers can have incentive to take more expensive gaming actions (such as turning on more appliances) during the days preceding high probability of event days. This results in even greater apparent DR, higher customer profits, and 70% higher utility cost per unit true DR compared to scenario 1 with no gaming.

Scenario 4, with a "5 in 5" baseline, has nearly identical values of true DR, apparent DR, and utility cost, to scenario 2, with a "5 in 10" baseline. Despite these similarities, utilities may still prefer the "5 in 10" baseline, because it can yield less demand inflation on non-event days.

Scenario 5 shows that varying default load can reduce incentives for both event response and gaming action. This is because when payments are capped above zero and default load is much higher than the baseline, customer actions may not be able to increase the payment above zero. Alternatively, variable load does not have a major impact on true or apparent DR quantities when negative payments are allowed (not the similar results for scenarios 4 and 6 in Table I).

Scenarios 7 and 8 show a similar effect as scenarios 5 and 6: not allowing negative payments reduces incentives for both event response and gaming. However, we observe a change in cost of customer actions changes the relative impacts to event response versus gaming. With lower load shifting costs (scenarios 5 and 6) capping payments at zero reduces event response more than gaming, resulting in 18% higher utility costs per unit true DR. Alternatively, scenarios 7 and 8 show an 32% lower utility cost per unit true DR because gaming is reduced much more than event response.

V. CONCLUSION

Our baseline-based DR model can identify optimal customer behaviors under a variety of baseline-based DR program parameters and a wide variety of customer parameters, thus revealing customer incentives to artificially inflate baselines. Analytical results provide some fundamental insights into the drivers of optimal customer decisions. Numerical results show that incentives for gaming may exist in real-world DR programs, and that the level and impact of gaming can depend on a number of factors, in ways that would be difficult to predict without our model. Since the impacts of gaming baselines can have significant and non-obvious impacts on the cost and effectiveness of baseline-based DR programs, gaming incentives should be considered during DR program design and in utility planning and operations.

REFERENCES

[1] J. S. Vardakas, N. Zorba, and C. V. Verikoukis, "A Survey on Demand Response Programs in Smart Grids: Pricing Methods and Optimization Algorithms," *IEEE Communications Surveys and Tutorials*, vol. 17, no. 1, pp. 152–178, 2015.

[2] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE transactions on Smart Grid*, vol. 1, no. 3, pp. 320–331, 2010.

[3] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *2011 IEEE power and energy society general meeting*. IEEE, 2011, pp. 1–8.

[4] Y. Xu, N. Li, and S. H. Low, "Demand response with capacity constrained supply function bidding," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1377–1394, 2016.

[5] P. Samadi, A. H. M. Rad, R. Schober, and V. W. Wong, "Advanced demand side management for the future smart grid using mechanism design." *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1170–1180, 2012.

[6] D. Muthirayan, D. Kalathil, K. Poolla, and P. Varaiya, "Mechanism design for demand response programs," *arXiv preprint arXiv:1712.07742*, 2017.

[7] B. Chew, B. Feldman, D. Ghosh, and M. Surampudy, "2018 Utility Demand Response Market Snapshot," Smart Electric Power Alliance, Tech. Rep., 2018. [Online]. Available: https://sepapower.org/resource/2018-demand-response-market-snapshot/

[8] P. Antunes, P. Faria, and Z. Vale, "Consumers performance evaluation of the participation in demand response programs using baseline methods," in *2013 IEEE Grenoble Conference PowerTech, POWERTECH 2013*.

[9] J. Vuelvas, F. Ruiz, and G. Gruosso, "Limiting gaming opportunities on incentive-based demand response programs," *Applied Energy*, vol. 225, pp. 668–681, sep 2018.

[10] J. Vuelvas and F. Ruiz, "Rational consumer decisions in a peak time rebate program," feb 2018. [Online]. Available: http://arxiv.org/abs/1802.08112

[11] D. G. Dobakhshari and V. Gupta, "A Contract Design Approach for Phantom Demand Response," *IEEE Transactions on Automatic Control*, p. 1, 2018.

[12] D. Muthirayan, E. Baeyens, P. Chakraborty, K. Poolla, and P. P. Khargonekar, "A Minimal Incentive-based Demand Response Program With Self Reported Baseline Mechanism," 2019. [Online]. Available: https://arxiv.org/pdf/1901.02923.pdf

[13] K. Khezeli and E. Bitar, "Risk-sensitive learning and pricing for demand response," *IEEE Transactions on Smart Grid*, vol. 9, no. 6, pp. 6000–6007, 2018.

[14] F. E. R. Commission, "Order No. 745, Demand Response Compensation in Organized Wholesale Energy Markets," 2011.

[15] F. Wolak, "Residential Customer Response to Real-Time Pricing: The Anaheim Critical-Peak Pricing Experiment," 2006. [Online]. Available: https://escholarship.org/uc/item/3td3n1x1

[16] J. Pierobon, "Two FERC settlements illustrate attempts to 'game' demand response programs," 2013. [Online]. Available: https://www.theenergyfix.com/2013/07/25/two-ferc-settlements-illustrate-attempts-to-game-demand-response-programs/

[17] S. Borenstein, "Money for Nothing?" 2014. [Online]. Available: https://energyathaas.wordpress.com/2014/05/12/money-for-nothing/

[18] Con Edison, "Commercial Demand Response Program Guidelines," 2019. [Online]. Available: https://www.coned.com/-/media/files/coned/documents/save-energy-money/rebates-incentives-tax-credits/smart-usage-rewards/smart-usage-program-guidelines.pdf?la=en

[19] D. Ellman and Y. Xiao, "Online Appendix for CDC 2019 Submission," Tech. Rep., 2019. [Online]. Available: http://yuanzhangxiao.com/papers/CDC2019Appendix.pdf

[20] Consolidated Edison Company of New York Inc., "Schedule For Electricity Service," 2019. [Online]. Available: https://www.coned.com/_external/cerates/documents/elecPSC10/electric-tariff.pdf

[21] Con Edison, "Demand Response Forum Presentation," 2018. [Online]. Available: https://www.coned.com/-/media/files/coned/documents/save-energy-money/rebates-incentives-tax-credits/smart-usage-rewards/demand-response-forum.pdf

[22] Consolidated Edison Company of New York Inc., "Energy Efficiency and Demand Management Procedure - General: Calculating Customer Baseline Load," 2018. [Online]. Available: https://www.coned.com/-/media/files/coned/documents/save-energy-money/rebates-incentives-tax-credits/smart-usage-rewards/customer-baseline-load-procedure.pdf?la=en