Personalized Feedback-based Customer Incentives in Automated Demand Response

Thanasis G. Papaioannou, George D. Stamoulis and Marilena Minou
Department of Informatics
Athens University of Economics & Business (AUEB), Greece
Email: {pathan, gstamoul, minou}@aueb.gr

***Abstract*—Automated Demand Response (ADR) can facilitate residential customers to effectively reduce their energy demand and make savings in a simple way, provided that appropriate incentives are offered to them. Most often, incentives involved in ADR contracts are statically defined and assume full customer rationality, thus hindering sustained customer enrollment to them of customers with other characteristics (e.g. altruism). In this paper, we derive appropriate (and personalized) incentives for ADR contracts, so that non-fully rational customers are compensated even when information for consumer utilities is not available. In case such information is hidden, we assume that customers provide feedback on their satisfaction from direct endowments, albeit sustaining energy-consumption reduction. Moreover, we consider the case where customers may strategically lie on their satisfaction from ADR incentives, so as to self-optimize. We mathematically model the customer and the utility company's problems and solve them algebraically or in a distributed manner. Furthermore, based on customer feedback on appropriate endowments for different energy-consumption reductions, we propose an algorithm that can find the optimal set of satisfied targeted customers, which achieve the total desired energy-consumption reduction at the minimum endowment cost. Based on numerical evaluation and simulation experiments, we showcase the validity of our analytical framework in realistic scenarios and that, for the case of hidden information, customer feedback is adequate for calculating incentives that can lead to successful DR campaigns.**

I. Introduction

Demand Response (DR) programs for curtailing energy consumption in critical times for the operation of the grid are becoming common. Automated DR (ADR) automates the response process of the customer to the DR signals by means of electric controls installed at the customer premises. An ADR contract usually predefines a financial reward for the customer as a compensation for the user utility losses due to her curtailed energy consumption. ADR rebates are defined mostly statically and based either on the costs of ADR equipment [1] or the cost per unit of energy at peak times.

However, this contractual form of financial endowment should not be based solely on the market value of the conserved energy for two reasons. First, the utility loss for the customer in the time periods that energy consumption is curtailed may not be linked to the market value of that energy, but may include other aspects, such as actual needs, sensitivity to personal-comfort loss, etc. In such a case, the endowment may fall short as means for customer engagement into the ADR program. Second, ADR programs currently do not take into account the customer satisfaction from the provided endowment for an energy-consumption reduction. As a result, an unsatisfied customer may not renew her ADR contract after it expires. In [2], we first considered the problem of calculating uniform satisfactory ADR incentives for uniform load curtailment based on explicit, yet anonymous, customer feedback, in the cases of full and hidden information on user utility functions.

In this paper[1], we focus on hidden information on user utility functions and investigate personalized ADR incentives that can ensure wide customer acceptance, for uniform and non-uniform energy-consumption reductions by customers. We build upon the model of [2] and consider that customers are not solely driven by financial motives, but also by a number of behavioral factors, such as *altruism*. Altruistic values in energy can complement or even dominate the narrow self-interest presumed by a standard rational choice theory of decision making [3], yet they can be measured in monetary terms [4]. We identify two optimization problems for the DR designer, in order to achieve a desired load curtailment: (i) find minimum ADR incentives that satisfy a lower bound on the percentage of customers, and (ii) maximize customer satisfaction ratio within an upper bound on the budget for endowments. We analytically derive that hidden information on partially-known user utilities can be approximated based on customer feedback on the acceptance of ADR incentives. We also introduce two variants of a feedback-based distributed iterative algorithm for solving the aforementioned optimization problems of the DR designer. We separately address the cases of observable (i.e., signed) and unobservable (i.e., anonymous) individual feedback. We also consider strategically-lying customers on their satisfaction feedback and their impact to the convergence of the distributed algorithm when a proposed *mitigation policy* against lying is employed by the utility company. Furthermore, based on customer feedback, we deduce satisfactory endowments for different energy-consumption reductions sustained by different customers and propose an algorithm for optimal customer selection, i.e., targeting. This algorithm can find the optimal set of customers and their respective energy-consumption reductions that achieve a total energy-consumption reduction at the minimum total endowment cost, while keeping targeted users satisfied. Based on numerical evaluation and simulation experiments, we show that hidden information on user utility, although costly, can be overcome by means of customer feedback at a low endowment cost. Finally, we demonstrate that optimal customer targeting based on individual feedback on customer satisfaction can further reduce the total endowment

[1]This research has been partly co-funded by EU's Horizon 2020 Programmes under grant agreements no 649796 OPTi and no 696170 ChArGED.

cost for keeping customers satisfied, even in the presence of strategic liars. To the best of our knowledge, our work is the first one that aims to tackle uncertainty in acceptance of ADR contracts based on customer feedback.

The remainder of this paper is organized as follows: In Section II, we overview our context and system model. In Section III, we overview the two aforementioned DR designer's optimization problems in the cases of full and hidden information on user utility functions. In Section IV, we discuss the problem of strategic lying and propose a mitigation policy for the DR designer. In Section V, we present our algorithm for finding the optimal set of targeted customers based on individual feedback. In Section VI, we perform numerical evaluation and simulation experiments with our proposed solutions. In Section VII, we review the related work, and finally, in Section VIII, we conclude our paper and indicate some future work.

II. System Model

We consider a district of residential buildings served by a utility company. (Alternatively, we could consider additional contexts, such as the residents of a social-housing establishment.) The utility company offers ADR contracts to the residents of the district (or of the social housing establishment). Denote $\mathcal{N}$ the set of residential houses that enroll into the ADR programs. According to the ADR contract, the utility company curtails the total energy consumption of the house of a customer in peak periods by a desired level. A customer i enjoys net benefit U_i (i.e., user satisfaction minus energy cost) from consuming baseline energy q_i^0 and an energy-consumption reduction ΔQ_i in specific time periods according to an ADR contract results to a net-benefit loss $\Delta U_i = -\eta_i U_i$. In return, the customer i receives an endowment b_i by the utility company. Throughout the paper, whenever the endowment is the same for all customers, it is denoted as b, otherwise as b_i.

Moreover, we consider that the user utility function of the consumer does not solely depend on her total energy consumption, but additionally on other socio-demographic or psychological factors [5], such as social norms, altruism, needs, habits, etc. In general, the utility function of customer i depends on her net benefit U_i from energy consumption q_i^0 and the net benefit of of all N customers except for customer i from their own energy consumption U_{-i}, i.e.,

$$u_i = g(U_i, U_{-i}) \,, \tag{1}$$

where $g(\cdot)$ is an arbitrary function.

Overall, in general, the user utility difference for customer i due to the ADR contract is given by:

$$\Delta u_i = \hat{g}(\Delta U_i, \Delta U_{-i}) + b_i \,, \tag{2}$$

where $\hat{g}(\cdot)$ is a function of net-benefit loss that depends on $g(\cdot)$.

For instance, we consider the utility model of [2] that incorporates altruism. According to [2], an individual's overall utility function is given by a convex combination of her own monetary payoff and of the sum of the payoffs of others. In the same spirit, our user utility model that incorporates altruism is given by:

$$u_i = (1 - \gamma_i) U_i + \gamma_i \overline{U}_{-i} \,, \tag{3}$$

where $\gamma_i \in [0, 1]$ is the degree of altruism of customer i and $\overline{U}_{-i}$ is the mean net benefit from consuming energy for all N customers except for i. Note that a customer i with $\gamma_i = 0$ is considered fully selfish, while with $\gamma_i = 1$ is considered completely "disinterested" in her own net benefit from consuming energy in the sense that she cares only for the net benefits of others from energy consumption. Henceforth, for easiness in the calculations and without loss of generality, we assume the net benefit U_i of customer i from her baseline energy consumption to be *normalized* by the maximum net benefit $\hat{U}$ from consuming energy of all customers, i.e., $U_i \in [0, 1]$. Then, the endowment b_i for customer i is also normalized by the maximum utility of all customers, e.g., $b_i = 0.3$ means that the endowment for customer i equals the 30% of the maximum utility value.

Then, the user utility difference for customer i due to the ADR contract is given by:

$$\Delta u_i = -(1 - \gamma_i)\eta_i U_i + \gamma_i \overline{\Delta U}_{-i} + b_i \,, \tag{4}$$

We define that when the endowment covers the loss of a customer, i.e., $\Delta u_i \geq 0$, then the customer is considered to be *satisfied* by the ADR contract; otherwise, i.e., $\Delta u_i < 0$, the customer is *unsatisfied*.

III. The Problem

The DR designer needs to construct ADR contracts appropriately, so that customers to (a) enroll in them in the first place, (b) extend/renew their ADR contracts. The former can be achieved if the reduction in the customer benefit due to the lower energy consumption, as specified in the ADR contract, is *expected* by the customer to be compensated for by the associated endowment. The latter necessitates that the *materialization* of the ADR contract is indeed satisfactory for the customer. Obviously, an indefinitely high endowment would achieve both aforementioned goals, but that would be prohibitively costly. There is a *trade-off* among the value loss for the customer due to the energy-consumption reduction as defined in her contract, the associated endowment to the customer and the customer satisfaction by the ADR contract. Within this trade-off, the objective of the DR designer can be one of the following:

(I) Maximize customer satisfaction α for a specific net-benefit reduction η_i for each customer i due to lower energy consumption within a budget limit B for endowments.

(II) Minimize total endowment cost for a lower-bound $\underline{\eta}$ in the net-benefit loss due to energy-consumption reduction of each customer and a lower-bound $\underline{\alpha}$ in customer satisfaction.

In case of full-information on user utilities, since the customer satisfaction ratio α is monotonic in the uniform endowment b, problem (I) can be solved by sorting all consumers with respect to Δu_i of each customer i after using maximum endowment $b = B/N$ and count how many of them are positive. Also, problem (II) can be solved by sorting all consumers with respect to Δu in a descending order for b=0 and then solve the equation $\Delta u_k = 0$ to find b, with $k = \lceil \alpha \cdot N \rceil$. This is the value of b that makes α% customers having $\Delta u \geq 0$.

If personalized incentive b_i per customer i is employed, then problem (I) is again solved as described above, while problem (II) is solved as follows: For each customer i, calculate the personalized incentive that renders $\Delta u_i = 0$. Sort the list of customers based on their personalized incentive in ascending order. The minimum total incentive required for satisfying $\alpha \cdot 100\%$ customers is given by summing the top-(αN) personalized incentives.

A. Anonymous Feedback

Assume now that user utility functions are private, thus $\Delta U_i, \Delta u_i$ are private and known only to customer i. Each customer i *honestly* provides *binary feedback* on her respective satisfaction Δu_i. Customer feedback may be provided on an anonymized manner (e.g., through a ballot) or it may be attributed to specific individuals. In this subsection, we consider anonymized customer feedback, while observable individual feedback is considered in the next subsection.

Consider a distributed algorithm where, at each round t, the DR designer sets a b_t and each customer i responds to it with feedback $v_{i,t+1}$, which collectively result to a mean satisfaction level α_{t+1} for the received incentives at the next round. The feedback $v_{i,t+1}$ of customer i at round $t+1$ is determined by the sign of:

$$\Delta u_{i,t+1} = \hat{g}(\Delta U_i, \widetilde{\Delta U}_{-i,t}) + b_t \, , \tag{5}$$

where $\widetilde{\Delta U}_{-i,t}$ is an estimate by customer i on the net benefit losses of other customers at round t. We show how such an estimate can be calculated later in this subsection.

For solving problem (I), the DR designer needs to update b_t at each round t, so as to maximize customer satisfaction within in the feasible set of direct incentives, i.e. $b_t \leq B, \ \forall t$. Employing gradient ascent, the DR designer selects b_{t+1} for the round $t+1$ as follows:

$$b_{t+1} = \begin{cases} \max\{b_t + \Delta\alpha \cdot \kappa, 0\} \, , & \text{when } \Delta\alpha \neq 0 \\ b_t + \kappa, & \text{when } \Delta\alpha = 0 \end{cases} \tag{6}$$

where $\Delta\alpha = \alpha_{t+1} - \alpha_t$ and $0 < \kappa << 1$ the step size of the gradient ascent algorithm.

If $\Delta\alpha/\alpha_t < \Delta b/b_t$ or $b_t \geq B$, then stop iterations. In plain words, if the increase rate of b is greater than the increase rate of α or if we exceeded the available endowment budget B, then we stop iterations.

For solving problem (II), the DR designer again employs the aforementioned gradient ascent algorithm, but the stopping criterion is $\alpha_t \geq \underline{\alpha}$.

We now show how $\widetilde{\Delta U}_{-i,t}$ at round t can be estimated by customer i for a specific instance of user utility functions. Assume the user utility model in (3) and Δu_i defined by (4). Also, assume that U_i and γ_i are private information known only to customer i; hence, Δu_i is private. Then, the feedback $v_{i,t+1}$ of customer i at round $t+1$ is determined by the sign of the following:

$$\Delta u_{i,t+1} = -(1-\gamma_i)\eta_i U_i + \gamma_i \overline{\Delta U}_{-i,t} + b_t \tag{7}$$

We need to calculate $\overline{\Delta U}_{-i,t}$ at round t. Note that there holds

$$-1 \leq -U_i \leq 0 \, . \tag{8}$$

Since $\eta_i \in [0,1]$, it follows that

$$-\eta_i \leq \Delta U_i \leq 0 \, , \tag{9}$$

$$-\eta_{max} \leq \overline{\Delta U} \leq 0 \, , \tag{10}$$

where $\eta_{max} \in [0,1]$ is the maximum fraction of resulting net benefit loss for customers from energy-consumption reduction. Therefore and since $\gamma_i \in [0,1]$, it is true that

$$b_t - \eta_{max} \leq \Delta u_{i,t+1} \leq b_t \, , \forall i \in \mathcal{N} \, . \tag{11}$$

Recall that

$$\overline{\Delta U}_{-i} = -\frac{\sum_{\substack{j=1 \\ j \neq i}}^{N} \eta_j U_j}{N} \, . \tag{12}$$

Adding down (7) for all customers and dividing by N and taking that $\Delta U_{-i} \approx \Delta U$ for large N, we have that

$$\overline{\Delta u}_{t+1} = -\frac{\sum_{i=1}^{N} (1-\gamma_i)\eta_i U_i}{N} + \bar{\gamma}\overline{\Delta U}_t + b_t \, . \tag{13}$$

We can assume the following approximation

$$\frac{\sum_{i=1}^{N} (1-\gamma_i)\eta_i U_i}{N} \approx (1-\bar{\gamma})\overline{\Delta U} \, , \tag{14}$$

which, when γ_i values for each customer i are close to each other, is a rather accurate. Based on this approximation, equation (13) becomes

$$\overline{\Delta u}_{t+1} = \overline{\Delta U}_t + b_t \, . \tag{15}$$

Also, according to the user feedback at each round t, the mean satisfaction a_{t+1} of the customers can be calculated. Recall that according to (11), $\Delta u_{i,t+1}, \overline{\Delta u}_{t+1} \in [b_t - \eta_{max}, b_t]$. Since by definition $a_{t+1} \in [0,1]$, it follows that $\eta_{max}(a_{t+1}-1) + b_t \in [b_t - \eta_{max}, b_t]$ as well, i.e., $\overline{\Delta u}_{t+1}$ and $\eta_{max}(a_{t+1}-1)+b_t$ have the same codomains (i.e., sets of outputs). Moreover, observe that

$$\eta_{max}(a_{t+1}-1)+b_t = \eta_{max}\left(\frac{\sum_{i\in\mathcal{N}} \mathbb{1}(\Delta u_{i,t+1} \geq 0)}{N} - 1\right)+b_t,$$

where $\mathbb{1}(\cdot)$ is the indicator function that equals 1 when its argument is true and 0 otherwise. Therefore, $\eta_{max}(a_{t+1}-1)+b_t$ and $\overline{\Delta u}_{t+1}$ increase or decrease in the same way according to the values $\Delta u_{i,t+1}$ of each customer i. To this end, it can be taken as approximation that:

$$\overline{\Delta u}_{t+1} = \eta_{max}(a_{t+1}-1) + b_t \tag{16}$$

Employing (16) in (15), we derive that

$$\overline{\Delta U}_t \approx \eta_{max}(a_{t+1}-1) \, . \tag{17}$$

Then, employing (17) in (7), we have that for each customer i at round $t+1$ the following is true:

$$\Delta u_{i,t+1} = -(1-\gamma_i)\eta_i U_i + \gamma_i \eta_{max}(a_t - 1) + b_t \tag{18}$$

Customer i provides her satisfaction feedback at round $t+1$ according to the sign of (18). Note that η_{max} can be guessed by the DR designer based on the desired energy-consumption reduction and the consumption baselines of customers, through which shiftable and non-shiftable loads can be estimated. Even if η_{max} cannot be estimated, it can be set equal to 1 for a more loose, yet adequate, approximation of $\overline{\Delta U}_t$, as shown in [2].

B. Named Feedback

If the origin of individual votes can be observed (i.e., if feedback is signed), then b can be decided on a personalized basis. This situation can arise in various realistic cases of feedback provision, e.g., through a mobile app, through a web interface with authorized access, etc. In the algorithm of Subsection III-A, the personalized direct incentive b_i could be selected by the DR designer at round t, so as $\Delta u_{i,t} \approx 0$ by observing when the individual feedback of customer i changes. Then, b_i is kept constant in the subsequent rounds until target satisfaction is achieved.

IV. Strategic Feedback

So far, we have assumed that the customers report honestly their satisfaction from the direct incentive that is provided to them. However, in the aforementioned setting, customers have indeed incentive to lie; untruthfully reporting to be unsatisfied, results in higher direct incentive for them. To this end, we consider the following *mitigation policy* against strategic liars in the ADR contract: the DR designer is willing to provide a direct incentive b to the customers as an endowment for their net benefit losses due to the lower energy consumption as long as b does not exceed a budget upper bound B that is unknown to the customers. If b_t becomes infeasible, then the DR designer provides no direct incentives to customers, while they still sustain the desired energy-consumption reduction as a *penalty*. We assume that the DR provider selects the upper bound on endowment budget B based on pilot experiments with a trusted subset of customers or by having an estimation of the maximum net benefit $\hat{U}$ and knowing the forcible energy consumption decrease. Then, considering again the iterative approach for finding b, the expected utility difference for customer i can be calculated by:

$$\Delta u_{i,t+1} = -(1-\gamma_i)\eta_i U_i + \gamma_i \overline{\Delta U}_{-i,t} + Pr[b_t < B | z_i \textit{ lies}] b_t \,, \quad (19)$$

where $Pr[b_t < B | z_i \text{ lies}]$ is the probability that b_t does not exceed the budget upper bound B given that the customer i has already lied z_i times. It can be assumed that the probability of overshooting the upper endowment bound B increases exponentially with the number of iterations that the user sends untruthful feedback. Equation (19) is employed by strategically lying customers for determining their satisfaction feedback (and whether to lie about it or not) in the various algorithms run by the DR designer to find appropriate endowment values.

V. Feedback-based Customer Targeting

So far, we have assumed that the desired energy-consumption reduction ΔQ is equally shared by all customers. We now consider customer targeting policies for achieving ΔQ, in which, different customers undertake different burden for achieving the energy reduction objective. Recall that the same energy-consumption reduction results to a different net benefit loss fraction η_i for each customer i. In general, there is an increasing function $h_i(\Delta Q_i)$ for each customer i that determines the percentage of net benefit loss for different energy-consumption reduction values ΔQ_i incurred by i. In this case, equation (4) is rewritten as follows:

$$\Delta u_i = -(1-\gamma_i) h_i(\Delta Q_i) U_i + \gamma_i \overline{\Delta U}_{-1} + b_i(\Delta Q_i) \,, \quad (20)$$

where b_i is the endowment that needs to be given to customer i for being indifferent for the energy-consumption reduction ΔQ_i.

We now address the problem of finding the specific energy-consumption reduction that needs to be incurred by each individual customer, so that the total desired energy-consumption reduction is achieved at the minimum endowment cost. Assume a discrete set of energy-consumption reductions that may be enforced to different customers $\mathcal{H} = \{\frac{1}{N}\Delta Q, \frac{2}{N}\Delta Q, \ldots, \theta_{max}\Delta Q\}$, where θ_{max} is the maximum fraction of energy-consumption reduction to be incurred individually by a customer. In general, θ_{max} can be personalized per customer, but we assume it here to be the same for all customers for the sake of simplicity. Then, in case of full information on user utility functions, the DR designer can calculate for each customer i and each different energy-consumption reduction $\Delta Q_j \in \mathcal{H}$ the appropriate endowment $b_{i,j}$, so that $\Delta u_i = 0$. In case of hidden information on user utilities, yet with observable individual feedback on user satisfaction by the provided endowment, the DR designer can consecutively enforce different $\Delta Q_j \in \mathcal{H}$ to all customers and determine respective $b_{i,j}$ for each customer i by employing the approach of Subsection III-B. Strategic feedback, if present, is deterred according to the approach of Section IV. However, note that, both in the case of full information and in that of hidden information on user utilities, this approach entails an approximation on the value of $\overline{\Delta U}_{-i}$. Specifically, while each $(\Delta Q_j, b_{i,j})$ pair for $\Delta Q_j \in \mathcal{H}$ and customer i is calculated based on $\overline{\Delta U}_{-i}$ resulting from applying the consumption reduction ΔQ_j to all customers, in fact, $\overline{\Delta U_{-i}}$ should have been the one resulting by targeting only a specific subset of users, which is unknown. We experimentally assess the accuracy of this approximation in Subsection VI-B.

Having determined pairs of $(\Delta Q_j, b_{i,j})$ for each customer i and different energy reductions j with $\Delta Q_j \in \mathcal{H}$, the DR designer needs to select the targeted users and their respective energy-consumption reductions, so as to cover energy ΔQ at the minimum total endowment cost. This problem resembles bounded Knapsack with item weight ΔQ_j and item value $b_{i,j}$, and it can be solved as follows (see Algorithm 1): Sort pairs $(\Delta Q_j, b_{i,j})$, $\forall i, j$ according to their $\Delta Q_j / b_{i,j}$ value in decreasing order and add them to list $\mathcal{L}$. Start from the top of the sorted list $\mathcal{L}$ and add top item to the bag. When a pair belonging to customer i is added to the bag, remove from the sorted list $\mathcal{L}$ all other pairs of customer i. Assume current energy-consumption reduction of the bag is W. Continue adding pairs to the bag until $W + \Delta Q_j \geq \Delta Q$ for next pair $(\Delta Q_j, b_{k,j})$ of customer k. Then, from the remaining pairs in the list $\mathcal{L}$ select the first one (w, v) that gives $W + w = \Delta Q$ if exists. Otherwise, among remaining pairs in $\mathcal{L}$, select the one (w, v) with the lowest v, while $W + w > \Delta Q$. For J different energy-consumption reduction steps for each customer and N customers the complexity of this algorithm is $O(NJ \log(NJ))$ for sorting the list and $O(NJ)$ for filling the bag, i.e., $O(NJ \log(NJ))$ overall.

Algorithm 1 Feedback-based Customer Targeting

Input: $\mathcal{P} \leftarrow \{(\Delta Q_j, b_{i,j})\}\ \forall i \in \mathcal{N}, \Delta Q_j \in \mathcal{H}$
Output: Set $\mathcal{S}$ of targeted customers

$\mathcal{L} \leftarrow$ sorted $\mathcal{P}$ in descending order based on $\Delta Q_j / b_{i,j}\ \forall i \in \mathcal{N}, \Delta Q_j \in \mathcal{H}$
$W \leftarrow 0$, $v_{extra} \leftarrow +\infty$
while $\mathcal{L} \neq \emptyset$ **do**
  $(w_i, v_i) \leftarrow \mathcal{L}.pop()$ ▷ Get first item from list $\mathcal{L}$. It belongs to customer i.
  if $W + w_i > \Delta Q$ **then**
    break ▷ Exit the while loop.
  end if
  $\mathcal{S}$.add((w_i, v_i))
  $\mathcal{L}$.remove(other pairs of customer i)
  $W \leftarrow W + w_i$
end while
$v_{extra} \leftarrow +\infty$
$p \leftarrow$ **None**
for $(w, v) \in \mathcal{L}$ **do**
  if $W + w = \Delta Q$ **then**
    $\mathcal{S}$.add((w, v))
    $W \leftarrow W + w$
    break ▷ Exit the for loop, optimal solution found!
  else ▷ Find item (w, v) with smallest v in the remaining items, so that $W + w > \Delta Q$.
    if $W + w > \Delta Q \wedge v_{extra} > v$ **then**
      $v_{extra} \leftarrow v$
      $p \leftarrow (w, v)$
    end if
  end if
end for
if $W < \Delta Q$ **then** ▷ If optimal solution not found.
  $\mathcal{S}.add(p)$ ▷ Add last item to $\mathcal{S}$.
end if

Theorem 1. *When it is possible to fill the bag with exactly ΔQ, then Algorithm 1 finds an* optimal *solution.*

Proof. Assume that $(w_1, v_1), (w_2, v_2), \ldots, (w_K, v_K)$ are the items selected by Algorithm 1, so that $\sum_{k=1}^{K} w_k = \Delta Q$ with total endowment cost $V = \sum_{k=1}^{K} v_k$. If it was not optimal, then there would exist a pair (w_e, v_e) that replacing an item m from $1, \ldots, k$ would give $V' < V$. Since after the item replacement the total energy-consumption reduction should still be ΔQ, it is derived that $w_m = w_e$. Then, for $V' < V$, it should be that $v_e < v_m$. However, if this were true, then pair (w_e, v_e) should have been before (w_m, v_m) in the sorted list $\mathcal{L}$, which is not true. Thus, when $W = \Delta Q$, the solution found by the Algorithm 1 is optimal. □

Note that Algorithm 1 is more likely to find the optimal solution, the larger the J of discrete steps of energy-consumption reduction fractions for which customer feedback is collected, the smaller the ΔQ and/or the larger the number N of customers.

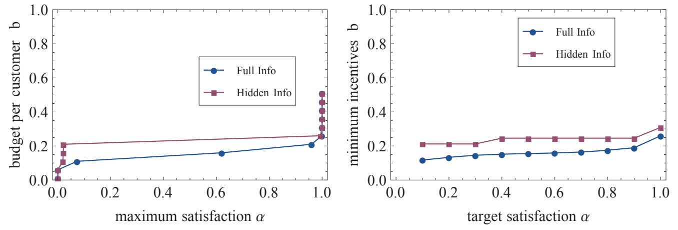

(a) (b)

Fig. 1: (a) The maximum customer satisfaction per customer budget B/N and (b) the minimum individual endowment that achieves a desired customer satisfaction.

VI. Evaluation

We consider a system of $N = 1000$ customers. The degree of altruism γ_i for customer i follows Uniform distribution in $[0, 1]$, while the fraction η_i of her net benefit loss due to energy-consumption reduction is assumed to be uniformly distributed in $[0.1, 0.3]$, unless otherwise specified. The DR designer is assumed to have guessed semi-correctly $\eta_{max} = 0.5$. The satisfaction of customer i for her nominal energy consumption is assumed to follow Normal distribution with mean 0.8 and standard deviation 0.1, i.e., $U_i \sim N(0.8, 0.1)$.

A. Uniform Load Reduction

We assume that the DR designer does not know the user utility functions of customers and that customers provide feedback on their satisfaction for the endowment provided to each of them in an iterative manner according to (6) with κ=0.1. Solving problem (I) for the cases of full and hidden information, the maximum achievable customer satisfaction within a certain budget B is depicted in Fig. 1a. Also, solving problem (II) for the cases of full and hidden information, the minimum customer endowment that achieves a specific minimum customer satisfaction is depicted in Fig. 1b. As it can be seen from Fig. 1a and Fig. 1b, hidden information creates an extra incentive cost, yet, very low, thanks to our feedback-based approximation and distributed algorithm. Our distributed algorithm converges very fast, requiring only 5 and 6 iterations for these instances of problems (I) and (II) respectively.

Next, we consider observable individual (i.e., named) feedback and solve problems (I) and (II), as described in Section III for the full and hidden info cases. As depicted in Fig. 2, tailored individual endowment b_i for keeping each customer i that has sustained an energy-consumption reduction satisfied are more economic for the utility company, as compared to a uniform endowment b for all customers, for both optimization problems.

We now assume that a fraction of the customers strategically lie on their reported feedback. We assume that $Pr[b_t < \hat{b}|z_i \text{ lies}]$ is modeled by an exponential distribution with $\lambda = 1$, the random variable of which is the number of times that the customer is lying. A strategic liar does so opportunistically: if lying at an iteration reduces her user utility difference at the next iteration, she stops lying from then on. We present results from the solution to problem (II) for different fractions of liars $\{20\%,\ 50\%\}$ in Fig. 3. Observe therein that, the presence of liars increases the amount of incentives to be paid for the same

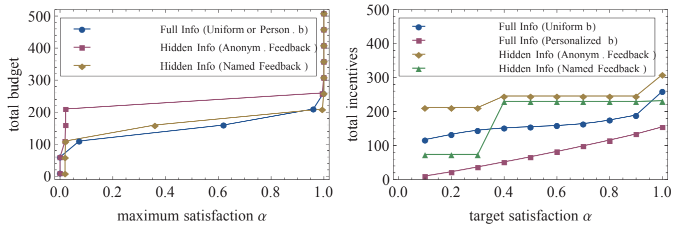

(a) (b)

Fig. 2: Observable individual feedback: (a) The maximum customer satisfaction with respect to the total incentive budget and (b) the minimum total endowments that achieve a desired customer satisfaction.

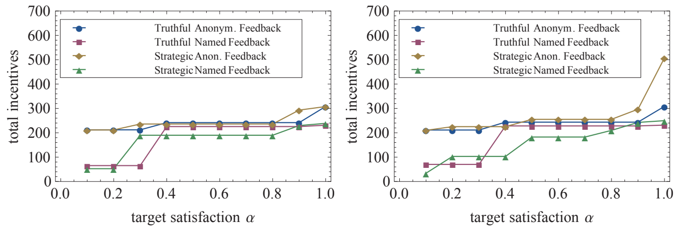

(a) (b)

Fig. 3: Strategic feedback: minimum endowment per targeted customer satisfaction in presence of (a) 20% and (b) 50% liars.

targeted customer satisfaction when the feedback is anonymous. However, for named feedback, strategically lying does not pay off for liars: In fact, since targeted customer satisfaction can be achieved mostly by satisfied honest customers (especially true when targeted satisfaction ratio is smaller than the fraction of liars) and since incentives are given to satisfied customers only, the total minimum amount of endowments is smaller.

B. Customer Targeting

In this subsection, we deal with customer targeting for achieving a desired total energy-consumption reduction. We assume that the net benefit loss of customer i for a specific individual energy-consumption reduction ΔQ_i is given by $h_i(\Delta Q_i) = \Delta Q_i^{1+n_i}$, with $n_i \sim U(0.1, 0.3)$ expressing the sensitivity of customer i to energy-consumption reduction. We set $\mathcal{H} = \{0.05\Delta Q, 0.1\Delta Q, 0.15\Delta Q, 0.2\Delta Q, 0.25\Delta Q, 0.3\Delta Q\}$ to be the set of different individual energy-consumption reductions considered. For simplicity, we set the nominal energy-consumption reduction of each customer equal to 1. We consider both full and hidden information on user utility functions with anonymous or named individual feedback and find $(\Delta Q_j, b_{i,j})$ pairs for each customer i and $\Delta Q_j \in \mathcal{H}$ according to the approach of Section V. We also consider the presence of 30% strategic liars for either anonymous or named user feedback. Employing Algorithm 1, we find the set of targeted users and the energy-consumption reduction that should be applied to each user at this set, in order to achieve different total energy-consumption reductions at the minimum total endowment cost for all the cases considered. As shown in Fig. 4a, targeting results in lower total endowment cost for achieving a specific total energy-consumption reduction, as compared to those

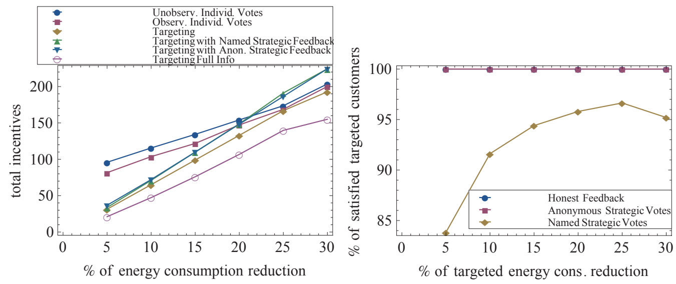

(a) (b)

Fig. 4: (a) Total incentives for keeping customers satisfied for different energy-consumption reductions, either sharing the reduction among all customers for anonymous feedback and named feedback, or allotting the reduction specific customers by means of targeting. (b) The number of satisfied targeted customers for different total energy-consumption reductions.

found by applying the same individual energy-consumption reduction to all users, even when full information on user utility functions is not available and individual feedback is observable or not. However, as expected, the benefit obtained by targeting decreases as the aimed total energy-consumption reduction increases (and so do the set of targeted users and their respective individual consumption reductions). Also, observe in Fig. 4a that, hidden information on user utility functions indeed increases total endowment cost, while strategic feedback (either anonymous or named) creates an additional cost burden for the DR designer, yet a very limited one for low desired energy reductions.

Since the calculation of the respective necessary individual endowments for keeping targeted customer i satisfied for different individual consumption reductions entails an approximation on the value of $\overline{\Delta U}_{-i}$, as explained in Section V, one may think that targeted users may not be satisfied with the provided endowments after all. We experimentally study this case as follows: We employ equation (18) without t indices and calculate the resulting satisfaction of the targeted users by the provided endowment. In equation (18), we initially consider the fraction α of customers that are satisfied to be equal to the fraction of customers that are not targeted. We then repetitively adjust α to include the number of targeted customers that are satisfied by their provided endowments and re-employ (18) to find the number of satisfied targeted customers, until the number of satisfied targeted customers remains unchanged, in a *fixed-point* manner. As depicted in Fig. 4b, *all* targeted users are satisfied by their provided endowments for all different desired total energy-consumption reductions considered both for honest and for anonymous strategic feedback. However, for named strategic feedback, a small fraction of targeted customers may remain unsatisfied. Thus, in presence of strategic liars, anonymous feedback is preferable for targeting purposes.Fig. 4b depicts the number of (satisfied) targeted users for achieving the different total energy-consumption reductions considered.

VII. Related Work

There exists an extensive literature in incentive-based DR and DLC programs [6], [7] that perform social welfare maximization, while minimizing electricity generation/supply. Below, we briefly discuss certain articles that are more closely related to our work. In [8], the problem of optimal incentive design for voluntary participation of electricity customers in a Direct Load Scheduling (DLS) program is studied. The incentives are posted by an aggregator in dynamically updated and publicly available tables for all users, which would then respond by deciding whether they want to participate or not, and how much laxity they wish to offer to the aggregator for scheduling their appliances. In the same direction, in [9], they study the design of an optimal contract between a DR aggregator (DRA) and a user for incentive-based demand response. They aim to maximize the utility of the DRA by incentivizing users both to exert maximal effort in reducing the load and to avoid falsifying their consumption baseline by compensating users in proportion to their reported energy-consumption reduction. In [10], the utility company proposes dynamic contracts to the users for load curtailment. The utility company employs a double-edged incentive, i.e., provide reward to the users that curtail their load and issue a fine to those that do not, in order to select the optimal set of targeted users and the curtailment amounts of their loads. Also, different incentives and targeting policies are proposed in [11]. They focus on ADR programs and propose an approach to determine the value of incentives to be offered to a user based on the individual assessment of her net benefit loss by energy-consumption reduction. In [11], they propose a targeting algorithm to select the optimal set of users to be targeted for DR and two accompanying policies that restrict (in different ways) the user discomfort caused. However, in [11], neither hidden information on user utility functions nor feedback on customer satisfaction by ADR incentives are considered. In [12], they aim to design an optimum scheme for achieving the maximum benefit out of a DR program and not only reduce costs and improve reliability, but also increase customer acceptance of the DR program by limiting price volatility. We also aim at increasing customer acceptance of ADR contracts, yet, based on customer feedback, as opposed to [12].

Non-monetary incentives can be strong motivators in some contexts and may be less expensive than monetary ones that would be required to generate a similar degree of behavioral change, as dictated by behavioral economics [13]. Towards social pressure, understanding the impact of altruism is an important topic in many research areas beyond computer science, such as economics, psychology and biology. The experiments designed by Leider et al. [14] show that directed altruism strongly impacts people's behavior in an allocation game, where players are allocated some total quantity. The effect of directed altruism is explored in Incentive Networks [15]. Participants in [15] are asked to make a contribution towards a global task and receive some sort of reward for it. The contribution of each person in [15] is a function of her expected reward and the expected rewards of others related to her.

To sum up, there exists no prior approach on the design of appropriate incentive-based ADR contracts that take into account behavioral characteristics of users and their explicit feedback on contract acceptance, as opposed to our work.

VIII. Conclusion

In this paper, we studied the calculation of ADR endowments for customers that sustain energy-consumption reduction based on anonymous and named feedback on customer satisfaction, which may be strategic or not, and proposed respective algorithms for the different cases. We also introduced an algorithm for selecting the optimal set of customers to sustain specific individual energy-consumption reductions, while being provided with satisfactory ADR incentives based on customer feedback, so as to achieve an overall objective on energy-load curtailment. This algorithm works for both known and hidden user utility functions. Our evaluation has shown the effectiveness of our various algorithms for the calculation of appropriate ADR incentives, in the cases of either full or hidden information on user utility functions, even in the presence of high fractions of strategic liars among customers. Also, we found that customer targeting is preferable for low ($<20\%$) desired energy consumption reductions, even for hidden information on user utility functions and even in the presence of 30% strategic liars. Overall, we have explained and showcased how satisfactory endowments (individual or not) for uniform or personalized load curtailment can be accurately calculated by the utility company based on customer feedback. Our formulation and approach are generic-enough to consider different user utility functions; yet, the consideration of different behavioral factors [5] is left for future work.

References

[1] C. Riker, K. Wang, and F. Yoo, "Energy Efficiency and Automated Demand Response Program Integration: Time for a Paradigm Shift," in *ACEEE Summer Study on Energy Efficiency in Buildings*, 2016.

[2] T. G. Papaioannou, G. D. Stamoulis, and M. Minou, "Adequate feedback-based customer incentives in automated demand response," in *ACM e-Energy*, June 2018.

[3] T. Dietz, "Altruism, self-interest, and energy consumption," *Proc. of the National Academy of Sciences*, vol. 112, no. 6, pp. 1654–1655, 2015.

[4] B. M. Hannon, "Energy, growth, and altruism," *Technological Forecasting and Social Change*, vol. 20, no. 3, pp. 173–197, 1981.

[5] E. Frederiks, K. Stenner, and E. Hobman, "The Socio-Demographic and Psychological Predictors of Residential Energy Consumption: A Comprehensive Review," *Energies*, vol. 8, no. 1, pp. 573–609, 2015.

[6] F. A. Qureshi, T. T. Gorecki, and C. N. Jones, "Model predictive control for market-based demand response participation," in *IFAC Proceedings Volumes*, vol. 19, 2014, pp. 11 153–11 158.

[7] P. Samadi, A.-H. Mohsenian-Rad, R. Schober, V. W. S. Wong, and J. Jatskevich, "Optimal Real-Time Pricing Algorithm Based on Utility Maximization for Smart Grid," in *IEEE SmartGridComm*, 2010.

[8] M. Alizadeh, Y. Xiao, A. Scaglione, and M. Van Der Schaar, "Incentive design for Direct Load Control programs," in *Allerton '13*, 2013.

[9] D. G. Dobakhshari and V. Gupta, "Optimal contract design for incentive-based demand response," in *American Control Conference*, 2016.

[10] A. Anastopoulou, I. Koutsopoulos, and G. D. Stamoulis, "Optimal targeting and contract offering for load curtailment in nega-watt markets," *IEEE TCNS*, vol. 4, no. 4, pp. 805–815, 2017.

[11] M. Minou, G. D. Stamoulis, G. Thanos, and V. Chandan, "Incentives and targeting policies for automated demand response contracts," in *IEEE SmartGridComm*, 2015.

[12] A. Asadinejad and K. Tomsovic, "Optimal use of incentive and price based demand response to reduce costs and price volatility," *Electric Power Systems Research*, vol. 144, pp. 215–223, 2017.

[13] D. Kahneman, "Maps of bounded rationality: Psychology for behavioral economics," *American Economic Review*, vol. 93, no. 5, pp. 1449–1475, 2003.

[14] S. Leider, M. M. Möbius, T. Rosenblat, and Q.-A. Do, "Directed Altruism and Enforced Reciprocity in Social Networks," *Quarterly Journal of Economics*, vol. 124, no. 4, pp. 1815–1851, 2009.

[15] Y. Lv and T. Moscibroda, "Incentive networks," in *AAAI*, 2015.