Flexibility Management for Residential Users Under Participation Uncertainty

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

Abstract— **Demand flexibility management, often by means of Demand Response (DR), can significantly enhance the stability of the electric grid and reduce the investment cost for infrastructure upgrades in case of dynamic energy mix with renewable sources. However, uncertainty in the consumer response to the DR signals may disrupt this goal. In this paper, we deal with the optimal management of the flexibility offered by residential users under uncertainty. We develop a probabilistic user model to account for the uncertainty in the actual provision of the flexibility by a user in conjunction with incentives' offered thereto, which we subsequently introduce in the Demand Response (DR) targeting process. We consider a suitable optimization framework to enable flexibility maximization and budget minimization as separate single-objective expressions with the appropriate constraints. We define representative problems and solve them numerically for a wide range of user parameters, in order to illustrate the applicability and accuracy of our method, and to extract valuable insights. Finally, we develop techniques to resolve practical issues and to enable real-world implementation of the proposed scheme in pilot sites; namely, a mathematical expression to estimate the confidence intervals of the attained flexibility and a learning algorithm for extracting the individual user parameters according to their participation patterns.**

Keywords—demand response, uncertainty, probabilistic user model, demand response targeting, optimization

I. INTRODUCTION

Effective demand-side flexibility management is key to the robust transition to a decentralised and decarbonised electric energy mix. Electric energy from renewable sources should be appropriately matched to the demand to ensure the stability of the electric grid, postpone huge investments in the grid infrastructure, and enable the effective use of renewable resources. Demand response (DR), i.e., demand management by means of consumer response to received signals, is often employed for flexibility management. Residential consumers are considered a prime source for flexible demand due to their varying consumption schedules. Incentivization of the household users is a key factor for the success of this approach since the user will have to tolerate a discomfort due to his modified consumption schedule. However, uncertainty on whether the user will indeed provide the requested flexibility even if being offered incentives may exist.

Relevant literature on uncertainty regarding the flexibility provision is rich. In [2], a programming model is presented that enables the delivery of DR by means of hot water storage. Optimization considers the lack of confidence regarding many factors including outdoor temperature and electricity consumption. The optimization is either constrained directly by a limit in the temperature deviation or indirectly by a price in the thermal discomfort, calculated by a proposed metric. In [3], DR is extracted by a Time-of-Use (TOU) method and an incentive-based scheme and is offered to the market. The Conditional Value-at-Risk (CVaR) measure is employed for the load uncertainty. It is found that increasing the risk factor leads to reduced consumer participation rates. In [4], the reaction of users to a flexibility signal is examined according to the context of the event by employing statistical methods. Certain methods are then proposed to maximize consumer participation based on the previous findings. The scope of [5] is the planning of users for load shifting according to their house and load characteristics. During the optimization phase of the DR amounts to be requested, the uncertainties due to different sources are also considered, such as thermodynamic conditions, occupant conditions, etc. In [7], a data-driven method is proposed to exploit the DR capabilities of a smart grid in a sustainable way. The willingness of users can be captured via a metric according to the historical DR rewards. Trade-offs are then considered between the system benefit and the sustainability of the DR program. In [8], the problem of hidden information regarding the consumer utilities is overcome by estimating the user satisfaction according to feedback from his reactions in previous DR offers. The problem is dealt with indirectly in [6], in the initial stages of a DR scheme, i.e., the user selection. Interest for DR and asset potential were deemed as the most influential factors for alleviating uncertainty. Some approaches consider the DR uncertainty due to the environmental conditions [2], [4], [5], or due to user consumption behavior [2], [3], [5], whereas others focus on the willingness of the users [7]. Others try to bypass the lack of information on the user premises [8] and others move a step back and take into account the user selection best practices to maximize involvement in DR.

The authors of [1] focus on the design of contract-based automated DR programs by energy providers and establish the formula for minimum the amount of incentives that should be offered to a consumer to accept such a contract. In this paper, we go one step further, by considering uncertainty in the provision of the actual flexibility by the user without focusing on the specific underlying sources of uncertainty. For example, if Direct Load Control (DLC) is applied, it is possible that a user does not accept the incentives offered thereto, and thus does not participate in DR; or if the user controls his own loads, it is possible that he does not manage to provide the flexibility promised, e.g., due to miscalculation of the necessary actions. Henceforth, we assume, for simplicity and clarity reasons, that DLC is applied. We

This work has been funded by the EU project iFLEX (grant no. 957670).

978-1-6654-3254-2/22/$31.00 ©2022 IEEE

approach the matter of uncertainty, by linking the probability of user participation with the amount of incentives provided. Initially, we define a user model that captures (1) the minimum acceptable incentivization for the user to seriously consider the provision of flexibility, and (2) how responsive the user is to the provision of DR incentives to actually decide to participate. Next, we propose an optimization framework that incorporates maximization of the total flexibility to be extracted from the users, under a budget constraint, and minimization of the expenditure for DR incentives under a minimum total flexibility constraint. Then, we define special cases of the optimization problems on the basis of simplifying assumptions, and we solve them numerically, to: (1) characterize the optimal targeting policies and the impact of certain parameters to the optimal values of the objective function under certain assumptions on the user model, (2) evaluate the positive effect that the flexibility resale may bring, and (3) assess the extent of the positive impact of approaches, revealing information that was not completely known to the provider, to the optimal value of the flexibility and to the user's well-being.

Finally, we explore several practical issues for real-world implementation of the schemes proposed, namely in one of the pilots of the EU-funded project iFLEX (https://www.iflex-project.eu/). In particular, we propose a method to calculate the confidence intervals of the expected flexibility, we apply it and we evaluate numerically a practical approach of ensuring the desired flexibility target is achieved with high probability, despite uncertainty. Moreover, we propose a practically applicable algorithm that monitors the positive or negative reactions of the users and gradually learns their profile by determining the respective parameter values for each user separately, as well as a scheme that further accelerates this learning procedure. Fig. 1 below illustrates in a compact manner the respective logical flow of our work, which is built gradually from modeling to optimization, then to analysis of the flexibility management approach developed, and finally to issues related to its practical applications.

II. User Modelling

We consider a provider/aggregator targeting users for DR and offering incentives to them. We define a model regarding the uncertainty (which was motivated in Section I) on whether the desired flexibility will indeed be attained by each of the targeted users. This model pertains particularly to the selection of DR incentives thereto for motivating a user to indeed refrain from using certain electrical devices. Since we have taken that DLC is applied, this amount to participating in the DR event; namely, accepting the DR incentives offered and allowing control of the respective loads by the provider.

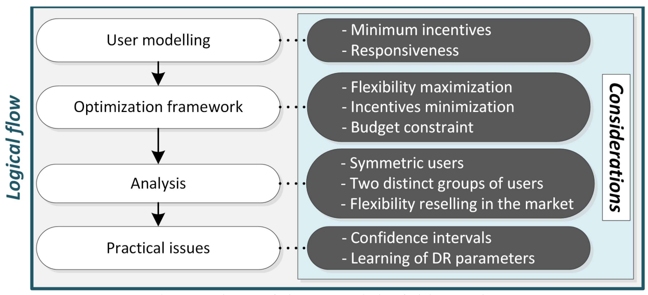

Fig. 1: Flow of the material of this work

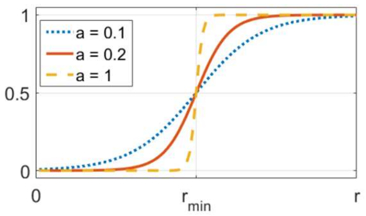

Fig. 2: Modified sigmoid function for various values of *a*

We assume that there are N users, indexed $1, \dots, N$. A subset of them will be targeted (i.e., selected) for DR. Thus, we define a binary variable y_n per user, where $y_n = 1$ if user n is targeted and $y_n = 0$ otherwise.

If user n is indeed targeted for a certain DR event concerning a particular time slot, then he is offered incentives r_n, in order to meet a demand flexibility (i.e., reduction of consumption in a particular time-zone) equal to x_n, which is different per customer n and depends on his load consumption profile. We assume that the demand flexibility x_n requested by user n (e.g., turn-off an electric device for a time period) does not depend on the offered incentives, but on his baseline consumption schedule.

In this model, due to our assumption about DLC, the user is faced with two choices: either to accept the offered incentives r_n and indeed provide the flexibility x_n, or to reject the incentives and maintain his consumption schedule. According to [1], if the DR incentives cover user's discomfort (i.e., loss of utility) due to not using certain loads, discounted by the savings in the energy bill, if $r_n \geq NBloss(n)$ (i.e., loss of Net Benefit due to DR), then participation in the DR event is the optimal decision for the user. Therefore, ideally, we can define the minimum acceptable incentives of user *n*, henceforth denoted as $r_{\min(n)}$, and take that the probability $p_n(r_n)$ for user n to participate in the DR event is a step function, rising from 0 to 1 at $r_n = r_{\min(n)}$.

In order to incorporate uncertainly in our analysis, and thus make it more general, we take the outcome of DR as an outcome of a Bernoulli trial, with a success probability $p_n(r_n)$ that depends on the economic incentives. In case of failure in this trial, we take that user n does not participate in DR, and thus he does not attain any flexibility and he is not paid the incentives r_n initially offered to him, and vice versa. This participation-probability function should have the following properties: (1) $p_n(r_n)$ is increasing, continuous and differentiable in the incentives r_n, (2) $p_n(0) = 0$ and $p_n(\infty) = 1$, and (3) $p_n(r_n)$ ascends steeply from low to high values around $r_{\min(n)}$.

Therefore, $p_n(r_n)$ constitutes a smooth approximation of the step function discussed above. Employing such a function rather than the unit-step function also allows for cases where the user can accept (resp. reject) somewhat lower (resp. higher) incentives than $r_{\min(n)}$ since his discomfort by not using the electrical device at the specific time slot can occasionally be slightly lower (resp. higher). A function possessing all of the above properties is the sigmoid function. For a user *n* with minimum acceptable incentives $r_{\min(n)}$, we can take that:

$$p_n(r_n) = \frac{1}{1 + e^{-a_n(r_n - r_{min,n})}} \tag{1}$$

which is a slightly modified version of the sigmoid function, for which $p_j(r_{\min(n)}) = 1/2$ regardless the value of a_n, which however determines how steeply the function rises. In particular, the larger a_n, the steeper the function. Fig.2 illustrates the function shape according to various a_n values.

III. Optimization Framework

The optimization problems of the flexibility aggregator can now be specified. We take that the latter has a total budget B available for DR incentives. His objective is to maximize the expected total flexibility X_E, without exceeding the total DR incentives budget:

$$\begin{gathered} \max \sum_n y_n \cdot x_n \cdot p_n(r_n) \\ s.t.\ y_n \in \{0,1\}\ and\ \sum_n y_n \cdot r_n \le B \end{gathered} \tag{2}$$

where y_n for $n = 1, \dots, N$ constitute binary decision variables for targeting user n and r_n are the incentives offered to user n. For obvious reasons, we assume that a user n that is not targeted (i.e., if $y_n = 0$) is offered no incentives (i.e., $r_n = 0$). By monotonicity, the budget constraint will be met with equality under the optimal solution.

Extracting the optimal targeting by solving the above problem may prove a conservative and "generous" approach for the provider. Indeed, certain targeted users may ultimately decide not to accept the incentives offered and/or not meet their DR objective and thus rightly not be rewarded by the provider. In such a case, part of the total DR incentives budget would be left unused. To improve on this, we consider alternatively a looser constraint, restricting the expected total DR incentives actually paid. For simplicity, we refer to this metric as the expected total DR incentives, or the expected total reward. The optimization problem becomes:

$$\begin{gathered} max \sum_n y_n \cdot x_n \cdot p_n(r_n) \\ s.t. \sum_n y_n \cdot r_n \cdot p_n(r_n) \le B \end{gathered} \tag{3}$$

The optimal solution of the previous problem is clearly a feasible solution of the problem with the less tight constraint. Therefore, a higher X_E can now be attained. However, it is possible that the total reward actually paid occasionally exceeds the threshold B.

An alternative (dual) optimization problem is to seek for the minimum expected total DR incentives (reward) that is necessary for the flexibility aggregator to meet a particular threshold X for X_E. This problem is formulated as follows:

$$\begin{gathered} \min \sum_n y_n \cdot r_n \cdot p_n(r_n) \\ s.t. \sum_n y_n \cdot x_n \cdot p_n(r_n) \ge X \end{gathered} \tag{4}$$

Similarly, to the above, it is possible that X_E occasionally is lower than the desired level X.

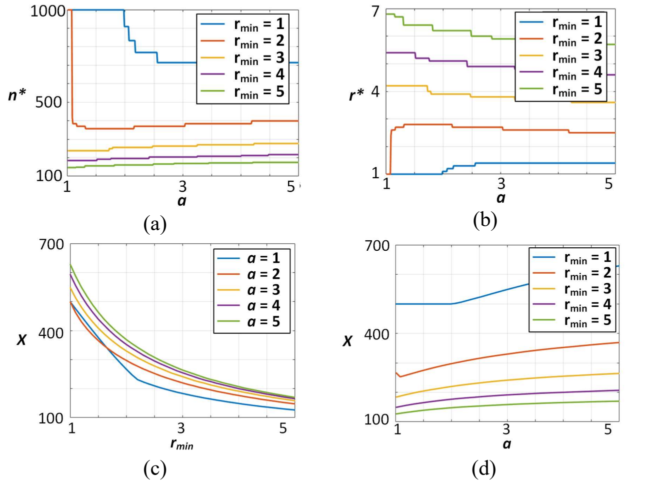

Fig. 3. Relationship of (a) n^* and (b) r^* with a for various values of r_{min}. Relationship of X_E with (c) r_{min} for various values of a, and (d) a for various values of r_{min}.

IV. Analysis

Below, we study three special case studies of the above problems. Each case study has a different scope. In particular, the first one involves symmetric users, and mostly concerns the influence of user parameters and different constraints to the decision variables and the values of the metrics of the DR objectives. The assumption of symmetric users leads to the simplest possible set of parameters and nullifies any influence of the degree of user heterogeneity. This assumption is relaxed in the second case study. This examines the positive effect that successful discrimination of users belonging to two (for simplicity) different groups may bring to the provider in terms of flexibility. Both of the above case studies assume that the optimization objective is to maximize the expected flexibility. Thus, the third one explores the implications and the benefits of flexibility resale for the user and the provider. Due to their different underlying assumptions, it is hard to compare these case studies, which however, all provide fruitful insights regarding the various factors that make up the success of a DR program.

A. Case Study 1: Symmetric users

In the first case study, users are considered symmetric, i.e. $x_n = x$, $r_{\min(n)} = r_{min}$, $a_n = a$, which implies that $p_n(.) = p(.)$. The problem of maximizing X_E for a given incentives budget amounts to deriving the optimal number n^* of users to be targeted and the optimal reward r^* to be offered to each of them, which is given by:

$$\begin{gathered} \max\ x \cdot n \cdot p(r) \\ \text{s.t. } n \cdot r \le B, \quad n \in \mathbb{N} \end{gathered} \tag{5}$$

because, at the optimal point, due to symmetry, all targeted users should be offered the same incentives r. Clearly, n^* and r^* are constrained by B and depend on the parameters r_{min} and a of the sigmoid participation probability function. Next, we investigate the dependence and the monotonicity properties of n^* and r^* on r_{min} and a.

Fig.3a and Fig.3b depict the relationship of n^* and r^* with a, respectively, for different values of r_{min}. Fig.3a shows that when r_{min} and a both have low values, the max. X_E value is attained by targeting all the users. The number of

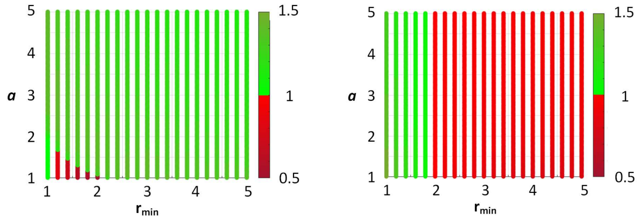

Fig. 4. (a) r/r_{min} ratio with respect to r_{min} and a, under the constraint $n \cdot r \leq B$, (b) r/r_{min} ratio with respect to r_{min} and a, under the constraint $n \cdot r \cdot p(r) \leq B$.

targeted users decreases, however, as the probability function becomes steeper (a rises) and r_{min} increases. Fig.3b shows that the increase of r_{min} also boosts r^* for each user. As a increases, r^* drops. Fig.3c and Fig.3d exhibit how X_E relates to r_{min} and a. Smaller values of r_{min} and greater values of a lead to superior flexibility. The value of r_{min} has a deeper effect on X_E, than a. When a is big, lower incentives are required to achieve the same amount of flexibility.

Next, we consider the problem of maximizing X_E with the alternative constraint concerning the expected total reward. The optimization problem is now defined as follows:

$$\begin{gathered} \max \ x \cdot n \cdot p(r) \\ \text{s.t. } n \cdot r \cdot p(r) \leq B, n \in \mathbb{N} \end{gathered} \tag{6}$$

The relation of the ratio r^*/r_{min} with r_{min} and a, is shown in Fig.4 for the two types of constraints regarding B. It can be seen that the $n \cdot r \cdot p(r) \leq B$ constraint involves optimal incentives that are less than the minimum acceptable ones ($r^* < r_{min}$) for the majority of considered parameters. In other words, this constraint leads to the counterintuitive and "risky" policy of offering low incentives to many users. This may also lead to significant budget overshoots in case more users than those expected actually participate in DR. On the other hand, under the $n \cdot r \leq B$ constraint, the optimal incentives are in general higher than the minimum acceptable ones ($r^* > r_{min}$). This constitutes a robust policy, since each targeted user now has a high participation probability.

B. *Case Study 2: Two distinct groups of users*

In this problem, users are considered to belong to two types (of identical users each), namely the "small" users with a low flexibility capability x_{low} that also require low incentives r_{low}, and the "big" ones with x_{high} and r_{high} respectively; e.g. single-person households and family-households. For simplicity, we can take that a is common regardless the type. We consider the case where the provider-side cannot distinguish to which type each user belongs. To make a fair comparison with the case where the group to which each user belongs can be identified, we take that in this case the provider knows the proportion of users in each type, namely q_{low} and $q_{high} = 1 - q_{low}$, the average flexibility $x_{mean} = x_{low} \cdot q_{low} + x_{high} \cdot (1 - q_{low})$.

Note that for simplicity we take $N_{low} = N \cdot q_{low}$ and thus $N_{high} = N \cdot (1 - q_{low}) = N - N_{low}$ are both integers. Due to not knowing the exact type per user, the provider solves the optimization problems by considering that all users are identical, i.e., he takes that $x_n = x_{mean}$ and offers each targeted user the average of the optimal pair of incentives, i.e., $r_n = r_{mean} = r_{low} \cdot q_{low} + r_{high} \cdot (1 - q_{low})$, of the case that the user types can be distinguished. The benefit for the provider of being able to distinguish the user types can be quantified by comparing X_E to that of the previous case.

TABLE I. MEAN FLEXIBILITY GAINS FOR SPECIFIC PARAMETER VALUES

x_{high}/x_{low} ratio	N_{low}	Mean X_E gain in %
3	200	10
6	200	13
3	400	26
6	400	35
3	600	49
6	600	73
3	800	76
6	800	142
3	200	10

TABLE II. MEAN FLEXIBILITY GAINS FOR SPECIFIC PARAMETER VALUES

$r_{min,high}/r_{min,low}$ ratio	a	Mean X_E gain in %
<2	<1	32
>2	<1	50
<2	>1	33
>2	>1	51

A numerical verification is conducted, regarding the comparison specified above. The parameter space is swept, and the flexibility gains are observed as a percentage of the improvement of X_E when there is lack of knowledge of the user groups, assuming that the respective user groups can indeed be identified. Table 1 shows how the mean X_E gain varies according to the x_{high}/x_{low} ratio and N_{low}, and Table 2 according to the $r_{min,high}/r_{min,low}$ ratio and the parameter a respectively. It can be seen that the $r_{min,high}/r_{min,low}$ ratio and N_{low} have the highest impact on the flexibility gains when the user groups are known. It should be noted that the unused budget is always less than 2%, thus implying that high enough incentives (leading to high participation probability) are offered in general. Overall, the average X_E improvement for the provider is 43% on the complete parameter map exploration, which implies that extraction of the per user type by the provider is a very beneficial ability for him.

C. *Case Study 3: Flexibility reselling in the market*

Next, we investigate the case study where the flexibility attained is resold in the market; then, the provider earns some revenue and achieves an economic profit P. This profit equals the revenue from this resale minus the expected total incentives paid. If pricing of flexibility is linear, i.e., each flexibility unit is resold at a price q, then the relevant optimization problem is as follows:

$$\begin{gathered} \max q \sum_n [y_n \cdot x_n \cdot p_n(r_n)] - \sum_n [y_n \cdot r_n \cdot p_n(r_n)] \\ \text{s.t. } \sum_n y_n \cdot r_n \leq B \end{gathered} \tag{7}$$

Initially, the effect of the total X_E target and the selected B is studied with respect to P. Similarly to the previous case studies, we sweep the parameter space and optimize for P in every set of parameter values. The first set of results (Fig. 5a) show how P relates to B for various values of r_{min}. It can be

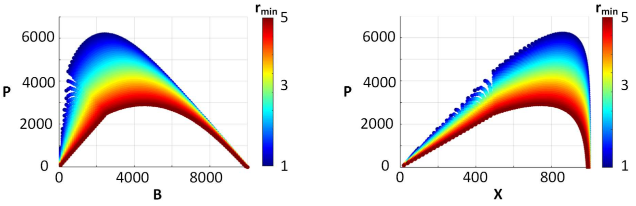

(a) (b)

Fig. 5 Profit of the provider P with respect to (a) B (b) X_E, for various values of the minimum acceptable incentives r_{min} of the users.

seen that the profit curve comprises three segments: (1) A linear segment, for small values of B. In this region, P increases almost linearly with the increase in the number of targeted users. Every single new user entering the DR event is a new source of profit and the profit equals practically the product of the profit from each user with the number of the participants. (2) A saturated segment, where the profit margin begins to shrink until it reaches its peak value. In this area, the maximum number of users is targeted, and incentives are increased across all users to increase their participation probability. (3) A decreasing segment, where regardless of the increase in the incentives, the probability of the users participating is already very close to 1, and thus any further increase in the budget does not provide any gain. It can also be seen in Fig. 5a that the r_{min} value strongly affects the slope of the linear part and the maximum profit, but not significantly the profit in the highly saturated area (decreasing part).

Fig. 5b shows how P relates to X_E for various values of r_{min}. The optimal profit is maximized for a flexibility value that is smaller than the maximum that can be attained. Indeed, the max. X_E is accompanied by the saturation of the participation probability p, which is achieved by offering high incentives, thus leading to low profits. The value of r_{min} has a significant impact on the value of P.

In Fig.6, we turn our attention to the Net Benefit (N.B.) of the user and examine how this relates to $P$and X_E for various values of r_{min}. In Fig. 6a, three regions can be distinguished: (1) A constant N.B. region, where its value almost coincides with the x-axis and remains constant, while P increases. The targeted number of users here is less than maximum. The N.B. is the same for all the selected users and does not depend on their number. (2) A region of mutual profit (for the DR aggregator and the user), where the increase in P also leads to an increase in user N.B. This is the case where incentives improve across all users, after they have been all targeted. The provider gains in P due to the higher probability of DR

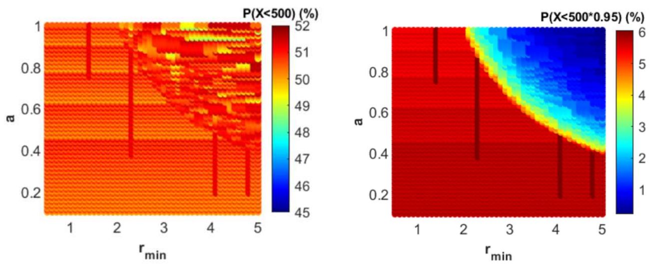

(a) (b)

Fig. 7 (a) Possibility that the actual flexibility value obtained is below its nominal value according to the binomial distribution with respect to the r_{min} and a parameters. (b) Similarly, with the 95% of the nominal value.

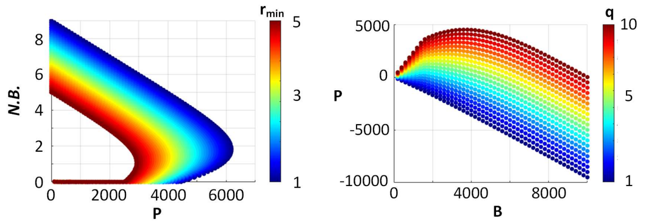

(a) (b)

Fig. 6 (a) User Net Benefit (N.B.) plotted against the profit of the provider P for various values of the minimum acceptable incentives rmin of the users, (b) Sensitivity analysis of the profit of the provider P with respect to B for various values of the resale price q.

participation, and users benefit as well. (3) The almost-maximum flexibility region, where P decreases, N.B. continues to grow indefinitely and X_E is asymptotically maximized.

Next, we investigate the effect of the reselling price q per flexibility unit to the required B and P. The corresponding results are depicted in Fig. 6b. We can observe that the profit of the provider is increasing in q. This is expected since $P = q \cdot X_E - B$, for a given B. To achieve the maximum P for a specific value of q, the budget B should be adjusted accordingly. In general terms, as the resale price increases, the budget should also increase, in order to attain a higher profit. Moreover, there is a region for certain ranges of B and q where the profit is negative, i.e., it corresponds to monetary loss for the provider. In such a case, the DR program is not beneficial for the provider and cannot even be sustained.

V. Practical Issues

A. Confidence intervals

An important aspect of the optimization problem is defining the confidence that the targeted X_E will indeed be obtained. In the case of identical users, the total number of users participating successfully in DR follows the binomial distribution. Calculation of bounds in the possibility that a random binomial variable deviates by a percentage from its mean value can be calculated by the probability mass functions of the binomial distribution, which is an accurate, although computationally intensive way.

We shall run an optimization problem to track the optimal number n of targeted users and the incentives r to be offered to each one of them under the criterion of minimizing the total incentives offered. At the same time, the respective bounds from the binomial distribution will be calculated. We assume

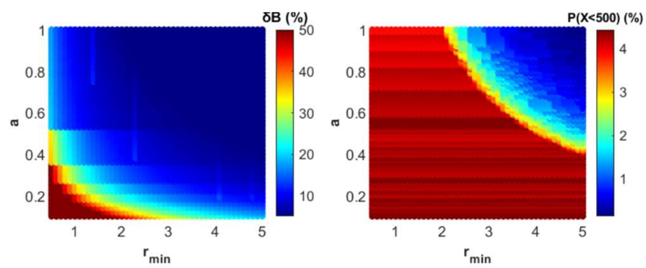

(a) (b)

Fig. 8 Differentiation of various quantities when optimizing with a X_E target of 500/0.95=526.3 instead of 500. (a) Difference of total budget expenditure in % and (b) probability that $X_E < 500$ with respect to r_{min} and a parameters.

a homogeneous population where all users are characterized by the same a and r_{min} parameters. We take that the aggregator budget is not constrained and that the minimum total flexibility to be attained is 500. In this optimization case study, the possibilities of X_E deviation from two types of boundaries is also calculated and presented. It can be seen in Fig. 7a that the possibility of the actual value being less than the nominal is significant and is between 44% and 49%. As we begin lowering the minimum bounds that we are willing to accept then the possibility steeply drops, e.g., if we are willing to accept a $0.94 \cdot 500 = 470$ value then the likelihood of a lower value than that is below 3%.

Now, let's suppose that we desire a low probability of the flexibility dropping below the nominal value, e.g. 5%. We can run the optimization again for a greater flexibility target that is selected, for example as $500/0.95 = 526.3$. We could see how the new incentives' offers and the rest of the quantities differ with respect to the previous optimization. We can see in Fig.8a that B is not significantly affected, except for the cases of highly unresponsive users (low a). Most importantly, we see in Fig. 8b that the likelihood of X_E dropping below 500 is always lower than 5%. If flexibility to be obtained is found to be unacceptable then the target could be adjusted upwards in order to achieve the desirable confidence, according to, e.g., the penalty.

B. Learning of DR parameters

In this section, we introduce and evaluate an algorithm to identify the user DR parameters of each user without previous knowledge whatsoever (Parameter Learning Algorithm – PLA). With those parameters captured, the p of each user will be fully known, the real user will thus be modelled, and the optimization algorithms of the provider will be in position to provide the optimal targeting.

The basic concept of the algorithm can be outlined as follows. Initially, the provider begins offering DR incentives randomly. This approach can be employed for all users at the beginning of the DR program, and also upon entry of every new customer in the program. While incentives are offered, the participation or not of the user in each DR session is observed and recorded (as a binary variable), along with the value of the respective incentives. For each incentives' value, the ratio of the number of times that the user participated to the number of times the user was targeted, constitutes the participation rate. The participation rates for each value of incentives make up the user participation rate pattern. If the user parameters were known, a similar pattern could be calculated with the probability function and this would be equal to the user participation rate pattern, at least theoretically for infinite DR attempts. We will show, however, that only a relatively small number of attempts suffices for the two patterns to be almost equal. Continuing with the description of the algorithm, in every DR attempt, the user participation rate pattern is calculated. The respective patterns for all the possible a and r_{min} values have been also calculated and the two are continuously compared. The pair of parameter values that constitutes the best match for the two is the solution.

The respective mathematical representation follows. To simplify the formulation, the expressions below refer to a single user n. The participation rate pattern s for the value r of DR incentives can be defined as:

$$s_r = \frac{q_{r,k(r)}}{k(r)} \tag{8}$$

where $q_{r,k(r)}$ is the number of times the user participated in DR and $k(r)$ the number of times the user was targeted with DR incentives equal to r. Since $q_{r,k(r)} \leq k(r)$ there follows $0 \leq s_r \leq 1$. This rate is different for each incentives' value r. This pattern is updated after each DR session. The respective pattern, when calculated by the user model is provided by (1). Thus, the goal of the algorithm is then to specify the optimal values for parameters a and r_{min} that minimize the Root Mean Square Error (RMSE) between s and p.

To illustrate the method attractiveness, we employ the algorithm to identify 3 random users, A, B and C with a 0.1, 0.5 and 1 and r_{min} 50, 80, 20 respectively. Parameter values have been selected in a such way that different shapes of the probability function are acquired. Regarding the incentives, we have assumed two schemes: random offers (open loop procedure) and a predictive method (closed loop procedure). In the open loop scheme the incentives are provided randomly. In the closed loop scheme, the current estimation of r_{min} and a values in each iteration are employed to predict more relevant incentive values for the next DR event. More specifically, for the current r_{min} and a values, the incentive values corresponding to probabilities 10% and 90% are calculated. The next offered incentive will be a random number, uniformly chosen between these two values.

The procedure is repeated 10 times, the results are averaged and illustrated in Fig.9. Blue color corresponds to random offering and red color to the predictive method. The following conclusions can be extracted: (1) It can be seen that only a relatively small number of attempts (i.e., between 5-20 and sometimes even smaller) suffice to approach the r_{min} parameter with remarkable accuracy. (2) Prediction of responsiveness a is more challenging. However, precise

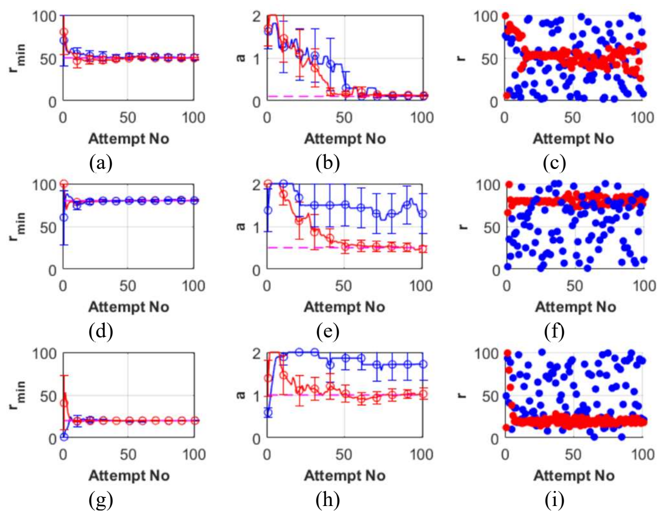

Fig. 9 Parameter identification efficiency test for the 3 random users (user A: a, b, c – user B: d, e, f – user C: g, h, i). The first column (a, d, g) illustrates the r_{min} convergence, the second column (b, e, h) the a convergence and the third column (c, f, i) shows the values of the r offerings. Blue color corresponds to the random method and red color to the predictive one.

identification of the value of this parameter is of considerably lower importance than the identification of r_{min}. (3) The random offers' method is efficient towards capturing r_{min}, but is rather poor regarding α. (4) The predictive method improves even further the speed of convergence of r_{min} and enhances drastically the speed of identifying parameter a.

VI. Practical Implementation and Implications

The proposed methodology can be directly applied in practice as follows. We may assume that we only have a single DR slot every 24 hours. Initially, a random set of users can be selected and be provided with personalized arbitrary incentive values. Using the PLA (Parameter Learning Algorithm), their parameters will be discovered on a per user basis with significant accuracy within a number of different DR events. Afterwards, the optimization scheme from equation (2) can be employed to maximize the total flexibility under a budget constraint, or vice versa as prescribed in equation (4). In the latter scheme, where a flexibility constraint is considered, the confidence intervals' considerations (Section V) should be employed to avoid violation of the constraint. The optimization calculations have to be repeated in every DR event, while the PLA shall keep gathering results to update user parameters. The scheme can be extended to account for estimations concerning multiple slots per day.

The practical implementation is also fueled by the insights of the analysis of Section IV. In particular, the symmetric users' case study provides an overview of what results and trade-offs are to be expected when the mean values only of the users' parameters are known. Moreover, the case study with the two distinct user groups has proved that it is advantageous for the provider to be able to identify two different user groups, instead of one (symmetric users). It is reasonable to claim that the more user groups identified the better. PLA was therefore designed to discover user parameters in the highest resolution, i.e. on a per user basis. The case study involving flexibility reselling in the market has shown that with careful selection of parameters, i.e. incentives budget and flexibility unit resale price (if negotiable) can lead to "all-win" situations which are beneficial for all players involved in the DR value chain. For example, the Transmission Network Operator (TSO) and the Distribution Network Operator (DSO) suffer from the fluctuations of the transmitted and the distributed power. Steady flexibility provision on the basis of this framework shall facilitate a steadier power supply leading to less faults, less expansion costs and less congestion issues. Balancing Service Providers (BSPs) and Balance Responsible Parties (BRPs) would also employ this new source of flexibility that can be cheaper than the existing ones, such as employing the production of expensive units to meet an unexpected demand increase. Incentivization of users to lower their consumption can be significantly less costly than operating e.g., a thermal power plant or investing in a huge battery array. Finally, the self-consumption of an energy community can also be optimized by means of the flexibility management mechanism, and thus promoting energy self-sufficiency of the community.

VII. Conclusion

We defined an optimization framework that considers uncertainty in the user responsiveness in DR for flexibility management subject to budget constraints for incentives. We numerically analyzed the influential factors for flexibility management under such uncertainty and we proposed a practical, yet effective, learning algorithm to identify the minimum acceptable incentives and the DR responsiveness of users in real deployments. Our main findings are as follows: The aggregator should locate and engage the highly responsive users, since the DR program will be more successful in terms of flexibility and budget expenditure. When users are heterogeneous, the aggregator should classify users in groups of similar characteristics for maximizing his revenue profit without increasing his budget. The more diverse the user groups are, the stronger the advantage of discriminating them. The "strict" constraint on the total budget for DR incentives should be considered in practical cases, as opposed to the "loose" one, because it leads to a robust incentive policy with acceptable rewards for the users, despite the somewhat lower total flexibility attainable. When the flexibility is resold in the market, the maximum provider profit can be accomplished with a significant benefit for the user, so that both sides are satisfied. As a future work, we intend to validate our findings and the proposed learning approach for flexibility management with real users.

References

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

[2] N. Good, E. Karangelos, A. Navarro-Espinosa and P. Mancarella, "Optimization Under Uncertainty of Thermal Storage-Based Flexible Demand Response With Quantification of Residential Users' Discomfort," IEEE Transactions on Smart Grid, vol. 6, no. 5, pp. 2333-2342, 2015.

[3] M. Vahid-Ghavidel, M. S. Javadi, S. F. Santos, M. Gough, M. Shafie-khah and J. P. S. Catalão, "Optimal Stochastic Conditional Value at Risk-based Management of a Demand Response Aggregator Considering Load Uncertainty," IEEE International Conference on Environment and Electrical Engineering and IEEE Industrial and Commercial Power Systems Europe, 2021.

[4] A. Bortolato, P. Faria and Z. Vale, "Probabilistic Determination of Consumers Response and Consumption Management Strategies in Demand Response Programs," IEEE/PES Transmission and Distribution Conference and Exposition (T&D), 2020.

[5] K. Paridari, L. Nordstrom and C. Sandels, "Aggregator strategy for planning demand response resources under uncertainty based on load flexibility modeling," IEEE SmartGridComm 2017.

[6] M. Curtis, "Demand side response aggregators: How they decide customer suitability," International Conference on the European Energy Market (EEM), 2017.

[7] B. Zeng, X. Wei and J. Feng, "A Data-Driven Dispatching Approach for Sustainable Exploitation of Demand Response Resources," IEEE SmartGridComm 2018.

[8] T. G. Papaioannou, G. D. Stamoulis, and M. Minou, "Adequate Feedback-based Customer Incentives in Automated Demand Response," ACM International Conference on Future Energy Systems (e-Energy '18), Karlsruhe, Germany, 2018.