The Effect of Altruism in Automated Demand Response for Residential Users

Marilena Minou, George D. Stamoulis, Thanasis Papaioannou
Department of Informatics, Athens University of Economics and Business, Greece
{minou, gstamoul, pathan}@aueb.gr

Abstract— Automated Demand Response (ADR) programs play key role in alleviating the energy supply and demand imbalances by (i) controlling user loads either directly or indirectly, and (ii) economically mitigating the uncertainties that impact power system operations in an automated (pre-contracted) way. In general, users are assumed to act rationally, i.e., optimize their decision-making process so as to maximize their financial net benefit. However, an extensive literature on behavioural economics (BE) contends that the decision-making process of users is far more complex, not always self-interested and depends on a number of individual factors, such as altruism. Our work aims to advance our knowledge on how to engage in ADR contracts users that may exhibit different degrees of altruism and motivate them effectively to ultimately optimize the overall use of energy. We investigate the impact of altruism on the total financial incentives to be offered by the provider and on the social welfare, and identify the optimal demand reduction and user targeting strategies for performing ADR in such populations. Based on experiments with real and synthetic data, we find that appropriate targeting policies and demand reduction strategies that take advantage of altruism can be beneficial for the social welfare of the users and the incentive costs of the provider. However, leveraging altruists should be performed carefully, since saddling them with high power reductions, although yielding lower total incentives, can prove inefficient for the social welfare of the system.

***Index Terms*-- Altruism; Behavioural Economics; Demand Side Management; Incentives**

I. INTRODUCTION

Demand Response (DR) encompasses the modification of the normal demand patterns of end-users in response, either to changes in the price of electricity over time, or to incentive payments designed to induce lower consumption, when system reliability is jeopardized. The success of a DR program greatly depends on users' participation, which is highly affected by the level of discomfort caused during a DR event due to the modification of users' consumption patterns and the rewards offered by the provider. Traditionally, users are assumed to act rationally by choosing that consumption schedule that maximizes their total net benefit, i.e., the utility gained minus the monetary charge. Behavioural factors, though, play critical role in shaping users' behavior and decisions. Our contribution constitutes a first, yet innovative endeavor to consider one aspect of behavioural traits, namely altruism, in DR programs. According to social experiments in [1], some degree of altruism, i.e., the unselfish concern for the welfare of others, is common in the society. In technical terms, altruism means that the first derivative of the utility function of an individual with respect to the material resources received by any other agent is always strictly positive. Leveraging it in our work aims to advance our knowledge of how to engage users that exhibit some degree of altruism towards other users and motivate them effectively, in order to ultimately optimize their use of energy. We focus on the case of contract-based automated DR (ADR) programs while taking into account users' preferences and the external context. Our approach presumes consumption data gathering and processing. This data is used to reveal a user's preferred consumption pattern and her discomfort by altering it. It reckons historical demand and other data revealing the context that affected the consumption, e.g., weather.

In our work, we assume that the provider has estimated accurately the user utility functions with respect to energy consumption as well as the baseline demand. This can be accomplished by assuming a generic multi-parametric utility model and find its parameters for each individual user based on a machine-learning algorithm. This process falls beyond the scope of the present work. We also consider that the provider has an upper threshold on the daily energy demand that may be served, above which the marginal energy production cost becomes very high due to the activation of costly supplementary energy generators. Whenever it is predicted that the total demand will exceed that daily threshold, the provider activates the ADR programs in order to restrict the total demand below the threshold while minimizing the total incentives offered. Our objective is to investigate the impact of altruism on the total incentives offered and the social welfare of the users after ADR. Moreover, we identify the optimal demand reduction and user targeting strategies for performing ADR in populations where some degree of altruism may be anticipated. We assume that the energy provider is aware of the degree of altruism in the user community through data mining techniques, surveys, social experiments [1], etc. We adopt in part the model of altruism presented in [1], based on which a user's overall utility function is given by a combination of his own monetary payoff and a "disinterested social welfare function", and propose different ADR policies for exploiting this aspect. We evaluate our model using real and synthetic energy consumption data corresponding to user trials in Lulea, Sweden. Our findings suggest that appropriate targeting policies and demand reduction strategies that exploiting altruism can be beneficial for the users, in terms of social welfare losses, and for the ADR provider, in terms of incentive costs. However, leveraging of altruists should be performed carefully, as their saddling with high reductions of power although yielding in low total incentives, can yet prove inefficient for the social welfare of the system.

II. BACKGROUND ON BE AND RELATED WORK

While the literature on DR is extensive, it is almost always assumed that users act rationally by choosing that consumption schedule that maximizes their total net benefit. In stark contrast, many scientific works as well as empirical evidence from the application of BE demonstrate that people are rarely the rational decision-makers envisaged by the traditional economics [2]. In particular, [2] claims that psychological factors, such as altruism, and intrinsic and extrinsic rewards, play key role in decision making and result in users showing non-individually-optimizing behaviour in their decisions. Various experimental studies aim to prove that behavioural interventions influence user consumption through increasing awareness of social norms [3]. It is argued in [4] that users contribute to social welfare even though they are better-off without doing so. Several studies have shown that social incentives can be more effective compared to monetary ones. If a person is intrinsically prompted to be altruist, offering a monetary incentive for achieving the desired behaviour can have a counteractive effect of "crowding out" the reward [5].

The impact of DR programs on the peak reduction considering the loss-aversion on the perception of users is studied in [6]. Also, the factors of influence between self-reported distrust and users' willingness to participate in direct load control programs are examined in [7]. Our aim is to explore the impact of altruism within energy consumption of residential users by means of ADR programs accompanied with demand reduction and targeting policies, which intend to leverage altruists in the interest of provider's objective regarding the total incentives offered and the social welfare achieved after the ADR.

III. The Model

In this work, we focus on an energy provider that owns supportive generators to meet excess demand and employs ADR in different population mixes as a means to avoid their costly activation. We build upon and extend our prior work in [1] on optimal ADR policies for rational users. In order to represent the altruistic nature of users and its effect on the decision making process, we tailor the models presented in [1]. We first describe briefly the system model and its main participants maintaining the notation of [1]. We consider a set of N households served by a single energy provider. Households have already signed contracts, according to which, during a DR event they give up control of specific appliances if targeted by the provider. We henceforth view each household as a single user that collaborates with the utility company (energy provider) as agreed. Note that, throughout the entire section, we utilize the notions of user and user alternately in order to refer to a single household. To make things more concrete, one can think that our basic time frame in which DR is applied is again a single day divided into timeslots (say 24 hours or 4 6-hour periods), indexed as $t \in T \coloneqq \{1, 2, \ldots, T\}$. This day corresponds to a known context; e.g. warm summer weekday.

A. The Users

According to [1], each user i operates a set of appliances A_i such as air conditioning, refrigerator, etc. For each appliance $a \in A_i$ of user i we denote by $q_{i,a,t}$ its power consumption during timeslot t and by $\overrightarrow{q_{\iota,a}}$ the vector $(q_{i,a,t}, \forall t \in T)$ of power consumptions, all applying for the day considered. Each user i is characterized by an optimal consumption Q_i and the associated optimal daily consumption vector for the day considered and is charged according to a given price p_t, which corresponds to the per unit of consumption charge and is previously announced by the provider. For simplicity we assume that p_t is common for all users, but may depend on the time-of-day. In each timeslot t, user i is assumed to attain a utility $U_{i,a,t}(q_{i,a,t})$ from consuming $q_{i,a,t}$ on appliance a. This utitlity function would fully characterize the decision making of a rational user. Nevertheless, [1] proposes a model that integrates social preferences, according to which a user's overall utility function is given by a combination of his own monetary payoff and a "disinterested social welfare function". The latter is a combination of the maximin or "Rawlsian" criterion and the total surplus maximization criterion, i.e. the sum of the monetary payoffs of all players:

$$U_i(x_1, x_2, \ldots, x_n) = (1 - \gamma_i)x_i + \gamma_i\,[\delta_i \min\{x_1, \ldots, x_n\} + (1 - \delta_i)(x_1 + \cdots + x_n)] \quad (1)$$

where $x = x_1, x_2, \ldots, x_n$ denotes an allocation of the payoffs belonging to some set X of feasible payoffs, $\delta_i \in (0,1)$ is a parameter reflecting the weight that is put on the maximin criterion and $\gamma_i \in (0,1)$ indicates how user i cares pursuing the overall social welfare rather than his own self-interest [1]. The first part of the second term of the equation represents Rawlsian inequity aversion, while the second part reflects altruism based on the idea that each user's payoff has the same weight.

In our work, users are considered to be homogeneous as regards the set of the appliances that they can possibly own but differ depending on their level of altruism. We assume that each user i can only belong to one of two categories, i.e. (i) Rational: those interested only in their own well-being and are more inflexible in tolerating any discomfort and (ii) Altruists: those who are willing to endure higher reductions in their utility (as compared to rational ones) in favor of the well-being of the users in the system. This renders them being more flexible in changes in their consumption pattern. For a rational user there is an optimal level of comfort, which is usually associated with that consumption pattern that maximizes her total net benefit; and any deviation from this pattern, leading to a lower net benefit is not acceptable. On the other hand, users characterized by some degree of altruism are distinguished by greater flexibility concerning the levels of comfort, i.e. their comfort threshold can be stretched out according to the comfort levels of other people. In essence, altruists are willing to experience some discomfort if this results in higher average utility for the rest of the users. Building on (1), we formulate the utility function for each user i as a combination of the utility stemming from the operation of her appliances $U_{i,t}(Q_{i,t})$ and possibly of the average utility of all users. That is, we assume that for all users $\delta_i = 0$. For $\gamma_i = 0$, user i is rational, caring only for his self-interest; hence, the utility function remains as defined in [1].

$$U_{i,t} = U_{i,t}(Q_{i,t}) \quad (2)$$

In general $\gamma_i \in (0,1)$ meaning that the user is interested both in her own utility as in the average utility of others, thus the utility function takes the following form:

$$U_{i,t} = (1 - \gamma_i)U_{i,t}(Q_{i,t}) + \gamma_i\,\bar{U} \quad (3)$$

The case of $\gamma_i = 1$ corresponds to a user with purely "disinterested" preferences. In fact, there are little or no purely rational or altruistic users; all users are actually featured by altruism but in a different scale. [8] states that the level of altruism is bimodal, i.e. some users express little altruistic behaviour, while others tend to decide and act with high levels of altruism. This distribution outlines user behaviour in real environments with great accuracy. We adopt this representation, i.e. the value of γ_i is drawn from the range (0.0, 0.1] and [0.7, 0.9] for rational and altruists respectively.

B. The Energy Provider

As already mentioned, we follow the modeling of the provider as in [1], i.e. the provider imposes a maximum percentage reduction in comfort, by exploiting either Policy 1 or Policy 2 presented in [1] and offers incentives to each user that are at least equal to the resulting net benefit loss. The main objectives of the provider are (i) to apply ADR in order to restrict the demand at/or below Q^s while offering the least total incentives and (ii) to leverage the existence of altruists to meet its overall objective. Therefore, whenever a DR event is necessary, the provider has to solve the optimization problem formulated in [1], i.e. the following:

$$\arg\max_{\overrightarrow{q_{\iota,a}}} \sum_{i\in N}\sum_{t\in T} p_t\widehat{Q_{\iota,t}} - C\Big(\sum_{i\in N}\widehat{Q_\iota}\Big) - \sum_{i\in N} I_i \quad (4)$$

such that
$$\sum_{i\in N}\widehat{Q_\iota} \le Q^s$$

where $\widehat{Q_\iota} = \sum_{t\in T}\widehat{Q_{\iota,t}}$. Note that $Q_{i,t}$ denotes the unconstrained total consumption of user i if no reduction is imposed, while we use "hat" to denote the new consumption schedules after the reductions in demand are enforced by the provider to users recruited for ADR. Also, I_i denotes the incentives offered to user i if participating in DR. In cases where the constraint on the total demand (5) holds with equality - which we will assume when applying our targeting policies - both the total revenue and the total cost of the provider are fixed. Therefore, as shown in [1], the optimization problem of equation (4) is equivalent to meeting the required total demand Q^s with the least total incentives, i.e.:

$$\arg\min_{I_i} \sum_{i\in N} I_i \quad (5)$$

Note that the incentives are already prescribed in the contracts, and in essence they should be defined such that users accepting the contracts are adequately compensated whenever targeted for DR. Therefore, the above problem amounts to targeting those users needing the least total incentives to meet provider's threshold. In fact, the provider should verify that the cost savings attained due to DR exceed the losses due to incentives and selling of less energy. To benefit from the existence of altruists, the provider can employ different demand reduction and targeting policies. In this way, the provider can meet her objectives, while impacting the social welfare of the system either positively or negatively. It should be emphasized that the maximisation of the social welfare is not part of this optimisation problem; although our work can be applied to address it. The demand reduction and targeting policies implemented by the provider besides the total incentives can also affect it.

C. Incentives for DR

In this section, we specify the DR incentives to be included as term in the contracts of the users engaged in DR. We adopt the assumptions and the approach described in [1], according to which the incentives offered to users should equal their total net benefit loss due to the reduction in convenience. More specifically, by assumption users are characterized by an optimal consumption per day, which they choose for themselves as a result of maximizing their net benefit NB_i. Hence, any deviation from this consumption is bound to reduce the net benefit obtained by the user [1]. In this direction, in the case of DR, a user is mandated to consume less than his optimal daily quantity, thus resulting in a lower net benefit, as well as in lower utility and charge values. To this end, in order for such a user to be convinced to participate in an ADR program, the provider should offer monetary incentives that correspond to its net benefit loss.

$$I_i = NB_i - \widehat{NB_\iota} \quad (6)$$

Note that $\widehat{NB_\iota}$ (where $\widehat{NB_\iota} < NB_i$) expresses the net benefit obtained by user i under the new consumption schedule, that is imposed after the reduction in demand by the provider. However, this can be interpreted in a different manner for each type of users. In particular, if users are only self-interested, by receiving a compensation they would become indifferent to the changes in their consumption; thus the incentives to be offered to each user account for the same value as in (6). On the contrary, if users are characterized by a flavor of altruism, they would be reimbursed by incentives that are smaller than their net benefit loss, since their valuation of utility depends on the average utility of others as per (3), i.e. the smaller the average utility the less incentives offered to the altruists. In essence, altruists are "punished" for any negative impact on the average utility of the society. Therefore, in order for ADR to be attractive to users, the contract terms should specify beforehand the incentives' calculation methodology, and particularly that whenever a DR event is issued the user will be compensated the amount equal to the loss of net benefit caused as per (6). Hence, whenever it is predicted that the total demand will exceed the threshold Q^s, the provider activates the ADR programs to restrict total demand at or below Q^s, while also abiding with the terms of the signed contracts. To do so, the provider can employ one of the demand reduction policies either coupled with a targeting policy or not. The decision is based on whether it is beneficial, in terms of incentives and the social welfare accomplished after the ADR, to impose certain reductions to all users without selecting only a few of them, and if saddling altruists with very high reductions, thus easing rational users, further improves the total incentives offered and the social welfare achieved after the ADR.

D. A Simple Model

Intuitively, the existence of altruists in the population tends to yield positively in the welfare of the system. This allows for their optimal treatment, so that users that are more sensitive to changes can be mitigated in terms of the reduction in their consumption. To attest the validity of such an argument, we present in this section a simple model. We take N users denoted as $R, A, 3.., N$, where we emphasize that the first user belongs to the category of rational users and her utility is modelled based on (2), while the second one belongs to altruists with utility expressed by (3). These users are assumed to have identical daily consumption (Q_t) and utility obtained from it $\big(U_t(Q_t)\big)$. Hence the average utility of the system is:

$$\overline{U} = \frac{2U_t(Q_t) + \sum_{i=3}^{N} U_{i,t}\big(Q_{i,t}\big)}{N} \quad (7)$$

The provider wishes to restrict the total demand and applies ADR by imposing a particular reduction δ to only one user each time. If only the rational user R is targeted for ADR, then her new utility is expressed by

$$\widehat{U_{R,t}} = U_t(Q_t - \delta) \quad (8)$$

and due to (6) the corresponding incentives are estimated as:

$$I_{R,t} = U_t(Q_t) - U_t(Q_t - \delta) - \delta p_t \quad (9)$$

As this is the only user participating in the ADR the total incentives are equal to those awarded to the rational user in Eq. (9). In the case where only the altruist A is targeted for ADR, her utility is expressed by the following form:

$$\widehat{U_{A,t}} = (1-\gamma_A)U_t(Q_t - \delta) + \gamma_A\,\widehat{\overline{U}}, \quad (10)$$

where

$$\widehat{\widehat{U}} = \frac{U_t(Q_t - \delta) + U_t(Q_t) + \sum_{i=3}^{N} U_{i,t}(Q_{i,t})}{N}. \quad (11)$$

Recall that since one user (namely A) is targeted, for all other users, including rational R, the consumption remains unchanged. The resulting incentives offered now to user A are:

$$I_{A,t} = (1-\gamma_A)(U_t(Q_t) - U_t(Q_t - \delta)) + \gamma_A(\overline{U} - \widehat{\widehat{U}}) - \delta p_t \quad (12)$$

Therefore, the total incentives to be offered by the provider are expressed by (12), which based on Eq. (7) and (11) becomes:

$$I_t = (1 - \gamma_A + \frac{\gamma_A}{N})\,(U_t(Q_t) - U_t(Q_t - \delta)) - \delta p_t \quad (13)$$

The expression in the first bracket has crucial impact in the total incentives as it shows that the greater the sense of altruism the lower the impact of the loss in the utility obtained from the consumption; hence leading to lower incentives than in the previous case as shown in (9). Eq. (11) and (13) clearly demonstrate the effect of parameter γ_i in the estimation of incentives and the social welfare of the system i.e. that in a society comprising homogeneous users, burdening only rational users with the total required reduction in consumption leads to higher total incentives and social welfare loss compared to the case of targeting only altruists for ADR. This motivates for employment of various demand reduction and targeting policies, in order to make the most of altruists in favor of the society.

E. Targeting Policies

1) Optimal Targeting

Our findings infer that γ_i plays crucial role in the assessment of both the overall incentives offered by the provider and the social welfare achieved. In fact, saddling altruists does not guarantee the lowest total incentives for the provider, as the impact of a unit of demand reduction to the discomfort of these users may be high. An optimal combination of a demand reduction and a targeting policy is rather bound to serve the provider's objectives more effectively. Whenever it is predicted that the total demand will exceed Q^s by an excess demand Q^e, the provider selects the reduction δ_i for each user i, so as to reduce demand up to Q^s, while offering the minimum possible incentives, i.e. the provider has to solve the following problem:

$$\max_{\delta_i} \sum_{i=1}^{N} \left(\delta_i p_t + (1-\gamma_i)\left(U_{i,t}(Q_{i,t} - \delta_i) - U_{i,t}(Q_{i,t})\right) + \gamma_i \frac{\sum_{j=1}^{N} U_{j,t}(Q_{j,t} - \delta_j) - U_{j,t}(Q_{j,t})}{N} \right) \quad (15)$$

such that $0 \le \delta_i \le \eta_{max_q} Q_{i,t}$

$\sum_{i=1}^{N} \delta_i = Q^e$

where δ_i is the reduction induced to each user i. Notice that the provider maximizes the negative total incentives to users for a consumption reduction δ_i from each user i. Assuming that the necessary and sufficient Karush-Kuhn-Tucker (KKT) conditions for optimality hold (i.e. the negative total incentives function is concave, the inequality constraints are continuously differentiable convex functions and the equality constraints are affine functions), then the problem can be optimally solved by the method of KKT multipliers. The partial derivatives of the KKT function with respect to the problem variables and the KKT multipliers are taken equal to 0 and the solution to this system of equations solves the original optimization problem. However, the system of equations corresponding to the KKT conditions is usually not solved directly, except in special cases where a closed-form solution can be derived analytically. Next, we propose some approximation algorithms to this problem.

2) Heuristic Policies for Targeting

Here, we devise a user selection policy associated with each of two policies described in [1] to restrict total demand at or below Q^s, while also abiding with the terms of the signed contracts. We should denote that for the sections to follow we refer to the selection algorithm introduced in [1] as NBIADR. The provider can opt to apply one of the following targeting policies, i.e. (i) **NwR**: normal application of NBIADR, where all users are considered as rational and are imposed the same proportional reduction of consumption and comfort, (ii) **NwA**: normal application of NBIADR with altruists included in the population, which are targeted in the same way as defined by the NBIADR algorithm, i.e. according to a unified ranking containing all users. Users are either imposed the same proportion of reduction in consumption and comfort (NwA) or altruists are imposed a higher proportion than rational ones (NwAsort/targ) and (iii) **PwA**: preselected altruistic based application of NBIADR, where the population consists of both rational and altruistic users, and altruists are targeted first and are imposed a higher proportion of consumption and comfort reduction than rational ones

All these policies utilise the NBIADR algorithm of [1] and can be implemented jointly with both policies of constraining the reduction in either the consumption or the utility gained by each user and by applying various demand reduction strategies. According to NBIADR the provider's objective is to limit the amount of incentives needed to reduce the total demand below or at Q^s. One important issue that arises in the case of a population mix with both rational and altruistic users concerns the estimation of $\overline{U}$. According to the original definition of NBIADR, users are being sorted based on their valuation of the incentive per unit of reduction in their consumption. In our context and under any of the targeting policies defined above, the calculation of incentives entails $\widehat{\widehat{U}}$, which in turn depends on the new consumption schedules imposed after the reduction, meaning that the estimation of $\widehat{\widehat{U}}$ is approximate, as finally not all users are targeted for ADR. This adds inaccuracy in the solution of the optimization problem expressed by Eq. (4) and (5), rendering our targeting policies suboptimal.

IV. Experimental Evaluation

To experimentally evaluate our approaches, we use real consumption data from 6 households in Lulea, Sweden. The data constitutes of sensor readings for only one appliance, i.e. space heating, at a granularity of one hour. In order to expand our dataset to target more households (20 in total), we created synthetic data based on the real data; in essence, we created groups of households with similar physical characteristics and consumption patterns. Recall that a household is taken to correspond to a single user, although actually multiple users can reside in it and that, throughout this section, we utilize the notions of user and user alternately in order to refer to a single user. The readings are extracted for a given context, that is weekday in May 2015, and for each day and user, the recorded data is used to obtain consumption in Wh for each time slot with a duration of 6 hours during the day. So, we obtain the optimal consumption Q_i of user i and the utility function of (2) takes the form:

$$U_i(T_i(t)) \coloneqq C_i - b_i(T_i(t) - T_i^{com})^2$$

Basically, the utility obtained is modelled as the distance of the current temperature ($T_i(t)$) from the temperature that is comfortable for the user (T_i^{com}) and is borrowed from [9]. The current temperature depends on the current power draw as well as the temperature in the previous timeslot. The parameters C_i and b_i are positive constants and can be inferred by the methodology described in [10]. For each household, we assume that the comfortable temperature range is [19C, 21C]; the household cares about the inside temperature through the whole day.

Say that the energy provider wishes to narrow the total demand by a value ΔQ that amounts to 10% of the unconstrained total optimal consumption and define the value of η_{max_q}, by constraining the reduction in consumption, i.e. applying Policy 1. Following the feasibility condition of (9) in [1], the value of η_{max_q} must be greater or equal to 10%. All users are characterized by some level of altruism, which is assumed to be known to the provider and is reflected by the parameter γ_i. Rational users tend to invest very little in altruism in comparison to altruists; hence for the former the value of γ_i is taken randomly within the range of $(0, 0.1]$, whereas the latter fall within [0.7, 0.9]. For both user types, the value of γ_i is considered in the calculation of incentives. Note, here, that both the total amount of incentives (I) to be offered and the social welfare achieved after DR ($\widehat{SW}$) are expressed in Swedish Crone (SEK), while the reductions imposed are notated as [% reduction to rational, % reduction to altruistic], e.g. [10,10].

A. ADR without Targeting Policies

This case can be viewed as a variation of NwA algorithm, wherein the provider is only interested to apply ADR without performing any targeting and hence all users are imposed a percentage reduction in their consumption. Initially, the provider imposes the same reduction in all users [10,10]. Then, the provider applies lower reduction to rational users to be favored over altruists. The reduction imposed on altruists is calculated by the algorithm, so that the total reduction obtained covers the remaining amount of reduction required to meet provider's threshold. Normally the exploitation of altruists would reduce the total incentives to be offered, as inferred by Eq. (14); surprisingly, though, this is not the case. Table 1 shows that saddling altruists with high reductions does not limit the total incentives offered. This is due to the fact that some users, even if rendered with high sense of altruism, may value a unit of reduction rather significant for their comfort, thus experiencing great loss by any modification in their consumption pattern and in turn affecting the total average utility of the society in a negative manner. Note that this fact is not captured in the simple model of Section 3.4, because the rational and the altruistic users compared were assumed to have identical daily consumption and utility. On the contrary, altruists and their proper exploitation appear in our numerical results to have a sound effect on the social welfare after the ADR. Overall, these outcomes highlight the importance of altruists in the society and motivates the implementation of demand reduction and targeting policies by the energy provider taking into account altruism.

B. ADR with various Targeting Policies

Intrigued by the results of the previous case, the provider wishes to apply ADR by targeting a set of users to impose reduction to, while offering the least total incentives yet without suffering significant social welfare loss. To achieve that, we propose to couple ADR with various demand and targeting policies, while leveraging the existence of altruists. The results are rewarding; more specifically, considering altruists during the sorting and targeting process (NwA) results in targeting less users (15 users are targeted, where 6 users are altruists) with quite high levels of social welfare (3721.9SEK) and low incentives (718.1SEK) (Figure 1).

TABLE 1. APPLYING ADR WITHOUT PERFORMING TARGETING

Applying ADR without Targeting		
% Reduction [R,A]	$\boldsymbol{I}$(SEK)	$\boldsymbol{SW_{ADR}}$(SEK)
[10,10]	748,7	5485,7
[8,12.02]	749,3	5491,1
[5,13.6]	750	5498,9
[3,18.4]	750,5	5503,9
[0,21.9]	751,5	5511,1

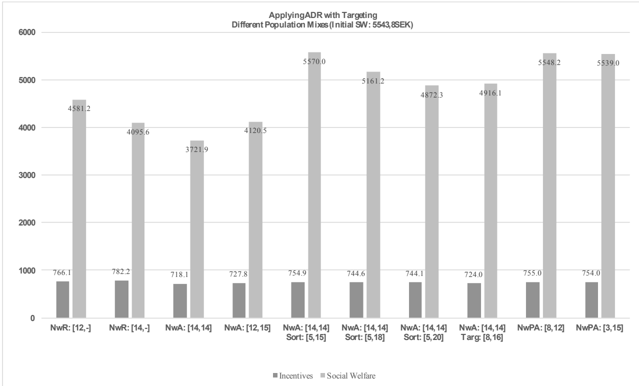

Figure 1. Results for different targeting policies

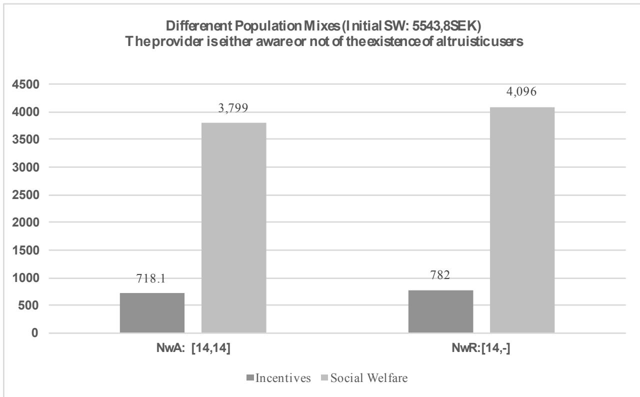

Figure 2. Applying different targeting policies based on provider's knowledge about the existence of altruists

On the other hand, applying different percentage reductions for rational and altruistic users ([12,15]) and performing the selection algorithm as normally can lead to increased social welfare but with also slightly higher incentives compared to NwA (see Figure 1). The targeting set consists of 15 users, where 8 users are altruists. Exploiting the altruists in favor of the rational users by imposing lower reduction in the consumption of the latter may lead to better results in social welfare after the ADR. This can be achieved in two ways, i.e. via (i) NwAsort: modifying the percentage of reduction during the sorting procedure of users and (ii) NwAtarg: by performing the targeting process normally; after identifying the optimal set of users to target, by imposing different reductions to the targeted users and performing the targeting process in this subset again, thus resulting in an updated targeting subset.

Note that the reductions per category of users are set in such a way that the total reduction required is met.

According to NwAsort, the effect of altruists is straightforward pertaining to the total incentives offered, i.e. for different combinations of percentages of reduction in consumption of altruists the total incentives decrease (See Figure 1, NwA:[14,14], Sort: [5,15],[5,18],[5,20]). This is not the case, though, for the social welfare after ADR; it is not monotonically increasing as a gradually higher reduction in consumption is imposed to altruistic than to rational users. On the contrary, increasing the consumption reduction imposed to altruists leads to reduced social welfare and total incentives to be offered. Thus, saddling altruists with very high reductions, while favoring the rational users with small reductions, may prove economically beneficial for the provider but inefficient for the social welfare of the system. Although it seems as a more complex procedure, the results for NwAtarg can be compensatory - Figure 1 shows that 18 users are finally targeted, 8 of them being altruists, with low total incentives and high social welfare after ADR. Thus, under certain circumstances, a provider can leverage the presence of altruists by imposing to them a higher reduction in their consumption to achieve its goal with low total incentives, while reckoning the tradeoffs stemming from the non-monotonicity of the social welfare in γ_i.

Finally, the application of PwA with various combinations of percentage reduction for users, both after the sorting and targeting procedure, reveals that a smaller reduction to altruists leads to targeting more rational users (19 users are targeted – 10 of them are altruists) (Figure 1). The fact that the preselection of altruists leads to higher social welfare after ADR with marginally higher total incentives than under the previous policies is a noteworthy benefit for the provider. All the aforementioned approaches require detailed knowledge by the provider in relation to special demographic characteristics of the users, e.g. level of altruism, etc. Nevertheless, the acquisition of such knowledge not only raises privacy concerns but even if it is feasible it might require a learning period of users' behavior.

In this context, PwA outperforms the other approaches in terms of simplicity of implementation, the total incentives to be offered and the social welfare achieved after ADR. As a result, it is beneficial for the provider to identify the set of altruists and the related values of parameter γ_i, in order to exploit the benefits of the aforementioned policies as well as assess the trade-offs regarding the knowledge acquisition. Moreover, if the provider wishes to influence the value of γ_i, e.g. by means of careful messaging, environmental campaigns, etc., it should be performed carefully as the results may be inconsistent with the objectives of the system. To attest the validity of the above statement, we next look into the case where the provider is not aware of the existence of altruists in the population mix; thus, assuming all participants to behave rationally, the provider applies ADR by targeting users in a normal fashion via the NwR algorithm. In this way, users are compensated by incentives that exceed their net benefit loss and also comply with the contract terms with regard to the maximum percentage reduction, so that they are accepted by users. As expected, the total incentives to be offered increase in the case of any knowledge. Despite the fact that the social welfare after ADR increases, the difference is not large compared to applying NwA. The same does not hold, though, for the total incentives offered, which are provider's main objective and are increased in the case of no information (NwR). To conclude, our outcomes reinforce the argument on the benefits of the awareness of altruists in the society and the related values of parameter γ_i accordingly.

V. Conclusions and Future Work

In this paper, we explored the impact of altruism to the energy consumption behaviour of residential users within ADR programs. Our experimental findings with real and synthetic data clearly show that the employment of targeting policies appropriately coupled with demand reduction strategies are able to utilize altruism in favor of the community as a whole and the providers' objectives. However, we found that leveraging of altruists should be performed carefully, since saddling them with very high reductions of power, although yielding lower total incentives to be given to users, can yet prove inefficient for the social welfare of the system. Preselecting altruists improves the social welfare after ADR, but with higher total incentives, as compared to the other policies considered. Nonetheless, the selection of the optimal targeting approach is a decision to be made by the provider taking into consideration the information available with regard to the utility functions and the demographic characteristics, the trade-offs between the total incentives to be offered and the social welfare achieved. It is clear, though, that it is beneficial for the provider to identify the set of altruists and their levels of altruism, in order to exploit the outcomes of the demand reduction strategies and targeting policies effectively to meet its goals, ensuring at the same time the system welfare. As a future work, we aim to extend our model to accommodate feedback of users on their net satisfaction from comfort losses due to a specific reduction from their baseline consumption and its associated compensation in terms of incentives.

Acknowledgment

This research has been partly funded by EU's Horizon 2020 Programme under grant agreement no 649796 OPTi.

References

[1] Charness, G. and Rabin, M., 2002. Understanding social preferences with simple tests. Quarterly journal of Economics, pp.817-869.

[2] Minou, M., Stamoulis, G.D., Thanos, G. and Chandan, V., 2015, November. Incentives and targeting policies for automated demand response contracts. In *IEEE SmartGridComm 2015* (pp. 557-562).

[3] Frederiks, E.R., Stenner, K. and Hobman, E.V., 2015. Household energy use: Applying behavioural economics to understand user decision-making and behaviour. Renewable and Sustainable Energy Reviews, 41, pp.1385-1394.

[4] Allcott, H. and Rogers, T., 2014. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *The American Economic Review*, *104*(10), pp.3003-3037.

[5] Faruqui, A., 2010. The ethics of dynamic pricing. *The Electricity Journal*, *23*(6), pp.13-27.

[6] Deci, E.L., Koestner, R. and Ryan, R.M., 1999. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological bulletin*, *125*(6), p.627.

[7] Carbajal, J.C. and Ely, J.C., 2016. A model of price discrimination under loss aversion and state-contingent reference points. *Theoretical Economics*, *11*(2), pp.455-48

[8] Stenner, K., Frederiks, E.R., Hobman, E.V. and Cook, S., 2017. Willingness to participate in direct load control: The role of user distrust. *Applied Energy*, *189*, pp.76-88.

[9] Koella, J.C., 2000. The spatial spread of altruism versus the evolutionary response of egoists. *Proceedings of the Royal Society of London B: Biological Sciences*, *267*(1456)

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

[11] Chandan, V., Ganu, T., Wijaya, T.K., Minou, M., Stamoulis, G., Thanos, G. and Seetharam, D.P., 2014, June. idr: User and grid friendly demand response system. In *Proc. of the 5th international conference on Future energy systems* (pp. 183-194).