ORIGINAL RESEARCH

Prediction and assessment of demand response potential with coupon incentives in highly renewable power systems

Hao Ming[1], Bainan Xia[2], Ki-Yeob Lee[3], Adekunle Adepoju[3], Srinivas Shakkottai[1] and Le Xie[3*]

Abstract

Demand Response (DR) provides both operational and financial benefits to a variety of stakeholders in the power system. For example, in the deregulated market operated by the Electric Reliability Council of Texas (ERCOT), load serving entities (LSEs) usually purchase electricity from the wholesale market (either in day-ahead or real-time market) and sign fixed retail price contracts with their end-consumers. Therefore, incentivizing end-consumers' load shift from peak to off-peak hours could benefit the LSE in terms of reducing its purchase of electricity under high prices from the real-time market. As the first-of-its-kind implementation of Coupon Incentive-based Demand Response (CIDR), the **EnergyCoupon** project provides end-consumers with dynamic time-of-use DR event announcements, individualized load reduction targets with EnergyCoupons as the incentive for meeting these targets, as well as periodic lotteries using these coupons as lottery tickets for winning dollar-value gifts. A number of methodologies are developed for this special type of DR program including price/baseline prediction, individualized target setting and a lottery mechanism. This paper summarizes the methodologies, design, critical findings, as well as the potential generalization of such an experiment. Comparison of the EnergyCoupon with a conventional Time-of-Use (TOU) price-based DR program is also conducted. Experimental results in the year 2017 show that by combining dynamic coupon offers with periodic lotteries, the effective cost for demand response providers in EnergyCoupon can be substantially reduced, while achieving a similar level of demand reduction as conventional DR programs.

Keywords: EnergyCoupon, Incentive-based demand response, Critical peak pricing, Electricity market, Prospect theory

1 Introduction

During the past decade, there has been an increasing penetration of renewable energy resources (such as wind and solar generation) in the power grid. For instance, the Electric Reliability Council of Texas (ERCOT) wind and solar generation has more than doubled in their fuel mix in the past decade, from 7.5% in 2008 to 18.6% in 2018 [1]. On the other hand, demand response (DR) has been identified as having the potential to become a flexible resource to solve the reliability and efficiency issue of the power grid incurred by renewable penetration [2]. Demand response is defined as "the changes of end-consumers' electricity consumption in peak hours from their normal patterns" [3]. Many independent system operators in the U.S. including the Electric Reliability Council of Texas (ERCOT), New York ISO (NYISO), California ISO (CAISO) and ISO New England, already have a number of ongoing day-ahead and real-time DR programs in their operating areas for providing energy reserve and auxiliary services [4–6].

The story of demand response in the U. S begins in the 1970s, growing with the popularity of household air

* Correspondence: le.xie@tamu.edu
[3]Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, USA
Full list of author information is available at the end of the article

conditioning [7]. A large number of DR programs have been designed and implemented since then. With development over almost 40 years, it is generally accepted that DR programs can be categorized in two dimensions by: 1) the subject who takes control over devices (direct load control vs. self-controlled and market-based programs), and 2) the scale of target end-consumers (large industrial/commercial customers vs. small residential customers). The term "direct load control" indicates that the DR operator (such as the utility) can remotely turn on/off or modify the setpoint of customers' equipment. The amount of load shedding can be precisely controlled at the expense of customers' comfort and satisfaction (for instance, an air conditioner might be turned off for some hours during a hot summer day). In contrast, market-based DR programs tend to use price signals or other incentives to encourage customers' self-motivated load-control behaviors. Such programs usually have less impact on customer comfort and satisfaction, but are less precise and effective when a specified target of demand reduction goal needs to be achieved.

Because of their profit-seeking characteristics and higher electricity usage than small residential customers, industrial and commercial customers usually have more self-motivation and better performance than small residential customers in participating in DR programs. Energy management systems have been developed to help increase energy efficiency in data centers, retail stores, telecoms providers etc., and to coordinate with market-based signals (such as real-time electricity and gas prices) [12]. On the other hand, residential customers are often more concerned with their personal comfort. Their acceptance of price-based mechanisms (such as time-of-usage (TOU), critical peak pricing (CPP) [14], and market-index retail plans offered by the utility) still remains at a low level, with the majority of residential end-consumers choosing fixed-rate electricity retail plans. Given the fact that residential electricity consumption leads electricity usage in the U.S (38%, as compared with commercial at 37% and industry at 25%) [12], the potential of residential DR is far from fully explored. Table 1 summarizes some recent research and operational programs using different approaches to demand response.

Table 1 Classification of demand response programs

Customer type	Direct load control (centralized)	Market-based
Large commercial and industrial	Research papers [8–10]	Energy-management systems [12]
Small residential	Direct load control programs [11]	Variable rate retail plans [13], CPP [14], EnergyCoupon [15]

There has been some academic research [14, 16] and commercial implementation (e.g. ENERNOC [17], Ohm-Connect [18]) of market-based DR; however, as an alternative approach to current existing market-based solutions, the efficiency gain of Coupon Incentive-based Demand Response (CIDR) is still underexplored. CIDR aims at providing coupon-based incentives to reduce the electricity consumption of residential end-consumers during peak hours [19–21]. Compared to traditional DR programs, this mechanism has the following advantages: it is purely voluntary, penalty-free to customers, and compatible with the fixed-rate electricity retail plans which are most popular among residential end-consumers. A program named ***EnergyCoupon*** is the first-of-its-kind implementation of CIDR, with further additional innovations such as: 1) dynamic DR events to end-consumers with *individualized reduction targets*; and 2) *periodic lotteries* designed to convert coupons earned in DR events into dollar-value prizes. A small-scale pilot experiment was conducted in 2016, with substantial load profile changes of the residential participants in the posterior analysis [15].

In terms of 2) the *periodic lottery*, many academic, as well as commercial, studies have shown how "nudge engines", such as games and lotteries could help to encourage the desired behavior of human. Reference [22] tries to discover the social value of energy saving, [21] models the CIDR system as a two-stage Stackelberg game, and [23–25] use the "mean field games" framework to describe end-consumers' behaviors in the DR program with lottery-based incentives. Furthermore, the lottery-based incentive scheme has already been implemented in a platform that aims at encouraging uniform temporal demand on public transportation [26], and relieving congested roadways [27]. However, apart from some ongoing experiments [28, 29], there is not much work trying to adapt the lottery idea to electricity DR programs.

Built upon our previous studies in 2016, a larger-scale experiment was conducted in 2017, with much more comprehensive designs and critical assessment.[1] The improvements of experiment ('17) include but are not limited to: 1) an extra comparison group for data analysis; 2) an improved baseline prediction algorithm (named as the "similar day" algorithm); 3) two subgroups divided from the treatment group facing fixed and dynamic DR events separately. More facts and comparisons between two experiments are listed in Table 2. We will show in later sections that these changes help to analyze end-consumers' behaviors in-depth.

The main contributions of the EnergyCoupon program are as follows:

---

[1]Unless otherwise specified, in the remaining part of this paper, "experiment ('16)" refers to our previous study conducted in the year 2016 and "experiment ('17)" refers to the new one in 2017.

**Table 2** Overview of EnergyCoupon experiments in year 2016 and 2017

Year	2016	2017
Experiment length (weeks)	12	12
Treatment group size	8	29
Comparison group existence	No	Yes
Baseline algorithm	"Hybrid"	"Similar day"
Active subjects defined by	Energy saving	Lottery participation
Number of active subjects	3	7

1. Providing price and baseline prediction algorithms suitable for DR programs;
2. Systematically documenting the experimental design, data collection, and posterior analysis for the selected residential customers;
3. Experimental result showing load shedding/shifting effects, different behaviors over fixed/dynamic coupon targets, financial benefits of the LSE and end-consumers, impact of periodic lotteries on human behaviors, as well as the effective cost saving of *EnergyCoupon* over traditional DR programs.

This paper is organized as follows: Section 2 introduces the system architecture and the interface of the EnergyCoupon App. Key algorithms including price prediction, baseline prediction, individualized target setting and periodic lottery are explained in Section 3. Experimental design is described in Section 4, and data analysis is shown in Section 5. We finally conclude our findings in Section 6.

2 System overview

The EnergyCoupon system is designed to inform end-consumers of an upcoming DR event along with individualized targets, measure the demand reduction within the DR event, provide statistics and tips for energy saving, as well as conduct periodic lotteries. Figure 1 exhibits the system architecture of EnergyCoupon. As the core component in the architecture, an SQL database is hosted on
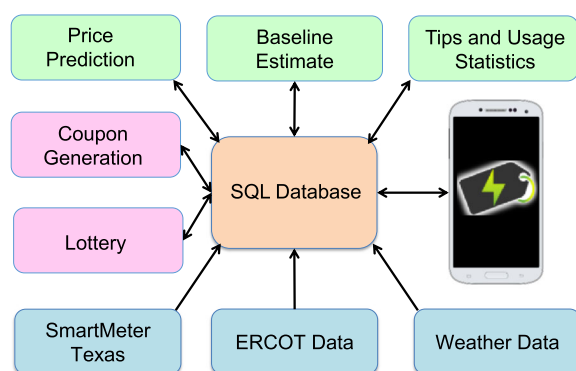
a server running 24/7, interacting with the data resources (shown in blue blocks), mathematical algorithms (green blocks), and the lottery scheme (pink blocks). The EnergyCoupon App (both Android/IOS versions available) is developed and installed in the mobile phones of the treatment group. The app (interface shown in Fig. 2) receives and shows coupon targets, tips and statistics from the server, and also enables the app user to participate in periodic lotteries. A brief overview of all the other crucial components in Fig. 2 is as follows:

1) *SmartMeterTexas*: This is the source of the electricity consumption of all end-consumers at 15-min resolution [30]. In our study, we received this information each day from a collaborating retail provider. Data is used in both baseline prediction and coupon target generation algorithms, which will be introduced in Section 3.2 and 3.3.

2) *ERCOT Data*: This is the source of day-ahead and real-time market prices, as well as the system load in the ERCOT area [1]. Pricing data is used in the price prediction algorithm described in Section 3.1.

3) *Weather Data*: This is the source of weather information used in the price (Section 3.1) and baseline prediction algorithms (Section 3.2). Weather information was pulled from the website of *Weather Underground* (a commercial weather service provider) [31].

4) *Price Prediction*: This is an algorithm whose purpose is to predict in advance whether a dynamic DR event should be announced. Our goal is to ensure that this can be done with a lead time of at

Fig. 1 System architecture

Fig. 2 EnergyCoupon app interface (The EnergyCoupon app is still available on Play Store (Andriod) when this paper is published (2020). Upon agreeing to the terms and conditions, Android mobile users in Texas, US are able to explore some of the functions (such as monitoring real-time energy consumption and submit coupon in lottery) although no EnergyCoupon event and lottery is held right now.). **a** Main page, coupon targets and tips. **b** Usage statistics. **c** Lottery interface

least 2 h in advance of the event, so as to provide the participants enough time to respond to DR events. This algorithm is introduced in detail in Section 3.1.

5) *Baseline Estimate*: This is an algorithm whose purpose is to predict the "normal consumption" of the end-consumer without considering the impact of DR. This algorithm is designed to eliminate the gaming effects described in [32], and tries to balance accuracy of the predication and its computational cost. Details are included in Section 3.2.

6) *Tips and Usage Statistics*: The following types of usage statistics and personalized tips are randomly shown on the user app interface such as: *(a) High price alert* based on the price prediction algorithm (Section 3.1) for the upcoming hours. *(b) Coupons* acquired every day and total coupons acquired last week. *(c) Energy consumption* for the users in the past week, and estimated electricity bill based on retail price. *(d) Gold, silver or bronze Medal* as an indicator of a user's saving behavior in the past week compared with other participants. All the above statistics as well as a *figure* showing the detailed energy consumption curve were included in an email and sent to the user every week, which further helps the user to better engage in the demand response program.

7) *Coupon Generation*: DR events are determined according to price prediction, and personalized targets are generated based on the user's predicted baseline at the time interval when a DR event is triggered. See Section 3.3 for details.

8) *Lottery*: Periodic lotteries enable the end-consumers to convert his/her coupons earned into dollar-value gifts. See Section 3.4 for more details.

3 Methods
In this Section, we elaborate on the key analytics behind the experiment ('17). The methodologies introduced include price prediction, baseline estimate, coupon generation, and lottery. These analytics are important not only for this experiment, but also for designing other possible demand response mechanisms.

3.1 Price prediction
In demand response, end-consumers are incentivized to perform a load shedding or load shift from peak hours to off-peak hours (measured by wholesale electricity price). In order to run our EnergyCoupon system in real-time, we must be capable of predicting the high price occurrences ahead of time. A lot of research has been carried out on the topic on electricity price prediction. For example, time series models have been used to predict day-ahead electricity prices in [33, 34]. A combination of wavelet transform and an ARIMA model is used in this context in [35]. A hybrid solution method using both time series and a neural network is presented in [36]. In [37], spot price prediction is discussed, when both load prediction and wind power generation are involved.

However, our goals for price prediction are to some extent different from previous work. Since our question is whether or not to trigger the DR event for potential peak prices, the precise prediction of the market price will be less important; instead we only want to predict if the 30-min average wholesale market price 2 hours later is likely to be higher than a certain threshold (In Energy-Coupon "high price" is defined as greater than or equal to $50 per MWh). Furthermore, time series techniques show good performance in handling data with repeating periods, such as 24-hour period, and achieve high accuracy in predicting the following successive samples. While the high prices that we target in our scenario have some relation to time of day (typically, the late afternoon), they do not have a precise correlation at a 24-hour period, and are more related to events of that day (such as the ambient temperature). Last but not least, for the app such as EnergyCoupon, an online algorithm with low computational complexity is preferred. Accounting for all these concerns, we design and deploy a customized decision tree to deal with price prediction in our system.

The decision tree is a well-known classifier, with selected features in non-leaf nodes and labels in leaf nodes. An advantage of the decision tree is the fact that it allows for easy interpretability, which enables one to identify which features are most relevant and why. Different from the traditional approach, we have unbalanced error concerns in our EnergyCoupon system, since a false high price alert which might trigger more DR events will not induce much loss to the EnergyCoupon program because of the fixed budget for weekly lottery prizes (while the coupons issued during the DR event might be slightly depreciated). However, a failure to catch an actual high market price may have a more significant opportunity cost for a potential saving in demand response. Hence, our decision tree should have higher tolerance to false positive errors than false negatives. This requirement can be captured by adjusting the penalty ratio between two kinds of errors in the training stage, though one must be careful while doing so because of the risk of overfitting the training set. An exhaustive search was conducted in the plane of two parameters, minimum leaf size and penalty ratio between the two types of errors, to address this trade-off and set the values at 70 and 1:8 respectively. Details are presented in [15] and omitted here given the focus of this paper.

Considering the DR procedure conducted in our system, we believe that a 2-hour-advance notification is a reasonable time window for participants to react. Given this goal, we need to select features for our classifier from a large body of data and possible features. Since weather determines air-conditioning usage that dominates household electricity consumption in Texas, and it also has an crucial impact on renewable energy availability, five fundamental feature classes are chosen: Price(π), Demand(P), Temperature(T), Humidity(H) and Wind Speed(W). Furthermore, we choose the temporal offsets in each feature class according to the self and cross-correlation between the feature and the price label. In addition, a numerical study was carried out to choose a proper threshold for our field experiments, so as to label data (price) samples. Table 2 in reference [15] shows a prediction accuracy of over 90% in the validation data set. Full details on training data preparation, feature selection and performance evaluation are beyond the scope of this paper. Readers may refer to [15] for more information.

3.2 Baseline estimate

As defined by the U.S. Department of Energy, the *baseline* is the "normal consumption pattern" by end-consumers without the impact of DR [3]. A daily baseline prediction algorithm is of crucial importance to our EnergyCoupon program, since it affects the energy reduction measurement, as well as the number of coupons the participant earns during a DR event. Energy reduction for an end-consumer i on interval k in a particular day D, $P_{DR,i}^D(k)$ is calculated as the difference between the consumer's predicted baseline $P_{base,i}^D(k)$ and his/her real electricity consumption $P_{real,i}^D(k)$ (as shown in eq. (1)); $P_{real,i}^D(k)$ can be measured by the smart meter installed in the his/her household with high reliability.

$$P_{DR,i}^D(k) = P_{base,i}^D(k) - P_{real,i}^D(k). \tag{1}$$

There are two major concerns within the design of a baseline algorithm, namely (i) *Baseline Manipulation*: the end-customer may intentionally increase thier usage during certain periods of time in advance in order to fabricate a reduced load appearance during a DR event, and (ii) *User's Dilemma*: if targets are set for the end-customer with a baseline that depends on a short window of time in the past (such as the previous few days/weeks), the baseline for a responsive user will continuously reduce in the future, resulting in potentially unattainable reduction targets with the progress of experiment. There exist several works [15, 32, 38] that discuss these issues that pertain to conventional baseline estimate algorithms widely used by some major independent system operators (ISOs) in the U.S. [4, 5].

As one candidate solution to these concerns, the "hybrid" method adopted in our previous experiment ('16) computes a *weighted average* of the consumer's own recent consumption and the whole group's consumption [15]. However, in post-experiment analysis we discovered that this algorithm neither (i) eliminated gaming effects, nor (ii) provided good baseline prediction because of the large diversity among residential end-consumers. However, we could not address them during the experiment in 2016, and had conjectured that a "similar day" algorithm might be a better solution [15].

The proposed "similar day" algorithm derives from the k-nearest neighbors algorithm (k-NN) and kernel regression [39, 40]. The main idea behind the algorithm is to build up a statistical model of a particular home by using a consumption data set of that particular end-consumer for a year in advance of the experiment. Since we empirically observe that the feature that best correlates with energy usage is the ambient temperature, the algorithm focuses on finding a window of temperature that has a close fit with the temperature profile of the target time window to predict in the following manner.

For instance, to predict a given 6-hour time window of a certain end-consumer in the future, the "similar day" algorithm first obtains the historical consumption for 1 year of the same user **before the experiment begins**. Then candidate "similar" time windows are selected based on the following criteria:

1) Selected time window(s) should have the same length (6-hour) and time-of-day with the predicted time period. Weekday/weekend days are treated separately, e.g. only time windows on weekdays can be selected when the target time window is on weekday.
2) Selected time window(s) should have the similar ambient temperature with the predicted time period (measured by Euclidean distance in eq. (2))

$$T_{MSE}^{D,l,t} = \frac{1}{N_t} \sum_{k=1}^{N_t} \left(T^{D,t}(k) - T^{l,t}(k) \right)^2. \tag{2}$$

D, l represent the index of the target day and a particular historical day, $t \in \{1, 2, 3, 4\}$ is the index of the time window representing the time period of hour ending 1–6, 7–12, 13–18 or 19–24, and N_t is the number of samples in each section. Therefore, the day-ahead baseline in this section is calculated as the average consumption of all corresponding N_s similar time windows for the **same** end-consumer,

$$P_{base,i}^{D,t}(k) = \frac{1}{N_s}\sum_{k=1}^{N_s} P_{base,i}^{l,t}(k),\qquad(3)$$

Therefore, the "similar day" algorithm 1) predicts the baseline calculating the average consumption of the similar 6-hour time window in the history, and 2) effectively eliminates the gaming effect of participants since no recent behavior (consumption data after the experiment begins) of the consumer is considered. Because of the benefits mentioned above, the "similar day" algorithm was implemented in *both baseline estimate and data analysis* in the recent EnergyCoupon experiment ('17) in 2017.

The accuracy of the "similar day" algorithm in baseline prediction has been evaluated right before the experiment started in 2017. Results shows the average mean absolute percentage of error (MAPE) was around 20% on average for all participants, which is about the same level (15% - 30%) with other machine-learning methodologies applied to individual households [41]. The "similar day" algorithm has been extended to other areas such as non-intrusive load monitoring (NILM) and has achieved over 80% accuracy in all tested datasets [42].

3.3 Individualized target settling and coupon generation

In the EnergyCoupon program, there are two types of DR events: "fixed" and "dynamic" events. Both types of event last for 30 min, and can only be triggered between 1 and 7 pm each day. However, these two types of event follow quite different triggering methodologies:

Fixed DR events: We conduct a statistical analysis of historical prices in ERCOT's real-time market [1], and observe that high wholesale market prices more often occur at certain hours in the day than others, and the "high risk" hours vary over the month of the year. By following this discovery, no more than three "fixed" DR events will show at the fixed "high risk" hours every day, and the fixed hours may be different from month to month, and from weekdays to weekends.

Dynamic DR events: These are DR events that are triggered when the 2-hour ahead price prediction algorithm (introduced in Section 3.1) indicates that price is likely to be higher than \$50/MWh. There is no restriction on the number of "dynamic" events in a day.

Sometimes we use the term "**hybrid event**" to denote the situation when both types of DR events can be triggered for the user.

After the time period of a DR event is determined by either methodology, a *multi-layer coupon target* is generated based on the individual predicted baseline (as shown in Fig. 3). Based on reaching the different level of reduction (such as 30% and 70%) from the baseline, the participant will be given a different number of coupons. In the EnergyCoupon app, this procedure is visualized as
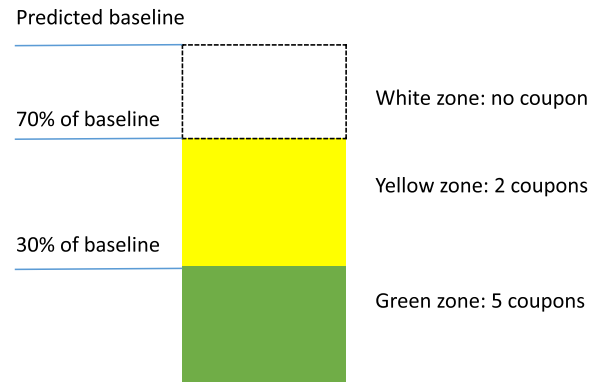
Fig. 3 Individual target setting

comparing the participant's real-time consumption with different colored areas (white, yellow and green) of the baseline. When the consumer's consumption lies between 70% of baseline or above, no EnergyCoupon is earned for this event; otherwise, the consumer will be awarded 2 EnergyCoupons when his/her consumption lies between 30% and 70% (the yellow area) and 5 EnergyCoupons if under 30% of the predicted baseline (the green area). We will use "*coupon*" as the synonym of EnergyCoupon in the rest this paper.

Figure 4 summarizes the logic flow of a coupon target generated based on algorithms introduced in Section 3.1 to 3.3.

3.4 Lottery algorithms

We use a lottery system to convert EnergyCoupons into monetary rewards. There is much work that has developed the concept of "prospect theory" to model the behavior of humans when exposed to lottery schemes [43–46]. The general finding is that humans are much more risk-seeking under larger low-probability rewards engendered by using a lottery system. Hence, lotteries have the potential of attaining larger reductions from the user population than a fixed reward. We observed this same effect during earlier numerical studies [25], and hence employed a lottery-based reward system in all our field trials.

In our experiment, *weekly* lotteries are conducted to convert end-consumers' coupons earned during DR events into dollar-value prizes. In each lottery, a participant is allowed to bid any number of coupons between zero and the total number of coupons in his/her account; the more coupons he/she bids, the higher probability he/she will win the prize. A pyramidal lottery scheme is designed, with three Amazon gift cards of face value of \$20, \$10 and \$5 as the first, second and third prizes each week. The brief lottery procedure is to conduct a top-down drawing at each level of the pyramid, remove the coupons of the winning user, and move to the lower level and continue. Hence, each participant
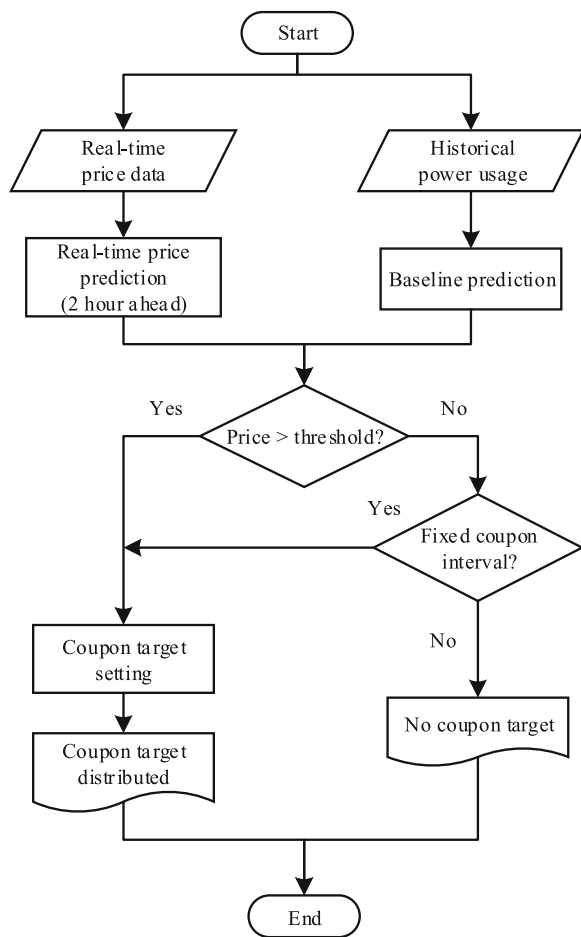
Fig. 4 EnergyCoupon algorithm flow chart

will have at most three chances of winning a prize with progressively smaller rewards at each drawing. Note that if a participant chooses to only bid a portion of his/her coupons at a particular lottery game, the remaining coupons can be saved in his/her account for future use. Therefore, a participant can be strategic in choosing the number of coupons that he/she bids in each game.

4 Experimental design

4.1 Brief summary of experiment ('16)
A small-scale preliminary EnergyCoupon experiment was conducted between June and August in 2016, with 7 end-consumers in a residential area in Cypress, Texas, United States enrolled. During the 12-week experiment, each participant received a number of 30-min-length DR events along with individualized coupon targets, between 1 and 7 pm every day, and was allowed to participate in lotteries with total prize of $35 Amazon gift cards each week. Peak time estimate, individualized target settling, coupon generation and lottery scheme followed the algorithms described in Section 3. A Hybrid baseline prediction method was used for the baseline estimate,

and the "similar day" algorithm was used in posterior data analysis.

The experiment revealed a load shifting effect from peak to off-peak hours; it yielded substantial savings for the LSE, about $0.44/(week·user) on average, and $1.15/ (week·user) per active user. Readers can refer to [15] for more details.

4.2 Subject in experiment ('17)
A larger-scale EnergyCoupon experiment was conducted in the summer of 2017, with 29 anonymous residential end-consumers in The Woodlands, TX recruited to form the treatment group. All participants were the customers of a local retail electric provider. Their participation was purely voluntary, and participants were free to quit the experiment at any time (though there was no one that actually quitted).

In addition, the retail electric provider also provided us with some residential electricity consumption data of another 16 anonymous households for the same period of time. These end-consumers formed the *comparison group*, and they neither participated in the DR event nor the periodic lotteries. The relationship between the treatment and comparison group is in Fig. 5a.

4.3 Procedure in experiment ('17)
All the treatment and comparison group participants had a smart meter installed in their household before the experiment, which made their 15-min interval electricity consumption data available on *SmartMeterTexas. com*, a websited endorsed by the Public Utility Commission of Texas [30]. With the permission of all participants, we were able to obtain their ESIID, register an account for them, and download their historical and real-time electricity consumption data periodically through the secure backend server located on the campus of Texas A&M University.

In test Week 0 (Jun 10-Jun 16, 2017), the treatment group subjects were asked to download and install the EnergyCoupon App, get familiar with the interface, practice how to undertake energy reduction by following individualized coupon targets, and participate in a trial lottery. The electricity consumption data during this period of
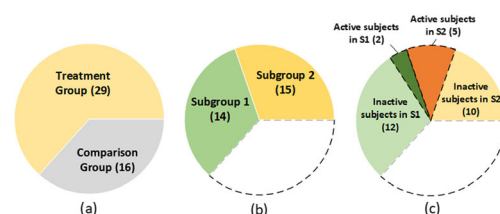

Fig. 5 Subjects in experiment ('17). **a** Treatment vs. comparison group. **b** Subgroup 1 vs. Subgroup 2. **c** Active vs. Inactive subgroups. Numbers in brackets are group sizes

time was neither considered as experimental data, nor used as the historical data in the baseline estimate.

During the experiment, the treatment group subjects were able to see the all daily "fixed" coupon targets at the beginning of each day, or "dynamic" coupon targets at least 2 hours prior to the DR event. A subject who wanted to save energy and earn coupons could turn off, or change the setpoints of his/her appliances during the 30-min-length DR event period without the need of notifying the organizer. The subject's electricity consumption would be recorded by the smart meter installed in the house, and data would become available and downloaded to the server within 36 h after the DR event. Thereafter each subject would be awarded coupons based on his/her coupon target achievement during the DR events.

In the first 3 weeks (Jun 17, 2017 to Jul 7, 2017), all the subjects in the treatment group were faced with "hybrid" coupon targets for their demand response. Starting from Week 4 (Jul 8, 2017), and till the end of the experiment, subjects were randomly assigned to two subgroups (**Subgroup 1 and 2**, or **S1** and **S2** for short) of almost the same size (14 subjects in S1 and 15 in S2). Subjects in S1 only received "fixed" coupon targets, while those in S2 only received "dynamic" coupon targets (Fig. 5b).

The "similar day" algorithm was used in the baseline estimate, and coupon target generation followed the algorithm in Section 3.3. DR events can only be triggered between 1 and 7 pm each day.

Weekly lotteries were conducted during the experiment, with each lottery cycle beginning at 12:00 am on Saturday and ending at 11:59 pm on the Friday of the following week. Lotteries are designed according to the schema explained in Section 3.4. In the posterior analysis at the end of the whole experiment, we further categorized all the subjects into another two subgroups according to their lottery engagements: "Active" subgroup contains *"active"* subjects who participated in *at least 5 out of a total of 11 lotteries*, and the remaining treatment group subjects are regarded as *"inactive"* subjects and are assigned to the "Inactive" subgroup. Figure 5c shows the relationship between two dimensions of categorization (based on coupon targets and lottery engagement), as there are a total of 7 active subjects among all participants, with 2 belonging to S1 and 5 belonging to S2. In contrast, among the remaining 22 inactive subjects, 12 of the inactive subjects belong to S1 and 10 belong to S2.

As we have briefly described in Section 1, some major differences exist between the designs of the EnergyCoupon experiment ('16) and ('17). The change of the algorithm from "hybrid" to "similar day" and the removal of normalization in the baseline estimate help to increase the baseline prediction precision, as well as eliminate the gaming effect. The availability of the comparison group provides an alternative means of measuring energy saving for the treatment group, and the assignment of S1 and S2 helps to reveal more intricate behavior of the treatment group subjects.

5 Results and discussion

In this section, we present an analysis of the data collected in our experiment ('17).

5.1 Energy saving for the treatment group

There are two ways to measure the electricity reduction for the treatment group during the experiment by means of comparing their electricity consumption with (i) the comparison group, or (ii) their own predicted baseline. Figure 6 exhibits the *energy consumption ratio* (we will call it "ratio" for short in Section 5.1 and 5.2) of the treatment and comparison groups following method (i). The ratio is defined as the group's weekly consumption **between 1 and 7 pm** divided by their own historical consumption during the same period in the previous year (2016). A lower ratio indicates a relatively higher behavior change of making more energy reduction during the experiment than that in the previous year.

Figure 6 also shows the energy saving for the active subjects during the experiment. While the ratios for inactive and comparison groups are overall close to each other in most weeks during the experiment, there is clear gap between the active subjects (red curve) and these two groups. It seems that active subjects who have more lottery engagement also have more significant better-than-average energy saving behaviors, with the maximum of around 40% in Week 8.
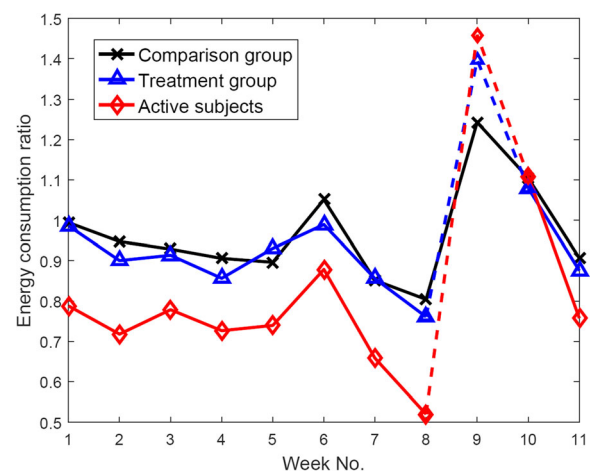

Fig. 6 Energy consumption ratio at 1–7 pm for experiment ('17) based on the consumption of same days in 2016 (Historical consumption data for some treatment group subjects in Week 9, 2016 is not available. We use dashed lines to show the less reliable trend between Week 8–9 and 9–10.)

The disadvantage of method (i) is that multiple variables between 2 years such as temperature are not well-controlled. Therefore, energy saving for the treatment group cannot be characterized precisely.

5.2 Comparison between active and inactive subjects in treatment group

As introduced above, method (ii) calculates *energy consumption ratio* using the subject's own estimated baseline as the denominator. Figure 7 shows ratios of active and inactive subgroups and the whole treatment group. We will show in the following paragraphs that the observation of energy saving using method (ii) is similar to method (i) in some sense.

The performance of the inactive subgroup is quite consistent, with the ratio around 1.0 in most weeks, and never falls below 0.9. This is in line with our intuition that *less engagement in the lottery is a sign of lack of enthusiasm about energy saving* via *the EnergyCoupon program*. Since inactive subjects form the majority of the treatment group (as shown in Fig. 5), the gap between the inactive subjects and the average value is minor, and this is partially because the majority of subjects in the treatment group are inactive subjects.

In contrast, the curve for the active subgroup is far below the other two curves, indicating a significant energy saving and load pattern change for active subjects during the experiment. Energy savings for the active subgroup gradually increase in the first few weeks and reaches a peak at about 40% in Week 8. After Week 9, the saving begins to decline, until being only 10% in Week 11. The rebound of the ratio can be explained by the arrival of Hurricane Harvey, which was in the area for the end of Week 10 and whole of Week 11. Flooding and potential house repair likely distracted many of the
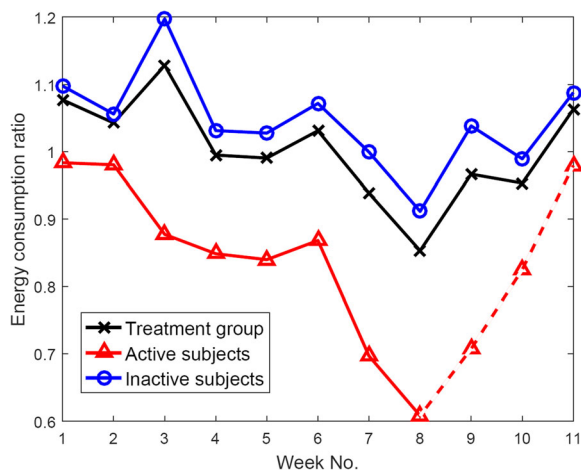
subjects from participating in the DR program during that time.

To better visualize the load pattern change for the active subjects, 1 week during the experiment (7/29–8/4/2017) was selected as an example, and the daily average of electricity consumption vs. baseline is illustrated for both active and inactive subgroups (Fig. 8). It can be calculated that for this particular week, energy saving during 1–7 pm for active subjects was 28.9%, while that of inactive subjects was only – 0.2%. The close-to-zero energy saving for inactive subjects is unsurprising, and it also supports the precision of our baseline estimate algorithm to some extent. However, the surprising finding from Fig. 8a is the load shedding effect in non-peak hours (25.0%). This observation clearly conflicts with the assumption of pure load shifting in our previous paper [15]. Therefore, an assumption could be created that there is some "inertia" in demand response; incentivized energy reduction in peak hours would influence that of off-peak hours.

5.3 Comparison between subjects in treatment group facing fixed/dynamic coupons

Starting from Week 3, and until the end of the experiment, the treatment group subjects were randomly assigned into two subgroups S1 and S2 facing "fixed" and "dynamic" coupon targets, respectively. We aim to discover how different types of coupon targets could have an impact on end-consumers' energy saving. The energy savings for two subgroups S1 and S2 during 1–7 pm are exhibited in Fig. 9a.

As observed from Fig. 9a, the subjects in S1 and S2 cannot be considered homogeneous, as the energy saving (35% vs. -5%) was quite different in Week 1–3 when they are facing the same "hybrid" coupon targets. For the following weeks (3–10) when subjects were separated with different coupon targets, we see an "activation" phenomenon by the dynamic coupon targets, as S2's saving jumps from – 5% to 15%, while no such effect is observed for S1 subjects. In week 11, the energy saving for S2 returns to the initial level. This can be attributed to the hurricane, as mentioned before.
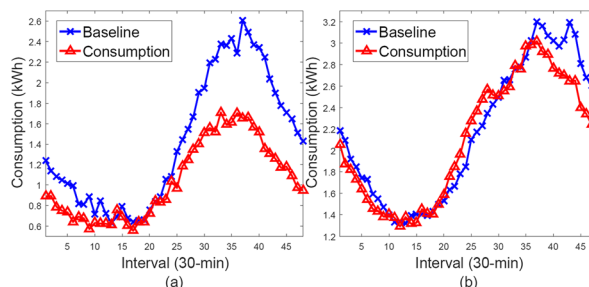
Fig. 7 Energy consumption ratio at 1–7 pm for active/inactive subjects by week, based on their baseline

Fig. 8 Daily consumption vs. baseline for active/inactive subjects during the week (7/29–8/4/2017). **a** active subjects, **b** inactive subjects
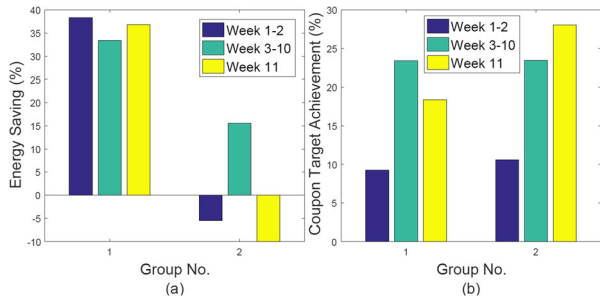
Fig. 9 Behavior comparisons between subjects in Subgroup S1 and S2. **a** Average energy saving at 1–7 pm. **b** Coupon target achievement percentage

Figure 9b illustrates the coupon target achievement ratios for active subjects in two subgroups. The ratio is defined as the proportion of DR events that the subjects at least earn one coupon (which is equal to a reduction of at least 30% energy saving from their baseline). Comparing Fig. 9a and b, an interesting finding is that *although S1 has overall higher energy saving than S2 in all periods, both subgroups reach a similar level of coupon achievement.*

One possible explanation for this observation is that S1 subjects facing "fixed" DR events would prefer to program their home appliances (such as AC) in advance to hit all coupon targets, and do not change their setpoints frequently, while S2 subjects facing "dynamic" DR events tended to check the app and DR events more frequently and tried to "play" to catch the coupon targets which only started to appear 2-hours before real-time. Figure 10 shows the load patterns for two active subjects in S1 and S2 as an example, and how the subject in S1 reducing the consumption for the whole afternoon vs. S2 moving his/her consumption to catch the yellow targets.

5.4 Financial benefit analysis

In our earlier analytical model and numerical studies, we assumed that all the subjects perform *pure load shifting* from peak to off-peak hours [15]. With such an
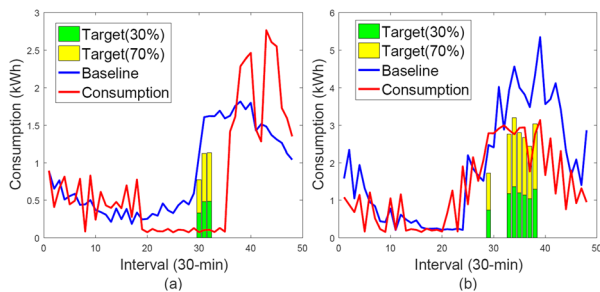

Fig. 10 Energy consumption curve for two active subjects on 7/19/2017. **a** Subject No. 19 (in Subgroup 1). **b** Subject No. 18 (in Subgroup 2)

assumption, the DR program would lead to a win-win situation with positive financial benefits to both the retail provider and active end-consumers. The brief explanation of this is that with pure load shifting in DR, the retailer will not lose its retail revenue, and can purchase electricity at the time period with lower wholesale market prices. At the same time, end-consumers earn reward for their energy saving behavior. However, our finding of load shedding behaviors in Section 5.2 conflicts with this *pure load shifting* assumption. Therefore, it is not obvious that the retail provider and end-consumers can still reach a win-win situation as described before. Below is the analysis we have using the newly gathered data in the experiment ('17).

The net benefit for the retail provider consists of three parts: (i) the saving in (wholesale) electricity purchase in high-price hours, (ii) the decrease of sales revenue due to the load shedding effect, and (iii) the cost of rewards issued to lottery winners. Our calculation shows that the saving in three parts is $2.6, $-2.7 and $-4.0(week•subject). Because of the load shedding effect, the benefit in (a) is not enough to cover the loss in (b) and the retail provider suffers a net loss of around -$4.0 for each active user per week. Note that this loss is localized to the year 2017, because of the low oil prices, and consequent low electricity peak prices in the summer of that year. The same DR program conducted in other years such as 2019 might have yielded substantial benefits because of the record-breaking high prices [47] (Table 3).

An active subject, in contrast, on average receives $4.0 lottery rewards per week from the retail provider; at the same time, the load shedding effect leads to the decrease of his/her electric bill by around $2.7 per week. Therefore, our EnergyCoupon program brings positive financial benefit to active subjects.

Although the DR program may not bring a win-win situation to both the retail provider and end-consumers, it still does increase the social welfare on the demand side, as the summation of benefits is positive ($2.6/(week•active subject)). We can also conclude that the financial benefit of the retail provider in the DR program is closely related to the load shifting /load shedding pattern of each subject; in experiment ('17), the load shift is minor and load shedding is major, and cost saving from its wholesale electricity purchase may not cover the loss of its retail revenue, which leads to a net financial loss to the retail provider.

The profit of the retail provider can also be affected by other factors such as 1) real-time electricity price. The

Table 3 Financial benefit of the retail provider and active subjects

Subjects	Retail provider	Active subjects
Savings ($/(week·subject))	$2.6 - 2.7 - 4.0 = -4.1$	$4.0 + 2.7 = 6.7$

high real-time electricity price would increase the value of demand reduction and therefore increase the profit for the retailers; 2) the management of the lottery budget, with a proper choice of prize, could possibly decrease the total cost of the lottery while keeping demand reduction at an acceptable level; 3) possible subsidies on demand response programs in some countries or grids.

It is worth noticing that since factors mentioned above could vary with different human participant groups, in different physical areas, or even in different years in the same area, there is no general conclusion as to whether Energy-Coupon (or other DR programs) would help the retail provider to save money or not.

5.5 Influence of the lottery on human behaviors
As discussed in Section 3.4, the lottery scheme is considered to provide an incentive to promote desirable behaviors (such as more energy saving and participation) of the treatment group. Table 4 lists some numbers showing the influence of periodic lotteries on participant behavior.

The first column in Table 4 shows that winning a lottery prize has a positive impact on future energy saving, as lottery winners make an average of energy saving improvement of 10.7% in the next lottery cycle.[2] In contrast, the average energy saving improvement for participants who win nothing is close to zero (– 0.03). The second and third columns clearly demonstrate that lottery winners on average tend to have higher engagements than other participants in the next lottery (56.6% vs. 40.0%), and the next three lotteries (80.5% vs. 70.0%). Therefore, we can summarize that the lottery prize has a positive impact on both energy-saving and lottery engagement in future lottery cycles.

5.6 Comparison with previous CPP experiment
In this subsection, we compare our EnergyCoupon experiment with a typical price-based DR experiment conducted in Anaheim, California, United States in the year 2005 [14]. Critical peak pricing (CPP) was used in this experiment. Here, CPP days are selected based on a price prediction algorithm, and on CPP days during noon-6 pm. Subjects in the treatment group receive $0.35 for every *kWh* reduction from the baseline. Some comparisons of these two experiments are listed in Table 5.

We observe that our experiment reached a similar level of energy reduction to that of the CPP experiment (10.7% to 12%). Since our EnergyCoupon provides DR events every day compared to only 12 CPP days in the CPP experiment, the EnergyCoupon project helps to save a much higher amount of energy in total.

Table 4 Subjects' behavior change due to lottery

Subjects	Energy saving improvement (1–7 pm)	Next lottery participation prob. (%)	Prob. of at least 1 participation in next 3 lotteries (%)
All winners	10.7	56.6	80.5
Other participants	−0.03	40.0	70.0

In addition, effective cost is calculated as an indicator of cost saving efficiency for each experiment, as it is defined as on average the money the retailer has to pay for participants' reducing 1 *kWh* of electricity during peak hours. This value in our experiment is calculated by the total value of lottery prizes divided by the energy reduction for all treatment group subjects. Data analysis shows an effective cost in our experiment of $0.053/kWh, which is only 1/7 of that in the CPP experiment (which is directly given as $0.35 in the experimental design).

5.7 Cost saving decomposition
Table 5 shows the significant difference between the effective costs of two demand response experiments. As an indicator of relative cost effectiveness, *effective cost saving ratio* (ECSR) is defined as the ratio of effective cost (normalized by retail price) between two experiments. If we use the CPP experiment [14] as the reference case, the ECSR of EnergyCoupon is calculated as

$$ECSR = \frac{C_{CPP}}{C_{EnergyCoupon}} = \frac{368\%}{58\%} = 6.34 \qquad (4)$$

$ECSR > 1$ indicates that EnergyCoupon is more cost-effective than CPP.

There are different factors that may contribute to the high ECSR value shown in eq. (4), such as (i) the innovative coupon design in CIDR mechanism, (ii) the development of the EnergyCoupon mobile app that improves

Table 5 Comparison between EnergyCoupon and previous experiment

Comparison item	Critical peak pricing	EnergyCoupon
Category	Price-based DR	Incentive-based DR
Treatment group size	71	29
Experiment length	Jun - Oct	Jun – Aug
Number of DR days	Certain CPP days (12)	Daily (77)
Peak hours	Noon – 6 pm	1–7 pm
Energy reduction	12%	10.7%[a]
Effective cost compared with retail price[b]	368%	58.8%

[a]Electricity reductions for active and inactive and inactive subjects are 34.8% and 7.36%, respectively
[b]We choose retail price in Anaheim in 2005 as $0.095/kWh [48], and average retail price in Woodland, TX in 2017 as $0.090/kWh [13]

---

[2]As an example, 1% improvement means if this week's saving is 10%, next week will be 11%.

the communication with participants as well as (iii) the lottery scheme that encourages more participation because of human risk seeking behavior.

In this section, we are interested in how the lottery scheme (factor (iii)) contributes to the cost saving in our EnergyCoupon program. If we assume that all the factors listed above contribute independently to ECSR and can be measured by multipliers, then

$$\alpha\beta = ECSR = \frac{C_{CPP}}{C_{EnergyCoupon}} = 6.34 \tag{5}$$

where α, β are multipliers representing the contribution of the lottery scheme (factor (iii)) and other factors ((i),(ii)...) respectively.

The value α can be estimated using cumulative prospect theory. As a behavioral game theory, this theory describes the individual choice between risky probabilistic alternatives [44]. It models the probability weighting and loss aversion, which lead to the overweighting of small probabilities and underweighting of moderate and high probabilities. In the game with potential outcomes $x_1, x_2, ..., x_n$ and respective probabilities $p_1, p_2, ..., p_n$, a **gain prospect** $f = (x_1, p_1; x_2, p_2; ...; x_n, p_n)$ describes a prospect results in the multiple outcome x_i with probability p_i, $i \in \{1, 2, ..., n\}$, and

(i) $x_i < x_j$ iff $i < j, i, j \in \{1, 2, ..., n\}$
(ii) $\sum_{i=1}^{n} p_i = 1$.

For instance, in EnergyCoupon, on average each active subject has an approximately 7.0% chance to win each prize ($20, $10 and $5) in the weekly lottery; the probability for an inactive user to win each prize is around 2.3%. Therefore, the prospect of each active/inactive subject faces (f_a and f_b) can be described as

$$\begin{aligned} f_a &= (0, 0.79; 5, 0.07; 10, 0.07; 20, 0.07) \\ f_b &= (0, 0.931; 5, 0.023; 10, 0.023; 20, 0.023) \end{aligned} \tag{6}$$

and $n = 4$ for both prospects.

The prospect theory defines the utility of a certain prospect f as

$$V(f) = \sum_{i=0}^{N} \pi_i V(x_i) \tag{7}$$

where V is the utility function, π_i are *decision weights* calculated as

$$\pi_i = \omega(p_i+, ..., +p_n) - \omega(p_{i+1}+, ..., +p_n), 0 \le i \le N \tag{8}$$

and ω is the *probability weighting function*. Equation (7) can be explained as the utility of a prospect f equalling the sum of all decision weights π_i times the utility of the corresponding outcomes x_i. It is worth noting that decision weight π_i has close correlation with the probability p_i but they may have different values. Deviation of π_i from p_i

represents the way the lottery scheme "distorts" human beings' feeling for the probabilities.

Furthermore, we introduce the *equivalent* of prospect f as c, which can also be described as

$$V(c) = V(f) \tag{9}$$

Therefore, the equivalent for prospect f_a as c_a represents the fixed-return an active participant receives that would make him/her indifferent between choosing the fixed return c_a or play in the lottery f_a. The same explanation applies to c_b (inactive users). Given the total number of active users N_a and inactive users N_b, the total equivalent

$$c = c_a N_a + c_b N_b \tag{10}$$

would be an estimate of total direct cash needed in the experiment to maintain the same level of incentive to the treatment group, if no lottery scheme is adopted. In the EnergyCoupon experiment, $N_a = 7$ and $N_b = 22$, which reflect the number of active/inactive participants.

As the next step, we would like to get an estimate of value c_a and c_b. Combining the definition of fixed-return equivalent (9) with eq. (10), we have

$$V(c) = \sum_{i=0}^{N} \pi_i V(x_i). \tag{11}$$

Since in our experiment each lottery prize is relatively small ($x_i < \$200$, $i \in \{1, 2, 3, 4\}$), the utility function V is linear and can be removed from both sides of (7) [46] as

$$c = \sum_{i=0}^{N} \pi_i x_i. \tag{12}$$

By combining eq. (12) and (8) we can calculate the equivalent per active/inactive user. We take a typical active user as an example. The equivalent c_a of prospect f_a can be calculated as

$$\begin{aligned} c_a &= \pi_2 \times 5 + \pi_3 \times 10 + \pi_4 \times 20 \\ \pi_2 &= \omega(0.07 \times 3) - \omega(0.07 \times 2) = \omega(0.21) - \omega(0.14) \\ \pi_3 &= \omega(0.07 \times 2) - \omega(0.07) = \omega(0.14) - \omega(0.07) \\ \pi_4 &= \omega(0.07) \end{aligned} \tag{13}$$

The value of ω can be estimated from Fig. 1 in reference [44], as median c/x in prospect $(0, 1 - p; x, p)$ is an estimate of $\omega(p)$. From the curve $x < 200$ we get the value $\omega(0.21) = 0.26$, $\omega(0.14) = 0.22$ and $\omega(0.07) = 0.16$. Therefore

$$\begin{aligned} c_a &= (0.26 - 0.22) \times 5 + (0.22 - 0.16) \times 10 \\ &\quad + 0.16 \times 20 = 4.0 \end{aligned} \tag{14}$$

Similarly we can calculate the equivalent $c_b = 2.4$. According to eq. (10), the total equivalent for all

participants $c = 4.0 \times 7 + 2.4 \times 22 = \$80.8$. Total equivalent c shows the estimate of direct cash needed in our experiment to maintain the same level of incentive to the treatment group, if no lottery scheme is adopted. Therefore, multiplier α is estimated as the ratio of equivalent cash divided by total weekly lottery prizes $\alpha = 80.8/35 = 2.3$. According to eq. (5), $\beta = 2.75$ and we can conclude that the lottery scheme and other EnergyCoupon designs have *similar* levels of contribution to reducing the effective cost in our experiment.

6 Concluding remarks

This paper presents the design and critically assesses the empirical experiment of a coupon incentive-based demand response program for end-consumers over a two-year period in the Houston area, United States. Different from traditional price-based DR programs, EnergyCoupon has the following features: (1) Dynamic time-of-use DR events and individualized coupon targets; (2) End-consumers receive coupon targets and usage statistics through a mobile app; (3) Voluntary participation in demand response events, and (4) Periodic lottery that allows the participant to convert their coupons into dollar-value prizes.

Data analysis shows the significant load shedding effect for the treatment group; however, not much load shifting effect is observed. In addition, we observe the positive impact of lottery prizes on the growth of desirable behavior, such as energy saving improvement and lottery participation. Our posterior analysis also shows that EnergyCoupon has much lower effective cost (¢5.3/ kWh) compared to previous CPP projects (¢35.0/kWh); Using prospect theory we estimate that the design of system architecture and lottery scheme have equal-level contributions to the cost saving.

This paper is generalizable towards other Internet-of-Things-enabled demand response activity, and could shed light on the overall discussion of incentive-based versus price-based demand response. Future work would examine the value added by obtaining consumer behavior data in this experiment. Another possible avenue of future work is to further develop a platform that allows for the end-consumers to aggregate and participate in wholesale-level ancillary services.

Abbreviations
CAISO: California ISO; CIDR: Coupon incentive-based demand response; CPP: Critical peak pricing; DR: Demand response; ERCOT: Electric Reliability Council of Texas; ISOs: Independent system operators; NYISO: New York ISO; TOU: Time-of-use

Authors' contributions
The authors all participated in the EnergyCoupon experiment. L. Xie and S. Shakkottai conceived the idea and supervised the research. Hao Ming, Bainan Xia, Srinivas Shakkottai and Le Xie developed algorithms in EnergyCoupon and write the manuscript. Ki-Yeob Lee, Adekunle Adepoju was in charge of mobile app development and data management. The author(s) read and approved the final manuscript.

Authors' information
Not applicable.

Availability of data and materials
Reader only has limited access to the data and materials due to the policy in SmartMeterTexas.com. Reader may need to contact the corresponding author and get permission from the retail provider to have full access to the data.

Competing interests
The authors declare that they have no competing interests.

Author details
[1]Electric Reliability Council of Texas, 800 Airport Rd, Taylor, TX 78664, USA. [2]Intellectual Ventures, Bellevue, WA, USA. [3]Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, USA.

References
1. Electric Reliability Council of Texas. (2020). ERCOT.com, [online] Available: http://www.ercot.com/.
2. Green the Grid, Green the Grid: the role of storage and demand response. (2015) [online] Available: https://www.nrel.gov/docs/fy15osti/63041.pdf.
3. U.S. Department of Energy. (2006). Benefit of demand response in electricity market and recommendations for achieving them, [online] Available: http://www.caiso.com/Documents/DemandResponseandProxyDemandResourcesFrequentlyAskedQuestions.pdf.
4. New York Independent System Operator (NYISO). (2009). Day-ahead demand response program manual, [online] Available: http://www.nyiso.com/public/webdocs/markets_operations/documents/Manuals_and_Guides/Manuals/Operations/dadrp_mnl.pdf.
5. California ISO. (2015). Demand response user guide, [online] Available: http://www.caiso.com/Documents/July20_2009InitialCommentsonPDAdoptingDRActivities_Budgets_2009-2011inDocketNos_A_08-06-001_etal_.pdf.
6. California ISO. (2011). Demand response & proxy demand resource-frequently asked questions, [online] Available: http://www.caiso.com/Documents/DemandResponseandProxyDemandResourcesFrequentlyAskedQuestions.pdf.
7. D. Hurley, P. Peterson, M. Whited. Demand response as a power system resource. (2013). Synapse Energy Economics Inc.
8. Ruiz, N., Cobelo, I., & Oyarzabal, J. (2009). A direct load control model for virtual power plant management. *IEEE Transactions on Power Apparatus and Systems, 24*(2), 959–966.
9. Gholian, A., Rad, H. M., & Hua, Y. (2016). Optimal industrial load control in smart grid. *IEEE Transactions on Smart Grid, 7*(5), 2305–2316.
10. Parvania, M., Fotuhi-Firuzabad, M., & Shahidehpour, M. (2014). ISO's optimal strategies for scheduling the hourly demand response in day-ahead markets. *IEEE Transactions on Power Apparatus and Systems, 29*(6), 2636–2645.
11. ClearlyEnergy.com. (2018). Residential demand response programs, [online] Available: https://www.clearlyenergy.com/residential-demand-response-programs.
12. Airedale. (2018). ACIS building energy management system, [online] Available: http://airedale.com.
13. Power2Choose.org. (2018). [online] Available: http://www.powertochoose.org.

14. F. A. Wolak. (2007). Residential customer response to real-time pricing: The anaheim critical peak pricing experiment. Center for the Study of Energy Markets.

15. Xia, B., Ming, H., Lee, K.-Y., Li, Y., Zhou, Y., Bansal, S., Shakkottai, S., & Xie, L. (2017). Energycoupon: a case study on incentive-based demand response in smart grid. In *Proceedings of the Eighth International Conference on Future Energy Systems, ACM* (pp. 80–90).

16. U.S. Department of Energy. (2015). The pecan street project: developing the electric utility system of the future, Ph.D. thesis.

17. ENERNOC: An Enel Group Company. (2020). https://www.enernoc.com/resources/datasheets-brochures/get-more-commercial-and-industrial-demand-response.

18. OhmConnect. (2020). https://www.ohmconnect.com/.

19. Zhong, H., Xie, L., & Xia, Q. (2012). Coupon incentive-based demand response (CIDR) in smart grid. In *2012 IEEE Power and Energy Society General Meeting, IEEE* (pp. 1–6).

20. Zhong, H., Xie, L., & Xia, Q. (2013). Coupon incentive-based demand response: Theory and case study. *IEEE Transactions on Power Apparatus and Systems, 28*(2), 1266–1276.

21. Ming, H., & Xie, L. (2014). Analysis of coupon incentive-based demand response with bounded consumer rationality. In *North American Power Symposium (NAPS), 2014, IEEE* (pp. 1–6).

22. G. A. Schwartz, H. Tembine, S. Amin, S. S. Sastry. (2012). Electricity demand shaping via randomized rewards: A mean field game approach. Allerton Conference on Communication, Control, and Computing.

23. Li, J., Xia, B., Geng, X., Ming, H., Shakkottai, S., Subramanian, V., & Xie, L. (2015). Energy coupon: a mean field game perspective on demand response in smart grids. *ACM SIGMETRICS Performance Evaluation Review, 43*(1), 455–456.

24. Schwartz, G. A., Tembine, H., Amin, S., & Sastry, S. S. (2014). Demand response scheme based on lottery-like rebates. *IFAC Proceedings Volumes, 47*(3), 4584–4588.

25. Li, J., Xia, B., Geng, X., Ming, H., Shakkottai, S., Subramanian, V., & Xie, L. (2018). Mean field games in nudge systems for societal networks. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS), 3, 15*(4).

26. Merugu, D., Prabhakar, B. S., & Rama, N. (2009). An incentive mechanism for decongesting the roads: a pilot program in Bangalore. In *Proc. of ACM NetEcon workshop*.

27. NuRide, Inc. (2020). [online] Available: https://www.nuride.com.

28. A. Ju. (2014). Experiment makes energy savings a game. Cornell Chronicle.

29. Pan, S., Shen, Y., Sun, Z., Mahajan, P., Zhang, L., & Zhang, P. (2013). Demo abstract: saving energy in smart commercial buildings through social gaming. In *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication, ACM* (pp. 43–46).

30. Public Utilities Commission of Texas. (2017). https://www.smartmetertexas.com/home.

31. Weather Underground. wunderground.com. [online] Available: https://www.wunderground.com.

32. Chao, H.-p. (2011). Demand response in wholesale electricity markets: the choice of customer baseline. *Journal of Regulatory Economics, 39*(1), 68–88.

33. Nogales, F. J., Contreras, J., Conejo, A. J., & Espinola, R. (2002). Forecasting next-day electricity prices by time series models. *IEEE Transactions on Power Apparatus and Systems, 17*(2), 342–348.

34. Contreras, J., Espinola, R., Nogales, F. J., & Conejo, A. J. (2003). Arima models to predict next-day electricity prices. *IEEE Transactions on Power Apparatus and Systems, 18*(3), 1014–1020.

35. Conejo, A. J., Plazas, M. A., Espinola, R., & Molina, A. B. (2005). Day-ahead electricity price forecasting using the wavelet transform and Arima models. *IEEE Transactions on Power Apparatus and Systems, 20*(2), 1035–1042.

36. Wu, L., & Shahidehpour, M. (2010). A hybrid model for day-ahead price forecasting. *IEEE Transactions on Power Apparatus and Systems, 25*(3), 1519–1530.

37. Jonsson, T., Pinson, P., Nielsen, H. A., Madsen, H., & Nielsen, T. S. (2013). Forecasting electricity spot prices accounting for wind power predictions. *IEEE Transactions on Sustainable Energy, 4*(1), 210–218.

38. Muthirayan, D., Kalathil, D., Poolla, K., & Varaiya, P. (2016). Mechanism design for self-reporting baselines in demand response. In *American Control Conference (ACC), 2016, American Automatic Control Council (AACC)* (pp. 1446–1451).

39. Wikipedia. (2019). K-nearest Neighbors Algorithm. [online] Available: https://en.wikipedia.org/wiki/K-nearest_neighbors_algorithm.

40. Wikipedia. (2019). Kernel Regression. [online] Available: https://en.wikipedia.org/wiki/Kernel_regression.

41. Edwards, R. E., New, J., & Parker, L. E. (2012). Predicting future hourly residential electrical consumption: a machine learning case study. *Energy and Buildings, 49*, 591–603.

42. Shi, X., Ming, H., Shakkottai, S., Xie, L., & Yao, J. (2019). Nonintrusive load monitoring in residential households with low-resolution data. *Applied Energy, 252*, 113283.

43. Kahneman, D., & Tversky, A. (2013). Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: part I, world scientific* (pp. 99–127).

44. Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297–323.

45. Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology, 38*(1), 129–166.

46. Abdellaoui, M., Bleichrodt, H., & L'Haridon, O. (2008). A tractable method to measure utility and loss aversion under prospect theory. *Journal of Risk and Uncertainty, 36*(3), 245.

47. C. Martin, N. S. Malik. (2019). Power blows past $9,000 cap in Texas as heat triggers emergency, [online] Available: https://www.bloomberg.com/news/articles/2019-08-13/texas-power-prices-briefly-surpass-9-000-amid-searing-heat.

48. C. E. Commission. (2018). Sta_ forecast: average retail electricity prices 2005 to 2018. [online] Available: http://www.energy.ca.gov/2007publications/CEC-200-2007-013/CEC-200-2007-013-SD.PDF.