Title
Social Game for Building Energy Efficiency: Utility Learning, Simulation, and Analysis

Authors
Konstantakopoulos, Ioannis C
Ratliff, Lillian J
Jin, Ming
et al.

Publication Date
2014-11-17

Social Game for Building Energy Efficiency: Utility Learning, Simulation, and Analysis

Ioannis C. Konstantakopoulos⋆⋆, Lillian J. Ratliff⋆, Ming Jin⋆⋆, S. Shankar Sastry⋆, and Costas Spanos⋆⋆

Department of Electrical Engineering and Computer Sciences
University of California, Berkeley,
Berkeley, California, 94720
{ioanniskon,ratliffl,jinming,sastry,spanos}@eecs.berkeley.edu

Abstract. We describe a social game that we designed for encouraging energy efficient behavior amongst building occupants with the aim of reducing overall energy consumption in the building. Occupants vote for their desired lighting level and win points which are used in a lottery based on how far their vote is from the maximum setting. We assume that the occupants are utility maximizers and that their utility functions capture the tradeoff between winning points and their comfort level. We model the occupants as non-cooperative agents in a continuous game and we characterize their play using the Nash equilibrium concept. Using occupant voting data, we parameterize their utility functions and use a convex optimization problem to estimate the parameters. We simulate the game defined by the estimated utility functions and show that the estimated model for occupant behavior is a good predictor of their actual behavior. In addition, we show that due to the social game, there is a significant reduction in energy consumption.

Keywords: Utility Learning, Energy Efficiency, Game Theory

1 Introduction

Energy consumption of buildings, both residential and commercial, accounts for approximately 40% of all energy usage in the U.S. [16]. Lighting is a major consumer of energy in commercial buildings; one-fifth of all energy consumed in buildings is due to lighting [23].

⋆ NSF CPS:Large:ActionWebs award number 0931843, TRUST (Team for Research in Ubiquitous Secure Technology) which receives support from NSF (award number CCF-0424422), and FORCES (Foundations Of Resilient CybEr-physical Systems) which receives support from NSF (award number CNS-1239166)

⋆⋆ This research is funded by the Republic of Singapores National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore.

There have been many approaches to improve energy efficiency of buildings through control and automation as well as incentives and pricing. From the meter to the consumer, many control methods, such as model predictive control, have been proposed as a means to improve the efficiency of building operations (see, e.g., [3],[5],[6],[13],[18],[12]). From the meter to the energy utility, many economic solutions have been proposed, such as dynamic pricing and smart meter technology, to reduce consumption by providing economic incentives (see, e.g.,[14],[21]).

Many of the past approaches to building energy management only focus on heating and cooling of the building. We are advocating that due to new technological advances in building automation, incentives can be designed around more than just heating, ventilation and air conditioning (HVAC) systems. In particular, our experimental set-up allows us to design incentives based on lighting and individual plug-load in addition to HVAC and interact with occupants through a social game.

Social games have been used to encourage energy efficient behavior in transportation [17] as well as in the healthcare domain for understanding the tradeoff between privacy and desire to win by expending calories [4].

There are many ways in which a building manager can be motivated to encourage energy efficient behavior. The most obvious is that they pay the bill or, due to some operational excellence measure, are required to maintain an energy effcient building. Beyond these motivations, recently demand response programs have begun to be implemented by utility companies with the goal of correcting for improper load forecasting (see, e.g., [1],[15], [11]). In such a program, consumers enter into a contract with the utiltiy company in which they agree to change their demand in accordance with some agreed upon schedule. In this scenario, the building manager may now be required to keep this schedule.

Our approach to efficient building energy management focuses on office buildings and utilizes new builidng automation products such as the Lutron lighting system[1]. We design a social game aimed at incentivizing occupants to modify their behavior so that the overall energy consumption in the building is reduced. The social game consists of occupants logging their vote for the lighting setting in the office and they win points based on how energy efficient their vote is compared to other occupants. The average of the votes is what is actually implemented in the office. The points are used to determine an occupants likelihood of winning in a lottery. We designed an online platform so that occupants can login and vote, view their points, and observe all occupants consumption patterns and points. This platform also store all the past data allowing us to use it for estimation of the behavior of the occupants.

In this paper we present the results of a social game focused only on the encouraging more energy efficient lighting usage; however, we emphasize that the framework is easily adapted to incorporate the full capabilities of the automation installed in our experimental set-up (i.e. lighting, HVAC, and plug-load). The occupants are modeled as utility maximizers who engage in a non-cooperative

[1] http://www.lutron.com/en-US/Pages/default.aspx

game with all occupants. We parameterize their utility functions in such a way that we capture the tradeoff between the desire to win and comfort. Using data from the social game that occured over the period of roughly three months, we formulate the utility learning problem as a convex optimization problem and form estimates of each occupants utility function. We simulate the game using the estimated utility functions and show that the Nash equilibrium from the simulations is a good predictor of occupant behavior. Our results are compared to other estimation techniques.

A major advantage of modeling occupants as utility maximizers competing in a game and using the Nash equilibrium concept is this game theoretic model fits in the Stackelberg framework for incentive design in which the builiding manager performs an online estimation of occupant's utility function and designs incentives for behavior modification. This, in essence, is a problem of closing-the-loop around the occupants so that the building manager acheives sustained energy savings. We leave this as future work.

The rest of the paper is organized as follows. In Section 2, we start with the game theoretic framework for modeling the competitive environment between the non-cooperative occupants. We formulate the utility estimation problem as a convex optimization problem and take a dynamical systems perspective for developing a method of computing the Nash equilibrium of the estimated game. In Sections 3 and 4 we describe the experimental set-up and report on our findings including utility estimation results as well as simulation of the game corresponding to the estimated utilities. Finally, in Section 5, we make concluding remarks and comment on future research directions.

2 Game Formulation

We begin by describing the game theoretic framework used for modeling the interaction between the occupants. We remark that the use of game theory for modeling the behavior of the occupants has several advantages. First, it is a natural way to model agents competiting over scarce resources. It can also be leveraged in the design of incentives for behavioral change in that it incorporates the ability to model the occupants as strategic players.

Let the number of occupants participating in the game be denoted by n. We model the occupants as utility maximizers having utility functions composed of two terms that capture the tradeoff between comfort and desire to win. We model their comfort level using a Taguchi loss function which is interpreted as modeling occupant dissatisfaction as increasing as variation increases from their desired lighting setting. In particular, each occupant has the following Taguchi loss function as one component of their utility function:

$$\psi_i(x_i, x_{-i}) = -\left(\bar{x} - x_i\right)^2 \tag{1}$$

where $x_i \in \mathbb{R}$ is occupant i's lighting vote, $x_{-i} = \{x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$, and

$$\bar{x} = \frac{1}{n} \sum_{i=1}^{n} x_i \tag{2}$$

is the average of all the occupant votes and is the lighting setting which is implemented.

Each occupant's desire to win is modeled using the following function

$$\phi_i(x_i, x_{-i}) = \ln \left(\rho \frac{x_b - x_i}{n x_b - \sum_{j=1}^{n} x_j} \right) \tag{3}$$

where ρ is the total number of points distributed by the building manager and x_b is the baseline setting for the lights. The term inside the natural log function is how the points are distributed; ρ, being the total number of points, is multiplied by the distance an occupant's vote is from the baseline and then normalized by the sum of the differences of all occupants' votes from the baseline.

Hence, each occupant's utility function is given by

$$f_i(x_i, x_{-i}) = \psi_i(x_i, x_{-i}) + \theta_i \phi_i(x_i, x_{-i}) \tag{4}$$

where θ_i is an unknown parameter.

The occupants face the following optimization problem

$$\max_{x_i \in S_i} f_i(x_i, x_{-i}) \tag{5}$$

where $S_i = [0, 100] \subset \mathbb{R}$ is the constraint set for each x_i.

Note that each occupant's optimization problem is dependent on the other occupant's choice variables. Thus, the occupants are non-cooperative agents in a continuous game with constraints. We denote the joint strategy space by $\mathcal{C} = S_1 \times \cdots \times S_n \subset \mathbb{R}^n$. We model their interaction using the Nash equilibrium concept.

Definition 1. *A point $x \in \mathcal{C}$ is a* **Nash equilibrium** *for the game $(f_1, \ldots, f_n)$ on $\mathcal{C}$ if*

$$f_i(x_i, x_{-i}) \geq f_i(x_i', x_{-i}) \quad \forall\ x_i' \in S_i \tag{6}$$

for each $i \in \{1, \ldots, n\}$.

The interpretation of the definition of Nash is as follows: no player can unilaterally deviate and increase their cost.

If the parameters $\theta_i \geq 0$, then the game is a concave n-person game on a convex set.

Theorem 1 **([22]).** *A Nash equilibrium exists for every concave n-person game.*

We can define

$$\omega(x) = \begin{bmatrix} D_1 f_1(x) \\ \vdots \\ D_n f_n(x) \end{bmatrix} \tag{7}$$

where $D_i f_i$ denoets the derivative of f_i with respect to x_i.

It is the local representation of the differential game form [19] corresponding to the game between the occupants.

Definition 2 ([19]). *A point $x \in \mathcal{C}$ is a* **differential Nash equilibrium** *for the game $(f_1, \ldots, f_n)$ on $\mathcal{C}$ if $\omega(x) = 0$ and $D_{ii} f_i(x) < 0$.*

A sufficient condition guaranteeing that a Nash equilibrium x is isolated is that the Jacobian of $\omega(x)$, denoted $D\omega(x)$, is invertible [19],[22].

2.1 Utility Estimation

We formulate the utility estimation problem as a convex optimization problem by using first-order necessary conditions for Nash equilibria. In particular, the gradient of each occupant's utility function should be identically zero at the observed Nash equilibrium. This is the case since the observed Nash equilibria are all inside the feasible region so that none of the constraints are active, i.e. we do not have to check the derivative of Lagrangian of each occupant's optimization problem.

In particular, for each observation $x^{(k)}$, we assume that it corresponds to occupants playing a strategy that is approximately a Nash equilibrium where the superscript notation $(\cdot)^{(k)}$ indicates the k-th observation.

Definition 3. *A point $x \in \mathcal{C}$ is a* **ε-Nash equilibrium** *for the game $(f_1, \ldots, f_n)$ on $\mathcal{C}$ if*

$$f_i(x_i, x_{-i}) \geq f_i(x_i', x_{-i}) - \varepsilon \quad \forall\ x_i' \in S_i \tag{8}$$

for each $i \in \{1, \ldots, n\}$.

Thus, we can consider first-order optimality conditions for each occupants optimization problem and define a residual function capturing the amount of sub-optimality of each occupants choice $x_i^{(k)}$ [10],[20]. Note that all our observations are on the interior of the constraint set so we need only consider the following residual defined by the stationarity condition for each occupant's optimization problem:

$$r_i^{(k)}(\theta_i) = D_i f_i(x_i^{(k)}, x_{-i}^{(k)}) = D_i \psi_i(x_i^{(k)}, x_{-i}^{(k)}) + \theta_i D_i \phi_i(x_i^{(k)}, x_{-i}^{(k)}) \tag{9}$$

Define $r^{(k)}(\theta) = [r_1^{(k)}(\theta_1) \ \cdots \ r_n^{(k)}(\theta_n)]^T$.

Given observations $\{x^{(k)}\}_{k=1}^K$ where each $x^{(k)} \in \mathcal{C}$, we can solve the following convex optimization problem:

$$\min_\theta \left\{ \sum_{k=1}^K \chi(r^{(k)}(\theta)) \middle| \ \theta_i \geq 0 \ \ \forall\ i\{1, \ldots, n\} \right\} \tag{10}$$

where $\chi : \mathbb{R}^n \to \mathbb{R}_+$ is a nonnegative, convex penalty function satisfying $\chi(z) = 0$ if and only if $z = 0$, i.e. any norm on $\mathbb{R}^n$.

With a specific choice of χ we can explicitly write the estimation problem as follows. Let

$$\Psi_i = \begin{bmatrix} D_i\psi_i(x_i^{(1)}, x_{-i}^{(1)}) \\ \vdots \\ D_i\psi_i(x_i^{(K)}, x_{-i}^{(K)}) \end{bmatrix}, \ \Phi_i = \begin{bmatrix} D_i\phi_i(x_i^{(1)}, x_{-i}^{(1)}) \\ \vdots \\ D_i\phi_i(x_i^{(K)}, x_{-i}^{(K)}) \end{bmatrix} \tag{11}$$

for each $i \in \{1, \dots, n\}$ and denote $\theta = [\theta_1 \ \cdots \theta_n]^T$. Then, we can formulate the following convex optimization problem to solve for θ:

$$\min_\theta \left\{ \sum_{i=1}^{n} \|\Psi_i + \theta_i \Phi_i\|_2^2 \;\middle|\; \theta_i \geq 0 \ \ \forall \ i \in \{1, \dots, n\} \right\} \tag{12}$$

Note the constraint that the θ_i's be non-negative. This is to ensure that the estimated utility functions are concave. We add this restriction so that we can employ techniques from simulation of dynamical systems to the computation of the Nash equilibrium in the resulting n-person concave game with convex constraints.

2.2 Dynamical Systems Perspective

We can take a dynamical systems perspective in order to come up with a method for computation of the Nash equilibrium (see, e.g. [8], [19], [22]). We first write down a reasonable set of dynamics, then we show that a Nash equilibrium is a stable fixed point of these dynamics, and finally we suggest a subgradient projection method for computation.

It is natural to consider computing Nash equilibria by following the gradient of each occupant's utility function. Hence, we consider the dynamical system obtained by taking the derivative with respect to their choice variable of the Largrangian's for each occupant's optimization problem.

Indeed, let $h_{i,j}(x_i, x_{-i})$ for $j \in \{1, 2\}$ denote the constraints on occupant i's optimization problem. In particular, following Rosen [22], for occupant i, the constraints are

$$h_{i,1}(x_i) = 100 - x_i \tag{13}$$

$$h_{i,2}(x_i) = x_i \tag{14}$$

so that we can define $\mathcal{C}_i = \{x_i \in \mathbb{R} |\ h_{i,j}(x_i) \geq 0,\ j \in \{1, 2\}\}$ and $\mathcal{C} = \mathcal{C}_1 \times \cdots \times \mathcal{C}_n$. Due to the fact that our constraint set is convex, closed and bounded in $\mathbb{R}^n$ and there is a point in its strict interior, we satisfy a constraint qualification condition which is a sufficient condition for the Karush-Khun-Tucker (KKT) conditions for each occupant's optimization problem [2]. It is known that for concave games, i.e. concave player utility functions constrained on a convex set, given that the problem satisfies a constraint qualification condition, then a point satisfying KKT conditions for each player's optimization problem is a Nash equilibrium [22].

In addition, it is clear that each $h_{i,j}$ is smooth. Thus, the KKT conditions are given as follows. Let $x^* = (x_1^*, \dots, x_n^*)$ be a Nash equilibrium. Then, $h_{i,j}(x_i^*) \geq 0$ for each $i \in \{1, \dots, n\}$ and $j \in \{1, 2\}$. Further, for each $i \in \{1, \dots, n\}$ there exists $\mu_{i,j}^* \geq 0$, for $j \in \{1, 2\}$ such that $\mu_{i,j}^* h_{i,j}(x_i^*) = 0$ and

$$0 = D_i f_i(x^*) + \mu_{i,1}^* D_i h_{i,1}(x_i^*) + \mu_{i,2}^* D_i h_{i,2}(x_i^*). \tag{15}$$

We remark that the KKT conditions are necessary for optimality of each occupant's individual optimization problem and for a Nash equilibrium.

We can study the continuous-time dynamical system generated by the gradient of the Lagrangian of each occupant's optimization problem with respect to her own choice variable; we let

$$\dot{x}_i = D_i f_i(x_i, x_{-i}) + \sum_{j=1}^{2} \mu_{i,j} D_i h_{i,j}(x_i) \tag{16}$$

for $i \in \{1, \dots, n\}$ and where $\mu_{i,j}$ is the j-th dual variable for occupant i's optimization problem. The first term is the derivative of occupant i's utility with respect to her own choice variable x_i. The second term, with the appropriate dual variables $\mu_{i,j}$, ensures that for any initial condition in the feasible set $\mathcal{C}$, the trajectory solving (16) remains in $\mathcal{C}$. The right-hand side of (16) is the projection of the psuedogradient on the manifold formed by the active constraints at x [22].

We can rewrite the dynamics in a compact form as follows. Let $H(x) = [Dh_1 \ Dh_2]$ where $h_j(x) = [h_{1,j} \ \cdots \ h_{n,j}]^T$ for $j \in \{1, 2\}$ and D is the Jacobian operator. Also, let $\mu = [\mu_{1,1} \ \cdots \ \mu_{n,1} \ \mu_{1,2} \ \cdots \ \mu_{n,2}]^T$. Define $F(x, \mu) = \omega(x) + H(x)\mu$. Then, the dynamics can be written as

$$\dot{x} = F(x, \mu), \ \mu \in U(x) \tag{17}$$

where

$$U(x) = \left\{ \mu \,\middle|\, \|F(x, \mu)\| = \min_{\substack{\nu_j \geq 0, j \in J(x) \\ \nu_j = 0, j \notin J(x)}} \|F(x, \nu)\| \right\} \tag{18}$$

and $J(x) = \{j| \ h_j(x) \leq 0\}$. This formulation is given in the seminal work by Rosen [22] along with the theorem that states that for any initial condition in $\mathcal{C}$, a continuous solution $x(t)$ to (17) exists such that $x(t) \in \mathcal{C}$ for all $t > 0$. Thus, we have the following results.

Proposition 1 (Theorem 8 [22]). *The dynamical system* (17) *is asymptotically stable on* $\mathcal{C}$ *if* $D\omega(x)$ *has eigenvalues in the open left-half plane for* $x \in \mathcal{C}$.

Further, if $x^* \in \mathcal{C}$ is a differential Nash equilibrium, we can linearize ω around x^* and get the following sufficient condition guaranteeing x^* attracts nearby strategies under the gradient flow $F(x, \mu)$.

Proposition 2. *If $x^* \in \mathcal{C}$ is a differential Nash equilibrium, i.e. $\omega(x^*) = 0$ and $D_{ii}\ f_i(x^*) < 0$, and the eigenvalues of $D\omega(x^*)$ are in the open left-half plane, then x^* is an exponentially stable fixed point of the continuous-time dynamical system* (16).

Note that since in our estimation, we restrict $\theta_i \geq 0$, the f_i will be concave; hence, Nash equilibria of the game will be differential Nash equilibria.

These results imply that we can simulate the dynamical system in (17) in order to compute Nash equilibria of the game. Using a forward Euler discretization scheme and a subgradient projection method, we can compute Nash equilibria of the constrained game. The subgradient projection method is known to converge to the unique Nash equilibrium of the constrained n-person concave game [8].

3 Experimental Set-Up

The social game for energy savings that we have designed is such that occupants in an office builing vote according to their usage preferences of shared resources and are rewarded with points based on how *energy efficient* their strategy is in comparison with the other occupants. Having points increases the likelihood of the occupant winning in a lottery. The prizes in the lottery consist of three Amazon gift cards.

We have installed a Lutron[2] system for the control of the lights in the office. This system allows us to precisely control the lighting level of each of the lights in the office. We use it to set the default lighting level as well as implement the average of the votes each time the occupants change their lighting preferences.

We have divided the office into five lighting zones and each zone has four occupants. Thus, there are 20 occupants who participate in the social game. In addition, we have two heating, ventalating and air conditioning (HVAC) zones and each zone has ten occupants (see figure 1(a)).

We have developed an online platform in which the occupants can login and participate in the game. This includes the ability for the occupants to vote on their lighting and heating, ventilating and air conditioning (HVAC) preferences as well as view all occupant point balances and all occupant consumption patterns including the ability to monitor individual occupant plug-load consumption. Figure 1(b) shows a display of how an occupant can select their lighting preference and Figure 1(c) shows a sample of how occupants can see their point balance.

In this paper, we focus on a game focused on encouraging occupants to select lower lighting settings in exchange for a chance to win in a lottery. An occupant's vote is for the lighting level in their zone as well as for neighboring zones. The lighting setting that is implemented is the average of all the votes.

Each day when an occupant logs into the online platform the first time after they enter the office, they are considered present for the remainder of the day. There is a default lighting setting. An occupant can leave the lighting setting

[2] `http://www.lutron.com/en-US/Pages/default.aspx`

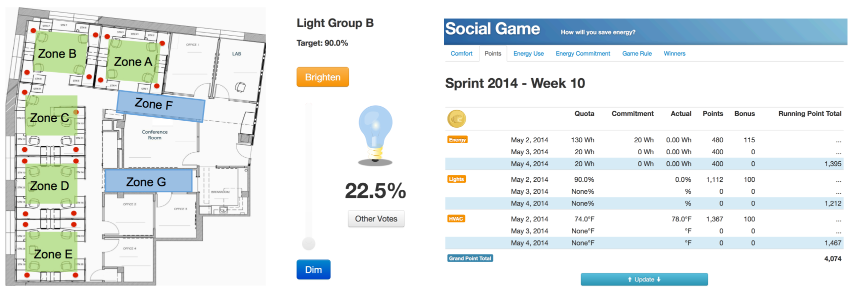

(a) (b) (c)

Fig. 1. (a) Display of HVAC and lighting zones. Zones A-B are the five lighting zones and zones F-G are the two HVAC zones. (b) Display of how occupants can log their lighting vote. (c) Display of an occupant's point balance.

as the default after logging in or they can change it to some other value in the interval $[0, 100]$ depending on their preferences.

Some of the energy savings we acheive is due to the default setting and some due to the social game. We are currently conducting experiments to determine how much savings is due to the social game. It is the building managers duty to ensure that the occupants are satisfied (via appropriate lighting level) and the building is operating in an energy efficient way. We believe that through optimal design of the incentives, we will be able to achieve greater energy savings than would be possible by only adjusting the default lighting setting. We leave this for future work.

4 Results

In this section, we report the results on the savings acheived through the game, the utility learning problem as well as simulation of the estimated utilities.

We use the data collected over the period from Mar. 3, 2014 to Jun. 5, 2014 when occupants have regular working schedules in the office. The baseline lighting, x_b, is 90%, which is the standard lighting level prior to the beginning of the experiment. Throughout this period, we have changed the default lighting level three times (see Table 1). We divide each day into four regions based on the outside lighting in Berkeley during the summer, namely from 5 to 10am (Dawn), 10am to 5pm (Daylight), 5pm to 8pm (Dusk), and 8pm to the next day 5am (Night). The data is further processed by taking the average of votes in each region of the day for each user.

4.1 Savings

First, we highlight the savings that is achieved as a result of instituting the social game. In Figure 2, we report the savings per day in KWh for the four

Period	Default Level
March 3–April 10	20 %
April 11–May 1	10 %
May 2–May 23	60 %
May 24 – June 5	90 %

Table 1. Default levels for four periods during the experiment. By changing the default setting to 90% we isolate the savings due to the social game from those achieved by changing the default setting.

periods in which the default varied. We remark that in the last period in which the default setting was set to 90% (which is the baseline line setting), we still acheived a savings of 3KWh on average per day. Using the mean savings for each of the periods and a rate of \$0.12/KWh, we estimate that we saved \$73. In addition, over the period of 101 days that the experiment was conducted the office consumed 2,185 KWh for lighting and we saved approximately 601KWh. That is a 27.5% reduction in energy. This savings is just due to a change in lighting usage behavior for one small portion of a building.

Our platform has the capability of including HVAC and plug-load in addition to lighting. We plan to implement a similar social game in Singapore and we expect much greater savings. This current experiment shows that a social game is a viable way to engage building occupants and induce behavioral change toward more energy efficient behaviors.

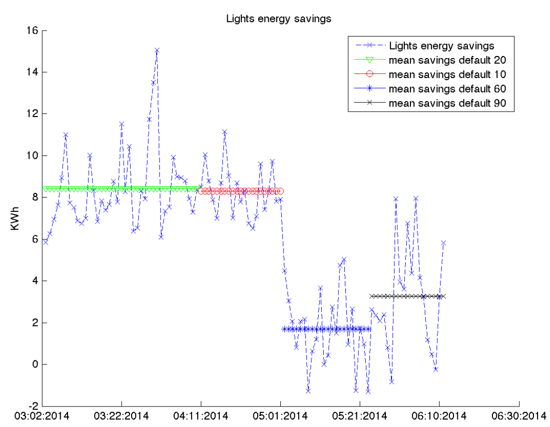

Fig. 2. Savings acheived per day (KWh). The mean savings over the four periods in which the default varied. Notice that in the period during which the default setting was at the baseline, there is still a savings of around 3KWh per day.

4.2 Estimation

The estimation proposed is Section 2.1 is performed for each user in each day interval and default lighting interval. Only true votes, not the default votes, are considered. We apply the bootstrapping method to obtain the empirical distribution of θ_i for $i \in \{1, ..., 20\}$ by randomly sampling a subset from the data [7]. The mean and standard deviation for the users which are the most active are reported in Table 2.

Table 2. Estimated utility parameter for selected set of active users. A, B, C, D stand for the periods Dawn, Daylight, Dusk, and Night respectively. The standard deviation is indicated inside the parentheses and the mean is given outside of the parentheses. 'NaN' indicates that the occupant did not vote during that period. Hence, they have no estimated utility. Occupants' whose number is in boldface have won at least once in the lottery.

		Active users (selected)					
		2	**6**	**8**	10	**14**	20
Default 20	A	.26(.25)	4657(481)	3671(126)	4658(0)	3054(141)	3857(110)
	B	.41(.09)	2932(99)	3386(73)	271(307)	3350(96)	691(473)
	C	.00(.01)	1808(1008)	3290(194)	1220(0)	3332(164)	1222(259)
	D	.00(.00)	822(1465)	3700(461)	3420(295)	3756(575)	1095(285)
Default 10	A	.96(.39)	294(759)	2923(215)	2446(0)	2971(508)	195(335)
	B	.24(.60)	833(796)	2847(320)	2042(0)	3219(339)	258(339)
	C	.07(.12)	0(0)	2924(224)	3485(0)	670(441)	643(469)
	D	.09(.25)	625(816)	3542(474)	3305(0)	1793(1187)	824(534)
Default 60	A	.28(.59)	469(1717)	6790(1267)	NaN	3180(827)	504(940)
	B	.07(.19)	1062(1135)	5741(734)	6327(199)	6180(881)	104(484)
	C	.00(.00)	1146(1927)	6166(502)	3752(0)	7856(1728)	588(903)
	D	.12(.18)	3947(2434)	6670(0)	5296(0)	3628(3394)	881(4)
Default 90	A	.01(.01)	9045(1562)	7835(2465)	NaN	NaN	3333(0)
	B	.00(.01)	7624(1699)	9479(926)	NaN	NaN	1923(2010)
	C	.02(.03)	8962(947)	8761(983)	NaN	NaN	3333(0)
	D	NaN	NaN	5000(461)	5000(0)	NaN	NaN

We remark that occupant 2 has a very low mean for the parameter θ_2 as compared to the other active occupants. By examining the ground truth values (red dots) in Figure 3, we see that occupant 2 often votes for a lighting setting around 60-70%. On the other hand, in Figure 4, we can see the ground truth of occupant 14 who often votes for a lighting setting of 0%. This player is more aggressive than occupant 2 and this behavior is reflected in the mean of the estimate for parameter θ_{14}.

4.3 Simulation

To capture the working schedules of each user, we employ a simple probabilistic model which determines the probability of individual user being absent, p_i^{absent},

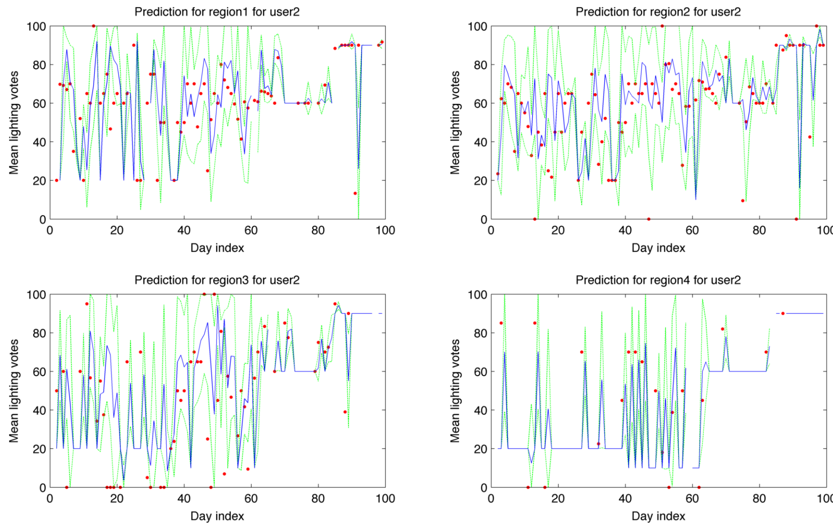

Fig. 3. One day ahead prediction by the Nash equilibrium algorithm for occupant 2. Red: ground truth, i.e. occupant 2's actual votes. We sample from the distributions across the events *absent*, *active*, *default* for each occupant and simulate the game with the set of *active* and *default* players. We repeat this 20 times for each day and generate a distribution for the predictions of each occupant's behavior. Blue: mean of prediction. Green: prediction within 1 standard deviation of the prediction mean. Gaps in the plots indicate that the occupant was not present on that day.

present and playing default, $p_i^{\text{present, default}}$, and present and actively playing, $p_i^{\text{present, active}}$. By assumption the sample space Ω includes the above three outcomes, and the probability mass functions should sum to unity. This probability is estimated by $p_i^E = \frac{N_{i,E}}{N_i}$, where E is the event of one of the three outcomes, $N_{i,E}$ is the number of event E for user i, and N_i is the number of total events.

For the prediction of the next day lighting votes, we randomly sample from this distribution to determine the set of active, default, and absent users, then obtain a local Nash equilibrium for them. This step is performed 20 times for each day to predict the distribution of votes, as shown in Figure 5.

As can be seen, the Nash equilibrium captures substantial variations in the data. We also compared the results of prediction with the autoregressive integrated moving average (ARIMA) model [9], constant model which uses the default lighting for prediction, and the persistent model which uses the previous day value for prediction. The mean squared errors (MSE) of the models are summarized in Table 3. The Nash equilibrium achieves a prediction that is the most accurate as compared with other models, which presents it favorably for leaders in the Stackelberg game to design optimal incentives to motivate energy saving behaviors. Indeed in the Stackelberg framework, the leader (building manager) assumes that the agents (occupants) are utility maximizers and play

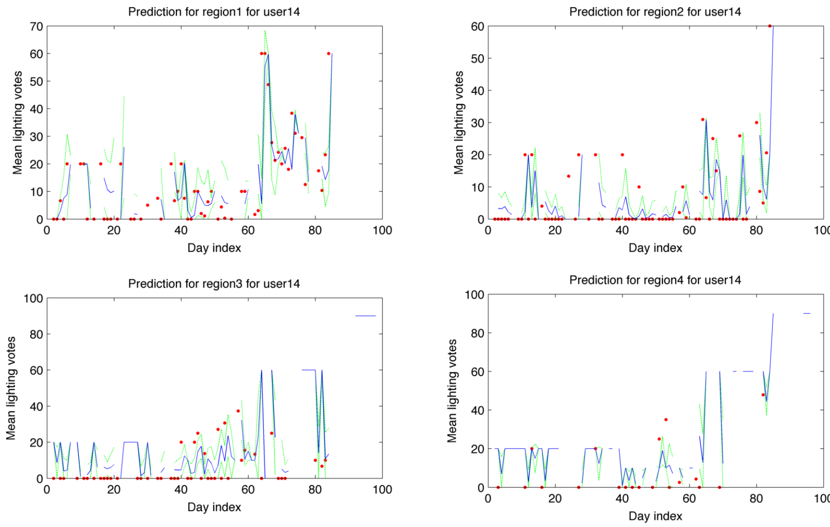

Fig. 4. One day ahead prediction by the Nash equilibrium algorithm for occupant 14. Red: ground truth, i.e. occupant 14's actual votes. We sample from the distributions across the events *absent*, *active*, *default* for each occupant and simulate the game with the set of *active* and *default* players. We repeat this 20 times for each day and generate a distribution for the predictions of each occupant's behavior. Blue: mean of prediction. Green: prediction within 1 standard deviation of the prediction mean. Gaps in the plots indicate that the occupant was not present on that day.

Nash. Hence, we will be able to integrate our estimation algorithm into an online algorithm for designing incentives.

Table 3. Mean square error (MSE) of four algorithms that predict the one day ahead occupant behavior over the period of study (101 days): ARIMA(1,0,1) (we use one autoregressive term, zero nonseasonal differences, and one lagged forecast error in the prediction equation) [9], Nash, a model which uses the default as the prediction, and a model which uses occupants' previous votes as the prediction. Nash out performs each of the other methods.

Model	ARIMA(1,0,1)	Nash	Constant	Persistent
MSE	13.9183	**12.4561**	16.9585	13.4247

5 Discussion and Future Work

We have designed and implemented a social game for inducing building occupants to behave in an energy efficient manner. We presented data and results pertaining to the game in which occupants select their lighting preferences and

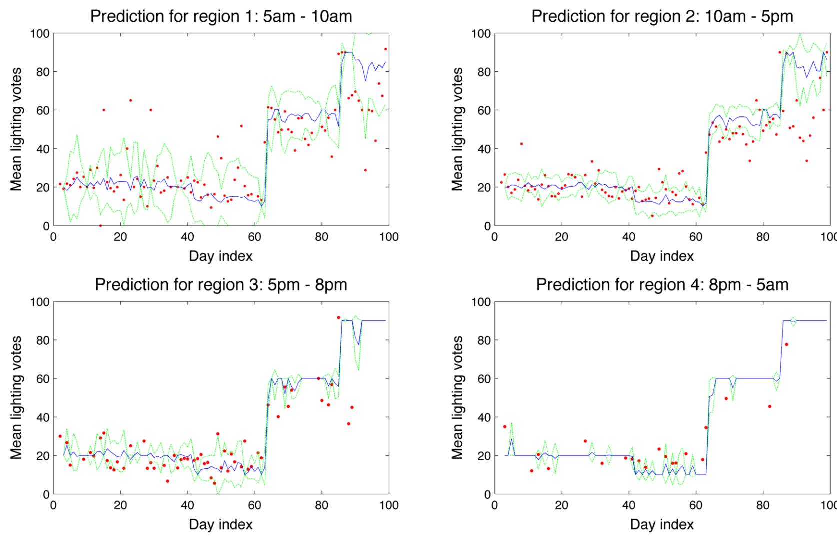

Fig. 5. One day ahead prediction by the Nash equilibrium algorithm for the average implemented lighting setting in the office each day. We sampled from the distribution over the events *absent*, *active*, *default* for each of the players and simulated Nash given the sample of *active* and *default* players. We repeated this 20 times for each day and generated a distribution for the prediction of the implemented lighting setting. Red: ground truth, i.e. the average lighting setting that is implemented per day. Blue: mean of prediction as given by the Nash simulation. Green: prediction within 1 standard deviation of the prediction mean.

win points depending on how far their vote is from the baseline lighting setting and proportional to other occupants' votes distances from the baseline. As a result, the occupants are interacting in a competitive environment which we model as a non-cooperative game. We show that we get significant savings as compared to usage prior to the implementation of the social game. This savings is due to both a change in the default setting as well as due to the incentives offered in the social game.

We described the experimental set-up which includes an online platform for the implementation of the social game as well as the use of a Lutron lighting system for percise control of the lighting setting. Our platform also includes the ability to implement a social game centered around HVAC settings as well as occupant plug-load consumption. We leave exploring these additional features as future work.

We have formulated the problem of estimating the occupant utility functions as a convex optimization problem and estimated occupant utilities in a 20 player social game. We simulated the game using the estimated utility functions and showed that our model is a good predictor for occupant behavior. It out performs a number of other estimation techniques including ARIMA.

There are several ways in which we believe we can improve our estimate of the utility functions of the occupants. We did not consider the environmental noise such as variations in natural light. We instead used a heuristic to capture this variation by breaking the day into intervals in which the natural light entering the office is most consistent. In addition, we did not consider any information on the occupants' schedules or location in the office with respect to windows. We could incorporte these aspects into our estimation as priors on the parameters of the occupants utility function or as a noise process in the estimated behavior model. We leave this as future work.

In the experiments used for this paper, we selected the value of ρ based on heuristics. Our goal is to design ρ in an optimal way. We can leverage the fact that we have modeled occupants as utility maximizers who play in a non-cooperative game by considering the design of ρ by the building manager. In particular, we can model this interaction between the building manager and the occupants as a Stackelberg game. In this framework, the building manager would perform an online estimation of the occupants' utility functions and update ρ accordingly [20]. We believe that by optimizing the incentive ρ, we can acheive greater savings. We are currently implementing such a scheme in our experimental platform.

References

1. Albadi, M., El-Saadany, E.: A summary of demand response in electricity markets. Electric Power Systems Research 78(11), 1989 – 1996 (2008)
2. Arrow, K.J., Hurwicz, L., Uzawa, H.: Constraint qualifications in maximization problems. Naval Research Logistics Quarterly 8(2), 175–191 (1961)
3. Aswani, A., Master, N., Taneja, J., Smith, V., Krioukov, A., Culler, D., Tomlin, C.: Identifying models of HVAC systems using semi-parametric regression. In: Proceedings of the American Control Conference (2012)
4. Bestick, A., Ratliff, L.J., Yan, P., Bajcsy, R., Sastry, S.S.: An inverse correlated equilibrium framework for utility learning in multiplayer, noncooperative settings. In: Proceedings of the 2nd ACM international conference on High confidence networked systems. pp. 9–16. ACM (2013)
5. Boman, M., Davidsson, P., Skarmeas, N., Clark, K., Gustavsson, R.: Energy saving and added customer value in intelligent buildings. In: Third International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology. pp. 505–517 (1998)
6. Bourgeois, D., Reinhart, C., Macdonald, I.: Adding advanced behavioural models in whole building energy simulation: A study on the total energy impact of manual and automated lighting control. Energy and Buildings 38(7), 814 – 823 (2006), special Issue on Daylighting Buildings
7. Efron, B., Tibshirani, R.J.: An introduction to the bootstrap, vol. 57. CRC press (1994)
8. Flåm, S.D.: Solving non-cooperative games by continuous subgradient projection methods. In: System Modelling and Optimization, pp. 115–123. Springer (1990)
9. Hannan, E.J., Deistler, M.: The statistical theory of linear systems, vol. 70. SIAM (2012)

10. Keshavarz, A., Wang, Y., Boyd, S.: Imputing a convex objective function. In: IEEE International Symposium on Intelligent Control. pp. 613–619. IEEE (2011)
11. Lee, M., O.Aslam, Foster, B., Kathan, D., Kwok, J., Medearis, L., Palmer, R., Sporborg, P., Tita, M.: Assessment of demand response and advanced metering. Tech. rep., Federal Energy Regulatory Commission (2013)
12. Lovett, T., Gabe-Thomas, E., Natarajan, S., O'Neill, E., Padget, J.: 'just enough' sensing to enliten: A preliminary demonstration of sensing strategy for the 'energy literacy through an intelligent home energy advisor' (enliten) project. In: Proceedings of the Fourth International Conference on Future Energy Systems. pp. 279–280. e-Energy '13, ACM, New York, NY, USA (2013), `http://doi.acm.org/10.1145/2487166.2487206`
13. Ma, Y., Anderson, G., Borrelli, F.: A distributed predictive control approach to building temperature regulation. In: American Control Conference (ACC). pp. 2089–2094. IEEE (2011)
14. Mathieu, J., Dyson, M., Callaway, D., Rosenfeld, A.: Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations. In: ACEEE Summer Study on Energy Efficiency in Buildings (2012)
15. Mathieu, J.L., Price, P.N., Kiliccote, S., Piette, M.A.: Quantifying changes in building electricity use, with application to demand response. IEEE Transactions on Smart Grid, 2(3), 507–518 (2011)
16. McQuade, J.: A system approach to high performance buildings. United Technologies Corporation, Tech. Rep (2009)
17. Merugu, D., Prabhakar, B.S., Rama, N.: An incentive mechanism for decongesting the roads: A pilot program in bangalore. In: Proceedings of ACM NetEcon Workshop (2009)
18. Oldewurtel, F., Parisio, A., Jones, C., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Wirth, K.: Energy efficient building climate control using stochastic model predictive control and weather predictions. In: American Control Conference (ACC), 2010. pp. 5100–5105. IEEE (2010)
19. Ratliff, L.J., Burden, S.A., Sastry, S.S.: Characterization and computation of local nash equilibria in continuous games. In: Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (2013)
20. Ratliff, L.J., Dong, R., Ohlsson, H., Sastry, S.S.: Incentive design and utility learning via energy disaggregation. In: Proceedings of the 19th World Congress of the International Federation of Automatic Control (2014)
21. Roozbehani, M., Dahleh, M., Mitter, S.: Dynamic pricing and stabilization of supply and demand in modern electric power grids. In: First IEEE International Conference on Smart Grid Communications (SmartGridComm),. pp. 543 –548 (oct 2010)
22. Rosen, J.B.: Existence and uniqueness of equilibrium points for concave n-person games. Econometrica 33(3), 520 (1965)
23. Swenson, A.: Lighting in commercial buildings. Tech. rep., U.S. Energy Infomation Administration (2009), `http://www.eia.gov/consumption/commercial/data/archive/cbecs/cbecs2003/lighting/lighting1.html`