

A Robust Utility Learning Framework via Inverse Optimization

Ioannis C. Konstantakopoulos, Lillian J. Ratliff *Member, IEEE*, Ming Jin,
S. Shankar Sastry *Fellow, IEEE*, and Costas J. Spanos *Fellow, IEEE*

Abstract—In many smart infrastructure applications, flexibility in achieving sustainability goals can be gained by engaging end users. However, these users often have heterogeneous preferences that are unknown to the decision maker tasked with improving operational efficiency. Modeling user interaction as a continuous game between noncooperative players, we propose a robust parametric utility learning framework that employs constrained feasible generalized least squares estimation with heteroskedastic inference. To improve forecasting performance, we extend the robust utility learning scheme by employing bootstrapping with bagging, bumping, and gradient boosting ensemble methods. Moreover, we estimate the noise covariance, which provides approximated correlations between players, which we leverage to develop a novel *correlated utility learning* framework. We apply the proposed methods both to a toy example arising from Bertrand–Nash competition between two firms and to data from a social game experiment designed to encourage energy efficient behavior among smart building occupants. Using occupant voting data for shared resources such as lighting, we simulate the game defined by the estimated utility functions to demonstrate the performance of the proposed methods.

Index Terms—Game theory, inverse optimization, smart building energy efficiency.

I. INTRODUCTION

DUE TO pervasive utilization of Internet of Things and cyber-physical systems sensing/actuating platforms, we are increasingly observing human decision makers being integrated into operational and managerial decisions in infrastructure systems. Their actions can be leveraged to increase both resilience and sustainability thereby making smart infrastructure a worthwhile investment. Smart buildings, being no exception, are a fundamental component of emerging *smart cities*; their efficient design and operation enables

Manuscript received December 28, 2016; accepted April 7, 2017. Manuscript received in final form April 25, 2017. Recommended by Associate Editor R. N. Shorten. This work was supported by the Republic of Singapore’s National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore for the Singapore–Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. The work of I. C. Konstantakopoulos was supported by a scholarship of the Alexander S. Onassis Public Benefit Foundation. (Ioannis C. Konstantakopoulos and Lillian J. Ratliff contributed equally.) (*Corresponding author: Ioannis C. Konstantakopoulos.*)

I. Konstantakopoulos, M. Jin, S. Sastry, and C. Spanos are with the Electrical Engineering and Computer Sciences Department, University of California, Berkeley, Berkeley, CA 94720 USA (e-mail: ioanniskon@eecs.berkeley.edu; jinming@eecs.berkeley.edu; spanos@eecs.berkeley.edu; sastry@eecs.berkeley.edu).

L. Ratliff is with the Electrical Engineering Department, University of Washington, Seattle, WA 98195 USA (e-mail: ratliff@uw.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCST.2017.2699163

flexibility, e.g., by automatically shifting or curtailing demand during peak hours, in making urban spaces sustainable. More abstractly, in many infrastructure systems, there is often an entity acting as a *planner* (e.g., facility managers and departments of transportation) that introduces incentives or control policies to coordinate autonomously acting agents in the system (e.g., selfish human decision makers) so that their collective behavior leads to system-level efficiency gains.

One approach to designing such policies is to leverage game-theoretic models of decision making in an optimization framework to produce policies that encourage or induce behavior that optimizes an objective [1], [2]. Often the planner has at best a prior on the decision making model of the individual agents. Such information asymmetries lead to inefficiencies [1], [3]. In this paper, we propose a framework for estimating decision making models of self-interested decision makers consuming a shared resource (e.g., lighting in a smart building) that can be leveraged in control or incentive design to aid in closing the efficiency gap.

To concretize ideas, consider a smart building—an example we will return to throughout the text. A facilities manager may be incentivized or even tasked to encourage energy efficient behavior if they are accountable for energy costs or are required, e.g., to maintain an operational excellence measure (see [4]). At the same time, the facilities manager generally must also ensure user comfort and productivity. Beyond these motivations, demand response (DR) programs are being rolled out by utility companies and third-party solution providers with the goal of correcting for improper load forecasting. Participating consumers decide to change their consumption when DR events are called. The facilities manager may be required to keep this schedule.

Smart building technologies enable new avenues for facilities managers to keep such a prescribed schedule via automation or integrating the end user. Yet, in office buildings, the occupants, as employees, are not typically responsible for paying for the energy resources they consume. Hence, there is often a misalignment between the incentives of the facilities manager and the occupants. Social games are a means to engage the occupants to address these inefficiencies. In Section VI, we describe one such social game that we designed and implemented on the UC Berkeley campus, aimed at incentivizing energy efficient consumption of shared resources by leveraging building automation.

The broader purpose of this paper is to present a general framework that leverages game-theoretic concepts to learn models of players’ decision making in competitive

environments such as the building energy social game described above. The framework supports learning agents' preferences over shared resources as well as understanding how preferences change as a function of external stimuli such as physical control or incentives. Such a framework can be used in the design of incentive mechanisms that realign agents' preferences with those of the planner, which often represent system-level performance criteria, through fair compensation.

More concretely, we model decision making agents as *utility maximizers*, and using inverse optimization and game-theoretic techniques, we derive a robust scheme to infer their utility functions. At the core of our approach is the fact that we model the agents as noncooperative players in a game playing according to a *Nash equilibrium strategy*. From this point of view, agents are strategic entities that make decisions based on their own preferences despite others. The game-theoretic framework both allows for qualitative insights to be made about the outcome of such a selfish behavior—more so than a simple prescriptive model—and, more importantly, can be leveraged in designing mechanisms for incentivizing agents.

We assume a parametric form of utility function for each player that is dependent on the decisions of others. Correlations between players' decisions are not known *a priori*. Assuming observations are approximately Nash equilibria, we use first- and second-order conditions on player utility functions to construct a constrained regression model. The result is as a constrained generalized least squares (cGLS) problem with nonspherical noise error terms. Using constrained feasible generalized least squares (cFGLS), an implementable version of cGLS, we utilize heteroskedastic inference to approximate the correlated errors.

Noting that data sets of observed decisions often may be small relative to the number of model parameters in practice, we employ bootstrapping to generate pseudodata from which we learn additional estimators. The bootstrapping process allows us to derive an asymptotic approximation of the bias and standard error of an estimator. We utilize ensemble methods such as bagging, bumping, and gradient boosting to extract an estimator from the pseudodata generated estimators that results in a reduced forecasting error. The ensemble methods are robust under noise and autocorrelated error terms. We apply the robust utility learning framework to a model of Bertrand–Nash competition between firms in order to illustrate the framework and its performance.

Building on the robust utility learning framework, we use the approximated standard error to derive an innovative utility learning method in which we modify players' utility functions to create a *correlated game*. The resulting correlated utility learning method leverages correlations between players and the ensemble estimators to minimize the estimation error by optimizing scaling coefficients that appear in the correlated game utility functions. Applying this method results in a significant improvement over the constrained ordinary least squares (cOLS) estimations and outperforms many of the ensemble methods. It also provides insights into how players interact with one another and indicates which players are potentially forming coalitions. Moreover, this technique is amenable to online implementation after an initial training

period so that using cOLS estimators in the correlated utility learning framework, our adaptive incentive design schemes, introduced in [3] and [9], can be made robust.

To demonstrate the efficacy of both the robust and correlated utility learning frameworks, we apply them to data generated from the smart building social game experiment we conducted. We show that estimating the players' utility functions via the proposed methods results in a predictive model that outperforms several other standard techniques such as ordinary least squares (OLS).

The rest of this paper is organized as follows. We describe the abstracted game framework for modeling the interaction of agents as well as define equilibrium concepts in Section II. In Section III, we formulate the robust utility learning framework and provide an algorithm for implementing it. Section IV contains the Bertrand–Nash competition example, and we present the correlated utility learning framework in Section V. In Section VI, we describe the social game experimental setup on the UC Berkeley campus within the CREST Center,¹ provide a brief literature review, and present the results of both proposed utility learning methods applied to data from the social game. We make concluding remarks and discuss future directions in Section VIII.

II. GAME FRAMEWORK

In this section, we abstract the agents' decision making processes in a game-theoretic framework.

A. Agent Decision Making Model

Consider p agents,² i.e., decision making entities, indexed by the set $\mathcal{I} = \{1, \dots, p\}$. Each agent is modeled as a *utility maximizer* that seeks to select $x_i \in \mathbb{R}$ by optimizing

$$f_i(x_i, x_{-i}) = f_i^{\text{nom}}(x_i, x_{-i}) + f_i^{\text{inc}}(x_i, x_{-i}) \quad (1)$$

where $f_i^{\text{nom}}(x_i, x_{-i})$ and $f_i^{\text{inc}}(x_i, x_{-i})$ are the nominal and incentive components, respectively, of agent i 's utility function and where $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \in \mathbb{R}^{n-1}$ is the collective choices of all agents excluding the i th agent.³

The choice x_i abstracts the agent's decision; it could represent how much of a particular resource they choose to consume. The nominal component of f_i captures the agent's individual preferences over x_i and may depend on the decisions of others x_{-i} . The incentive component models the portion of the agent's utility that can be designed by the planner; it also may depend on the decisions of other agents.

Agent i 's optimization problem is also subject to constraints; the constraint set is given by $\mathcal{C}_i = \{x_i \mid h_{i,j}(x_i) \geq 0, j = 1, \dots, \ell_i\}$, where each $h_{i,j}$ is assumed to be a concave function of x_i . Such constraints may encode cyber or physical constraints arising from the underlying system—in the social game example presented in Section VI-C, we will see that

¹<http://crest.berkeley.edu/>.

²We refer to the decision makers as *agents* and use the term interchangeably with *players*.

³Note that while for notational simplicity we assume that $x_i \in \mathbb{R}$, the work easily extends to a higher dimensional choice vector for each agent.

these constraints are physical bounds. Thus, given x_{-i} , agent i faces the following optimization problem:

$$\max\{f_i(x_i, x_{-i}) \mid x_i \in \mathcal{C}_i\}. \quad (2)$$

B. Game Formulation

The game (f_1, \dots, f_p) is a continuous game on a convex strategy space $\mathcal{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_p$. To model the outcome of the strategic interactions of agents, we use the *Nash equilibrium* concept.

Definition 1: A point $x \in \mathcal{C}$ is a Nash equilibrium for the game (f_1, \dots, f_p) on \mathcal{C} if, for each $i \in \mathcal{I}$

$$f_i(x_i, x_{-i}) \geq f_i(x'_i, x_{-i}) \quad \forall x'_i \in \mathcal{C}_i. \quad (3)$$

We say $x \in \mathcal{C}$ is an ε -Nash equilibrium for $\varepsilon > 0$ if the above inequality is relaxed

$$f_i(x_i, x_{-i}) + \varepsilon \geq f_i(x'_i, x_{-i}) \quad \forall x'_i \in \mathcal{C}_i. \quad (4)$$

We say a point is a *local Nash equilibrium* (a ε -local Nash equilibrium) if there exists $W_i \subset \mathcal{C}_i$ such that $x_i \in W_i$ and the above inequalities hold for all $x'_i \in W_i$.

If each f_i is concave in x_i and \mathcal{C} is convex, then the game is a p -person concave game. In [6], it was shown that a (pure) Nash equilibrium exists for every concave game.

The Lagrangian of agent i 's optimization problem is given by

$$L_i(x_i, x_{-i}, \mu_i) = f_i(x_i, x_{-i}) + \sum_{j \in \mathcal{A}_i(x_i)} \mu_{i,j} h_{i,j}(x_i) \quad (5)$$

where $\mathcal{A}_i(x_i)$ is the active constraint set at x_i and $\mu = (\mu_1, \dots, \mu_p)$ with $\mu_i = (\mu_{i,j})_{j=1}^{\ell_i}$ are the Lagrange multipliers. Assuming appropriate smoothness conditions on each f_i and $h_{i,j}$, the differential game form [3], [7], which characterizes the first-order conditions of the game, is given by

$$\omega(x, \mu) = [D_1 L_1(x, \mu_1)^\top \cdots D_p L_p(x, \mu_p)^\top]^\top \quad (6)$$

where $D_i L_i$ denotes the derivative of L_i with respect to x_i .

Consider agent i 's optimization problem (2) with x_{-i} fixed and where each f_i and $h_{i,j}$ for $j \in \{1, \dots, \ell_i\}$, $i \in \mathcal{I}$, are concave, twice continuously differentiable functions. Then, assuming an appropriate constraint qualification condition [8], the necessary and sufficient conditions for optimality of a point x_i are as follows: there exists $\mu_i \in \mathbb{R}_+^{\ell_i}$ such that: 1) $D_i L_i(x, \mu_i) = 0$; 2) $\mu_i h_{i,j}(x_i) = 0$ for each $j \in \{1, \dots, \ell_i\}$; and 3) $h_{i,j}(x_i) \geq 0$ for each $j \in \{1, \dots, \ell_i\}$. Regardless of the concavity assumption, the point x_i is a local maximizer if $\mu_{i,j} > 0$ and $z^\top D_{ii}^2 L_i(x, \mu_i) z < 0$ for all $z \neq 0$ such that $D_i h_{i,j}(x_i)^\top z = 0$ for $j \in \mathcal{A}_i(x_i)$. Such conditions motivate the following definition.

Definition 2 (Differential Nash Equilibrium): Consider a game (f_1, \dots, f_p) on \mathcal{C} where f_i and $h_{i,j}$ for each $j \in \{1, \dots, \ell_i\}$ and $i \in \mathcal{I}$ are twice continuously differentiable. A point $x \in \mathcal{C} \subset \mathbb{R}^p$ is a differential Nash equilibrium if there is a $\mu \in \mathbb{R}^{\sum_{i=1}^p \ell_i}$ such that the pair (x, μ) Satisfies: 1) $\omega(x, \mu) = 0$; 2) for each $i \in \mathcal{I}$, $z^\top D_{ii} L_i(x, \mu_i) z < 0$ for all $z \neq 0$ such that $D_i h_{i,j}(x_i)^\top z = 0$ and $\mu_{i,j} > 0$ for $j \in \mathcal{A}_i(x_i)$. If, for a given $\varepsilon > 0$, (i') $\omega(x, \mu) = \varepsilon$ with all

the other conditions being satisfied, then x is a ε -differential Nash equilibrium.

The above definition extends the definition of a differential Nash (if we restrict to Euclidean spaces), first appearing in [7], to constrained games on Euclidean spaces. Using this definition, we can also extend [7, Proposition 1], again where strategy spaces are restricted to be subsets Euclidean.

Proposition 1: A differential Nash equilibrium of the p -person concave game (f_1, \dots, f_p) on \mathcal{C} is a local Nash equilibrium.

The proof is straightforward and we leave it to the Appendix. The proposition says that the conditions of Definition 2 are sufficient for a local Nash. In contrast to single-agent optimization problems, for games, the second-order conditions do not imply the equilibrium is isolated [7]. A sufficient condition guaranteeing that a Nash equilibrium x is isolated is that the Jacobian of $\omega(x, \mu)$, denoted $D\omega(x, \mu)$, is invertible [3].

We use necessary and sufficient optimality conditions on individual player optimization problems holding other players' strategies fixed to formulate the utility learning framework.

III. ROBUST UTILITY LEARNING

In the previous works, we have explored utility learning and incentive design as a coupled problem both in theory [3] and in practice [9], [10]. In this paper, we reexamine the utility learning problem using statistical methods that serve to improve estimation and prediction accuracy.

Looking forward, our aim is to fold the new estimation scheme into the overall incentive design framework. This goal motivates why we are interested in learning more than a simple predictive model for agents, but rather a utility-based forecasting framework that accounts for individual preferences.

We parameterize f_i by $\theta_i = (\theta_{i,1}, \dots, \theta_{i,m_i}) \in \mathbb{R}^{m_i}$ and a finite set of basis functions $\{\phi_{ij}(x_i, x_{-i})\}_{j=1}^{m_i}$ such that

$$f_i(x; \theta_i) = \lambda \phi_i(x_i, x_{-i}), \quad \theta_i) + \bar{f}_i(x) \quad (7)$$

where $\phi_i = [\phi_{i,1} \cdots \phi_{i,m_i}]^\top$ and $\bar{f}_i(x)$ is a function that captures *a priori* knowledge of the agent's utility function (e.g., the incentive component designed by the planner).

A. Base Utility Estimation Framework

We start by describing the basic utility estimation framework using equilibrium conditions for the game played between the players. The utility learning framework we propose is quite broad in that it encompasses a wide class of continuous games. In [3], [12], and [13], we have shown that the utility learning problem can be formulated as a convex optimization problem by using first- and second-order conditions for Nash equilibria. Let us briefly review this formulation as it serves as the basis for the robust utility learning method.

Each observation $x^{(k)}$ is assumed to be an ε -approximate differential Nash equilibrium where the superscript notation $(\cdot)^{(k)}$ indicates the k th observation. For each observation $x^{(k)}$, it may be the case that only a subset of the players, say $\mathcal{S}^k \subset \mathcal{I}$ at observation k , participate in the game. Then notationally each observation is such that

$$x^{(k)} = (x_j^{(k)})_{j \in \mathcal{S}^k}. \quad (8)$$

If player i participates in n_i instances of the game, then there are n_i observations for that player. Let $n = \sum_{i=1}^p n_i$ be the total number of observations.

We can consider first-order optimality conditions for each player's optimization problem and define a residual function capturing the degree of suboptimality of $x_i^{(k)}$ [9], [11]. Indeed, for player i 's optimization problem, let the residual of the stationarity condition be given by

$$r_{s,i}^{(k)}(\theta_i, \mu_i) = D_i f_i(x_i^{(k)}, x_{-i}^{(k)}) + \sum_{j=1}^{\ell_i} \mu_i^j D_i h_{i,j}(x_i^{(k)}) \quad (9)$$

and the residual of the complementary conditions be given by

$$r_{c,i}^{j,(k)}(\mu) = \mu_i^j h_{i,j}(x_i^{(k)}), \quad j \in \{1, \dots, \ell_i\}. \quad (10)$$

Define

$$r_{c,i}^{(k)}(\mu_i) = [r_{c,i}^{1,(k)}(\mu_i) \cdots r_{c,i}^{\ell_i,(k)}(\mu_i)]. \quad (11)$$

Using data from the players' decisions (e.g., lighting votes from the social game experiment, which we describe in Section VI-A), the base utility learning framework consists of solving the optimization problem given by

$$\begin{aligned} \min_{\mu, \theta} \quad & \sum_{i=1}^p \sum_{k=1}^{n_i} \chi_i(r_{s,i}^{(k)}(\theta, \mu), r_{c,i}^{(k)}(\mu)) \\ \text{s.t.} \quad & \theta_i \in \Theta_i, \mu_i \geq 0 \quad \forall i \in \{1, \dots, p\} \end{aligned} \quad (\text{P})$$

where Θ_i is a constraint set on the parameters θ_i that captures prior information about the objective, $\chi : \mathbb{R}^p \times \mathbb{R}^{\sum_{i=1}^p \ell_i} \rightarrow \mathbb{R}_+$ is a non-negative, convex penalty function satisfying $\chi(z_1, z_2) = 0$ if and only if $z_1 = 0$ and $z_2 = 0$, i.e., any norm on $\mathbb{R}^p \times \mathbb{R}^{\sum_{i=1}^p \ell_i}$, and the inequality $\mu_i \geq 0$ is element wise.

The goal of this optimization problem, which is a finite dimensional optimization problem in the θ_i 's, is to find θ_i for each player such that $(\hat{f}_i)_{i \in \mathcal{I}}$ is consistent (or approximately consistent) with the data. As is noted in [11], we also remark that it is important that the sets Θ_i contain enough prior information about the objectives f_i in order to prevent trivial solutions. For example, if it is the case that $\hat{f}_i(x^{(k)}) = 0$ for each k and each $\Theta_i = \mathbb{R}^{m_i}$, then the trivial solution $\theta_i = \mathbf{0}_{m_i}$ is feasible. For many applications, some *a priori* knowledge on part of the utility functions of players may be encoded in each Θ_i (e.g., choosing Θ_i such that $\theta_{1i} = 1$ or similarly selecting the incentive component of the utility, a design possibility for the planner [3]) or through other normalization techniques to prevent such trivial solutions. In the context of the social game application (in Section VI-C), we explicitly discuss how to construct this constraint set in such a way that we ensure that the estimated utility functions are concave, which in turn guarantees that there exists a Nash equilibrium to the estimated game.

B. Robust Utility Learning

Let us now formulate a robust version of the utility learning framework that allows us to reduce our forecasting error and learn the noise structure that can be leveraged in extracting

pseudocoalitions between players, which we describe in the sequel. Define

$$X_i^{(k)} = \begin{bmatrix} D_i h_i(x_i^{(k)}) & D_i \phi_i(x^{(k)}) \\ \hat{h}_i(x_i^{(k)}) & \mathbf{0}_{\ell_i \times m_i} \end{bmatrix} \quad (12)$$

where

$$\hat{h}_i(x_i) = \text{diag}(h_{i,1}(x_i), \dots, h_{i,\ell_i}(x_i)) \quad (13)$$

$$D_i h_i(x_i) = [D_i h_{i,1}(x_i) \cdots D_i h_{i,\ell_i}(x_i)] \quad (14)$$

and $n_d = (\ell_i + 1)n$ is the total number of data points. The regressor matrix is then defined as $X = \text{diag}(X_1, \dots, X_p) \in \mathbb{R}^{n_d \times (\ell_i + 1)p}$, where $X_i = [(X_i^{(1)})^\top \cdots (X_i^{(n_i)})^\top]^\top$. Define the regression coefficient

$$\beta = [\mu_1^1 \cdots \mu_1^{\ell_1} \theta_1 \cdots \mu_p^1 \cdots \mu_p^{\ell_p} \theta_p]^\top \in \mathbb{R}^{(\ell_i + 1)p} \quad (15)$$

and the observation matrix $Y = [Y_1 \cdots Y_p]^\top \in \mathbb{R}^{(\ell_i + 1)p}$, where

$$Y_i = [\bar{f}_i(x^{(1)}) \mathbf{0}_{\ell_i} \cdots \bar{f}_i(x^{(n_i)}) \mathbf{0}_{\ell_i}]^\top. \quad (16)$$

Using the Euclidean norm for χ in (P) leads to an OLS problem with inequality constraints, i.e., a cOLS

$$\min_{\beta} \{\|Y - X\beta\|_2 \mid \beta \in \mathcal{B}\} \quad (\text{P1})$$

where $\mathcal{B} = \{\beta \mid \theta_i \in \Theta_i, \mu_i \geq 0, \forall i \in \mathcal{I}\}$. Enforcing that each of the constraint sets Θ_i is encoded by inequalities on θ_i , the above-stated problem can be viewed as a classical multiple linear regression model with inequality constraints described by the data generation process

$$Y = X\beta + \epsilon, \quad \beta \in \mathcal{B} \quad (17)$$

where $\epsilon = (\epsilon_1, \dots, \epsilon_p)$ is the error term satisfying: 1) $E(\epsilon|X) = 0^{n_d \times 1}$; 2) $\text{cov}(\epsilon|X) = \sigma^2 I^{n_d \times n_d}$; and 3) $\{\epsilon_i\}_{i=1}^p$ independent and identically distributed (i.i.d) with a zero mean and σ^2 variance. In addition, we assume that ϵ is nonspherical [12]. With this general statistical model, we are able to describe a data generation process in which the error terms are correlated or lack constant variance. This fact will be leveraged in creating coalitions between players, as we describe in Section V.

Mathematically, the nonspherical errors are modeled by

$$\text{cov}(\epsilon|X) = G \succ 0, \quad G \in \mathbb{R}^{n_d \times n_d}. \quad (18)$$

One drawback of this technique is that, given nonspherical standard errors, the cOLS estimator is biased, that is, it does not satisfy the best linear unbiased estimator (BLUE) property, a result of the Gauss–Markov theorem [12, Th. 1, Ch. 5]. However, we can derive an unbiased estimator by multiplying (17) on the left with $G^{-1/2}$. This leads to the cGLS statistical model given by

$$(G^{-1/2}Y) = (G^{-1/2}X)\beta + (G^{-1/2}\epsilon), \quad \beta \in \mathcal{B} \quad (19)$$

which now satisfies the BLUE property. In general, the explicit form of $\text{cov}(\epsilon|X) = G$ is unknown. We use the residuals (17) to infer the noise by imposing structural constraints on G .

We remark that there are many types of noise structures that can be used for imposing structure on G . We provide

two example noise structures that could be used. The first is the block diagonal structure [12, Ch. 5]; in particular, we impose that $G = \text{blkdiag}(K_1, \dots, K_p) \in \mathbb{R}^{n_d \times n_d}$ where $K_i = \text{blkdiag}(B_{i,1}, \dots, B_{i,n_i}) \in \mathbb{R}^{(\ell_i+1)n_i \times (\ell_i+1)n_i}$ with each $B_{i,k} \in \mathbb{R}^{(\ell_i+1) \times (\ell_i+1)}$. Estimating β with cOLS, we get $\hat{\beta}_{\text{cOLS}}$ with residual vector $e = Y - X\hat{\beta}_{\text{cOLS}} \in \mathbb{R}^{(\ell_i+1)n}$. The residual vector e can be decomposed into residuals for each player by writing $e = [e_1^\top \dots e_p^\top]^\top$. We use e_i to compute an estimate \hat{K}_i of K_i , which is in turn used to compute \hat{G} . The residuals come in triplets since for each k , $Y_i^{(k)} \in \mathbb{R}^{\ell_i+1}$. For ease of presentation and comprehension, we will use a paired index for the residuals instead of a single index. For example, for player i , there are n_i instances at which we have ℓ_i observations. Let $(e_i)_{k,j} = (e_i)_{(\ell_i+1)(k-1)+j}$, where $k \in \{1, \dots, n_i\}$ and $j \in \{1, \dots, (\ell_i+1)\}$. With the residuals, we can then form estimates $\hat{B}_{i,k} \in \mathbb{R}^{(\ell_i+1) \times (\ell_i+1)}$ of $B_{i,k}$, where $\hat{B}_{i,k}$ takes the form

$$\hat{B}_{i,k} = [(\hat{B}_{i,k})_{l,j}]_{l,j=1}^{\ell_i+1} \quad (20)$$

with $(\hat{B}_{i,k})_{j,j} = n_i^{-1} \sum_{t=1}^{n_i} e_{t,j}^2$ and $(\hat{B}_{i,k})_{l,j} = n_i^{-1} \sum_{t=1}^{n_i} e_{t,j} e_{t,l}$ for $j \neq l$. We provide this noise structure as an example because in our formulation, we allow for constraints on the players' optimization problems so that for each iteration k , we in fact have multidimensional observations as can be seen in (12).

The second noise structure we consider is adapted from the HC₄ estimator [13] and is given by

$$\hat{G} = \text{diag} \left(\frac{e_1^2}{(1-b_1)^{\delta_1}}, \frac{e_2^2}{(1-b_2)^{\delta_2}}, \dots, \frac{e_{n_d}^2}{(1-b_{n_d})^{\delta_{n_d}}} \right) \quad (21)$$

where $\delta_i = \min\{4, n_d b_i / (\sum_{i=1}^{n_d} b_i)\}$ and the b_i s are the diagonal elements of $B = X(X^\top X)^{-1}X^\top$. With this structure, the penalty for each residual increases with $b_i / \sum_{j=1}^{n_d} b_j$. As with the previous noise structure, we use the fitted cOLS estimator $\hat{\beta}_{\text{cOLS}}$ and residuals to get an initial \hat{G} . We selected to present this noise structure because it is computationally efficient compared with many other noise structures.

In both cases, we substitute the inferred noise \hat{G} into cGLS statistical model (19) to get the one-step cFGLS estimators. We iterate between the estimation of \hat{G} and $\hat{\beta}_{\text{cFGLS}}$ either until convergence or for a fixed number of iterations to prevent overfitting. To resolve this tradeoff and find the optimal iteration size, we adopt a simple cross validation method.

C. Boosting With Ensemble Methods

In this section, we describe several ensemble methods. Combined with a bootstrapping process, ensemble methods not only boost the size of what can often be a small data set in practice but also allow us to improve the estimator performance and explore the bias–variance tradeoff.

1) *Bootstrapping and Bagging*: Bootstrapping is a technique for asymptotic approximation of the bias and standard error of an estimator in a complex and noisy statistical model [12], [14]. We employ *wild bootstrapping* to generate a *pseudodata set* from which we generate several weak estimators that we then combine using *bagging*. While we assume

that $E(Y|X) = X\beta$, we also allow for heteroskedasticity by conditioning on the residual transformations that we imposed in the noise structure. Wild bootstrapping is a technique of parametric bootstrapping that is consistent with heteroskedastic inference and cFGLS data generation.

The bootstrapping process can be described in two steps: 1) we fit our cFGLS model, which gives us $\hat{\beta}_{\text{cFGLS}}$, and 2) then generate N replicates of *pseudodata* using the data generation process

$$\tilde{Y} = X\hat{\beta}_{\text{cFGLS}} + \Phi(e)\varepsilon \quad (22)$$

where $\tilde{Y} \in \mathbb{R}^{n_d}$ is the new observation vector (pseudoobservations), $\hat{\beta}_{\text{cFGLS}} \in \mathbb{R}^{n_d}$ is the cFGLS estimator, $\varepsilon \sim N(0, I^{n_d \times n_d})$, $e \in \mathbb{R}^{n_d}$ is the residual vector given by $e = \tilde{Y} - X\hat{\beta}_{\text{cFGLS}}$, and $\Phi : \mathbb{R}^{n_d} \rightarrow \mathbb{R}^{n_d}$ is a nonlinear transformation such that $\Phi(e) = \hat{G}^{1/2} \in \mathbb{R}^{n_d \times n_d}$. Since $E(\Phi(e)\varepsilon|X) = \Phi(e)E(\varepsilon|X) = \Phi(e)E(\varepsilon) = \mathbf{0}_{n_d \times n_d}$, using the data generation process in (24), we resample from i.i.d variables.

Bagging in regression models and trees is a technique for reducing the overall variance [14]. Using the N replicates of pseudodata generated by wild bootstrapping, we train N different models. We combine the resulting bootstrapped estimators by averaging

$$\hat{\beta}_{\text{bag}} = \frac{1}{N} \sum_{j=1}^N \hat{\beta}_{\text{cFGLS},j} \quad (23)$$

where $\hat{\beta}_{\text{cFGLS},j}$ is the estimator using the j th pseudo-data sample. Bagging works efficiently with high variance models and do not hurt the overall performance of the statistical model. We refer to the bagged estimates as bagged megalearners since they combine several weak learners/estimators. Using wild bootstrapping, the empirical covariance matrix of $\hat{\beta}$ is an asymptotic approximation of the covariance matrix and is given by

$$\hat{C}_\beta = \frac{1}{N} \sum_{j=1}^N (\hat{\beta}_{\text{cFGLS},j} - \hat{\beta}_{\text{bag}})(\hat{\beta}_{\text{cFGLS},j} - \hat{\beta}_{\text{bag}})^\top. \quad (24)$$

Asymptotic estimation of the empirical covariance matrix reveals hidden structures between players and is what we leverage in the correlation utility learning procedures.

2) *Bootstrapping and Bumping*: In a similar fashion as the bagging ensemble method, we combine bumping—a method for fitting cFGLS estimators using a random search over the model space [15]—with the wild bootstrapping generated pseudodata. In particular, we apply a *stochastic search* over several different statistical models coming from a similar data process, i.e., the data process in (24).

We add the original training data sample to the N replicates of pseudodata generated by the wild bootstrapping process and we use this data to estimate $N+1$ cFGLS estimators. We evaluate these estimators on the training set and select the one with the least training error. The cFGLS bumping estimator is given by

$$\hat{\beta}_{\text{bump}} = \arg \min_{\hat{\beta}_{\text{cFGLS},j}} \|\tilde{Y} - X\hat{\beta}_{\text{cFGLS},j}\|_2^2 \quad (25)$$

Algorithm 1 L_2 -Gradient Boosting With cFGLS

```

1: function cFGLSGRADBOOST( $X, Y, \nu$ )
2:    $\hat{H} \leftarrow X(X^\top X)^{-1} X^\top$   $\triangleright$  compute  $\hat{H}$  matrix
3:    $\nu \leftarrow s \in (0, 1]$   $\triangleright$  set shrinkage (updating) parameter
4:    $m_{\text{stop}} \leftarrow 1$   $\triangleright$  iteration number
5:   choose  $M_{\text{max}}$   $\triangleright$  upper iterations bound
6:    $\text{AIC}_{\text{list}} \leftarrow []$   $\triangleright$  create empty list
7:   Compute stopping iteration time  $m_{\text{stop}}$ :
8:   while  $m_{\text{stop}} < M_{\text{max}}$  do
9:      $R_{m_{\text{stop}}} \leftarrow (I_{n_d \times n_d} - \nu \hat{H})^{m_{\text{stop}}}$ 
10:     $B_{m_{\text{stop}}} \leftarrow (I_{n_d \times n_d} - R_{m_{\text{stop}}})$ 
11:     $\sigma_{m_{\text{stop}}}^2 \leftarrow n_d^{-1} \sum_{i=1}^{n_d} (Y_i - (B_{m_{\text{stop}}} Y)_i)^2$ 
12:     $\text{AIC}_{m_{\text{stop}}} \leftarrow \left( \log \sigma_{m_{\text{stop}}}^2 + \frac{1 + (\text{Tr}(B_{m_{\text{stop}}})) / n_d}{1 - (\text{Tr}(B_{m_{\text{stop}}}) + 2) / n_d} \right)$ 
13:     $\text{AIC}_{\text{list}} \cdot \text{append}(\text{AIC}_{m_{\text{stop}}})$ 
14:     $m_{\text{stop}} \leftarrow m_{\text{stop}} + 1$ 
15:   end while
16:    $\hat{M} \leftarrow \arg \min \text{AIC}_{\text{total}}$   $\triangleright$  find minimum point
17:    $\hat{\beta}_{\text{cFGLS}} \leftarrow$  estimate of  $\beta_{\text{cFGLS}}$  using cFLGS
18:    $e_{\text{FGLS}} \leftarrow Y - X \hat{\beta}_{\text{cFGLS}}$   $\triangleright$  residuals estimation
19:    $e \leftarrow e_{\text{cFGLS}}$   $\triangleright$  initialize residuals
20:    $k \leftarrow 1$   $\triangleright$  iteration index
21:    $\beta_{\text{boost}} \leftarrow \hat{\beta}_{\text{cFGLS}}$   $\triangleright$  initialize cFGLS boosted learner
22:   Compute boosted learner  $\beta_{\text{boost}}$ :
23:   while  $k < \hat{M}$  do
24:      $\beta_i \leftarrow (X^\top X)^{-1} X^\top e$   $\triangleright$  residuals fitting
25:      $\hat{\beta}_{\text{boost}} \leftarrow \hat{\beta}_{\text{boost}} + \nu \beta_i$   $\triangleright$  update formula
26:      $e \leftarrow Y - X \hat{\beta}_{\text{boost}}$   $\triangleright$  residuals update
27:      $k \leftarrow k + 1$ 
28:   end while
29: end function

```

where $\hat{\beta}_{\text{cFGLS},j}$ s are the cFGLS estimators from derived from the bootstrapped data.

3) *Gradient Boosting*: We combine L_2 -gradient boosting, which is a repeated least squares fitting of residuals [16], with cFGLS. Gradient boosting is a boosting technique that uses an L_2 loss function combined with a gradient descent update method for combining weak learners at each iteration. Boosting estimators are trained in sequence using a weighted version of the original data set. In general, boosting methods are extremely useful for combining models by incrementally training each new model by emphasizing the errors of the previous training instances. They are used extensively in classification methods such as logistic regression and support vector machines.

Repeated residual fitting is applied until we reach iteration m_{stop} , a stopping criterion selected using the Akaike information criterion (AIC) to avoid overfitting [17]. The procedure is detailed in Algorithm 1.

IV. APPLICATION TO BERTRAND–NASH COMPETITION

Let us illustrate the framework and its performance of the robust utility learning framework before moving on by applying it to estimate market demand functions under the Bertrand–Nash equilibrium (see [18]–[20]). The toy model can

TABLE I

MSE OF FORECASTING USING THE PROPOSED ROBUST UTILITY LEARNING METHODS VERSUS cOLS ESTIMATORS FOR THE BERTRAND–NASH COMPETITION. THE BEST PERFORMING METHOD IS INDICATED IN BOLD FOR EACH OF THE FIRMS

Firm 1	bagging	boosting	bumping	cOLS
<i>MSE</i>	0.05	0.51	0.65	1.62
Firm 2	bagging	boosting	bumping	cOLS
<i>MSE</i>	1.58	0.71	0.89	2.54

be thought of as an abstraction of Bertrand price setting for commodities such as oil, gas, and coal [21], [22].

Consider two firms competing to sell their product by setting the prices p_1 and p_2 for firm 1 and 2, respectively. The firms' utility functions are their revenue, i.e., $f_i(p_1, p_2) = p_i D_i(p_1, p_2, \xi)$, where D_i is the demand function for firm i and $\xi \sim \mathcal{N}(1.5, 0.5)$ is a random variable that captures the fact that demand is dependent on economic indicators in addition to the prices set by the firms. In this stylized example, we consider linear demand functions given by

$$D_i(p_1, p_2, \xi) = \theta_{i,1} + \theta_{i,2} p_1 + \theta_{i,3} p_2 + \nu \xi \quad (26)$$

where $\theta_i = (\theta_{i,j})_{j=1}^3$ are unknown parameters to be estimated and $\nu = 1.5$ is a known parameter. The prices are constrained to be in the interval $[0, \bar{p}]$, where $\bar{p} \in \mathbb{R}_+$ is the upper bound. We let $\theta_1 = (-1.0, 0.5, -1)$ and $\theta_2 = (0.3, -1, 0.3)$ be the ground-truth values for the parameters we wish to estimate. Thus, $\bar{f}_i(p_1, p_2) = \nu \xi$, and examining the marginal revenue functions $D_i f_i(p_1, p_2)$, we have that $\phi_1(p_1, p_2) = [1 \ 2p_1 \ p_2]^\top$ and $\phi_2 = [1 \ p_1 \ 2p_2]^\top$.

In order to generate the data set, we add a noise term $\varepsilon \sim \mathcal{N}(0, 0.5)$ to the marginal revenue functions, i.e., $D_i f_i(p_1, p_2) + \varepsilon$, and solve for the Bertrand–Nash equilibrium. We simulate the game between the firms 600 times. In the robust utility learning framework, for this example, we employ the HC₄ noise structure and compute the cOLS, cFGLS, bagging, boosting, and bumping estimators. We use a tenfold cross-validation procedure to prevent overfitting. Table I contains error using two metrics for both firms. Fig. 1 shows the forecast for part of the testing set using cOLS and each of the ensemble methods compared with the ground truth. While bagging performed best for firm 1 and boosting for firm 2 in the particular instantiation of this toy example, the performance more generally is dependent on the noise structure in the demand and marginal revenue functions, the sample size, and the dynamics between the two firms. However, it is interesting to point out that as we increase the variance on ξ , each ensemble method performance stays relatively the same, yet the cOLS error increases significantly.

V. CORRELATED UTILITY LEARNING

In this section, we describe how learned correlations between players can be leveraged to *boost* estimator performance. We add a second step to the estimation procedure in which we craft a new game where players' utilities are composed of their original estimated utility plus some combination of other players' utilities weighted by the estimated correlation between players.

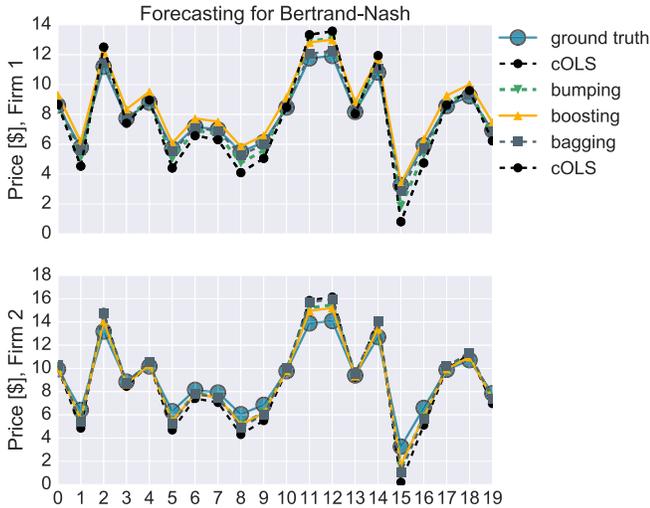


Fig. 1. Forecast for Firms 1 and 2 using cOLS and each of the ensemble methods. The ground-truth prices are depicted by the blue dashed lines; the cOLS forecasts are depicted by the black dashed lines, the bagging forecasts are depicted by the gray dashed lines, the bumping forecasts are depicted by the green dashed lines, and the boosting forecasts are depicted by the gold lines.

When the correlations between players are positive, we are creating what we refer to as *pseudocoalitions* since players are not *explicitly* agreeing to collude in the game but are rather doing so *implicitly*. The degree of *coalition* is discovered by the robust utility learning process through estimating the empirical covariance \hat{C}_β , i.e., asymptotic approximation of the covariance matrix, of $\hat{\beta}_{\text{est}}$, where we use the notation $\hat{\beta}_{\text{est}}$ to abstractly denote the estimator derived from whichever of the methods described in the previous section is employed. On the other hand, when the correlations between players are negative, by combining their utilities, we aim to take advantage of active players' richer data sets in predicting the behavior of players with less variation and frequency in their observed actions.

We refer to the learned utility— \hat{f}_i for player i —from the robust utility learning framework as the *base utility* and it is given by

$$\hat{f}_i(x_i, x_{-i}; \hat{\theta}_i) = \bar{f}_i(x_i, x_{-i}) + \langle \phi_i(x_i, x_{-i}), \hat{\theta}_i \rangle \quad (27)$$

where $\hat{\theta}_i$ is extracted from $\hat{\beta}_{\text{est},i}$.

Using the correlations we learn that when we estimate \hat{f}_i , we construct a new utility \hat{g}_i by combining scaled versions of a subset (potentially all) of the other agents' utilities that are correlated with agent i . We formulate an optimization problem to determine the scaling coefficients. The correlated utility \hat{g}_i for player i is given by

$$\hat{g}_i(x_i, x_{-i}) = \sum_{j \in \mathcal{K}_i} z_{i,j} \sigma_{i,j} \bar{f}_j(x_i, x_{-i}) + \sigma_{i,i} \langle \phi_i(x_i, x_{-i}), \hat{\theta}_j \rangle \quad (28)$$

where $\mathcal{K}_i \subset \mathcal{I}_i$ is a subset of the players correlated with player i , $\sigma_{i,i}$ is the estimated variance of player i determined by the empirical covariance matrix, $\sigma_{i,j}$ is the covariance between the parameter estimates for player i and j also

determined by the empirical covariance matrix, and $z_{i,j}$ are scaling constants to be optimized. We refer to the resulting game as an *approximated correlation game*.⁴

Given the form of \hat{g}_i , our goal is to select the scaling constants $z_{i,j}$ in order to reduce the forecasting error. Analogous to the base utility learning framework presented in Section III-A, using our training data, we formulate a convex optimization problem using optimality conditions on each player's individual optimization problem where we assume that player i is optimizing \hat{g}_i with respect to its own choice variable x_i . In particular, we solve a convex optimization problem formulated as follows. Define the vector $z_i \in \mathbb{R}^{|\mathcal{K}_i|}$ by $z_i = (z_{i,j})_{j \in \mathcal{K}_i}$ and let $z = (z_i)_{i \in \mathcal{I}}$. For player i 's optimization problem $\max\{\hat{g}_i(x_i, x_{-i}) \mid x_i \in \mathcal{C}_i\}$, let the residual of the stationarity condition be given by

$$r_{s,i}^{(k)}(z_i, \mu_i; \hat{\theta}) = D_i \hat{g}_i(x_i^{(k)}, x_{-i}^{(k)}) + \sum_{j=1}^{\ell_i} \mu_i^j D_i h_{i,j}(x_i^{(k)}) \quad (29)$$

and the residual of the complementary conditions be given by

$$r_{c,i}^{j,(k)}(\mu_i) = \mu_i^j h_{i,j}(x_i^{(k)}), \quad j \in \{1, \dots, \ell_i\}. \quad (30)$$

As before, let $r_{c,i}^{(k)}(\mu_i) = [r_{c,i}^{1,(k)}(\mu_i) \dots r_{c,i}^{\ell_i,(k)}(\mu_i)]$. Define $Q_i \in \mathbb{R}^{n_i \times |\mathcal{K}_i|}$ by

$$Q_i = [\sigma_{i,j} D_{i,i}^2 \bar{f}_i(x^{(k)})]_{k=1, j \in \mathcal{K}_i}^{n_i} \quad (31)$$

and $q_i \in \mathbb{R}^{n_i}$ by

$$q_i = \left[\sum_{j \in \mathcal{K}_i} \sigma_{i,j} (D_{i,i}^2 \phi_i(x^{(k)}), \hat{\theta}_j) \right]_{k=1}^{n_i}. \quad (32)$$

Then, we have the following convex optimization problem to determine the scaling factors $z_{i,j}$:

$$\begin{aligned} \min_{z, \mu} \quad & \sum_{i=1}^p \sum_{k=1}^{n_i} \chi_i(r_{s,i}^{(k)}(z_i, \mu_i; \hat{\theta}), r_{c,i}^{(k)}(\mu_i)) \\ \text{s.t.} \quad & Q_i z_i + q_i \leq 0, \quad \mu_i \geq 0 \quad \forall i \in \mathcal{I} \end{aligned} \quad (\text{P}')$$

Solving (P') gives us estimated correlated utilities \hat{g}_i for each $i \in \mathcal{I}$ that we then use to forecast the players' decisions.

VI. APPLICATION TO SMART BUILDING SOCIAL GAME

We now specialize the robust and correlated utility learning frameworks to the smart building social game.

A. Social Game Experimental Set-Up

Our experimental setup is in a collaborative space—an open shared work space with cubicles—within the CREST Center on the UC Berkeley campus. We crafted a social game such that occupants in this collaborative freely vote according to their usage preferences of shared resources and

⁴We remark that there exists an equilibrium concept called *correlated equilibrium* [23], which generalizes a Nash equilibrium by characterizing correlations between randomized strategies; we mention this only to alleviate any potential confusion. The equilibrium concept we utilize for the approximated correlation game is still a pure Nash equilibrium and there is no coordinating mechanism.



Fig. 2. Graphical user interface (GUI) for energy-based social game. (a) Display, in table form, of points and votes for energy consumption, HVAC, and lights. (b) Display of the GUI for logging lighting setting preferences.

are rewarded with points based on how energy efficient their strategy is in comparison with the strategy of the other occupants. We employ a lottery mechanism consisting of three Amazon gift cards executed biweekly to reward occupants; occupants with more points are more likely to win the lottery.

The office is divided into five lighting zones and two heating, ventilating, and air conditioning (HVAC) zones. In this space, there are a total of 20 occupants who are eligible to participate in the social game. If the occupants are not present in the office, they are excluded from the game at that time instant. When they arrive at the office, they can rejoin the game. To enforce the rule that those who are not present in the space cannot vote remotely, we executed a simple presence detection algorithm based on their power usage [24].

We have installed a Lutron⁵ system for precise control of the lighting setting (dim level of the lights) in the office as well as desk-level energy monitoring devices (i.e., ACME wireless sensors [25]) to meter the energy usage of each occupant. In addition, we have modified the HVAC system so that it can be precisely controlled. We have verified prior to our experiment that the implemented control of these systems results in expected performance.

We have developed a platform to interface with the occupants as well as manage and process collected data. The platform includes a Web portal and a mobile app that the occupants may use to participate in the game. It also allows occupants to visualize different aspects of the social game, e.g., the lighting setting and the energy efficiency level of different occupants or the entire building, as well as view the point level and historical voting record of other occupants among many other statistics. Fig. 2 shows the user interface for viewing points and logging votes. Fig. 3(a) shows a visualization of the current light level using a green-to-red scale with green being more energy efficient. The current temperature is also displayed. Fig. 3(b) shows a visualization of each present and participating occupant's energy efficiency level.

In this paper, we report on a social game experiment conducted based only on the lighting shared resource.⁶ Prior to the start of the social game experiment, the lighting setting was 90% of the maximum possible lighting setting. At the start of the social game experiment, we set a default lighting setting, which acts as the suggested lighting setting and is the dim level setting in the office if no occupants are participating in the game. Throughout the game, we adjust the default lighting setting as well as the points. The lottery mechanism coupled with the points we distribute comprises the *incentive component* of the feedback to the participants while the default lighting level is the *physical control component* of the feedback. These two mechanisms act as our control inputs and our feedback mechanism to the participants. We seek to design them by taking into consideration the preferences of the participants. In this way, these mechanisms close the loop around the participant, and with our proposed utility learning scheme, these mechanisms can be modified to encourage more energy efficient resource consumption.

The game is designed to leverage interactions among occupants, who win points based on how energy efficient their lighting vote is compared with others. An occupant's vote is for the lighting setting in their zone as well as neighboring zones. The occupants select their desired lighting setting in the continuous interval $[0, 100]$, where each value represents the percentage of the maximum lighting setting possible in the space. The occupants can vote as frequently as they like and the average of all the occupants' current votes sets the implemented lighting setting in the collaboratory. An occupant can leave the lighting setting as the default level after logging in or they can change it depending on their preferences and other environmental factors that may affect their choice.

The experimental trials reported on in this paper were conducted over the period of 285 days.⁷ Experiments with four different default levels, $\{10\%, 20\%, 60\%, 90\%\}$, were

⁶We remark that while our experimental platform is capable of conducting a social game that includes lights, HVAC, and personal energy consumption, we report only on an experiment that focuses on lighting in order to isolate combined effects from these different resources. In on-going experiments, we are examining all aspects jointly.

⁷The period of the experiment was March 3, 2014 to December 14, 2014.

⁵<http://www.lutron.com/en-US/Pages/default.aspx>.



Fig. 3. Occupants can access a variety of information when they log into the social game portal, including various displays of energy consumption by other participants in the game. (a) Display of current light level and temperature in the collaboratory space; energy efficiency of the lights is coded by color, where the light green strip indicates *higher energy efficiency*. (b) Display of collaboratory floor plan with dots indicating where present and participating players sit. Players not in the office are excluded from the game. The color of the dot indicates the level of energy efficiency of the player compared with the other participants; the green light indicates higher efficiency, while the red light indicates lower efficiency.

conducted, covering a spectrum of lighting conditions. Since occupants were allowed to vote whenever they chose, their response rate per day varies. The data set we collected consists of occupant votes (meaning the lighting level they select) over the period of investigation as well as the points that were distributed to each occupant. We collected 6885 votes over the period of the experiment.

B. Brief Background

In order to place the work pertaining to building energy efficiency in the context of the state of the art, we briefly overview existing approaches.

Recognizing that HVAC systems are responsible for a large portion of building energy consumption, many control theoretic approaches such as those in [30] and [31] derive model predictive and distributed control policies for HVAC systems. While these control theoretic approaches make efforts to account for the presence of occupants, they tend to ignore occupant behaviors and, more importantly, their heterogeneous preferences.

There are other works that make strides toward incorporating behavioral models of occupants, Boman *et al.* [28] employ a multiagent system approach to develop a framework

for incorporating occupant comfort preferences and Bourgeois *et al.* [29] develop behavioral models for lighting usage. In a more active approach, Song [30] develops a *collaborative setting definition paradigm* in which occupants and facilities managers submit preferences and requirements and a rule engine tries to resolve them in order to create a universal control policy. While occupants' preferences are taken as inputs to the building control design, it is not clear that it is possible to satisfy all the occupants' comfort preferences simultaneously with those of the facilities manager; hence, the misalignment between preferences and incentives remains.

In our approach, on the other hand, we leverage a social game that creates a friendly competition between users and employs incentives to resolve conflicting preferences by compensating users. Within the energy application domain, gamification has been largely used for education or awareness (see [31], [32]). There are works that are closely related to ours in the sense that they also recognize that occupants are self-interested participants in smart buildings and try to account for their strategic behavior. For example, Li *et al.* [33] develop an interesting scheme for engaging occupants directly in DR. Analogous to our approach, occupants are modeled as utility maximizers in a game theoretic context where they are incentivized to curtail their consumption in response to an event. Our approach differs in that we focus on shared resources such as lighting and HVAC instead of personal devices (e.g., desk appliances). Furthermore, it is assumed in [33] that the type space (i.e., their preferences) of the users is a known finite set of two possible values. We do not assume that the facility manager knows the utility function or the type of the users and we propose an algorithm for learning this utility function from the observations of decisions.

While incorporating occupant preferences into building automation is not novel in and of itself, we propose an innovative algorithm for learning occupant preferences in competitive environments and, moreover, learning how their actions are correlated. Such correlations can be leveraged in improving incentive mechanisms to shape users' preferences thereby providing more flexibility. Our method is applied to real-world data from experimental trials we conducted as opposed to simulations as is the case with many existing works. Furthermore, it is agnostic to the application and could be applied in general to other scenarios in which users are competing for constrained but shared resources. For example, the utility learning method can be easily adapted to learning preferences of individual buildings interacting with an aggregator or learning preferences of drivers seeking on-street parking. In each of these cases, there exists a planner—the aggregator or department of transportation—tasked with managing a resource being consumed by self-interested users.

C. Occupant Decision Making Model

Each agent's vote x_i is constrained to be in the interval $[0, 100] \subset \mathbb{R}$. Let \bar{x} denote the average of the lighting votes and the setting that is implemented, e.g., at observation instance indexed by k , $\bar{x}^{(k)} = (1/|S^k|) \sum_{j \in S^k} x_j^{(k)}$. We model each agent's utility as being composed of two basis functions

that capture the tradeoff between desired lighting (satisfaction) and desire to win. The lighting satisfaction an occupant feels may be a function of several factors including their productivity (ability to perform their job) and physical comfort. We abstractly model their desired lighting level using a Taguchi loss function, $\psi_i(x_i, x_{-i}) = -(\bar{x} - x_i)^2$, which is interpreted as modeling occupant dissatisfaction in such a way that it is increasing as variation increases from their reported desired lighting setting (their vote) [34].

We acknowledge that an agent may have some internal desired lighting level that is different from its vote, e.g., the agent may realize that voting an extreme value pushes the average toward a more desirable setting. This type of *gaming* results in *moral-hazard*-type issues that can be addressed in the incentive design step [1], [2]. Thus, we set this type of gaming aside for the time being and focus instead on the unknown preferences—a different kind of asymmetric information that leads to *adverse selection*—between *lighting* and *winning*.

Points are distributed by the planner using the relationship $\rho(x_b - x_i)(p(x_b - \bar{x}))^{-1}$, where x_b is the baseline setting for the lights. For the experiment $x_b = 90\%$, i.e., the lighting setting used before the implementation of the social game. However, we model each occupant as having a *winning* basis function given by $\phi_i(x_i, x_{-i}) = -\rho c(x_i)^2$, where ρ is the total number of points distributed by the planner and c is a scaling factor that is used primarily to scale the two terms of the utility function, given that we artificially inflate the points offered in order to increase their appeal to players and thus induce greater participation.⁸ The form of the winning function can be interpreted as capturing the perception that by voting zero, the occupant is selecting the action that will provide the greatest return of points, given that points are awarded based on how energy efficient their vote is compared with others.⁹

Hence, the utility functions for the social game are modeled as $f_i(x_i, x_{-i}; \theta_i) = \theta_i \phi_i(x_i, x_{-i}) + \psi_i(x_i, x_{-i})$. The constraint sets \mathcal{C}_i for each player are determined by the box constraints on the lighting vote for that player, i.e., $\mathcal{C}_i = \{x_i \in \mathbb{R} \mid h_{i,j}(x_i) \geq 0, j \in \{1, 2\}\}$, where $h_{i,1}(x_i) = 100 - x_i$ and $h_{i,2}(x_i) = x_i$.

In order to formulate (P) for the social game application, we need to determine the admissible parameter sets $\Theta_i, i \in \mathcal{I}$ in such a way that we ensure the estimated utility functions are concave and such that equilibria of the estimated game are isolated. We derive a lower bound θ_{LB} such that all $\theta_i \in \Theta_i = \{\theta_i \in \mathbb{R} \mid \theta_i > \theta_{\text{LB}}\}, i \in \mathcal{I}$, induce games with these characteristics. To this end, we utilize the second derivative condition on players' utility functions, that is, if for each

⁸Inflating the points is a process of *framing* [35], that is, dependent on how the reward system presented to agents greatly impacts their participation. Framing is routinely used in reward programs for credit cards among many other point-based programs. The scaling factor c in the winning function removes the framing effect from the estimation procedure. It is selected to ensure that the scale of the two basis functions are similar.

⁹We explored other forms of the winning function including the log function, a quasi-concave function that is typically used to represent how individuals value money since it represents the diminishing returns property well [9]. However, the quadratic form of the function we report on here significantly outperformed other choices so that, for the purpose of a prescriptive model, it captures the agents' perceptions about the point distribution mechanism and their value more accurately.

$i \in \mathcal{I}, D_{i,i}^2 f_i(x) < 0$, then the game is concave. Computing $D_{i,i}^2 f_i$ and using some algebra, we have that $\theta_i > -(c\rho)^{-1}(1-p^{-1})^2$, where the right-hand side is a negative nonincreasing function of p . Thus, concavity is ensured regardless of the number of players by setting $p = 2$, the minimum number of players in a noncooperative game. Then, given fixed ρ and $0 < \zeta \ll 1$, the lower bound $\bar{\theta}_{\text{LB}} = -(4c\rho)^{-1} + \zeta$ will guarantee that the estimated game is concave.

If $D\omega(x, \mu)$ is invertible, we know that differential Nash equilibria are isolated [7]. Hence, we can augment the constraint sets Θ_i to encode this condition. Given the structure of the utility functions, $D\omega(x, \mu)$ is simply the game Hessian $H = [H_{i,j}]_{j,i=1}^p$ with $H_{i,i} = D_{i,i}^2 f_i$ and $H_{i,j} = D_{i,j}^2 f_i$. Hence, if H is invertible, then the differential Nash are isolated; this is guaranteed for $p \geq 4$ provided the constraint defined by $\bar{\theta}_{\text{LB}} = -(4c\rho)^{-1} + \zeta$ using $\zeta = 10^{-2}$. Indeed, let $H(p)$ denote the game Hessian as a function of the number of players and note that for a particular p , with some simple algebra, it is easy to write $H(p)$ as an off-diagonal matrix constant matrix such that $H_{i,i} = d_i + \alpha$ and $H_{i,j} = \alpha$, where $d_i = -2(1 - 1/p) - 2c\rho\theta_i$ and $\alpha = 2(p - 1)/p^2$. It is straightforward to verify by determining the eigenvalues of H as p varies via the method described in [36] that for $p \geq 4$, H will be invertible. For the social game data, at each observation indexed by k , the number of participating players is at least 4. Thus, to ensure concavity and isolated equilibria of the estimated social game, we define $\Theta_i = \{\theta_i \in \mathbb{R} \mid \theta_i > \bar{\theta}_{\text{LB}}\}$ with $\bar{\theta}_{\text{LB}} = -(c\rho 4)^{-1} + \zeta$ with $\zeta = 10^{-2}$.

VII. UTILITY LEARNING RESULTS

We now present the results of the proposed robust utility learning method applied to data collected from the social game experiment.

As we previously described, our data set consists of the votes logged by the players who vote throughout the day. We present estimation results for the complete data set of all the votes, which we refer to as the *dynamic data set*, and estimation results for an aggregated data set constructed by taking the average of a players' votes over the course of each day in the experiment—this is referred to as the *average data set*. While this aggregation significantly reduces the size of our data set, it smooths the players' voting profiles and increases the size of active players in each game—occupants may arrive or leave the office when they so choose. This *average data set* also reduces the computational load, which may be beneficial to a facilities manager in the incentive design process, especially if the incentive scheme is quasi-static and uses historical data to generate the next incentive. The dynamic data set is much richer, being composed of every vote (a total of 6885 votes) the occupants made throughout the duration of the experiment (285 days). The time from one vote to the next may be several minutes to hours depending on the activity of the occupants. This data set is much larger and thus increases the computational load. However, it allows us to extract more distinct player profiles and can support real-time incentive design schemes.

We present results for both data sets using data from the period of the experiment in which the default lighting setting

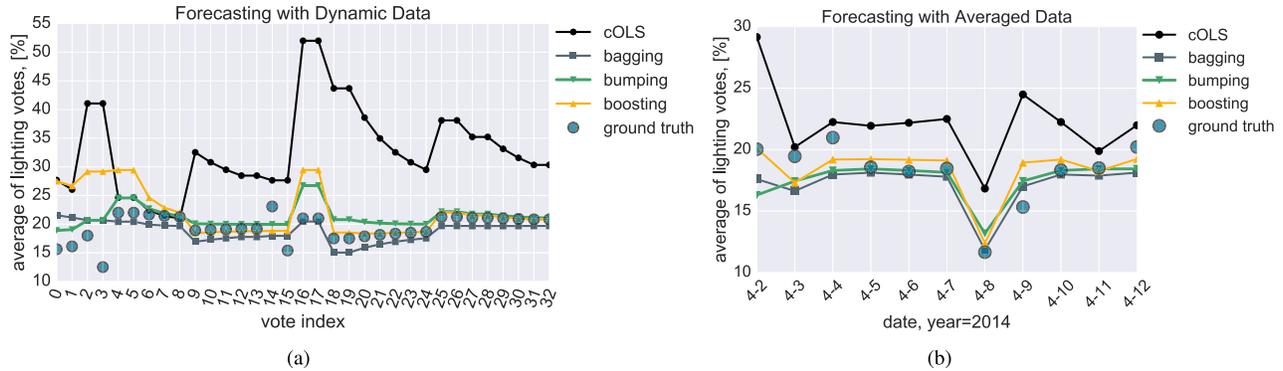


Fig. 4. Forecasting results for (a) dynamic data and (b) averaged data for the default lighting setting 20: For the dynamic data, the x -axis values indicate the index of when a choice was made by one or more of the occupants (i.e., when the implemented lighting setting is changed); the time from one index to the next may be several minutes to hours depending on the activity of the occupants. For the averaged data, the x -axis values are dates (month and day). The ground-truth average of the lighting votes is depicted by the blue curves, the forecast for cOLS is depicted by the black curves, the forecast for bagging is depicted by the gray curves, the forecast for bumping is depicted by the green curves, and the forecast for boosting is depicted by the gold curves. The forecast for the robust utility learning methods is approximately near the ground truth for both data sets, while the cOLS estimates produce Nash equilibria with a large error.

was 20%—the results for the other default lighting settings are similar. The period of the experiment where the default lighting setting was 20% consisted of 42 days and thus the size of the averaged data set is 42. Over this period, there were 220 votes by occupants, which is the size of the dynamic data set. We divide each of the data sets into training (80% of the data) and testing (20% of the data) sets and apply each of the methods discussed in Section III. We apply a tenfold cross validation [14] procedure to limit overfitting.

A. Forecasting via Robust Utility Learning

We estimate the parameters using cFGLS and the ensemble methods bagging, bumping, and boosting for both the average and dynamic data sets. For gradient boosting, we use the HC₄ noise structure (21) since the *leverage* values b_{ii} of B are larger [13]; in each of the other methods, we used the block diagonal noise structure (20).

Using the estimated utility functions, we simulate the game using a projected gradient descent algorithm, which is known to converge for concave games [38]. In Fig. 4(a) and (b), we compare the ground-truth voting data to the predictions for each of the learning schemes using the dynamic and averaged data sets, respectively. Our proposed robust models, i.e., using the estimated parameters obtained via bagging, bumping, and boosting, capture most of the variation in the true votes (in both data sets) and significantly outperform cOLS. In Table II, using three metrics—root-mean-square error (RMSE), mean absolute error (MAE), and mean absolute scaled error (MASE)—we report the forecasting error for each of the methods.

The estimated models using our robust utility learning methods significantly reduce the forecasting error compared with cOLS. The cOLS method has particularly poor forecasting performance on the dynamic data set since it does not capture the correlated error terms describing the interactions between users. Moreover, our robust methods perform better than cOLS with the averaged data set even though the sample size is small.

TABLE II
RMSE, MAE, AND MASE [37] OF FORECASTING USING THE PROPOSED ROBUST UTILITY LEARNING METHODS VERSUS cOLS ESTIMATORS FOR BOTH DATA SETS IN DEFAULT LIGHTING SETTING 20. THE BEST PERFORMING METHOD IS INDICATED IN BOLD FOR EACH OF THE DATA SETS, DYNAMIC AND AVERAGE

<i>Dynamic, \hat{f}_i</i>	bagging	boosting	bumping	cOLS
RMSE	8.31	10.11	12.56	22.53
MAE	5.20	6.55	6.38	18.35
MASE	2.08	6.38	2.55	7.34
<i>Averaged, \hat{f}_i</i>	bagging	boosting	bumping	cOLS
RMSE	2.05	1.68	1.96	9.36
MAE	1.58	1.31	1.48	6.01
MASE	0.71	0.59	0.67	2.69

As for the ensemble methods, bagging outperforms the other three methods when using the dynamic data set. On the other hand, for the averaged data set, gradient boosting gives the least forecasting error. This is in large part due to the fact that we use the HC₄ noise structure. Since the average data set has been smoothed, we expect less correlation between players and the HC₄ noise structure captures this.

B. Estimated Utility Functions

Fig. 5 shows the estimated utility functions and their contour plots for occupants 2 and 8—passive and aggressive occupants, respectively—using the parameters obtained via the bagging ensemble method with the dynamic data set. We remark that we do not observe the actual value of agents’ utilities; we instead observe only the agents’ decisions. The purpose of the figures is to show the estimated utility shapes for players with significantly different voting profiles (the observable we have). The particular occupants we selected represent players that prefer *winning* to lighting satisfaction (occupant 8) and players that prefer *lighting satisfaction* to winning (occupant 2). In particular, occupant 2’s estimated utility function appears to be higher at greater lighting settings. Exactly the opposite occurs for occupant 8 whose estimated utility function indicates that

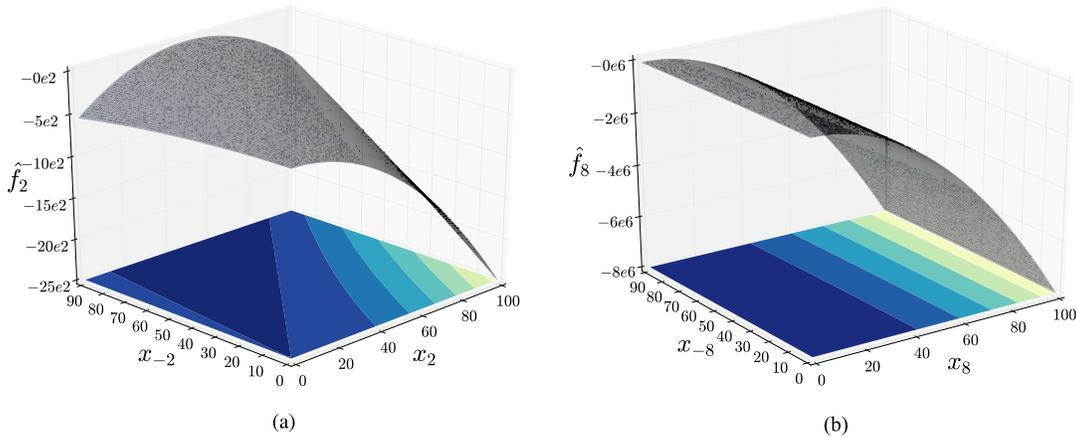


Fig. 5. Bagging estimated utility functions—using the dynamic data set—of (a) agent 2 and (b) agent 8. The functions are plotted as a function of each agent’s own vote x_2 (x_8) and other players’ votes x_{-2} (x_{-8}). Note that agent 8, a very aggressive player, is indifferent to the choices of the other agents as indicated by the fact that its utility is maximized in the same location given any value of x_{-8} . On the other hand, occupant 2 responds to changes in the other agents’ votes and appears to prefer a greater lighting setting (more illumination). This indicates that there are different types of players, and thus, incentives may need to be designed individually for these player types in order to elicit the desired response.

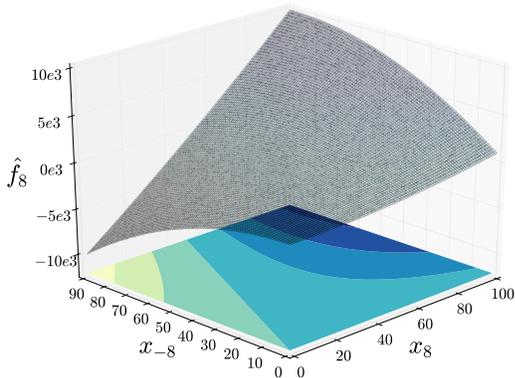


Fig. 6. Agent 8’s cOLS estimated utility function—using the dynamic data set—plotted as a function of (x_8, x_{-8}) . This figure demonstrates that using cOLS (the worst performing estimator) results in learning a utility function that is not representative of this type of player’s behavior [as can be seen by comparing with Fig. 5(b)]. Incentives or control designed using this function may result in performance.

despite changes in the average lighting vote of other players, occupant 8 aggressively votes for a zero lighting setting, which returns the most points.

For comparison and to highlight the improvement that the robust utility learning framework offers, in Fig. 6, we show the estimated utility function for occupant 8 using cOLS. What we see is a very different utility function that indicates occupant 8 cares more about lighting satisfaction than winning—indicated by the fact that its utility is not maximized at zero. This is misleading since occupant 8 predominately votes for zero. This is significant since incentive/control design based on such an erroneous utility function may lead to very poor performance and occupant dissatisfaction.

C. Bias Approximation and Bias–Variance Tradeoff

Forecasting accuracy can be enhanced by allowing for a small amount of bias if it results in a large reduction in variance. For a process $Y = X\theta + \epsilon$, the mean square

TABLE III

cFGLS ESTIMATOR VALUE AND THE BAGGING, GRADIENT, BOOSTING, AND BUMPING ENSEMBLE METHODS BIAS APPROXIMATION FOR THE MOST ACTIVE USERS. WE UTILIZED THE DYNAMIC DATA SET FROM THE PERIOD IN WHICH THE DEFAULT LIGHTING SETTING WAS SET TO 20. WE DENOTE THE OCCUPANTS WITH NEARLY UNBIASED ESTIMATORS IN BOLD

Id	cFGLS	Bagging Bias	Boosting Bias	Bumping Bias
2	-0.7	0.11	0.17	0.02
6	0.5	1.12	1.77	0.93
8	298.1	-176.9	-370.3	120.5
14	337.5	-186.3	-400.2	149.7
20	-0.8	0.07	0.21	-0.53

error (MSE) characterizes the *bias–variance tradeoff*

$$\text{MSE}(x) = E[(Y - \theta_{\text{est}}^\top x)^2] \quad (33)$$

$$= \underbrace{(E[\theta_{\text{est}}^\top x] - Y)^2}_{\text{bias}} + \underbrace{E[(\theta_{\text{est}}^\top x - E[\theta_{\text{est}}^\top x])^2]}_{\text{variance}}. \quad (34)$$

Introducing bias in exchange for reduced variance is widely used in ridge regression and in lasso techniques in the form of *a priori* knowledge [14]. In our robust utility learning framework, we introduce noise structures that approximate the true data process so that we can fit cFGLS estimators that are nearly unbiased for those players whose historical voting record has a large amount of variation.

We approximate the bias for each of the estimators. In Table III, we present cFGLS estimates obtained using the dynamic data during the time window in which the default lighting setting was 20%¹⁰ for selected occupants—the most active players—as well as the approximated bias for the estimates generated by bagging, bumping, and boosting.

Figs. 7 and 8 contain histograms of the cFGLS estimators obtained using the dynamic data and bootstrapped average, respectively. In each of these histograms, we also indicate

¹⁰The results for the other default lighting settings are similar.

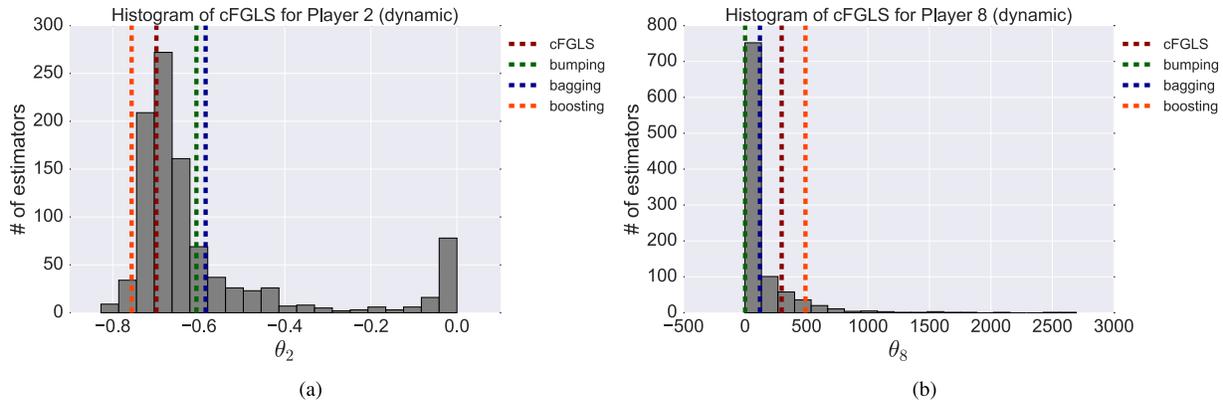


Fig. 7. Histograms depicting the estimates generated with the wild bootstrapping technique using the *dynamic data set* for (a) player 2 and (b) player 8. The vertical lines mark the value of the cFGLS (red dashed line), bumping (green dashed line), bagging (blue dashed line), and boosting (orange dashed line) estimators. The histogram for player 2 is approximately normally distributed around the initial cFGLS estimator, indicating that it is unbiased. On the other hand, this is not the case for player 8. Thus, its cFGLS estimator is biased. Overall, the majority of the proposed ensemble methods result in a significant reduction in variance in exchange for a small increase in bias and a greater forecasting accuracy. In [39], we develop a hierarchical mixture model that considers both bias and variance.

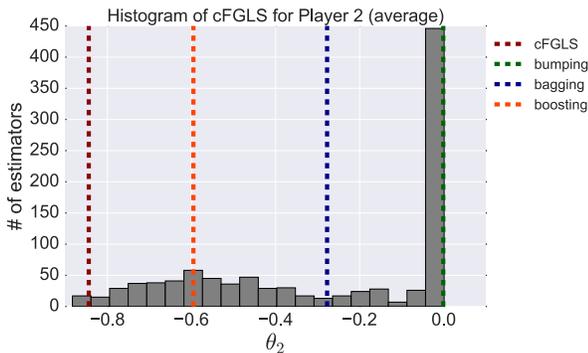


Fig. 8. Histogram depicting estimator values for player 2 using the wild bootstrapping technique using the *average data set*. The vertical lines mark the value of the cFGLS (red dashed line), bumping (green dashed line), bagging (blue dashed line), and boosting (orange dashed line) estimators. We remark that the estimators are all biased. This is expected due to the limited sample size of the average data set. Thus, the average data set cannot be used for optimizing the bias–variance tradeoff.

the original cFGLS¹¹ (indicated by the red dashed line), bagging (indicated by the blue dashed line), bumping (indicated by the green dashed line), and boosting (indicated by the orange dashed line) estimators with dashed vertical lines.

The histogram in Fig. 8 contains the cFGLS estimators for occupant 2. This histogram is representative of the other occupants for the average data set. We see that the original cFGLS, bagging, bumping, and boosting estimators each show some amount of bias. This is largely due to the fact that the average data set has a small sample size.

On the other hand, in Fig. 7(a), we show the histogram of cFGLS estimators for occupant 2 produced via bootstrapped dynamic data and we can see that the original cFGLS estimator (vertical red line) is nearly unbiased, indicated by the approximate Gaussian distribution around the cFGLS estimate. This is generally true for the occupants with the most variation

¹¹This is the cFGLS estimator produced using the original average and dynamic data sets and not the bootstrapped data sets.

and frequency in their voting record. However, bagging, bumping, and boosting produce estimates that are slightly biased in exchange for a reduction in estimator variance (33).

Occupant 2 is representative of players who prefer to focus on lighting satisfaction as opposed to winning, whereas occupant 8 is representative of players who prefer winning to lighting satisfaction. While a very active voter, frequently participating in the game, occupant 8’s voting record has little variation (the majority of the time $x_8 = 0$). Fig. 7(b) contains the cFGLS estimators for occupant 8 and we see that each of the estimators are slightly biased. Again, these estimators introduce bias in exchange for a reduction in variance.

D. Forecasting via Approximated Correlated Game

We now show the results for the correlated utility learning method. Let us use the notation

$$\hat{g}_i(x_i, x_{-i}; \{\hat{\theta}_j\}_{j \in \mathcal{K}_i}) = \sum_{j \in \mathcal{K}_i} z_{i,j} \sigma_{i,j} \psi_i(x_i, x_{-i}) + \sigma_{i,j} \hat{\theta}_j \phi_i(x_i, x_{-i}) \quad (35)$$

where recall that $\mathcal{K}_i \subset \mathcal{I}$ is the index set for the players whose parameters are used to modify player i ’s utility function in generating the correlated game and $\hat{\theta}_j$ is the estimated parameter from the utility learning methods including cOLS, cFGLS, bagging, bumping, and boosting. We use the notation $\hat{g}_i(\cdot; \{\hat{\theta}_j\}_{j \in \mathcal{K}_i})$ as shorthand.

In Table IV, we show a subset of the estimated covariance matrices obtained using the dynamic and average data sets. Using these values, we construct the following correlated game. Player 2’s utility function is modified by player 20’s

$$\hat{g}_2(x_2, x_{-2}; \mathcal{K}_2) = (z_{2,2} \sigma_{2,2} + z_{2,20} \sigma_{2,20}) \psi_2(x_2, x_{-2}) + (\sigma_{2,2} \hat{\theta}_2 + \sigma_{2,20} \hat{\theta}_{20}) \phi_2(x_2, x_{-2}) \quad (36)$$

where $\mathcal{K}_2 = \{2, 20\}$. Player 2 and 20 are passive players in that their votes tend to be strongly related to their lighting satisfaction as opposed to increasing their chances of winning.

TABLE IV

ESTIMATED COVARIANCE MATRIX FOR THE MOST ACTIVE PLAYERS USING THE (A) DYNAMIC DATA SET AND (B) AVERAGE DATA SET. THE COLORED COLUMN-ROW PAIRS INDICATE THE AGENTS WHOSE UTILITIES WE MODIFY TO CREATE THE CORRELATED GAME; THE COLUMN INDICATES THE AGENT(S) WHOSE ESTIMATED PARAMETER IS USED TO MODIFY THE ROW AGENT'S UTILITY. IN PARTICULAR, AGENT 2'S UTILITY FUNCTION IS MODIFIED BY AGENT 20'S ESTIMATED PARAMETER (IN RED), AGENT 8'S UTILITY FUNCTION IS MODIFIED BY AGENT 14'S ESTIMATED PARAMETER (IN GREEN), AND AGENT 14'S UTILITY FUNCTION IS MODIFIED BY AGENT 2'S AND AGENT 8'S ESTIMATED PARAMETER (IN BLUE). NOTE THAT AGENTS 2 AND 14 ARE ANTICORRELATED, WHERE AGENTS 8 AND 14 (AGENTS 2 AND 20) ARE POSITIVELY CORRELATED. AGENTS 2 AND 20 ARE PASSIVE PLAYERS, VOTING MORE FOR COMFORT THAN WINNING, WHERE AGENTS 8 AND 14 VOTE MORE AGGRESSIVELY

Id	2	6	8	14	20
2	0.086	0.080	-0.190	-0.248	0.059
6	0.080	7.56	8.64	9.02	0.028
8	-0.190	8.64	170.98	44.29	-0.337
14	-0.248	9.02	44.29	87.34	-0.312
20	0.059	0.028	-0.337	-0.312	0.063

(a) Average Data

Id	2	6	8	14	20
2	0.044	0.059	-2.805	-5.191	0.031
6	0.059	7.836	-16.82	0.844	-0.016
8	-2.805	-16.82	6.43×10^4	4.28×10^4	-7.60
14	-5.191	0.844	4.28×10^4	8.84×10^4	-12.59
20	0.031	-0.016	-7.60	-12.59	0.073

(b) Dynamic Data

They are also very active players, having a lot of variation in their voting record. These two players are positively correlated with one another (see the red cells in Table IV).

On the other hand, players 8 and 14 are aggressive players in that their votes tend to be much lower indicating a greater desire to win points. These players are also positively correlated (see the green cell's in Table IV). With this in mind, we modify player 8's utility function by player 14's

$$\hat{g}_8(x_8, x_{-8}; \mathcal{K}_8) = (z_{8,8}\sigma_{8,8} + z_{8,14}\sigma_{8,14})\psi_8(x_8, x_{-8}) + (\sigma_{8,8}\hat{\theta}_8 + \sigma_{8,14}\hat{\theta}_{20})\phi_8(x_8, x_{-8}) \quad (37)$$

where $\mathcal{K}_8 = \{8, 14\}$.

Player 14 is also negatively correlated with player 2. Hence, player 14's utility function is modified by player 2's and 8's utilities. That is, with $\mathcal{K}_{14} = \{2, 8, 14\}$, we have

$$\hat{g}_{14}(x_{14}, x_{-14}; \mathcal{K}_{14}) = \sum_{i \in \mathcal{K}_{14}} (z_{14,i}\sigma_{14,i}\psi_{14}(x_{14}, x_{-14}) + \sigma_{14,i}\hat{\theta}_i\phi_{14}(x_{14}, x_{-14})). \quad (38)$$

All the other players' utilities in the correlated game remain unchanged, that is, they are taken to be $\hat{g}_i = \hat{f}_i$, $i \in \mathcal{I}/\{2, 8, 14\}$.

These player combinations were selected since, through the correlated game, we aim to improve our estimators by leveraging correlations between players. In particular, the goal is to utilize information learned from players with the most variation in their votes in improving the estimates of players who consistently vote the same value or have a limited participation record.

In Table V, we present the RMSE, MAE, and MASE for the estimated correlated game $\{\hat{g}_i(\cdot; \{\hat{\theta}_j\}_{j \in \mathcal{K}_i})\}_{i \in \mathcal{I}}$ where the $\hat{\theta}_j$ s are taken to be the cOLS, bagging, boosting, and bumping estimators. Comparing these results with those in Table II, we see that correlated estimation schemes applied to the dynamic data set reduce the estimation error for almost every method. Moreover, correlated bagging outperforms bagging, the best performing ensemble method, by all three metrics. For the average data set, correlated boosting outperforms the best performing ensemble method, boosting, again by all three metrics.

TABLE V

RMSE, MAE, AND MASE OF FORECASTING USING THE ESTIMATED CORRELATED UTILITY FUNCTIONS. WE ESTIMATED CORRELATED UTILITY FUNCTIONS $\hat{g}_i(\cdot; \{\theta_j\}_{j \in \mathcal{K}_i})$ USING PARAMETERS FROM THE BAGGING, BUMPING, BOOSTING, AND cOLS METHODS FOR BOTH DATA SETS IN THE DEFAULT LIGHTING SETTING 20

<i>Dynamic</i> , \hat{g}_i	bagging	boosting	bumping	cOLS
RMSE	6.38	9.58	8.82	8.44
MAE	4.59	6.81	5.52	5.58
MASE	1.84	2.72	2.21	2.23
<i>Averaged</i> , \hat{g}_i	bagging	boosting	bumping	cOLS
RMSE	2.18	1.63	2.36	2.83
MAE	1.75	1.27	1.92	2.30
MASE	0.78	0.56	0.86	1.03

In Fig. 9, we show the forecast produced by the correlated utility learning method using the cOLS, bagging, bumping, and boosting estimators and the ground-truth test data. Fig. 9(a) and (b) shows the forecasts for the dynamic and average data sets, respectively.

What is perhaps most interesting is that for both data sets, the correlated cOLS results improve the forecasting error compared with those of cOLS and the results are not significantly different than the other ensemble methods. This can be seen in Table V and Fig. 9. The importance of this finding is that correlated cOLS has the potential to be integrated into an online algorithm. The classical cOLS can be performed online and is, thus, amenable to an online incentive design framework [3], [9]. However, as we have seen, the ensemble methods significantly outperform cOLS. Determining the estimated covariance matrix requires solving a GLS and noise covariance estimation problem [40]. Given that the estimated correlated game using cOLS parameters provides nearly the same estimation error as the ensemble methods, these methods can be adapted to estimate the correlated game parameters and then introduced into an adaptive incentive design framework. We are currently exploring this extension as the ultimate objective is to utilize the learned utilities in an incentive design framework, preferably one that can be executed in an adaptive/online manner. This will support a more robust online utility learning and incentive design algorithm.

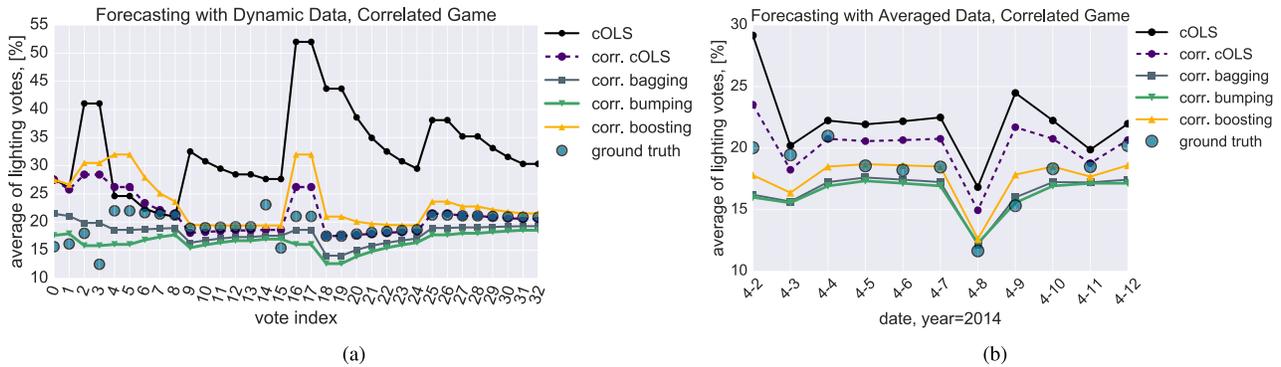


Fig. 9. Forecasting results for the correlated game using (a) dynamic data and (b) averaged data for the default lighting setting 20. For the dynamic data, the x -axis values indicate the index of when a choice was made by one or more of the occupants (i.e., when the implemented lighting setting is changed); the time from one index to the next may be several minutes to hours depending on the activity of the occupants. For the averaged data, the x -axis values are dates (month and day). The ground truth of the average of the lighting votes is depicted by the blue dots, the forecast for cOLS is depicted by the black curves, the forecast for correlated cOLS is depicted by the purple curves, the forecast for correlated bagging is depicted by the gray curves, the forecast for correlated bumping is depicted by the green curves, and the forecast for correlated boosting is depicted by the gold curves. The forecast for the robust utility learning methods is approximately near the ground truth for both data sets, while the cOLS estimates produce Nash equilibria with a large error. However, the correlated cOLS forecast significantly improves on the cOLS forecast.

VIII. DISCUSSION AND CONCLUSION

We presented a general framework for robust utility learning using a heteroskedastic inference adaptation to cGLS and we leveraged learned correlations between players in constructing a correlated utility learning framework that matches the robust utility learning errors while also being amenable to online implementation. The latter is important for integrating the proposed utility learning techniques with adaptive control or online incentive design. For example, it has been shown that static programs for encouraging energy efficiency are subject to the rebound effect in which participants often return to less efficient behavior after some time [41], [42]. By integrating our utility learning framework with incentive design, we will be able to create an adaptive model that learns how users' preferences change over time and thus generate the appropriate incentives to ensure active participation.

To demonstrate the utility learning methods, we applied them to data collected from a smart building social game we conducted where occupants vote for shared resources and participate in a lottery. We were able to estimate nearly unbiased estimators for several agent profiles and significantly reduce the forecasting error compared with cOLS. The robust utility learning framework enables us to effectively *close the loop* around smart building occupants by providing the foundation for learning a decision making model that can be integrated into the incentive or control design process. While we apply the method to smart building social game data, it can be applied more generally to scenarios with the task of inverse modeling of competitive agents and provides a useful tool for many smart infrastructure applications where learning decision making behavior is crucial.

ACKNOWLEDGMENT

The authors would like to thank C. Hsu, the applications programmer at CREST laboratory, who developed and deployed the Web portal application of the social game at UC Berkeley.

APPENDIX PROOF OF PROPOSITION 1

Proof of Proposition 1: Suppose the assumptions hold. The constraints for each player do not depend on other players' choice variables. We can hold x_{-i}^* fixed and apply [8, Proposition 3.3.2] to the i th player's optimization problem $\max\{f_i(x_i, x_{-i}^*) \mid x_i \in \mathcal{C}_i\}$. Since each f_i is concave and each \mathcal{C}_i is a convex set, x_i^* is a global optimum of the i th player's optimization problem under the assumptions. Since this is true for each of the $i \in \{1, \dots, n\}$ players, x^* is a Nash equilibrium. ■

REFERENCES

- [1] P. Bolton and M. Dewatripont, *Contract Theory*. Cambridge, MA, USA: MIT Press, 2005.
- [2] J.-J. Laffont and D. Martimort, *The Theory Incentives: The Principal-Agent Modeling*. Princeton, NJ, USA: Princeton Univ. Press, 2002.
- [3] L. J. Ratliff, "Incentivizing efficiency in societal-scale cyber-physical systems," Ph.D. dissertation, Elect. Eng. Comput. Sci., Univ. California, Berkeley, Berkeley CA, USA, 2015.
- [4] A. Aswani and C. Tomlin, "Incentive design for efficient building quality of service," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput.*, 2012, pp. 90–97.
- [5] L. J. Ratliff, R. Dong, H. Ohlsson, and S. S. Sastry, "Incentive design and utility learning via energy disaggregation," *IFAC Proc. Vols*, vol. 47, no. 3, pp. 3158–3163, 2014.
- [6] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n -person games," *Econometrica*, vol. 33, no. 3, pp. 520–534, Jul. 1965.
- [7] L. J. Ratliff, S. A. Burden, and S. S. Sastry, "On the characterization of local Nash equilibria in continuous games," *IEEE Trans. Autom. Control*, vol. 61, no. 8, pp. 2301–2307, Aug. 2016.
- [8] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA, USA: Athena Scientific, 1999.
- [9] L. J. Ratliff, M. Jin, I. C. Konstantakopoulos, C. Spanos, and S. S. Sastry, "Social game for building energy efficiency: Incentive design," in *Proc. 52nd Allerton Conf. Commun., Control, Comput.*, 2014, pp. 1011–1018.
- [10] M. Jin, L. J. Ratliff, I. Konstantakopoulos, C. Spanos, and S. Sastry, "Rest: A reliable estimation of stopping time algorithm for social game experiments," in *Proc. ACM/IEEE 6th Int. Conf. Cyber-Phys. Syst.*, 2015, pp. 90–99.
- [11] A. Keshavarz, Y. Wang, and S. Boyd, "Imputing a convex objective function," in *Proc. IEEE Int. Symp. Intell. Control*, Sep. 2011, pp. 613–619.
- [12] D. A. Freedman, *Statistical Models: Theory and Practice*. Cambridge, U.K.: Cambridge Univ. Press, 2009.

- [13] F. Cribari-Neto, "Asymptotic inference under heteroskedasticity of unknown form," *Comput. Statist. Data Anal.*, vol. 45, no. 2, pp. 215–233, 2004.
- [14] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning, Data Mining, Interference, and Prediction*. New York, NY, USA: Springer-Verlag, 2009.
- [15] R. Tibshirani and K. Knight, "Model search and inference by bootstrap 'bumping,'" *J. Comput. Graph. Statist.*, vol. 8, no. 4, pp. 671–686, 1999.
- [16] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.*, vol. 29, no. 5, pp. 1189–1232, 2001.
- [17] C. M. Hurvich, J. S. Simonoff, and C.-L. Tsai, "Smoothing parameter selection in nonparametric regression using an improved Akaike information criterion," *J. Roy. Statist. Soc., Ser. B*, vol. 60, no. 2, pp. 271–293, 1998.
- [18] H. Prakash, R. Narayanam, D. Garg, and Y. Narahari, *Game Theoretic Problems in Network Economics and Mechanism Design Solutions*. London, U.K.: Springer, 2009.
- [19] S. Berry, J. Levinsohn, and A. Pakes, "Automobile prices in market equilibrium," *Econometrica*, vol. 63, no. 4, pp. 841–890, 1995.
- [20] D. Bertsimas, V. Gupta, and I. C. Paschalidis, "Data-driven estimation in equilibrium using inverse optimization," *Math. Program.*, vol. 153, no. 2, pp. 595–633, 2015.
- [21] S. Gabriel, A. Conejo, J. Fuller, B. Hobbs, and C. Ruiz, *Complementarity Modeling in Energy Markets* (International Series in Operations Research & Management Science). New York, NY, USA: Springer-Verlag, 2013.
- [22] A. Ledvina and R. Sircar, "Dynamic bertrand oligopoly," *Appl. Math. Optim.*, vol. 63, no. 1, pp. 11–44, 2011.
- [23] R. J. Aumann *et al.*, "Subjectivity and correlation in randomized strategies," *J. Math. Econ.*, vol. 1, no. 1, pp. 67–96, 1974.
- [24] M. Jin, R. Jia, Z. Kang, I. C. Konstantakopoulos, and C. Spanos, "Presencesense: Zero-training algorithm for individual presence detection based on power monitoring," in *Proc. 1st ACM Conf. Embedded Syst. Energy-Efficient Buildings*, 2014, pp. 1–10.
- [25] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler, "Design and implementation of a high-fidelity AC metering network," in *Proc. Int. Conf. Inf. Process. Sensor Netw.*, Apr. 2009, pp. 253–264.
- [26] A. Aswani *et al.*, "Identifying models of HVAC systems using semi-parametric regression," in *Proc. Amer. Control Conf.*, 2012, pp. 3675–3680.
- [27] M. Jin, N. Bekiaris-Liberis, K. Weekly, C. Spanos, and A. Bayen, "Sensing by proxy: Occupancy detection based on indoor CO₂ concentration," in *Proc. UBIComm*, 2015, pp. 1–14.
- [28] M. Boman, P. Davidsson, N. Skarmeeas, K. Clark, and R. Gustavsson, "Energy saving and added customer value in intelligent buildings," in *Proc. 3rd Int. Conf. Pract. Appl. Intell. Agents Multi-Agent Technol.*, 1998, pp. 505–517.
- [29] D. Bourgeois, C. Reinhart, and I. Macdonald, "Adding advanced behavioural models in whole building energy simulation: A study on the total energy impact of manual and automated lighting control," *Energy Buildings*, vol. 38, no. 7, pp. 814–823, 2006.
- [30] Z. Song, "Collaborative building control to optimize energy saving and improve occupants' Experience," *ASHRAE Trans.*, vol. 119, no. 1, pp. 1–8, 2013.
- [31] M. Bang, C. Torstensson, and C. Katzeff, *The PowerHouse: A Persuasive Computer Game Designed to Raise Awareness of Domestic Energy Consumption*. Berlin, Germany: Springer, 2006.
- [32] J. Simon, M. Jahn, and A. Al-Akkad, "Saving energy at work: The design of a pervasive game for office spaces," in *Proc. 11th Int. Conf. Mobile Ubiquitous Multimedia*, 2012, Art. no. 9.
- [33] S. Li, K. Deng, and M. Zhou, "Social incentive policies to engage commercial building occupants in demand response," in *Proc. IEEE Inter. Conf. Autom. Sci. Eng.*, Aug. 2014, pp. 407–412.
- [34] G. Taguchi, E. A. Elsayed, and T. C. Hsiang, *Quality Engineering in Production Systems*. New York, NY, USA: McGraw-Hill, 1989.
- [35] A. Tversky and D. Kahneman, "The framing of decisions and the psychology of choice," *Science*, vol. 211, no. 4481, pp. 453–458, 1981.
- [36] M. Gendreau, "On the location of eigenvalues of off-diagonal constant matrices," *Linear Algebra Appl.*, vol. 79, pp. 99–102, Jul. 1986.
- [37] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *Int. J. Forecasting*, vol. 22, no. 4, pp. 679–688, 2006.
- [38] S. D. Flåm, "Solving non-cooperative games by continuous subgradient projection methods," in *System Modelling and Optimization*. Berlin, Germany: Springer-Verlag, 1990, pp. 115–123.
- [39] I. C. Konstantakopoulos, L. Ratliff, M. Jin, C. Spanos, and S. S. Sastry, "Inverse modeling of non-cooperative agents via mixture of utilities," in *Proc. IEEE Conf. Decision Control*, Dec. 2016, pp. 6327–6334.
- [40] R. Isermann and M. Münchhof, *Identification of Dynamic Systems: An Introduction with Applications*. Berlin, Germany: Springer-Verlag, 2011.
- [41] J. A. Laitner, "Energy efficiency: Rebounding to a sound analytical perspective," *Energy Policy*, vol. 28, nos. 6–7, p. 471–475, 2000.
- [42] L. Schipper and M. Grubb, "On the rebound? feedback between energy intensities and energy uses in IEA countries," *Energy Policy*, vol. 28, nos. 6–7, pp. 367–388, 2000.



Ioannis C. Konstantakopoulos received the Diploma (Hons.) in electrical and computer engineering from the University of Patras, Patras, Greece, in 2012, and the M.S. degree in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 2014, where he is currently pursuing the Ph.D. degree.

His current research interests include statistical learning, machine learning, game theory, and optimization.



Lillian J. Ratliff (S'08–M'15) received the B.S. degree in mathematics and the B.S. and M.S. degrees in electrical engineering (EE) from the University of Nevada, Las Vegas, NV, USA, in 2008 and 2010, respectively, and the Ph.D. degree from the University of California, Berkeley, CA, USA, in 2015.

She was a Post-Doctoral Researcher in Electrical Engineering and Computer Sciences with the University of California, Berkeley. She is an Assistant Professor in EE with the University of Washington, Seattle, WA, USA. Her current research interests

include the intersection of game theory, optimization, and statistical learning.

Dr. Ratliff is a recipient of the National Science Foundation Graduate Research Fellowship and the CISE Research Initiation Initiative Award.



Ming Jin received the B.Eng. (Hons.) degree in electronic and computer engineering from the Hong Kong University of Science and Technology, Hong Kong. He is currently pursuing the Ph.D. degree in electrical engineering and computer science with the University of California, Berkeley, CA, USA.

His current research interests include the societal-scale cyber physical systems, from buildings to the microgrid, by employing modeling, optimization, and control theories to design new algorithms

and mechanisms.

Mr. Jin was a recipient of the Best Paper Award in the 2015 International Conference on Mobile Ubiquitous Computing, Systems, Services, and Technologies.



S. Shankar Sastry (S'79–M'80–SM'95–F'95) received the B.Tech. degree from IIT Bombay, Mumbai, India, the M.S. degree, the M.A. degree in mathematics, and the Ph.D. degree in engineering from University of California at Berkeley, Berkeley, CA, USA.

He is the Dean of Engineering with the University of California, Berkeley, and the Faculty Director of the Blum Center for Developing Economies, Berkeley. He has served on the faculties of MIT, Cambridge, MA, USA, and Harvard University,

Cambridge. His current research interests include the design of resilient network control systems, autonomous systems, computer vision, nonlinear and adaptive control, and hybrid and embedded systems.

Dr. Sastry is a member of the National Academy of Engineering, Washington, DC, USA, and the American Academy of Arts and Sciences, Cambridge. He received the honorary M.A. from Harvard University and honorary doctorates from the Royal Swedish Institute of Technology, Stockholm, Sweden, and the University of Waterloo, Waterloo, ON, Canada. He has been a member of the Air Force Scientific Advisory Board and the Defense Science Board, among other national boards. He is currently on the Corporate Boards of C3-IoT, Redwood City, CA, USA, and HCL Technologies, India and on the Scientific Advisory Boards of Interwest Partners, GE, and Lockheed Martin. He is also a member of the UN Secretary General's Scientific Advisory Board.



Costas J. Spanos (S'77–M'85–SM'92–F'98) received the Diploma in electrical engineering from the National Technical University of Athens, Athens, Greece, in 1980, and the M.S. and Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, USA, in 1981 and 1985, respectively.

In 1988, he joined the faculty with the Department of Electrical Engineering and Computer Sciences (EECS), University of California, Berkeley, CA, USA, where he has served as the Director of the Berkeley Microlab, the Associate Dean for Research in the College of Engineering, and as the Chair of EECS. He is involved in statistical analysis in the design and fabrication of integrated circuits, on novel sensors and computer-aided techniques in semiconductor manufacturing, and on statistical data mining techniques for energy applications. He has participated in two successful startup companies, Timbre Technologies Inc., Fremont, CA, USA, (acquired by Tokyo Electron) and OnWafer Technologies, San Jose, CA, USA, (acquired by KLA-Tencor). He is currently the Director of the Center of Information Technology Research in the Interest of Society and the Chief Technical Officer for the Berkeley Educational Alliance for Research in Singapore.