Modeling Thermostat Adjustment Behavior in Residential Communities During Eco-feedback Energy Interventions

Jaehyun GO[1, 2], Huijeong KIM[1, 2], Ilias BILIONIS[3], Panagiota KARAVA[1, 2*]

[1] Lyles School of Civil Engineering, Purdue University,
West Lafayette, IN, USA

[2] Center for High Performance Buildings,
Ray W. Herrick Laboratories, Purdue University,
West Lafayette, IN, USA
go16@purdue.edu, kim2683@purdue.edu, pkarava@purdue.edu

[3] School of Mechanical Engineering, Purdue University,
West Lafayette, IN, USA
ibilion@purdue.edu

* Corresponding Author

ABSTRACT

Eco-feedback interventions have been proven to positively impact user behavior toward heating and cooling energy conservation in the residential sector. Understanding human decision-making in thermostat adjustment during interventions enables policy makers (building operators or community managers) to anticipate occupant behaviors and corresponding outcomes when planning and managing the intervention strategies. This paper presents a modeling methodology for deriving household decisions using MySmartE, an eco-feedback and social gaming platform designed to engage residents in understanding and reducing their home energy use. The proposed model is a utility function that captures preferred heating setpoints in different smart thermostat modes (i.e., 'home', 'sleep', 'away', and 'hold') for different households. In this study, a social game with three levels of rewards was implemented in three different communities over a heating season. The parameters of the utility function were inferred using hierarchical Bayesian calibration with the collected household data. The results show that the developed model effectively captures the behavioral responses of residents for different reward sizes. It is proposed that this model serves as a foundation for decision simulation and energy intervention policy design.

1. INTRODUCTION

In the U.S., average space heating and cooling energy consumption in residential buildings accounts for 42% and 9% of the annual energy usage (EIA, 2023). A significant portion of the energy used for heating and cooling in residential buildings is related to user behavior and it has been incorporated in energy efficiency interventions (Pan et al., 2017). Normative feedback effectively induces behavioral changes and experiments have demonstrated that personalized comparisons of energy use led to reduced consumption in high-energy homes (Wolske et al, 2020). Dolan and Metcalfe (2015) found that social norms could reduce energy consumption by approximately 6%. Kim et al. (2022) developed MySmartE, a platform that provides residents with an energy score for their heating and cooling setpoints and enables peer comparisons via a tablet display. This form of normative feedback was proven effective during the cooling season. When social norms are combined with gamification strategies, the effect of norms on user engagement is enhanced. Kim et al (2024) also deployed social games in a multi-family residential community where users collaborate to meet weekly goals and receive rewards, thereby fostering further behavioral changes.

Although these strategies can be effective, a quantitative approach is needed to evaluate and redesign the interventions (Anderson and Lee, 2016). Kim et al (2024) used utility functions to simulate the influence of game rewards and game design elements on indoor temperatures, which serve as decision variables of individual households. The response of

residents in multi-family apartments to norms or games was quantitatively assessed using the coefficients of these utility functions. It should be noted, however, that while indoor temperature serves as an indicator of the overall impact of behavioral changes, it is the heating and cooling setpoints that residents can actually adjust. Furthermore, when evaluating these strategies from a cost perspective, it is crucial to quantify how behavioral changes convert into energy consumption. Building energy simulation (BES) programs require heating/cooling setpoint temperatures for target spaces. A thermostat adjustment behavior model capable of outputting these setpoint temperatures can be beneficial in terms of adaptability when used as input for BES programs.
Meanwhile, defining the setpoint as the decision variable becomes more intricate with the adoption of smart thermostats. These devices, which feature various schedule modes (e.g., 'home', 'sleep', 'away', and 'hold'), allow residents to operate systems efficiently but also introduce challenges in identifying adjustments to heating and cooling setpoints under different modes following an energy intervention.

To address these challenges, this paper presents a new model that accurately represents residents' decisions and facilitates comprehensive evaluation and analysis through user-specific and mode-specific utility parameters. More specifically, the model has been developed with the following considerations:

- The decision variable of the utility function is defined as the heating and cooling setpoint.
- A hierarchical utility function is formulated to estimate the setpoints across different thermostat modes.
- A single model structure calibrates the utility function parameters for multiple users in three residential communities.

In this paper, Section 2 introduces a field experiment of a competitive lottery-style game as an energy intervention strategy implemented in three actual communities. Section 3 elaborates on the proposed utility function, pertaining to its structure and calibration methods for inferring individual and mode-specific utility parameters. Section 4 presents the analysis and comparison of model-calibrated parameters across various households and communities. Section 5 concludes the paper presenting the main findings along with limitations and future directions.

2. FIELD EXPERIMENT OVERVIEW

2.1 Energy Intervention Strategy

To implement human behavior interventions toward heating/cooling energy conservation in residential communities, we designed a competitive community game offering various reward sizes. We utilized an eco-feedback and gaming platform developed by our team, MySmartE (Kim et al, 2023). As shown in Figure 1, MySmartE features an interface for adjusting the thermostat setpoint and a layout for checking the game status, thus, allowing data collection on resident behavior changes and preferred setpoint temperatures/schedules. The game implemented in this study is a competitive lottery game, and the procedure is as follows. Households are assigned an energy score with a range of 0 to 100 based on their heating and cooling setpoints (Kim et al, 2022). This score helps residents recognize how efficiently they set heating/cooling setpoints. The energy score is continuously updated daily through the accumulated thermostat setpoint settings during each calendar week. Each day, the residents receive lottery tickets based on their energy scores. Every Sunday, we spin the roulette with the tickets accumulated over the week, and the winning resident receives the prize for that week. This reward mechanism encourages residents to reduce their energy consumption by adjusting their behaviors. During the game period, we examined whether varying the reward sizes ($25, $50, and $100) induces behavioral changes.

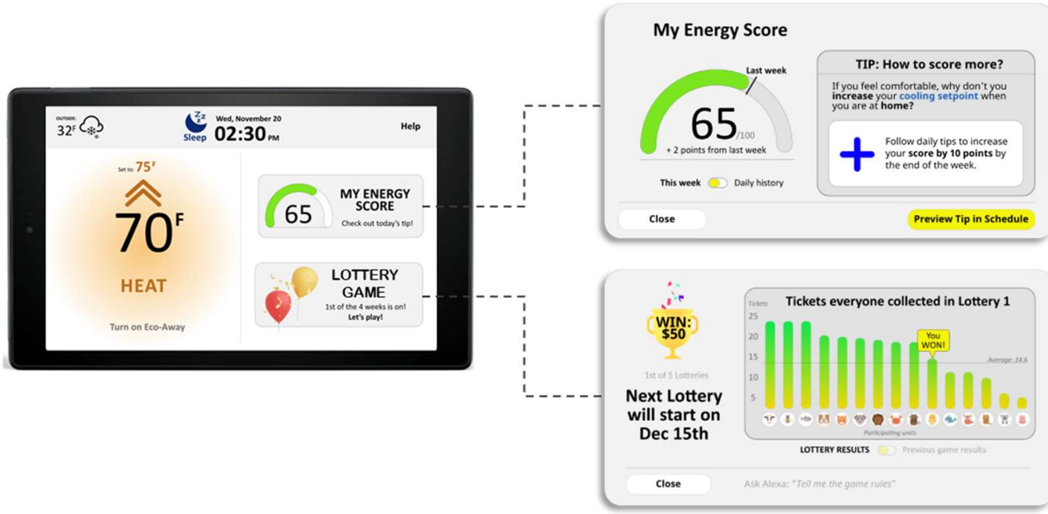

Figure 1: MySmartE resident interface

2.2 System Deployment

In this paper, MySmartE was implemented in three distinct communities during the heating period from December 18, 2023, to March 10, 2024. The first community (Community A), located in Fort Wayne, IN, is composed of two multi-family apartment buildings. In this community, we recruited 23 households. Each unit has an electric heat pump with a backup heater for space heating and residents are subsidized for their heating costs. In the second community (Community B) with single-family detached houses located in South Bend, IN, we recruited 10 participants. Each house is heated using a gas furnace, and the residents are responsible for their heating costs. In Community C, located in New Albany, IN, we recruited 24 households in multi-family apartment buildings that use furnaces and tenants pay their electricity bills but receive energy assistance for heating.

A smart thermostat and a wall-mounted tablet were installed in each unit. The thermostats control the heating devices and collect data on temperature, occupancy, and equipment operation in a 5-minute time interval. The thermostats provide a scheduling feature that allows residents to set different temperatures for modes such as 'home', 'away', and 'sleep'. Additionally, residents can manually override their scheduled modes by using a 'hold' mode to maintain specific setpoints for specific durations. The tablet, which is connected to the smart thermostat, allows residents to easily adjust the set temperature, receive personalized eco-feedback, and access information about the game (Kim et al, 2022). The field experiment was carried out in accordance with the ethical guidelines of Purdue's Institutional Review Board (IRB Protocol #: 1702018811), ensuring that all research involving human subjects was conducted ethically.

Figure 2 illustrates the weekly average heating setpoints associated with different reward levels for all communities. For this paper, we included households with data availability above 95% who participated in the game at least four times. This is 13 out of 23 users in Community A, 6 out of 10 users in Community B, and 13 out of 24 users in Community C, resulting in 32 users in total. As shown in Figure 2, residents 108 and 129 displayed a noticeable decrease in heating setpoints during periods with higher rewards, while other users showed only slight reductions or no changes. This variability in behavior highlights the importance of behavior modeling in designing intervention strategies.

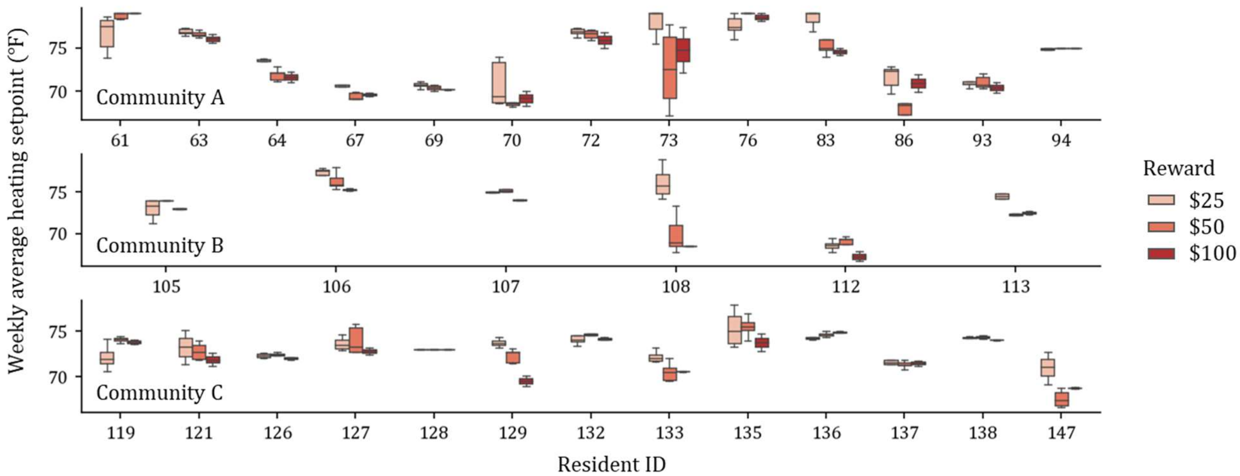

Figure 2: Households' weekly average heating setpoints during the lottery game

3. MODELLING METHODOLOGY

Residents participate in games with weekly lotteries offering rewards that encourage them to adjust their thermostat setpoints towards energy conservation. In this study, we conducted a weekly lottery game to examine the influence of expected rewards on residents' thermostat adjustment behaviors, with three reward levels: $25, $50, and $100. The game was implemented via the MySmartE platform, which facilitated data collection on game details, residents' thermostat mode schedules, and specific heating setpoints for each mode. As the potential rewards increased, participants were hypothesized to decrease their heating setpoints to accumulate more lottery tickets, thereby enhancing their probability of winning. Residents can change the heating setpoints in scheduled modes ('home', 'sleep', and 'away') or override the 'hold' mode to adjust the setpoints. This paper models the behavioral changes of residents as illustrated in Figure 3. The model is a utility function of heating setpoints across different thermostat modes. This utility is then used to derive the heating setpoints that a resident is likely to select for each thermostat mode. Additionally, this model allows inferring household responses to various rewards in the lottery game.

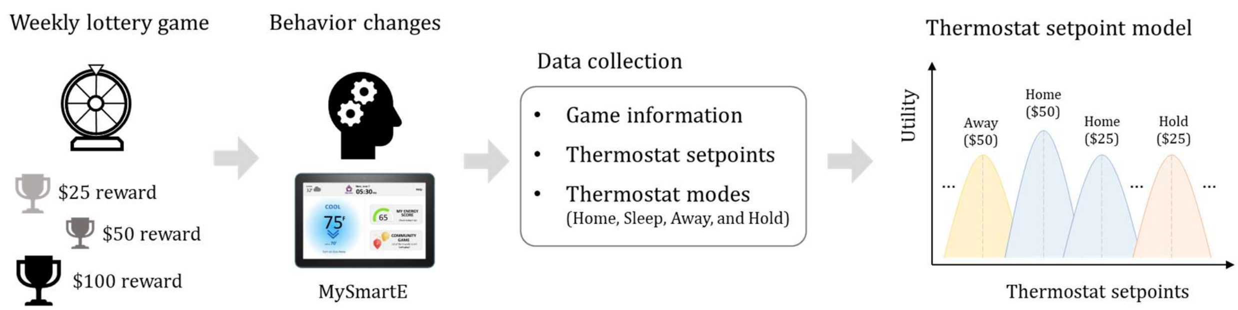

Figure 3: Modeling framework for the energy intervention strategy

3.1 Utility Function

In this paper, it is assumed that the total utility is composed of satisfaction with the thermal environment and satisfaction from the expected rewards of the game, as expressed in Equation (1).

$$U(T^s) = U^C(T^s) + U^G(T^s) \tag{1}$$

T^s indicates the heating setpoint and U^C represents the resident's satisfaction with the thermal environment considering energy costs, thermal comfort, and individual feedback (e.g., energy score). This is represented by a quadratic function, as shown in Equation (2), where satisfaction increases as it approaches T^{cs}.

$$U^C(T^s) = -0.5(T^s - T^{cs})^2 \tag{2}$$

T^{cs} represents the hypothesized setpoints, which are assumed to be determined by residents based on factors such as energy costs, thermal comfort, and satisfaction derived from energy scores, rather than the rewards from the game. The satisfaction from the expected rewards of the game (U^G) is expressed as shown in Equation (3)

$$U^G(T^s) = \chi I p_w(T^s) \tag{3}$$

The satisfaction from game rewards is calculated by the product of the size of the game reward (I), the probability of winning (p_w), and the individual's response coefficient to game rewards (χ). The probability of winning is determined by the resident's accumulated number of tickets relative to the total accumulated tickets on the last day of the week. The accumulated number of tickets is negatively correlated with the heating setpoint; the lower the heating setpoint, the higher the probability of winning. In this paper, it is assumed that reducing the heating setpoint proportionally increases the probability of winning, and Equation (3) is modified to be Equation (4).

$$U^G(T^s) = \chi I(\alpha T^s + \beta) \tag{4}$$

The coefficients α and β describe the linear relationship between the heating setpoint and the probability of winning. We assume that the residents set the heating setpoint to maximize the entire utility function as shown in Equation (5).

$$T^{s*} = \arg\max_{T^s} U(T^s) \tag{5}$$

Since U^C in Equation (2) is a convex function, the heating setpoint (T^{s*}) that sets the derivative of the utility function in Equation (1) to zero is the one that maximizes the utility function in Equation (6).

$$T^{s*} = T^{cs} + \hat{\chi} I \tag{6}$$

$\hat{\chi}$ is the product of χ and α. Residents will decide to increase or decrease the heating setpoint by an amount equal to the product of $\hat{\chi}$ and I added to T^{cs}. Since the game was implemented in the heating season, a decrease in the heating setpoint as game rewards increase represents a positive effect on the game. Therefore, a negative value for $\hat{\chi}$ indicates a positive response to the game, while a value of zero signifies no response to the game. A positive value may seem counterintuitive, as it does not make sense to increase heating setpoints to obtain game rewards. This suggests a need to improve the distribution assumptions for this parameter in future work.

3.2 Hierarchical Structure

Households have different preferred heating setpoints, which may vary according to the thermostat mode (i.e., 'home', 'sleep', 'away', and 'hold'). Thus, the adjustments made to the heating setpoints in any mode to increase the probability of winning the game can also vary. To reflect this, the utility parameters should be varied depending on the resident and the thermostat mode as shown in Equation (7)-(9).

$$U_{i,m}(T^s_{i,m}) = U^C_{i,m}(T^s_{i,m}) + U^G_{i,m}(T^s_{i,m}) \tag{7}$$

$$U^C_{i,m}(T^s_{i,m}) = -0.5(T^s_{i,m} - T^{cs}_{i,m})^2 \tag{8}$$

$$U^G_{i,m}(T^s_{i,m}) = \chi_{i,m} I(\alpha T^s_{i,m} + \beta) \tag{9}$$

$T^s_{i,m}$ and $U_{i,m}$ represent the heating setpoints and the corresponding satisfaction for resident i in thermostat schedule mode m. Although individual households may have different preferred setpoints for each mode, there is usually a common order of preference for the heating setpoints across the modes, with 'home' mode having the highest setpoint and 'away' having the lowest. Therefore, this paper assumes there are common heating setpoint/behavior levels across the different modes. This similarity within the same mode and differences between different modes are modeled using non-centered parameterization.

$$T_{i,m}^{cs} = \mu_{T_m^{cs}} + \sigma_{T_m^{cs}} \cdot \zeta_{T_{i,m}^{cs}} \tag{10}$$

$$\hat{\chi}_{i,m} = \mu_{\hat{\chi}_m} + \sigma_{\hat{\chi}_m} \cdot \zeta_{\hat{\chi}_{i,m}} \tag{11}$$

$\mu_{T_m^{cs}}$ and $\mu_{\hat{\chi}_m}$ represent the average common setpoints and game reward response coefficient that users have. $\zeta_{T_{i,m}^{cs}}$ and $\zeta_{\hat{\chi}_{i,m}}$ serve as offset variables that allow individual residents to have their own values deviating an amount $\sigma_{T_m^{cs}}$ and $\sigma_{\hat{\chi}_m}$ from the mean. With $T_{i,m}^{cs}$ and $\hat{\chi}_{i,m}$, we can estimate $\hat{T}_{i,m}^{s*}$, which maximizes the utility function as shown in Equation (12).

$$\hat{T}_{i,m}^{s*} = T_{i,m}^{cs} + \hat{\chi}_{i,m} I \tag{12}$$

3.3 Bayesian Calibration

We can determine the utility function representing specific user behaviors by estimating unobserved parameters from data on households' heating setpoints. Specific users' thermal preferences and their responsiveness to game rewards can vary due to a variety of factors. In this regard, we estimated the coefficients as probabilistic distributions using Bayesian inference. To reflect the discrepancy between the actual residents' $T_{i,m}^{s}$ and the estimated $\hat{T}_{i,m}^{s*}$, we have defined the likelihood of $T_{i,m}^{s}$ as a normal distribution with $\sigma_{T_m^s}$ in Equation (13), as follows:

$$T_{i,m}^{s} \sim N(\hat{T}_{i,m}^{s*}, \sigma_{T_m^s}) \tag{13}$$

The structure of the hierarchical Bayesian model is depicted in Figure 4. In this paper, the priors for μ and ζ are set to a normal distribution (i.e., $\mu \sim N(0, 10)$ and $\zeta \sim N(0, 10)$), and σ to a half-normal distribution (i.e., $\sigma \sim \mathrm{HalfNormal}(2)$). The data used for calibration includes heating setpoint ($T_{i,m,d}^{s}$) and the size of game rewards ($I_{i,d}$) for *i* in {1, …, 32} and *m* in {1, …, 4} corresponding to a resident identification number and thermostat mode index, respectively. Bayesian model references these indexes to estimate the parameters of the utility function for a particular user's specific thermostat mode. To avoid the substantial computational load that comes with calibrating using 5-minute interval data, we aggregated the data to hourly intervals. We represented each hour by the heating setpoint temperature of the most frequently occurring thermostat mode. Since rewards are set on a weekly basis, aggregating the data to hourly intervals did not affect the reward information. The Bayesian inference was performed using the Python PyMC package (Abril-Pla O et al, 2023). Latent parameters of utility function were estimated using Markov chain Monte Carlo (MCMC) with NumPyro JAX NUTS sampler. We used four MCMC chains, and each chain ran 4000 total simulations, with the first 2000 simulations used for parameter tuning.

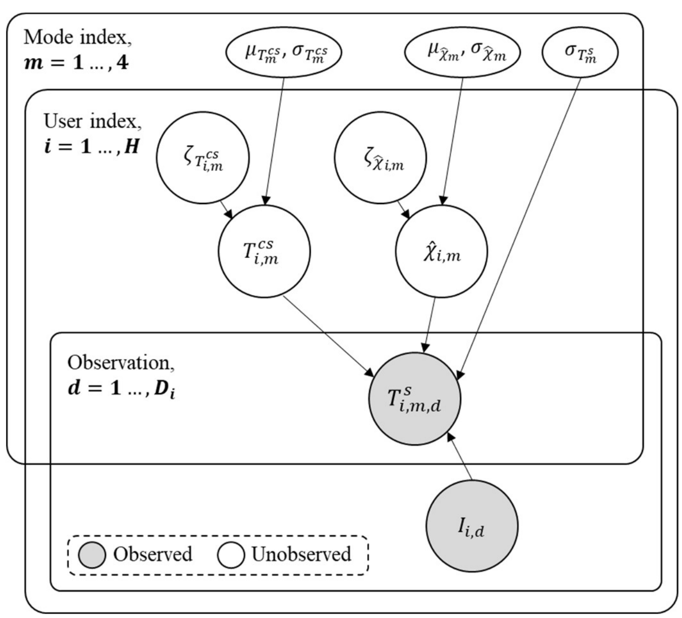

Figure 4: Hierarchical structure of the Bayesian model

4. MODEL CALIBRATION RESULTS

4.1 Model Validation

The fit of the hierarchical Bayesian model was validated using posterior predictive checks and the Bayesian p-value (Gelman et al, 2013). The posterior predictive distribution is as follows:

$$p(y_{rep}|y_{obs}) = \int p(y_{rep}|\theta)p(\theta|y_{obs})d\theta \tag{13}$$

y_{obs} represents the observed data, θ the model parameters, $p(y_{rep}|\theta)$ the likelihood of the data sampled given the prior parameter distributions, and $p(\theta|y_{obs})$ the posterior distribution of the model parameters based on the observed data. This represents the posterior predictive distribution, allowing for a comparison of how well the replicated posterior predictive distribution with the calibrated parameters explains the observed data. Figure 5-(a) displays the replicated distribution for an example draw. The replicated distribution mostly covers the observed data range. However, the observed data shows a significantly higher frequency at certain values like 70 and 80 °F, which the replicated distribution does not cover precisely. This could be because it is difficult to estimate a normal distribution with a single mode for users who have more than one preferred heating setpoint.
A quantitative comparison can be further captured through the Bayesian p-value. The equation for the Bayesian p-value (p_B) is:

$$p_B = P(T(y_{rep}, \theta) > T(y_{obs}, \theta)|y_{obs}) \tag{14}$$

$T(y_{rep}, \theta)$ and $T(y_{obs}, \theta)$ represent specific statistics (i.e., mean and standard deviation) of y_{rep} and y_{obs}. The model fit is considered more accurate when the probability of $T(y_{rep}, \theta)$ being greater than $T(y_{obs}, \theta)$ approaches 0.5. If it is close to 0 or 1, it indicates that the model may not be accurate. Figure 5-(b) compares the distribution of statistics (blue histogram) for 4,000 draws of an example MCMC chain against the observed data's statistics (black line). The mean and standard deviation values indicate that the model captures them well.

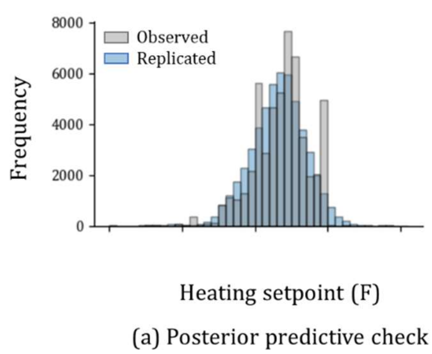

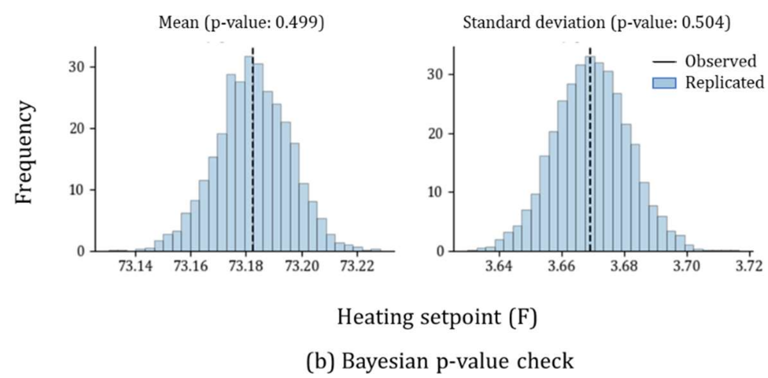

Figure 5: Posterior predictive check and Bayesian p-value

Figure 6 presents the disaggregated posterior predictive distribution for 10 example residents for each thermostat mode. Each subplot displays the observed data alongside the posterior predictive distribution, represented by histograms and kernel density estimate plots, respectively, each of which represents a probability density. The absence of a distribution indicates that the user did not use that specific mode. The model appears to capture the heating setpoints that were set by most residents across different modes. In the scheduled modes ('home', 'sleep', and 'away'), the residents show little variation in setpoints. This is because the scheduled modes often maintain the set heating setpoints with few changes. In contrast, in the 'hold' mode, a variety of values are being used, and the posterior predictive distribution is well aligned with the observed data. This variability arises because each use of the 'hold' mode involves overriding the schedule and resetting the setpoint.

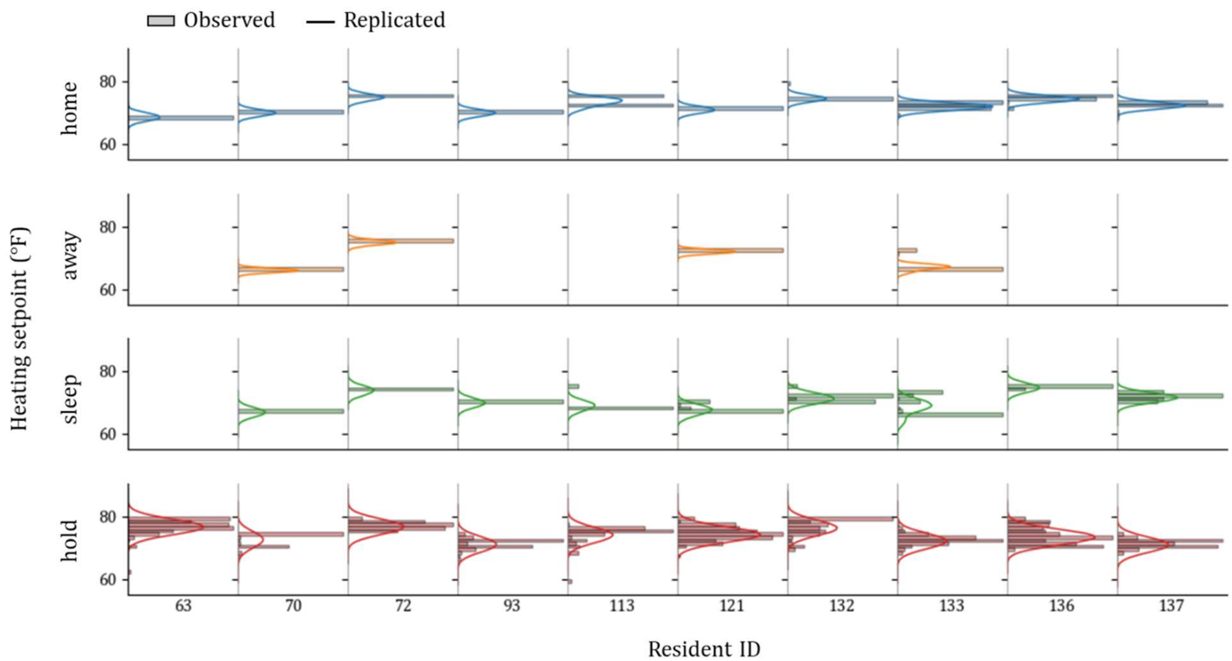

Figure 6: Posterior predictive check for individual residents with different thermostat modes

4.2 Model Parameters

In this paper, two major latent parameters were estimated (i.e., $\hat{\chi}$ and T^{cs}). $\hat{\chi}$ measures residents' willingness to adjust heating setpoints in response to expected game rewards, while T^{cs} is assumed to be the heating setpoint that maximizes user satisfaction when no game reward is offered. The analysis in this paper focused on $\hat{\chi}$ since the calibration data only involved scenarios with game rewards. A negative value of $\hat{\chi}$ means residents tend to reduce the heating setpoint in response to expected rewards, whereas a value close to zero suggests the resident is unresponsive to the game. Figure 7 shows the distribution of $\hat{\chi}$ for all residents in three communities. The absence of a boxplot for any thermostat mode indicates that the resident did not use that particular mode.

In 'home' mode, no residents in Community A responded to the game, and responses in Communities B and C were not substantial. This may suggest that in the 'home' mode, the change of game rewards may not be enough to encourage them to adjust heating. In 'away' mode, there was no noticeable response across all communities. This could be because MySmartE assists in setting back the heating setpoint when residents are away. Residents might not respond to game rewards since they are already practicing good behavior (lower heating setpoints). However, in 'sleep' mode, some residents showed a willingness to reduce the heating setpoint. This could be because residents are less sensitive to thermal comfort in 'sleep' mode compared to 'home' mode and have more room to reduce the heating setpoint compared to 'away' mode. In 'hold' mode, a significant number of households across all communities showed a positive response to the game. This is potentially because 'hold' mode allows residents to override the schedule and set the temperature each time, making it relatively easy for them to respond to game rewards. The results showed that thermal preferences and the willingness to adjust heating setpoints can vary across different modes, each providing different levels of adjustability for user actions. The developed model can effectively identify user interaction characteristics between the system and residents, which need to be carefully considered when designing and improving intervention strategies in future research.

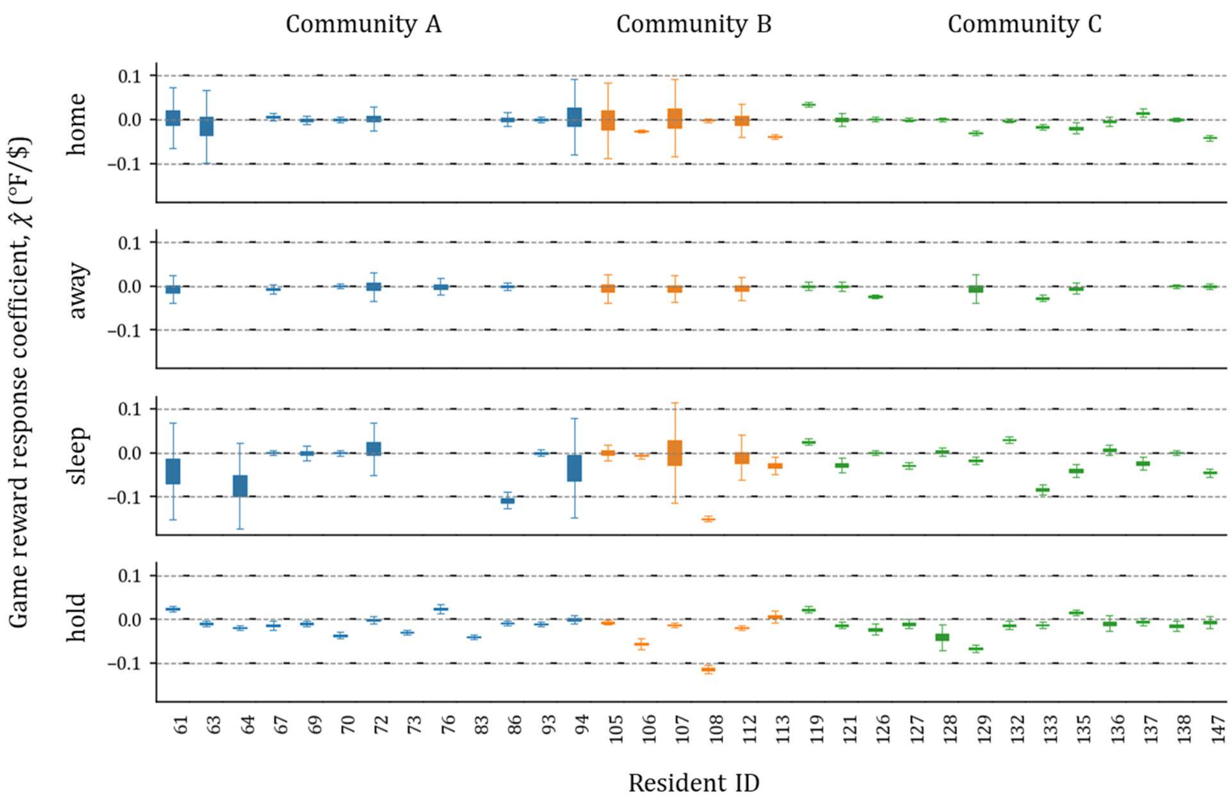

Figure 7: Game reward response coefficient for all residents

5. CONCLUSIONS

In this study, we developed a model to examine user behavior in response to a lottery-based social game. The model employs a utility function to systematically capture the preferred heating setpoints and uses a hierarchical Bayesian calibration to infer its unobserved parameters. This model effectively characterizes residents' responses to the game reward in terms of their heating setpoints with various thermostat modes. However, the model did not account for the impact of external climate conditions on residents' utility.

The model developed in this paper can not only quantify but also simulate resident behavior on the impact of energy intervention strategies. In building energy simulation, the operation of HVAC equipment requires the time-series heating/cooling setpoint temperature. The output of the developed model can be assigned to thermostat mode schedules such as 'home', 'sleep', 'away', and 'hold', enabling the derivation of time-series heating/cooling setpoint profiles. We expect this model can serve as a foundation for a methodology capable of analyzing residents' thermostat adjustment behaviors in response to energy intervention strategies. Quantifying the environmental and cost impacts of the intervention strategies enables informed decisions for policy design and evaluation.

REFERENCES

Abril-Pla O, Andreani V, Carroll C, Dong L, Fonnesbeck CJ, Kochurov M, Kumar R, Lao J, Luhmann CC, Martin OA, Osthege M, Vieira R, Wiecki T, Zinkov R. (2023)

Anderson, K., & Lee, S. (2016). An empirically grounded model for simulating normative energy use feedback interventions. Applied energy, 173, 272-282.

Dolan, P., & Metcalfe, R. (2015). Neighbors, knowledge, and nuggets: two natural field experiments on the role of incentives on energy conservation. Becker Friedman Institute for Research in Economics Working Paper, (2589269).

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). Bayesian Data Analysis. Taylor & Francis Group, LLC (3rd ed.). Chapman and Hall/CRC.

Kim, H., Bilionis, I., Braun, J. E., & Karava, P. (2024). Social Energy Games for Smart and Energy-Efficient Multi-Unit Residential Buildings: Mechanism Design. ASHRAE Transactions - 2024 Winter Conference, Chicago, IL 130.

Kim, H., Bilionis, I., Karava, P., & Braun, J. E. (2023). Human decision making during eco-feedback intervention in smart and connected energy-aware communities. Energy and Buildings, 278, 112627.

Kim, H., Ham, S., Promann, M., Devarapalli, H., Bihani, G., Ringenberg, T., ... & Karava, P. (2022). MySmartE–An eco-feedback and gaming platform to promote energy conserving thermostat-adjustment behaviors in multi-unit residential buildings. Building and Environment, 221, 109252.

Pan, S., Wang, X., Wei, S., Xu, C., Zhang, X., Xie, J., ... & De Wilde, P. (2017). Energy waste in buildings due to occupant behaviour. Energy Procedia, 105, 2233-2238.

U.S. Energy Information Administration. (n.d.). Use of energy explained: Energy use in homes. Retrieved 04/18/2024, from https://www.eia.gov/energyexplained/use-of-energy/homes.php

Wolske, K. S., Gillingham, K. T., & Schultz, P. W. (2020). Peer influence on household energy behaviours. Nature Energy, 5(3), 202-212.

ACKNOWLEDGEMENT

This work was funded by the National Science Foundation under grant number 1737591.