

Thermostats for the Smart Grid: Models, Benchmarks, and Insights

Yong Liang*, David I. Levine**, and Zuo-Jun (Max) Shen***

We model two existing thermostats and one novel thermostat to see how well they operate under dynamic pricing. The existing thermostats include a traditional thermostat with set temperature goals and a rigid thermostat that minimizes cost while always keeping temperature within a rigid predetermined range. We contrast both with a novel optimizing thermostat that finds the optimal trade-off between comfort and cost. We compare the thermostats' performance both theoretically and via numerical simulations. The simulations show that, under plausible assumptions, the optimizing thermostat's advantage is economically large. Importantly, the electricity demand of the rigid thermostat (but not the optimizing thermostat) ceases to respond to electricity prices on precisely the days when the electricity grid tends to be near capacity. These are the times when demand response is the most socially valuable to avoid massive price spikes. The social benefits of the optimizing thermostat may provide incentives for utilities and regulators to encourage its adoption.

Keywords: Thermostat, Smart Grid, Dynamic Pricing

<http://dx.doi.org/10.5547/01956574.33.4.4>

1. INTRODUCTION

The electricity grid that we have been using for more than a century is being challenged. Occasional peaks in demand (for example, from air conditioning on hot afternoons) require costly year-round ancillary services. Declines in supply due to maintenance at traditional plants and (increasingly) due to the variability of renewable sources such as wind exacerbate the challenge (Joskow,

* Corresponding author. Ph.D. Candidate of Department of Industrial Engineering and Operation Research, U.C. Berkeley, CA, 94720; (510)-708-9909; E-mail: yongliang@berkeley.edu.

** Chair of the Economic Analysis and Policy Group, Chair of the Advisory Board of the Center on Evaluation for Global Action and Eugene E. and Catherine M. Trefethen Professor of Walter A. Haas School of Business, U.C. Berkeley.

*** Chancellor's Professor of Department of Industrial Engineering and Operations, U.C. Berkeley.

2011). Moreover, because users are charged flat-rate prices, they have little incentive to reduce demand when the grid is nearing capacity. As a result, there has been a great desire to move to a *Smart Grid* that can better coordinate the supply and demand of electricity (Ipakchi, 2009).

Most visions of the Smart Grid include time-varying prices (also known as real-time-pricing (RTP) or dynamic pricing) to balance supply and demand. Dynamic prices provide users with an incentive to use less electricity when the price is high, helping match demand and supply. To quote from Schweppe et al. (1980): “Homeostatic Utility Control can offer a set of advantages of both *supply follows demand* and *demand follows supply* while avoiding the majority of their pitfalls.” Although some emerging technologies allow for efficient data collection and for bi-directional communication between users and the grid, dynamic prices are only likely to have much effect on demand if users have automatic systems to manage their electricity demand.

Heating, venting, and air conditioning (HVAC) is a prime candidate to provide the needed price-elastic demand. HVAC accounts for about 31% of the total electricity usage in U.S. homes (DOEEIA, 2005). Deschênes and Greenstone (2008) predict a 33% increase in electricity consumed by HVAC appliances in the U.S. due to climate change.

Meanwhile, demand for HVAC is more flexible than other residential demand because it is often possible to shift heating or cooling away from periods of peak demand with little or no reduction in comfort. With this motivation, we look at the design of smart thermostats that can react to dynamic prices and schedule the use of HVAC appliances to maximize users’ total utility. While we focus on HVAC, this optimal control strategy can easily accommodate other electric appliances.

Unfortunately, the comprehensive survey conducted by Walker and Meier (2008) found no thermostat on the market that could react to dynamic prices. Recently, Livengood and Larson (2009) propose a conceptual energy-management system called “Energy Box” for residential and small-business buildings. They show by simulations that using the Energy Box to control room temperature enhances users’ total utility in terms of both cost and comfort.

Early work on modeling users’ comfort involved purely empirical relations between physical variables and comfort (also known as thermal sensation rating). For instance, the Predicted Percent Dissatisfied and Predicted Mean Vote proposed by Fanger (1972), which has been largely accepted for design and field assessment, measure how various factors such as air humidity, clothing, and users’ activities affect the perceived comfort of a large group of users. However, it is difficult to incorporate this model into the design of an optimal control strategy that takes into account both cost and comfort. Alternatively, some recent work models comfort/discomfort by using functions of the parameters of interest. For instance, Livengood and Larson (2009) uses a discretized function that maps the room temperatures at discretized time points to numbers quantifying comfort; Liang and Shen (2011) model discomfort as two parts: delay in meeting demand

Thermostats for the Smart Grid / 63

and failure to meet demand. We notice that users are usually risk-averse over the deviation of room temperature from the preferred one; for example, a 2°F move away from the preferred temperature is barely noticeable, but a similar increase when the temperature is already 8°F causes a noticeable increase in discomfort. This paper attempts to use a comfort function that captures that characteristic.

While one of the objectives of the studies cited above and other related work is to use price to reduce demand in peak hours, little has been done to analyze what would happen in the off-peak periods. Black and Tyagi (2010) study the rebound effect under critical-peak-pricing (CPP), which is a variant of time-of-use (TOU) pricing. They notice that when CPP is applied, a significant portion of the load is shifted from peak-pricing hours to subsequent hours. With deep penetration of CPP in pilot studies, this effect is enlarged and results in new peaks right after the peak-pricing periods. Fortunately, the optimizing thermostat proposed in this paper mitigates the rebound effect.

Because HVAC appliances operate more efficiently when temperature differences between inside and outside are small, in some cases pre-cooling can reduce costs even without dynamic pricing (Braun, 2003). Braun (2003) notes that better control strategies are needed to optimize pre-cooling. Our optimizing thermostat provides optimal pre-cooling under flat-rate prices (see section 4.1).

This paper focuses on analyzing the economic impacts of having smarter thermostats for Smart Grid users. It contributes to the current literature in the following ways. We first model the control problems of three types of thermostat; namely, the traditional thermostat, the rigid thermostat, and the optimizing thermostat. In particular, we formulate both the thermodynamics of room temperatures and users' comfort. We further show by both analytical results and controlled simulations that the optimizing thermostat gives users maximum overall benefit. We also illustrate the social value of the optimizing thermostat and discuss drawbacks of the traditional and the rigid thermostats.

The remainder of this paper is organized as follows. In section 2, we formulate the control problems of the three types of thermostat and analyze their performance. In section 3, we conduct computational studies to compare these thermostats and obtain managerial insights. In section 4, we discuss an improved thermal model and the uncertainties in forecasts. In section 5, we conclude the paper and discuss the limitations and future research directions.

2. MODEL

2.1 Three Types of Thermostat

Users in most American homes program a traditional thermostat with a target room temperature (or temperature range), which may vary by the time of day. The traditional thermostat uses a trigger policy: HVAC is activated every time the room temperature deviates from the target temperature (or falls out of the target range). Unfortunately, the traditional thermostat gives the same response regardless of energy prices.

To enable automatic management of electricity usage, a new generation of somewhat intelligent—or “rigidly smart”—thermostats have been proposed, which we name rigid thermostats. A rigid thermostat runs an optimization routine in the background that, given dynamic prices, finds the lowest-cost path to keep room temperatures within a rigid range set by the user (such as the thermostat studied by Oldewurtel et al. (2010)). The range, for instance, may be what the users find tolerable. Unlike the traditional thermostat, the rigid thermostat shifts electricity demand from high-price peak hours to lower-price hours. However, the rigid thermostat’s price-responsiveness disappears precisely when the social value of lowering peak demand is highest. In addition, how users set their minimum and maximum room temperatures is subtle. For instance, to save cost, users may set a temperature range broad enough to lead to extreme discomfort, such as overcooling in summer.

To overcome the drawbacks of the rigid thermostat, we propose a novel optimizing thermostat. The optimizing thermostat trades off comfort for cost saving based on users’ preferences. Instead of indicating the acceptable range of room temperature, users set parameters to reflect their preferences on room temperatures and cost savings. The optimizing thermostat runs an optimization routine in the background to maximize users’ total utility.

Both the rigid and the optimizing thermostats pre-cool, but the latter is more flexible in how it treats high and low temperatures. For example, on a hot afternoon, the rigid thermostat will typically let the afternoon temperature rise to its rigid maximum. The optimizing thermostat might keep inside temperatures lower than that level if electricity prices are low, but permit the afternoon temperatures to rise higher if electricity prices are extremely high. Although the optimizing thermostat sometimes costs its users more than the rigid one does, it always gives users the highest total utility in cost and comfort.

All three types of thermostat first ask users to input the necessary parameters: a target temperature or range for the traditional thermostat, a temperature range for the rigid thermostat, and preference parameters for the optimizing thermostat. Presumably, the users of the traditional and rigid thermostats would like to set the inputs to maximize their total utility, but the control strategies of those thermostats do not help them achieve that goal.¹ We discretize time into periods and assume that, at the beginning of each period, each thermostat runs its control algorithm to decide whether or not to run HVAC and, if so, what the target room temperature should be.² The detailed model formulations for the control problems are described in the following subsection.

1. To numerically compare the performance of the thermostats, we tweak the inputs so the thermostats generate almost identical comfort and compare the costs.

2. In Appendix A, we study how the length of each period, the number of look-ahead periods, the length of the planning horizon, and the aggregation level of temperatures affect computational requirements (measured by CPU time) and the performance of the control algorithms.

2.2 Thermal Model and Users' Utility

First we define the following notation:

- Parameters:
 - n : index of periods, $n = 1, 2, \dots, N$.
 - N : the number of (look-ahead) periods.
 - t_n : time at the beginning of period n and the end of period $n-1$.
 - h : the length of each period.
 - p_n : price of electricity in period n .
 - $T_{in,s}$: indoor/room temperature at time s .
 - $T_{ini,n}$: initial indoor/room temperature in period n .
 - $T_{out,n}$: ambient temperature in period n . We assume the ambient temperature remains unchanged throughout any period.
 - $T_{pfr,n}$: user-preferred room temperature in period n . This is the room temperature at which users feel most comfortable.
- User's decision variables:
 - $I_{tra,n}$: the traditional thermostat users' target room temperature in period n .
 - $I_{min,n}$: the rigid thermostat users' minimum acceptable room temperature in period n .
 - $I_{max,n}$: the rigid thermostat users' maximum acceptable room temperature in period n .
- Thermostat's decision variables:
 - $X_{tar,n}$: the target room temperature for HVAC at the beginning of period n . ($X_{tar,n} = I_{tra,n}$ for the users of the traditional thermostat.)

There have been various models for how room temperature evolves. These models are typically functions of initial room temperatures, ambient temperatures and other factors such as construction materials, sources of internal gain (e.g., appliances and lighting), and the behavior of residents. Different factors and parameters can lead to quite different models; for instance, models for houses in California can be quite different from those for apartments in New York.

However, because the focus of this paper is to investigate the performance of the optimizing thermostat in general, we use a simple but general model (as in Reddy et al. (1991)) instead of those ad-hoc models. Our model can describe the temperature-evolving process, which is necessary in order to measure users' comfort.

We model a building's thermodynamics in the following way. The whole building is treated as one object with room temperature $T_{in,s}$ at time s . Let k_1 denote the building's projected heat capacity—the amount of electricity required by HVAC to change the building's temperature by 1°F; k_1 is smaller if HVAC is

66 / *The Energy Journal*

more efficient. At the beginning of period n (that is, time $s = t_n$) $T_{in,s} = T_{in,t_n} = T_{ini,n}$. The amount of electricity Q_n required to change the temperature of the building in period n is:

$$Q_n = k_1 |X_{tar,n} - T_{ini,n}| = k_1 \Delta T \quad (1)$$

Let k be the building material's conductance and A be the cross-sectional area of the conducting surface. By Newton's law of cooling, the rate of heat flux from the ambience to the internal space of the building at time s ($s \in [t_n, t_{n+1}]$), is described by the following Fourier differential equation:

$$\vec{q}_s = \frac{dQ}{ds} = kA(T_{out,n} - T_{in,s}) = k_2(T_{out,n} - T_{in,s}) \quad (2)$$

We call $k_2 = kA$ the projected conductance rate, the rate at which the building loses heat to or gains heat from the outside. A smaller k_2 indicates that the building is better insulated. Note that the heat transfer rate is proportional to the temperature difference. Therefore, even if we allow air exchange between the inside of the building and its surroundings, the same model still works after a small modification of k_2 .

When the HVAC is off, we assume that the only driving force for temperature change is conduction through the walls, windows, and ceilings. According to Equations (1) and (2), the rate of change of room temperature $T_{in,s}$ at time s ($s \in [t_n, t_{n+1}]$) without HVAC is:

$$k_1 \frac{dT_{in,s}}{ds} = k_2(T_{out,n} - T_{in,s}) \quad (3)$$

Let $t_{stop,n}$ be the time at which HVAC is stopped and $X_{tar,n}$ be the room temperature at the end of the previous HVAC run, then the function $f^T: \mathbb{R}_+^3 \rightarrow \mathbb{R}$ calculates the room temperature $T_{in,s}$ at time s ($s \in [t_{stop,n}, t_{n+1}]$):

$$\begin{aligned} T_{in,s} &= f^T(T_{out,n}, X_{tar,n}, s - t_{stop,n}) \\ &= T_{out,n} + (X_{tar,n} - T_{out,n})e^{-\frac{k_2}{k_1}(s - t_{stop,n})} \end{aligned} \quad (4)$$

Remark: In practice, the amount of energy required to cool the building is determined by outside temperatures, solar gain (and associated solar load), time-variant internal gain, the characteristics of the building (such as size, structure, building material, and furnishings), etc. In our model, we assume that k_1 captures all the factors that affect the difference in a building's internal energy at different equilibrium room temperatures. In addition, by treating the entire building as one object, we actually assume that the heat transfer processes among the objects inside the building happen instantly. Thus, the electricity required by HVAC is

Thermostats for the Smart Grid / 67

proportional to the change in room temperature. We further assume that k_2 is constant over time, while in reality, k_2 increases when windows or doors are open. Fortunately, relaxing these assumptions does not affect the main results of this paper (see section 4.1).

Let U_n be users' total utility in period n . We assume that it is separable into (1) the cost of heating and cooling, C_n , and (2) the comfort of having an inside temperature different from the users-preferred one, D_n , both of which reduce total utility (thus both take negative values):

$$U_n = C_n + D_n$$

Without loss of generality, we convert comfort D_n to dollar values. In particular, we assume that users have a total utility function that depends on comfort and on heating and cooling cost. The unit of the utility function is arbitrary up to a positive affine transformation. To save one more parameter, we rescale comfort to be in dollar values. This normalization permits us to describe the comfort for which a user is willing to pay \$100 as "\$100 worth of comfort". In particular, function $F^C: \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ calculates C_n , the cost of changing the room temperature by HVAC from $T_{ini,n}$ to $X_{tar,n}$:

$$C_n = F^C(T_{ini,n}, X_{tar,n}) = -p_n Q = -p_n k_1 |T_{ini,n} - X_{tar,n}| \quad (5)$$

For simplicity, we assume that changing room temperature by HVAC is instantaneous: $t_{stop,n} = t_n$. For two reasons, the error generated by this assumption is believed to be small. First, because the conduction process that gradually gains or loses heat to ambience after HVAC stops is much slower than the convection process through which HVAC changes room temperature, users' experience in the former process dominates. Second, the cost is not related to how quickly the energy is consumed, but to the total amount of energy used. Results are similar without these simplifying assumptions.

Because users usually become increasingly uncomfortable when the room temperature deviates further from $T_{pfr,n}$, we model their utility associated with comfort by a negative-valued concave function. In particular, we choose $f^D: \mathbb{R}_+^4 \rightarrow \mathbb{R}_-$ to be concave quadratic to calculate $d(s)$, the comfort gain per unit time due to the temperature deviation at time s in period n :

$$\begin{aligned} d(s) &= f^D(T_{out,n}, T_{pfr,n}, X_{tar,n}, s - t_n) \\ &= a(T_{in,s}^{-T_{pfr,n}})^2 + b |T_{in,s} - T_{pfr,n}| \\ &= a(f^T(T_{out,n}, X_{tar,n}, s - t_n) - T_{pfr,n})^2 \\ &\quad + b |f^T(T_{out,n}, X_{tar,n}, s - t_n) - T_{pfr,n}| \end{aligned} \quad (6)$$

where the parameters a and b are the coefficients of the second-order and first-order deviation terms. Specifically, b measures how much the users' comfort

decreases with small deviations from the preferred room temperature, while a accounts for the increasing decrement in comfort with each increment in the deviation. For example, imagine that two users prefer the same room temperature of 73°F. Let $a_1 = -1$, $b_1 = 0$ for the first user and $a_2 = 0$, $b_2 = -1$ for the second user. According to Equation (6), the first user finds 4 hours of 74°F neither better nor worse than 3 hours of 73°F plus 1 hour of 75°F. On the other hand, the second user has the same level of comfort with 4 hours of 74°F, and finds that neither better nor worse than 3 hours of 73°F plus 1 hour of 77°F. Hence, the second user will be less uncomfortable by having 3 hours of 73°F plus 1 hour of 75°F, which results from the fact that the first user is more sensitive to greater deviation in room temperature.

Because room temperatures keep changing between two consecutive HVAC runs, let function $F^D: \mathbb{R}_+^3 \rightarrow \mathbb{R}_-$ calculate users' utility associated with comfort, D_n , by integrating Equation (6):

$$D_n = F^D(T_{out,n}, T_{pfr,n}, X_{tar,n}) = \int_0^h f^D(T_{out,n}, T_{pfr,n}, X_{tar,n}, s) ds \quad (7)$$

$$= \int_0^h \{ a [f^T(T_{out,n}, X_{tar,n}, s) - T_{pfr,n}]^2 + b |f^T(T_{out,n}, X_{tar,n}, s) - T_{pfr,n}| \} ds$$

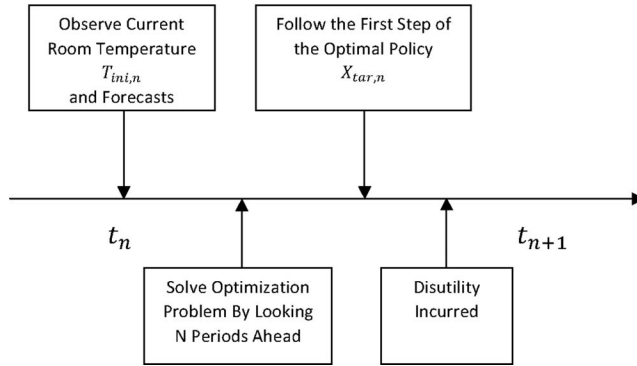
Total utility U_n is the sum of Equation (5) and Equation (7):

$$U_n(T_{out,n}, T_{ini,n}, T_{pfr,n}, X_{tar,n}) = -p_n k_1 |T_{ini,n} - X_{tar,n}|$$

$$+ \int_0^h a [f^T(T_{out,n}, X_{tar,n}, s) - T_{pfr,n}]^2 ds \quad (8)$$

$$+ \int_0^h b |f^T(T_{out,n}, X_{tar,n}, s) - T_{pfr,n}| ds$$

Unless otherwise noted, we assume in the remainder of this paper that the users of the traditional thermostat set $I_{tra,n} = T_{pfr,n}$ for all n . In addition, rather than using the trigger policy of the traditional thermostat, the other two thermostats work in the following way. As shown in Figure 1, at the beginning of period n , a thermostat reads the current room temperature and obtains forecasts on electricity prices and weather condition for the next N periods. Then, by solving an optimization problem outlined in the next subsection, the thermostat generates a sequence of target room temperatures and executes the first step. The room temperature is adjusted by the HVAC and the total utility U_n is incurred. The thermostat repeats this procedure in every period.

Figure 1: Sequence of Events

2.3 Rigid Thermostat

The rigid thermostat minimizes cost while keeping the temperature within a given range. We model the rigid thermostat's decisions as a dynamic programming problem. Denote the total cost and total utility from period n to the end of the planning horizon as $J_{\cdot,n}$ and $V_{\cdot,n}$, respectively; then the Bellman equation for the rigid thermostat is:

$$\mathbf{P1}: J_{rigid,n}(T_{ini,n}) = \max_{X_{tar,n}} F^C(T_{ini,n}, X_{tar,n}) + \alpha J_{rigid,n+1}(T_{ini,n+1}) \quad (9)$$

$$s.t. \quad T_{ini,n+1} = \left\lceil f^T(T_{out,n}, X_{tar,n}, h) + \frac{1}{2} \right\rceil \quad (10)$$

$$X_{tar,n}, T_{ini,n+1} \in [I_{min,n}, I_{max,n}] \quad (11)$$

where $n \in \{0, 1, 2, \dots, N\}$ and $J_{rigid,N+1}(T_{ini,N+1}) \equiv 0$. In Equation (9), the first term is the one-period cost for electricity. The second term denotes the future value, given that the initial room temperature is $T_{ini,n}$ and the decision in period n is $X_{tar,n}$. α is the discount factor, which is set to be 1 for all the simulations in this paper. In order to solve the problem efficiently, we confine the number of states by aggregating the temperatures into integer values. Constraint (10) rounds the end temperature of period n . This rounding will bring error into the optimization, but we can reduce the error by having finer discretization over temperature. Constraints (11) are the rigid bounds on the feasible variable.

Theorem 1 For any fixed initial room temperature and a sequence of ambient temperature forecasts, let $X: \mathbb{R}_+^{N+1} \rightarrow \mathbb{R}_+^{N+1}$ be the function that maps the sequence of electricity prices forecasts to the sequence of control signals obtained by solving **P1**. Then X is homogeneous of degree zero; that is,

$$X(p) = X(\lambda p), \quad \forall p \in \mathbb{R}_+^{N+1}, \quad \lambda > 0$$

where $p \in \mathbb{R}_+^{N+1}$ is the vector of electricity prices.

As will be shown later, the rigid thermostat often reduces electricity usage if the price rises during one particular hour (relative to other hours that day). This theorem suggests that the rigid thermostat does not reduce energy use when prices rise proportionately for an entire day (relative to other days). Conversely, even if the price profile is smooth, the demand in each period may vary a lot, as long as the prices in different periods are not equal.

Let the initial room temperature be $T_{ini,0}$. Denote the optimal policy, corresponding cost, and total utility as $\{X_{rigid,n}^*(T_{ini,n})\}_{n=0}^N$, $\{J_{rigid,n}^*(T_{ini,n})\}_{n=0}^N$, and $\{V_{rigid,n}^*(T_{ini,n})\}_{n=0}^N$, respectively. Similarly, denote by $\{X_{tra,n}^*(T_{ini,n})\}_{n=0}^N$, $\{J_{tra,n}^*(T_{ini,n})\}_{n=0}^N$, and $\{V_{tra,n}^*(T_{ini,n})\}_{n=0}^N$ the optimal control policy, corresponding cost and total utility of the traditional thermostat.

Our next theorem compares the costs of the rigid and traditional thermostats:

Theorem 2 (Weak Duality of Cost) *If $I_{tra,n} \in [I_{min,n}, I_{max,n}]$ for all $n \in \{0, 1, 2, \dots, N\}$, then: $J_{tra,n}^*(T_{ini,n}) \leq J_{rigid,n}^*(T_{ini,n})$ for all $n \in \{0, 1, 2, \dots, N\}$.*

That is, when the traditional thermostat is set within the bounds of the rigid thermostat, the cost of running HVAC is at least as high as it would be with the rigid thermostat. The intuition here is that the rigid thermostat reduces cost by allowing the room temperatures to be different from $I_{tra,n}$, whereas the traditional thermostat triggers HVAC whenever $T_{ini,n}$ deviates from the predetermined $I_{tra,n}$. It is worthwhile to point out that, with a flat-rate electricity price, the rigid thermostat reduces to the controllers that minimize energy usage by pre-cooling, the same as the one studied in Braun (2003).

2.4 Optimizing Thermostat

Like the rigid thermostat, the optimizing thermostat solves a control problem that evolves dynamically. However, the objective function of the optimizing thermostat's problem includes both D_n , the utility associated with comfort, and C_n , the cost. Moreover, there are no rigid bounds on $X_{tra,n}$ for the optimizing thermostat. This feature, along with our use of a concave function to represent comfort, differentiates our model from similar work (such as Livengood and Larson (2009)). Although having no bounds on the feasible region greatly increases the dimension of the action space and leads us to the widely known difficulties of dynamic programming—the “curses of dimensionality”—it allows us to find the solution that maximizes users' total utility. In particular, the optimizing thermostat solves the following Bellman equation:

Thermostats for the Smart Grid / 71

$$\begin{aligned} \mathbf{P2}: V_{opt,n}(T_{ini,n}) = & \max_{X_{tar,n}} F^C(T_{ini,n}, X_{tar,n}) + F^D(T_{out,n}, T_{pfr,n}, X_{tar,n}) \\ & + \alpha V_{opt,n+1}(T_{ini,n+1}) \end{aligned} \quad (12)$$

$$s.t. \quad T_{ini,n+1} = \left[f^T(T_{out,n}, X_{tar,n}, h) + \frac{1}{2} \right] \quad (13)$$

where, in (12), the second term accounts for the utility associated with comfort. To make the problem computationally tractable, we add lower and upper bounds on the feasible room temperatures, which admit a different range than that of a rigid thermostat. The modified problem ($\tilde{\mathbf{P2}}$) is stated below:

$$\begin{aligned} \tilde{\mathbf{P2}}: \tilde{V}_{opt}(T_{ini,n}) = & \max_{X_{tar,n}} F^C(T_{ini,n}, X_{tar,n}) \\ & + F^D(T_{out,n}, T_{pfr,n}, X_{tar,n}) + \alpha \tilde{V}_{opt,n+1}(T_{ini,n+1}) \\ s.t. \quad T_{ini,n+1} = & \left[f^T(T_{out,n}, X_{tar,n}, h) + \frac{1}{2} \right] \\ X_{tar,n}, T_{ini,n+1} \in & [B_n^l, B_n^u] \end{aligned} \quad (14)$$

The major differences between $[I_{min,n}, I_{max,n}]$ (in (11)) and $[B_n^l, B_n^u]$ (in (14)) are: For the rigid thermostat, users set $[I_{min,n}, I_{max,n}]$ to prevent room temperatures from falling outside the rigid range, while for the optimizing thermostat, $\tilde{\mathbf{P2}}$ has $[B_n^l, B_n^u]$ to make the optimization problem solvable. As before, let the optimal policy, corresponding cost, and total utility of $\mathbf{P2}$ be $\{X_{opt,n}^*(T_{ini,n})\}_{n=0}^N$, $\{J_{opt,n}^*(T_{ini,n})\}_{n=0}^N$, and $\{V_{opt,n}^*(T_{ini,n})\}_{n=0}^N$, respectively, while letting those of the modified problem $\tilde{\mathbf{P2}}$ be $\{\tilde{X}_{opt,n}^*(T_{ini,n})\}_{n=0}^N$, $\{\tilde{J}_{opt,n}^*(T_{ini,n})\}_{n=0}^N$, and $\{\tilde{V}_{opt,n}^*(T_{ini,n})\}_{n=0}^N$.

Intuitively, if the optimizing thermostat's bounds are chosen such that taking an action that violates the bounds costs more than maintaining the preferred temperature throughout the planning horizon, then the bounds will never be violated. The following proposition generates tighter bounds, as it gives a sufficient condition for $\tilde{\mathbf{P2}}$ to have the same solution as $\mathbf{P2}$. The tighter bounds allow us to significantly save computational time.

Proposition 1 Define ranges $[\tilde{B}_n^l, \tilde{B}_n^u]$ ($n \in \{0, 1, 2, \dots, N\}$) as following:

• for the last period N :

$$\begin{cases} \tilde{B}_n^u = \min_{x_N} \{x_N: D_N(T_{out,N}, T_{pfr,N}, x_N + 1) - D_N(T_{out,N}, T_{pfr,N}, x_N) > k_1 p_N, \\ \quad f^T(T_{out,N}, x_N, s) \geq T_{pfr,N} \quad \forall s \in [0, h]\} \\ \tilde{B}_n^l = \max_{x_N} \{x_N: D_N(T_{out,N}, T_{pfr,N}, x_N - 1) - D_N(T_{out,N}, T_{pfr,N}, x_N) > k_1 p_N, \\ \quad f^T(T_{out,N}, x_N, s) \leq T_{pfr,N} \quad \forall s \in [0, h]\} \end{cases}$$

- for the other periods n ($n \in \{0, 1, 2, \dots, N-1\}$):

$$\left\{ \begin{array}{l} \tilde{B}_n^u = \min_{x_n} \{x_n: D_n(T_{out,n}, T_{pfr,n}, x_n + 1) - D_n(T_{out,n}, T_{pfr,n}, x_n) > k_1 p_n, \\ f^T(T_{out,n}, x_n, s) \geq T_{pfr,n} \quad \forall s \in [0, h], \\ f^T(T_{out,n}, x_n, h) \geq B_{n+1}^u \} \\ \tilde{B}_n^l = \max_{x_n} \{x_n: D_n(T_{out,n}, T_{pfr,n}, x_n - 1) - D_n(T_{out,n}, T_{pfr,n}, x_n) > k_1 p_n, \\ f^T(T_{out,n}, x_n, s) \leq T_{pfr,n} \quad \forall s \in [0, h], \\ f^T(T_{out,n}, x_n, h) \leq B_{n+1}^l \} \end{array} \right.$$

If the feasible region of **P2** satisfies $[\tilde{B}_n^l, \tilde{B}_n^u] \subseteq [B_n^l, B_n^u]$, then the optimal solutions of **P2** are interior points of the feasible region of $\tilde{\mathbf{P2}}$.

The following theorem relates the total utilities of using the three types of thermostat:

Theorem 3 (Weak Duality) If $I_{min,n} \geq B_n^l$ and $I_{max,n} \leq B_n^u$ for all $n \in \{0, 1, 2, \dots, N\}$, then the total utilities of using the three types of thermostat satisfy:

$$V_{tra,n}^*(T_{ini,n}) \leq \tilde{V}_{opt,n}^*(T_{ini,n}) \leq V_{opt,n}^*(T_{ini,n})$$

and

$$V_{rigid,n}^*(T_{ini,n}) \leq \tilde{V}_{opt,n}^*(T_{ini,n}) \leq V_{opt,n}^*(T_{ini,n})$$

The second inequalities are binding if the condition in Proposition 1 is met.

Note that the conditions in Theorem 3 are sufficient conditions. The intuition behind Theorem 3 is that the optimizing thermostat penalizes deviation from the preferred temperature; in contrast, the rigid thermostat treats the preferred and the other temperatures in $[I_{min}, I_{max}]$ as equally comfortable and the temperature outside $[I_{min}, I_{max}]$ as infinitely uncomfortable. Furthermore, it is straightforward to extend the results of Theorem 3 to obtain the following useful corollary.

Corollary 1 Compared with the rigid thermostat, the optimizing thermostat:

- Creates more comfort for users if room temperatures are constrained within $[I_{min}, I_{max}]$ (or if the generated room temperatures are bounded by $[I_{min}, I_{max}]$), and
- Costs less when it generates lower or equal comfort.

The above results do not rely on the specific model for room temperatures that we assumed. More accurate models (for room temperatures and the

associated comfort calculations) are able to better describe room temperatures but do not change the relationship among the regions of the feasible solutions in the Bellman equations of the thermostats; thus, Theorem 2 and Theorem 3 still hold. Theorem 1 also holds because scaling the price does not affect the relationship among different solutions. In fact, as long as the model of interest possesses the Markovian property, we can model the control problems of the rigid and the optimizing thermostats by similar Bellman equations and the above theorems still hold.

Our models also work when there are forecasting uncertainties. In particular, the *certainty equivalence control* approach (Bertsekas, 1995) is applied. Although applying this approach is penalized when “bad” scenarios with low probability occur (Papavasiliou, 2011), the worst-case one-period total utility of setting $X_{tar,n}$ at X (X equals to $X_{rigid,n}^*$ or $X_{opt,n}^*$) is bounded below by the summation of $F^D(T_{out,n}, T_{pfr,n}, X)$, $F^C(T_{ini,n}, T_{pfr,n})$, and $F^C(f^T(T_{out,n}, X, h), f^T(T_{out,n}, T_{pfr,n}, h))$. Simulations that study the performance of this CEC approach are presented in section 4.2. An alternative approach that also preserves the weak duality is to relax the full-information assumption by setting prices and ambient temperatures as Markov processes. This approach slightly increases the computational requirements, as the deterministic value-to-go functions need to be replaced by their expectations.

3. SIMULATION AND MANAGERIAL INSIGHTS

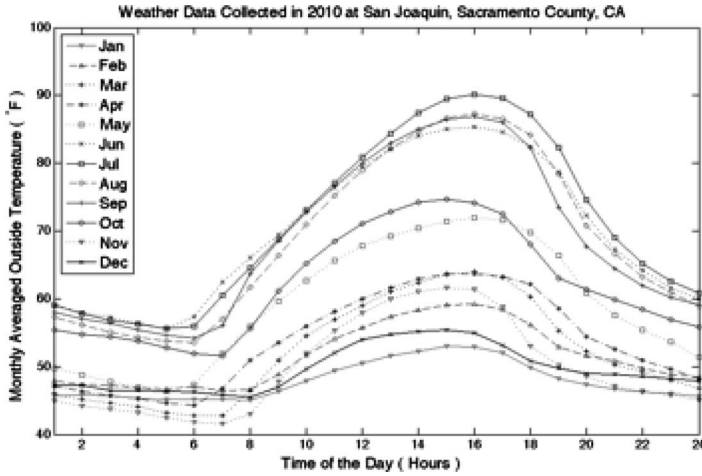
Theorem 3 implies that a user’s total utility when using the optimizing thermostat, $V_{opt,n}^*$, is always better than or equal to those of using the traditional and the rigid thermostats. In this section, we conduct numerical studies to learn the magnitude of the increase in comfort and/or reduction in cost of the optimizing thermostat relative to the other two. Then, we perform several controlled numerical experiments to examine how the different thermostats affect peak demand and how the different parameters influence the performance of these thermostats.

3.1 Summer Savings of the Rigid and Optimizing Thermostats

We first investigate how the optimizing and rigid thermostats benefit Smart Grid users in the summer. In order to get a realistic simulation, we choose the ambient temperature data collected at San Joaquin in Sacramento County, California in 2010.³ The hour-by-hour monthly average ambient temperature profiles are plotted in Figure 2.⁴ We simulate one entire summer (24 hourly data points per day on the 93 days from June 21 to September 21).

3. Source: <http://www.ipm.ucdavis.edu/WEATHER/wxretrieve.html>

4. All figures use only data collected at the beginning and the end of each hourly period. In later figures, data from distinct simulations are plotted with different markers and we connect the data points to make the results that come from different settings more distinguishable.

Figure 2: Monthly Averaged Temperatures in 2010 in San Joaquin, CA

Note: Data are from San Joaquin in Sacramento County, CA

Because there is currently no real-time-price data for retail markets available to us, we generate simulated dynamic prices by the following sine function:

$$p(t_i) = p_b - p_s \sin\left(\frac{2\pi i}{N}\right) \quad (15)$$

where p_b and p_s are the baseline and scale of the prices. We set $p_b = 0.2\$/\text{kWh}$ and $p_s = 0.11\$/\text{kWh}$, such that the electricity prices vary from $0.09\$/\text{kWh}$ in the midnight to $0.31\$/\text{kWh}$ in the afternoon. Our p_b and p_s approximate the average and half of the difference between the highest and lowest prices of PG&E's tier 1 TOU price structure (for residential users).⁵ Although this assumption on prices does not necessarily represent future pricing strategies for the SmartGrid, it allows us to simulate the savings for users under a real-time pricing scheme that is similar to PG&E's TOU structure.

We further assume that there is no uncertainty in the price and weather forecast.⁶ The advantage of using deterministic prices and ambient temperatures in simulations is that it will not bring error into the evaluation due to lack of sample paths. Motivated by the possibly more volatile prices in the SmartGrid, such as PG&E's pilot SmartRate (George et.al, 2009), we study the impact of having different p_b and p_s in the next subsection.

We assume that the users-preferred temperature T_{pfr} is 73°F . Parameters

5. http://www.pge.com/notes/rates/tariffs/electric.shtml#RESELEC_TOU

6. We run simulations to test the impact of forecasting uncertainties in section 4.2

Thermostats for the Smart Grid / 75

Table 1: Performances of the Thermostats in the Summer of 2010

	Traditional	Rigid	Optimizing
Settings (Inputs)	if $T_{ini} > 77^{\circ}\text{F}$, $I_{tra} = 76^{\circ}\text{F}$ if $T_{ini} < 76^{\circ}\text{F}$, $I_{tra} = 77^{\circ}\text{F}$ otherwise $I_{tra} = T_{ini}$	$I_{min} = 68^{\circ}\text{F}$ $I_{max} = 77^{\circ}\text{F}$	$a = -0.0075 \text{ } \$/(\text{F}^2 \cdot \text{hr})$ $b = -0.005 \text{ } \$/(\text{F} \cdot \text{hr})$
Average Daily Cost(\$)	<i>9.68</i>	<i>4.93</i>	<i>4.02</i>
Average Daily Comfort(\$)	-2.29	-2.22	-2.24
Mean Abs. Temp. Gap ($^{\circ}\text{F}$)	3.04	2.8	2.54
Root Mean Squared Deviation: ($^{\circ}\text{F}$)	3.19	3.12	3.15
% of Days at/below 68°F	0	0.98	0.23
% of Days at/above 77°F	0	0.83	0.83
Hours per day on 68°F	0	2.72	0.33
Hours per day on 77°F	0	1.4	1.17
Hours per day below 68°F	0	0	0
Hours per day above 77°F	0	0	3.1

Notes: Temperature data are from San Joaquin, CA. G_e and R_e are the average and root mean squared temperature deviation from T_{pfr} . We adjust the inputs of the traditional and the rigid thermostats such that all the thermostats generate almost identical total comfort (in bold), which allows us to compare their costs (in italics).

in Equation (6) are $a = -7.5 \times 10^{-3} \text{ } \$/(\text{F}^2 \cdot \text{hr})$ and $b = -5 \times 10^{-3} \text{ } \$/(\text{F} \cdot \text{hr})$, both of which are fixed throughout the year. These parameters indicate that users would be willing to pay an extra \$0.0125 per hour to lower the room temperature from 74°F to 73°F or an extra \$0.0425 per hour for a decrease from 76°F to 75°F .

We further assume that $k_1 = 1(\text{kWh}/^{\circ}\text{F})$, meaning that it takes 1 kWh of electricity to change the room temperature by 1°F . We assume $k_2 = 0.2(\text{kWh}/(^{\circ}\text{F} \cdot \text{hr}))$, which implies that the room temperature will rise from 73°F to 75°F in 1 hour when the outside temperature is 83°F . This adjustment speed is close to that assumed by others, such as Braun (2003), although in our model the room is slightly less insulated. Note that less insulation reduces the benefit of pre-cooling and, thus, favors the traditional thermostat.

The settings on the traditional and the rigid thermostats are adjusted to give almost the same average daily comfort over the summer as the optimizing thermostat does. Specifically, the input of the traditional thermostat is a fixed range $[76^{\circ}\text{F}, 77^{\circ}\text{F}]$, a bit above our assumed user-preferred temperature of 73°F . The rigid thermostat is set to keep the indoor temperature between 68°F and 77°F . This is by no means the only way to let all three thermostats yield almost the same average daily comfort, but it is the best we have found.

The main results are summarized in Table 1, where G_e and R_e are the mean absolute temperature gaps and the root mean squared deviations from

$T_{pfr} = 73^\circ\text{F}$, respectively. In this simulation, the rigid thermostat costs \$458.5 per summer, (49%) less than the traditional thermostat to achieve the same average comfort. Consistent with Theorem 3, the optimizing thermostat does even better, saving \$526.4 per summer (58%) relative to the traditional thermostat and \$84.63 (18%) relative to the rigid thermostat. To put these savings in perspective, an Internet-enabled thermostat costs less than \$100 at Home Depot. With the simulation results, we believe that, even if we ignore the costs of the other two thermostats, the setup cost of the optimizing thermostat can be more than compensated for by the cost saving in the first summer of use.

The advantage of the optimizing thermostat over the rigid thermostat comes in part from keeping the indoor temperature on average closer to the users-preferred temperature; the mean absolute deviation G_e is 2.54°F for the optimizing versus 2.80°F for the rigid thermostat. The rigid thermostat treats all temperatures between its lowest (I_{min}) and highest (I_{max}) permitted temperatures as equally valuable. Thus, if the indoor temperature is anywhere within that range, the rigid thermostat makes no effort to keep the temperature any nearer to the users-preferred temperature. Therefore, on any fairly hot day, the thermostat pre-cools all the way down to the lowest rigid constraint I_{min} , and also spends the warm part of the afternoon at the highest rigid constraint I_{max} . In our simulations, the rigid thermostat pushes the temperature down to I_{min} on 91 of the 93 days over the summer, while the optimizing thermostat only pushes temperatures that low on 21 days. Similarly, the rigid thermostat keeps the inside temperature within a degree of I_{min} for 2.72 hours per day over the summer, while the optimizing thermostat only does so for 0.33 hours per day.

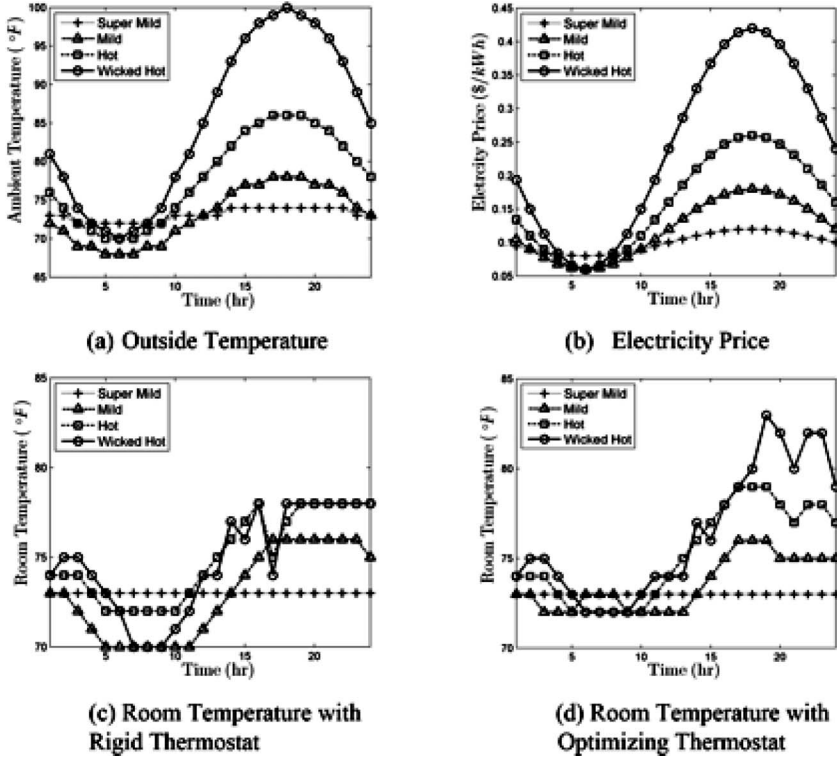
At the same time, the rigid thermostat never permits temperatures above I_{max} , while the optimizing thermostat balances cost and comfort with few extreme room temperatures. For instance, in the above simulation, the optimizing thermostat rises over I_{max} for 3.1 hours per day during the simulated summer.

3.2 Four Basic Scenarios

To understand the behavior of each thermostat, we run simulations on four scenarios: a “super mild” day, a “mild” day, a “hot” day, and a “wicked hot” day. To be consistent with the curves in Figure 2, we use the following sine function to generate ambient temperature profiles:

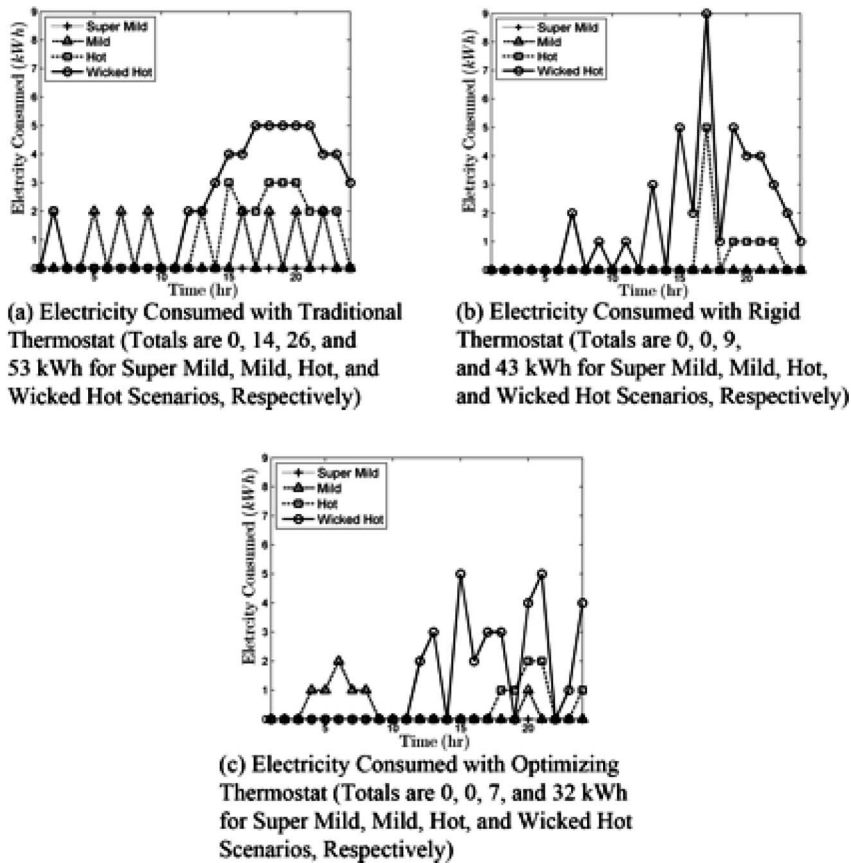
$$T(t_i) = T_b - T_s \sin\left(\frac{2\pi i}{N}\right) \quad (16)$$

where T_b and T_s are the baseline and scale of ambient temperatures. In addition, we generate prices from Equation (15), which implies that the prices and ambient temperatures are perfectly positively correlated, as shown in Figures 3(a) and 3(b). For the rigid and the optimizing thermostats, the look-ahead period is 24 hours. The users-preferred temperature, T_{pfr} , is assumed to be 73°F . In addition, $I_{tra} = T_{pfr}$ and $[I_{min}, I_{max}] = [68^\circ\text{F}, 78^\circ\text{F}]$.

Figure 3: Four Basic Scenarios

Figures 3(c) and 3(d) show the room temperature curves when the rigid and the optimizing thermostat are used, respectively. Because the traditional thermostat runs HVAC whenever the room temperature deviates from T_{pfr} , it produces similar room temperatures in all four scenarios and consumes much more electricity on hotter days, as shown in Figure 4(a). In contrast, as shown in Figure 4(b), the rigid thermostat only runs HVAC in “hot” and “wicked hot” conditions. In addition, it runs the HVAC in periods with low ambient temperature to pre-cool the room; that is, to reduce the room temperature below T_{pfr} . Because the electricity prices in those periods are lower, the rigid thermostat reduces cost by shifting the peak load to low-price periods to save cost (as long as there is need for HVAC during the planning horizon). Figure 4(c) shows that the optimizing thermostat also pre-cools, but not as extensively as the rigid thermostat.

Remark As shown in Figure 3(c), the room temperatures in the “super mild” case remain constant while the ambient temperatures change slightly over time; this is because the continuous temperatures are rounded to integer values.

Figure 4: Electricity Consumed in Four Basic Scenarios

3.3 Benchmark Simulations

To study the effect of parameters such as prices, ambient temperatures, building insulation, and users' sensitivities to room temperatures, we conduct nine controlled numerical experiments, listed in Table 2. For all nine cases, the ambient temperature profiles are generated from (16) with $T_b = 83^\circ\text{F}$, $T_s = 5^\circ\text{F}$, and the preferred room temperature $T_{pr} = 73^\circ\text{F}$. Both the planning horizon and look-ahead periods are set to 24 hours.

Different price profiles are tested to study the price sensitivity of the thermostats. We restrict the price range to be no greater than that of PG&E's SmartRate (George et.al, 2009), which ranges from 0.085\$/kWh to 0.85\$/kWh in summer. Specifically, we group the cases into a low-price group (cases 1 through 5) and a high-price group in which the prices are both higher and more volatile (cases 6 through 9).

Thermostats for the Smart Grid / 79

Table 2: Parameter Values for Sensitivity Tests

	Average p_b (\$/kWh)	Scale p_s (\$/kWh)	Conductance k_2 (kWh/(°F·hr))	Comfort Coefficient a (10 ⁻² \$/(°F ² ·hr))	Rigid's LB I_{min} (°F)
Case 1	0.1	0.03	0.2	-0.75	68
Case 2	0.1	0.08	0.2	-0.75	68
Case 3	0.1	0.08	0.2	-1.5	68
Case 4	0.1	0.08	0.2	-3	68
Case 5	0.1	0.08	0.1	-0.75	68
Case 6	0.3	0.24	0.2	-0.75	68
Case 7	0.3	0.28	0.2	-0.75	68
Case 8	0.3	0.28	0.2	-0.75	56
Case 9	0.3	0.28	0.2	-1.5	68

Notes: Highlights are key changes in parameters. The other parameters, T_b , T_s , k_1 , b and I_{max} , are fixed across all scenarios.

For the parameters reflecting building properties, $k_1 = 1$ (kWh/°F) is fixed and k_2 is varied to investigate the effect of better insulation (case 5). Similarly, for the parameters associated with comfort, $b = -5 \times 10^{-3}$ \$/(°F·hr) is fixed and a is adjusted to represent different users' sensitivities to the squared deviation from their preferred room temperature. Three values of a are tested. They are -7.5×10^{-3} \$/(°F²·hr), -1.5×10^{-2} \$/(°F²·hr) and -3×10^{-2} \$/(°F²·hr) for low sensitivity, medium sensitivity, and high sensitivity, respectively. In addition, I_{max} of the rigid thermostat is fixed at 78°F and I_{min} is lowered to study the impact of allowing extensive pre-cooling (case 8).

The daily cost (C_{day}), daily comfort in dollars (D_{day}), and daily total utility in dollars (U_{day}) are summarized in Table 3. Table 4 shows the mean absolute temperature gap and the root mean squared deviation from the users-preferred temperature by the three thermostats in all nine cases.

Price Elasticity and Social Surplus

In many regions, electricity demand peaks on hot summer afternoons, largely due to air conditioning. Borenstein (2005) shows that when supply is near capacity at these times, prices can spike to extremely high levels if demand is very inelastic. His simulation shows that the price spikes are much lower if a modest share of users have a modest price elasticity. Avoiding the price spikes, in turn, reduces generators' investments in high-cost and socially inefficient "peaker" plants used only a few hours a year. Finally, Borenstein (2005) shows that many of the gains from higher price elasticity are reaped by users (and there is higher average user's surplus gain for the price-elastic users), as they no longer

Table 3: Sensitivity Tests Results

	Traditional			Rigid			Optimizing		
	$C_{day}(\$)$	$D_{day}(\$)$	$U_{day}(\$)$	$C_{day}(\$)$	$D_{day}(\$)$	$U_{day}(\$)$	$C_{day}(\$)$	$D_{day}(\$)$	$U_{day}(\$)$
Case 1	5.39	-0.29	-5.69	2.62	-2.62	-5.25	3.5	-0.81	-4.3
Case 2	6.04	-0.29	-6.34	3.08	-2.19	-5.27	3.68	-0.88	-4.57
Case 3	6.04	-0.48	-6.53	3.08	-4.04	-7.11	3.86	-1.33	-5.19
Case 4	6.04	-0.87	-6.91	3.08	-7.73	-10.8	4.98	-0.99	-5.97
Case 5	3.68	-0.4	-4.08	0.78	-2.33	-3.11	1.81	-0.27	-2.08
Case 6	18.13	-0.29	-18.42	9.23	-2.19	-11.42	3.52	-5.7	-9.22
Case 7	18.65	-0.29	-18.95	9.27	-2.42	-11.7	3.47	-5.57	-9.04
Case 8	18.65	-0.29	-18.95	8.44	-8.65	-17.08	3.47	-5.57	-9.04
Case 9	18.65	-0.48	-19.14	9.27	-4.47	-13.74	8.87	-3.65	-12.52

Notes: The cost, normalized value of discomfort, and total utility each day (C_{day} , D_{day} and U_{day}) are all expressed in dollars.

Table 4: Thermostat Performance in the Scenarios

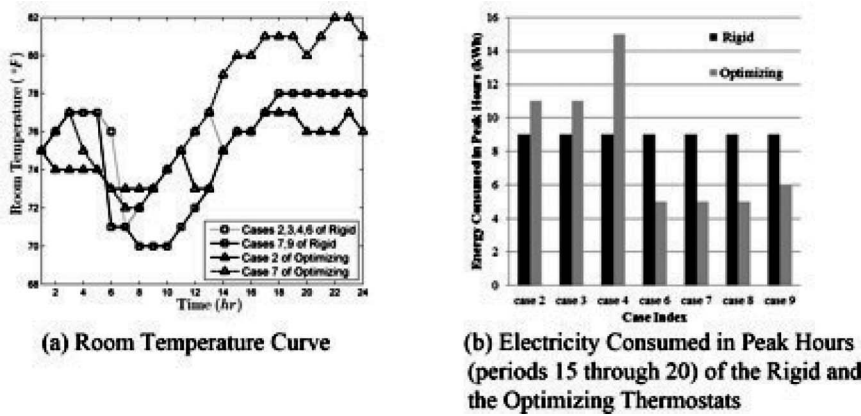
	Traditional		Rigid		Optimizing	
	$G_e(^{\circ}\text{F})$	$R_e(^{\circ}\text{F})$	$G_e(^{\circ}\text{F})$	$R_e(^{\circ}\text{F})$	$G_e(^{\circ}\text{F})$	$R_e(^{\circ}\text{F})$
Case 1	0.58	3.74	3.75	19.18	1.88	11.27
Case 2	0.58	3.74	3.33	17.83	1.88	11.62
Case 3	0.58	3.74	3.33	17.83	1.75	10.86
Case 4	0.58	3.74	3.33	17.83	1.25	7.48
Case 5	0.38	3.00	2.92	16.31	0.83	4.90
Case 6	0.58	3.74	3.33	17.83	4.63	27.26
Case 7	0.58	3.74	3.33	17.72	4.54	27.04
Case 8	0.58	3.74	5.25	30.43	4.54	27.04
Case 9	0.58	3.74	3.33	17.72	2.58	16.43

Notes: Denote the Mean Absolute Temperature Gap and the Root Mean Squared Deviation from T_{pfr} , as G_e and R_e , respectively.

suffer from very high price spikes. At the same time, Joskow (2011) shows that, in a situation with low demand elasticity, renewable sources of power may not be as socially valuable as they claim to be if their power does not help smooth price spikes and avoid the need for costly peaker plants. Conversely, increasing demand elasticity can increase the social value of renewable sources.

These results from Borenstein (2005) and Joskow (2011) have implications for the social values of the three thermostats. The traditional thermostat

Figure 5: Temperature and Electricity Consumed in Peak Hours

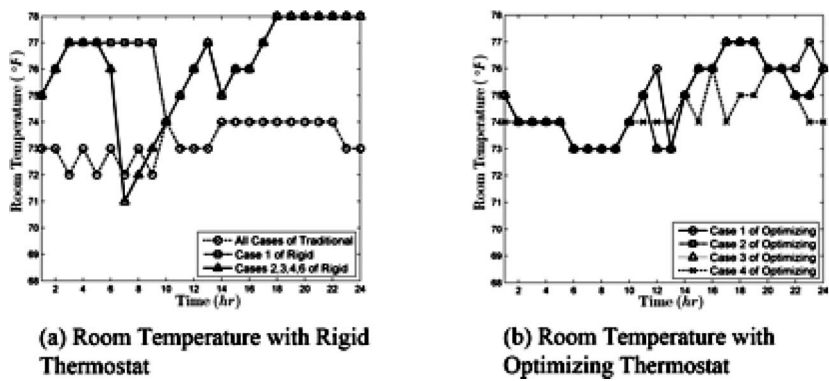


has zero price elasticity, making the price spikes emphasized in Borenstein (2005) larger and more likely. In contrast, the optimizing thermostat saves cost in peak hours by reducing comfort when energy prices are high and by shifting demand to lower-cost times of the day. The rigid thermostat usually has higher price elasticity than the optimizing thermostat, as it only aims to reduce cost and gives no attention to comfort as long as the inside temperature is within the rigid bounds. Unfortunately, on hot days which are typically the days with the highest social value of reducing peak demand, the price elasticity of the rigid thermostat falls to almost zero. Thus, the optimizing thermostat provides more social benefits than the rigid thermostat on precisely the hot days when price spikes are more likely to occur.

The price elasticity of the rigid thermostat breaks down for two reasons. First, consistent with Theorem 1, the optimal policy of the rigid thermostat is homogeneous of degree zero in prices. For instance, case 6 is identical to case 2 except that the prices in the former case are 3 times higher. The rigid thermostat consumes exactly the same electricity in each case, but has 3 times the cost (C_{day}) in case 6 as in case 2. In contrast, the optimizing thermostat reduces electricity consumption in case 6 relative to case 2.

Second, on hot days the rigid thermostat hits its upper bound I_{max} even after pre-cooling to its lower bound I_{min} . Thus, no matter how high electricity prices rise the rigid thermostat will not lower its peak demand. Moreover, its price elasticity can drop to near zero even if the weather is not severe. For example, in section 3.3, the simulated ambient temperatures are moderate, ranging from 78°F to 88°F. However, as shown in Figure 5(a), in cases 6,7,and 9 (of the high-price group), the rigid thermostat generates room temperature profiles similar to those in cases 2, 3, and 4 (of the low-price group). In contrast, the optimizing thermostat generates higher room temperatures in peak hours with higher peak

Figure 6: Room Temperatures under Different Price Functions



Notes: Prices are less volatile in case 1. Users are more sensitive to comfort in cases 3 and 4.

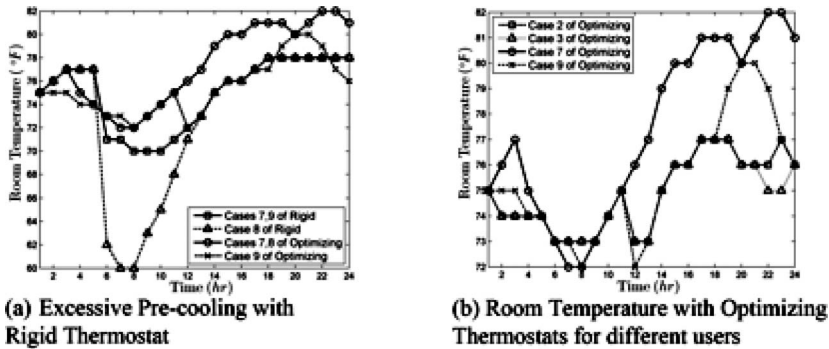
prices. In addition, Figure 5(b) shows that in the peak hours (periods 15 through 20), the rigid thermostat consumes the same amount of electricity across the simulated cases, whereas the optimizing thermostat consumes less in case 7 than in case 2.

More Concerns with the Rigid Thermostat

We further analyze the controlled experiments listed in Table 2 and obtain the following important observations. First, the rigid thermostat makes room temperatures reach I_{max} in the peak hours of summer. Moreover, it may make its users uncomfortably cold if I_{min} is too low or if the system does not have a minimum for hot days.

As shown in Figure 6(a), the control signals and the resulting room temperatures of the traditional thermostat, which applies a trigger policy, are only functions of weather; thus, the room temperatures remain the same in all the nine cases. If users of a traditional thermostat want to trade off comfort for cost saving, they have to monitor prices, perform complex calculations, and manually adjust the target room temperature. It is implausible that many users will undertake this real-time optimization. The other two types of thermostat admit price elasticity. In particular, the rigid thermostat always saves cost by making room temperatures reach I_{max} during summer afternoons when prices are high. Thus, the price elasticity of the rigid thermostat depends on the selection of the interval $[I_{min}, I_{max}]$; larger intervals lead to higher price elasticity. For example, in cases 2, 3, and 4 of the low-price group (which differ only in the users' sensitivity to room temperatures), the rigid thermostat makes the room temperatures more volatile than the optimizing thermostat does, as shown in Figure 6.

The dependence of the rigid thermostat's price elasticity on $[I_{min}, I_{max}]$ sometimes leads to extreme room temperatures. Figure 7(a) shows that, in case

Figure 7: Impacts of Other Parameters on Room Temperatures

7 (with high price and low sensitivity), the rigid thermostat reaches I_{max} in peak hours, during which HVAC needs to be run continuously. When I_{min} is dropped, as in case 8, the rigid thermostat pre-cools more aggressively to achieve more savings. However, this extensive pre-cooling will make users extremely uncomfortable in off-peak hours.

In contrast, the optimizing thermostat tries to avoid excessive pre-cooling or pre-heating, adjusting its usage according to users' preferences (a and b). For example, the optimizing thermostat keeps the room temperature closer to T_{pfr} when the user puts more value on comfort relative to cost (compare the low-price case 2 with the similar case 3, which has more sensitivity to temperatures, or compare the high-price case 7 with case 9), as shown in Figure 7(b). In cases in which the optimizing thermostat leads to very low or very high room temperatures, the discomfort will always be more than repaid by the lower cost.

As discussed in Appendix B, under some parameter values and settings the rigid thermostat sometimes gives users lower total utility than the traditional one. In addition, although better insulation is in general beneficial, without careful adjustment of the temperature range, better insulation can reduce comfort for the users of the rigid thermostat.

3.4 Another Impact on the Grid Side

Critical-peak-pricing (CPP) is a variant of dynamic pricing with a specific time period that has a discretely higher price. An undesirable consequence of CPP is that when the penetration of demand-responsive systems is deep, there may be another peak demand right after the CPP ends (Black and Tyagi, 2010). Intuitively, users delay their demand until the peak pricing ends, leading to the *rebound effect* as all the accumulated demand hits the grid.

One way to avoid this rebound effect would be to have the SmartGrid provide dynamic prices instead of discrete jumps in price. However, even with dynamic pricing, the rigid thermostat tends to run HVAC appliances at the max-

Table 5: Maximum energy use versus usage during peak hours

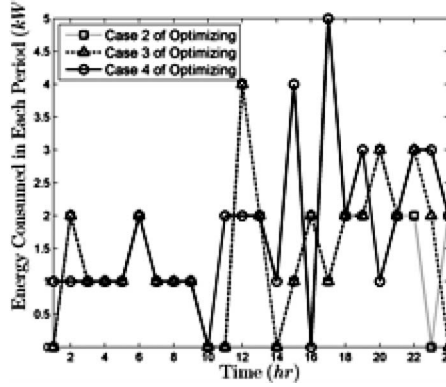
	Traditional		Rigid		Opt	
	E_{max}	E_{peak}	E_{max}	E_{peak}	E_{max}	E_{peak}
Case 1	3	3	4	1.5	5	1.83
Case 2	3	3	6	1.5	4	1.83
Case 3	3	3	6	1.5	4	1.83
Case 4	3	3	6	1.5	5	2.5
Case 5	2	2	5	0	2	0.67
Case 6	3	3	6	1.5	4	0.83
Case 7	3	3	7	1.5	3	0.83
Case 8	3	3	18	1.5	3	0.83
Case 9	3	3	7	1.5	5	1

Notes: Comparison of E_{max} (the maximum electricity demand in kWh per period) and E_{peak} (the average electricity consumed in kWh per period in periods 15 through 20, the peak periods, with the traditional thermostat)

imum level during periods when the prices reach minimum; this maximizes cost saving, but may lead to an undesirable secondary peak load. To illustrate this point, we summarize in Table 5 the maximum electricity consumed per period (E_{max}) and the average electricity consumed per period in the peak hours with the traditional thermostat (E_{peak}).

As shown in Table 5, although the E_{peak} of the rigid thermostat is lower than that of the traditional thermostat, the E_{max} of the rigid thermostat is usually much greater than that of the traditional thermostat. This implies that, with high penetration of the rigid thermostat, if a substantial share of users uses the settings $[I_{min}, I_{max}]$, a much more significant peak load may be created. A price schedule that ignores this peak demand prior to (or after) peak prices will not balance supply and demand (for example, the morning demand spike from the rigid thermostat's pre-cooling may overwhelm the grid). To avoid the morning spike, price schedules have to be smoothed within the day. Unfortunately, this smoothing reduces the energy saving and reductions in peak load that the SmartGrid hopes to achieve.

In contrast, simulation results suggest that the optimizing thermostat keeps the maximum electricity consumed per period much lower than the rigid thermostat does, although the reduction in peak load is not as significant as that of the rigid thermostat in some cases. Thus, the rebound effect is mitigated. Moreover, with the help of the optimizing thermostat, we can take advantage of the heterogeneity among users. For instance, as shown in Figure 8, although in cases 2, 3, and 4 all settings except a and b (parameters associated with comfort) remain the same, users with different temperature sensitivity have their peak demand

Figure 8: Electricity Consumed with Optimizing Thermostats

Notes: Total daily consumption is 33, 34, and 42kWh for Cases 2, 3, and 4, respectively.

shifted to different periods. Hence, the optimizing thermostat helps further smooth the aggregate demand.

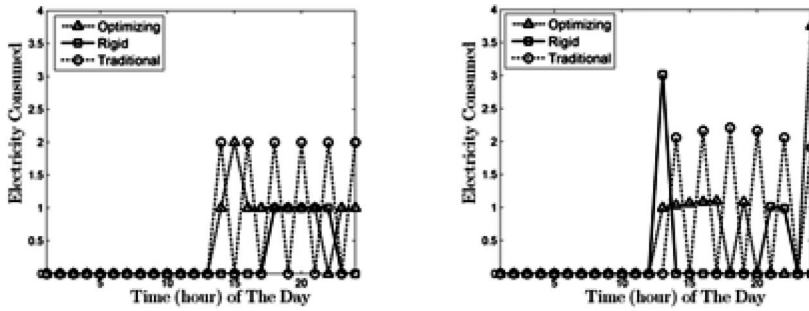
4. EXTENSION OF THE MODELS

4.1 Incorporating a More Accurate Thermal Model

As noted above, replacing the simple model for room temperatures with more accurate ones does not change the results of Theorems 1, 2, and 3. In this subsection, we test a more realistic thermal model in which the *coefficient of performance* (COP) of HVAC decreases in the temperature difference between inside and outside. This property of HVAC appliances actually favors the rigid and the optimizing thermostats because pre-cooling becomes more efficient. Therefore, as pointed out at the end of section 2.3, the rigid and optimizing thermostats are expected to pre-cool even if electricity prices are constant.

The new thermal model with a variable COP is as follows. We observe from the study of Payne and Domanski (2002) that a linear regression fits well with their calibration data for an air-conditioner. Thus, we replace Equation (5) with the following one:

$$C_n = \begin{cases} -p_n \times \frac{COP_o k_1 |X_{tar,n} - T_{ini,n}|}{COP_{ub} + \frac{COP_{lb} - COP_{ub}}{COP_{span}} \left(T_{out,n} - \frac{X_{tar,n} + T_{ini,n}}{2} \right)} & \text{if cooling} \\ -p_n \times \frac{COP_o k_1 |X_{tar,n} - T_{ini,n}|}{COP_{ub} + \frac{COP_{lb} - COP_{ub}}{COP_{span}} \left(T_{out,n} - \frac{X_{tar,n} + T_{ini,n}}{2} \right) + 1} & \text{if heating} \end{cases}$$

Figure 9: Energy Consumption Incorporating COP

where COP_{lb} , COP_{ub} , and COP_{span} are the HVAC's lowest, highest COPs, and the its span of working temperatures, respectively. We also assume, for consistency, that COP_o is the COP that has been used in previous simulations.

Simulation parameters are chosen to test performance under flat-rate electricity prices. We let the price be 0.3\$/kWh and simulate ambient temperatures according to Equation 16 with $T_b = 83^\circ\text{F}$ and $T_s = 15^\circ\text{F}$. We increase the incentive for using pre-cooling to save energy usage by decreasing k_2 to 0.05. The remaining parameters are the same as in case 7 in section 3.3.

The simulation results are summarized in Figure 9. With variable COP, both the rigid and the optimizing thermostats pre-cool to save energy, although the prices are constant. Furthermore, with variable COP, the daily cost with the rigid thermostat is the lowest (\$3.76, \$2.09, and 3.03 for the traditional, rigid, and optimizing thermostat, respectively). Consistent with Theorems 2 and 3, the optimizing thermostat gives the highest total utility ($-\$3.84$, $-\$4.52$, and $-\$3.42$ for the traditional, rigid, and optimizing thermostat, respectively). Finally, it can be shown that Proposition 1 still holds with some modifications, as explained in the proof of Proposition 1.

4.2 Incorporating Forecasting Uncertainties

The rigid and the optimizing thermostats schedule energy usage based on the forecasts of future weather conditions and electricity prices, both of which are difficult to predict. In this subsection, we test the performance of all three types of thermostat in the presence of forecasting uncertainties.

Forecasting uncertainties complicate the optimization for both rigid and optimizing thermostats by (1) creating the need to estimate the distribution of the uncertainties, and (2) increasing the computational cost of the problem. Furthermore, more accurate forecasts usually lead to the well-known "curses of dimensionality" of dynamic programming, because they rely heavily on both the current conditions and past history.

We consider the uncertainties in forecasting both ambient temperatures and electricity prices. In particular, we assume that the forecasts are able to deliver

Table 6: Simulation Results with Uncertain Forecasts

Case	Traditional			Rigid			Optimizing		
σ_{ε_p} and $\sigma_{\varepsilon_{T_{out}}}$	C_{day}	D_{day}	U_{day}	C_{day}	D_{day}	U_{day}	C_{day}	D_{day}	U_{day}
$\sigma_{\varepsilon_p} = \sigma_{\varepsilon_{T_{out}}} = 0$	18.65	-0.29	-18.95	9.27	-2.42	-11.70	3.47	-5.57	-9.04
$\sigma_{\varepsilon_p} = \sigma_{\varepsilon_{T_{out}}} = 1$	17.90	-0.10	-18.01	10.19	-3.10	-13.29	4.46	-5.78	-10.24
$\sigma_{\varepsilon_p} = \sigma_{\varepsilon_{T_{out}}} = 2$	16.63	-0.12	-16.75	9.62	-3.07	-12.70	5.18	-5.77	-10.95
$\sigma_{\varepsilon_p} = \sigma_{\varepsilon_{T_{out}}} = 3$	17.87	-0.20	-18.06	10.51	-3.06	-13.56	7.11	-5.73	-12.84
$\sigma_{\varepsilon_p} = \sigma_{\varepsilon_{T_{out}}} = 5$	17.55	-0.27	-17.82	8.56	-3.03	-11.59	5.90	-4.36	-10.27

Notes: σ_{ε_p} and $\sigma_{\varepsilon_{T_{out}}}$ are the standard deviation of the noise of the forecasts of prices (\$/kWh) and ambient temperatures (°F). The cost, normalized value of discomfort, and total utility each day (C_{day} , D_{day} and U_{day}) are all expressed in dollars.

the accurate expected values ($\hat{T}_{out,n}$ and $\hat{p}_n \forall n = 1, 2, \dots, N$). The realized values ($T_{out,n}$ and p_n) are the summation of expected values and some random noise ($\varepsilon_{T_{out,n}}$ and ε_{p_n}). That is:

$$T_{out,n} = \hat{T}_{out,n} + \varepsilon_{T_{out,n}}, \text{ and } p_n = \hat{p}_n + \varepsilon_{p_n}$$

We run simulations with the following parameters. To study the effects of the uncertainties, we assume $\varepsilon_{T_{out,n}}$ and ε_{p_n} are normal random variables with mean zero and that their variances are varied. The rest of the parameters are the same as in case 7 in section 3.3.

In the presence of forecasting uncertainties, our algorithm works following the certainty equivalence control approach. In particular, the future prices and ambient temperatures are replaced with their expectations when running the optimization routine and the system evolves stochastically after executing the optimal decisions. The simulation results are presented in Table 6, which shows that, even with relatively poor forecasts (when the standard deviations of the noise are comparable to half of the span of the forecasts' values), the optimizing thermostat still dominates the other two thermostats in terms of total utility.

5. CONCLUSION, LIMITATIONS AND EXTENSIONS

This paper analyzes strategies for managing U.S. households' major source of electricity demand: HVAC. Because this demand is flexible compared with other demand and often has synchronized peaks, the potential of shaving peak demand through load shifting is substantial. However, users will require tools to facilitate load shifting as the SmartGrid and its dynamic pricing become more common.

We model three types of thermostat—*traditional*, *rigid*, and *optimizing*—and the corresponding control strategies. We compare the performance of these thermostats both theoretically and via numerical studies. Both the traditional and the rigid thermostats have serious drawbacks and both deliver lower total utility than the optimizing thermostat. The optimizing thermostat creates more value for its users because it balances cost saving and comfort. In addition, the optimizing thermostat creates more value to society because it lowers peak electricity usage on the hottest days.

Our model captures the main elements of the problem, but makes a number of simplifying assumptions. It is straightforward to extend the current model to incorporate more complex circumstances, such as users' adaption to sustained high or low outdoor temperature (Angilletta (2009), see Appendix C for a preliminary analysis), forecasts (with uncertainty) of the number of people who will soon be home (perhaps using input from GPS-enabled phones, online calendars, motion sensors in burglar alarms, etc.), and multiple residents with different preferences. Parameters such as the users-preferred temperature and willingness to pay for comfort (T_{pfr} , a and b) can be learned by statistical learning modules yet to be developed. It is also straightforward (1) to replace the model for room temperatures with one that accounts better for lags, humidity, solar gain, wind, and so forth, (2) to utilize an algorithm that uses each day's data to update the parameters describing the building, and (3) to optimize multiple zones in a home, each with different uses and occupancy. Other extensions include modeling the time that HVAC appliances take to change inside temperatures, adding in other appliances that permit load shifting (such as dishwashers and clothes washers), and dealing with forecasting uncertainties by more efficient techniques. None of these extensions should change the main results of the paper.

In addition, our simulations can be replicated with data on weather, price variation, and building parameters that are appropriate for different structures, regions, and (real and potential) electricity markets. Simulations in this paper can help users determine the potential savings from an optimizing thermostat.

When coupled with data on the elasticity of supply and demand for electricity, these simulations can also inform policy makers about the potential social value of optimizing thermostats. If these social gains are large, policy-makers may want to encourage adoption of optimizing thermostats and related products that increase the responsiveness of electricity demand to prices.

REFERENCES

- Angilletta, M.J.(2009). "Thermal Adaptation: A Theoretical and Empirical Synthesis." *Oxford University Press*, Oxford.
- Borenstein, S.(2005). "The Long-Run Efficiency of Real-Time Electricity Pricing." *The Energy Journal*, 26(3): 93–116. <http://dx.doi.org/10.5547/ISSN0195-6574-EJ-Vol26-No3-5>.
- Bertsekas, D.P.(1995). *Dynamic Programming and Optimal Control*. Athena Scientific, New Hampshire.
- Black, J.W. and R. Tyagi (2010). "Potential Problems with Large Scale Differential Pricing Programs." *Transmission and Distribution Conference and Exposition*. IEEE PES.

Thermostats for the Smart Grid / 89

- Braun, J.E. (2003). "Load Control Using Building Thermal Mass." *Journal of Solar Energy Engineering* 125: 292–302. <http://dx.doi.org/10.1115/1.1592184>.
- Department of Energy Energy Information Administration (DOE/EIA) (2005). "U.S. Household Electricity Report." Accessed at http://www.eia.doe.gov/emeu/repse/enduse/er01_us.html.
- Deschenes, O., and M. Greenstone (2008). "Climate Change, Mortality and Adaptation: Evidence from Annual Fluctuations in Weather in the U.S." Massachusetts Institute of Technology Department of Economics Working Paper Series.
- Fanger, P.O. (1972). *Thermal Comfort: Analysis and Applications in Environmental Engineering*. McGraw-Hill, New York.
- George, S.S., J. Bode, M. Perry, Z. Mayer and Freeman, Sullivan & Co. (2009). "2009 Load Impact Evaluation for Pacific Gas and Electric Company's Residential SmartRate Peak Day Pricing and TOU Tariffs and SmartAC Program Volume 1: Ex Post Load Impacts."
- Ipakchi, A. and F. Albuyeh (2009). "Grid of the Future." *Power and Energy Magazine, IEEE* 7(2): 52–62. <http://dx.doi.org/10.1109/MPE.2008.931384>.
- Joskow, P.L. (2011). "Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies." *American Economic Review Papers and Proceedings* 101(3): 238–241. <http://dx.doi.org/10.1257/aer.101.3.238>.
- Liang, Y., and Z.-J. M. Shen (2011). "Stochastic Control for Smart Grid Users with Flexible Demand." working paper, University of California, Berkely, September.
- Livngood, D., and R. Larson (2009). "The Energy Box: Locally Automated Optimal Control of Residential Electricity Usage." *Service Science* 1(1): 1–16.
- Oldewurtel, F., A. Ulbig, A. Parisio, G. Andersson, M. Morari (2010). "Reducing Peak Electricity Demand in Building Climate Control Using Real-time Pricing and Model Predictive Control." *Decision and Control*, 2010 49th IEEE Conference: 1927–1932.
- Papavasiliou, A., and S.S. Oren (2011). "Integration of Contracted Renewable Energy and Spot Market Supply to Serve Flexible Loads." 18th World Congress of the International Federation of Automatic Control, August 28–September 2, 2011, Milano, Italy.
- Payne, W.V., and P.A. Domanski (2002). "A Comparison of an R22 and an R410A Air Conditioner Operating at High Ambient Temperatures, R2-1." *Proceedings of the International Refrigeration and Air Conditioning Conference*, Purdue, West Lafayette, IN, 16–19 July 2002.
- Reddy, T.A., L.K. Norford and W. Kempton (1991). "Shaving Residential Air-conditioner Electricity Peaks by Intelligent Use of the Building Thermal Mass." *Energy* 16(7): 1001–1010. [http://dx.doi.org/10.1016/0360-5442\(91\)90060-Y](http://dx.doi.org/10.1016/0360-5442(91)90060-Y).
- Schweppe, F.C., R.D. Tabors, J.L. Kirtley, H.R. Outhred, F.H. Pickel and A.J. Cox (1980). "Homeostatic Utility Control." *IEEE Transactions on Power Apparatus and Systems* vol. PAS-99 (3): 1151–1163. <http://dx.doi.org/10.1109/TPAS.1980.319745>.
- Walker, I. S. and A. K. Meier (2008). "Residential Thermostats: Comfort Controls in California Homes." Lawrence Berkeley National Laboratory.

APPENDIX

A. On CPU Time and Aggregation

In our simulations, we aggregate temperatures into discretized values and make only one decision per period. In this appendix, we study how the length of each period, the number of look-ahead periods, the length of the planning horizon, and the aggregation level of temperatures affect computational requirements (measured by CPU time) and performance of the control algorithms. Let ΔT denote the greatest common divider of the differences between any two discretized temperatures and let Δt denote the length of each period (in hours). N and T denote the number of hours of the planning horizon and the look-ahead

Table 7: CPU Time Tests

ΔT (°F), Δt (hr)	N , T (hr)	Rigid Thermostat		Optimizing Thermostat	
		t_{table} (sec)	t_{total} (sec)	t_{table} (sec)	t_{total} (sec)
$\Delta T = 0.5$ $\Delta t = 0.5$	$N = 24$, $T = 12$	0.017	0.032	132.781	132.799
	$N = 24$, $T = 24$		0.043		132.811
	$N = 24$, $T = 36$		0.056		132.821
$\Delta T = 1.0$ $\Delta t = 1.0$	$N = 24$, $T = 12$	0.008	0.01	30.109	30.112
	$N = 24$, $T = 24$		0.011		30.113
	$N = 24$, $T = 36$		0.012		30.114

Notes: CPU Time Required at Each Step (ΔT is the length of the each interval of which the temperature is aggregated and Δt is the length of each period. N and T are the number of periods simulated and the number of look-ahead periods for making each decision. For all previous simulations, $\Delta T = 1.0$, $\Delta t = 1.0$, $N = 24$, and $T = 24$.)

periods at each step, respectively. t_{total} summarizes the CPU time required to obtain control signals at each step.

In general, the CPU time consumed by the traditional thermostat is negligible, for only one operation is required at each step. The other two control algorithms are coded in MATLAB R2008a and tested on a computer with Intel Core 2 Duo CPU at 2.26GHz and 3 GB of RAM running Windows XP. The CPU time required for obtaining the control signals at the beginning of each period is summarized in Table 7. As expected, finer aggregation of temperatures and time, more look-ahead periods, and longer planning horizon increase CPU time. In our algorithms, we calculate look-up tables consisting of state transitions and the corresponding comfort levels before applying backward recursion. Computing comfort involves computing an integral, which takes much more effort than the backward recursion does. t_{table} records the time required to build the look-up table.⁷ Hence, if the look-up tables are calculated in advance, the real-time computational requirement can be substantially reduced.

The aggregation of temperatures and the assumption of one decision per period reduce computational time, but at the cost of suboptimality, for two reasons. First, the feasible region may exclude the optimal solution. Second, rounding state transitions may increase the error in approximating the value-to-go function. Table 8 summarizes the total utilities of the three types of thermostat with finer aggregation: $\Delta T = 0.5^\circ\text{F}$ and $\Delta t = 0.5$ hr (compared to Table 3, in which $\Delta T = 1.0^\circ\text{F}$ and $\Delta t = 1.0$ hr). Surprisingly, the daily total utilities for the users of the traditional and rigid thermostats drop under finer aggregation. The daily total

7. The rigid thermostat's control algorithm only needs state transition, so constructing its look-up table does not involve computing comfort.

Table 8: Simulation Tests Results with Finer Discretization

	Traditional			Rigid			Optimizing		
	C_{day}	D_{day}	U_{day}	C_{day}	D_{day}	U_{day}	C_{day}	D_{day}	U_{day}
Case 1	5.68	-0.42	-6.1	2.55	-2.45	-5	3.52	-0.51	-4.03
Case 2	6.3	-0.42	-6.72	2.9	-2.41	-5.31	3.76	-0.55	-4.3
Case 3	6.3	-0.7	-7.01	2.9	-4.45	-7.35	4.07	-0.68	-4.75
Case 4	6.3	-1.27	-7.57	2.9	-8.53	-11.43	4.66	-0.65	-5.31
Case 5	3.4	-0.45	-3.85	0.72	-2.35	-3.07	1.78	-0.24	-2.02
Case 6	18.91	-0.42	-19.33	8.71	-2.41	-11.11	3.91	-5.16	-9.07
Case 7	19.41	-0.42	-19.83	8.58	-2.66	-11.25	5.06	-3.88	-8.94
Case 8	19.41	-0.42	-19.83	7.88	-10.89	-18.77	5.06	-3.88	-8.94
Case 9	19.41	-0.7	-20.11	8.58	-4.94	-13.52	9.67	-2.26	-11.93

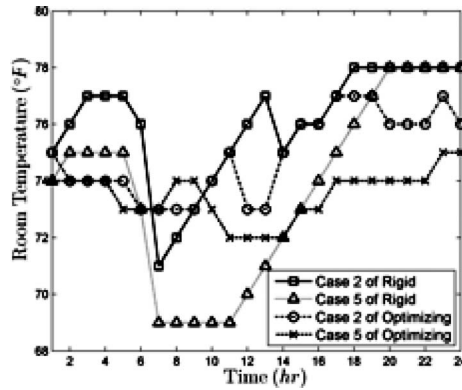
Notes: Simulation Results with Finer Discretization for $\Delta T = 0.5^\circ\text{F}$ and $\Delta t = 0.5$ hr. The cost, normalized value of discomfort, and total utility each day (C_{day} , D_{day} , and U_{day}) are all expressed in dollars.

utility decreases for the users of the traditional thermostat because it consumes more electricity in order to confine the room temperature to a tighter range around T_{pfr} , whereas the gain in comfort is small. Meanwhile, the rigid thermostat consumes less energy (6.5%) by pushing room temperatures closer to the rigid bounds, but the comfort (D_{day}) drops more than the cost saved. In contrast, finer aggregation and shorter period length improve the total utility for the users of the optimizing thermostat. The maximum and average relative improvements are 11.1% and 4.3%. However, according to Table 7, the CPU time for the optimizing thermostat is increased by over 300%. Therefore, it is important in practice to find a balance between the increased computational cost and the improved total utility from finer aggregation and shorter period length.

B. More Sensitivity Analysis and Insights

1. The rigid thermostat sometimes generates the lowest total utility. If price varies a lot within a day or if the room is very well insulated—both of which encourage extensive morning cooling—then the rigid thermostat holds the room temperature on the rigid lower bound for longer. As a result, users will experience chilly summer mornings. Thus, sometimes, as in cases 3 and 4, the rigid thermostat generates the lowest total utility. In all the other cases of our simulation, the traditional thermostat has the lowest total utility.

2. Better insulation is in general beneficial, but it can create discomfort for the users of the rigid thermostat. All three thermostats generate higher total utility with better insulation, as in case 5, than they do in the otherwise similar

Figure 10: Room Temperature with Better Insulation

case 2 (Table 3). Intuitively, better insulation preserves indoor temperature by slowing down heat loss; thus, electricity consumption is reduced. However, the resulting room temperatures are not necessarily more comfortable. For instance, Figure 10 suggests that the rigid thermostat leads to colder and more volatile room temperature because better insulation makes pre-cooling more effective. As a result, the rigid thermostat pre-cools more. On the contrary, because canceling heat loss takes less electricity with better insulation, the optimizing thermostat keeps the room temperature closer to the users' preference, T_{pfr} . Therefore, users with the optimizing thermostat are better off in terms of both cost and comfort.

C. Incorporating the Acclimation Effect by Using the Optimizing Thermostat

Humans acclimate (adapt) to recent temperatures (Angilletta, 2009), and effect that can be exploited by the optimizing thermostat. Due to acclimation, users who come into a room at 77°F when the ambient temperature is over 110°F will most likely feel just as comfortable as they would if the room temperature were 73°F, the most preferred one (T_{pfr}), and the ambient temperature were 80°F. By setting T_{pfr}^* as a function of a user-specified "normal" preferred temperature and of the recent history of ambient temperatures, the optimizing thermostat can easily incorporate acclimation. For example:

$$T_{pfr,n}^* = g(T_{pfr,n}, T_{out,n})$$

Then we could utilize $T_{pfr,n}^*$ instead of $T_{pfr,n}$ in the control algorithm of the optimizing thermostat to capture the acclimation effect. In addition, comfort depends on factors such as humidity as well as temperature. Learning the function $g(\cdot)$ is critical and general statistical learning models could be applied, but that

discussion is beyond the scope of this paper. Nonetheless, so long as we have a model for $T_{pfr,n}^*$, acclimation effect can be adapted by using the optimizing thermostat. On the contrary, it is more difficult to let a rigid thermostat adjust its control according to ambience, as users' comfort is not in the objective function.

D. Proofs

Proof of Theorem 1:

Note that P1 is equivalent to the following problem:

$$\begin{aligned} \max_{\{X_{tar,n}\}_{n=0}^N} \quad & \sum_{n=0}^N \alpha^n F^C(T_{ini,n}, X_{tar,n}) \\ \text{s.t.} \quad & T_{ini,n+1} = \left[f^T(T_{out,n}, X_{tar,n}, h) + \frac{1}{2} \right], \quad \forall n = 0, 1, 2, \dots, N \\ & T \in [I_{min,n}, I_{max,n}] \end{aligned}$$

where the objective function can be expanded as:

$$- \sum_{n=0}^N \alpha^n p_n k_1 |T_{ini,n} - X_{tar,n}|$$

Suppose the electricity prices $\{p_n\}_{n=0}^N$ are replaced with $\{\lambda \bar{p}_n\}_{n=0}^N$, where $\lambda > 0$ is the scale and $\{\bar{p}_n\}_{n=0}^N$ is the normalized price sequence. Then dividing the objective function by λ will affect neither the feasible region nor the ordering of candidate solutions. Hence, under price sequences that share the same normalized $\{\bar{p}_n\}_{n=0}^N$, the optimal solutions are the same.

Proof of Theorem 2.

- The proof of Theorem 2 utilizes the same approach as that of Theorem 3. Please refer to the proof of Theorem 3.

Proof of Proposition 1.

First, setting a different target temperature will result in a different total utility for the same user. This difference in utility can be obtained by combining Equations (4) and (7). Let \tilde{X}_{tar} and \hat{X}_{tar} be two proposed target temperatures and let $\tilde{X}_{tar} = \hat{X}_{tar} + \Delta T$; then the difference $\Delta F^D = F^D(T_{out}, T_{pfr}, \tilde{X}_{tar}) - F^D(T_{out}, T_{pfr}, \hat{X}_{tar})$ is:

94 / The Energy Journal

$$\begin{aligned}
\Delta F^D &= \int_0^h \{ b[|f^T(T_{out}, \tilde{X}_{tar}, t) - T_{pfr}| - |f^T(T_{out}, \hat{X}_{tar}, t) - T_{pfr}|] \\
&\quad + a[(f^T(T_{out}, \tilde{X}_{tar}, t) - T_{pfr})^2 - (f^T(T_{out}, \hat{X}_{tar}, t) - T_{pfr})^2] \} dt \\
&= \int_0^h \{ b[|f^T(T_{out}, \tilde{X}_{tar}, t) - T_{pfr}| - |f^T(T_{out}, \hat{X}_{tar}, t) - T_{pfr}|] \\
&\quad + a(\tilde{X}_{tar} - \hat{X}_{tar})e^{-kt} [2(T_{out} - T_{pfr}) + (\tilde{X}_{tar} - T_{out})e^{-kt} + (\hat{X}_{tar} - T_{out})e^{-kt}] \} dt \\
&= \int_0^h b[|f^T(T_{out}, \tilde{X}_{tar}, t) - T_{pfr}| - |f^T(T_{out}, \hat{X}_{tar}, t) - T_{pfr}|] dt \\
&\quad + a \int_0^h \{ (\tilde{X}_{tar} - \hat{X}_{tar})e^{-kt} [2(T_{out} - T_{pfr}) + (\tilde{X}_{tar} + \hat{X}_{tar} - 2T_{out})e^{-kt}] \} dt \\
&= \int_0^h b[|f^T(T_{out}, \tilde{X}_{tar}, t) - T_{pfr}| - |f^T(T_{out}, \hat{X}_{tar}, t) - T_{pfr}|] dt \\
&\quad + a(\tilde{X}_{tar} + \hat{X}_{tar} - 2T_{out})(\tilde{X}_{tar} - \hat{X}_{tar}) \int_0^h e^{-2kt} dt + 2a(T_{out} - T_{pfr}) \\
&\quad (\tilde{X}_{tar} - \hat{X}_{tar}) \int_0^h e^{-kt} dt
\end{aligned}$$

Obviously, the last term is fixed because it is only a function of ΔT . The middle term is a function of both ΔT and \tilde{X}_{tar} . The first term satisfies the following inequality:

$$\int_0^h b[|f^T(T_{out}, \tilde{X}_{tar}, t) - T_{pfr}| - |f^T(T_{out}, \hat{X}_{tar}, t) - T_{pfr}|] dt \leq \int_0^h b|\tilde{X}_{tar} - \hat{X}_{tar}| e^{-kt} dt$$

hence the upper bound on the first term is also a function of ΔT . When B_n^u and B_n^l satisfy the conditions as described in the proposition, we can always be better off by switching the target temperature from $X_{tar,N} > B_n^u$ to B_n^u (or $X_{tar,N} < B_n^l$ to B_n^l) because the decrease in comfort from this adjustment is bounded below by the increase in cost. Therefore, the optimal solutions at time T of **P2** are interior points of feasible region $[B_n^l, B_n^u]$. Moreover, $V_{opt_N}^*(x)$ is non-decreasing for $\{x: x \geq B_n^u\}$ and non-increasing for $\{x: x \leq B_n^l\}$.

It can then be easily verified by mathematical induction that (1) we can always be better off by switching the target temperature from $T > B_n^u$ to B_n^u (or $X_{tar,n} < B_n^l$ to B_n^l) for the same reason and (2) $V_{opt_n}^*(x)$ is non-decreasing for $x: x \geq B_n^u$ or $x \leq B_n^l$. As a result, the optimal solutions of **P2** are interior points of the region defined by the conditions of Proposition 1, which is the feasible region of **P2**.

Thermostats for the Smart Grid / 95

If we adopt the COP-based model for room temperatures, this proposition still holds after we replace k_1 with $\frac{COP_o k_1}{COP_{lb}}$. In addition, if COP_x is decreasing at \tilde{B}_n^u and increasing at \tilde{B}_n^l , then replacing k_1 with $\frac{COP_o k_1}{COP_{B_n^j}}$ ($j \in \{u, l\}$ and $n = 1, 2, \dots, N$) is sufficient to guarantee that the optimal solutions of **P2** are interior points of the region defined by the modified conditions.

Proof of Theorem 3.

- The following inequality:

$$V_{tra,n}^*(T_{ini,n}) \leq \tilde{V}_{opt,n}^*(T_{ini,n})$$

that is, the first inequality in Theorem 3 holds, because the modified optimizing thermostat optimizes total utility from both cost and comfort over a broader feasible region, while the policy of the traditional thermostat is a feasible solution to the former one.

- The other inequality:

$$V_{rigid,n}^*(T_{ini,n}) \leq \tilde{V}_{opt,n}^*(T_{ini,n})$$

holds for the same reason as explained above.

- The second inequalities in both equations comes from the fact that the feasible region of the modified problem $\tilde{\mathbf{P2}}$ is a subset of that of the original problem **P2**.

Proof of Corollary 1.

The optimizing thermostat costs less than the rigid and traditional thermostats when it generates lower or equal comfort because the former creates the highest total utility for its users, as stated in Theorem 3.

Compared with the rigid thermostat, the optimizing thermostat creates more comfort for users when the obtained room temperatures are within $[I_{min}, I_{max}]$. Although not readily seen, this is also consistent with Theorem 3. If the obtained room temperatures of the optimizing thermostat are bounded by $[I_{min}, I_{max}]$, the cost of using the optimizing thermostat is greater or equal to that of using the rigid thermostat, as the optimal solution generated by the optimizing thermostat is also a feasible solution to the rigid thermostat's problem. Therefore under this condition, Theorem 3 implies that the optimizing thermostat is able to provide greater comfort, as the optimizing thermostat creates the highest total utility for its users.

