

Heat Pump Water Heater Electric Load Shifting: A Modeling Study



June 8, 2018

Prepared by:

Nick Carew, Ben Larson, Logan Piepmeier, and
Michael Logsdon
Ecotope, Inc.

Table of Contents

Introduction	3
Background	3
Methodology	4
Simulation	4
Simulation Inputs	4
Hot Water Draw Patterns	4
Lab Testing	5
Test Plan	6
Testing Outcomes.....	6
Simulations	9
Price Signals	9
TDV+NEM2	9
Utility Marginal Costs	10
Residential Time-of-Use Electricity Rates	10
Control Strategies.....	10
Simplest: Allow/Block.....	11
Smarter: Load-Up & Shed.....	11
Smartest: Optimal Price	14
Bi-Directional: State-of-Charge	17
Results	20
Sensitivity of Customer Savings	24
CO ₂ HPWH Discussion.....	24
Conclusions	25
References	28

Introduction

The National Resources Defense Council (NRDC)-Ecotope HPWH Load Flexibility project aims to assess the potential for load flexibility via heat pump water heaters (HPWH). The concept is simple: to shift load away from unfavorable times, create hotter water when favorable conditions exist and store the energy in the HPWH tanks to be used later. This study primarily focuses on shifting load away from times when electricity is expensive, although the developed strategies could be used to avoid any unfavorable condition with an hourly schedule (e.g. greenhouse gas emissions).

For the purposes of this study, a series of price schedules were analyzed representing the costs as seen by the various stakeholders (i.e. utilities & consumers). These price schedules, along with the HPWH simulation engine (HPWHsim¹) developed by Ecotope, provide the basis for the optimization. The objective is to find the lowest cost without running out of hot water. Several control strategies were developed, with various levels of complexity, to transform these price schedules into control signals that dictate the operation of the HPWHs.

The strategies are tied to the functionality of the equipment. All but one of the strategies require additional control functions that do not currently exist in HPWHs, making this a future-looking exercise. The strategies range from creating a simple on/off signal to dynamically changing the HPWH setpoint temperature. The resulting control signal schedules were then fed into HPWHsim to investigate the effects on overall cost of operation, energy consumption, hot water delivery, and greenhouse gas (GHG) emissions.

Background

Electric storage water heaters can act as thermal batteries. They can be controlled to shift operation from peak hours to off-peak hours, and to heat water at a higher temperature than the conventional setpoint to maximize storage capacity. This demand flexibility can help reduce peak demand and better utilize renewable energy when it is abundant, helping balance the grid, integrate high levels of renewable energy, and reduce GHG emissions.

This practice is already in limited use with electric resistance water heaters in some U.S. regions, but it is just emerging for HPWHs due to their smaller, yet growing, market share. Given the growing interest in the efficiency and emissions reduction benefits of HPWH, this project aims to assess the potential for demand flexibility by HPWH in California, in support of the state's 2019 building energy code, integrated resource planning (IRP), and utility incentive programs. The simulation is focused on California, and uses state-specific climate, water usage, and grid conditions. However, the findings can be informative to other states and regions. Further, the same methods developed here, could be employed in other regions given those state-specific circumstances.

¹ <https://github.com/EcotopeResearch/HPWHsim>

Methodology

Simulation

To help understand and estimate the amount of energy savings that can be achieved through heat pump water heating, a HPWH simulation (HPWHsim) was written by Ecotope beginning in 2012 (Ecotope Inc., 2015). The simulation was designed to run quickly, as the typical use case would see many simulations run, each representing a year of activity. The HPWHsim engine is currently used by the California Building Energy Compliance Calculator (CBECC-Res²) and by the Regional Technical Forum to determine energy use associated with mass deployment of HPWHs in the residential sector (Regional Technical Forum, 2018).

The simulation was designed to model storage tank water heaters, specifically HPWHs. The configuration of a HPWH varies between different models so the simulation must accommodate these variations. Common variables are the number and position of electric resistance elements, the arrangement of the condensing coils, and the performance of the compressor system, among others. In addition to the physical properties of the HPWH and the topology of its heat sources, each model of HPWH has a unique set of criteria which direct it to engage or disengage its various heat sources. This control logic is specifiable as well, by choosing from a set of standard decision criteria, such as the temperature of the top third of the tank, and supplying setpoints for those criteria.

The simulation takes a hot water set-point, inlet water temperature, and ambient space temperature and steps through a draw schedule at one-minute increments, tracking tank temperature and activating heating components accordingly. The inlet water and ambient space temperatures are dependent on the climate zone and the location of the water heater in the house. It is assumed for this project that all water heaters are located in the garage since this is the case for the overwhelming majority of residential water heaters in California. The draw profile is supplied to the simulation by the user.

Simulation Inputs

An important guideline for the project was to align as closely as possible with existing modeling tools and assumptions in California. Therefore, the important inputs to the simulation, including inlet water temperature, ambient space temperature, and draw profile all come from CBECC-Res. Specifically for this project, the CBECC-Res software team provided hourly values for inlet water temperature, garage temperature, outside air temperature, and draw profiles for 1-, 2-, 3-, 4-, and 5-bedroom prototype houses across California's 16 climate zones. These values are either direct inputs to CBECC-Res from weather files or calculations the software performs. By using these hourly values, we are able to run HPWHsim outside of CBECC-Res which was necessary due to the sheer number of simulations and computation time required.

Hot Water Draw Patterns

The hot water draw patterns used in this study warrant special mention. As a general comment, the operation of HPWHs, and especially those with resistance elements, is highly complex, interdependent,

² <http://www.bwilcox.com/BEES/cbecc2016.html>

and non-linear. The solution for overcoming the interdependent complexity of draws, control logic, and operating conditions, is to simulate using draw patterns with as much randomness as possible. Simulating the same daily draw profile over and over again for an entire year creates an estimate that is extremely fragile and sensitive to operating conditions, such as ambient temperature, inlet water temperature, and setpoint. It is preferable to either define a unique draw profile for each day of the year based on real-life water usage data, or for a more compact representation repeating a given daily/weekly draw pattern with added stochastic variations. The draw profiles used in CBECC-Res follow the latter approach and provide exactly what we need.

The draw schedules used in CBECC-Res are based off real hot water usage data that is linked to house occupancy.³ Briefly described, an annual draw pattern consists of discrete, real, weekday, weekend, and holiday days. The total annual draw volume is determined by the number of bedrooms selected for a house so that 1-bedroom households have the lowest use and volume increases with bedroom count. Within the annual patterns, weekday draw patterns, which typically peak in the morning and evening, are assigned to real, calendar weekdays. Likewise, for weekends, which tend to have more spread-out use, and also for holidays. In all, there are a total of 48 unique draw pattern days that are mixed and matched to create each annual draw pattern. Further, for draw events like showers, where a user balances hot and cold water to mix to a set temperature, the hot water volumes change over the year. This is due to the changing cold water temperature – more hot water is needed in the winter to mix with cold to create the same shower temperature. All of this variation is very effective in reflecting the complex reality of hot water usage. Consequently, the simulations in this project have a reliable foundation of inputs to rest on.

Lab Testing

To both inform and ensure the accuracy of the simulation, laboratory testing was completed for a range of current product offerings from Bradford White, Rheem, AO Smith, and Sanden. We investigated a range of nominal tank sizes spanning 50 to 80 gallons. The Bradford White, Rheem, and AO Smith products are all integrated, hybrid units using R-134a refrigerant. That is, they have the heat pump components sitting atop the tank and have two electric resistance elements. The Sanden product has the heat pump components, using CO₂ refrigerant, located remote from the tank in a package designed for outside installations.

A substantial body of testing data was available as a result of work funded by the Northwest Energy Efficiency Alliance to study most resident HPWHs on the market (Kvaltine & Larson, 2015a, 2015b; Larson & Kvaltine, 2015). That previous work was more general in nature whereas input needs for the load shifting aspects of the simulation required specific additional information. In particular, previous lab work had characterized compressor efficiency only up to typical tank water operating temperatures of 125°F-135°F. The heat pump efficiency, output capacity, and input power all vary substantially over the range of water temperatures encountered in the tank. Consequently, the simulations in this project explored set points as high as 160°F so we needed lab measurements in that range to accurately predict performance.

³ <http://www.bwilcox.com/BEES/docs/Kruis%20-%20Dhw%20Analysis%205.docx>

Under the default controls for these HPWH products, there is a fixed, upper limit set point. The value varies slightly by the equipment model but is typically 145°F. The limit is not due to physical limitations of the R-134a vapor compression cycle so we asked the product manufacturers for special controls on the water heaters to observe operation at even higher temperatures. Two of the three integrated product manufacturers were able to provide tanks that allowed the setpoint to reach in excess of 165°F which proved extremely beneficial to the testing. Note that previous lab measurements of the split-system, CO₂ Sanden HPWH had characterized performance well in to the high water temperature range so further detailed measurements were not explored for this product.

Test Plan

The laboratory testing, specific to the needs of this project, was completed at the Pacific Gas & Electric (PG&E) Applied Technology Services test facility in San Ramon (CA). A test matrix was developed by Frontier Energy and Ecotope consisting of three different test types:

- **Coefficient of Performance (COP)** – The relatively simple goal of the COP tests was to map the heat pump performance as a function of water and air temperature.
- **Behavior/Changing Controls** – The goal of the Behavior testing was to identify the behavior of the HPWH in response to changes in setpoint, and in response to potential changes in control logic. These were used to explore which heating components, under the manufacturer’s current control scheme, turn on or off after the setpoint was dramatically increased or decreased. We would expect future products, designed with full load shifting capability in mind, to have different control responses. Further, the water heater’s default controls may be moot because a fully functioning load shifting program would likely implement direct control over the heat pump and resistance element operation, rather than setpoint control.
- **Draw Profiles** - The Draw Profile testing serves two purposes. The first is to demonstrate, with actual equipment, as opposed to a computer simulation, that the load shifting we are considering is possible. Second is to provide a data set that we can validate the simulation against. For this task, two demanding draw patterns, each comprising a day, from the CBECC-Res draw profiles were selected. Those two days we explicitly selected to demonstrate the ability of the water heater to ride out the late afternoon / early evening peak time under challenging circumstances. The original project design had hoped to control each water heater with a remote signal. Due to limitations in both the products themselves and the lab environment, this was not possible. Instead, control was exercised over the equipment by manual changes in set point. A “load-up” event was simulated by increasing the tank set point. Then, during the peak price period, the tank set point was lowered.

Testing Outcomes

COP Mapping

The final output of the COP testing is both a plot and calculation of the COP of the compressor as a function of the average temperature of water in the tank. This piece of the testing is critical to understanding the effects on performance of raising the setpoint beyond typical temperatures; a sample finding is presented in Figure 1. The testing is carried out at three ambient air conditions: 50°F, 67.5°F, and 95°F. For ambient air conditions between or beyond those measured, the performance data are interpolated or extrapolated. Note that for garages with water heaters located in them, the typical range of air temperatures is 35°F-110°F. In terms of water temperature, Figure 1, representative of integrated

HPWHs, shows that the heat pump COP drops from ~2.8 at 125°F to ~1.8 at 155°F with an ambient temperature of 67°F. This data is encoded and exercised in the simulation which allows us to dynamically evaluate the tradeoff between the benefits of higher temperature storage and costs of lower efficiency.

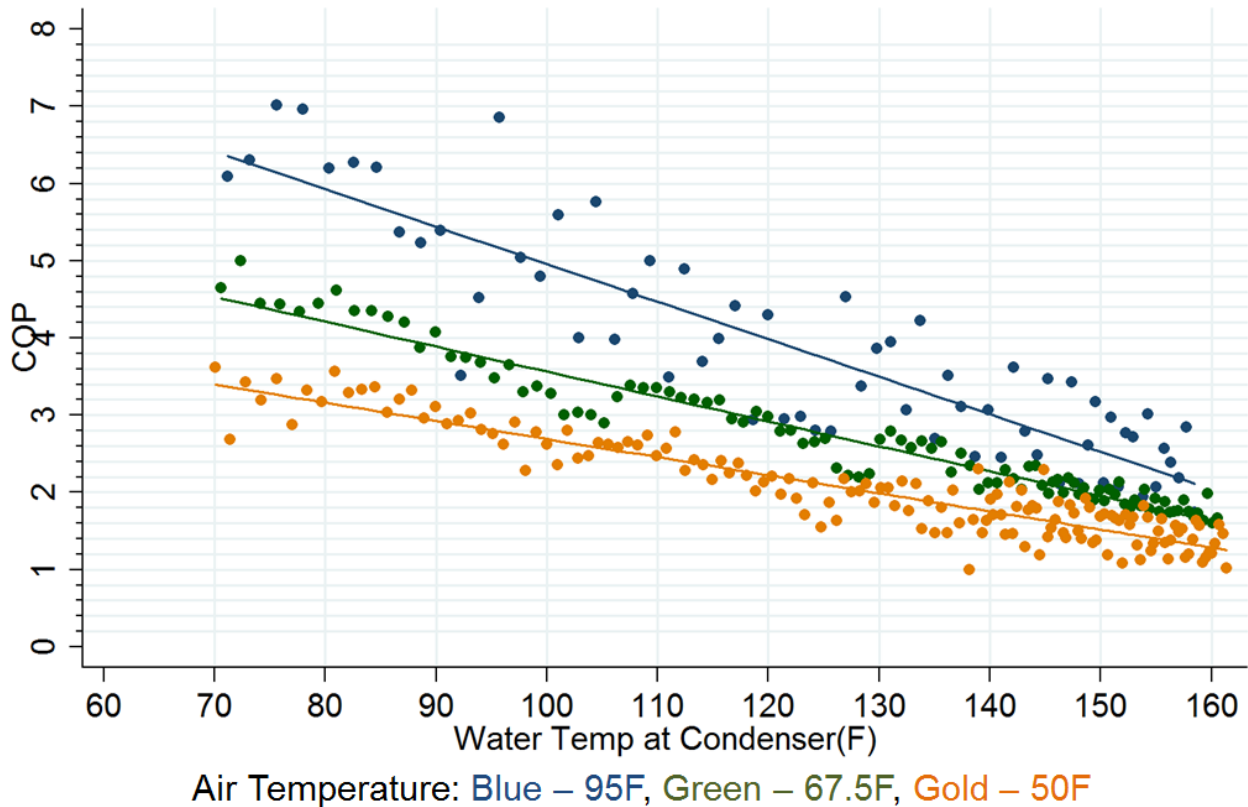


Figure 1. Measured Compressor Efficiency

Behavior Observations

The Behavior observation tests showed that for the default, hybrid mode controls, our previous understanding of when, and in what sequence, the compressor and resistance elements operate still applied to the higher water temperature setpoints. In particular, we were able to see that, even for a water heater already at 120°F, a large increase in the setpoint usually triggered the resistance element to engage. Similarly, starting with a 150°F tank, decreasing the setpoint to 120°F and conducting a large hot water draw triggered the resistance elements on as if there wasn't an added benefit of having 150°F water at the top of the tank. In the end, we were able to use these test results to inform how we would conduct the Draw Profile tests. The plan was to step the setpoint up in smaller increments to avoid the resistance heat use.

The observations also informed how we conducted the simulations. For one, we reasoned future equipment, with specific demand response controls would be designed to minimize the use of resistance heat. Therefore, we changed element temperature controls to be on an absolute basis and not one relative to setpoint. This made effective use of the extra hot water storage in the tank. Second, much like the way we conducted the Draw Profile tests in the lab, we deployed a setpoint increase in smaller steps to avoid unnecessary resistance heat use when loading up the tank.

Draw Profile Validation

The simulation controls have been previously calibrated and validated under baseline operation – that is when the tanks operate at a constant set point and are not controlled to shift load (Ecotope 2015 and RTF 2018). This project took the validation one step farther to assess how the simulation compared to measured load shifting draw profiles in the lab. In picking draw profiles to use, we turned to the CBECC-Res collection and selected two out of a possible forty-eight daily profiles. Those profiles are identified as 2D1 and 4E1. They were selected because they are the largest profiles from the collection that coincide with peak electricity draw. In other words, they have large afternoon/evening water uses. They also total 95 and 114 gallons per day respectively. Importantly, these are not the average draw profiles found in CBECC-Res which are much lower. Again, these were purposefully selected for their heavy use. To replicate a load-up and shed scenario, the lab test elevated the tank set point at 9am and then returned it to 125F at 5pm so the water heater could coast through the evening peak. The water heaters were left in hybrid mode.

The validation process was one of seeing how closely the simulation could replicate the observed behavior. The simulation was provided with the same draw profile, inlet water temperature, ambient air temperature, and set point schedule over the 24 hours of the lab test. Then, the resulting simulated energy was compared to the measured energy.

The starting point for of the calibrated control parameters were those used in CBECC-Res for the water heaters in question. Those were calibrated to ~5 lab tests – mostly available from compliance testing with the Advanced Water Heater Specification. Next, before load shifting simulations were conducted, we had an additional, limited set of lab tests from PG&E ATS facility which gave us more information with which to tweak the control parameters. These were the behavior tests but not the full suite of draw profile results. We used these because, generally, more information is better. Then, the load shifting simulation work was conducted. Subsequently, with the draw profile tests complete from the lab, we assessed the ability of the simulation to replicate the measured behavior (validation step).

The comparison showed the simulation over-predicted energy use of the 2D1 and 4E1 profiles in both the fixed and variable set point cases by 15-20% for two of the integrated HPWHs and by >35% for a third. The finding does not necessarily indicate a flaw in the simulation that needs correction, but rather, serves to highlight the difficulties with both calibrating and validating a simulation. To do both, one needs copious data points which is generally hard to come by in the water heating realm. First, the dataset to which the simulation was calibrated was relatively small. Second, the dataset the simulation was validated against (the 2D1 and 4E1) was not representative of typical water use. Those draw profiles are skewed to high demand situations. Overall, we are not so much interested in how the simulation compares to those scenarios but, instead, how it compares with the average water heating use. In retrospect, it is clear a more thorough validation dataset would be one that involved 10-20 days of lab tests, which enforced variability of low, average, and high draw days or 30+ days of field data. Such datasets would likely contain enough diversity as well as average behavior to provide a useful platform for calibration.

As it stands, the validation work shows the simulations could be over-predicting energy use somewhat in load shifting scenarios. If that is the case, we could expect real world installations to offer slightly higher energy and cost savings from load shifting than modeled here. Further, the simulations for the third unit, which showed over-prediction of more than 35% were excluded from our ultimate findings summaries. Last, the split-system HPWH was excluded from the validation analysis due to project budget and time

constraints, however, that water heater has no resistance element whose operation was the focus of the validation efforts.

Simulations

With the independent simulation inputs of the draw patterns, inlet water temperature, and ambient air conditions established and the simulation further informed by the lab testing results, we move on to consider the price signals sent to the water heater and possible ways to respond to them (algorithms or control strategies). Since we did not know *a priori* the best control algorithms, it was important to consider a wide range of parameters to identify the most effective load shifting method. Additionally, we needed to consider how well the methods worked across a wide range of climate zones and household sizes. Doing both implies exploring a parameter space over many simulation runs. Table 1 displays the parameters of the optimization and exploration process. It is worth noting that the baseline case was taken to be uncontrolled water heaters with a 125°F setpoint.

Table 1. Optimization Parameters

Parameter	# of values
<i>Load shaping metric to optimize (Price Signals):</i> TDV&NEM2, Utility Marginal Costs, Time-of-Use	3
<i>Units:</i> Generic Resistance Heat Tanks [50, 65, & 80 gallons], Bradford White [50, 80, 80(hp-only)], AOSmith HPTU [50, 66, 66(hp-only), 80], Rheem Gen4 [50, 50(hp-only), 65, 80], Sanden Gen3 [80]	15
<i>Max temperature:</i> 125, 135, 145, 155	4
<i>Climate zone:</i> CZ1-CZ16	16
<i>House size:</i> 1-5 bedrooms	5
<i>Control algorithms:</i> Allow/Block, Load-Up/Shed, Optimal Price, Optimal Price w/ Cold Weather Load-Up, State-of-Charge	5
Total	72,000

Price Signals

The control signals provided to the HPWHs are based solely on the price signals provided as inputs to the algorithms. Three different price signals were considered in this study as described below.

TDV+NEM2

The California Energy Commission (CEC) has developed a time-dependent valuation (TDV) of energy for use in California's building energy codes. The concept behind TDV is that energy efficiency measure savings should be valued differently depending on which time of day and day of the year the savings occur, to better reflect the actual costs of energy to consumers, to the utility system, and to society (California Energy Commission, 2017). The TDV price signal is essentially a 30-year present value projection of grid energy costs, calculated for each climate zone in California. The values of TDV are constructed from a long-term forecast of hourly electricity, natural gas, and propane costs to building owners consistent with the latest CEC forecasts and outlook for California's energy sectors. The time dependent nature of TDV reflects the underlying marginal cost of producing and delivering an additional unit of energy, similar to a time of use retail tariff, and the resulting economic signal aligns energy savings with the cost of producing and delivering energy to consumers (California Energy Commission, 2017).

The *NEM2* component of the price signal stands for Net Energy Metering version 2 which compensates user-produced electricity exports at the retail electricity rate less 2-3 cents of non-bypassable charges, or the *NEM2*-adjusted retail rate. This component of the price signal was included to reflect the 2019

building code for homes with solar production, as it is expected that most new buildings will have solar starting in 2020 as a result of the 2019 building code update.

The resulting hourly price signal is dynamic relative to a fixed time-of-use rate schedule, however, still relatively flat for much of the year. Large price spikes are relatively few thus limiting the opportunity for significant amounts of load shifting. And roughly half of the annual TDV value comes from a “retail adder”, which limits how low TDV values go, even when wholesale marginal prices are negative.

Utility Marginal Costs

The source of this price schedule is the PG&E 2024 hourly marginal costs. The utility marginal costs factor in the costs of energy, emissions, capacity, transmission and distribution of electricity. These costs do not include a retail rate adder, they reflect the actual marginal costs to the utility, which may (and does) dip to negative prices when production exceeds demand in the middle of the day when solar peaks and inflexible baseload resources cannot ramp down sufficiently. This provides an opportunity for the utility to actually get paid to use connected HPWHs to store energy during the surplus power hours and then lower peak demand later.

Residential Time-of-Use Electricity Rates

Time-of-use (TOU) retail rates better represent the value of load shifting to the consumer. These rates can change hourly but, in practice, are set for several specific blocks of time each day for a given season. Table 2 details the breakdown of TOU prices by day type.

Table 2. Time-of-Use Price Schedule Detail

Season	Day Type	2hr Morning Peak	2hr Pre-Peak Shoulder	4hr Afternoon Peak	3hr Post-Peak Shoulder
Winter	weekday	X		X	X
	weekend			X	
Summer	weekday	X	X	X	X
	weekend			X	X

The 4-hour afternoon peak period is always the most expensive time of day. Shoulder periods are sometimes present on either side of the peak when the price is elevated but not as expensive as during the peak.

TOU prices are much more stable than the TDV+NEM2 and Utility Marginal Costs price signals. This repeating pattern is the only price schedule that lends itself to the simple, timer-based control strategies (discussed subsequently), or to a local, non-grid connected, control module using a fixed TOU price schedule.

Control Strategies

The control algorithms can be broken down into several levels of complexity (or “smartness”) as described below.

Simplest: Allow/Block

This strategy was designed to respond to the fixed price TOU price signal. The idea is that the user can install a timer-based control to prevent the water heater from running in the most expensive price windows. Figure 2 shows the concept functioning under various seasons and day types.

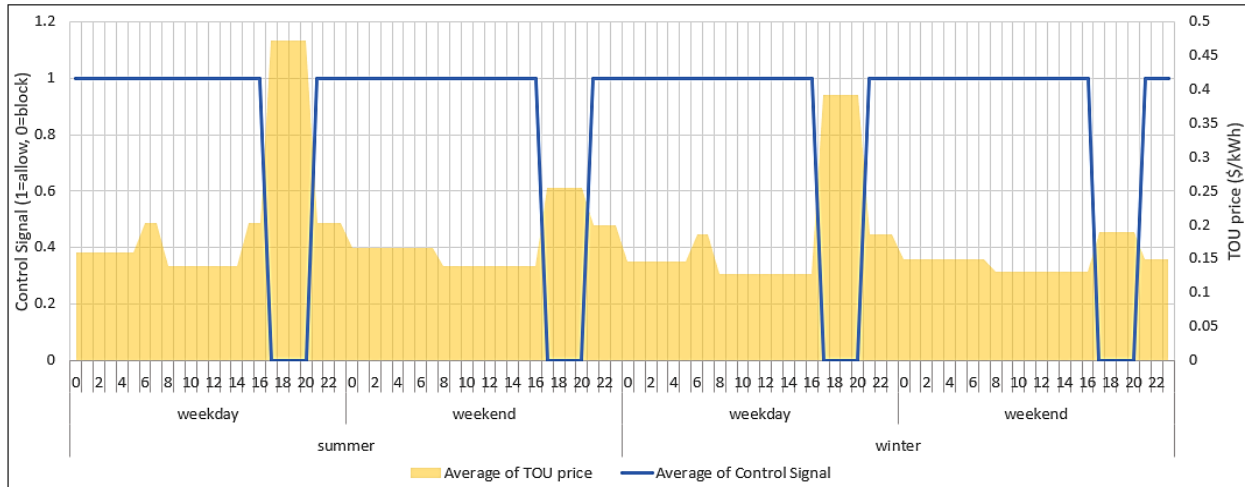


Figure 2. Simple allow/block strategy

The benefits of this strategy are that it requires no change to current technology, it uses a simple low-cost external timer and can be installed by the user (or an electrician/plumber), and it could have an immediate benefit. The downside being that the user may run out of hot water in the time that the water heater is forced off.

Smarter: Load-Up & Shed

The “load-up” strategy was also designed to respond to the fixed price TOU signal, however, an additional *engage* control signal was added. The goal of this strategy is to reduce the frequency of hot water runout events by “loading up” the energy stored in the tank prior to the peak price periods when the water heater will be blocked from operating. This is where the engage control signal is required to force the tank to charge up, via the heat pump, when it would not normally run. Loading up just before peak periods, as opposed to a constantly elevated setpoint, also reduces efficiency losses from super-heating water. While this strategy will reduce runout events and efficiency losses compared to the simpler strategy, it is still only applicable to the highly predictable TOU price signal and therefore has limited capabilities.

The initial approach for this strategy accomplished this by simply increasing the setpoint for a set number of hours immediately preceding the blocked off time window. This essentially forces the water heater on since the water temperature will suddenly be well under setpoint. If the new setpoint is reached prior to the conclusion of the load-up period, then the heat pump will shut off until the water temperature drops enough to trigger the heat pump once again. A variety of setpoints and hours of load-up were investigated to identify the solution with that minimized runouts without causing excess energy consumption. Figure 3 provides a visual of how the algorithm functioned.

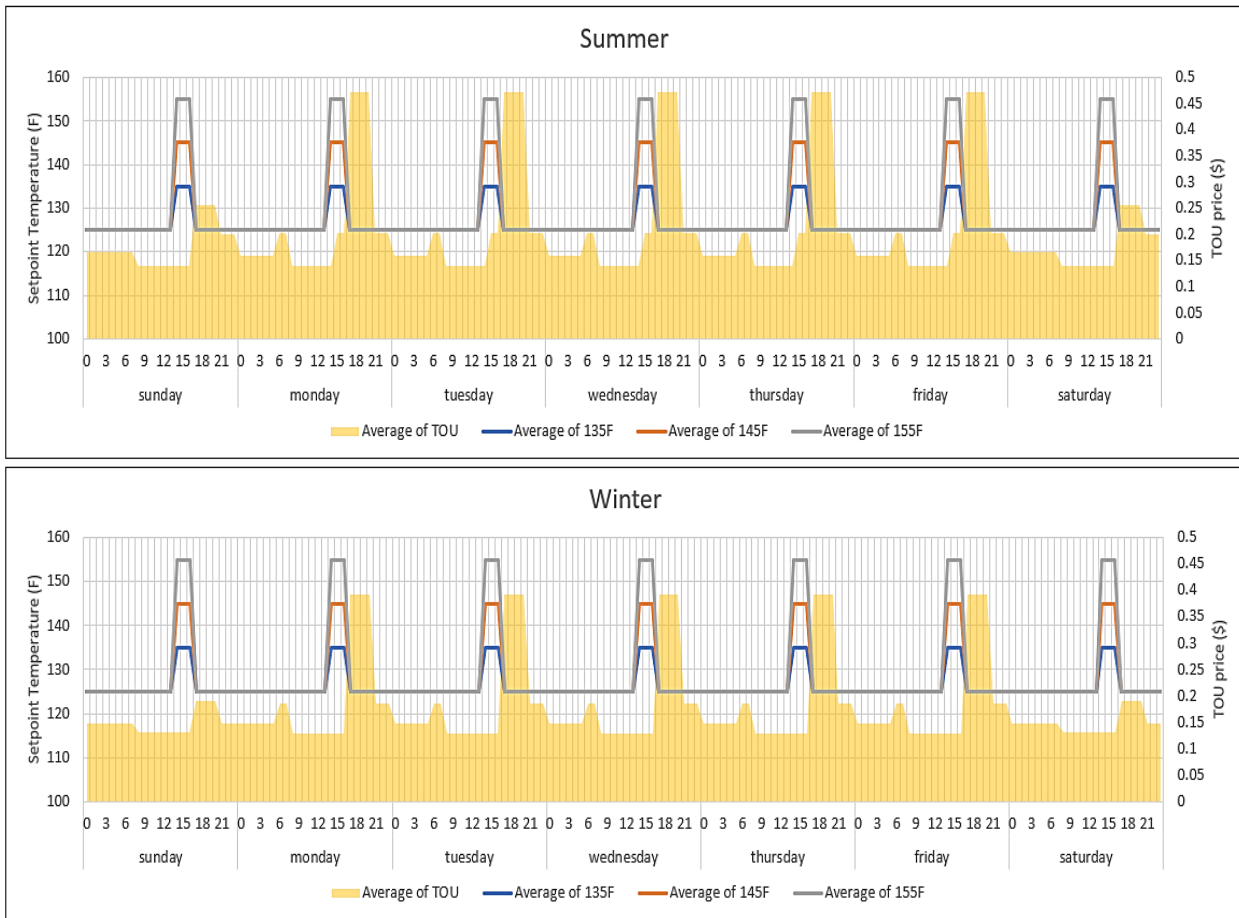


Figure 3. Load-up/shed strategy

Through analysis of the initial results, it was determined that additional complexity could be added to this strategy that would better optimize the cost without requiring additional control functions. The improvements to be made to the algorithm included: more sophisticated version of load-up/shed periods in both the evening and the morning, minimization of charge on the shoulder periods, and more efficient recovery periods to avoid the use of resistive elements as much as possible.

A *shed* period was introduced to reduce the amount of time that the water heater runs during periods of high prices, both during the peak and the shoulders. The idea again is to coast through the high-price periods without needing the water heater to run by loading up (i.e. raising the setpoint) before the peak and shedding (i.e. dropping the setpoint) during the peak. As mentioned the TOU price schedule differs both seasonally and between weekdays and weekends. For this reason, unlike in the previous Load-up strategy, it was necessary to build some conditional logic into the algorithm. A separate setpoint schedule was necessary for each of the following: Winter weekdays, Winter weekends, Summer weekdays, Summer weekends. Additionally, the algorithm needed to be able to transition easily between these different day types.

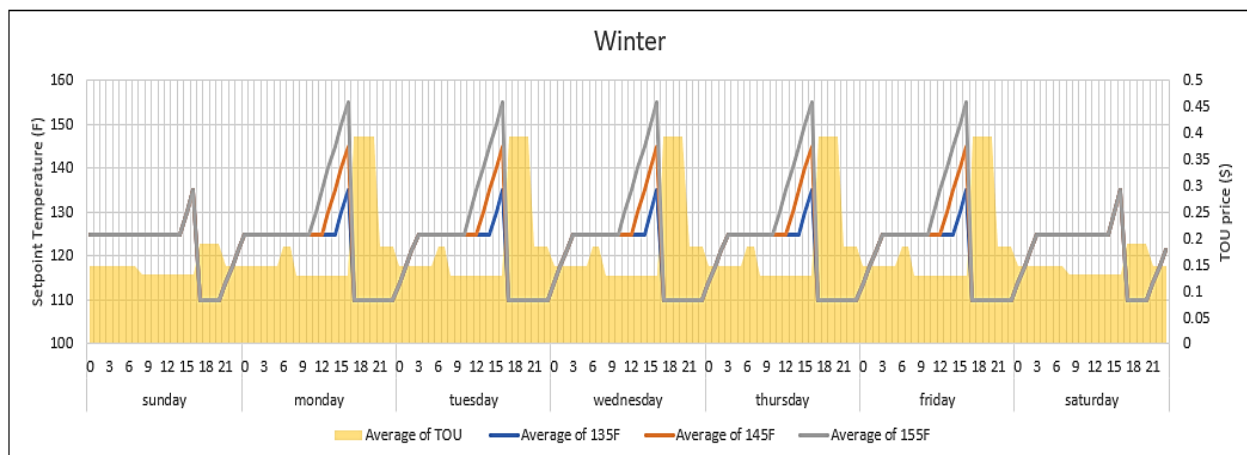
It was quickly determined through simulation that the water heater could not coast through the entire 9-hour afternoon/evening peak and shoulder period that occurs on summer weekdays, regardless of the energy level in the tank prior to the shed period. These attempts led to frequent runout events due to the relatively high water draws that tend to occur within this same window. It was feasible, however, to coast

through the peak and post-peak shoulder periods after loading up to a higher setpoint. This 7-hour period occurs on all day types except for the Winter weekend.

Simply jumping from the default setpoint of 125°F to the higher setpoint for the load-up period may trigger undesirable electric resistance usage, so a progressive ramp up to the target setpoint was incorporated. The higher the target setpoint, the earlier this ramp-up started to ensure that the setpoint could be reached using the compressor only. To deal with the 2-hour pre-peak shoulder in the summer, the ramp-up was set to reach the maximum setpoint prior to the start of the afternoon shoulder; the setpoint would then stay at this maximum setpoint until the peak price period began and the setpoint was lowered. The idea behind this strategy was to minimize the run time during the pre-peak shoulder while ensuring that the tank was fully charged before the 7-hour peak/post-peak shoulder period began.

On weekends, the afternoon price peak is significantly cheaper and the hot water draws are less concentrated around these times. For this reason, the setpoint is only ever raised to 135°F prior to the peak, which only requires a ramp-up period of 1 hour. A similar concept was used to handle the 2-hour morning peak that occurs on all weekdays. The setpoint was raised progressively from the 110°F shed temperature (that was set during the previous evening's peak) to the 135°F load-up setpoint at a rate that the compressor could handle without triggering the electric resistance element. After 1 hour at the 135°F load-up temperature, a 2-hour shed period was implemented during the morning peak hours. This morning shed period called for a longer ramp leading into the afternoon load-up since the temperature now needed to be raised from the 110°F rather than the default 125°F. This approach was originally applied to both Winter and Summer weekdays, however, it was determined through analysis of simulation results that the morning load-up in the winter months was leading to higher levels of resistance heating due to cold inlet water and cold ambient air temperatures. For this reason, the morning load-up strategy was only applied in the Summer months.

Figure 4 provides a visual representation of how the algorithm functions.



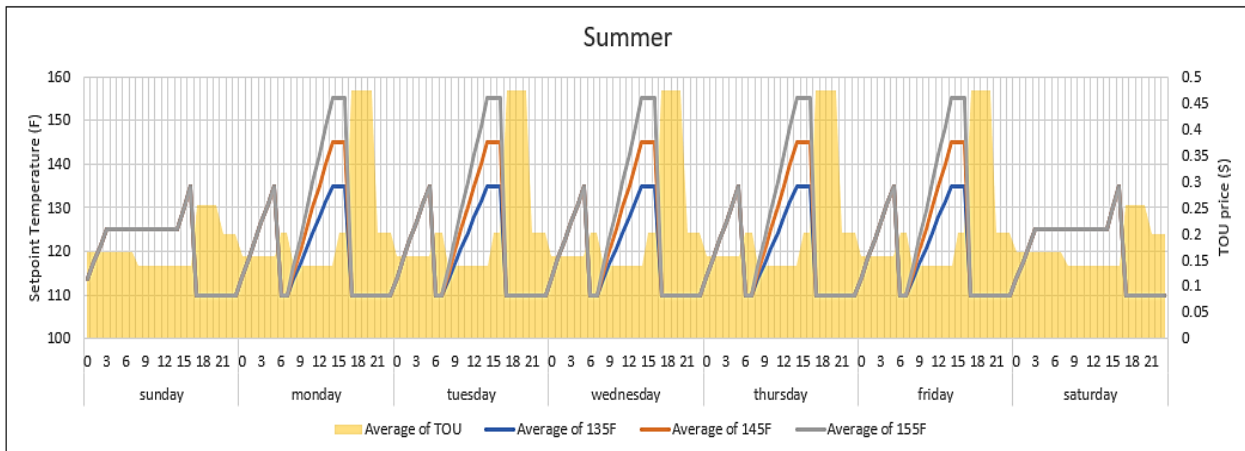


Figure 4. Load-up/Shed strategy.

Smartest: Optimal Price

The concept behind this control strategy is similar to that of the load-up approach with the goal being to make hot water at times of low price, accomplished by raising the setpoint, and to avoid energy consumption at times of high price, accomplished by lowering the setpoint. This approach, however, is more sophisticated. The algorithm looks ahead at the price signal, received via remote communication from the utility or third-party aggregator, and optimizes for lowest cost while minimizing runouts. While this strategy requires many new components and control functions, it can respond to any price signal, however dynamic. The results from this strategy most closely reflect the theoretically optimal solution: least cost and fewest runouts. An in-depth analysis of how this algorithm works is warranted.

The general mathematical problem is to define a mapping from continuous, real-valued prices to a fixed interval of real-valued setpoints in such a way as to achieve optimal cost. For example, the TDV&NEM2 prices range from roughly 15 to 3,000 kBtu/kWh; the TOU prices from 0.13 to 0.47 \$/kWh; the Utility Marginal Cost prices from -0.07 to 5.01 \$/kWh; and a range of plausible setpoints for the task range from 100°F to 160°F. This implies the need for a mapping that transforms essentially arbitrary continuous values into the fixed range [100, 160]. There are many mathematical ways to define such a mapping, and the first attempt was by way of a logistic function. Examining the output, though, indicated that the preference of the logistic function to force values away from the middle and to the extremes led to, absent intervention to the contrary, many hours at maximum setpoint and many hours at minimum setpoint.

This laid plain the tradeoff required for this task with heat pump water heaters. Loading up the tank comes at its own price. Standby losses increase but more significantly, the heat pump COP declines with tank temperature. Loading the tank up to 160°F may represent a loss of a full COP point compared to the standard operating regime at 125°F. Furthermore, a quirk of loading the tank to higher temperatures is that it may result in additional energy being delivered, and hence an even higher load than the base case. For example, machine draws (e.g. washing machine, dishwasher) will consume a fixed volume of water. With a hotter tank equipped with a mixing valve, this fixed volume of water would always be delivered at 125°F and never below that setpoint. With a traditional tank, floating within a temperature deadband, that fixed volume draw may average 119-121°F outlet temperature. Human draws which mix hot and cold at the point of use (e.g. showers), would likely consume an equivalent amount of energy as they would in the base case. Because of these factors, loading up the tank imposes its own penalty on kWh, and so an optimal solution must weigh the advantage of hitting the price schedule most favorably against the

disadvantage of loading up the tank.

The logistic function's preference for minimum and maximum setpoint temperatures tended to err on the side of too much load shifting, where the disadvantages in COP, standby, and delivered energy frequently outweighed the advantages of favorable pricing. Efforts were made to ameliorate this problem, but as it was endemic to the nature of the mapping a second approach was proposed. As an alternative mapping between prices and setpoint, we consider a partitioned Gaussian distribution, with variance scaled by a “load shifting severity” factor.

The current version of the Optimal Price algorithm works as follows. Once per day, the subsequent 24 hours of the price schedule is considered. The dynamic range is calculated by taking the ratio of the maximum and minimum price. The “load shifting severity” is calculated as a piecewise linear function of that value, where if it is below a certain limit then no load shifting will be performed, otherwise it is a linear function of the dynamic range. Then, the 24 prices are transformed into “z-scores” by subtracting the mean and dividing by the standard deviation, then scaled by the load shifting severity and translated by a setpoint bias factor. Finally, those resultant values are binned by the following thresholds: [less than -1.5, -1.5 to -0.5, -0.5 to 0.5, 0.5 to 1.5, greater than 1.5] which maps to five different setpoints. For example, the five different setpoints could be 100, 112.5, 125, 137.5, and 150 for a 150°F max temperature scenario. Although this may sound complicated, the result is that the specifics of the algorithm are controlled by only five parameters that can be adapted for various price schedules: the threshold to invoke load shifting, the slope and intercept of the load shifting severity function, the setpoint bias factor, and the maximum setpoint.

Figure 6 graphically shows the relationship between the z-score of the price at a given hour and the setpoint. In cases where there is little variation in the price over the next 24 hours, the load shifting severity will be low (blue lines in Figure 5). Those low values will encourage the water heater to remain at the default setpoint. For example, if the TDV ranges from 23 to 25 over the next time period, it is not worth the energy penalty to load up the water heater to a 150°F setpoint. A small fraction of price could be gained by running the HPWH at TDV of 23 compared to 25 but the energy losses of the high temperature would outweigh the benefits. Likewise, a proportionally incremental change in setpoint (to, say, 127°F) in the low TDV time has such an infinitesimal impact on stored energy availability, it's not worth it. In the opposite case, were the TDV is forecast to range from 20 to 600 over the next 24 hours, the load shifting severity will be high (green lines in Figure 5). Those high values will encourage the water heater to drastically alter its setpoint over the time period to only run at the cheap times and avoid the peaks.

If the ability exists to load up the tank every day, it can also be beneficial to operate the water heater at a generally lower set point at all other time periods. For example, it can be useful to bias the tank temperature low overnight when little water is used knowing that it's possible to load up in the middle of the day. This is the setpoint bias factor which shifts the curves in Figure 5 right or left.

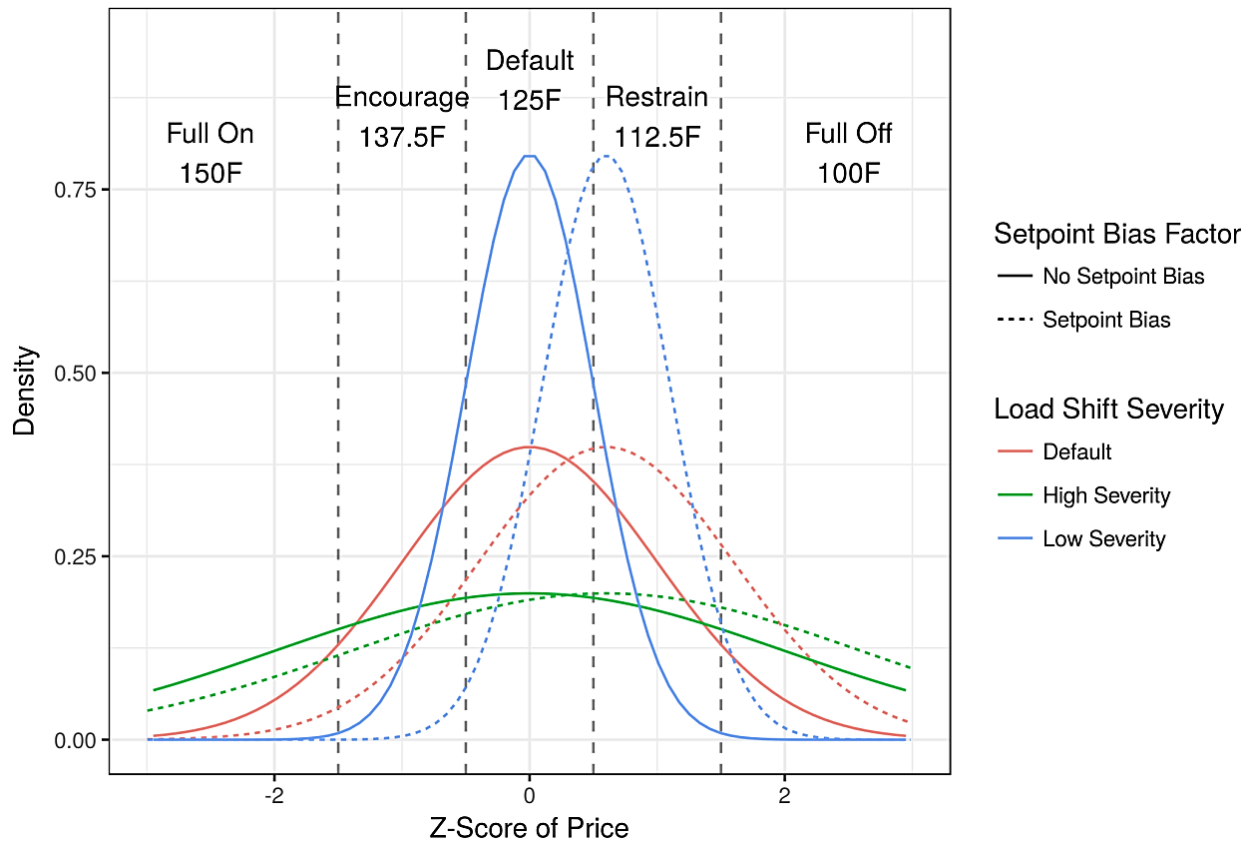


Figure 5. Partitioned Gaussian Setpoint Shifting Concept

Figure 6 provides an example of the setpoint response to the TDV&NEM2 schedule. It shows six days starting July 8th. For the first two days, the TDV is relatively flat and it is not worthwhile to shift load so the setpoint remains constant. For the next four days, there are spikes in the TDV which, at values over 2,000 are literally off the chart. In response to those spikes, on the afternoon of day three, the setpoint is dropped down to 100°F ensuring the HPWH will not run. Then, as soon as the price spike is over, the setpoint returns to normal. In day four, the algorithm sees that the price spike is tall and wide so it runs the water heater at a high setpoint of 160°F during the cheap hours. Days five and six continue in much the same way as day three. The takeaway is that the algorithm can handle varying circumstances.

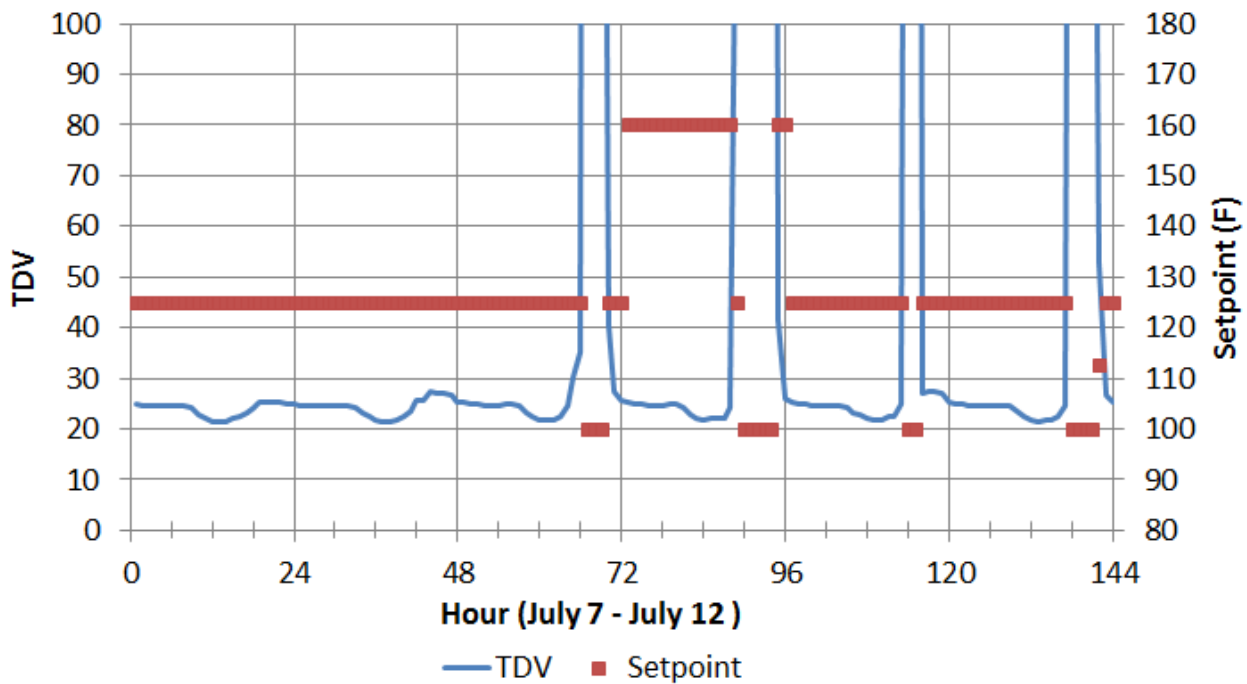


Figure 6. Setpoint Response to TDV Over Six Days

In examining all the output, it appears as though the load shift severity correction, as a function of the dynamic range of prices, is most useful only for the highly skewed TDV&NEM2 prices. Sufficient solutions were found with the TOU and Utility Marginal Costs schedules in which the scaling factor was a constant unrelated to the magnitude of the daily price swing. In the case of TDV&NEM2, the unadjusted z-scores would be sufficient to hit the “full off” condition during the peak, but not to enforce full on in anticipation of the peak. Similarly, the setpoint bias factor was found useful mainly for the TOU prices. Sufficient solutions for the Utility Marginal Costs and TDV&NEM2 were found that did not shift the Gaussian as shown on the plot.

Although the results indicated significant cost savings, it seemed that there was still room for improvement. Analysis showed that a significant amount of the electric resistance energy was being consumed when the ambient temperature was below the cut off temperature for the compressor to run during the colder winter months in certain climates. Often, this energy was being consumed in the early morning hours when temperatures were coldest. For this reason, we attempted to target these times of high resistance heating by loading-up using the compressor in the late-night hours before temperatures dropped. Without being able to incorporate weather forecasting, a general approach had to be applied to all winter days. This meant that we were unnecessarily loading-up on some days when the temperatures never fell below the cut-off temperature. The additional exercise showed that there is still room for improvement in minimizing electric resistance usage and overall power consumption, however, it would require more advanced knowledge of temperature forecasts to see real improvements to the optimized algorithm.

Bi-Directional: State-of-Charge

All the previous strategies involved sending a signal to the water heater from the controller telling it how to operate for the subsequent hour. The final strategy that was investigated requires two-way communication between the water heater tank and the controller, allowing the controller to know the

quantity of hot water available in the tank at each hour. The concept being that the controller can assess whether it has enough hot water to meet an expected demand for a fixed number of hours and, if not, determine the cheapest times to add energy to the tank while avoiding runout events.

As mentioned, this algorithm operates using an expected draw schedule. The expected draw is either the 50th, 75th, or 90th percentile, calculated over all the draw patterns used in CBECC-Res spanning all house sizes and climate zones.

It also uses three setpoints, which are referred to as minimum, default, and catch-up. The minimum setpoint is the minimum allowable temperature for water delivered to the fixtures, to be used whenever the tank contains enough hot water to meet the expected demand. The catch-up setpoint is used when expected demand is larger than the tank's current charge plus its future capacity at the default setpoint. The setpoints serve as a crude stand-in for "do nothing", "add to state of charge with compressor", and "add to state of charge by any means necessary", respectively. In this simulation we used a minimum setpoint of 110°F, a default setpoint of 125°F, and a catch-up setpoint of 140°F.

Variables used in the algorithm are defined as follows:

Capacity: We calculate the capacity of the heat pump as the number of liters of water at the minimum deliverable temperature that can be produced over some period of time by running the heat pump given the current inlet water temperature and tank temperature near the coil. This relationship is estimated using lab testing data.

Lookahead period: The lookahead period refers to the number of hours of future expected draws that are considered when making a decision about setpoint.

State-of-charge: The state-of-charge is defined as the volume of water known to be above the minimum deliverable water temperature (defined as 110°F for this simulation). In the "ideal" case, we calculate this using all six temperature nodes in the HPWHsim model. Therefore, our state-of-charge has a resolution of 1/5 the volume of the tank (there are nodes at the very top and bottom of the tank). In the "realistic" case, we assume there are only two thermocouples in the tank, located at nodes 1 and 4 (bottom and 3/5 from the bottom). In this case, we only know if the tank is completely uncharged, less than 3/5 full, or at least 3/5 full. Note that the state-of-charge is assumed to only be reduced by draw (that is, we do not account for heat loss). This results in a slightly optimistic prediction of whether we can meet demand, however we will generally make a pessimistic prediction about what that demand will be.

The algorithm checks at each timepoint whether the state-of-charge will meet demand without adding additional energy. If not, it decides whether it can find cheaper times to run before demand overtakes capacity. It also decides whether an unexpected draw has forced it to trigger the catch-up setpoint, which will presumably trigger the electric resistance element, but may help prevent a runout.

At the beginning of every hour, the simulation calculates the following:

- Current state of charge tank in liters (either "ideal" or "realistic")
- Expected demand over the lookahead period (4, 8, or 12 hours) in liters

If the state-of-charge exceeds the expected demand, we can meet demand without additional hot water and can therefore save energy by lowering to the minimum setpoint. The next step checks for a catch-up

condition (i.e. we expect to run out of hot water in the next hour) and, if present, increase the setpoint to the maximum to immediately add energy to the tank.

If the expected demand exceeds the state-of-charge, but there is no threat of run out in the following hour, a decision must be made to minimize energy consumption while avoiding the potential run out. We now look at the following: number of hours until the tank will run out of hot water, the expected demand until runout, and the expected hourly capacity to produce hot water. One of three scenarios will then play out depending on the current conditions:

- Running compressor at all hours until runout that are cheaper than the present will meet or exceed demand
 - We don't need to add charge now and can drop the setpoint to the minimum
- Running compressor at all hours until runout will meet or exceed demand
 - We need to start adding charge now and therefore choose the default setpoint
- Running compressor at all hours until runout will not produce enough water to meet demand
 - We need to charge the tank by any means necessary to avoid runout and therefore must increase to the maximum setpoint

This method utilizes knowledge of the energy stored in the tank to inform decisions of when to run and when not to run. The concept should save energy by not running unnecessarily, allowing the tank temperature to drift down when no large draws are expected in the near-term. It should also help avoid run-out events by adding energy to the tank when needed.

In practice, the simulation runs that utilized this strategy led to higher costs and more runouts than both the Optimal Price and Load-up/Shed strategies. This can mostly be attributed to the divergence of the draw patterns from the average draw profile that was used as a basis of decision making in the algorithm. This draw profile is the average across all house sizes and climate zones. As discussed previously, the simulations were run using the 50th, 75th, and 90th percentiles of this profile.

In an extreme case, the 50th percentile expected load is much too low for a 5-bedroom house. This leads to more runouts than in the base case because the tank isn't expecting a big enough need for hot water. Conversely, in the 1-bedroom case, the 50th percentile expected load is too large, so the tank is kept unnecessarily warm. If a large water draw was anticipated and the price was relatively inexpensive, the tank was heated. If this hot water draw never materialized, then the tank was heated unnecessarily.

In testing, it was also found that the “realistic” cases generally performed worse than the “ideal” cases. In the former, there is simply not enough information about the state-of-charge inside the tank. Most integrated HPWHs have two temperature measurements; one nearly the bottom of the tank and one near the top. Therefore, because the tanks are always stratified, the tank knows three states. The thermocline is: below the bottom temperature sensor, between the two sensors, or above the top sensor. There is a large volume between the two sensors where most of the important action occurs. Not knowing if you have 15 gallons or 40 gallons of hot water is a big problem for predicting what to do. The upshot is that, in order to get the state-of-charge algorithm to work in the physical world, you need more temperature sensors (or a flow meter on the tank). A total of 6 sensors were used in the ideal case for this simulation.

Results

Cost savings are highly dependent on climate zone, house size, and water heater tank size; however, averaged results across these variables are useful to assess the potential of load shifting using residential HPWHs. Figure 7 shows the average customer cost savings as a box-and-whisker plot displaying the range of possible outcomes for each of the control strategies.

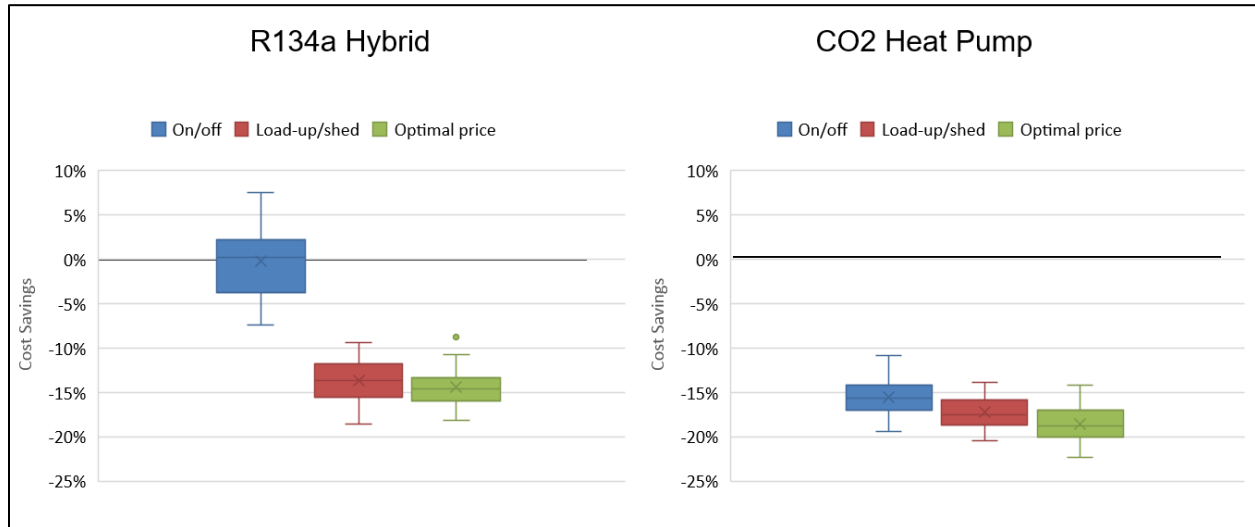


Figure 7. Simulation results for various control strategies

Focusing first on the R-134a integrated hybrid HPWHs, it is clear the on/off (i.e. allow/block) strategy yields limited savings and, in fact, leads to increased energy consumption on occasion. Both the Load-up/Shed and the “Advanced” algorithm (a.k.a. Optimal Price) yield significant cost reductions of roughly 10-20%. The CO₂-based HPWHs do an even better job of saving energy which can be attributed to several factors that will be discussed subsequently. It is important to note that these cost savings do not reflect every possible combination of climate zone, tank size, and house size as some combinations fell outside of the acceptable tolerance (<0.3%) of hot water delivered below 105°F. These scenarios would not be implemented in the real world and therefore should not be included in the results (e.g. a 50-gallon tank would not provide enough hot water for a 5-bedroom house in CZ16).

Figure 8 shows the cost and energy savings for the Load-up Shed strategy as a function of the maximum setpoint (see Figure 4) for all applicable integrated HPWHs.

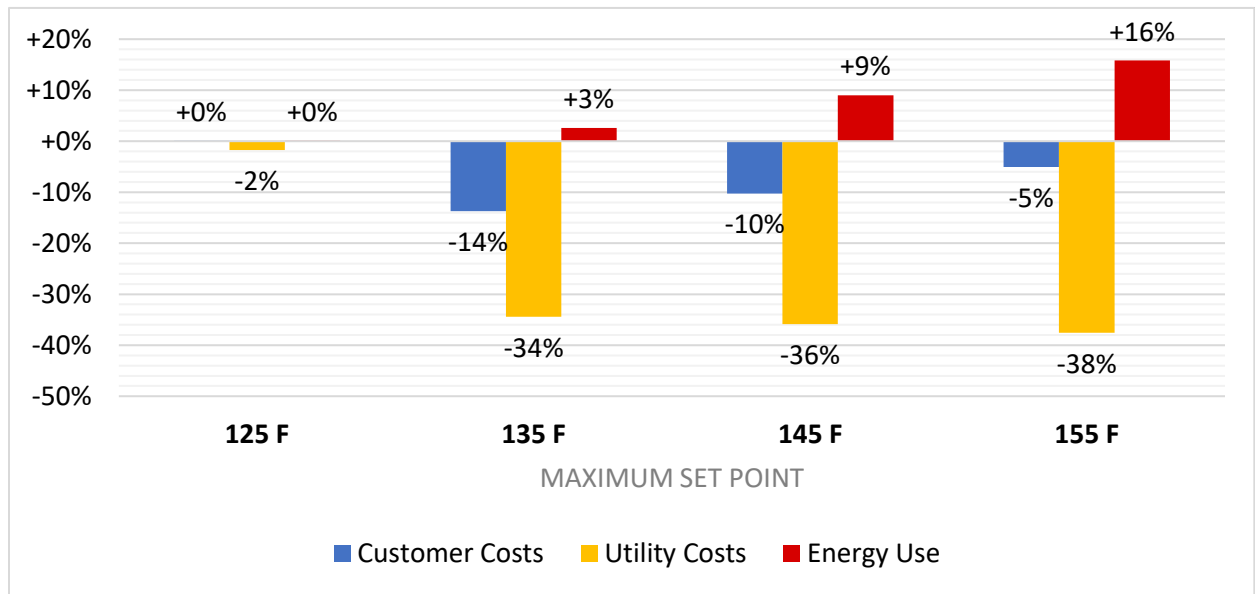


Figure 8. Cost and energy savings by setpoint (for Load-up/Shed strategy, TOU price signal)

Shifting from the uncontrolled baseline scenario at a constant setpoint of 125°F to a controlled load-shifting strategy with a maximum temperature of 135°F yields significant cost savings to both the customer and the utility at 15% and 34%, respectively. These savings come at the cost of a slight increase in energy consumed, however, this energy consumption has shifted to more favorable times of the day from a utility perspective. The customer also saves on their electricity bill by avoiding the peak prices dictated by the TOU schedule.

As you increase the maximum setpoint, the utility sees marginal increases in cost savings while the customer savings diminish at a quicker rate. Additionally, the energy consumption increases significantly to raise the water temperatures to the higher setpoints and keep them there. The takeaway from Figure 8 is that, according to our findings, 135-140°F is the optimal setpoint temperature for cost and efficiency.

Figure 9 shows solar and grid peak coincidence for managed and unmanaged hybrid HPWH. It turns out that even without load management, HPWHs are already a good, natural remedy for California's duck curve. With load management, peak coincidence is virtually eliminated, and solar off-peak is increased from 50% to 70%.

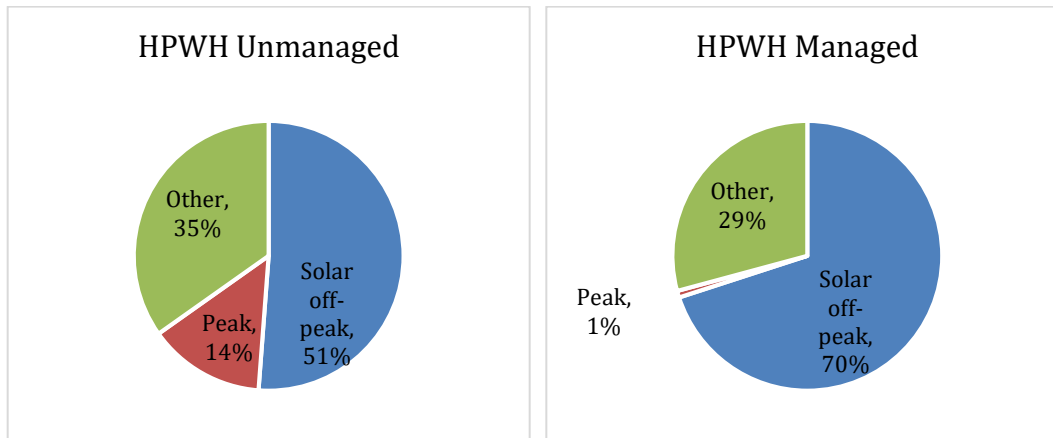


Figure 9: Off-Peak Solar and Peak Coincidence. Off-Peak Solar is defined as 8am-3pm, Peak is defined as 5pm-9pm.

One additional finding is that the operational savings to each stakeholder (i.e. the utility and the customer) is highly dependent on what the control strategy optimizes for. The Optimal Price control strategy, as its name suggests, optimizes for one of the price schedules, each reflecting the costs to either the consumer (TOU) or the utility (UMC). While the algorithm is effective at lowering costs of that particular price schedule, it is important to look at the overall effect. Table 3 summarizes the operational cost savings to each stakeholder group when optimizing for a specific price schedule.

Table 3. Operational cost savings by stakeholder group using the Optimal Price control strategy

	Customer bill savings	Utility marginal cost savings
Optimizing for customer costs (TOU)	-15% to -20%	-35%
Optimizing for grid marginal costs (UMC)	0% to +5%	-60%

Optimizing for TOU rates leads to similar savings as the Load-up/Shed strategy achieved. The utility and the customer achieve significant savings, motivating all parties to participate in such a measure. When optimizing for the UMC price schedule, utility cost savings of roughly 60% are feasible, however, this may lead to a slight increase in customer costs. This strategy may be appealing to the utility, but it would most likely require another mechanism to compensate the customers such as a free or discounted water heater or an annual cash payment.

It is worth taking a closer look at a specific scenario to help quantify the actual costs and energy consumption that could be expected for an actual house. Table 4 displays the results of the Load-up/Shed strategy for a 3-bedroom house in CZ12. The ERWH is 50 gallons and has a maximum setpoint temperature of 155°F in the *managed* scenario. The HPWH is 65 gallons and has a maximum setpoint of 135°F in the managed scenario. The ERWH has a smaller tank due to the higher heating capacity and the maximum setpoint is higher to minimize on-peak operation.

Table 4. Detailed Results Analysis

	ERWH Unmanaged	ERWH Managed	HPWH Unmanaged	HPWH Managed*
Peak coincidence (5pm-9pm)	20%	0%	14%	1%
Off-peak solar coincidence (8am-3pm)	40%	60%	50%	70%
Effective storage capacity / peak event	-	1.3-1.8 kWh	-	0.5-0.6 kWh
Energy use (kWh/y)	2,570	2,640 (+3%)	1,070 (-58%)	1,090 (-58%/+2%)
Resistive kWh	100%	100%	16%	14%
Consumer bills	\$500	\$380 (-25%)	\$190 (-60%)	\$160 (-70%/-15%)
Utility marginal costs	\$180	\$80 (-55%)	\$60 (-70%)	\$40 (-80%/-35%)
CO ₂ e (kg)	706	660 (-7%)	281 (-60%)	279 (-61%/-1%)

*HPWH Managed savings are presented as (% change from ERWH Unmanaged / % change from HPWH Unmanaged)

In this scenario, the consumer would save \$30 in annual water heating energy costs while the utility saved \$20 relative to an uncontrolled HPWH. The coincidence of energy consumption with the solar production window increases by 10% and the on-peak consumption is essentially eliminated.

The low CO₂e reductions are an artifact of the emissions factors used to calculate these emissions. The study used emissions factors from the California Public Utilities Commission's (CPUC) 2017 Avoided Cost Model. These emissions factors may be appropriate for valuing CO₂ reductions from energy efficiency measures, but the low peak/off-peak differentiation does not value emissions reductions from load shifting significantly. This seems due to two methodological assumptions: 1) The use of short-run (dispatch) vs. long-run (build) marginal emissions; 2) The use of California's renewables portfolio standard policy as both a floor and a ceiling. These methodological assumptions need to be reviewed to more appropriately value the emissions reduction potential of demand flexibility.

Sensitivity of Customer Savings

Sensitivity analysis was conducted by climate zone, water draws/occupancy, and HPWH model (hybrid technology only), with a base case of a 65-gallon HPWH in a 3-bedroom household, in climate zone 10 (Riverside) which represents the median savings case. For water draws, the number of bedrooms (1 to 5) was used as a proxy for different draw levels. The analysis showed least sensitivity to HPWH models with a range of 14% to 16%, medium sensitivity to climate zone with a range of 13% (climate zone 13 - Fresno) to 17% (climate zone 3 – Oakland), and the highest sensitivity to occupancy with a range of 11% (5 bedrooms) to 17% (4 bedrooms).

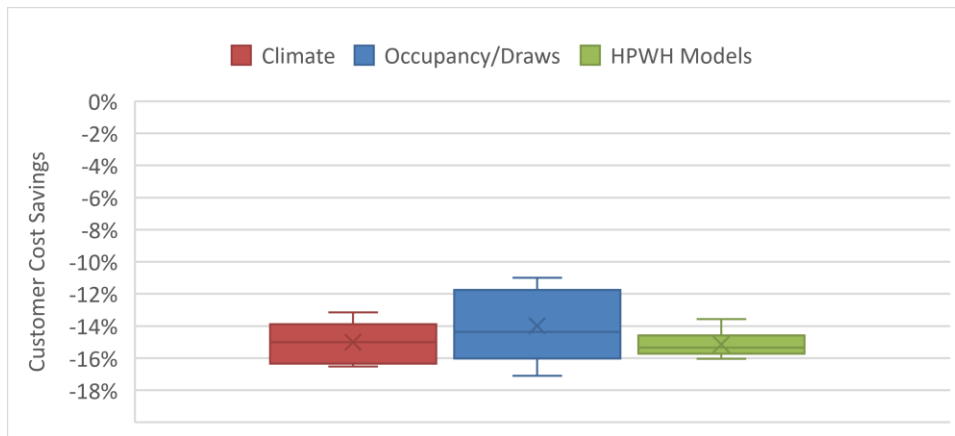


Figure 10: Customer savings sensitivity analysis

We did not analyze sensitivity by electricity price because savings obviously depend directly on the price schedule, particularly the differential between peak and off-peak prices. The price schedule developed for this study is reflective of utility marginal costs and therefore a reasonable modeling assumption, but rate design will be critical to incentivize customers to participate in load shifting programs.

CO₂ HPWH Discussion

As a refrigerant, CO₂ uses a transcritical cycle rather than the traditional vapor-compression refrigeration cycle. In operation, these HPWHs don't need to stay below the critical point and can therefore make much hotter water (130-176°F).

Although the cost savings seen in Figure 7 are higher for each control strategy, the most notable improvement is the allow/block strategy. Since CO₂ HPWHs produce 150°F water (in this study), they have a significantly higher energy content than that of a tank of 125°F water of the same volume (it should be noted that a mixing valve is assumed installed in all the elevated set point temperature cases). This higher energy content allows for the HPWH to coast through the peak price times without running and not risk hot water runout. There are also no resistance elements present meaning that all recoveries will be achieved by the compressor and there is no risk of unintentionally triggering electric resistance heating as a result of the internal logic as in the hybrid units.

Considering the Load-up/Shed and Optimal Price strategies, the dynamic setpoint strategy does not apply for CO₂ HPWHs. Instead, a control signal can be sent (as in the allow/block scenario) to provide an additional *engage* signal. This can be used to force the HPWH on at times when the internal logic may

normally be satisfied with the tank conditions. While the more complex control strategies do lead to additional savings relative to the allow/block strategy, it is not nearly as significant as a jump in savings.

Using the Allow/Block strategy, average consumer cost savings of 16% and average utility cost savings of 24% were observed with the CO₂-based units. When using the Load-up/Shed algorithm, average consumer cost savings of 17% and average utility cost savings of 65% were achieved. Lastly, the Optimal Price algorithm yielded average consumer cost savings of 19% and average utility cost savings of 31% when optimizing for TOU rates and average consumer cost savings of 17% and average utility cost savings of 73% when optimizing for the UMC price schedule. Interestingly, the results were even better when optimizing for the TDV&NEM2 price schedule, with average consumer cost savings of 20% and average utility cost savings of 75%. These results are presented in Table 5.

Table 5. CO₂-based HPWH Results

Control Strategy	Optimized for:	Average customer bill savings	Average utility savings
Optimal Price	UMC	17%	73%
	TOU	19%	31%
	TDV&NEM2	20%	75%
Load-up/Shed	TOU	17%	65%
On/Off	TOU	16%	24%

Conclusions

The work to-date on the NRDC-Ecotope HPWH Load Flexibility study has shown the potential of shifting load away from peak price periods through changes to residential HPWH operation. A series of price schedules, representing the costs as seen by the various stakeholders, were transformed into water heater control signals using several different control algorithms. The goal was to find the lowest cost of energy without running out of hot water. The HPWHsim engine was used to simulate the performance of the water heater under the various control schemes. The simulated performance was verified through laboratory testing conducted at the PG&E Applied Technology Services testing facility.

The simplest control strategy (Allow/Block) operated by simply disabling the HPWH during the 4-hour peak price window dictated by a time-of-use (TOU) schedule. In the case of the R-134a integrated HPWH units, this led to a significant amount of hot water runout events and minimal (sometimes negative) cost savings. The peak price window coincides with relatively high hot water consumption which contributed to the increase in runouts. Additionally, the higher setpoint during all off-peak hours led to significant efficiency losses. The stored hot water in the tank would often be depleted to such a degree after block periods that, when power was restored, the electric resistance element would be triggered due to the HPWH's internal logic. The cost savings of this strategy were significantly higher in the CO₂ HPWH units, as they can store hotter water and do not need resistance elements thanks to their ability to induce large temperature lifts. Average consumer cost savings of 16% and average utility cost savings of 24% were observed with the CO₂-based units.

The more sophisticated Load-up/Shed strategy attempted to resolve the issues of the Allow/Block strategy without adding too much additional complexity. Since runouts during the "Block" periods were an issue in the previous strategy, a "Load-up" period was added prior to the peak price times to add energy to the tanks. To accomplish this, a setpoint schedule was created that progressively transitioned between high and low setpoints to trick the HPWH to run at less expensive times and not run at peak times. Three

different maximum setpoints were investigated in an attempt to identify the point at which raising the setpoint becomes counterproductive.

Under the Load-up/Shed control strategy, customer energy costs were reduced by roughly 10-20% for appropriately sized integrated HPWHs with a maximum setpoint temperature of 135°F. Above this setpoint, marginal reductions in utility costs were observed, but at the expense of significantly lower customer savings and higher energy consumption. While reducing customer costs was the primary goal in this scenario (due to the TOU rate schedule that it was based off), the utility also saw a reduction in costs of approximately 34-35%. This makes sense since the TOU rate schedule is designed to help avoid producing power at peak times and to utilize times when generation is less expensive. Improvements from the Allow/Block strategy were also seen in the CO₂ HPWH units, showing average consumer cost savings of 17% and average utility cost savings of 65%.

The Optimal Price control strategy, as its name suggests, optimizes for one of the price schedules, each reflecting the costs to either the consumer (TOU) or the utility (utility marginal costs). While the previous two strategies can only be applied to a fixed TOU-style rate schedule, this algorithm works with any price schedule, regardless how dynamic it is. Utility marginal costs (UMC) factors in the costs of energy, emissions, capacity, transmission and distribution of electricity and are quite sensitive to these parameters, leading to a fairly dynamic schedule. The Optimal Price algorithm looks ahead at the upcoming UMC prices and determines a setpoint based on the “load shifting severity” (or spread) of the given prices. If the upcoming prices are relatively similar, the water heater is encouraged to remain at the default setpoint since there is little cost benefit to load shifting and potentially heavy energy impacts. Conversely, when a significant price spike is forecasted, load shifting becomes more desirable and the water heater is encouraged to vary its setpoint to avoid peak times and take advantage of inexpensive times.

Optimizing for customer savings (TOU) leads to similar savings as the Load-up/Shed strategy achieved. The utility and the customer achieve significant savings, motivating all parties to participate in such a measure. When optimizing for the UMC price schedule, utility cost savings of approximately 60% are feasible, however, this may lead to a slight increase in customer costs. This strategy may be appealing to the utility, but it would most likely require another mechanism to compensate the customers such as a free or discounted water heater or an annual cash payment. Again, the CO₂ HPWH units achieved greater savings with the more advanced control strategy, achieving average consumer cost savings of 19% and average utility cost savings of 31% when optimizing for TOU rates. Optimizing for the UMC price schedule yields average consumer cost savings of 17% and average utility cost savings of 73%.

A fourth control strategy deemed the State-of-Charge algorithm was investigated. The idea of this strategy was to utilize knowledge of the energy stored in the tank and a predicted water draw profile to inform decisions of when to run and when not to run. Theoretically, this concept would save energy by not running the HPWH unnecessarily, allowing the tank temperature to drift down when no large draws were expected. It would also help avoid run-out events by adding energy to the tank when needed. In practice, however, the concept was limited by the fact that real draw patterns are unpredictable and often stray from the average profile. A predictive algorithm that learns the occupants’ behavior may help overcome this issue and increase the load shifting potential.

This study was not intended to find the absolute maximum benefits to load shifting HPWHs, but to demonstrate that the benefit potential is significant. We trust that manufacturers and load management aggregators will create even more effective algorithms to maximize customer and societal benefits. To

this end, the study was successful in demonstrating the savings potential of load shifting to both the utility and the customer.

References

- California Energy Commission. (2017). *Time Dependent Valuation of Energy for Developing Building Efficiency Standards*. Retrieved from http://www.energy.ca.gov/title24/2013standards/prerulemaking/documents/general_cec_documents/Title24_2013_TDV_Methodology_Report_23Feb2011.pdf
- Ecotope Inc. (2015). *Heat Pump Water Heater Model Validation*. Retrieved from <https://neea.org/resources/heat-pump-water-heater-model-validation-study>
- Kvaltine, N., & Larson, B. (2015a). *Laboratory Assessment of A.O. Smith HPTU Series Heat Pump Water Heaters*. Retrieved from http://neea.org/docs/default-source/lab-test-reports/hpwh-lab-report_ao-smith_hptu_12-09-2015.pdf?sfvrsn=2
- Kvaltine, N., & Larson, B. (2015b). *Laboratory Assessment of Stiebel-Eltron Accelera 220 E Heat Pump Water Heater*. Retrieved from http://neea.org/docs/default-source/lab-test-reports/hpwh-lab-report_stiebel_220e_07-15-2015.pdf?sfvrsn=4
- Larson, B., & Kvaltine, N. (2015). *Laboratory Assessment of GE GEH50DFEJSRA Heat Pump Water Heater*. Retrieved from http://neea.org/docs/default-source/lab-test-reports/hpwh-lab-report_ge_h50dfejsra_04-09-2015.pdf?sfvrsn=2
- Regional Technical Forum. (2018). Residential Heat Pump Water Heater UES Update. Retrieved from <https://nwcouncil.app.box.com/v/20180130ResHPWHPres>