

Digital Dispatch and Demand Response during Grid Emergencies: Evidence from Household Cooling in California's Flex Alerts*

Soren Anderson[†] Dylan Brewer[‡] Maghfira Ramadhani[§]

December 9, 2025

Abstract

We study the interaction between moral suasion and automation in managing resource scarcity. Using data from smart thermostats during a California heatwave, we exploit a natural experiment involving voluntary conservation requests (Flex Alerts) and a statewide emergency phone alert. We document three findings. First, standard moral suasion suffers from rapid habituation. Second, high-salience signals (emergency alerts) reverse this habituation. Third, and most importantly, we identify a novel complementarity between salience and automation. High-salience alerts reduce the rate at which users override automated thermostat adjustments, tripling the efficacy of demand response technology. These results suggest that 'human frictions' limit the scalability of smart technologies, but crisis salience can mitigate these frictions.

JEL Codes: D12, D60, C93, L94, L98, Q41, Q48

Keywords: behavioral economics, nudges, welfare, energy use, extreme events, natural field experiments

*We gratefully acknowledge the Ecobee and its customers for supporting this research. We thank Bobby Harris, Matt Oliver, Laura Taylor, Casey Wichman, and participants at the CU Environmental and Resource Economics Workshop, Southeastern Workshop on Energy & Environmental Economics & Policy, Empirical Methods in Energy Economics Workshop, Midwest Economics Association (MEA) Annual Conference, AERE Summer Conference, and Southern Economics Association (SEA) Annual Meeting for valuable feedback.

[†]Michigan State University. Email: sta@msu.edu

[‡]School of Economics, Georgia Institute of Technology. Email: brewer@gatech.edu

[§]School of Economics, Georgia Institute of Technology. Email: maghfira.ramadhani@gatech.edu

1 Introduction

When climate-driven extreme events create resource scarcity, policymakers face the dual challenge of allocating the resource efficiently while minimizing welfare losses. Such challenges arise in many settings, such as during water scarcity (Wichman et al., 2016; Mahadevan and Shenoy, 2023), floods (Taylor and Druckenmiller, 2022), wildfires (Baylis and Boomhower, 2021), and hurricanes (Deryugina, 2017; Dinerstein et al., 2025; Strömberg, 2007). Electricity markets face an analogous challenge as climate change brings more extreme heat. Extreme heat increases both the frequency and intensity of peak electricity demand (Auffhammer et al., 2017), primarily as people adapt by staying indoors (Zivin and Neidell, 2014) and increasing air conditioning use (Davis and Gertler, 2015; Barreca et al., 2016). When supply margins are tight, this creates a scarcity where policymakers must curb electricity consumption strategically to prevent costly blackouts.¹

Economists have extensively studied price-based demand response, including dynamic pricing (Blonz, 2022; Fu et al., 2024; Burkhardt et al., 2023; Harding and Sexton, 2017; Ito et al., 2018), automation (Blonz et al., 2025; Bailey et al., 2025; Bollinger and Hartmann, 2020), and combinations of pricing with information provision (Prest, 2020; Jessoe and Rapson, 2014) to ensure conservation when electricity is scarce. However, during high-stakes emergencies, it is not politically feasible for utilities to impose an immediate and massive price change to reflect the actual scarcity cost. Instead, utilities resort to voluntary appeals for conservation with mixed evidence on their effectiveness (Brewer and Crozier, 2025; He and Tanaka, 2023; Holladay et al., 2015).² These appeals represent a distinct, non-pecuniary instrument, relying on the salience of the message and behavioral mechanisms such as warm glow (Andreoni, 1989), social pressure (DellaVigna et al., 2012), or moral payoff of contributing to public goods (Levitt and List, 2007; Ferraro and Price, 2013; Allcott and Kessler, 2019) to induce immediate demand reduction. Yet, while Bailey et al. (2025) and Blonz et al. (2025) find that automation bypasses human inattention, it is still unclear how this automation performs alongside moral suasion during an actual emergency.

We study California’s Flex Alert program, a state-wide emergency energy conservation campaign triggered when the California Independent System Operator (CAISO) forecasts critical grid conditions. Flex Alerts are disseminated through social media, utility communication channels, and via private email and text for customers who signed up to receive notifications. The Flex Alerts encourage customers to save electricity, and provide specific guidance to increase thermo-

¹ Blackouts can be deadly to individuals with certain health conditions, as Barreca et al. (2016) documents that diffusion of residential air conditioning significantly reduces temperature-related mortality in the US. In other settings, He and Tanaka (2023) documents that the energy conservation campaign in Japan following a nuclear plant accident caused an increase in temperature-related mortality.

² As additional examples, in June 2021, New York City sent an emergency alert requesting customers to conserve energy to prevent outages during intense heat. See <https://www.nytimes.com/2021/06/30/nyregion/nyc-energy-alert-heatwave.html> (last accessed June 15, 2024). Texas customers also received conservation requests from the Electric Reliability Council of Texas (ERCOT) during the extreme heat waves that struck in July 2022. See <https://www.reuters.com/business/energy/texas-grid-operator-asks-users-conserve-energy-amid-scorching-heat-2022-07-11/> (last accessed June 15, 2024). In a winter setting, Brewer and Crozier (2025) studied an emergency request caused by a supply-side energy emergency caused by a fire incident at a natural gas plant in Michigan, which coincides with extreme demand for heating.

stat settings to above 78°F (25.5°C) during the peak period, and precool to 70°F (21.1°C) beforehand. Compliance is voluntary and encouraged via appeals to prosocial preferences.

Our empirical strategy exploits a natural field experiment created by a ten-day heatwave that led to ten consecutive Flex Alerts from August 31 to September 9, 2022. CAISO issued the first Flex Alert on August 31st at 12:48 p.m. followed by a State of Emergency declaration from the governor at 3:15 p.m.³ CAISO then announced more Flex Alerts repeatedly each day. As system stress intensified, the governor issued another conservation request on September 6, followed by a statewide digital emergency alert sent by the California Office of Emergency Services (CalOES) to 27 million cell phones—an unusually salient intervention.⁴ During these ten days, in parallel, CAISO called several demand response events within the peak period for customers enrolled in automated demand response programs. This sequence of events provides rich variation in both policy instruments across households (voluntary conservation vs. automated demand response) and salience over time (low-salience standard Flex Alerts vs. the high-salience statewide phone alert).

To measure household responses, we use smart thermostat data from Ecobee’s Donate-Your-Data program, focusing on cooling setpoints and compressor run-time as our measure of cooling behavior and energy use. The data allow us to separately identify households participating in automated demand response, whose thermostats can be adjusted automatically. We compare California households to non-demand response households in neighboring states using a difference-in-differences design. During Flex Alert hours, non-demand response households received only the moral suasion message, while demand response households also experienced automated setpoint overrides and monetary incentives when demand response events were called. Our findings highlight the importance of moral suasion in crises and reveal strong interactions between behavioral responses and automated technology.

First, we find that salience plays a central role in voluntary conservation. Standard Flex Alerts produce modest but fading adjustments in cooling behavior. During the first six days of the Flex Alert series, households increase cooling setpoints by 0.04°F on average, with a peak increase of 0.2°F at 7 p.m. This peak response reduces compressor run-time by at most 2 minutes per hour. Our event study reveals that even these small effects weaken across days. By the third day, households exhibit clear habituation to repeated low salience requests, which is consistent with Ito et al. (2018). The statewide phone alert, which provided an abrupt and highly salient communication of the importance of the energy emergency, generated a sharp and immediate increase in conservation that reversed habituation. After the alert, cooling setpoints increased by up to 0.4°F on average, and compressor run time declined by up to 3 minutes per hour, with reductions beginning earlier at 5 p.m. Elevated salience is therefore a powerful amplifier of moral suasion (Ferraro and Price, 2013). After the emergency ends, households do not fully return to

³See California governor’s proclamation of a state of emergency <https://www.gov.ca.gov/wp-content/uploads/2022/08/8.31.22-Heat-Proclamation.pdf>.

⁴See the Executive Order N-15-22 <https://www.gov.ca.gov/wp-content/uploads/2022/09/9.6.22-Labor-Day-Heat-Event-EO.pdf> (last accessed June 15th, 2024).

baseline behavior, resembling a hysteresis pattern (Costa and Gerard, 2021).

Second, automated demand response (DR) consistently outperforms voluntary conservation. Across both low salience and high salience periods, households enrolled in automated DR exhibit substantially larger reductions in cooling demand than households that receive only voluntary conservation requests. During low salience periods, DR participants increase their cooling setpoints by 0.4°F in response to DR events with peak reductions above 0.5°F in some hours, which is significantly larger than the response to voluntary appeals. Automated overrides and monetary incentives ensure meaningful reductions even when voluntary conservation is weakened by habituation or low perceived urgency. These results complement earlier evidence from California’s 2000-01 energy crisis (Reiss and White, 2008) and the later Flex Alert program (Peplinski and Sanders, 2023) showing that public appeals can reduce consumption at the aggregate level, but our microdata allow us to observe the behavioral mechanisms that underlie these changes.

Third, salience and automation interact in a complementary way. During standard Flex Alerts, many DR participants override automated setpoint controls, which reduces the effectiveness of DR events. When the statewide phone alert increased the salience of the grid emergency, override behavior declined sharply, and DR participants became much more responsive. Their behavioral response to DR events increased by more than a factor of three, resulting in reductions in 1.1°F average reductions in thermostat setpoints and reduced the compressor run times by as much as 7 minutes per hour at peak. Households appear more willing to retain automated defaults when blackout risk becomes salient. Prior work examining smart thermostats has found that people often override energy-efficiency settings, eliminating potential energy savings (Brandon et al., 2022). The pattern we document reveals a novel mechanism: salience reinforces automation, which reduces behavioral crowd-out and significantly increases the effectiveness of automated DR in emergencies.

We estimate that the 2022 Flex Alerts reduced peak electricity demand by 800–1,300 MW and generated \$69.8 million in net welfare gains, primarily by avoiding costly blackouts that would have affected millions of California customers. To arrive at these welfare estimates, we extend the framework of Ito et al. (2018) following Brewer (2022) and Brewer and Crozier (2025). In our model, households choose a baseline cooling setpoint by equating the marginal benefit of additional cooling with the marginal cost, given retail electricity prices. Excessive cooling during emergencies generates welfare losses when the social marginal cost of electricity exceeds the private marginal cost reflected in retail prices. Moral suasion helps correct this distortion by imposing a psychological cost that encourages higher setpoints, while automated DR raises the effective marginal cost of cooling and automatically overrides the baseline setpoint for those who would otherwise be non-compliers. We translate changes in cooling setpoints to reductions in air conditioning compressor run-time, then scale these household-level reductions to aggregate electricity demand following the approach in Blonz et al. (2025).

Our welfare calculations reveal three important patterns. First, the demand reductions are substantial, equivalent to the output of a medium-sized power plant. Second, despite their high

per-household effectiveness, automated DR programs contributed less than 10 percent of total reductions because enrollment rates remain low. This implies very high returns to expanding DR enrollment: even modest increases in participation could generate millions in additional welfare gains. Third, the primary value of emergency conservation is grid reliability in avoiding blackouts.

This paper makes three primary contributions. First, we identify distinct behavioral dynamics: habituation, reactivation, and inattention. We build on research examining household responses to resource scarcity (Deryugina, 2017; Dinerstein et al., 2025) and public appeals for energy conservation (Reiss and White, 2008; He and Tanaka, 2023; Costa and Gerard, 2021). While these studies show appeals can reduce consumption, our high-frequency data allow us to observe mechanisms previously unseen in behavioral interventions (Ito et al., 2018; Allcott and Rogers, 2014). We document that households habituate rapidly to repeated low-salience requests but “reactivate” conservation when crisis salience increases sharply. Furthermore, we observe post-crisis inattention resembling hysteresis or default effects (Costa and Gerard, 2021; Fowlie et al., 2021). These insights deepen our understanding of how households process and respond to emergency communication under real-time scarcity.

Second, we demonstrate that crisis salience and automation operate as complements during grid emergencies. Prior research suggests that automated demand response produces larger, more reliable reductions than manual adjustments (Bollinger and Hartmann, 2020; Bailey et al., 2025; Blonz et al., 2025). The most similar paper to ours in this literature is Bailey et al. (2025), which finds that automated DR outperforms active response. However, we document imperfect compliance with automated defaults in high-stakes settings. We extend this literature by showing that high-salience moral suasion reduces the rate at which users override automated defaults, thereby tripling the effectiveness of automated demand response. We contribute to the literature on the interaction between behavior and smart technology (Brandon et al., 2022; Prest, 2020). While prior work highlights how human interference can limit the scalability of smart technologies (Brandon et al., 2022), we show that crisis salience mitigates these “human perils” by reducing override behavior, demonstrating that awareness is critical for realizing the full potential of automated DR during grid emergencies.

Third, we provide a comprehensive welfare analysis of emergency conservation in this setting. We contribute to the broader literature on communication strategies during energy emergencies (Brewer and Crozier, 2025; Holladay et al., 2015) and the welfare effects of energy conservation (Ito et al., 2018; Allcott and Kessler, 2019; Jacob et al., 2023; Bollinger and Hartmann, 2020). The most closely related work in the energy emergency literature is Brewer and Crozier (2025), which studies governor requests and phone alerts during a winter energy emergency. We advance this work by leveraging richer variation in policy instruments (voluntary appeals vs. automated DR), studying dynamic responses to an emergency spanning multiple days, and by estimating welfare impacts. Using a framework that links behavioral responses to electricity demand when the social marginal cost is high, we provide new welfare estimates for both voluntary conservation and automated DR during grid emergencies.

The rest of the paper is organized as follows. Section 2 provides background on California Flex Alerts. Section 3 describes our theoretical model and testable hypothesis. Section 4 details the data and the empirical strategy, while Section 5 presents our results. Section 6 presents our welfare analysis. Section 7 discusses policy implications. Section 8 concludes the paper.

2 Background

In this section, we provide a brief background on the Flex Alert program in California. A Flex Alert is issued by CAISO, which manages the high voltage-electricity grids that serve more than 80 percent of California. A Flex Alert calls for consumers to voluntarily conserve electricity when there is a predicted energy supply shortage, especially if CAISO needs to use reserve to maintain grid reliability. Flex Alerts are typically issued in the summer when extreme heat puts upward pressure on electricity consumption. It is typically issued from 4 p.m. to 9 p.m. when solar generation is declining and the electricity demand remains high.⁵




Figure 1. Routine Flex Alert recommendation infographics

During a Flex Alert, CAISO notifies the public of the Flex Alert via social media and the websites of CAISO and the utility. Private emails and texts are only sent to households that subscribe to receive Flex Alert notifications. However, CAISO may sometimes issue an alert with little or no advance notice.⁶ CAISO also provides suggested behaviors for conserving electricity during the Flex Alert hours. These include setting the thermostat to 78°F or higher, avoiding using major appliances, and turning off unnecessary lights. CAISO also recommends that consumers pre-cool their house to 70°F and use major appliances during the off-peak period before the Flex Alert hours. Figure 1a shows the recommended action before a Flex Alert that is typically posted around noon and Figure 1b shows the recommended action during a Flex Alert that is announced around 4 p.m. These two infographics are tweeted routinely by the Twitter account. Next, we specifically discuss the ten consecutive days of Flex Alerts that happen in 2022.

⁵ See <https://www.flexalert.org/> for more details (last accessed June 15th, 2024).

⁶ Most Flex Alerts are announced on the same day, as shown in Appendix Figure A1c.

The September 2022 Flex Alerts: During extreme heatwaves that happened from August 31st to September 9th, 2022, CAISO issued Flex Alerts for ten consecutive days. There are intervals within the peak period when CAISO also calls a demand response event for demand response participants. On the seventh day, there is also a phone alert issued by the California Office of Emergency Services (CalOES), which increases the salience of the Flex Alerts. Studying this Flex Alert series allows us to measure how households respond to Flex Alerts and demand response events with different levels of salience. It also allows us to study how households' responses evolved in repeated Flex Alerts.

The official Flex Alert Twitter account announced the first statewide Flex Alert of the series on Wednesday, August 31st, at 12:48 p.m. Later at 3:15 PM, the Governor of California declared a state of emergency due to the heatwave which is broadcasted statewide, and relayed the voluntary conservation request to the public.⁷ The Flex Alerts were then extended based on CAISO's forecast of the grid conditions on a daily basis. The average day-ahead market prices during the peak period have been over 400 \$/MWh, with higher price spikes to around 1500 \$/MWh between September 5th and 9th. The real-time market price also experienced an extremely high spike, nearly reaching 2000 \$/MW on September 5th and 6th, reflecting the opportunity cost to CAISO for procuring marginal electricity generation. Figure A3 shows the trends of high electricity prices and high risk of outage during these Flex Alerts.

The most critical grid condition occurred on Tuesday, September 6th, 2022. At 5:48 pm, CalOES sent a phone alert to all mobile phones in California requesting for energy conservation. These phone alert announcements reach almost all customers, unlike the standard announcements. The text of the phone alert read "CalOES, Conserve energy now to protect public health and safety. Extreme heat is straining the state energy grid. Power interruptions may occur unless you take action. Turn off or reduce nonessential power if health allows, now until 9 p.m." Media claimed that the phone alert successfully reduced the total electricity demand by 2,000 MW.⁸

Figure 2 shows the detailed timeline of the 2022 Flex Alert series. In summary, households in California can be categorized into demand-response households and non-demand response (voluntary) households. Throughout this paper, we refer to voluntary and non-demand response interchangeably. During Flex Alerts hours (the red-highlighted intervals in Figure 2), **voluntary households** are treated purely through the moral suasion from Flex Alerts, while **demand response households** additionally receive automation and conservation incentives when a demand response event is called in parallel with Flex Alerts (the green-highlighted intervals in Figure 2). All households in California received a phone alert on September 6th, which **increased the**

⁷ See the Proclamation of a State of Emergency <https://www.gov.ca.gov/wp-content/uploads/2022/08/8.31.22-Heat-Proclamation.pdf> (last accessed June 15th, 2024).

⁸ See <https://www.canarymedia.com/articles/grid-edge/californians-saved-the-grid-again-they-should-be-paid-more-for-it> (last accessed June 15th, 2024). It reports that 50,000 Ecobee thermostats voluntarily enrolled in several utility programs that responded to 37 demand response events across California during the emergency. It also reports smart thermostat programs/providers, other than Ecobee, that participate in demand reduction, such as Google Nest, OhmConnect, Leap, and Honeywell. Also see <https://www.honeywellhome.com/us/en/demand-response/> and <https://www.ecobee.com/en-us/citizen/top-5-reasons-to-take-advantage-of-community-energy-savings/> (last accessed June 15th, 2024) for examples of demand response programs offered by thermostat providers.




Figure 2. Timeline of 2022 Flex Alerts

Note. This figure shows the timeline of the 2022 Flex Alerts series. The blue line shows the maximum outdoor temperature across California which fluctuates from 90 to 110 degrees. The first Flex Alert is announced at noon on August 31st. The red dashed vertical lines show when the Flex Alerts notifications were released to the public. The red-highlighted intervals show when the Flex Alerts were in effect, which are typically from 4 to 9 pm, except for September 5th and 8th. The green-highlighted intervals show when a demand response event is called by CAISO, in parallel with Flex Alerts. Table A1 and A2 summarize the timing of Flex Alerts and demand response events on each date. Table A3 summarizes the information posted by the Flex Alert Twitter account and its timing in the September 2022 Flex Alerts.

salience of the energy conservation requests.

3 Conceptual Model

This section develops a theoretical model of household thermostat settings in the presence of automation and moral suasion. Our model extends the model of Ito et al. (2018) to model thermostat setpoint behavior as in Brewer (2022) and Brewer and Crozier (2025). In our model, households derive consumption utility from the thermostat setpoint, which incurs a cost of energy consumption. Households choose their preferred baseline thermostat setting ahead of time by equating the marginal benefits of an additional degree of cooling with the marginal costs due to the retail price of electricity. During an energy emergency, the social marginal cost of energy consumption exceeds the retail price of energy, resulting in welfare losses. Households can deviate from their baseline thermostat setpoint, but if they choose to do so, they incur an inertia cost of taking action, similarly to the cost of deviating from a default option as in the “Optimal Defaults” literature (Choi et al., 2003). We then model how households respond to moral suasion and an automated demand response program.

Baseline thermostat setpoint: Our model considers household i choosing the cooling setpoint T_{idh} on date d and hour h . We suppress the subscripts i , d , and h to simplify notation. In the first stage, households choose baseline thermostat settings ahead of time that they program into the thermostat. The household's baseline utility function U_0 is additively separable in the consumption benefits and the opportunity cost of cooling:

$$U_0(T, p) = u(T) - px(T). \quad (1)$$

The first component, $u(T)$, denotes utility or comfort from choosing a cooling setpoint. Since our empirical setting is the summertime, we assume locally that households weakly prefer lower temperatures ($\partial u / \partial T \leq 0$) and have a concave preference ($\partial^2 u / \partial T^2 < 0$). Households pay the retail electricity rate p in dollars per kilowatt hour for the amount of electricity consumed for cooling $x(T)$ in kilowatt hours. The energy required for cooling is decreasing in thermostat setting so that ($\partial x / \partial T < 0$) and ($\partial^2 x / \partial T^2 \leq 0$).

Households choose a baseline cooling setpoint that solves $T^0 = \arg \max \{u(T) - px(T)\}$, which is characterized by the following first-order condition:

$$\frac{\partial u(T^0)}{\partial T} - p \frac{\partial x}{\partial T} = 0. \quad (2)$$

This baseline thermostat setting T^0 is the cooling setpoint the household will have in a normal non-emergency hour, or during an energy emergency if they do not take conservation action. This choice may be explicit, such as when a household chooses a schedule that a programmable thermostat acts upon, or it may be implicit in that the household has previously chosen a setpoint that remains through inertia. In either case, the existing thermostat setpoint serves as a baseline behavior that the household will have to deviate from in the second stage if they wish to respond to any change in incentives, such as moral suasion or demand response during an energy emergency.

Second stage: In the second stage, we model how households respond to an energy emergency when there is a conservation appeal from moral suasion and an automated demand response program. During an energy emergency, the social marginal cost of electricity rises above the original retail price of electricity p that determined the household's baseline thermostat setpoint T^0 . We illustrate this case in figure 3, where we plot the household's marginal willingness to pay for indoor cooling. The socially optimal thermostat setpoint is at $T^* > T^0$ during the energy emergency, leading to a social welfare loss of area $A + B + C + D + E$ if the household maintains the baseline thermostat setpoint during the energy emergency.

When energy consumption incentives change in the second stage, a household can update its chosen thermostat setpoint from the baseline setpoint, but it will only do so when the household's cost to deviate from its baseline thermostat setting is low enough. We model the choice to deviate from default behavior so that the household reoptimizes the thermostat setting only when the

random variable $Z > 0$. Thus, in the second stage, the thermostat setting is:

$$T = \begin{cases} T' & \text{if } Z > 0 \\ T^{default} & \text{otherwise} \end{cases} \quad (3)$$

where $T^{default}$ is the default thermostat setpoint and T' is the reoptimized thermostat setpoint. The baseline thermostat setpoint then acts as a default choice if the household does not take action. For a typical household not enrolled in an automated demand response program, the default thermostat setting is the baseline thermostat setting chosen in the first stage so that $T^{default} = T^0$. In contrast, automation overrides the default thermostat setpoint so that $T^{default} = T^A$. This override can take advantage of the household's inertia to bias default behavior in favor of conservation if $T^A > T^0$. The choice to change the thermostat setting is determined by the difficulty for the household to adjust its thermostat based on both physical costs and informational barriers. For example, an individual not at home or who does not view a notification has a very high cost of adjusting a thermostat setpoint either to comply with an emergency appeal or respond to an automated thermostat override. We describe the components of the inertia variable Z in more formal detail later in this section.

If the household deviates from the default thermostat setpoint, they will reoptimize their thermostat setpoint. The household's utility in the second stage U_2 is a function of a potentially new price of energy and the impact of moral suasion:

$$U_2(T, p', s) = u(T) - p'x(T) - \mu(T, s). \quad (4)$$

This utility function differs from the first stage in two ways. First, p' is the second stage price, which may differ from the first stage price if a demand response program compensates households for reductions in consumption at a rate different from the retail rate, and second $\mu(T, s)$ is a moral payoff term that depends on the household's chosen thermostat setpoint and salience of moral suasion s . The moral payoff can be thought of as either a warm glow for conservation (Andreoni, 1989) or a moral cost for consumption (Levitt and List, 2007; Ferraro and Price, 2013) that acts effectively as a tax on consumption so that $\partial\mu/\partial T > 0$. We assume the marginal moral payoff increases in salience so that $\partial^2\mu/\partial T\partial s > 0$.

If the household deviates from the default thermostat setpoint, they will choose the setpoint that solves $T' = \arg \max_T u(T) - p'x(T) - \mu(T, s)$. The first order condition is

$$\frac{\partial u(T')}{\partial T} - p' \frac{\partial x}{\partial T} - \frac{\partial \mu}{\partial T} = 0. \quad (5)$$

For households that take action, higher emergency prices p' and stronger moral suasion $\partial\mu/\partial T$ increase conservation. Pure moral suasion relies solely on the moral cost term, while a pure demand response program relies solely on the change in price to achieve conservation.

The inertia variable Z determines which households take action and is comprised of the factors

that motivate or deter a household from adjusting the thermostat setpoint. We assume that a household is motivated to take action when the cost of inaction is large due to a high energy price, strong and salient moral suasion, or a more elastic demand for cooling, which is true when $\Delta U_2 = U_2(T', p', s) - U_2(T^{default}, p', s)$ is large. In addition, physical costs of adjusting the thermostat and informational barriers to receiving the emergency conservation request make action less likely, which we denote by ξ . These costs may include, for example, the cost of adjusting the thermostat when not at home or the barrier to adjusting the thermostat when an emergency conservation request is missed. We represent this by denoting $Z = Z(\Delta U_2, \xi, s)$ where $\partial Z / \partial \Delta U_2$ and $\partial Z / \partial s$ are positive and $\partial Z / \partial \xi$ is negative.

The characteristics of emergency conservation programs can affect the likelihood households take action by either affecting the incentives to conserve energy, improving messaging, or affecting baseline behavior via automation. For example, increasing moral suasion increases ΔU_2 , making it more likely a household will respond. Improving the reach of messaging or targeting times when households are home will increase ξ , also increasing the likelihood a household will respond. Automation can reduce the likelihood a household takes action if it changes $T^{default}$ to an automated setpoint T^A that is close enough to the household's reoptimized T' . Intuitively, the household may not feel the need to respond if it has the sense that the automated conservation program is responding. In contrast, if the automated setpoint is too aggressive and decreases comfort too much, it can encourage the household to take action to reduce conservation.

Hypotheses: We use our theoretical model to characterize the relative average treatment effects of these conservation programs, which will serve as testable hypotheses in the empirical portion of the paper. The average treatment effect of an emergency conservation program on thermostat settings is the average effect for households that take action and those that do not take action:

$$E[T - T^0] = E[T' - T^0 | Z > 0]P(Z > 0) + E[T^{default} - T^0 | Z \leq 0]P(Z \leq 0). \quad (6)$$

The first term is the difference between the chosen alternative thermostat setting and the baseline thermostat setting, multiplied by the probability the household takes action, while the second term is the difference between the default thermostat setting (T^0 or alternatively T^A for households in the automated demand response program) and the baseline thermostat setting, multiplied by the probability of accepting the default thermostat setting.

This implies that different emergency conservation programs can act on three margins. On the intensive margin, features of the program can increase the amount of conservation for people who take action $E[T' - T^0 | Z > 0]$. For example, increased moral suasion or a price increase raises the number of degrees a household will increase the thermostat setpoint for compliers. On the extensive margin, features of the program can increase or decrease the likelihood a household takes positive action $P(Z > 0)$. For example, increasing the salience of an emergency conservation request can independently increase the likelihood a household takes action. Finally, on the passive margin, automation can increase the baseline level of conservation by changing the default behav-




Figure 3. The marginal willingness to pay for cooling, with the baseline thermostat setpoint T^0 , conservation thermostat setpoint T' , and automation thermostat setpoint T^A .

ior from the baseline level, increasing $E[T^{default} - T^0 | Z \leq 0]$ for non-compliers. This channel will also affect the probability of taking action as it affects ΔU_2 .

During the California Flex Alerts, households were exposed to moral suasion of low and high salience, while another subset of households were additionally exposed to automated demand response programs that adjusted the thermostat setpoint and provided a monetary incentive to conserve. We construct two testable hypotheses from our model.

Hypothesis 1: Moral suasion and demand response increase conservation with or without automation present. Increasing the strength of moral suasion by increasing μ or implementing demand response by increasing the price p' will increase conservation on the intensive margin for compliers by increasing T' . In addition, this will increase the likelihood that households take action on the extensive margin by increasing ΔU_2 . This effect is true either with or without automation. When the default thermostat setpoint is adjusted using automation, moral suasion and demand response can result in the household contributing additional effort to conserve more than the automated level, or it can provide an incentive that prevents the household from overriding the automated default to return to the baseline thermostat setting.

Figure 3 illustrates this graphically. The household chooses baseline thermostat setting T^0 prior to the energy emergency. Moral suasion or demand response increases the cost of cooling, resulting in choosing T' if the household takes action. If the household does not take action, the household's default thermostat setting is the baseline thermostat setting or the automated thermostat setting T^A . The incentive to adjust the thermostat setpoint depends on ΔU_2 , which corresponds to triangle $D + E$ and is the household's surplus loss from choosing to not take action.

Increasing the strength of the moral suasion or demand response incentive will increase $D + E$, resulting in a higher likelihood the household exerts effort to respond to the energy emergency.

In the context of the California Flex Alerts, households experience low-salience requests to increase the thermostat setpoint, followed by high-salience requests. Hypothesis 1 suggests that when we move from low-salience to high salience moral suasion, we should see increased conservation by California households relative to control households. This salience effect should apply to non-demand-response and demand-response households. For demand response households, this hypothesis suggests that some of the increased treatment effect should come via reducing override behavior.

Hypothesis 2: Automation can increase or decrease conservation, depending on the presence of moral suasion and demand response. Automation that increases the default thermostat setting increases $E[T^{default} - T^0 | Z \leq 0]$, increasing conservation for non-compliers. At the same time, automation reduces ΔU_2 , which reduces the probability of taking action because the household knows that automation has reduced the benefits of the costly conservation action. Thus, in the presence of moral suasion or demand response, automation can crowd out additional conservation effort from would-be compliers.

In Figure 3, automation overrides the default thermostat setpoint from T^0 to T^A . If the household was not going to comply, this results in conservation. At the same time, this reduces the surplus loss created from inaction by reducing ΔU_2 from $D + E$ to D . On the margin this makes the household less likely to induce effort to change the thermostat setting to T' , crowding out conservation behavior. The relative magnitude of these effects is unclear, so automation may serve as either a substitute or complement to moral suasion and demand response.

The natural experiment we analyze allows us to resolve the theoretically ambiguous prediction within the context of the California Flex Alerts and automated thermostat demand response program. This hypothesis suggests that the average treatment effect for demand-response and non-demand-response households will differ. We will evaluate this by estimating distinct treatment effects for households participating and not participating in the automated demand response program.

4 Data and Empirical Strategy

In this section, we start by describing the data used for the analysis. We then describe our empirical strategy to estimate the effect of the Flex Alerts on the primary outcome variables. First, we estimate the treatment effect of the standard Flex Alerts and the Flex Alerts after the phone alert for different periods of the day. Second, we decompose the treatment effect for each hour of the day. We then estimate the dynamic treatment effect for each subsequent day of the Flex Alerts and after the conclusion of the Flex Alerts.

4.1 Data and Descriptive Evidence

Data source We use two primary data sources for this analysis. The first data source is information on the timing of Flex Alerts and Demand Response events from CAISO which is available in the Grid Emergencies History Report and CAISO Today’s Outlook.⁹ The Grid Emergencies History Report contains information on when the Flex alert was announced and the starting and ending hours of the Flex Alert. The CAISO Today’s Outlook contains information on the timing of the demand response event. We use this information to identify the timing of the Flex Alerts and demand response events in our data.

The second data source is the smart thermostat data from Ecobee’s Donate-Your-Data (DYD) program. The data is available at five-minute intervals, containing the household thermostat set-point (both heating and cooling), indoor temperature measurement from sensors, indoor humidity levels, number of minutes that the fan was running, movement indicators from sensors, the mode that the thermostat is on, thermostat event name, and household characteristics. A potential concern is that smart thermostat users may not be representative of the general population; however, previous research shows that these users have statistically similar characteristics, which rules out the concern of selection.¹⁰ The thermostat event name variable allows us to identify if households ever received any demand response event in their thermostat. We describe the procedure in Appendix D. We call these household demand response participants. The data also contains information on the household, which includes the number of occupants, size, age, and number of floors of the house, as well as the location of the households at the self-reported city level. We aggregate the data to hourly levels for the analysis in this paper, mainly to reduce noise and increase computation speed. The data is available from August 1, 2022, to September 25, 2022.

We complement the smart thermostat data with historical hourly weather data obtained from Visual Crossing.¹¹ The API processes multiple sources of weather data, including the National Oceanic and Atmospheric Administration (NOAA) Integrated Surface Database (ISD) and Meteorological Assimilation Data Ingest System (MADIS) database, which observes historical weather at multiple weather stations and allows us to observe hourly weather data for a specific location at a particular hour. The data contains hourly temperature, relative humidity, wind speed, precipitation, and cloud cover information. Since we rely on Ecobee’s household self-reported city data to match observations with the hourly weather data, we exclude households with unidentified or missing location data.

We limit our sample to households in California and four surrounding states: Arizona, Nevada, Oregon, and Utah. We use non-demand response households from other states as controls. In our treatment group, we have non-demand response households and demand response participant

⁹ See <https://www.caiso.com/Documents/Grid-Emergencies-History-Report-1998-Present.pdf> (last accessed June 15th 2024).

¹⁰ For example, Meier et al. (2019) find that the characteristics of households in the Ecobee sample are comparable to a sample of households in EIA 2015 Residential Energy Consumption Survey, while Brewer and Crozier (2025) also shows that early- and late-adopters of smart thermostats respond similarly to conservation requests.

¹¹ See visualcrossing.com/resources/documentation/weather-data/weather-data-documentation for more information (last accessed June 15th, 2024).




Figure 4. Map of households in the final sample

households from California. We exclude household extreme thermostat settings that set their cooling setpoint below 50 °F and above 104 °F. Our final sample consists of 3,706 control households, 5,180 California non-demand response households, and 3,274 California demand response participant households. Figure 4 shows the spatial distribution of the households in the final sample. For our main analysis, we use data from August 18th to September 9th, 2022, while for the event study, we extend our sample to include all available data back from August 1st, 2022, to September 23rd, 2022.

Outcome variables We measure household responses using two primary outcome variables. The first outcome variable is the thermostat cooling setpoint in degrees Fahrenheit which represents the household’s preferred indoor temperature. This cooling setpoint can be programmable—which the household sets to follow a schedule for each hour—or adjusted by the user in real time. The feature of the smart thermostat allows households to override the setting remotely. CAISO suggests precooling houses before the peak period and increasing their cooling setpoint during the peak period. Thus, the change in the household’s cooling setpoint after receiving the Flex Alert is a measure of the household’s response to the recommendation. The second outcome variable is the compressor run-time in minutes per hour. We use compressor run-time as a proxy for electricity consumption for cooling following the literature (Blonz et al., 2025; Fu et al., 2024). The compressor run-time is the duration the HVAC system runs to cool the house within an hour. We present our analysis using alternative outcome variables separately in Appendix E and F.




Figure 5. Hourly mean of outcome variables before and after the Flex Alert

Note. This figure shows the hourly average of the main outcome variables before and after the first Flex Alert announcements on August 31st, 2022, at 12:48 p.m. using observation ranging from August 18th to September 9th, 2022. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.

Descriptive evidence Figure 5 shows the hourly mean of the four main outcome variables for households in California and the controls before and after the treatment. Average households in California set their cooling setpoint around 77 °F while control households, on average, set their cooling setpoint around 75 °F. California households receive the Flex Alerts recommendation while the control households do not. The data show a cooling setpoint response during the peak period from California households after Flex Alerts. We also observe a change in compressor run-time pattern in the peak period after the Flex Alerts. We also observe that the hourly trends in outcomes variables between California and control households are similar before the Flex Alert events, which motivates our difference-in-differences approach.

4.2 Estimating household responses to Flex Alerts

For our empirical analysis, we employ a generalized difference-in-differences design. The intuition is to compare the outcome of California households to control households, allowing for different responses for demand response participants when a demand response event is called.

We estimate the effects of Flex Alerts by period of the day, distinguishing low-salience periods (after the initial Flex Alerts) from high-salience periods (after the statewide phone alert). Simultaneously, we identify the effect of demand response events using within-household variation by comparing the demand response participants' behavior during event hours and non-event hours within the same Flex Alert window. In doing so, we estimate the following difference-in-

differences specification.

$$\begin{aligned}
y_{idh} = & \sum_{k=0}^2 \beta_{FA,k} (D_{idh} \times \mathbb{1}[Period_{dh} = k]) + \delta_{FA} (D_{idh} \times \mathbb{1}[Period_{dh} = 1] \times \mathbb{1}[DRevent_{idh}]) \\
& + \sum_{k=0}^2 \beta_{PA,k} (P_{idh} \times \mathbb{1}[Period_{dh} = k]) + \delta_{PA} (P_{idh} \times \mathbb{1}[Period_{dh} = 1] \times \mathbb{1}[DRevent_{idh}]) \\
& + \theta X_{idh} + \alpha_{idh} + \epsilon_{idh},
\end{aligned} \tag{7}$$

where y_{idh} is the outcome variable of household i at date d and hour of day h . D_{idh} is a treatment indicator that equals one for California households at day t and hour of day h after the initial Flex Alert was issued and before receiving the phone alert. Similarly, P_{idh} is a treatment indicator that equals one for California households at hour h of the day after receiving the phone alert. We interact the treatment indicator with indicators for each of the three response periods, $\mathbb{1}(Period_{dh} = k)$, capturing separate treatment effects for the period $k = \{0, 1, 2\}$ which consecutively represents before-peak, peak, and after-peak period. We define the peak period following CAISO's definition of the Flex Alert hours for the day. Specifically for peak periods, we allow for different responses for demand response participants when they are in a demand response event. To do so, we interact the treatment indicator during the peak period with indicators for when a demand response event is called, $\mathbb{1}[DRevent_{idh}]$, which equals one for demand response participant when a demand response event is called at date d and hour of day h . The control variables, X_{idh} , include daily maximum outdoor temperature, hourly outdoor temperature, relative humidity, precipitation, wind speed, and cloud cover.

The combinations of fixed effects, α_{idh} , include hour-of-sample indicators and hour-by-day-of-week-by-household fixed effects. The hour-of-sample indicator picks up unobserved common shocks across households. The hour-by-day-of-week-by-household fixed effects absorb unobserved time-variant characteristics within the household, such as household commuting patterns, programmable thermostat settings, and their electricity pricing regime.¹² The treatment variation is assigned at the state-hour level, which necessitates two-way clustering of standard errors, ϵ_{idh} , at the state and hour-of-sample level (Abadie et al., 2023). However, since we have only five states, with a small number of clusters, the cluster-robust standard error estimates may be biased (Cameron et al., 2008). To overcome this problem, we use Driscoll and Kraay (1998) standard error that is valid for small clusters and long panels. This inference procedure accounts for unknown spatial and serial correlations in the treatment by relying on the length of the panel. We use a bandwidth of 24, which includes 23 lags, allowing the residuals to be auto-correlated within the past 24 hours.

The identifying assumption underlying the causal interpretation of the difference-in-differences

¹² Considering significant households in California are solar rooftop owners, these fixed effects can also capture the effects of having solar rooftop owners if we assume that the hourly solar production is relatively similar across weeks in our samples

approach is that the outcome variables for the treatment and control group will evolve similarly conditional on the controls and the fixed effects absent of Flex Alerts. Figure 5 provide a graphical evidence of parallel trend from the raw data. We also provide further statistical evidence of conditional parallel trends in the pre-treatment period from our event study estimates in Section 5.3. Given the set of fixed effects, the average treatment effect on the treated estimates identifies the difference in the change of the outcome variables using within-household variation before and after Flex Alert events in a given hour of the day and day of the week relative to the households in the unaffected state. For example, this specification allows us to compare the difference between changes in cooling setpoint for a household in California after receiving the Flex Alert nudges on Monday at 4 p.m. (relative to the same time window the week before) and the change in thermostat setting for a household in the control states, conditional on the controls.

In the specification, we estimate four sets of different coefficients, which allow us to distinguish heterogeneous average treatment effects on the treated by salience level of Flex Alerts and by whether or not households are treated with demand response events. Conditional on the controls and the set of fixed effects, the coefficient $\beta_{FA,k}$ captures the average treatment effect on California households of the standard Flex Alerts at period k of the day and the coefficient δ_{FA} captures the difference in the average treatment effect of the standard Flex Alerts between California demand response participants and California non-demand response households at the peak period of the day. Likewise, the coefficient $\beta_{PA,k}$ captures the average treatment effect on California households of the Flex Alerts after the phone alert at period k of the day, and the coefficient δ_{PA} captures the difference in average treatment effect of the Flex Alerts after the phone alert between California demand response participants and California non-demand response households at the peak period of the day.

We estimate equation (7) for cooling setpoint and compressor run-time. The specification allows us to identify the treatment effect for each period of the day: during the before-peak period when the recommendation is to precool, the peak period when the recommendation is to raise the cooling setpoint, and after-peak period when the recommendation is lifted. If households precool their house before the Flex Alert hours, we expect the treatment effect on the cooling setpoint to be negative before the peak period. Meanwhile, if households follow the recommendation to increase their cooling setpoint during the Flex Alert hours, we expect the treatment effect for this period to be positive. To evaluate the impacts on electricity consumption, we look at the treatment effects on compressor run-time. We expect the treatment effect to be positive before the peak period as a consequence of households following precooling recommendations. In the peak period, we expect the treatment effect on compressor run-time to be negative, reflecting energy savings.

4.3 Estimating hourly responses to Flex Alerts

To decompose household hourly response in a Flex Alert day, we estimate the effect of the Flex Alerts on our outcome variables for each hour of the day using the following specifications.

$$\begin{aligned}
y_{idh} = & \sum_{h=1}^{24} (\beta_{FA,h} D_{idh} \times \mathbb{1}[Hour_h = k]) + \sum_{h=1}^{24} \delta_{FA,h} (D_{idh} \times \mathbb{1}[Hour_h = k] \times \mathbb{1}[DRevent_{idh}]) \\
& + \sum_{h=1}^{24} (\beta_{PA,h} P_{idh} \times \mathbb{1}[Hour_h = k]) + \sum_{h=1}^{24} \delta_{PA,h} (P_{idh} \times \mathbb{1}[Hour_h = k] \times \mathbb{1}[DRevent_{idh}]) \\
& + \theta X_{idh} + \alpha_{idh} + \epsilon_{idh}.
\end{aligned} \tag{8}$$

The treatment indicator D_{idh} and P_{idh} and demand response event indicators $\mathbb{1}[DRevent_{idh}]$ are defined similarly to those in equation (7). We interact the treatment indicator with indicators for each hour of the day, $\mathbb{1}[Hour_h = k]$, capturing separate treatment effects for each hour. We use a similar set of fixed effects and inference procedure as in equation (7). The coefficient $\beta_{FA,h}$ measures the average hourly effect for California households at hour h in a standard Flex Alert day relative to the control group. Similarly, the coefficient $\beta_{PA,h}$ measures the average hourly effect for California households at hour h in Flex Alert days after the phone alert relative to the control group. The coefficient $\delta_{FA,h}$ measures the difference in the average hourly effect at hour h between California demand response participants and California non-demand response households in a standard Flex Alert day, while $\delta_{PA,h}$ measures the same effect but for the Flex Alert days after the phone alert. Note that $\delta_{FA,h}$ and $\delta_{PA,h}$ are identified using within household variation in demand response event timing in the peak period for demand response household, thus, the coefficients are only identified for the hours in the peak period.

The coefficient estimates of equation (8) characterize households' hourly response to the Flex Alert. We expect households to set their cooling setpoint lower before the peak period and increase their cooling setpoint after 4 p.m. when the Flex Alert hours start. We expect that the salience of Flex Alerts affects the magnitude of the responses and households' responsiveness in following the recommendation. After the largest treatment effect is reached, we expect the effects to diminish towards the end of the Flex Alert hours.

4.4 Estimating dynamic treatment effect

We estimate the dynamic treatment effect on the main outcome variables for each day using the following generalized event study specification:

$$y_{idh} = \sum_{t \in [-23, 23], t \neq -1} \sum_{k=0}^2 \beta_{tk} D_{itk} + \sum_{t \in DRdays} \delta_{tk} D_{it1} + \theta X_{idh} + \alpha_{idh} + \epsilon_{idh}. \tag{9}$$

where t is the lead or lag day relative to the first day of the Flex Alerts series (i.e. August 31st, 2022), indicator D_{itk} equals one for the period k of day t relative to the first day of the Flex Alerts, and $DRdays$ is the set of days when there is a demand response event called during the peak period (i.e. $DRdays = \{-15, 0, 1, 5, 6, 7, 8\}$). To validate our parallel trend assumption, we extend the event study sample from a month before to two weeks after the Flex Alert series. To identify the model, we bin all periods in the first week of August to the 23rd lead period. Even though there was a Flex Alert on August 17th, 2022, extending the start of the sample beyond the start of our difference-in-differences sample allows us to look at longer periods of pretrends.¹³ The event study includes 22 days of lead coefficients and 24 days of lag coefficients. The lead coefficients allow us to check for parallel trends in the outcome variables by comparing the evolution of the outcome variable separately for the before-peak, peak, and after-peak periods relative to the day before the Flex Alert. On the other hand, the lag coefficient gives us the dynamic treatment effects for each period of the day throughout and after the ten consecutive days of Flex Alerts. We use a similar set of fixed effects as in equation (7). To reduce computational run-time, we estimate equation (9) separately for each of the three periods of the day and use two-way cluster robust inference at the state and hour-of-sample level.

The coefficient estimates of equation (9) characterize how households respond differently on each day. Our causal interpretation relies on the conditional parallel trend assumption, thus we expect the lead coefficient estimates for the period where there are no Flex Alerts to be statistically zero. Looking at the dynamic treatment effect, in repeated treatment, the household may habituate in their response to the Flex Alert requests.¹⁴ We hypothesize that household responses in the first few days of the Flex Alert are small due to the low salience of standard Flex Alerts. We expect a more prominent treatment effect after the customer receives the phone alert from the CalOES. However, due to potential habituation, we expect the treatment effect to decline in the days following the phone alert. We also expect that the treatment effect will be higher during the peak period compared to the before-peak and after-peak periods. In the longer run, repeated Flex Alerts could lead to households changing their cooling setpoint, which could be a sign of habit formation. If that is the case, we expect the treatment effect to stay significant even after the Flex Alerts series ends.

5 Results

5.1 Household responses to Flex Alerts

Table 1 shows the effect of Flex Alert on each outcome variable for the before-peak, peak, and after-peak periods. The first four rows show the treatment effect of low salience Flex Alerts,

¹³ Using only difference-in-differences sample period limits our ability to only identify seven lead coefficient since we have to accumulate the first week of data. Using only difference-in-differences sample period yields quantitatively similar results as shown in Figure A4.

¹⁴ We direct readers to Ito et al. (2018) for a brief review about habituation, dishabituation, and habit formation in behavioral interventions.

and the next four rows show the treatment effect of high salience Flex Alerts. The third row, Peak $\times$ 1 (DR Event), shows the treatment effect for demand response participants for being in a demand response event during the peak period.

Table 1. Households Responses to Flex Alert

	(1) Cooling Setpoint	(2) Compressor Run-Time
After First Tweet		
Before-Peak	-0.086*** (0.021)	0.210 (0.224)
Peak	0.035 (0.029)	-0.266 (0.641)
Peak $\times$ 1 (DR Event)	0.358*** (0.120)	-0.182 (0.661)
After-Peak	-0.085*** (0.024)	0.736 (0.501)
After Phone Alert		
Before-Peak	0.001 (0.025)	-0.808 (0.627)
Peak	0.304*** (0.077)	-0.780 (0.658)
Peak $\times$ 1 (DR Event)	1.069*** (0.128)	-1.991** (0.854)
After-Peak	0.043 (0.035)	2.238** (0.993)
Pre-treatment Mean		
Before-Peak	76.35	8.41
Peak	76.27	20.31
After-Peak	76.01	12.94
No. of Household	11,807	12,135
Observations	6,342,016	6,632,642

Note. This table reports regression coefficients from difference-in-differences regression estimated using equation (7). The first panel show the effect of low salience Flex Alerts, while the second panel show the effect of high salience Flex Alerts. Each row shows the effect at different period of the day. The peak period is defined as following CAISO's Flex Alert hours of the day. Column (1) shows the effect on the household cooling setpoint in °F. Column (2) shows the effects on the compressor in minutes per hour. The control variables include daily maximum temperature, hourly outdoor temperature, outdoor relative humidity, precipitation, wind speed, and cloud cover. The fixed effects include hour-of-sample indicators and hour-by-day-of-week-by-household fixed effects. The sample period is from August 18th to September 9th, 2022. Standard errors reported follow the [Driscoll and Kraay \(1998\)](#) inference.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column 1 shows the treatment effect for cooling setpoint. While standard Flex Alerts elicit negligible response, the high salience alert triggers a behavioral shift. The combination of high salience and automated DR generates a treatment effect roughly three times larger than automa-

tion under low salience. We find no economically meaningful effect of standard Flex Alerts on cooling setpoints in the first six days in all periods of the day. When a demand response event is called, the demand response participant increases their setpoint by 0.35 °F. After the phone alert on the seventh day, we find that households increase their cooling setpoint by 0.3 °F. High salience triples the effect of demand response events, leading to a 1.1 °F increase in setpoints.

Columns 2 report the treatment effect estimates for compressor run-time. The treatment effect estimates on compressor run-time during the peak period in both low and salience Flex Alerts have the expected negative sign but statistically insignificant. The hourly average treatment effect estimates presented in the next section reveal that while the average treatment effect is not statistically significant, this is due to the underlying heterogeneity resulting in high variance in the treatment effect over time. This is also true for the effect of demand response event under low salience. After the phone alert, demand response participant reduces their compressor run-time by 2 minutes per hour more when a demand response event is called in the peak period.

5.2 Hourly responses to Flex Alerts

Figure 6 shows the average treatment effects for each hour of the day. Figure 6a reports the hourly treatment effects on the cooling setpoint. We find that in a standard Flex Alert, households respond highest around 6 to 7 p.m., increasing their thermostat by 0.2 °F higher. Before and after the peak period, the treatment effect is negative and statistically significant, but the effect size is modest. After the phone alert, they respond much earlier, increasing their thermostat by 0.4 °F higher. We find no significant effect on the cooling setpoint before and after the peak period.




Figure 6. Hourly treatment effect estimates of the Flex Alerts

Note. This figure shows the estimates of the hourly responses to Flex Alert estimated using equation (8). The hour labels for each coefficient refer to the ending time of each one-hour interval. The dots correspond to the hourly treatment effect estimates for non-demand response households for the standard Flex Alerts ($\hat{\beta}_{FA,h}$) and the Flex Alerts after receiving the phone alert ($\hat{\beta}_{PA,h}$). The highlighted area shows the 95% confidence interval, which follows the Driscoll and Kraay (1998) inference. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.

Figure 6b shows the hourly treatment effects on compressor run-time. In the standard Flex

Alerts, the treatment effect is zero after midnight and goes down to about -1 minutes per hour at 7 a.m. The treatment effect then ramps up to about 2 minutes per hour from 1 p.m. to 3 p.m. During the peak period, the treatment effect dips to the lowest point of -2 minutes per hour between 6 p.m. to 7 p.m. Towards the end of the peak period, the treatment effect goes to zero at the end of the peak period. After the phone alert, the treatment effect starts at zero after midnight and then goes down to -3 minutes per hour at 7 a.m. It then ramps up to 2 minutes per hour from 1 p.m. to 3 p.m. but it is not statistically significant. The effect starts to decline in the peak period to about -2 minutes per hour from 5 p.m. to 7 p.m. The decline in the peak period after the phone alert is relatively faster and higher than in the standard Flex Alerts indicating a stronger impact on household electricity consumption. We also find a negative effect on the off-peak period, which we suspect is due to energy efficiency, considering houses in California have a higher building code relative to other states.

5.3 Event study estimates

Figure 7 reports the dynamic treatment effect estimates for each period of the day relative to the day before the first Flex Alert announcement. The day before our difference-in-differences sample, August 18th, there is a Flex Alert, as shown in Figure 7a. The lead coefficient estimate on the cooling setpoint for that day is positive and significant.




Figure 7. Dynamic treatment effect estimates of the Flex Alerts

Note. This figure shows the estimates of the event study regressions using equation (9). The dots correspond to the lead and lag coefficient estimates for non-demand response households for each day. The highlighted area shows the 95% confidence interval which is two-way clustered at state and hour-of-sample level. We extend the event study sample from a month before to two weeks after the Flex Alert series. We present estimates from using only the difference-in-differences sample period in Figure A4

In general, the lead coefficient estimates within the difference-in-differences sample period exhibit conditional parallel trends prior to the treatment. The lead coefficient estimates for the cooling setpoint in Figure 7a is statistically zero. The lead coefficient estimates for the compressor run-time in Figure 7b are also statistically zero, however, they are more noisy. Concisely, the event

study estimates suggest that the conditional parallel trend assumption holds for all periods of the day, which supports our causal claim.

We find suggestive evidence that households habituate to the standard Flex Alerts. Focusing on the peak period, we find the largest treatment effect on the cooling setpoint of 0.1 °F during the peak period on the first and second day. This effect then goes to zero after the third day. The no-effect results are not surprising, and there are several explanations. First, the salience of Twitter posts can be considered weak in delivering the conservation request to households. The second possible explanation is habituation which causes households to be less responsive when Flex Alert is repeated, which is documented in other energy settings (Ito et al., 2018; Allcott and Rogers, 2014). Lastly, September 3rd to 5th coincides with Labor Day weekend; households likely spent more time at home and may have been less willing to adjust cooling setpoints. The dynamic treatment effect estimates for compressor run-time are noisy and mostly insignificant, but they move in the opposite direction to the cooling setpoint estimates, as we hypothesized. The highest reduction of 5 minutes per hour of compressor run-time happens on the last day of the series.

The phone alert increases the salience of the Flex Alerts, making it more effective in getting the household to act. The day the phone alert is sent, the dynamic treatment effect goes up to around 0.4 °F during the peak period. The treatment effect declined slightly on the eighth day and then continued to increase to around 0.5 °F on the last day of the Flex Alerts series. These higher treatment effects in the peak period stay until the end of the Flex Alert series, indicating a more persistent treatment effect from the phone alert. This is suggestive evidence that households dishabituate when the phone alert is sent. In contrast to Ito et al. (2018), we find that households do not habituate to the repeated Flex Alerts after the phone alert. We suspect that the phone alert makes households internalize the value of following the Flex Alerts recommendation.

After the end of the Flex Alerts, we observe a lasting change in the household cooling setpoint, even though there is no longer a grid emergency. The treatment effect on the cooling setpoint for California households stayed around 0.4 °F higher for almost two weeks after the Flex Alerts series ended. Even though the pattern resemble hysteresis phenomenon as in Costa and Gerard (2021), event study estimates in Figure A9a indicates that indoor temperature return to baseline levels which suggest more of inattention or default effect (Fowlie et al., 2021). The treatment effect on the compressor run-time evolved around zero after the Flex Alerts series ended.

5.4 Extensive margin responses

We construct alternative outcome variables to measure the extensive margin responses of the households, that is, how the households are changing the mode of cooling in response to the Flex Alerts. The first outcome is whether or not the thermostat is on hold. The second outcome is whether or not the cooling setpoint is below 70°F, which measures compliance with the precooling recommendation. The third outcome is whether or not the cooling setpoint is above 78°F, which measures compliance with the peak period recommendation. The fourth outcome is whether or not the household turns its cooling system off. The household can turn their cooling system off by

setting the HVAC mode to heat or off. This outcome measures how likely households are to turn off their cooling system during the Flex Alerts. We present these results in Appendix F.

In standard Flex Alerts, we find no meaningful effect on the proportion of cooling systems put on hold. We also find no significant effect on the proportion of people with a cooling setpoint less than 70 °F throughout the day. During the peak hour, we find that 0.7 percent of households change their cooling to above 78 °F. We find that households turn off their cooling system less on Flex Alert days; however, during the peak period, more than one percent of households turn off their cooling system. For demand response participant, we find that they are 1.5 percent less likely to put their thermostat on hold in the peak hour.

After the phone alert, we find no economically meaningful effect on the proportion of cooling systems put on hold and on the proportion of households with cooling setpoint less than 70 °F. During the peak period, we find 3 percent of households changed their cooling setpoint above 78 °F. The effect on the proportion of households turning off the cooling system stays the same as in standard Flex Alerts. After the phone alert, demand response participant, are 5 percent less likely to put their thermostat on hold during the peak period and 10 percent more likely to change their cooling setpoint to above 78 °F. Increase salience cause demand response participant to be less likely to reverse the automated override in a demand response event. Enrolling in a demand response program makes the override a default option when there is a demand response event.

5.5 Comparing the effect of Flex Alerts and demand response event

In this section, we present the effect of being a demand response event. Demand response participants in our setting receive both moral suasion via the Flex Alerts and monetary incentives when a demand response event is called. In addition to that, the household also allows the demand response provider to adjust their thermostat cooling setpoint directly, often known as Air Conditioning (AC) load control during a demand response event ([The Brattle Group, 2024](#)). These demand response participants can override the determined cooling setpoint if households prefer comfort over monetary incentives from a demand response event. The feature of smart thermostats allows automated response to a demand response event, which lowers the effort that households use to comply with a request.

We find that a demand response event is a more effective nudge compared to the Flex Alerts. Our difference-in-differences estimates show that the treatment effect of being in a demand response event within a Flex Alerts is ten times the effect of Flex Alerts. Increased salience of conservation requests through phone alerts also increases the effect of being in a demand response event, which is three times the effect of Flex Alerts with the phone alert.

Figure 8 shows the effect of the demand response event for the demand response participant. We find that typically in the first hour of an event, demand response participants exhibit a pre-cooling behavior and start conserving after the first hour. The treatment effect on cooling setpoint is more persistent within these intervals. In the standard Flex Alerts, a demand response successfully increases the cooling setpoint by up to 0.5 °F that lasts for three hours. After the phone alert,




Figure 8. Hourly treatment effect estimates of demand response event

Note. This figure shows the effect of the demand response event for demand response household estimated using equation (8). These coefficients are only identified during demand response events. The hour labels for each coefficient refer to the ending time of each one-hour interval. The dots correspond to the hourly effect of being in a demand response event for demands response households during the standard Flex Alerts ($\hat{\delta}_{FA,h}$) and the Flex Alerts after receiving the phone alert ($\hat{\delta}_{PA,h}$). The highlighted area shows the 95% confidence interval, which follows the [Driscoll and Kraay \(1998\)](#) inference. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.


it increases the cooling setpoint by up to 1.3 °F, which declines slightly to 1 °F increase at the end of the peak period. The reduction of compressor run-time from demand response event in the standard Flex Alert is up to 3 minutes per hour during the peak period. After the phone alert, the reduction is up to 7 minutes per hour during the peak period.

On a per-household basis, the Flex Alerts, even with an increased salience, yield lower responses compared to the demand response event. However, according to EIA Form 861, in 2022, the number of residential households enrolled in a demand response program in California was 505,116 households, which is about four percent of residential customers. These low enrollments might limit the potential of utilizing only demand response households to save the grid during a Flex Alert. Although the effectiveness of the Flex Alert for non-demand response households is lower, the potential energy saving might be significant in aggregate.

6 Welfare Analysis

6.1 Reduction in electricity consumption

The main goal of the Flex Alert recommended cooling setpoint during the before-peak and peak periods is to adjust household electricity consumption for cooling, shifting the electricity consumption from the peak period to the before-peak period when scarcity of generation capacity is not an issue. In this section, we estimate the total demand reductions achieved through the Flex Alerts and demand response event.



(a) Aggregate impact of Flex Alert and DR Event

(b) Aggregate impact of DR Event only

Figure 9. Hourly aggregate demand reduction

Note. This figure shows the hourly aggregate demand reduction from cooling in the standard Flex Alerts and after the phone alert which is computed following [Blonz et al. \(2025\)](#) using previous results from estimating equation (8) for compressor run-time. Figure 9a shows the demand reduction from non-demand response households and demand response participants, while figure 9b shows only the demand reduction from demand response participants. The dots correspond to the estimated hourly average demand reduction. The highlighted area shows the 95% confidence interval of the hourly average demand reduction. The hour labels for each estimate refer to the ending time of each one-hour interval. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.

We follow [Blonz et al. \(2025\)](#) to convert the hourly average treatment effect on the compressor run-time to an average reduction in electricity consumption for a household, which is detailed in Appendix C. From our results in Section 5.2, the reduction of compressor run-time for non-demand response household after the phone alert ranges from 2 to 4 minutes per hour, which translates to 0.084 to 0.168 kW per household. For demand response household after the phone alert with the additional effect of being in demand response event, the reduction ranges from 3 to 7 minutes per hour, which is equivalent 0.126 to 0.294 kW per household.¹⁵ To scale the per-household estimates to aggregate impacts, we assume the total number of residential customers to be 13,550,586 households according to the total number of households in California in the 2022 American Community Survey 1-year Estimates. We then assume that 72 percent of households in California own an AC following 2022 Residential Energy Consumption Survey ([EIA, 2020](#)). This leaves us with approximately 9.75 million households. We assume 505,116 households, about 5% of AC owners, to be demand response participants, and the rest are non-demand response households.

Figure 9 reports the aggregate demand reduction for cooling realized by non-demand response households and demand response participants during the standard Flex Alert and phone alert treatment. In the standard Flex Alert during the peak period, the estimated aggregate demand reduction reached almost 800 MW of electricity between 6 p.m. and 7 p.m. The demand response

¹⁵ This number is slightly lower but comparable to the estimates from the SDGE ACSDA Evaluation a pilot study on SDGE customer who have smart thermostat, of which the average impact per household during September 7th, 8th, and 9th are consecutively 0.43, 0.36, and 0.25 kW per household. See the [Demand Side Analytics \(2023\)](#) (last accessed June 15th, 2024)

participant contributes up to 60 MW during these periods, a modest amount compared to the non-demand response household. After the phone alert treatment, the estimated aggregate demand reduction is significantly higher, reaching more than 1,300 MW, of which around 140 MW is contributed by demand response participants. These estimates suggest that non-demand response households contribute significantly to the aggregate demand reduction, even though their response is much smaller individually than demand response participants. The highest demand reduction happens between 5 and 6 p.m. in both treatments. Our aggregate impact estimates between 5 and 6 p.m. after the phone alert is more than half of the reported reduction of 2,000 MW from CAISO on September 6th. This suggests that household cooling behavior response plays a substantial role in reducing the total electricity demand.

6.2 Welfare simulation method

We examine the welfare effects of the Flex Alerts and demand response events that are aimed at reducing electricity consumption during grid emergencies in California. To estimate the welfare impact, we return to our model of thermostat settings developed in Section 3. Based on our model prediction, there are welfare implications of voluntary requests. When the social marginal cost of electricity is very high, electricity consumption for cooling can result in welfare losses because the retail rate is low. When households are treated via moral suasion, they incur an additional opportunity cost of cooling that can be interpreted as a moral cost, which causes the household to increase its cooling setting. If the moral suasion helps to reduce the gap between the retail and wholesale price of electricity, it is welfare improving. However, when the retail rate is already higher or close to the social marginal cost of electricity. In that case, conservation is unnecessary, and moral suasion to conserve electricity could result in a welfare loss. To operationalize this model, we need empirical estimates of private and social costs per degree Fahrenheit and assumptions on the shape of demand for cooling. We evaluate the welfare impact for each period of the day during the ten days of Flex Alerts.

For this analysis, we assume the hourly marginal cost of electricity to be the hourly average real-time wholesale price across the trading zone in CAISO. This gives a lower-bound estimate of welfare.¹⁶ We take the average values of the real-time wholesale price for each period of the day during the ten days of Flex Alerts as the social marginal cost of electricity in our welfare simulation. We also assume the baseline price of electricity to be 26 cents/kWh, the average retail electricity price in California in August and September 2022 (EIA, 2022). To capture the true scarcity value during grid stress, we follow Jacob et al. (2023) and incorporate the value of lost load into welfare gain only during the peak period when demand reduction could avoid potential blackouts. To compute welfare gain from avoided blackout, we multiply the demand reduction by the value of lost load of \$4,300/MWh for California following Brown and Muehlenbachs (2024).

¹⁶ Using wholesale electricity price could underestimate the value of demand reduction. Previous studies use capacity payment price (Boomhower and Davis, 2020), avoided cost of future power plant investment (Blonz, 2022) because in the long-run, regulators set a minimum reserve margin to reduce the risk of electricity shortage through a resource adequacy process (Joskow and Tirole, 2007).

First, we convert the price of electricity to the price of cooling following our assumption on the linear relationship between electricity consumption and compressor run-time. We do this by estimating the marginal effect of cooling setpoint on compressor run-time $\gamma = \frac{\partial \kappa}{\partial T}$ from our data in the pre-treatment period. In estimating this, we use the same set of control variables and fixed effects as in equation (7). Thus, the relationship between the social marginal cost of electricity c and retail electricity price p , and the value of lost load (VoLL) in per kWh to their cooling price equivalent in per one °F of lower cooling setpoint is given by

$$p^T = -0.0417 p \gamma, \quad (10)$$

$$c_{dk}^T = -0.0417 c_{dk} \gamma, \quad (11)$$

$$\text{VoLL}_{dk}^T = \begin{cases} -0.0417 \text{VoLL} \gamma & \text{if } k = \{\text{Peak}\}, \\ 0 & \text{if } k = \{\text{Before-Peak, After-Peak}\}. \end{cases} \quad (12)$$

Note that the baseline retail cooling price is fixed, while the social marginal cost of cooling and the value of lost load vary by date and period of the day.

Next, we consider the quasi-linear utility function in equation (4). For household i at date d and period k , we model demand for cooling setpoints using a semi-log specification: $T_{idk} = \alpha_{idk}^T + \beta_{dk}^T D_{idk} + \varepsilon_{idk}^T \ln p^T$, where $D_{idk} = 1$ if the household receives moral suasion, p^T is the cooling price, and ε_{idk}^T is the cooling price semi-elasticity. We estimate the parameters β_{dk}^T which is the effect of Flex Alert on the cooling setpoint from our event study in Section 5.3. We compute the cooling price semi-elasticity ε_{idk}^T by assuming an electricity price elasticity ε of - 0.1 following Ito et al. (2018), Ito (2014) and Wolak (2011) using the following equation

$$\varepsilon_{idk}^T = \frac{\kappa_{idk}^0}{\gamma} \varepsilon, \quad (13)$$

where the κ_{idk}^0 is the household baseline compressor run-time, i.e., the average compressor run-time for the period of the day k of at the same day of the week from the pre-treatment period.¹⁷ We assume the demand response event enters the household utility through a change in price from p^T to $p^{T,DR}$ and the effect of the demand response event on the cooling setpoint for the demand response participant is δ_{dk}^T , which we get from our event study estimates.¹⁸

¹⁷ Following our linear assumption on the relationship between electricity consumption, x , and compressor run-time, κ , we obtain

$$\varepsilon^T = \frac{\partial T}{\partial \ln p^T} = \frac{\partial T}{\partial \ln x} \frac{\partial \ln x}{\partial \ln p} \frac{\partial \ln p}{\partial \ln p^T} = \frac{x}{\frac{\partial x}{\partial p}} \varepsilon = \frac{\kappa}{\gamma} \varepsilon$$

We get our estimates of γ from estimating the marginal effect of cooling setpoint on compressor run-time from a linear specification. An alternative way of computing the cooling price semi-elasticity depends on our assumption of the form of relationship between compressor run time and cooling setpoint.

¹⁸ The marginal effect of being in a demand response event on the cooling setpoint is

$$\delta_{dk}^T = \varepsilon_{dk}^T \ln \frac{p_{dk}^{T,DR}}{p^T} \Leftrightarrow p_{dk}^{T,DR} = p^T \exp(\delta_{dk}^T / \varepsilon_{dk}^T) \Rightarrow p_{dk}^{T,DR} - p^T = p^T (\exp(\delta_{dk}^T / \varepsilon_{dk}^T) - 1).$$

We derive the welfare effect of Flex Alerts and demand response events in detail in Appendix G. Table 2 summarizes these welfare effects. As explained in Section 3, the welfare effect depends on the relative gap between social marginal cost and retail price. Change in producer surplus arises from the change in electricity consumption from receiving the Flex Alerts and being in a demand response event, which is going to cost the producer if the social marginal cost is higher than the retail price. Changes in consumer surplus arise from changes in cooling setpoint, which is going to incur household disutility from experiencing a higher cooling setpoint. The total welfare change is the sum of the change in producer surplus and the change in consumer surplus, which, on the net, could be positive or negative.

Table 2. Welfare effects of Flex Alerts and Demand Response Event

Welfare Effects	
Flex Alerts	
Change in Producer Surplus	$\Delta PS_{\text{Moral},dk} = \begin{cases} \beta_{dk}^T (c_{dk}^T - p^T) & \text{if } c_{dk}^T \geq p^T, \\ -\beta_{dk}^T (p^T - c_{dk}^T) & \text{if } c_{dk}^T < p^T. \end{cases}$
Change in Consumer Surplus	$\Delta CS_{\text{Moral},dk} = \begin{cases} -\frac{1}{2} (\beta_{dk}^T)^2 \frac{p^T}{\epsilon_{dk}^T} & \text{if } \beta_{dk}^T \geq 0, \\ \frac{1}{2} (\beta_{dk}^T)^2 \frac{p^T}{\epsilon_{dk}^T} & \text{if } \beta_{dk}^T < 0. \end{cases}$
Total Welfare Change	$\Delta W_{\text{Moral},dk} = \Delta PS_{\text{Moral},dk} + \Delta CS_{\text{Moral},dk}$
Demand Response Event	
Change in Producer Surplus	$\Delta PS_{\text{DR},dk} = \begin{cases} \delta_{dk}^T (p_{dk}^{T,DR} - p^T) & \text{if } p_{dk}^{T,DR} \geq p^T, \\ -\delta_{dk}^T (p^T - p_{dk}^{T,DR}) & \text{if } p_{dk}^{T,DR} < p^T. \end{cases}$
Change in Consumer Surplus	$\Delta CS_{\text{DR},dk} = \begin{cases} -\frac{1}{2} (\delta_{dk}^T)^2 \frac{p^T}{\epsilon_{dk}^T} & \text{if } \delta_{dk}^T \geq 0, \\ \frac{1}{2} (\delta_{dk}^T)^2 \frac{p^T}{\epsilon_{dk}^T} & \text{if } \delta_{dk}^T < 0. \end{cases}$
Total Welfare Change	$\Delta W_{\text{DR},dk} = \Delta PS_{\text{DR},dk} + \Delta CS_{\text{DR},dk}$
Value of Lost Load	
Gain from Avoided Outage in Peak Period	$\Delta \text{VoLL}_{\text{Moral},dk} = \beta_{dk}^T \text{VoLL}^T$
	$\Delta \text{VoLL}_{\text{DR},dk} = \delta_{dk}^T \text{VoLL}^T$

To summarize, we perform the following procedure to simulate the welfare effect of the Flex Alerts and demand response event:

1. We estimate $\widehat{\beta}_{dk}^T$ and $\widehat{\delta}_{dk}^T$ from our empirical analysis in Section 5.3.


2. We resample i from California households and simulate the welfare effect 1,000 times. For each iteration, we perform the following steps:
 - (a) We draw the β_{dk}^T and δ_{dk}^T from the distribution of $\widehat{\beta}_{dk}^T$ and $\widehat{\delta}_{dk}^T$ for each period of the day in the series.
 - (b) We estimate the marginal effect of cooling setpoint on compressor run-time ($\partial\kappa/\partial T$) using data from the pre-treatment period.
 - (c) For each date d and period of the day k in the series, we compute
 - i. Cooling price p^T and social marginal cost of cooling c_{dk}^T using equation (10) and (11),
 - ii. Cooling price semi elasticity ε_{dk}^T using equation (13), and
 - iii. The welfare effect $\Delta W_{\text{Moral},dk}$, $\Delta W_{\text{DR},dk}$ and its associated gain from avoided outage $\Delta VoLL$ using Table 2.
3. From the bootstrap, we collect our estimates of the welfare effect and its confidence interval and compute the aggregate welfare effect.

6.3 Welfare impact of Flex Alerts


We find that households' semi-elasticity for cooling setpoints ranges from 0.2 to 0.59. The interpretation of these estimates is that 100% of the price increase will increase the household cooling setpoint by 0.2 to 0.59 °F. We find that households are more elastic during the peak period as shown in Figure A5. Our estimates are comparable in magnitude to semi-elasticity estimates for winter heating from stated preference, which is in the range of -0.31 to -0.97 (Brewer, 2023).

Figure 10 presents our welfare effect estimates. Figure 10a shows the change in total welfare by day. Detailed change in producer and consumer surplus are given in Figure A6a to A6b. In the first six days, we find no significant change in producer surplus and consumer surplus during the peak period as household cooling setpoint response is very small during the peak period. From September 3rd to 5th, 2022, we find positive but modest producer surplus gain in the before-peak period as households reduced cooling setpoint when it is profitable for the utility to deliver more electricity. Consumers also gain utility from cooler setpoints during this period without experiencing a change in total cost due to the moral subsidy from precooling recommendation, resulting in a positive consumer surplus gain.

During the peak period on September 6th, 2022, when the social marginal cost of electricity is very high, the salient phone alert successfully induces conservation from cooling. The producer surplus increases significantly by about \$4 million as massive conservation is realized. However, by responding to the conservation requests, households experienced disutility from having a higher cooling setpoint, which resulted in about \$1 million loss in consumer surplus. From September 7th until the end of the series, this loss in consumer surplus grew to about \$2 million as households increase their cooling setpoint even higher during the peak period. As the social



(a) Change in Total Welfare by Day



(b) Overall Welfare Effect

Figure 10. Simulated welfare effect of Flex Alerts

marginal cost of electricity is relatively lower than on September 6th, the producer surplus gain in this period is modest and not enough to compensate for the loss in consumer surplus.

Figure 10b provide a detailed breakdown of the welfare effect from low salience and high salience Flex Alerts. Over the ten-day period, the Flex Alerts results in a total producer surplus gain of \$5.7 million and a total consumer surplus loss of \$3.5 million. Taking into account the additional benefit of \$67.5 million from the avoided outage, the net welfare effect of Flex Alerts is \$69.8 million. On a standard Flex Alert, the average welfare gain is \$2.3 million per day, while after the phone alert, this welfare gain increases to \$14 million per day. This estimate is robust to conservative assumptions regarding blackout probability. Even if the probability of an outage

were only 50%, the net benefits would remain positive and substantial.

7 Policy Implications

To assess the effectiveness of Flex Alerts, we compare our findings with previous research. Earlier studies show that governor’s requests during grid emergencies reduced heating setpoints by 1.1°F and fan run-time by 1.47 minutes per hour (Brewer and Crozier, 2025), while time-of-use pricing increased cooling setpoints by 1.04 °F and reduced compressor run-time by 3.2 minutes per hour (Fu et al., 2024). Time-of-use pricing combined with automated cooling schedules increased setpoints by 1.7°F and reduced compressor run-time by 7.3 minutes per hour (Blonz et al., 2025).

Our results indicate that standard Flex Alerts have minimal impact on cooling behavior, with no significant effect on setpoints or compressor run-time during peak periods. Only after a phone alert did households increase cooling setpoints by an average of 0.3°F, reducing compressor run-time by 1 minute per hour, significantly less effective than time-of-use pricing or automation. Our analysis identifies the phone alert as the primary channel that increases the salience of Flex Alert rather than the governor request.¹⁹ While highly effective, high-salience alerts are a finite resource best reserved for critical conditions considering the potential habituation trade-off.

We also observed household habituation to consecutive Flex Alerts, decreasing their effectiveness after the third day. In consecutive grid emergencies, scheduling sufficient time between requests to counter habituation as recommended by Ito et al. (2018) is not feasible. While incorporating phone alerts successfully increased the magnitude of the response, with households showing increasing engagement over time, this approach is not sustainable for every Flex Alert. Rather than the current approach where customers need to sign up to receive Flex Alert notifications, enrolling customers in utility-issued Flex Alert notifications by default could provide a low-cost alternative to ensure conservation requests reach more households. As documented by Fowlie et al. (2021) in other setting, the use of default provision to mitigate future grid emergencies is justified by the positive welfare implication.

For demand response participants, we observe more substantial effects. With low salience, demand response events increased cooling setpoints by 0.5°F and reduced compressor run-time by up to 3 minutes per hour. After the phone alert, these effects more than doubled: setpoints increased by up to 1.3F, and compressor run-time decreased by up to 7 minutes per hour. This suggests that combining moral suasion with automation yields effectiveness comparable to time-of-use pricing (Fu et al., 2024) and its combination with automation (Blonz et al., 2025; Prest, 2020). Our findings demonstrate the complementary between demand response programs and moral suasion. Demand response participants responded more robustly when automation was

¹⁹ On the first day of the series, the California governor issued a state of emergency, which includes a conservation request. Our estimates suggest that the effect only lasts hours and does not last until the next day. On September 6th, the governor made another request via an executive order with a phone alert, similar to the request in Michigan (Brewer and Crozier, 2025), which resulted in a higher and more persistent response.

combined with moral appeals, providing new evidence for the effectiveness of this combined approach. As energy systems increasingly incorporate smart technology, opportunities for automated demand-side management will expand.²⁰

Despite lower per-household responses to Flex Alerts compared to demand response events, their aggregate impact remains significant. When both programs operated simultaneously, demand response participants contributed only about 10% of total demand reduction. Our welfare analysis shows that these combined efforts generated a net welfare gain of \$69.8 million during the ten-day period, with Flex Alerts increasing producer surplus during peak periods when social marginal costs were high and generating benefits from avoided outage.

Although demand response events generate stronger per-household responses, their aggregate impact remains modest due to low enrollment rates. Currently, demand response participation among smart thermostat users in California is just above 25%. Our estimates suggest that enrolling all 1.9 million smart thermostat owners could yield a demand reduction of 600 MW from cooling behavior alone. Increasing participation through both higher enrollment of existing smart thermostat users and automatic enrollment of new adopters could substantially improve conservation efforts during grid emergencies.

8 Conclusion

Our findings reveal that smart technology is not a perfect substitute for human attention. In fact, we show they are complements during high-stakes grid emergency: high salience appeals reinforce compliance with automated demand response. We leverage high-frequency smart thermostat data from Ecobee’s Donate-Your-Data program with a total of 6.6 million household-hour observations across California and surrounding states. Utilizing a natural experiment created by ten consecutive Flex Alert days during an extreme heat wave in September 2022, we employ a generalized difference-in-differences research design to estimate the causal effect of Flex Alerts and automated demand response on household cooling behavior.

We find that Flex Alerts, even with phone alerts are less salient nudge of household cooling behavior compared to other voluntary requests using thermostat reference point (Brewer and Crozier, 2025) and time-of-use pricing (Fu et al., 2024; Blonz et al., 2025). Non-demand response households exhibit minimal changes in cooling behavior under standard Flex Alerts, with the treatment effects becoming statistically significant only after the issuance of a widespread phone alert on September 6th, 2022. Being in a demand response event is more effective in changing household cooling behavior. Specifically, our findings highlight that demand response participants who receive automated thermostat override in a demand response event exhibit more substantial and timely changes in their cooling behavior. This finding underscores the effectiveness of combining moral suasion with automated demand response in grid emergency management.

²⁰ Utilities also increasingly use behind-the-meter storage to provide reliability services, see <https://www.utilitydive.com/news/pge-tesla-launch-program-to-use-customers-powerwall-batteries-to-tackle/> (last accessed June 15th, 2024).

Moreover, this study documents the dynamic effects of repeated conservation requests. Our estimates indicate a quick habituation to standard Flex Alerts, with treatment effects diminishing after the third day. However, the phone alert on the seventh day of the series resulted in a significant increase in conservation efforts, with a continued increase in response through subsequent days. In contrast to previous literature on habituation, we do not observe household habituation to repeated Flex Alerts after the phone alert. Following the conclusion of the Flex Alert series, we find a permanent effect on household cooling setpoint that lasted for two weeks, even though there is no longer a grid emergency.

Our welfare analysis highlights the substantial aggregate impact on demand reaching 1,300 MW in the peak period in these Flex Alert series. This series of Flex Alerts resulted in a welfare gain of \$69.8 million. Notably, the phone alert increases the maximum welfare gain by more than six times. Despite the more pronounced individual responses from demand response events, our analysis reveals that the majority of savings are attributable to non-demand response households due to the low proportion of demand response enrollment.

By providing granular evidence on household-level responses to conservation requests and documenting the cooling behavior mechanisms, the study offers valuable insights for policymakers in dealing with grid emergencies. First, even though the phone alert successfully increases the salience of voluntary requests, frequent use is unsustainable due to habituation. Instead, enrolling households in Flex Alert notification by default from utility could be a way to increase salience. Second, there is an incentive for the utility to increase available demand response resources by promoting the adoption of smart technology and enrollment in demand response programs. With climate change driving more frequent and severe weather events and, consequently, higher peak electricity demand, the insights provided by this research are important for the development of more effective demand response strategies.

References

- Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge**, “When should you adjust standard errors for clustering?,” *The Quarterly Journal of Economics*, 2023, 138 (1), 1–35.
- Allcott, Hunt and Judd B Kessler**, “The welfare effects of nudges: A case study of energy use social comparisons,” *American Economic Journal: Applied Economics*, 2019, 11 (1), 236–276.
- **and Todd Rogers**, “The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation,” *American Economic Review*, 2014, 104 (10), 3003–37.
- Andreoni, James**, “Giving with impure altruism: Applications to charity and Ricardian equivalence,” *Journal of Political Economy*, 1989, 97 (6), 1447–1458.
- Auffhammer, Maximilian, Patrick Baylis, and Catherine H. Hausman**, “Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the United States,” *Proceedings of the National Academy of Sciences*, 2017, 114 (8), 1886–1891.
- Bailey, Megan R, David P Brown, Blake C Shaffer, and Frank A Wolak**, “Take the Load Off: Time and Technology as Determinants of Electricity Demand Response,” Technical Report, National Bureau of Economic Research 2025.
- Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro**, “Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century,” *Journal of Political Economy*, 2016, 124 (1), 105–159.
- Baylis, Patrick W and Judson Boomhower**, “Mandated vs. Voluntary Adaptation to Natural Disasters: The Case of U.S. Wildfires,” *NBER WORKING PAPER SERIES*, 2021.
- Blonz, Joshua A.**, “Making the Best of the Second-Best: Welfare Consequences of Time-Varying Electricity Prices,” *Journal of the Association of Environmental and Resource Economists*, 2022, 9 (6), 1087–1126.
- Blonz, Joshua, Karen Palmer, Casey J. Wichman, and Derek C. Wietelman**, “Smart Thermostats, Automation, and Time-Varying Prices,” *American Economic Journal: Applied Economics*, January 2025, 17 (1), 90–125.
- Bollinger, Bryan K. and Wesley R. Hartmann**, “Information vs. Automation and Implications for Dynamic Pricing,” *Management Science*, January 2020, 66 (1), 290–314.
- Boomhower, Judson and Lucas Davis**, “Do energy efficiency investments deliver at the right time?,” *American Economic Journal: Applied Economics*, 2020, 12 (1), 115–139.
- Brandon, Alec, Christopher M Clapp, John A List, Robert D Metcalfe, and Michael Price**, “The Human Perils of Scaling Smart Technologies: Evidence from Field Experiments,” *National Bureau of Economic Research Working Paper*, 2022.
- Brewer, Dylan**, “Equilibrium sorting and moral hazard in residential energy contracts,” *Journal of Urban Economics*, 2022, 129, 103424.
- , “Household responses to winter heating costs: Implications for energy pricing policies and demand-side alternatives,” *Energy Policy*, 2023, 177, 113550.

- **and Jim Crozier**, “Who heeds the call to conserve in an energy emergency? Evidence from smart thermostat data,” *Journal of the Association of Environmental and Resource Economists*. *Forthcoming*, 2025.
- Brown, David P and Lucija Muehlenbachs**, “The value of electricity reliability: Evidence from battery adoption,” *Journal of Public Economics*, 2024, 239, 105216.
- Burkhardt, Jesse, Kenneth T Gillingham, and Praveen K Kopalle**, “Field experimental evidence on the effect of pricing on residential electricity conservation,” *Management Science*, 2023, 69 (12), 7784–7798.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller**, “Bootstrap-based improvements for inference with clustered errors,” *The Review of Economics and Statistics*, 2008, 90 (3), 414–427.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick**, “Optimal Defaults,” *American Economic Review*, May 2003, 93 (2), 180–185.
- Costa, Francisco and François Gerard**, “Hysteresis and the welfare effect of corrective policies: Theory and evidence from an energy-saving program,” *Journal of Political Economy*, 2021, 129 (6), 1705–1743.
- Davis, Lucas W. and Paul J. Gertler**, “Contribution of air conditioning adoption to future energy use under global warming,” *Proceedings of the National Academy of Sciences*, May 2015, 112 (19), 5962–5967.
- DellaVigna, Stefano, John A List, and Ulrike Malmendier**, “Testing for altruism and social pressure in charitable giving,” *The Quarterly Journal of Economics*, 2012, 127 (1), 1–56.
- Demand Side Analytics**, “SDG&E ACSDA Evaluation,” Available at <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/demand-response/demand-response-workshops/2023-load-impact-protocol-workshops/20230501-sdge-acsd-drme.pdf> (last accessed June 15th, 2024), 2023.
- Deryugina, Tatyana**, “The Fiscal Cost of Hurricanes: Disaster Aid versus Social Insurance,” *American Economic Journal: Economic Policy*, August 2017, 9 (3), 168–198.
- Dinerstein, Michael, Nadia Lucas, Ishan Nath, and Johanna Rayl**, “Private Markets and Public Assistance for Natural Disaster Supplies,” *AEA Papers and Proceedings*, May 2025, 115, 385–390.
- Driscoll, John C. and Aart C. Kraay**, “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data,” *The Review of Economics and Statistics*, 1998, 80 (4), 549–560.
- EIA**, “Residential Energy Consumption Survey (RECS),” Technical Report, US Energy Information Administration 2020.
- , “Electric Power Monthly September 2022,” Technical Report, US Energy Information Administration 2022.
- Ferraro, Paul J and Michael K Price**, “Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment,” *Review of Economics and Statistics*, 2013, 95 (1), 64–73.
- Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C Anna Spurlock, Annika Todd-Blick, and Peter Cappers**, “Default effects and follow-on behaviour: Evidence from an electricity pricing program,” *The Review of Economic Studies*, 2021, 88 (6), 2886–2934.

- Fu, Zheng, Kevin Novan, and Aaron Smith**, “Do time-of-use prices deliver energy savings at the right time?,” *Journal of Environmental Economics and Management*, 2024, p. 103054.
- Harding, Matthew and Steven Sexton**, “Household Response to Time-Varying Electricity Prices,” *Annual Review of Resource Economics*, October 2017, 9 (1), 337–359.
- He, Guojun and Takanao Tanaka**, “Energy Saving May Kill: Evidence from the Fukushima Nuclear Accident,” *American Economic Journal: Applied Economics*, April 2023, 15 (2), 377–414.
- Holladay, J. Scott, Michael K. Price, and Marianne Wanamaker**, “The perverse impact of calling for energy conservation,” *Journal of Economic Behavior & Organization*, 2015, 110, 1–18.
- Ito, Koichiro**, “Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing,” *American Economic Review*, 2014, 104 (2), 537–563.
- , **Takanori Ida, and Makoto Tanaka**, “Moral suasion and economic incentives: Field experimental evidence from energy demand,” *American Economic Journal: Economic Policy*, 2018, 10 (1), 240–267.
- Jacob, Maria, Ryan Jenkinson, Daniel L. Garcia, Robert D. Metcalfe, Andrew R. Schein, Cohen R. Simpson, and Lucy Yu**, “The Impact of Demand Response on Energy Consumption and Economic Welfare,” *Working Paper*, 2023.
- Jessoe, Katrina and David Rapson**, “Knowledge is (less) power: Experimental evidence from residential energy use,” *American Economic Review*, 2014, 104 (4), 1417–1438.
- Joskow, Paul and Jean Tirole**, “Reliability and competitive electricity markets,” *The RAND Journal of Economics*, 2007, 38 (1), 60–84.
- Levitt, Steven D and John A List**, “What do laboratory experiments measuring social preferences reveal about the real world?,” *Journal of Economic Perspectives*, 2007, 21 (2), 153–174.
- Mahadevan, Meera and Ajay Shenoy**, “The political consequences of resource scarcity: Targeted spending in a water-stressed democracy,” *Journal of Public Economics*, April 2023, 220, 104842.
- Meier, Alan, L Rainer, A Daken, T Ueno, M Pritoni, and Dan Baldewicz**, *What can connected thermostats tell us about American heating and cooling habits?* 2019.
- Peplinski, McKenna and Kelly T Sanders**, “Residential electricity demand on CAISO Flex Alert days: a case study of voluntary emergency demand response programs,” *Environmental Research: Energy*, December 2023, 1 (1), 015002.
- Prest, Brian C**, “Peaking interest: How awareness drives the effectiveness of time-of-use electricity pricing,” *Journal of the Association of Environmental and Resource Economists*, 2020, 7 (1), 103–143.
- Reiss, Peter C and Matthew W White**, “What changes energy consumption? Prices and public pressures,” *The RAND Journal of Economics*, 2008, 39 (3), 636–663.
- Strömberg, David**, “Natural Disasters, Economic Development, and Humanitarian Aid,” *Journal of Economic Perspectives*, July 2007, 21 (3), 199–222.
- Taylor, Charles A. and Hannah Druckenmiller**, “Wetlands, Flooding, and the Clean Water Act,” *American Economic Review*, April 2022, 112 (4), 1334–1363.

The Brattle Group, , “California’s Virtual Power Potential: How Five Consumer Technologies Could Improve the State’s Energy Affordability,” Available at <https://www.brattle.com/wp-content/uploads/2024/04/Californias-Virtual-Power-Potential-How-Five-Consumer-Technologies-Could-Improve-the-States-Energy-Affordability.pdf> (last accessed June 15th, 2024), 2024.

Wichman, Casey J, Laura O Taylor, and Roger H Von Haefen, “Conservation policies: Who responds to price and who responds to prescription?,” *Journal of Environmental Economics and Management*, 2016, 79, 114–134.

Wolak, Frank A, “Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment,” *American Economic Review*, 2011, 101 (3), 83–87.

Zivin, Joshua Graff and Matthew Neidell, “Temperature and the Allocation of Time: Implications for Climate Change,” *Journal of Labor Economics*, 2014.

Appendix

A Additional tables

Table A1. Summary of the Flex Alerts timing

Date	Start of peak	End of peak
Wednesday, August 31, 2022	4 p.m.	9 p.m.
Thursday, September 1, 2022	4 p.m.	9 p.m.
Friday, September 2, 2022	4 p.m.	9 p.m.
Saturday, September 3, 2022	4 p.m.	9 p.m.
Sunday, September 4, 2022	4 p.m.	9 p.m.
Monday, September 5, 2022	4 p.m.	10 p.m.
Tuesday, September 6, 2022	4 p.m.	9 p.m.
Wednesday, September 7, 2022	4 p.m.	9 p.m.
Thursday, September 8, 2022	3 p.m.	10 p.m.
Friday, September 9, 2022	4 p.m.	9 p.m.

Source: CAISO's Grid Emergencies History Report.

Table A2. Summary of the demand response event timing

Date	Start of event	End of event	Demand response treatment
Wednesday, August 31, 2022	7:25 p.m.	7:40 p.m.	7-8 p.m.
Thursday, September 1, 2022	5:00 p.m.	8:25 p.m.	5-9 p.m.
Monday, September 5, 2022	6:40 p.m.	8:20 p.m.	6-9 p.m.
Tuesday, September 6, 2022	4:10 p.m.	9:05 p.m.	4-9 p.m.
Wednesday, September 7, 2022	4:10 p.m.	8:55 p.m.	4-9 p.m.
Thursday, September 8, 2022	5:05 p.m.	8:15 p.m.	5-9 p.m.

Note. We define an hour with at least 15 minutes of demand response event within the hour as treated with demand response. Source: CAISO's Today's Outlook.

Table A3. Summary of the September 2022 Flex Alert events

Date	Time posted	Announcement
Wednesday, August 31, 2022	12:48 p.m.	Flex Alert issued
Wednesday, August 31, 2022	5:40 p.m.	Flex Alert extended
Thursday, September 1, 2022	4:23 p.m.	Flex Alert extended
Saturday, September 3, 2022	8:21 a.m.	Flex Alert issued
Saturday, September 3, 2022	4:03 p.m.	Flex Alert extended
Sunday, September 4, 2022	5:05 p.m.	Flex Alert extended
Monday, September 5, 2022	4:28 p.m.	Flex Alert extended
Tuesday, September 6, 2022	9:10 p.m.	Flex Alert extended
Wednesday, September 7, 2022	9:12 p.m.	Flex Alert extended
Thursday, September 8, 2022	10:00 p.m.	Flex Alert extended

Source: @flexalert Twitter account posts.

B Additional figures



Figure A1. Trends in Flex Alert events from 2015 to 2022




Figure A2. Example of Flex Alert Notification




Figure A3. Electricity market condition in California during the September 2022 Flex Alert

Note. Panel A and B show the extremely high wholesale electricity prices during the peak period from August 31st to September 9th both in the day-ahead and real-time market in CAISO NP15 Zone. On the customer side, Panel C documents that households in California also experienced significantly more power outages during the peak period of September 5th and September 6th, with more than 70,000 customers affected.




Figure A4. Dynamic treatment effects of Flex Alert using difference-in-differences sample

Note. This figure shows the estimates of the event study regressions using equation (9) using only the difference-in-differences sample period. The dots correspond to the lead and lag coefficient estimates for non-demand response households for each day. The highlighted area shows the 95% confidence interval which is two-way clustered at state and hour-of-sample level.






Figure A5. Cooling price semi-elasticity during the Flex Alerts



(a) Change in Producer Surplus



(b) Change in Consumer Surplus

Figure A6. Detailed welfare effect of Flex Alerts

C Detail on energy conversion calculation

Since we do not observe the characteristics of each household's cooling system, we adopt similar engineering assumptions as [Blonz et al. \(2025\)](#) and assume that electricity consumed for cooling, $x(T)$, is a linear function of air conditioner compressor run-time, $\kappa(T)$. We assume a residential central air conditioner in the southern region following Appendix A of EIA Updated Buildings Sector Appliance and Equipment Costs and Efficiencies, with a Unit Capacity Ratio (UCR) of 36,000 BTU per hour per unit with a typical Seasonal Energy Efficiency Ratio (SEER) of 14.4 BTU per W per hour. Equation (A1) shows the energy conversion calculations to convert compressor run-time, κ , in minutes per hour per household, to electricity consumption for cooling per household, x in kWh per household per hour.

$$\begin{aligned} x(\text{kWh}/\text{HH}\cdot\text{hour}) &= \kappa(\text{min}/\text{hour}) \times \frac{\text{UCR (BTU/hour} \cdot \text{HH)}}{\text{SEER (BTU/W} \cdot \text{hour)}} \\ &\quad \times \frac{1 \text{ hour}}{60 \text{ min}} \times \frac{10^{-3} \text{ kW}}{\text{W}} \\ &= \kappa(\text{min}/\text{hour}) \times 0.0417(\text{kWh}/\text{HH} \cdot \text{min}). \end{aligned} \quad (\text{A1})$$

We also need estimates of changes in demand for cooling in MW per HH. Thus, we convert our treatment effect on compressor run-time, in minutes per hour per household, to electricity consumption for cooling per household, in MW per household. The following equivalent version of the equation (A1) shows the energy conversion calculations.

$$\begin{aligned} \Delta\text{Demand}_{ith}(\text{MW}/\text{HH}) &= \hat{\beta}_{ith}^{\kappa}(\text{min}/\text{hour}) \times \frac{\text{UCR (BTU/hour} \cdot \text{HH)}}{\text{SEER (BTU/W} \cdot \text{hour)}} \\ &\quad \times \frac{1 \text{ hour}}{60 \text{ min}} \times \frac{10^{-6} \text{ MW}}{\text{W}} \\ &= \hat{\beta}_{ith}^{\kappa}(\text{min}/\text{hour}) \times 0.0000417(\text{MWh}/\text{HH} \cdot \text{min}). \end{aligned} \quad (\text{A2})$$

One minute per hour of compressor run-time in our data translates to an electricity consumption of 0.042 kW per household. For convenience, we also define the cooling price in dollars per one additional °F of lower cooling setpoint as $p^T = p^{\frac{\partial x}{\partial T}}$.²¹ We also use $\frac{\partial x}{\partial T}$ to convert the social marginal cost of electricity to the social marginal cost of cooling.

²¹ Our linear assumption allows us to obtain $\frac{\partial x}{\partial T}$ which reflects the energy efficiency of the household from the data by directly estimating the marginal effect of cooling setpoint, T , on compressor run-time, κ .

D Identifying demand response event from Ecobee data

Using the thermostat event name variable in the Ecobee data, we identified several demand response event names. First, we identify a general demand response event name that contains Demand Response ("DR") and Precooling ("PC" or "PRC"), which is a common term in AC load control.²² The second one is the California Public Utility Commission pilot Power Saver Rewards Program that started in May 2022.²³ Participating customers receive a bill credit of \$2 per kWh of electricity savings in a Flex Alert during our sample period. The program incurs no penalty for the household when they are enrolled and decide not to respond to emergency requests. Third, we identify program names from SDGE, they are AC Saver DA ("ACSDA"), Bring Your Own Thermostat ("BYOT"), and Reduce Your Use ("RYU").²⁴ Lastly, we also identify a demand response event name from Portland General Electric of which they collaborate with PGE.²⁵

E Results for alternative outcome variables

We look at two alternative outcome variables: the indoor temperatures of the home in °F and the fan run-time in minutes per hour. We use the indoor temperature as a measure of the comfort level of the household. The change in indoor temperature is an estimate of the household's opportunity cost for following the Flex Alerts recommendation. The fan run-time is the duration that the fan is running to circulate the air in the house within an hour.

Figure 5 shows the hourly mean of the two alternative outcome variables for households in California and the controls before and after the treatment. California households, on average, experience higher indoor temperatures during Flex Alerts. We also observe a change in fan run-time pattern in the peak period after the Flex Alerts, similar to those of compressor run-time.

We also estimate equation (7) for indoor temperature and fan run-time. Following the Flex Alert recommendation, we expect that the treatment effect on realized indoor temperature to be positive during the peak period. For the fan run-time, we expect the treatment effect to be the same as the compressor run-time, considering how an HVAC system works.

Table 1 shows the effect of Flex Alert on indoor temperature and fan run-time for the before-

²² For example, in Minnesota, Xcel Energy enrolled more than 50% of their residential customer to voluntary AC load control. See Docket No. E002/RP-19-368 Appendix G1 of Minnesota Public Utility Commission <https://www.edockets.state.mn.us/edockets/searchDocuments.do?method=showPoup&documentId=10FBAE6B-0000-C040-8C1D-CC55491FE76D&documentTitle=20197-154051-03> (last accessed June 15th, 2024).

²³ See <https://www.cpuc.ca.gov/> (last accessed June 15th, 2024). This program is managed by the three private utilities: Southern California Edison (SCE), San Diego Gas & Electric (SDGE), and PG&E.

²⁴ The SDGE ACSDA Evaluation reported around 17,000 SDGE customers participated in this program in 2022. See the <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/demand-response/demand-response-workshops/2023-load-impact-protocol-workshops/20230501-sdge-acsd-drmeec.pdf> (last accessed June 15th, 2024). Meanwhile, the RYU programs named the demand response event a RYU day; this program applies to residential customers enrolled in a time-of-use rate. See <https://www.sdge.com/residential/pricing-plans/about-our-pricing-plans/whenmatters> (last accessed June 15th, 2024).

²⁵ See <https://portlandgeneral.com/smart-thermostat-enrollment> for more information (last accessed June 15th, 2024). They call the demand response event a Peak Time Event when they will adjust the household cooling setpoint by 1 to 3 degrees higher than the previous setpoint.




Figure A7. Hourly mean of indoor temperature and fan run-time before and after the Flex Alert

Note. This figure shows the hourly mean of the indoor temperature and fan run-time before and after the first Flex Alert announcements on August 31st, 2022, at 12:48 p.m. using observation ranging from August 18th to September 9th, 2022. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.

peak, peak, and after-peak periods. Column 1 shows the treatment effect estimates for the indoor temperature. On Flex Alert days, California households experience higher indoor temperatures than on normal days. In standard Flex Alerts, even though the treatment effect on the cooling setpoint is small, households experience 0.33 to 0.42 °F higher indoor temperatures throughout the day. Demand response participants experience 0.2 °F higher indoor temperatures during a demand response event. After the phone alerts, households experience 0.67 to 0.98 °F higher indoor temperatures throughout the day, while the demand response participant experiences 0.62 °F higher indoor temperatures in a demand response event. Higher indoor temperature is not solely related to the cooling setpoint; rather, it is affected by the physics of the cooling system. As the outdoor temperature for California households during a Flex Alert is extremely high, with a relatively similar cooling setpoint, it is harder for the system to achieve the target temperature. Some households might also turn off their systems rather than increase the cooling setpoint. As shown in Table A5 relative to before the peak period, about two to four percent more households are turning off their cooling system during the peak period. After the phone alert, the treatment effect on indoor temperature is also positive throughout the day and higher compared to the treatment effects in the standard Flex Alerts. Comparing the magnitude of response during the peak period for both treatments, the estimates suggest that households experience an even higher indoor temperature after the phone alert compared to when they receive standard Flex Alerts announcements.

Columns 2 report the treatment effect estimates for fan run-time. The coefficient estimates for the fan run time are relatively similar to the coefficient estimates for compressor run-time in Table 1 for the before-peak and peak periods but have a smaller effect size. We suspect this is due to the recommendation to turn the cooling system off but keep the fan on during the peak period to keep the air circulation.

Table A4. Households Responses to Flex Alert

	(1) Indoor Temperature	(2) Fan Run-Time
After First Tweet		
Before-Peak	0.312*** (0.104)	0.565** (0.239)
Peak	0.374*** (0.092)	0.232 (0.608)
Peak $\times$ 1 (DR Event)	0.143* (0.082)	0.056 (0.642)
After-Peak	0.239*** (0.077)	1.206** (0.495)
After Phone Alert		
Before-Peak	0.877*** (0.166)	-0.366 (0.625)
Peak	0.761*** (0.142)	-0.317 (0.653)
Peak $\times$ 1 (DR Event)	0.565*** (0.131)	-1.906** (0.874)
After-Peak	0.565*** (0.106)	2.597** (1.125)
Pre-treatment Mean		
Before-Peak	74.87	12.64
Peak	76.47	23.34
After-Peak	75.44	16.74
No. of Household	11,821	12,135
Observations	6,377,563	6,632,642

Note. This table reports regression coefficients from difference-in-differences regression estimated using equation (7) on the alternative outcome variables. The first panel show the effect of low salience Flex Alerts, while the second panel show the effect of high salience Flex Alerts. Each row shows the effect at different period of the day. The peak period is defined as following CAISO's Flex Alert hours of the day. Column (1) shows the effect on indoor temperature in °F. Column (3) shows the effect on fan run-time in minutes per hour. The control variables include daily maximum temperature, outdoor temperature, outdoor relative humidity, precipitation, wind speed, and cloud cover. The fixed effects include hour-of-sample indicators and hour-by-day-of-week-by-household fixed effects. The sample period is from August 18th to September 9th, 2022. Standard errors reported follow the [Driscoll and Kraay \(1998\)](#) inference.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A8 shows the average treatment effects for each hour of the day on indoor temperature and fan run-time. Figure A8a reports the hourly treatment effects on the indoor temperature. After the initial Flex Alert announcement, the treatment effect is positive throughout the day with the highest effect of about 0.5 °F between 8 a.m. to 12 p.m. After 12 pm, the effect goes down and




Figure A8. Hourly treatment effect estimates of the Flex Alerts

Note. This figure shows the estimates of the hourly responses to Flex Alert estimated using equation (8). The hour labels for each coefficient refer to the ending time of each one-hour interval. The dots correspond to the hourly treatment effect estimates for non-demand response households for the standard Flex Alerts ($\hat{\beta}_{FA,h}$) and the Flex Alerts after receiving the phone alert ($\hat{\beta}_{PA,h}$). The highlighted area shows the 95% confidence interval, which follows the [Driscoll and Kraay \(1998\)](#) inference. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.

ramps up slowly and reaches its peak at the end of the peak period. The treatment effect after the phone alert is higher throughout the day compared to the standard Flex Alert treatment. We observe a similar treatment effect pattern after the phone alert, a higher treatment effect of about 1.3 °F is observed between 8 a.m. and 9 a.m. The treatment effect ramps by a small amount at the end of the peak period at about 0.8 °F higher.

The hourly treatment effects on the fan run-time shown in Figure A8b have a similar hourly pattern as the compressor run-time in Figure 6b. However, the treatment effect on fan run-time during the peak period is not significant. This indicates that households might turn off the cooling system but keep their fan on during the peak period to keep the air circulation, which is shown by the treatment effects on the indicator for turning off the cooling system shown in Table A5.

The lead coefficient estimates on indoor temperature shown in Figure A9a also evolve around zero prior to treatment. The lead coefficient estimates for fan run-time in Figure A9b are also statistically zero, however, they are more noisy. The dynamic treatment effect for indoor temperature steadily increased and peaked on the seventh day. Even though the cooling setpoint response is minimal since the outdoor temperature is high during the peak period, households still experience 1 °F higher indoor temperature during the weekend as shown in Figure A9a. Looking at the indoor temperature, the treatment effect is going down after two weeks. We suspect this is due to the accumulation of cooler outdoor temperature exposure over time and the energy efficiency of the home.

Figure A10 shows the effect of demand response on indoor temperature and fan run-time. In a demand response event, demand response participants experience up to 0.4 °F higher indoor temperature during the peak period. After the phone alert, the treatment effect is higher, up to 0.9 °F higher indoor temperature during the peak period. The fan run-time reduction from demand




Figure A9. Dynamic treatment effect estimates of the Flex Alerts

Note. This figure shows the estimates of the event study regressions using equation (9). The dots correspond to the lead and lag coefficient estimates for non-demand response households for each day. The highlighted area shows the 95% confidence interval which is two-way clustered at state and hour-of-sample level. We extend the event study sample from a month before to two weeks after the Flex Alert series. We present estimates from using only the difference-in-differences sample period in Figure A4.

response is of similar magnitude as the compressor run-time reduction.

F Results for extensive margin responses

Figure A7 shows the hourly mean of the binary alternative outcome variables for households in California and the controls before and after the treatment within the sample periods. We clearly see a different pattern during the peak period from California households after Flex Alerts for all alternative outcomes, except for the precooling indicator. The trends in the outcome variables for control groups after Flex Alerts are similar to those before the Flex Alert events, which are ideal characteristics for a control group in a difference-in-difference approach.

We also estimate equation (7) for the alternative outcome variables. An increase in the probability that the event variables are on hold means households override their scheduled settings more. During Flex Alert days, there could be many reasons households are overriding their schedule, one of which could be discomfort or setting their cooling setpoint to follow the Flex Alert recommendation. An increase in the probability that households set their cooling setpoint below 70 °F means that households are following the recommendation to precool their house. So, we expect the sign of the second binary outcome variable to be positive before peak periods. Similarly, an increase in the probability that households set their cooling setpoint above 78 °F means that households are following the recommendation to increase their thermostat setting during peak periods. Hence, we expect the sign of the coefficient for the third binary outcome variable to be positive during the peak period. Suppose households think that turning off the cooling system is a way to conserve more energy and thus help the grids. In that case, we expect an increase in the probability that households turn their cooling system off during peak periods relative to the




Figure A10. Hourly treatment effect estimates of demand response event

Note. This figure shows the effect of the demand response event for demand response household estimated using equation (8). These coefficients are only identified during demand response events. The hour labels for each coefficient refer to the ending time of each one-hour interval. The dots correspond to the hourly effect of being in a demand response event for demands response households during the standard Flex Alerts ($\hat{\delta}_{FA,h}$) and the Flex Alerts after receiving the phone alert ($\hat{\delta}_{PA,h}$). The highlighted area shows the 95% confidence interval, which follows the [Driscoll and Kraay \(1998\)](#) inference. The two vertical lines at 4 p.m. and 9 p.m. show the typical start and end of the peak period.

non-peak period within Flex Alert days.

Table A5 shows the effect of Flex Alert on the alternative outcome variables for the before-peak, peak, and after-peak periods. Column 1 shows the effect of the treatments on the percentage of households overriding their scheduled thermostat settings. After the first announcement of the Flex Alert, only less than one percent of the households override their schedule in the before-peak and peak periods. After the phone alert treatment, there is a positive effect, but it is not significant. In both of the treatments, the effect on the number of people who put their schedule on hold is minimal. This rules out putting the schedule on hold as the mechanism through which household change their cooling setpoint.

Columns 2 and 3 show the treatment effects on the compliance of the precooling setpoint for the before-peak period and the Flex Alert setpoint for the peak period. After the first announcement of the Flex Alert, the effect on compliance with precooling before the peak period is below one percent, while there is no significant effect on compliance with the Flex Alert in the peak period. After the phone alert, there is zero effect on compliance with the precooling before the peak period. In the peak period, there are 2.6 percent more households who set their cooling setpoint above 78 °F. These estimates reveal that there are no significant changes in the number of household precooling before the Flex Alert hours. Contrasting the salient of the two different treatments, the effect on compliance with the peak period target setpoint is higher after the phone alert compared to the standard Flex Alert announcement.

Column 4 shows the treatment effect on whether or not households turn their cooling system off. In both of the treatments, in all periods of the day, the sign is negative; this indicates that




Figure A11. Hourly mean of binary outcomes variables before and after the Flex Alert

Table A5. Extensive Margin of Households Response to Flex Alert

	(1) 1 (On Hold)	(2) 1 (Cool. Setpoint $\leq 70^\circ$)	(3) 1 (Cool. Setpoint $\geq 78^\circ$)	(4) 1 (Cooling Off)
After First Tweet				
Before-Peak	0.008*** (0.003)	0.002*** (0.001)	-0.009*** (0.002)	-0.036*** (0.004)
Peak	0.005 (0.004)	0.001 (0.001)	0.007*** (0.003)	-0.024*** (0.004)
Peak $\times$ 1 (DR Event)	-0.015 (0.013)	-0.006*** (0.002)	0.033*** (0.013)	-0.020*** (0.007)
After-Peak	0.003 (0.002)	0.001 (0.001)	-0.011*** (0.003)	-0.039*** (0.004)
After Phone Alert				
Before-Peak	0.007* (0.004)	-0.002 (0.001)	-0.006** (0.003)	-0.045*** (0.004)
Peak	-0.003 (0.009)	-0.003** (0.002)	0.031*** (0.006)	-0.022*** (0.004)
Peak $\times$ 1 (DR Event)	-0.051*** (0.017)	-0.011*** (0.002)	0.107*** (0.012)	-0.026*** (0.005)
After-Peak	0.000 (0.005)	-0.000 (0.001)	-0.003 (0.004)	-0.042*** (0.006)
Pre-treatment Mean				
Before-Peak	0.21	0.08	0.41	0.23
Peak	0.25	0.07	0.39	0.20
After-Peak	0.24	0.10	0.38	0.22
No. of Household	12,135	12,135	12,135	12,135
Observations	6,632,642	6,632,642	6,632,642	6,632,642

Note. This table reports regression coefficients from difference-in-differences regression estimated using equation (7) on the four alternative binary outcome variables. The first panel show the effect of low salience Flex Alerts, while the second panel show the effect of high salience Flex Alerts. Each row shows the effect at different period of the day. The peak period is defined as following CAISO's Flex Alert hours of the day. Column (1) shows the change in the proportion of households who override their scheduled setting. Column (2) shows the change in the proportion of households who comply with the precooling recommendation. Column (3) shows the change in the proportion of households who comply with the peak period recommendation. Column (4) shows the change in the proportion of households who turn off their cooling system. The control variables include daily maximum temperature, outdoor temperature, outdoor relative humidity, precipitation, wind speed, and cloud cover. The fixed effects include hour-of-sample indicators and hour-by-day-of-week-by-household fixed effects. The sample period is from August 18th to September 9th, 2022. Standard errors reported follow the [Driscoll and Kraay \(1998\)](#) inference.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

during a Flex Alert, an additional two to four percent of the households turn on their cooling system. After the initial Flex Alert announcement, there are 3.3 and 2.5 percent fewer people turning off their cooling system before the peak period and during the peak period consecutively. Similar behavior is observed after the phone alert, there are 4 and 2.3 percent fewer people turning off their cooling system before the peak period and during the peak period consecutively. Within a Flex Alert day, less than one percent of households turn the cooling off during the peak period relative to the before-peak period.

G Detailed welfare effect derivation

In this section, we formalize the welfare effect that we show in Figure 3 in Section 3 using a linear-log demand parameterization in a similar fashion to Ito et al. (2018) and Allcott and Kessler (2019). Consider households with baseline cooling setpoint $T_{idk}^0 = \alpha_{idk}^T + \varepsilon_{idk}^T \ln p^T$. When households receive moral suasion, the cooling setpoint becomes $T'_{idk} = \alpha_{idk}^T + \beta_{dk}^T + \varepsilon_{idk}^T \ln p^T$. The average change in producer surplus (PS) for period k in date d of the Flex Alert series depends on the gap between social cost and retail price, which is given by

$$\Delta PS_{\text{Moral},dk} = \int_{T^0}^{T'} c_{dk}^T - p^T(T) dT \approx \begin{cases} \beta_{dk}^T (c_{dk}^T - p^T) & \text{if } c_{dk}^T \geq p^T, \\ -\beta_{dk}^T (p^T - c_{dk}^T) & \text{if } c_{dk}^T < p^T. \end{cases}$$

The average change in consumer surplus (CS) for period k in date d of the Flex Alert series is given by

$$\Delta CS_{\text{Moral},dk} = \int_{T^0}^{T'} p^T(T) dT \approx \begin{cases} -\frac{1}{2} \frac{(\beta_{dk}^T)^2}{\partial T / \partial p^T} = -\frac{1}{2} (\beta_{dk}^T)^2 \frac{p^T}{\varepsilon_{dk}^T} & \text{if } \beta_{dk}^T \geq 0, \\ \frac{1}{2} \frac{(\beta_{dk}^T)^2}{\partial T / \partial p^T} = \frac{1}{2} (\beta_{dk}^T)^2 \frac{p^T}{\varepsilon_{dk}^T} & \text{if } \beta_{dk}^T < 0. \end{cases}$$

The average total welfare effect of Flex Alerts is then given by the sum of the change in producer surplus and consumer surplus for period k in date d of the series, i.e. $\Delta W_{\text{Moral}} = \Delta PS_{\text{Moral}} + \Delta CS_{\text{Moral}}$.

Now, consider demand response households who receive both moral suasion and monetary incentives via demand response events. When households receive monetary incentives, they receive a price adjustment. We do not observe the price change. Instead, we observe the hour when the demand response event is called by CAISO within the peak period. We use our estimates on the effect of being in a demand response event on the cooling setpoint, δ_{dk}^T . We then infer the effective price for demand response households, $p^{T,DR}$, from these estimates. When demand response households receive moral suasion, the cooling setpoint becomes $T'_{idk} = \alpha_{idk}^T + \beta_{dk}^T + \varepsilon_{idk}^T \ln p^T$. When demand response households receive moral suasion and monetary incentives, the cool-

ing setpoint becomes $T_{idk}^{DR} = \alpha_{idk}^T + \beta_{dk}^T + \varepsilon_{idk}^T \ln p^{T,DR}$. The additional change in producer surplus (PS) for period k in date d of the Flex Alert when a demand response event is called is given by

$$\Delta PS_{DR,dk} = \int_{T'}^{T^{DR}} c_{dk}^T - p^T(T) dT \approx \begin{cases} \delta_{dk}^T (p^{T,DR} - p^T) & \text{if } p^{T,DR} \geq p^T, \\ -\delta_{dk}^T (p^T - p^{T,DR}) & \text{if } p^{T,DR} < p^T. \end{cases}$$

The average change in consumer surplus (CS) for period k in date d of the series is given by

$$\Delta CS_{DR,dk} = \int_{T'}^{T^{DR}} p^T(T) dT \approx \begin{cases} -\frac{1}{2} \frac{(\delta_{dk}^T)^2}{\partial T / \partial p^T} = -\frac{1}{2} (\delta_{dk}^T)^2 \frac{p^T}{\varepsilon_{dk}^T} & \text{if } \delta_{dk}^T \geq 0, \\ \frac{1}{2} \frac{(\delta_{dk}^T)^2}{\partial T / \partial p^T} = \frac{1}{2} (\delta_{dk}^T)^2 \frac{p^T}{\varepsilon_{dk}^T} & \text{if } \delta_{dk}^T < 0. \end{cases}$$

The total welfare effect of the demand response event is then given by the sum of the change in producer surplus and consumer surplus for for period k in date d of the events, i.e. $\Delta W_{DR,dk} = \Delta PS_{DR,dk} + \Delta CS_{DR,dk}$.

The total welfare effect for period k in date d of the series is given by

$$\Delta W_{Aggregate,dk} = N \times \Delta W_{Moral,dk} + N_{DR} \times \mathbb{1}(DRevent,dk) \times \Delta W_{DR},$$

where N is the total number of households in California, and N_{DR} is the total number of demand response participants.