



Energy Institute WP 275R

**Making the Best of the Second-Best: Welfare
Consequences of Time-Varying Electricity Prices**

Joshua Blonz

Revised June 2020

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board. The Energy Institute acknowledges the generous support it has received from the organizations and individuals listed at <https://haas.berkeley.edu/energy-institute/about/funders/>.

© 2020 by Joshua Blonz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

Making the Best of the Second-Best: Welfare Consequences of Time-Varying Electricity Prices

Joshua A. Blonz*

June 2020

Abstract

In electricity markets, the price paid by retail customers during periods of peak demand is far below the cost of supply. This leads to overconsumption during peak periods, requiring the construction of excess generation capacity compared to first-best prices that adjust at short time intervals to reflect changing marginal cost. This paper investigates the response of small commercial and industrial establishments to a four-hour increase in retail electricity prices invoked by individual utilities during peak demand periods 15 times per year. This policy is intended to reduce electricity consumption when generation costs are highest. I find that the approximately tripled prices reduce establishment peak electricity usage by 13.5 percent. Using a model of capacity investment decisions, I find the program delivers 44 percent of the benefits of the first-best policy of continuously varying prices and suggest two simple improvements in program design that could nearly double these welfare gains.

*Federal Reserve Board, 20th & C Street, NW, Washington, D.C. 20551. Email: joshua.blonz@frb.gov. The views in this paper are those of the author and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. I am deeply grateful to my advisors Severin Borenstein and Maximilian Auffhammer for their invaluable advice. I also thank Patrick Baylis, Peter Berck, Fiona Burlig, Lucas Davis, Meredith Fowlie, Sylvan Herskowitz, Koichiro Ito, Andrew Lee, Karen Palmer, Paula Pedro, Louis Preonas, Jim Sallee, Reed Walker, Matt Woerman, Catherine Wolfram, and seminar participants at the Energy Institute at Haas Energy Camp for their thoughtful and helpful comments. I benefited from conversations with many PG&E employees for assistance with program implementation details. The online appendix can be found at http://www.joshua.blonz.com/s/Joshua_Blonz_Best_of_the_Second_Best_appendix.pdf.

1 Introduction

Supplying electricity during periods of peak demand is expensive. Because electricity storage is not cost effective, sufficient generation capacity must exist to satisfy demand at all moments in time. To avoid blackouts, electricity providers regularly invest in power plants that operate only on the few highest demand days of the year. Electricity prices, however, do not reflect the high cost of meeting peak demand. Most retail prices reflect the average cost of providing power and do not vary based on when this power is consumed. As a result, most retail electricity customers are undercharged for their electricity at peak times, leading to inefficiently high consumption (Boiteux 1949; Steiner 1957).

In the long run, higher peak consumption necessitates additional generation capacity. In most U.S. electricity markets, capacity investment decisions are made by regulators through the “resource adequacy” process (Joskow and Tirole 2007), where the regulator uses past demand levels to determine generation capacity requirements for electricity providers. If retail prices were adjusted to reflect the full cost of generation during peak periods, this new price would reduce both peak demand and the regulator’s capacity requirement. Borenstein (2005) and Borenstein and Holland (2005) estimate the efficiency loss due to flat retail prices to be 5 to 10 percent of wholesale electricity costs.

The first-best policy is to charge a price that reflects the short-run marginal scarcity value during periods of peak demand. In the case of electricity, this policy is “real-time pricing,” under which the retail price changes hourly or more frequently to reflect the time-varying marginal cost of supply. Real-time pricing is technologically feasible at low cost for most commercial and industrial customers because of the wide-scale deployment of smart meters (Joskow and Wolfram 2012).¹ Despite its large potential benefits, however, real-time pricing remains politically infeasible. Because many customers receive large cross-subsidies under existing flat-pricing schemes relative to real-time pricing, mandatory real-time pricing would be difficult to implement without politically unpopular transfer payments (Borenstein 2007).

The inability to implement real-time pricing suggests two important questions. First, how large are the potential benefits of real-time pricing? These benefits depends on the extent to which customers would respond to short-run price changes. If demand is sufficiently price inelastic, then any potential costs of implementing real-time pricing could outweigh the benefits. Borenstein (2005), however, shows that, for most plausible elasticities, the benefits

¹Smart meter deployment may be financially justified because the meters eliminate the need to pay employees to manually check electricity usage every month. As of 2016, the smart meter penetration for C&I customers in California and the rest of the U.S. was 80 percent and 45 percent, respectively (Energy Information Administration 2016a).

are very likely to outweigh the costs. Second, to what extent can second-best policies achieve the benefits of real-time pricing? This paper addresses the second question by examining a common second-best policy that raises electricity prices on high-demand days and by measuring this program’s effectiveness compared with the first-best, real-time pricing policy.

I study the largest peak demand program to date in the U.S., which includes commercial and industrial (C&I) establishments in the Pacific Gas & Electric (PG&E) Northern California service territory. Programs like PG&E’s “Peak Pricing” are among the most common time-varying pricing policies in the U.S. The popularity of such programs has grown with the recent deployment of advanced smart meter technology. Peak pricing programs also have the potential to facilitate the integration of renewable resources such as wind and solar at lower costs by, for example, charging higher prices when renewable generation quickly changes (e.g., when the sun sets). They also could be used in times of emergency, such as during wildfires, to help balance supply and demand when transmission or generation infrastructure is offline.

The peak pricing implementation I study is focused on reducing peak demand. It gives PG&E the ability to declare up to 15 “event days” per summer, during which the retail electricity price more than triples between 2:00 p.m. and 6:00 p.m. Customers are notified 1 day before each event day, and they receive a small discount on all other summer consumption in exchange for their participation in the program. My analysis focuses on small C&I establishments because the way in which the program was implemented for these customers created a quite similar—and exogenous—control group to which the treated group could be compared.

To identify the impacts of this peak pricing program, I leverage the rules that governed its rollout. PG&E placed establishments on peak pricing by default only after they satisfied a set of non-manipulable eligibility criteria. The establishments are allowed to opt out at any time. I compare establishments that satisfied the eligibility criteria for the first wave of peak pricing to similar establishments that just missed being eligible by not satisfying the eligibility criteria. I provide supporting evidence that assignment to peak pricing is as-good-as-random, using data from before program implementation. I estimate program impacts using a panel fixed effects instrumental variables design with an eligibility instrument.

Using hourly electricity consumption data, I find that peak pricing reduces electricity consumption for non-coastal establishments by 13.5 percent on event days compared with a control group. This estimate is likely a lower bound on the long-run impact of peak pricing because establishments were given “bill protection” during the program window I evaluate. Bill protection guarantees that customers do not pay more than if they were not on peak pricing, suggesting establishments may have exerted less effort toward reducing consumption

knowing they would not be responsible for a higher bill. Even with this backstop in place, establishments in warmer regions significantly reduced their consumption.

I estimate that the program would reduce PG&E peak demand by 118 MW during summer months among small C&I customers when fully implemented by 2018, thereby reducing the need for one or more specialized power plants that are constructed with the sole purpose of generating electricity during the highest demand hours of the year. To evaluate the welfare impacts of peak pricing, I model the regulatory resource adequacy process that governs the amount of capacity built specifically to meet peak demand. I find that the estimated reduction in peak demand increases welfare on the PG&E grid by \$159 million over a 30-year period because of avoided generation capacity investments.

To put these welfare effects in perspective, I compare my estimated peak demand reductions with a first-best real-time price. Using the empirical estimates of demand response, I calculate that the current program recovers 44 percent of the first-best welfare gains. I then consider simple adjustments to the policy to better target the highest demand days and show that substantial welfare gains would likely result from reducing the number of event days and increasing the event-day price. This better-targeted peak pricing policy could achieve 83 percent of first-best welfare gains.

This paper makes two distinct contributions to the economics literature. First, no previous academic research has estimated the impacts of day-ahead peak pricing on the U.S. C&I sector, which is responsible for two-thirds of California and U.S. electricity demand (Energy Information Administration 2016b). The program rollout that I study caused more business establishments to permanently move to peak pricing than any similar program in the U.S. Previous empirical work has focused on temporary peak pricing pilot programs in the residential sector (Bollinger and Hartmann 2019; Burkhardt et al. 2019; Fowlie et al. 2017; Ito et al. 2018; Jessoe and Rapson 2014; Wolak 2007, 2010).

Second, I also contribute to the literature on long-run investment efficiency in electricity markets by estimating the welfare impacts of peak pricing. In particular, I model how demand reductions caused by peak pricing impact the long-run capacity construction decisions in electricity markets—a necessary input to any full welfare analysis of peak pricing. Previous papers studying time-varying electricity pricing have focused on real-time pricing using stylized models that assume prices, demand, and capacity construction can respond instantaneously to balance supply and demand (Borenstein 2005, 2012; Borenstein and Holland 2005; Holland and Mansur 2006). These assumptions preclude the existing models from evaluating the welfare consequences of peak pricing because in practice building additional power plants takes years of planning, which requires grid operators to keep excess capacity in reserve to avoid blackouts.

I develop a new model based on the regulatory mechanisms that actually drive power plant construction decisions in electricity markets, which enables me to estimate welfare impacts under a more realistic set of assumptions. This new model is needed because the existing models are not able to effectively evaluate the impacts of critical peak pricing in the presence of regulatory constraints that exist across most of the U.S.² My model is also able to decompose how welfare depends on specific program design features, allowing me to propose design improvements informed by my empirical estimates.

The rest of the paper is organized as follows: Section 2 discusses the electricity industry, related literature, and the peak pricing program in detail. Section 3 outlines the data used in the analysis. Section 4 describes the empirical strategy, and Section 5 presents results. Section 6 proposes a model for calculating the welfare impacts of peak pricing programs, discusses potential improvements, and benchmarks the outcomes to the first-best, real-time price. Section 7 concludes.

2 Background

Most electricity in the U.S. is sold to retail customers at a constant flat rate that does not reflect the time-varying marginal cost of producing another kilowatt-hour (kWh). In most cases, the marginal cost consists of two components. The first is the short-run marginal production cost (SRMC), which includes the fuel costs associated with producing an additional kWh. The second is due to a regulatory process in most states that requires an electricity supplier to demonstrate it has adequate capacity to meet the peak demand it serves. These resource adequacy requirements are generally based on previous peak demand quantities. As a result, each additional kWh consumed on the highest demand days of the year increases future capacity requirements, adding significant costs.

The welfare costs of the current system of flat-rate pricing and the benefits of real-time pricing are well studied in the economics literature. The research shows that real-time pricing could provide large, long-run efficiency gains compared with the current flat-rate pricing by reducing total quantity demanded (load) during high-demand hours and increasing load when generation costs are low (Holland and Mansur 2006). By reducing peak demand, real-time pricing reduces the need for costly power plants specifically built for the hottest few days of the year (Borenstein 2005, 2012; Borenstein and Holland 2005).

Despite the large potential welfare gains from real-time pricing, implementation is politically challenging. Less than 1 percent of C&I customers in the U.S. are on some form

²My approach complements the work of Boomhower and Davis (2020), who use capacity market payments to value the benefits of energy efficiency at peak hours.

of real-time pricing (Energy Information Administration 2018).³ In the absence of real-time pricing, policymakers have introduced a number of other policies that pass through some portion of time-varying prices to customers without unexpected volatility. Time-of-Use (TOU) pricing adjusts the price of electricity in a prescribed manner by hour, day, and season, but does not pass through high-price events. For example, a previously flat retail price of \$.20/kWh could be changed to a TOU rate of \$.25/kWh between noon and 6:00 p.m., when demand is generally high, and \$.15/kWh at night. A number of studies have examined how C&I electricity customers respond to TOU prices. Some have found little response to TOU prices (Jesso and Rapson 2015), while others have found small reductions in peak consumption or substitution to off-peak hours when prices are increased (Aigner et al. 1994; Aigner and Hirschberg 1985; Qiu et al. 2018). TOU prices can capture some of the average shape of marginal costs, but they do not adjust when wholesale costs spike on the highest demand days of the year. It is during these highest-demand hours that real-time pricing is the most valuable, and TOU pricing is not able to capture most of the benefits. Borenstein (2005) uses a simulation model to show that TOU captures only 20 percent of the efficiency gains of real-time pricing.

Peak pricing programs, like the one studied in this paper, are designed to address the costs associated with the highest demand days of the year. To date, however, research on consumer response to peak pricing programs has focused on the residential sector. Existing studies find that households reduce their energy consumption when facing high prices during peak hours, though the estimated response magnitude varies across studies and depending on the use of automation technology. There has been no published research to date on peak pricing in the C&I sector.⁴ Because firms are responsible for twice the electricity usage of the residential sector, this paper fills an important gap.

The existing residential peak pricing studies have been run as temporary utility experiments and pilot programs. Fowlie et al. (2017) partnered with the Sacramento Municipal Utility District in California to study the impacts of opt-in versus opt-out peak pricing programs. They find that households in the opt-out program reduce their electricity usage by 13.2 percent during peak pricing events. Households that chose to opt in to peak pricing reduced their usage by 26.5 percent. Other residential peak pricing research has focused on the importance of information and technology in responding to peak pricing. Jesso and

³Form EIA-861, which collects this data, records the number of customers on dynamic pricing but does not break out how many customers are on each type of time-varying pricing. To account for this data limitation in the calculation, I exclude customers from Southern California Edison, which has widespread time-varying pricing but low real-time pricing enrollment.

⁴The existing studies on TOU pricing are not predictive of how an establishment will respond to peak pricing. TOU prices are in effect every day, and the prices are significantly lower than the level used in peak pricing programs.

Rapson (2014) find that providing households with detailed usage data results in substantially larger reductions than just the price alone. Burkhardt et al. (2019) examine a residential field experiment in Texas and find a 14 percent reduction from peak pricing, with 74 percent of the response coming from air conditioning. Bollinger and Hartmann (2019) investigate how automation technology that adjusts consumption in response to higher prices affects the response to peak pricing. They find that households are more than twice as responsive when they are given automation technology and higher prices along with price information, compared with price information alone. My paper is the first to investigate whether a similar overall response to peak prices is also found among U.S. commercial and industrial customers.

2.1 PG&E’s Peak Pricing Program

The PG&E peak pricing program for small C&I customers raises the price of electricity from the normal price of \$.25/kWh to \$.85/kWh from 2:00 p.m. to 6:00 p.m. on 9 to 15 “event days” per year. The program runs between June 1st and October 31st each year. Enrolled establishments receive a discount of \$.01/kWh on all other consumption during the summer to compensate them for participating. PG&E determines when event days are called based on day-ahead weather forecasts.⁵ Establishments are notified by 2:00 p.m. the day before an event via e-mail, text message, and/or phone call. Establishments are told about Monday event days on the prior Friday.

Establishments are given bill protection for the first summer they are enrolled. This protection guarantees that customers do not pay more in their first summer as a consequence of the peak pricing rates. If their total utility bill is higher between June 1st and October 31st on peak pricing than it would have been if they had opted out, the customer is refunded the difference. PG&E sent establishments a letter in November 2015 informing of them of how much money they saved or would have lost during the first year of the program. The letter explained that the bill protection credit would be dispensed on their November 2015 bill and that they would no longer receive bill protection going forward. In Section 5, I discuss the potential impact of bill protection on the estimates of price response.

The enrollment data suggest that customers will remain in the peak pricing program after they no longer have bill protection. In the first summer of peak pricing, 89 percent of establishments in my sample would have lost money if not for bill protection. The average loss for these establishments was \$104 over the summer of 2015.⁶ Despite these losses, only an

⁵When the forecasted maximum temperature at a set of five specified weather stations exceeds a given “trigger” temperature, an event day is called. See Appendix Section A for specific details on this process.

⁶2015 was the first year that small C&I establishments were included in peak pricing. The program is designed to be revenue neutral with respect to enrollees, suggesting that the \$.01/kWh subsidy for non-event hours may need to be increased in future years.

additional 5.5 percent of establishments dropped out between their bill protection payment in November 2015 and the most recent data from October 2016. This low percentage suggests that even after the first summer, when establishments no longer have bill protection, they do not choose to leave the peak pricing program.

Unlike the existing peak pricing research that has focused on short-term pilots, the PG&E program is a permanent policy where all C&I customers in the service territory will eventually be enrolled. Establishments will remain on peak pricing indefinitely unless they opt out. I study the first wave of enrollments in which PG&E placed 29 percent of small C&I establishments on peak pricing. Enrolled establishments could opt out at any time using a simple web interface.⁷ Only 5.9 percent of the establishments in the first wave opted out before the first summer. An additional 5.3 percent of establishments dropped out during the first summer of the program. There was no pattern to the establishments that opted out, with both high and low peak usage establishments opting out at similar rates. The high number of people remaining in the program reflects both the role of default bias and the impact of bill protection. A large economics literature documents the impact that changing the default can have on choice, including Abadie and Gay (2006), Choi et al. (2004), Johnson et al. (2002), and Madrian and Shea (2001), among many others.

3 Data

I use confidential data provided by PG&E for this analysis. The data consist of hourly electricity usage data for 19,071 establishments for the summers of 2014 and 2015. These establishments are used in the analysis because their smart meter data started within 6 months of September 1, 2011, which is a key feature of the identification strategy and is described in the following section. I classify establishments in the sample as being in coastal or inland areas based on a PG&E designation. This classification is used frequently in my analysis of peak pricing because the two regions have vastly different climates. The coastal region, which runs the length of the coast in PG&E's service territory, has much milder summers compared to the inland region.⁸

To construct the final data set, I combine hourly usage data with establishment characteristics. Exact establishment latitude and longitude coordinates were provided by

⁷See Appendix Figure A1 for an example of the letter sent to establishments 30 days before the program started, with directions on how to opt out.

⁸Appendix Section B describes the creation of the dataset in detail. See Appendix Figure A2 for a map of the 7,435 establishments used in the primary specification and their region designation.

PG&E and are used to match establishments to hourly weather data obtained from Mesowest.⁹ I observe when a customer was placed on the opt-out tariff and whether they decided to opt out. I also observe industry classification in the form of North American Industry Classification System (NAICS) codes for 89.2 percent of establishments in the sample.

PG&E categorizes its C&I customers based on electricity consumption. This paper focuses on the smallest nonresidential PG&E rate, the A-1 tariff, because the peak pricing rollout for this group allows me to causally identify program impacts.¹⁰ I remove smaller individual meters that consumed below 800 kWh/month in the summer of 2014.¹¹ This filter leaves me with the 19,071 establishments used in the analysis. The average customer in the sample consumed 87 kWh/day and spent \$560/month on electricity in the summer of 2014. This amount is larger than the average residential household, which consumed 21 kWh/day. Figure 1 shows the average summertime hourly consumption profile of the establishments in the sample where the vertical lines indicate the peak pricing window.

There are approximately 283,000 C&I customers of this size profile in the PG&E service territory. These establishments make up 82 percent of the load of the small C&I class. In total, small C&I customers constitute about 2,000 MW of peak load, which is around one-tenth of PG&E's total peak load. The customers in my sample are typically smaller businesses for which energy is not a major input, including restaurants, barber shops, bakeries, corner stores, small retail shops, strip mall storefronts, law offices, and doctors' offices. Energy intensive establishments from industries such as food processing, cement manufacturing, aluminum smelting, or commercial establishments with large refrigeration needs are on different tariffs and are not studied because they face different electricity prices and event-day prices.¹²

4 Empirical Strategy

4.1 Natural Experiment in Peak Pricing Enrollment

The nature of PG&E peak pricing does not permit the use of an OLS selection-on-observables design to carry out a simple comparison between enrolled customers and those

⁹The hourly weather station data were cleaned to remove any weather stations with unreliable data and are matched to the closest establishment. The final dataset contained measurements from 297 weather stations over 2014 and 2015.

¹⁰Establishments are placed on the A-1 tariff if they consume less than 150,000 kWh/year and have peak usage of less than 75 kW. The average PG&E residential customer consumes around 8,000 kWh/year. PG&E imposes a demand charge for its larger non-residential customers. This charge is based on the customer's maximum flow of electricity in a given month. A-1 establishments do not pay a demand charge.

¹¹I drop low-usage meters because most are not associated with an establishment. A full accounting of how the final dataset was constructed and cleaned is provided in Appendix Section B.

¹²Most larger establishments were moved to peak prices using different criteria before 2015. As a result of how this was done, there is no way to reliably identify the impacts of peak pricing on their usage.

yet to be enrolled. Peak pricing is an opt-out program, allowing any customer to remove themselves from the program at any time. Even with the low opt-out rate, comparing enrolled customers to those not on peak pricing would likely result in a comparison between dissimilar establishments and therefore biased estimates of program impacts. To avoid potential bias, I use an instrument that leverages a natural experiment in the rollout of peak pricing for the summer of 2015.

PG&E used a set of rules to determine when an establishment would be placed on peak pricing. They evaluated their customer base once per year starting in November 2014 to determine which establishments were eligible. This paper examines the first wave of this rollout. The regulator required that an establishment had a history of high-frequency metering data before PG&E placed them on peak pricing so that customers could be informed about the potential price impacts and could make informed decisions.

Specifically, *establishments' smart meter data needed to have started before September 1, 2011, to be eligible for the 2015 rollover to peak pricing.* Figure 2 provides a timeline of this process. I classify establishments in two groups: those that were eligible for peak pricing in 2015 because their smart meter data started before September 1, 2011, and those that were not. PG&E deemed those establishments with high-frequency data starting after September 1, 2011, as ineligible to be placed on peak pricing in 2015. These ineligible establishments are the control group in my empirical strategy.

The impacts of this eligibility status can be seen more than three years after the September 1, 2011 threshold, when treatment started.¹³ In November 2014, PG&E moved a portion of the eligible establishments to peak pricing.¹⁴ In contrast, PG&E did not move any of the ineligible establishments, which had to wait for subsequent rollovers.

To illustrate the transition to peak pricing, Figure 3 breaks down the eligible and ineligible groups by the week their smart meter data were first collected. The horizontal axis shows weeks relative to the September 1, 2011 cutoff. The vertical axis displays the percent of each bin that was placed on peak pricing for the summer of 2015. PG&E moved a portion of the establishments to the left of the September 1, 2011 threshold to peak pricing, while no establishments to the right were moved.¹⁵

¹³The long time lag was due to a number of requirements that the regulator had given PG&E about the information that had to be available to an establishment before it was transitioned to peak pricing. See Appendix Section C for more details on these requirements.

¹⁴Which establishments were moved depends on technological factors, which are described later in this section.

¹⁵The proportion of establishments that were moved to peak pricing by PG&E in Figure 3 does not reflect any opt-out decision made by establishments. Establishments made their choice to opt out after PG&E placed them on peak pricing.

The date an establishment’s smart meter data began is based on when its smart meter was installed. PG&E started installing smart meters in 2008, long before planning began for the peak pricing program rollout. PG&E treated installations as general capital upgrades, with installation decisions based on factors such as labor availability and logistical constraints. Installations typically took 5 to 15 minutes and did not require the account holder to be present. The smart meter installation date was not related to consumption or to any observable characteristics of a given establishment.¹⁶

The nature of the smart meter rollout suggests that establishments on either side of the September 1, 2011 threshold are similar. The peak pricing eligibility cutoff was not known when the smart meters were installed, suggesting that establishments had no reason to strategically adjust their installation date. While the installations are as good as random over short periods of time, there are longer-term patterns to consider. PG&E installed smart meters across California during this time period, but certain areas of the state were emphasized earlier in the rollout compared to others. I select customers within an eight-week bandwidth of the September 1, 2011 threshold to avoid potential bias from long-term installation trends. This bandwidth is indicated as the dashed vertical lines in Figure 3 and cuts the sample to 7,435 establishments.¹⁷ Table 1 shows the summary statistics for a number of characteristics broken out by peak pricing eligibility. The table shows that establishments within eight weeks of the September 1, 2011 cutoff are observationally similar to each other.

One noteworthy feature of the September 1, 2011 cutoff is that eligible establishments closer to the threshold were less likely to be rolled over. This pattern is due to technical requirements that govern when high-frequency usage data is considered usable. PG&E requires that the “remote meter reads become stable and reliable for billing purposes” before they can be used for any official purpose (Pacific Gas & Electric 2010).¹⁸ The stability of the remote meter reads is not related to the consumption patterns of the establishment and only relates to quality of the data transmitted from the smart meter to the receiver. The validation process can be quick for some establishments but can take a number of months to complete for others. For this reason, establishments that had high-frequency data for longer (farther to the left in Figure 3) are more likely to be placed on peak pricing in the summer

¹⁶See Appendix Section C.1 for more details on the smart meter rollout, including quotes from annual reports describing the process.

¹⁷I consider alternate bandwidths in the results section as robustness checks.

¹⁸PG&E still sends employees to physically read the meters monthly until this validation is complete. See Appendix Section C.1 for more details on the validation process. Some establishments may not have reliable smart meter data during the validation process, but this has no impact on the analysis because it happened over two years before the study period.

of 2015.¹⁹ The final set of establishments that are placed on peak pricing both had their smart meter installed and their data satisfy the technical data stability requirements before September 1, 2011. The eligible establishments that missed peak pricing in the summer of 2015 due to technical requirements were scheduled to be moved over for the summer of 2016.²⁰

To identify peak pricing program impacts, I instrument for peak pricing participation with whether an establishment’s smart meter was installed before September 1, 2011, limiting my sample to establishments getting smart meters within eight weeks. This approach uses establishment fixed effects to control for time-invariant characteristics. Establishments that PG&E placed on peak pricing and opted out before or during the first summer remain in the data. Their treatment status reflects their enrollment status at each peak pricing event. The unit of observation is the establishment-hour. In most specifications, I limit the sample to 2:00 p.m. to 6:00 p.m. on event days in the summer of 2014 and 2015. In the summer of 2014, PG&E called event days using the same methodology, but they did not apply to the establishments I study. This makes the summer of 2014 an ideal set of pre-period control days that are similar to the 2015 event days.

4.2 Empirical Approach

I estimate the impact of peak pricing using the following two equations via 2SLS:

$$Q_{it} = \beta_1 \widehat{Peak}_{it} + \beta_2 Temp_{it} + \beta_3 Temp_{it}^2 + \zeta_t + \gamma_{ihd} + \epsilon_{it} \quad (1)$$

$$Peak_{it} = \alpha_1 \{Eligible \times Post\}_{it} + \alpha_2 Temp_{it} + \alpha_3 Temp_{it}^2 + \zeta_t + \gamma_{ihd} + \eta_{it} \quad (2)$$

Equation (1) is the second stage. \widehat{Peak}_{it} is an indicator of whether establishment i is on peak pricing in hour-of-sample t , which is predicted in the first stage regression (Equation 2) using the eligibility instrument ($Eligible$) interacted with the 2015 dummy ($Post$).

Q_{it} is the log of electricity consumption for establishment i in hour-of-sample t . Hourly temperature is controlled for with $Temp_{it}$ and $Temp_{it}^2$.²¹ Hour-of-sample fixed effects, which control for any contemporaneous shocks that affect all establishments, are captured with

¹⁹Engineers at PG&E confirmed that the data stability requirement generates the downward sloping pattern seen in Figure 3 independent of the location of the cutoff. This observation suggests that the control establishments would have experienced the same pattern of enrollment and that the research design compares similar establishments on either side of the threshold. There is not data available on when a meter satisfied the stability requirement.

²⁰I do not have access to current data to verify which customers were moved over in subsequent waves.

²¹The temperature controls are used to increase precision, but the results are robust to their omission.

ζ_t . γ_{ihd} is a set of establishment fixed effects that control for time-invariant factors. Each establishment has a separate establishment fixed effect for each hour of day (h) and day of week (d) combination because these are both significant dimensions across which establishments change their energy consumption. β_1 is the coefficient of interest and represents the average hourly reduction across peak event hours for treated establishments in 2015 compared with the control group. The identifying variation comes from within-establishment variation in peak electricity consumption following the implementation of the peak pricing program in 2015.

ϵ_{it} is the error term in the second stage, and η_{it} is the error term from the first stage. The panel nature of this analysis makes each of the errors potentially correlated both over time and across establishments. To account for this two-way errors dependence, I two-way cluster at the establishment and hour-of-sample level, as suggested by Cameron et al. (2011). As a result, the errors are robust to both within-establishment and within-hour-of-sample correlation.

The identifying assumption underlying the 2SLS estimation is that peak pricing eligibility is not correlated with peak electricity consumption, conditional on fixed effects and temperature controls, through any other mechanism than being placed on peak pricing. Formally, this is written as $\text{cov}(Peak\ Eligibility_{it}, \epsilon_{it} | X_{it}) = 0$, where X_{it} represents the covariates and fixed effects that are controlled for in Equation (1). The exclusion restriction could be violated if there are time-varying trends that differentially affect establishments in the two eligibility groups.²² The estimation also requires a valid first stage, for which I provide evidence in Section 5.1.

Evidence of the validity of the research design restriction is provided in Figure 4, which shows the average summer 2014 (pre-period) consumption by eligibility group on event days after controlling for establishment-level fixed effects. The figure shows similar consumption patterns on event days in the year before these establishments are on peak pricing, indicating that the eligible and ineligible establishments function as good comparison groups conditional on the establishment fixed effects like those used in Equation (1).²³ Table 1 shows summary statistics by eligibility group for establishments in the eight-week bandwidth on either side of

²²I also employ a regression discontinuity approach, which explicitly controls for an establishment’s distance in days from the September 1, 2011 discontinuity in the post period by using a trend line. The regression discontinuity specification is described in Appendix Section C.2. I prefer the approach in Equation (1), because the treatment effect is estimated across the entire 8-week sample, instead of focusing on establishments close to the September 1, 2011 cutoff.

²³The comparison shown in Figure 4 is the relevant comparison for the validity of Equation (1) because it reflects the remaining variation after the establishment fixed effects are absorbed. Appendix Figure A4 shows the comparison using raw electricity usage and shows that the pre-period consumption patterns are similar, particularly between 2:00 pm and 6:00 pm.

the September 1, 2011 threshold. I cannot reject that eligible and ineligible establishments are statistically the same across all observables.

5 Results

5.1 Main Results

Table 2 shows the first stage results from estimating Equation (2). Column (1) shows the results for the sample that spans the PG&E service territory. The first stage has a value of .223 and is significant with an F-statistic of 406, providing evidence of a valid first stage. Columns (2) and (3) report the first stages for the coastal and inland regions separately. The results show a significant impact of eligibility on peak pricing enrollment for all specifications.

Table 3 shows impacts of peak pricing on electricity consumption using the primary specification. The sample for analysis comprises the 7,435 establishments with high-frequency data starting within eight weeks of the September 1, 2011 cutoff. Column (1) shows that peak pricing causes a small reduction in consumption between 2:00 p.m. and 6:00 p.m. on event days but with a p-value of .10. Columns (2) and (3) split the results by region, showing that the impact of peak prices varies substantially by geography and temperature. Coastal regions, which are characterized by lower electricity usage and temperatures, show almost no response to peak prices. In contrast, inland establishments reduce peak usage by 13.5 percent, which is significant at the 5 percent level.²⁴ This corresponds to an elasticity of -.119 with a p-value of .002 when the same regression is estimated using log price instead of the peak pricing indicator. The results provide evidence that, in the warmer inland regions of California, peak pricing significantly impacts electricity usage. Coastal customers, however, do not seem to be as responsive. The regional nature of the results is consistent with Ito (2015), who finds that inland households are more price elastic than coastal customers. The size of the reduction for inland customers is similar to the 13.2 percent reduction in peak demand that Fowlie et al. (2017) find in a residential opt-out peak pricing program in a neighboring region with a similar climate.²⁵ My inland establishment elasticity estimate is slightly larger than the medium-long run price elasticity estimate of -0.088 for residential

²⁴Percent reductions reflect antilog transformed coefficients. See Appendix Section D.1 for the regression discontinuity specification results. It yields larger impacts of peak pricing on usage, but I cannot reject that the two results are statistically the same. Appendix Table A9 shows the intent-to-treat results of the 2SLS results estimated in Table 3.

²⁵The peak pricing program in Fowlie et al. (2017) is similar, but the utility they study charges a peak price that is \$.10/kWh lower.

households in Southern California found in Ito (2014). While not directly comparable, my findings suggest that C&I customers have similar responses to residential customers.²⁶

The role of bill protection is important to consider when interpreting the results in this paper. Establishments know they cannot lose money in the first year of the program. This creates incentives similar to those in Ito (2015) where establishments, far from making money under the program, may choose to “give up,” take the bill protection, and not respond to the price. The role of bill protection can be seen by examining the financial impacts of peak pricing in its first year. Only 11 percent of establishments in my sample saved money in the 2015 peak pricing program, with the remainder receiving the help of bill protection. As discussed in Section 2.1, only 5.5 percent of establishments dropped out between the time they received the bill protection credit in November 2015 and the end of the second year of the peak pricing program. The low dropout rate after most establishments would have lost money in the first year, combined with the lack of bill protection in future years, suggest that my results are a lower bound for future peak pricing impacts. If establishments are exposed to potential monetary losses, they have a larger incentive to reduce their usage. It is possible that additional establishments may opt out of peak pricing after losing money, which could reduce future aggregate impacts. However, the low observed opt-out rate after the first summer suggests that this impact may not be very large. Future years of program data are necessary to determine the impact that opt-out behavior might have on program impacts.²⁷

The results should be interpreted as local average treatment effects around the September 1st, 2011 cutoff. It is possible that the average treatment effect across all small C&I customers could be smaller or larger than those found here.²⁸ Observationally, the establishments in the eight-week bandwidth are similar to those in a 27-week bandwidth, which is my complete sample. Appendix Figure A8 compares the pre-period consumption of these two groups and shows similar load shapes.²⁹

The primary specification uses an eight-week bandwidth around the September 1, 2011 cutoff, but the results do not change substantially at different bandwidths, as shown in Panel A of Appendix Figure A7. The results in this section are robust to a number of other specification and clustering choices, as shown in Appendix Section D. Appendix Section D.2 shows the non-instrumented OLS results, which finds a smaller impact of peak pricing.

²⁶In the long run, C&I establishments may be more elastic if they invest in ways to further respond to peak pricing.

²⁷I am not able to estimate the causal impact of peak pricing without bill protection in future years because the control group used in my identifications strategy will have rolled onto peak pricing.

²⁸It is also not possible to know how the results would compare to a mandatory peak pricing program. However, given the low opt-out rates, it likely would be similar to the opt-out program I study.

²⁹I do not estimate treatment effects for customers outside the 8 week bandwidth because the longer-term geographic patterns in meter installations could introduce bias.

5.2 Spillovers to Non-Event Hours

The analysis to this point has only focused on the change in usage between 2:00 p.m. and 6:00 p.m. on event days. This time frame ignores the scope for establishments to re-optimize their usage during off-peak hours. Figure 5 shows the treatment effects for inland establishments by hour of day. The results suggest that establishments begin to reduce their energy usage around 11:00 a.m., with the reductions becoming statistically significant by 1:00 p.m. This pattern of reductions is consistent with establishments making event-day changes that spill over to non-event window hours. For example, an establishment may adjust its air conditioner set point from the normal 72 degrees up to 76 degrees on event days. This behavior would reduce the overall demand for cooling on event days, leading to the reductions seen before 2:00 pm. Immediately after the event window ends, usage returns to the level of the control group. Many small C&I businesses close around 6:00 pm, which might explain the return to control consumption levels.

It is also important to consider if peak pricing impacts usage patterns on non-event days. Fowlie et al. (2017) find significant reductions during the 4:00 p.m. to 7:00 p.m. event window period on non-event days for residential customers. The authors suggest that this could be due to habit formation, learned preferences, or a fixed adjustment cost. Table 4 reports the results from estimating Equation (2) on summer non-event weekdays between 2:00 p.m. and 6:00 p.m. The results show a small reduction in the inland region that is significant at the 10 percent level and is less than half the event day impact.³⁰ The lack of large and significant responses suggest that the peak pricing program is not having a noteworthy impact on non-event day electricity consumption between 2:00 p.m. and 6:00 p.m.

5.3 Impacts of Temperature

The outdoor temperature on event days is much higher in the inland regions of California than on the coast.³¹ This difference suggests that temperature could play a role in an establishment's demand elasticity. Reiss and White (2005) show that residential customers with air conditioners have more elastic demand than those without. Ideally, I would measure the impacts of peak pricing on establishments with air conditioning separately from those without, but this is not possible with the data available. Instead, I focus on the role that temperature plays in event-day reductions.

³⁰It is important to note that these insignificant impacts are estimated from 95 non-event days in 2015, which is a much larger sample than the 15 event days in the main analysis.

³¹See Appendix Figure A3 for a map showing temperatures on event days.

Table 5 presents the results from interacting temperature with the treatment effect in Equation (1). Temperature does not appear to play a role in how coastal establishments respond to peak pricing. For inland establishments, the interaction term in Column (3) is negative and significant, indicating that peak reductions get larger as temperature increases. The estimated impacts are relative to a 75 degree day.³² The results show that, on average, higher reductions for inland establishments come from higher outdoor temperatures. This finding is significant because the highest demand days of the summer typically fall on the hottest days of the summer. As a consequence, the peak pricing program may provide larger reductions on the hottest event days when the grid is most stressed. This finding is relevant to program design because, if event days occurred only on the hottest few summer days, then reductions might be higher than the average impacts measured in Section 5.1.³³

5.4 Firm Heterogeneity

Small C&I establishments use electricity to produce a wide range of goods and services in their day-to-day operations. For example, retail establishments have different patterns of electricity usage than office spaces or doctors' offices (Kahn et al. 2014). In this subsection, I use the industry classification information provided by PG&E to test how different types of establishments respond to peak pricing. Specifically, I test how customer-facing and non-customer-facing establishments each respond to peak pricing. I hypothesize that customer-facing businesses such as retail establishments or movie theaters may be less likely to reduce air conditioning usage if it affects business. Customers may choose a different movie theater or store if the indoor temperature is above expectations. On the other hand, non-customer-facing establishments, such as office spaces, may be more willing to reduce peak usage if it is easier for employees to adapt. For example, an employer could inform their staff of an event day in advance and encourage them to dress for a warm office.

I classify establishments as customer-facing or non-customer-facing using the first two digits of their North American Industry Classification System (NAICS) industry code. To determine which two-digit industries are customer-facing, I use the U.S. Bureau of Labor Statistics classification of service-providing industries.³⁴ From this list, I define the set of service industries that are customer-focused. This list includes retail trade (NAICS 44-45), health care (NAICS 62), leisure and hospitality (NAICS 71), and accommodation and food services (NAICS 72). All other NAICS codes are classified as non-customer-facing.

³²I re-center temperature at 75 degrees for ease of interpretation; this does not impact the peak pricing times temperature coefficient.

³³I do not estimate treatment effects for individual event days because the sample is not large enough to reliably estimate these impacts.

³⁴<http://www.bls.gov/iag/tgs/iag07.htm>.

These include industries such as goods manufacturing (NAICS 11-31), transportation and warehousing (NAICS 48-49), and office spaces (NAICS 52-56).³⁵

Table 6 shows the results separately for customer-facing and non-customer-facing industries. In all cases, the customer-facing industries do not show a significant response to peak pricing. This result is in contrast to the non-customer-facing industries, where the impacts are larger than previously found when considering all industries together in Table 3. Inland customer-facing establishments show the largest response to peak pricing, reducing their peak usage by 18.0 percent.

It is not possible to identify how establishments are reducing their peak demand. However, the large response in non-customer facing establishments in warmer regions suggests that establishments are reducing peak demand by using their air conditioners less on event days. This is supported by the findings in Section 5.3 that show event-day reductions increase with temperature because reducing air conditioner usage can provide relatively larger savings on the hottest days. The results also show that a similar peak pricing program will likely be effective in warmer regions of the country with high air conditioning penetration. In other states where peak pricing is not structured as an opt-out program, it may be optimal to target non-customer-facing establishments for enrollment to generate the largest program impacts.

5.5 Aggregate Impacts

The previous subsections estimated the impacts of peak pricing on a subset of small C&I establishments in the PG&E service territory. Importantly, these customers are part of a utility-wide rollout that will place all small C&I establishments on peak pricing by 2018, which has the potential to generate large peak reductions.

To better understand the impacts of the fully deployed peak pricing program, I extrapolate my savings to all small C&I customers. There are three main assumptions that I make for this calculation. First, as discussed in Section 5.1, the estimate is a local average treatment effect. It is not possible to prove that the local average treatment effects estimated in the previous sections reflect the behavior of all small C&I establishments in California, but the estimates are the best available and are useful for back-of-the-envelope calculations.

Second, these estimates capture only the short-run impacts of peak pricing in the summer of 2015. It is possible that establishment demand will become more elastic as peak pricing continues. For example, customer-facing establishments may be able to reduce

³⁵The NAICS codes that I have are often imprecise, which limits the ability to finely cut the data into many different industries. See Appendix Table A3 for a breakdown of establishments by two-digit NAICS code.

peak consumption by upgrading their air conditioners to more efficient models or improving insulation. Third, I assume the estimated savings reflect future program year savings when there is no bill protection. This assumption could result in my estimates understating aggregate impacts, as discussed in Section 5.1.

I extend the savings estimates for inland customers from Column (3) of Table 3. I focus on the inland establishments and assume that coastal establishments in the PG&E service territory will not provide any peak reductions. I assume that the establishments in the 8-week bandwidth are representative of all small inland C&I customers (e.g., the proportion of customer-facing versus non-customer-facing establishments is the same), and that long-run opt-out rates will be similar to those observed in the first two years of my sample. I combine this with customer-count information provided by PG&E to estimate the projected total impacts when the program is fully rolled out by the summer of 2018.³⁶ Using this technique, I find that small C&I establishments will, on average, provide reductions of 118 MW in peak load on peak pricing event days.

5.6 Environmental impact

The peak pricing program is designed to improve grid efficiency, but it also affects power plant emissions. Holland and Mansur (2008) found that real-time pricing has a relatively small impact on emissions. The authors highlight that the change in emissions depends on which type of generation was previously meeting peak demand. In California, peak demand is typically satisfied by natural gas turbine generators, which have a moderate CO₂ emissions rate and low SO₂ and NO_x emissions rates. To better understand the magnitude of the impacts of peak pricing on CO₂ emissions, I conduct a simple back-of-the-envelope calculation using my estimates of the impact of peak pricing on demand along with estimates from the California Energy Commission on emissions rates (California Energy Commission 2015). I find that the peak pricing program will cause a reduction of around 4,000 metric tons of CO₂ per summer when the policy is fully implemented.³⁷ The calculation also assumes that the peak pricing program will not increase consumption during non-event hours. The annual reduction in CO₂ emissions, while not insignificant, is only .11 percent of California’s daily electricity sector emissions (California Air Resources Board 2019). Using a \$50/ton social

³⁶A full accounting of the assumptions and calculations can be found in Appendix section E.1.

³⁷This is a relatively small CO₂ reduction compared to California’s total emissions of 440 million metric tons in 2015 (California Air Resources Board 2019). This calculation of the peak pricing program’s effect is made using the CO₂ emissions rate of 1,239 lbs/MWh from a conventional single cycle plant. This is a conservative assumption because the other generation options, such as a combined cycle plant or hydropower, have lower or zero emissions rates. The use of hydropower to meet peak demand, while causing no emissions at the time of generation, has an opportunity cost that likely will lead to non-zero emission impacts in the long run.

cost of carbon, the reductions translate to around \$200,000 per year in benefits (Revesz et al. 2017).³⁸ I include these benefits in the welfare calculations conducted in the subsequent section. Overall, the carbon reduction benefits are a small fraction of the overall value provided by the peak pricing policy.

6 Welfare Impacts of Peak Pricing

In this section, I first introduce a new model to evaluate the welfare impacts of the peak pricing program, which I calibrate using the empirical peak demand reductions from the previous section. Using this approach, I consider changes to the current peak pricing program to better target the long-run investment inefficiencies that result from flat-rate pricing. I find that, by changing when event days are called, and adjusting the event hour price, program outcomes can be greatly improved. I conclude by benchmarking the impacts of peak pricing against the first-best, real-time price using a simple theoretical energy-pricing model.

6.1 A Model of Welfare Impacts from Peak Prices

The model is based on the current regulatory process in California, which is responsible for capacity construction decisions. Most other states follow a similar process. In the model, peak pricing reduces the level of summer peak demand, which in turn reduces long-run capacity requirements and saves costs by avoiding power plant construction. This framework allows me to calculate the welfare benefits of peak pricing in a manner that reflects how capacity construction decisions in electricity markets are made. The existing literature typically calculates the welfare impacts of alternative pricing policies using a “stylized model” approach of electricity prices and power plant construction (Borenstein 2005, 2012; Borenstein and Holland 2005; Holland and Mansur 2006).³⁹ The stylized model provides insight on the welfare impacts of real-time pricing but uses assumptions that do not realistically portray the nature of the binding capacity constraint in electricity markets.

The structure of electricity markets is defined by the lack of cost-effective storage, which requires supply and demand to be balanced in real time. This feature introduces a capacity constraint equal to the total capacity of generators; blackouts will result if demand exceeds this constraint at any time. The stylized models of electricity markets do not consider this

³⁸The impacts on SO_2 and NO_X are small enough that I do not include them in the benefits calculations. For example, I find that peak pricing will reduce NO_X emissions by less than 1 ton per year and SO_2 emissions by .05 tons. California’s carbon emissions are capped, but the cap is not currently binding. As a result, any additional emissions reductions from peak pricing will reduce total emissions. If the cap is binding in future years, the welfare benefits from reducing carbon emissions will be lower.

³⁹I will refer to the general approach used in the literature as the stylized model going forward for simplicity.

constraint and assume that the price and demand for electricity are able to adjust quickly enough to avoid shortfalls.⁴⁰ Additionally, such models assume a cost to build new generators, but not construction time. In practice, it can take six years from the initial proposal for a power plant to begin generating electricity. Much of this process is governed by the regulator that sets the amount of generation capacity that a utility must have in reserve to avoid blackouts. The stylized model approach used in the existing literature does not include the need for reserves because they assume price can increase in an unconstrained manner to clear the market.⁴¹

The assumption in the stylized model that supply and demand can instantaneously balance (including capacity construction decisions) equates the wholesale price to the full long-run cost of providing a kWh of electricity. This setup is central to how papers like Borenstein and Holland (2005) model real-time pricing and estimate benefits.⁴² However, any estimates that use these simplified assumptions may not accurately capture the effects of time-varying pricing in existing electricity markets. In most wholesale electricity markets, capacity construction decisions are not based on generating companies responding to wholesale market prices and building capacity until a zero-profit condition binds, but are based on other mechanisms including resource adequacy requirements where the regulator mandates how much peak generation capacity the utility must have on hand (Joskow and Tirole 2007).⁴³

The regulatory process is an important consideration when evaluating time-varying pricing, because the majority of the welfare benefits of raising prices during periods of peak demand comes from the long-run reduction in power plant construction and maintenance. In particular, it is the reduction of “peaker” power plants, which have a relatively low capital cost but a high marginal cost of generation. Some of these plants run for only a few hours on the hottest day of the year. A large amount of peaker capacity is expensive to build and maintain. The conditions in existing electricity markets, including time-invariant retail pricing and a risk-averse regulatory process, result in the construction of more peak capacity

⁴⁰In the models, this is done by raising wholesale prices high enough that new generators are built to meet demand.

⁴¹Most electricity markets also have price caps that would prevent prices from increasing to the levels necessary to balance supply and demand without reserves. For example, almost half of the scenarios simulated in Borenstein and Holland (2005) have peak prices above the \$1000/MWh price cap in the California market.

⁴²Borenstein and Holland (2005) assume that generation capacity will enter the wholesale market as long as profits are positive.

⁴³Some markets use secondary mechanisms called capacity markets that pay generators to provide and maintain capacity. Capacity markets are a similar tool to resource adequacy requirements, but in many markets are less transparent and harder to evaluate. In many cases, capacity markets are used as a subsidy to generators to make it less expensive for a utility to reach its resource adequacy goal. See Boomhower and Davis (2020) for a discussion of capacity markets and how they can be used to measure the value of generation capacity.

than a stylized model predicts. Existing stylized models are likely to undervalue the benefits of peak pricing by not capturing these critical determinants of peak capacity construction.

I solve this problem by introducing a model that is based on the resource adequacy process that drives most capacity construction decisions.⁴⁴ Typically, the regulator forecasts future peak demand using historical data and load growth projections. Using this forecast and a valuation of blackouts, they set a resource adequacy level for the utility in the coming year. The model I introduce is based on how peak pricing affects the resource adequacy process. The model proceeds in three steps that happen yearly.

In step 1, no peak pricing program has been implemented. The regulator has information about the distribution of historical peak loads and temperatures, which also includes information about the peak load from the previous summer. I denote this information set as L_0 . The regulator uses this information to determine how much peak capacity is needed, using the decision function $F()$, which does not change over time.⁴⁵ I assume this is a well-defined process known to all market participants and that the regulator sets capacity high enough that there will be no generation shortages in the coming year.⁴⁶ I define this peak capacity requirement as $X_1 = F(L_0)$.

In step 2, the utility acquires capacity X_1 at cost $C(X_1)$. I assume that the utility must fulfill this requirement and that the regulator perfectly observes the utility's behavior. The cost function is linear and does not change over time. For simplicity, it reflects the utility's yearly cost to acquire capacity.

Step 3 is when the demand for the year is realized. In the absence of peak prices, demand would have reached a peak level of L_1 . However, with the implementation of peak prices and their corresponding demand reductions, the new load is \widetilde{L}_1 , such that $\widetilde{L}_1 < L_1$. I assume peak pricing reduces peak load on all event days in a summer by the same amount and that this amount remains constant from year to year. This is a conservative assumption with respect to program benefits for two reasons. First, it is likely that establishments will become more elastic in the long run, as they have more time to adjust to the new prices.

⁴⁴The process I model is based on the California resource adequacy process, but is representative of how capacity requirements are set in most regions.

⁴⁵If the regulator were an optimal social planner, the $F()$ decision function would balance the benefits of reliability against the costs of acquiring capacity, and pick an optimal capacity requirement for the utility. In practice, most regulators are risk averse and put a very high cost on supply shortfalls that result in localized blackouts. As a consequence, regulators typically set high reserve requirements for utilities. I do not take a stand on the exact approach the regulator should use. I simply assume they follow the same rule each year.

⁴⁶Blackouts from demand exceeding capacity are rare. The current California process requires capacity at 1.15 times the projected peak load. This level is sufficiently high to assure that capacity limits will never be reached. See Joskow and Tirole (2006) and Joskow and Tirole (2007) for a discussion of optimal capacity with the possibility of rationing. In California and states with restructured electricity markets where the utility does not directly own generation capacity, the regulator requires the utility to contract for sufficient capacity with independent power producers.

Second, Section 5.3 showed that peak pricing reductions are larger at higher temperatures, which typically happen on the highest demand days of the year. This finding suggests that peak pricing will cause the highest reductions when they are the most valuable.

The majority of the benefits from peak pricing come from this reduction in summer peak demand. This can be seen in Panel A of Figure 6, using a simplified peaker capacity supply curve. By reducing the total peak demand, peak pricing reduces the total generation capacity necessary to satisfy demand. This saving in capacity cost reflects the high costs associated with building generation capacity and is the main driver of savings under the peak pricing program.⁴⁷

The second impact from peak pricing is the net consumer surplus loss that results from changes in customer behavior when paying higher prices on event days.⁴⁸ For example, establishments may choose to run their air conditioner less, leading to a less comfortable indoor environment. It may seem counterintuitive to count a reduction in air conditioner usage as a cost during periods when the flat electricity price paid by the establishment is below the cost of providing that electricity. However, the (much larger) benefits from reducing peak demand are captured by the capacity cost savings (shown in Panel A of Figure 6). To properly calculate the full effect of peak pricing, I must separately account for the net consumer surplus loss.

The consumer surplus impact can be seen graphically in Panel B of Figure 6. To calculate this impact, I first recognize that the electricity still sold at the peak price (to the left of Q_1) induces no change in total surplus, as it is just a transfer from consumers to producers. For units that go unsold due to the price increase (to the right of Q_1 and to the left of Q_0), the change in surplus is the area under the demand curve minus the resource savings from not producing these units. In this case, the resource savings are equal to the fuel savings of the peaker plants that would otherwise be used to generate this electricity.

I value the reduction in fuel used to run a peaker plant at its short-run marginal production cost (SRMC), which I assume to be \$.102/kWh, based on current natural gas prices (California Energy Commission 2015). I use the SRMC for this calculation because I assume the regulatory process dictates that sufficient capacity is available at all hours of the year, meaning the surplus losses are net of the short-run costs associated with running a

⁴⁷In cases where there is older excess generation capacity available, the decision faced by regulators may be to prevent retirements instead of building new plants. This would make my estimates an upper bound on the benefits from peak pricing. In many cases older excess capacity is being forced to close for environmental reasons (e.g. the phase out of once-through cooling in California or local politics) or in the case of older technologies like coal is not able to be dispatched quick enough to meet peak demand.

⁴⁸I include the small increase in consumer surplus from the \$.01/kWh discount given in non-peak hours in this calculation to match the structure of the PG&E program. The benefits of the discount are much smaller than the costs of higher prices on event days.

peaker plant.⁴⁹ The existing capacity guaranteed by the resource adequacy process prevents me from using wholesale market prices to make this calculation.⁵⁰

The set of calculations leaves what I term the “net consumer surplus loss,” which is represented by the shaded triangle in Panel B of Figure 6. I use a linear demand curve for simplicity and because it provides a conservative upper bound on the net CS losses compared to other concave alternatives such as modeling demand as constant-elasticity, which I use as a robustness check. I define the net consumer surplus loss in year 1 as CS_1 .

The process now restarts at step 1 in year 2. In the world with peak pricing, the regulator observes peak load \widetilde{L}_1 and sets peak capacity requirements for the coming year $\widetilde{X}_2 = F(\widetilde{L}_1)$. In the non-peak pricing scenario, the regulator observes peak load L_1 , resulting in peak capacity requirement $X_2 > \widetilde{X}_2$. In step 2 of year 2, the utility must acquire capacity at cost $C(\widetilde{X}_2)$ and $C(X_2)$ for the peak and non-peak pricing scenarios, respectively. This process continues and repeats for both scenarios over time.

To calculate the welfare impacts of the peak pricing program, I subtract the costs in the peak pricing scenario from the costs in the non-peak pricing scenario. The benefits are calculated over T periods on N event days per year using discount rate r. The change in welfare from implementing peak pricing is defined as follows:

$$\Delta Welfare = \sum_{t=1}^T \frac{C(X_t) - C(\widetilde{X}_t) - N \times CS_t}{(1+r)^t} \quad (3)$$

This calculation compares the cost of peak generation capacity with standard pricing to the lower peak capacity needs when peak prices are used.

Equation 3 divides the effects of peak pricing into the benefits ($C(X_t) - C(\widetilde{X}_t)$) and the costs ($N \times CS_t$). Calculating the benefits using the resource adequacy process is the main innovation of this model. This approach allows me to estimate the benefits of peak pricing using the real-world process that determines capacity investment. It is applicable to most wholesale electricity markets because some form of resource adequacy process is typically used for capacity planning and single cycle natural gas peaker plants are used to

⁴⁹I assume the SRMC is flat, meaning all peaker plants have the same efficiency. It is possible that if $Q_0 - Q_1$ is large enough, less efficient peakers would be used and the curve could slope upward. This would likely be a relatively small effect. The flat SRMC means that there are no producer surplus impacts of peak pricing.

⁵⁰The price cap and the capacity market also affect the wholesale price, making it lower than the long-run costs of supplying a kWh of electricity.

meet peak demand. The approach could also be applied in vertically integrated electricity service territories where there is no wholesale market price.⁵¹

The cost term leans on a stylized formulation of net consumer surplus losses. The calculation only includes the surplus losses to establishments that occur between 2:00 p.m. and 6:00 p.m. on event days, when prices increase.⁵² It is possible that customers are responding to peak prices in ways that are not reflected in these hours. For example, Section 5.2 shows that peak pricing enrolled establishments reduce their usage before the 2:00 p.m. to 6:00 p.m. event window starts. The model presented here does not capture the net consumer surplus losses associated with this change in behavior before the event window or any other non-event window impacts. Bill protection could also impact the magnitude of the net consumer surplus losses. If the price signal to establishments is affected by bill protection, then the response to peak pricing may not reflect the true net consumer surplus impacts.

6.2 Calculating Welfare Impacts of PG&E’s Peak Pricing Program

In this section, I calculate the welfare impact of the PG&E peak pricing program for small C&I establishments using the model from the previous section and my empirical results. The calculations are summarized in Table 7. Some of the simplified assumptions in the model are adjusted to better reflect the PG&E service territory. In the model, the utility purchases capacity yearly at cost $C(X_t)$. In practice, the peaker plants that are used to satisfy peak demand typically last at least 30 years.

To approximate the cost function, I use the construction cost of a single cycle peaker plant. The California Energy Commission estimates it costs \$1,185,000/MW to build a natural gas combustion turbine peaker plant (California Energy Commission 2015).⁵³ Using these plant construction numbers and my empirical estimates, I find that the peak pricing program would provide a one-time saving of \$140 million in construction costs. I assume this cost savings occurs in year 1 of a 30-year program. To value the total impacts of the program, I include the discounted stream of annual costs and benefits. Reducing peaker

⁵¹The model may be less applicable in markets with significant hydro resources, such as in the Pacific Northwest. Some modifications to how the resource adequacy process is run in other wholesale electricity markets may be needed. In vertically integrated markets, the benefits would reflect how the utility makes capacity investment decisions in conjunction with the regulator.

⁵²The stylized model in the literature uses a similar assumption, but it is not explicitly broken out into its own term.

⁵³All values used in this paper are in 2016 dollars. Original 2011 values are inflated using the IHS North American Power Capital Costs Index.

capacity provides an annual benefit of avoided staffing and maintenance costs, which in this case totals \$3.07 million per year.

To make the CS loss calculation, I use a linear demand curve, as discussed in the previous section. One important difference between the model and PG&E prices is that retail electricity rates for small C&I customers are set at \$.25/kWh. Retail prices are higher than the short-run marginal cost of production because fixed costs are recovered volumetrically in PG&E. In the previous section, I set retail rates at the short-run marginal cost of production, which I assume to be \$.102/kWh for a peaker plant (California Energy Commission 2015). Establishments were willing to pay the \$.25/kWh price for their electricity during these peak periods, meaning economic surplus is lost on event days when prices are increased and consumption is reduced. Graphically, this impact is represented by a rectangle between the \$.25/kWh electricity price and the \$.102/kWh short-run marginal cost of production. The total CS loss from peak pricing is this rectangle plus the triangle under the demand curve, shown in Figure 6.⁵⁴ Using my empirical estimates, I find that the total net consumer surplus loss in 2015 equals \$3.16 million/year.

The PG&E peak pricing program gives enrolled establishments a \$.01/kWh discount on all non-event day electricity consumption. As a result, establishments will consume more electricity in off-peak hours, resulting in increased consumption across almost all summer hours. Importantly, this price reduction is welfare improving because the retail price of electricity for small C&I customers exceeds any reasonable social cost. The benefits of the \$.01/kWh discount are small but add up across all non-event hours during the summer. Using my elasticity estimates and linear demand, I calculate these welfare gains to be \$0.84 million/year.⁵⁵ The benefit of this small price decrease in non-event hours are added to the overall welfare benefits of the policy.

To come up with a total welfare value, I take the construction costs and add on the discounted stream of costs and benefits detailed above. I also include the benefits of reducing CO₂ emissions calculated in Section 5.6, which total \$.2 million/year.⁵⁶ This results in total welfare benefits of \$159 million (2016 dollars) using a 3 percent real discount rate and a 30-year horizon.⁵⁷ These numbers represent the welfare benefits of running the peak pricing program every summer for 30 years. Embedded in this back-of-the-envelope calculation is

⁵⁴Appendix Figure A9 shows the CS loss triangle and rectangle that reflect PG&E prices.

⁵⁵This is a strong assumption because I am applying my demand curve estimates, derived for the period between 2:00 pm and 6:00 pm on event days, to all other hours in the summer. Using the empirical analysis on non-event hours in the summer of 2015, I can reject the level of responsiveness I am using for this calculation. Ultimately, the response from the off-peak CS gains is small and does not significantly impact outcomes.

⁵⁶The local air pollution benefits are not significant and I do not include them in the calculation.

⁵⁷The results are not sensitive to discount rate assumptions because most of the benefit is incurred upfront with the avoidance of capital construction costs. The other annual costs and benefits are roughly offsetting.

the assumption that electricity supply and demand will not change in ways that affect the numbers calculated above. The assumptions needed to estimate the aggregate energy savings from peak pricing in Section 5.5 also apply. I also assume that the operation and maintenance costs stay constant over the life of the plant, which likely understates the costs as the plant ages. Furthermore, establishment demands are likely to become more elastic as they face peak prices over many summers. Bill protection may also lead to an underestimation of the welfare benefits if establishments reduced less in their first year, knowing they would not lose money because of peak pricing.

The above welfare calculations only capture the negative net consumer surplus impacts from peak pricing that occur between 2:00 p.m. and 6:00 p.m. Establishments may undertake behaviors that affect consumption outside of the event window, resulting in welfare impacts that are not captured by this model. These changes in consumption appear to be relatively limited compared to the reductions during an event. Figure 5 shows small reductions in the hours leading up to an event, and Table 4 shows small and insignificant at the 5 percent level reductions in usage on non-event days. For robustness, I consider a scenario where the net consumer surplus impacts are double what is calculated above. Using this assumption, I find the total welfare impacts of the program to be \$96 million. This shows that, even under a conservative set of assumptions, the welfare impacts of peak pricing remain positive. The net consumer surplus losses would need to be 3.4 times larger than what is estimated in the primary specification for the costs of peak pricing to outweigh the benefits, which is unlikely considering the limited consumption changes that are observed outside the event window.

6.3 Targeting the Capacity Constraint

The PG&E peak pricing program is designed in a manner similar to other peak pricing policies around the U.S. The utility has discretion over when to charge higher prices on 9 to 15 event days per summer. In this section, I consider the welfare implications of how event days are chosen and the price charged during event hours. I do this in the context of the current peak pricing program, where resource adequacy requirements guarantee that there will be sufficient capacity available to avoid blackouts.

PG&E calls event days using day-ahead weather forecasts. When the average forecasted temperature for inland California exceeds a trigger temperature of 96 degrees or 98 degrees, an event day is called. The trigger temperature is based on how many event days have been called so far in a given summer and on historical weather trends.⁵⁸ This approach is effective

⁵⁸The trigger temperature is adjusted every 15 days throughout the summer to hit the target number of 12 to 15 event days. See Appendix Section A for more details.

at selecting the top 12 to 15 demand days each summer, but it is not designed to maximize the net benefits of the peak pricing program.

The typical summer in California has a small number of days with very high demand that are responsible for peak load. For example, the difference between the demand on the highest event day and the median event day in 2015 was 1,220 MW, more than 14 percent of total peak load.⁵⁹ The few highest demand days each summer drive resource adequacy requirements and the long-run construction of peaker plants. I define as “super-peak” days the set of days each summer for which calling an event day reduces the total summer peak load. The number of super-peak days each summer depends on both the level of reduction due to peak pricing and the number of high-demand days. Most Northern California summers have between one and three super-peak demand days based on the estimated reduction due to the peak pricing program for small C&I customers.⁶⁰

The role of super-peak days in the 2015 program can be seen in the first two columns of Table 8. To project the impact that the peak pricing program for small C&I establishments will have once it is fully rolled out, I use the aggregate 118 MW reduction that projects outcomes for 2018. This reduction would lower the 2015 summer peak from 19,451 MW to 19,333 MW, which would become the new summer peak \tilde{L}_t . In 2015, no event days other than the one with the highest demand will affect \tilde{L}_t . For example, reducing the load on September 9, 2015, from 19,017 MW to 18,899 MW will not affect \tilde{L}_t and will not provide savings in long-run generation capacity investment.⁶¹

Table 8 shows the 2015 event days with the welfare impacts broken out. These values reflect the welfare impact associated with each event day in 2015, using the ex-post information about the realized demands. In practice, the peak pricing program is based on day-ahead forecasts, which introduces significant uncertainty about which event days will provide benefits when they are called. Column 3 shows the capacity value of reducing peak load. Only the highest demand event day of the summer provided capacity cost savings because none of the other event days affect summer peak load. Column 4 shows the net consumer surplus loss figure of \$209,000 per event day, which is reported in the same discounted manner to allow for easy comparison. Conditional on the super-peak event days being called, the non-super peak event days reduce the welfare impacts of the program without providing capacity cost savings.

⁵⁹The same pattern holds for all years 2010–2015, with the difference between maximum and median peak load of 1,600 MW.

⁶⁰Summers with one unusually hot day will usually have just one super-peak day. Summers with a long heat wave where there are many days close to the peak will have more super-peak days.

⁶¹A list of the top 20 demand days in 2015 can be seen in Appendix Table A4.

The cost of non-super peak event days quickly adds up. Each extra event day that does not provide capacity cost savings results in a loss of \$4.2 million of net consumer surplus over the life of the program. A refinement to the peak pricing program would call just the super-peak event days each summer.⁶² This approach is challenging with event day programs because it is not possible to forecast ex-ante which summer days will be super-peak (Borenstein 2012). Despite this limitation, there are a number of improvements that could be made to the current program using the day-ahead information that is available to PG&E.

One simple change to the peak pricing program is to tighten the criterion used to call an event day. The second-to-last column of Table 8 shows the day-ahead temperature forecasts for the inland region of California. An event day is called when this temperature equals or exceeds the “trigger temperature” set by PG&E, which is shown in the last column.⁶³ The current set of trigger temperatures typically calls the super-peak demand days each summer, but also includes a large number of additional days that do not provide capacity cost savings. A simple adjustment to the peak pricing program would move the trigger temperature to 101 degrees and remove the current 9 days per summer minimum. This approach uses the same day-ahead temperature forecast that PG&E currently uses to pick event days. It would result in a program that is better targeted at the super-peak event days, and would result in fewer low-demand event days each year.

In an electricity market with regulated resource adequacy requirements, the impacts of missing a super-peak event day are a higher summer peak \tilde{L}_t and the costs of building capacity in a future period. In most cases, the welfare loss from missing a super-peak event day is much higher than the benefit of avoiding a non-super peak event day. Any day-ahead program must take this tradeoff into account. The proposed 101 degree trigger temperature accurately selects the super-peak event days over the last five years using day-ahead temperature forecast data, but a different trigger may be preferred in the future. For example, a utility may prefer a 99 degree trigger if they are worried about the accuracy of their day-ahead temperature projections. I considered more complicated models using load forecasts for event-day targeting, but this adds unnecessary complexity without additional insight.

The second dimension of the peak pricing program that could be adjusted is the level of the event day price. Currently, small C&I establishments pay \$.85/kWh during event windows, which is \$.60 higher than their typical rate. Wholesale prices are routinely above \$.85/kWh, and the peak price for large C&I PG&E establishments is set at \$1.35/kWh. This level is designed to reflect the long-run value of capacity (not just the resource cost to run existing

⁶²It may be useful to set a minimum number of event days so that establishments do not forget they are on the program. I have not found any research that identifies the impact of using too few event days.

⁶³See Appendix Section A for more details on trigger temperatures.

plants) and is based on the regulator’s avoided cost of capacity (California Public Utilities Commission 2001).⁶⁴ There is no reason to charge different event-day prices to different customer classes, because both are subject to the same capacity constraint that drives system costs. If \$.85/kWh is below the efficient wholesale cost of electricity on event days, there are potential welfare gains from raising the event-day price for small C&I establishments.

I quantify the welfare benefits of changing the number of event days and level of peak prices in Table 9. Panel A uses a linear demand curve to calculate the demand at higher prices, and Panel B uses a constant elasticity demand curve. The current peak pricing program has an event price of \$.85/kWh on 15 days per year and is shown in the top-left entry of each panel. Column (2) shows outcomes if the small C&I peak price were raised to \$1.35/kWh, the level paid by large establishments.⁶⁵ Panel A shows that using the current 15 event days per summer and increasing the event price from \$.85/kWh to \$1.35/kWh would increase the welfare benefits from \$159 million to \$212 million using a linear demand curve. The third column shows the impacts of a peak price set at \$1.85/kWh.⁶⁶ Moving down the table decreases the number of event days per summer from 15 to 8 to just the 3 super-peak days.⁶⁷ The table shows that moving to a 101 degree trigger and using the large C&I peak price of \$1.35/kWh—both of which are realistic adjustments—could improve program outcomes by 87 percent.

Panel B of Table 9 reproduces the analysis in Panel A using a constant elasticity demand curve.⁶⁸ It shows small differences in the welfare benefits compared to the main results that are consistent with the properties of a constant elasticity demand curve relative to a linear demand curve. At the current price of \$.85/kWh, the concavity of the constant elasticity demand curve results in lower net consumer surplus losses, leading to higher estimated welfare benefits. At peak price levels above \$.85/kWh, the constant elasticity demand curve predicts lower peak demand reductions than the linear curve, yielding lower welfare benefits. For

⁶⁴An alternate trigger could use day-ahead market prices to determine when to call event days and set the price level. In theory, it could improve targeting and outcomes. However, in practice, it is not a good predictor of peak demand because of the regulatory rules that govern the California electricity market.

⁶⁵The estimated welfare effects at higher prices are out-of-sample projections, which may not reflect how establishments would respond. There are no empirical estimates in the literature to inform this estimate. It is possible that establishments would be more or less elastic at higher price levels, which I cannot predict using my estimates. The constant elasticity demand curve assumes that same elasticity at all levels of demand, similar to what is used in Borenstein (2005) and Borenstein and Holland (2005).

⁶⁶The calculations assume that peak wholesale prices are greater than or equal to the peak price in each column. If, for example, peak prices only reached \$1.50/kWh, then the results in Column (3) would overestimate the benefits.

⁶⁷One important caveat of this analysis is that the estimated impacts assume that the super-peak days are correctly called as event days under all three approaches. A utility that faces a high level of uncertainty in forecasting super-peak days may choose to use 15 event days per year. This may erode some of the program benefits but would avoid the large costs of missing a super-peak event day.

⁶⁸I use the form $Q = AP^c$ for this calculation. I anchor the demand curve at the control price of \$.25/kWh.

example, the welfare benefits at \$1.35/kWh on 8 event days per year is \$259 million, which is 13 percent lower than when a linear demand curve is used. The small differences between the outcomes with two different types of demand curves illustrate that the functional form has a limited impact on the overall welfare results.

6.4 Comparing Peak Pricing to First-Best Policy

To put the second-best peak pricing program in perspective, I compare outcomes to the first-best alternative. Real-time pricing has been shown by previous research to result in efficient long-run outcomes, making it a useful benchmark (Borenstein 2005; Borenstein and Holland 2005). The comparison between peak and real-time pricing conducted in this section is parameterized by the empirical estimates from the paper, but the findings are robust to a range of assumptions. The model highlights that the majority of the welfare benefits of real-time pricing are due to the reduced need to build expensive and infrequently used peaker generation. Peak pricing, if effectively designed, can capture a large portion of this value.

To compare peak pricing and real-time pricing, I consider a theoretical energy-only electricity market where electricity supply and demand are cleared continuously with a uniform price auction. As discussed in Section 6.1, this stylized model approach has shortcomings when calculating the welfare benefits of time-varying pricing. However, many of these shortcomings are minimized when I use a stylized model to compare real-time pricing to peak pricing. The resource adequacy process would likely have similar effects on peak pricing and real-time pricing, meaning most of the inaccuracies will cancel out in the comparison.⁶⁹

Under this stylized real-time price, customers face retail rates that change every five minutes to reflect the real-time wholesale cost of electricity. I assume customers are fully informed about the real-time price they are paying and that their usage reflects the five minute price.⁷⁰ To allow a simple comparison between peak pricing and a real-time price, I assume the wholesale price takes on two distinct values. The low price reflects the marginal cost of generation at high-efficiency natural gas power plants, which I set at \$.10/kWh. When demand exceeds the capacity of the low-cost plants, the price of electricity spikes to the high level.⁷¹ The high price reflects the long-run cost of generation, which includes the costs of building and running peaker power plants to meet demand. I assume a high real-time price

⁶⁹Other factors that could affect the net consumer surplus loss calculations, such as bill protection, will also cancel.

⁷⁰Real-time pricing programs could have prices vary as frequently as every minute or in larger 15 to 30 minute increments. Joskow and Tirole (2006) suggest that customers may not respond to short-run changes in electricity price if transaction costs are too high. They suggest that this cost will be reduced through the use of advanced technologies that can quickly take advantage of price variation.

⁷¹This pricing structure reflects a retail electricity model with fixed charges, where the retail rate reflects the marginal cost of generation. Depending on natural gas prices, the cost at a high-efficiency natural gas

of \$1.35/kWh, which corresponds to the peak price paid by large commercial and industrial customers in the existing program. When demand drops to a level where the base load capacity is sufficient to balance load, the price returns to the low price level.

The simplified version of real-time pricing facilitates a transparent comparison to peak pricing, but it does not capture all of the potential benefits of real-time pricing. In practice, real-time prices vary in the short run between the low and peak prices throughout the day and year, which can result in a more efficient usage of existing generation capacity. These benefits, however, are an order of magnitude smaller than the gains from more efficient long-run capacity investment. Holland and Mansur (2006) focus on the short-run impacts of real-time pricing variation and find welfare improvements of only .24 percent. In contrast, studies that consider long-run capacity investment find that the benefits of real-time pricing are 5 to 10 percent of wholesale energy costs (Borenstein and Holland 2005). By focusing on the long-run capacity effects of real-time pricing, the model used in this section is able to capture the majority of the benefits of the policy using a transparent framework.

I benchmark outcomes under the peak pricing program against the benefits under real-time pricing. I first consider the existing peak pricing program, where prices are increased to \$.85/kWh between 2:00 p.m. and 6:00 p.m. on 15 event days per summer. In non-peak hours, customers are charged the same \$.10/kWh under peak pricing and real-time pricing to simplify the comparison.⁷² The well-targeted peak pricing program uses the optimal event-hour price of \$1.35/kWh on only 8 event days per year.⁷³ In both scenarios, I assume there are 3 super-peak event days each summer that provide capacity savings, and that these will be called as event days under both systems. During the super-peak days, I assume the price is at the high level for an hour between 2:00 p.m. and 6:00 p.m.

To compare peak pricing to first-best, I use my empirical estimates to calculate the welfare gains under peak pricing and compare them to the outcomes under real-time pricing.⁷⁴ For the existing peak pricing program, the difference between first and second-best comes from two sources. First, by charging an event price below \$1.35/kWh, peak pricing will

power plant may be lower than \$.10/kWh. A full accounting of the assumptions can be found in Appendix Section E.2.

⁷²I do this so real-time pricing and peak pricing raise similar amounts of revenue and are comparable in off-peak hours. Under both pricing plans, I assume revenue shortfalls are recovered with a monthly fixed charge.

⁷³It may be ideal to set the peak price slightly below \$1.35/kWh due to the net consumer surplus loss caused by peak pricing. For simplicity, I assume the well-targeted peak price is set at \$1.35/kWh. I use the temperature trigger proposed in Section 6.3 to select eight event days per summer.

⁷⁴I use the same elasticity for both the peak pricing and real-time price reductions. I assume customer response to the high price will be the same whether they face the high price for a short time or the full peak window. Wolak (2011) showed that, for residential customers, the response to peak prices was similar using both a short and a long event window. Appendix Section F shows that the results are robust to assuming that the real-time pricing response will be lower than the peak pricing response.

generate lower capacity construction savings than will real-time pricing. Second, under the current peak pricing program, the event price will be charged for 60 hours per year compared to just three hours under the real-time price. I choose a short period of time during which real-time prices are at the high level in order to remain conservative in reporting the benefits of peak pricing compared to real-time pricing. The longer the high event price is charged while real-time prices are low, the lower the relative benefits that peak pricing provides. The well-targeted peak pricing program, by setting peak prices at \$1.35/kWh, provides the same capacity construction savings as the real-time price. The lower number of event hours each year (32) also reduces the extra net consumer surplus loss that comes from non-super peak event days.⁷⁵

Using this approach, I find that the current peak pricing program provides 44 percent of the welfare benefits of the first-best approach. The result illustrates that the current program is providing some value, but performs poorly compared to the first-best policy. The well-targeted peak pricing program is able to make significant welfare improvements, delivering 83 percent of the first-best outcome. This result underscores the value of targeting. In markets with a binding capacity constraint, directly targeting the distortion caused by the constraint is an effective tool to improve welfare.

The benchmarking model is useful in understanding the impacts of poorly targeting the peak pricing program. Table 11 shows a number of alternate scenarios that consider how an incorrectly targeted peak pricing program might perform. As before, I assume the high real-time price is \$1.35/kWh. Column (1) mirrors the current program, where peak prices are set at \$.85/kWh; Column (2) shows the results for the correctly chosen peak price; and Column (3) shows the impacts if peak prices were set too high, at \$1.85/kWh. The first row shows the outcome when eight event days are chosen per year using the 101 degree trigger of the well-targeted program. The second row shows the outcomes with 15 event days. The bold entries correspond to the current and well-targeted peak pricing programs discussed previously. The other entries show the consequences of poorly targeting the peak pricing program.⁷⁶ The benchmarking model shows that, while the returns to targeting can be large, the downsides to incorrectly targeting are also significant. Setting the wrong price or calling too many days reduces program effectiveness. For example, calling 15 event days per year at a price of \$1.85/kWh would capture only 25 percent of the first-best outcome.

⁷⁵I include the small CO₂ benefits in each scenario that come from peak pricing providing more overall kWh reductions each summer than real-time pricing. The benefits are not large enough to offset the cost of the peak pricing program being in effect for more hours each summer.

⁷⁶See Appendix Table A13 for a robustness check where prices hit the peak for the full four-hour period between 2:00 p.m. and 6:00 p.m.

The model outlined for this calculation is stylized in nature and makes a number of simplifying assumptions that could impact outcomes. A linear demand curve is used to both estimate the net consumer surplus loss and demand reductions at prices above the observed \$.85/kWh peak pricing level. To better understand the role of this assumption, I use a constant elasticity demand curve to make the same calculations, finding that the current and well-targeted peak pricing programs provide 63 percent and 87 percent of the first-best benefits, respectively.⁷⁷ The constant elasticity demand curve yields better relative outcomes for peak pricing because the curvature of the demand curve results in a smaller net consumer surplus loss than under a linear demand curve. Despite the change in functional form, the results using a constant elasticity demand curve show the same qualitative results that the current program can be greatly improved using the proposed changes.

A number of additional robustness checks are conducted in Appendix Section F. I consider what happens if establishments are more responsive to peak pricing than real-time pricing because they have day-ahead notice. I also examine how the results are impacted if establishments have larger net consumer surplus losses from peak pricing than real-time pricing. Changing these assumptions has small impacts on the relative benefits of the two pricing programs, but it does not change the overall finding that well-designed peak pricing program can achieve a large portion of the first-best benefits.

The ability to design a well-targeted program depends on the aggregate peak pricing reductions combined with knowledge of institutional details to value the costs and benefits of the program. Using these insights helps inform the best way to target the costly capacity constraints by observing the underlying structure of peak event days on the PG&E grid. The estimation of the short-run electricity demand curve allows me to parameterize and balance the net consumer surplus losses under peak pricing against the capacity cost savings from higher prices. The empirical results also make the findings directly applicable to the current policy landscape because they show how small C&I establishments respond to peak pricing. The empirical analysis also provides a realistic starting point for the welfare comparisons based on the current program. Taken together, these results suggest that it is possible to achieve over four-fifths of the first-best outcome using a well-targeted, second-best policy.

7 Conclusion

Retail electricity customers in the U.S. are typically charged a flat price per kWh consumed. This time-invariant price does not reflect the cost of capacity at peak demand

⁷⁷See Appendix Table A14 for all of the scenarios shown in Table 11 recalculated using a constant elasticity demand curve.

hours. This paper studies a policy, peak pricing, that charges higher prices to retail customers on high-demand days when it is more costly to supply marginal units of electricity. Using quasi-random variation in program implementation and two different identification strategies, I find that establishments reduce their usage between 2:00 p.m. and 6:00 p.m. by 13.5 percent. In the aggregate, the peak pricing program will provide 118 MW of peak demand reductions in the PG&E service territory when fully implemented. The peak savings reduce the amount of generation capacity required at peak, yielding \$159 million of welfare benefits. I compare outcomes to a theoretical first-best, real-time pricing policy, finding that the current program captures 44 percent of the benefits. I show that a well-targeted peak pricing program could provide greatly improved outcomes, equaling 83 percent of the first-best outcome.

This paper fills an important gap in the literature by providing evidence of how commercial and industrial customers respond to peak pricing. This research is particularly important as the popularity of peak pricing programs continues to grow, fueled by the installation of low-cost, advanced metering technology. Programs like peak pricing also have the potential to facilitate the low-cost integration of intermittent renewable resources such as solar and wind. Furthermore, targeted peak pricing programs could be used in emergency situations when natural disasters interfere with transmission or generation infrastructure. More work is necessary to learn how peak pricing can be used to improve electricity grid efficiency as the grid changes and adapts to future challenges and demands.


References

- Abadie, Alberto and Sebastien Gay (2006). “The Impact of Presumed Consent Legislation on Cadaveric Organ Donation: A Cross-Country Study.” *Journal of Health Economics* 25.4, pp. 599–620.
- Aigner, D. J., J. Newman, and A. Tishler (1994). “The Response of Small and Medium-Size Business Customers to Time-of-Use (TOU) Electricity Rates in Israel.” *Journal of Applied Econometrics* 9.3, pp. 283–304.
- Aigner, Dennis J. and Joseph G. Hirschberg (1985). “Commercial/Industrial Customer Response to Time-of-Use Electricity Prices: Some Experimental Results.” *The RAND Journal of Economics* 16.3, pp. 341–355.
- Boiteux, M. (1949). “La Tarification Des Demandes En Pointe: Application de La Théorie de La Vente Au Coût Marginal.” *place Henri-Bergson*. Translated into English in Boiteux, M (1960) ”Peak-Load Pricing.” *The Journal of Business* 33.2, pp. 157-179.
- Bollinger, Bryan K. and Wesley R. Hartmann (2019). “Information vs. Automation and Implications for Dynamic Pricing.” *Management Science* 66.1, pp. 290–314.
- Boomhower, Judson and Lucas Davis (2020). “Do Energy Efficiency Investments Deliver at the Right Time?” *American Economic Journal: Applied Economics* 12.1, pp. 115–139.
- Borenstein, Severin (2005). “The Long-Run Efficiency of Real-Time Electricity Pricing.” *The Energy Journal* 26.3, pp. 93–116.
- (2007). “Wealth Transfers Among Large Customers from Implementing Real-Time Retail Electricity Pricing.” *The Energy Journal* 28.2, pp. 131–149.
- (2012). “Effective and Equitable Adoption of Opt-In Residential Dynamic Electricity Pricing.” *Review of Industrial Organization* 42.2, pp. 127–160.
- Borenstein, Severin and Stephen Holland (2005). “On the Efficiency of Competitive Electricity Markets With Time-Invariant Retail Prices.” *The RAND Journal of Economics* 36.3, pp. 469–493.
- Burkhardt, Jesse, Kenneth Gillingham, and Praveen K Kopalle (2019). “Experimental Evidence on the Effect of Information and Pricing on Residential Electricity Consumption.” *NBER Working Paper No. 25576*.
- California Air Resources Board (2019). *California Greenhouse Gas Emissions for 2000 to 2017*.
- California Energy Commission (2015). *Estimated Cost of New Renewable and Fossil Generation in California*.
- California Public Utilities Commission (2001). *Rulemaking 02-06-001 - Adopting a Business Case Analysis Framework for Advanced Metering Infrastructure*.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2011). “Robust Inference with Multiway Clustering.” *Journal of Business & Economic Statistics* 29.2, pp. 238–249.
- Choi, James et al. (2004). *Perspectives in the Economics of Aging*. University of Chicago Press.
- Energy Information Administration (2016a). *Form EIA-861*.
- (2016b). *Monthly Energy Review*.
- (2018). *Form EIA-861*.
- Fowle, Meredith et al. (2017). “Default Effects and Follow-on Behavior: Evidence from an Electricity Pricing Program.” *NBER Working Paper No. 23553*.
- Holland, Stephen P. and Erin T. Mansur (2006). “The Short-Run Effects of Time-Varying Prices in Competitive Electricity Markets.” *The Energy Journal* 27.4, pp. 127–155.
- (2008). “Is Real-Time Pricing Green? The Environmental Impacts of Electricity Demand Variance.” *Review of Economics and Statistics* 90.3, pp. 550–561.

- Ito, Koichiro (2014). “Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing.” *American Economic Review* 104.2, pp. 537–63.
- (2015). “Asymmetric Incentives in Subsidies: Evidence from a Large-Scale Electricity Rebate Program.” *American Economic Journal: Economic Policy* 7.3, pp. 209–37.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka (2018). “Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand.” *American Economic Journal: Economic Policy* 10.1, pp. 240–267.
- Jessoe, Katrina and David Rapson (2014). “Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use.” *American Economic Review* 104.4, pp. 1417–38.
- (Sept. 1, 2015). “Commercial and Industrial Demand Response Under Mandatory Time-of-Use Electricity Pricing.” *The Journal of Industrial Economics* 63.3, pp. 397–421.
- Johnson, Eric J., Steven Bellman, and Gerald L. Lohse (2002). “Defaults, Framing and Privacy: Why Opting In-Opting Out.” *Marketing Letters* 13.1, pp. 5–15.
- Joskow, Paul L. and Catherine D. Wolfram (2012). “Dynamic Pricing of Electricity.” *The American Economic Review* 102.3, pp. 381–385.
- Joskow, Paul and Jean Tirole (2006). “Retail Electricity Competition.” *The RAND Journal of Economics* 37.4, pp. 799–815.
- (2007). “Reliability and Competitive Electricity Markets.” *The RAND Journal of Economics* 38.1, pp. 60–84.
- Kahn, Matthew E., Nils Kok, and John M. Quigley (2014). “Carbon Emissions from the Commercial Building Sector: The Role of Climate, Quality, and Incentives.” *Journal of Public Economics* 113, pp. 1–12.
- Madrian, Brigitte C. and Dennis F. Shea (2001). “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior.” *The Quarterly Journal of Economics* 116.4, pp. 1149–1187.
- Pacific Gas & Electric (2010). *Advanced Metering Infrastructure Semi-Annual Assessment Report SmartMeter™ Program Quarterly Report*.
- Qiu, Yueming, Loren Kirkeide, and Yi David Wang (2018). “Effects of Voluntary Time-of-Use Pricing on Summer Electricity Usage of Business Customers.” *Environmental and Resource Economics* 69.2, pp. 417–440.
- Reiss, Peter C. and Matthew W. White (2005). “Household Electricity Demand, Revisited.” *The Review of Economic Studies* 72.3, pp. 853–883.
- Revezs, R. et al. (2017). “Best Cost Estimate of Greenhouse Gases.” *Science* 357.6352, pp. 655–655.
- Steiner, Peter O. (1957). “Peak Loads and Efficient Pricing.” *The Quarterly Journal of Economics* 71.4, pp. 585–610.
- Wolak, Frank A. (2007). “Residential Customer Response to Real-Time Pricing: The Anaheim Critical Peak Pricing Experiment.” *Center for the Study of Energy Markets Working Paper* no. 151.
- (2010). “An Experimental Comparison of Critical Peak and Hourly Pricing: The PowerCents DC Program.” *Working Paper*.
- (2011). “Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment.” *American Economic Review* 101.3, pp. 83–87.


Figures

Figure 1: Average Consumption Profile of Small Commercial and Industrial Establishments in Sample



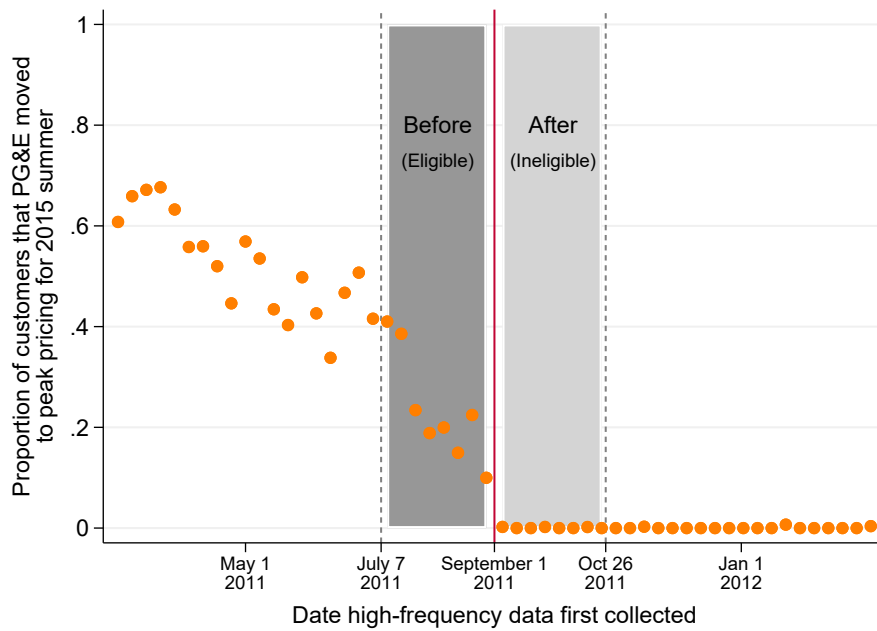
Note. — This figure shows the average consumption profile of the establishments in my analysis for all weekdays during the summer of 2014. The vertical lines signify the beginning (2:00 p.m.) and end (6:00 p.m.) of the peak event window. The system peak demand for the PG&E grid typically is between 4:00 p.m. and 6:00 p.m.

Figure 2: Timeline of Peak Pricing Rollout



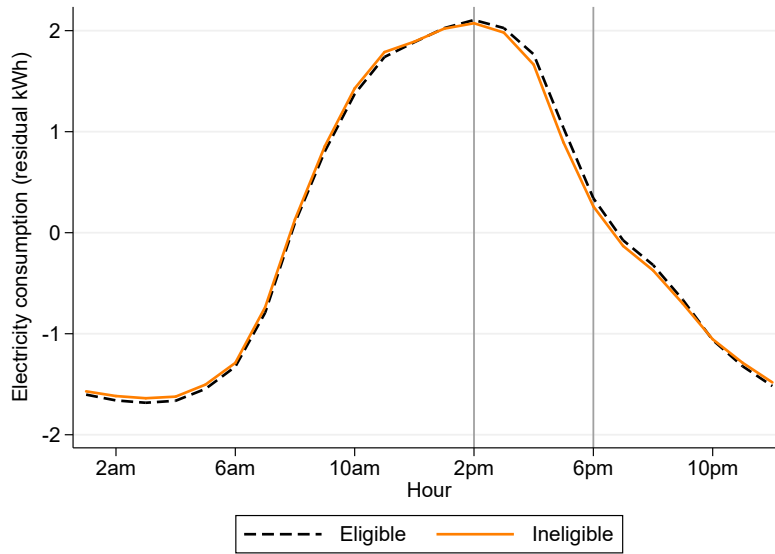
Note. — This figure shows the timeline of peak pricing implementation. I classify establishments as eligible for peak pricing in 2015 if their high-frequency metering data began before September 1, 2011 (eligibility cutoff). PG&E began placing some establishments on peak pricing in November 2014 (enrollment starts) based on eligibility and technical requirements, which are described in Section 4.1. Establishments may opt out of peak pricing at any time.

Figure 3: The Effect of Eligibility on Peak Pricing Treatment Status



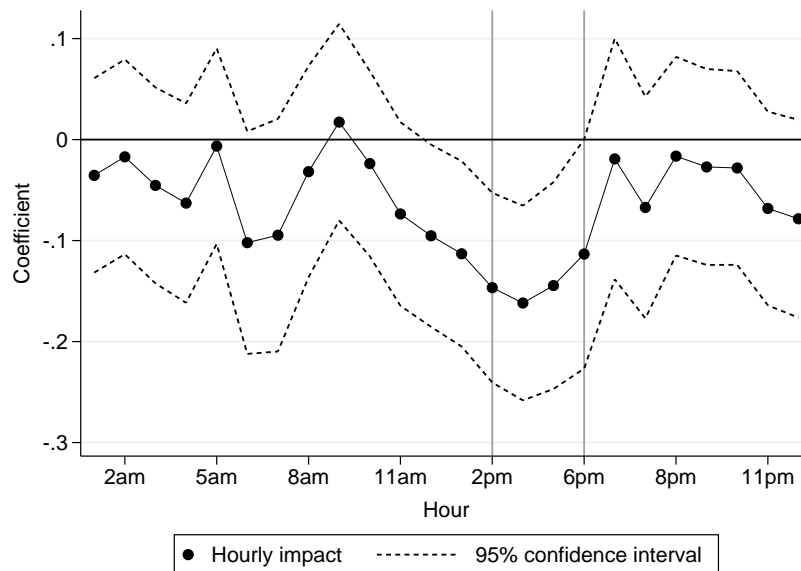
Note. — This figure shows the impact of peak pricing eligibility on PG&E moving an establishment to peak pricing. Establishments are binned by the week their high-frequency data began. Establishments to the left of the September 1, 2011 threshold are eligible to be moved onto peak pricing by PG&E. Establishments to the right of the threshold are the control group. There are around 500 establishments per bin. The figure shows 27 weeks in each direction from the threshold to show the larger default patterns. The vertical dashed lines represent the 8-week bandwidth used in the main specification.

Figure 4: Pre-Period Electricity Consumption on Event Days by Eligibility Group



Note. — This figure shows the summer 2014 pre-period average hourly consumption for peak pricing eligible and ineligible establishments. Electricity usage is on event days called in 2014 by PG&E the year before these establishments were enrolled in the program. Consumption is shown conditional on establishment fixed effects. I cannot statistically reject that the pre-period consumption is the same for both groups using hour-by-hour t-tests. The vertical lines signify the beginning (2:00 p.m.) and end (6:00 p.m.) of the peak event window.

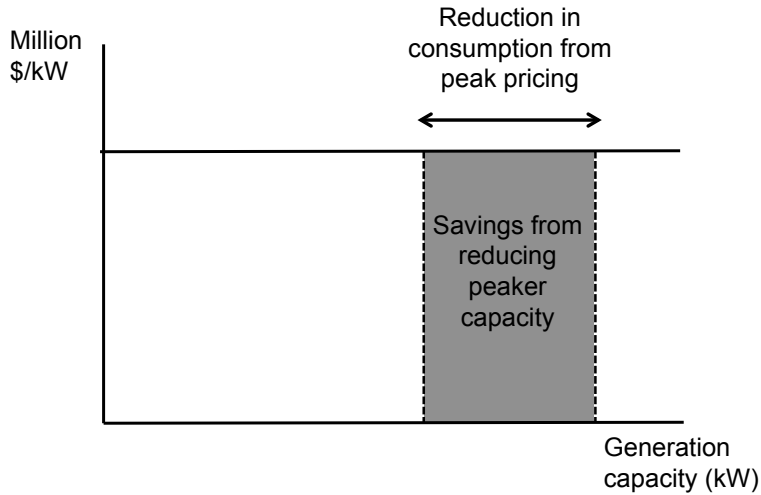
Figure 5: Effect of Peak Pricing on Inland Establishment Electricity Consumption




Note. — This figure shows the results of a regression estimating the hourly impacts of peak pricing for inland establishments on event days. Each dot corresponds to an hourly treatment effect comparing treated establishments with the control group. The dotted lines signify the 95 percent confidence interval. The vertical lines signify the beginning (2:00 p.m.) and end (6:00 p.m.) of the peak event window. The average impact between 2:00 p.m. and 6:00 p.m. reflects the coefficient in Column (5) of Table 3. The results show that establishments begin reducing their electricity usage in the hours before the event window starts.

Figure 6: Benefits and Costs of Peak Pricing

Panel A: Peaker capacity supply curve



Panel B: Demand curve



Note. — This figure graphically shows the benefits and costs of peak pricing. Panel A shows the capacity supply curve for fossil generation. Reducing peak demand lowers the need for peaker power plants. I assume a constant cost of \$1.2 million/MW to build a peaker plant, using California Energy Commission estimates to value the benefits. Using the demand reductions estimate for inland establishments, I find an aggregate reduction of 118 MW, which translates into a reduction of \$140 million in capacity costs. Panel B shows the hourly net consumer surplus (CS) loss from calling an event day. The horizontal axis is in kWh per hour (kWh/h), which is equivalent to kW. Short-run marginal production costs (SRMC) are \$.102/kWh and reflect the fuel cost at marginal power plants during peak hours. Q_0 is the quantity of electricity consumed during an event hour without peak prices, and Q_1 is the quantity consumed during an event hour with peak prices. I assume a linear demand curve and find that each event day reduces welfare by \$209,000. See Section 6.2 for a full discussion of the welfare impacts of peak pricing.

Tables

Table 1: Characteristics of Establishments by Peak Pricing Eligibility Status

Variable	Ineligible	Eligible	P value of difference
Summer 2014 avg peak hourly consumption (kWh)	5.17 (3.79)	5.19 (3.8)	.87
Summer 2014 max peak hourly consumption (kWh)	9.92 (6.82)	10.00 (6.86)	.61
Summer 2014 total event hours consumption (kWh)	219 (167)	219 (167)	.93
Summer 2014 total non-event hours consumption (kWh)	12,437 (8991)	12,267 (8612)	.40
Summer 2014 electricity expenditure	\$562 (396)	\$556 (385)	.52
Percent of establishments customer facing	.44 (.5)	.43 (.5)	.72
Money saved if program run on 2014 usage	-\$10 (59)	-\$12 (58)	.17
Average peak hour temperature (F)	73.24 (7.55)	73.38 (6.96)	.41
Establishment count	3,220	4,215	

Note. — This table shows the mean and standard deviation of the observable characteristics by peak pricing eligibility status for establishments within eight weeks of the September 1, 2011 threshold. Standard deviations are shown in parentheses. Customer-facing establishments are defined based on North American Industry Classification System codes, as discussed in Section 5.4.

Table 2: The Effect of Peak Pricing Eligibility on Enrollment (First Stage)

	(1) All PG&E	(2) Coastal	(3) Inland
Eligible \times Post	0.2230*** (0.0064)	0.1547*** (0.0068)	0.3654*** (0.0129)
Establishments	7,435	5,096	2,339
F statistic	406	174	268

Note. — This table reports regression coefficients from three separate first-stage regressions estimated using Equation (2). The dependent variable in all regressions is a binary indicator if an establishment is enrolled in the peak pricing program. Eligible \times Post is an interaction of an establishment's eligibility for peak pricing and 2015. The coefficients show the impact of peak pricing eligibility on program enrollment. All regressions control for temperature and include hour-of-sample fixed effects and establishment fixed effects. Standard errors are in parentheses and are two-way clustered at the establishment and hour-of-sample levels. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Table 3: The Effect of Peak Pricing on Peak Electricity Consumption (2SLS results)

	(1)	(2)	(3)
	All PG&E	Coastal	Inland
Peak pricing	-0.06974* (0.04139)	0.00842 (0.07095)	-0.14514*** (0.04575)
Temperature	-0.00499** (0.00236)	-0.01768*** (0.00347)	0.02517*** (0.00674)
Temperature squared	0.00008*** (0.00001)	0.00015*** (0.00002)	-0.00008** (0.00004)
Observations	742,939	509,245	233,694
Establishments	7,435	5,096	2,339
Event day kWh usage	5.55	5.03	6.70
Average temperature	78	71	92

Note. — This table reports regression coefficients from three separate 2SLS regressions estimated using Equation (1). The dependent variable in all regressions is the log of establishment hourly kWh consumption. Peak pricing is an indicator for enrollment in peak pricing, for which I instrument with eligibility status. The coefficients show the impact of peak pricing on peak consumption between 2:00 p.m. and 6:00 p.m. For inland establishments, the coefficient corresponds to a 13.5 percent reduction in usage. See Table A10 for the non-instrumented OLS regression. All regressions control for temperature and include hour-of-sample fixed effects and establishment fixed effects. Standard errors are in parentheses and are two-way clustered at the establishment and hour-of-sample levels. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Table 4: The Effect of Peak Pricing on Non-Event Day Peak Electricity Consumption

	(1)	(2)	(3)
	All PG&E	Coastal	Inland
Peak pricing	-0.0231 (0.0340)	0.0264 (0.0593)	-0.0718** (0.0364)
Observations	5,736,644	3,930,559	1,806,085
Establishments	7,435	5,096	2,339
Event day kWh usage	5.13	4.88	5.70
Average temperature	73	69	81

Note. — This table reports regression coefficients from three separate 2SLS regressions estimated using Equation (1) run with 2:00 p.m.–6:00 p.m. usage on non-event weekdays between June 1 and October 31, 2015. This approach is to test for spillovers to non-event days, as establishments do not face high peak prices during these hours. The results show no significant impacts at the 5 percent level on non-event days. The dependent variable in all regressions is the log of establishment hourly kWh consumption. Peak pricing is an indicator of enrollment in peak pricing, for which I instrument with eligibility status. All regressions control for temperature and include hour-of-sample fixed effects and establishment fixed effects. Standard errors are in parentheses and are two-way clustered at the establishment and hour-of-sample levels. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Table 5: The Effect of Peak Pricing on Peak Electricity Consumption for Inland Establishments: Temperature Interaction

	(1)	(2)	(3)
	All PG&E	Coastal	Inland
Peak pricing \times Temperature (F)	-0.00496* (0.00284)	0.00729 (0.00736)	-0.01128** (0.00535)
Peak pricing	-0.03885 (0.05568)	0.00616 (0.08190)	0.03957 (0.08833)
Temperature	0.00558*** (0.00060)	0.00397*** (0.00094)	0.01211*** (0.00159)
Temperature squared	0.00006*** (0.00002)	0.00008* (0.00005)	-0.00013*** (0.00004)
Observations	475,316	325,825	149,491
Establishments	7,435	5,096	2,339
Event day kWh usage	5.59	5.06	6.73
Average temperature	78	71	92

Note. — This table reports the regression coefficients from three separate 2SLS regressions where treatment is interacted with temperature. The dependent variable is the log of establishment hourly kWh consumption. Peak pricing \times Temperature (F) is the interaction between the treatment variable and hourly establishment temperature. Temperature has been re-centered at 75 degrees for scaling purposes. The coefficients show that peak pricing impacts are larger on hotter inland event days. The peak pricing impacts for inland establishments become positive around 79 degrees, which is lower than the temperature for all inland event days. The regression includes hour-of-sample fixed effects and establishment fixed effects. Standard errors are in parentheses and are two-way clustered at the establishment and hour-of-sample levels. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Table 6: The Effect of Peak Pricing on Peak Electricity Consumption for Inland Establishments: Industry Classification

	All PG&E		Coastal		Inland	
	(1) Customer facing	(2) Non-cust facing	(3) Customer facing	(4) Non-cust facing	(5) Customer facing	(6) Non-cust facing
Peak pricing	0.0352 (0.0557)	-0.1266** (0.0588)	0.0849 (0.0982)	-0.0477 (0.1021)	-0.0161 (0.0520)	-0.1984*** (0.0664)
Observations	288,797	374,130	213,208	246,580	75,589	127,550
Establishments	2,889	3,745	2,133	2,468	756	1,277
Event day kWh usage	6.34	5.15	5.66	4.61	8.25	6.19
Average temperature	76	78	71	72	92	92

Note. — This table reports regression coefficients from six separate 2SLS regressions estimated using Equation (1) broken down by industry type. The dependent variable in all regressions is the log of establishment hourly kWh consumption. Peak pricing is an indicator of enrollment in peak pricing, for which I instrument with eligibility status. The coefficients show the impact of peak pricing on peak consumption between 2:00 pm and 6:00 pm. Establishments are classified as customer-facing or non-customer-facing by their industry classification code, as described in Section 5.4. The regression is run on the subset of establishments for which I have an industry classification code. The estimates in Columns (1) and (2) are different with a p-value of .028 and the estimates in Columns (5) and (6) are different with a p-value of .052. For inland establishments, the non-customer-facing coefficient corresponds to a 18.0 percent reduction in usage. All regressions control for temperature and include hour-of-sample fixed effects and establishment fixed effects. Standard errors are in parentheses and are two-way clustered at the establishment and hour-of-sample levels. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Table 7: Welfare Impacts of PG&E’s Peak Pricing Program Components

Category	Net benefits of peak pricing (millions 2016\$)
Peak capacity construction costs (one time)	\$140 million
Avoided staffing and maintenance costs	\$3.07/year
Net consumer surplus losses	-\$3.16/year
Off peak \$.01/kWh discount	\$0.84/year
Environmental benefits	\$0.2/year
Total benefits	\$159 million

Note. — This table summarizes the welfare impacts of PG&E’s peak pricing program calculated in Section 6.2. All values are in millions of 2016 dollars. The total benefit of peak pricing is calculated using a 3 percent real discount rate and a 30-year horizon. The results are not sensitive to the discount rate because the yearly benefits and costs mostly offset each other.

Table 8: Welfare Impacts for 2015 Event Days

Event day	PG&E max load	Annual capacity cost savings (discounted)	Annual net consumer surplus loss (discounted)	NWS day ahead max temperature forecast	Trigger temperature
8/17/2015	19,451	\$10,000,000	-\$209,000	101	96
6/30/2015	19,320	\$0	-\$209,000	101	96
7/29/2015	19,248	\$0	-\$209,000	104	98
8/28/2015	19,233	\$0	-\$209,000	96	96
9/10/2015	19,230	\$0	-\$209,000	104	98
9/9/2015	19,017	\$0	-\$209,000	102	98
7/28/2015	18,403	\$0	-\$209,000	101	98
8/27/2015	18,328	\$0	-\$209,000	97	96
6/25/2015	18,114	\$0	-\$209,000	103	96
9/11/2015	18,019	\$0	-\$209,000	101	98
6/26/2015	17,950	\$0	-\$209,000	100	96
7/30/2015	17,750	\$0	-\$209,000	100	98
7/1/2015	17,734	\$0	-\$209,000	100	98
8/18/2015	17,372	\$0	-\$209,000	96	96
6/12/2015	17,275	\$0	-\$209,000	99	96

Note. — This table shows the two main welfare impacts of the 2015 event days. The annual capacity cost savings shows the benefits of reducing peak load. Annual capacity cost savings includes both the plant construction and operating costs, amortized over the assumed 30-year power plant life. There are non-zero savings numbers only for the super-peak event days of each summer. In 2015, only the highest load day was super-peak. The annual net consumer surplus loss shows the negative welfare consequences of charging higher prices during event hours and is displayed in the same units as capacity cost savings. The values are the same for all event days because the estimate is based on the average impact of peak pricing. NWS day-ahead maximum temperature forecast is the day-ahead temperature used by PG&E to call event days. It is based on the average of five National Weather Service weather stations. When the day-ahead maximum forecast equals or exceeds the trigger temperature, an event day is called.

Table 9: Welfare Impacts of Peak Pricing Under Alternate Scenarios

Scenario	(1)	(2)	(3)
	\$.85/kWh peak (current price)	\$1.35/kWh peak (large C&I peak price)	\$1.85/kWh peak (high price)
<u>Panel A: Linear demand</u>			
15 days/summer	\$159	\$212	\$212
101 degree trigger (8 days)	\$189	\$298	\$361
Super-peak days (3 days)	\$210	\$358	\$468
<u>Panel B: Constant elasticity demand</u>			
15 days/summer	\$189	\$211	\$212
101 degree trigger (8 days)	\$214	\$259	\$277
Super-peak days (3 days)	\$232	\$293	\$323

Note. — This table shows the welfare benefits (in millions of 2016 dollars) of the peak pricing program under different program design scenarios. Panel A shows the welfare calculations using a linear demand curve and Panel B does the same using a constant elasticity demand curve. Column (1) shows outcomes under the current \$.85/kWh peak price. Column (2) shows the estimated outcomes if the peak price were set at \$1.35, which is the level of large commercial and industrial customers and is based on a PG&E valuation of capacity at peak. Column (3) shows the impacts if the price was set at \$1.85/kWh. The first row of each panel reflects the current 15 event days per summer and the entry in the top left shows the welfare impacts estimated for the current program. The middle row of each panel reflects the proposed alternate 101 degree trigger for event days, and the bottom row of each panel shows the hypothetical scenario when only the three super-peak event days each year could be called. The welfare calculations assume that peak wholesale prices are greater than or equal to the peak price in each column.

Table 10: Welfare Impacts of Peak Pricing Compared to First-Best, Real-Time Price

Event days called per summer	(1)	(2)	(3)
	\$.85/kWh peak price (peak price < RTP)	\$1.35/kWh peak price (peak price = RTP)	\$1.85/kWh peak price (peak price > RTP)
8 event days (well targeted)	50%	83%	64%
15 event days (current)	44%	65%	25%

Note. — This table compares the peak pricing program to the first-best, real-time price across a number of scenarios. The percent values reflect the percent of the welfare benefits the peak pricing scenario can achieve compared to the first-best alternative. RTP is real-time price. For this table, the optimal peak price is set at \$1.35/kWh for an hour on three super-peak days per summer. Column (1) reflects the current program, where peak prices are set at \$.85/kWh, which is below the optimal level. Column (3) shows the impacts when prices are set above this level. The top row reflects the outcomes when eight event days are called per year. The bottom row shows the results for the current program, in which I assume 15 event days are used each summer. The current program achieves 44 percent of the first-best policy, while the well-targeted program could achieve 83 percent of the benefits.

Table 11: Welfare Impacts of Peak Pricing Compared to First-Best, Real-Time Price

Category	Net benefits of peak pricing (millions 2016\$)
Peak capacity construction costs (one time)	\$140 million
Avoided staffing and maintenance costs	\$3.07/year
Net consumer surplus losses	-\$3.16/year
Off peak \$.01/kWh discount	\$0.84/year
Environmental benefits	\$0.2/year
Total benefits	\$159 million

Note. — This table compares the peak pricing program to the first-best, real-time price across a number of scenarios. The percent values reflect the percent of the welfare benefits the peak pricing scenario can achieve compared to the first-best alternative. RTP is real-time price. For this table, the optimal peak price is set at \$1.35/kWh for an hour on three super-peak days per summer. Column (1) reflects the current program, where peak prices are set at \$.85/kWh, which is below the optimal level. Column (3) shows the impacts when prices are set above this level. The top row reflects the outcomes when eight event days are called per year. The bottom row shows the results for the current program, in which I assume 15 event days are used each summer. The current program achieves 44 percent of the first-best policy, while the well-targeted program could achieve 83 percent of the benefits.