An Experimental Study of Monthly Electricity Demand (In)elasticity

David P. Byrne,[a] Andrea La Nauze,[b] and Leslie A. Martin[a]

ABSTRACT

We document substantial rigidity in household electricity demand in response to large price shocks. We partnered with an electricity retailer to run a field experiment in which randomly-selected households received discounts of up to 50% on their total electricity bill or up to 95% off their per unit cost of electricity for a full month. We show that the quantity of electricity consumed was unaffected by these discounts. Exploiting rich billing, smart meter, and survey data, we document responses that are much more inelastic than previously observed in scenarios that raise prices for a few hours or raise or lower prices for indefinitely-long periods of time. Our results hold even among subgroups that we ex-ante believed were most likely to respond.

Keywords: Price Elasticity, Energy Demand, Field Experiment

https://doi.org/10.5547/01956574.42.2.dbyr

1. INTRODUCTION

The price elasticity of demand for electricity is a key parameter for analyzing the costs of climate change mitigation, the incidence of carbon pricing, market power, and electricity market design. This paper reports results from a field experiment designed to estimate this elasticity at the monthly level. In our experiment, we partnered with an electricity retailer to randomly provide households discounts of up to 95% off of their per-unit electricity price for two months. Combining our experiment with billing, smart meter, and survey data, we find that residential electricity demand is unresponsive to large reductions in both marginal and average prices. Our preferred own-price elasticity of demand estimate is –0.003. The estimate is statistically insignificant, but economically important as it points to perfectly inelastic demand. With 95% confidence, we can rule out households in our sample having a price elasticity more negative than –0.04.

In some ways our results are surprising. The non-experimental literature documents that households are relatively inelastic at this time horizon, but not as inelastic as we find them to be. Many of these studies struggle to control for endogeneously-set prices, endogeneous lagged consumption, and measurement error in prices (Alberini and Filippini, 2011). Some key exceptions are Ito (2014), who uses variation in relative prices along different increasing block rates across two adjacent utilities to find a monthly demand elasticity of –0.05 with respect to average price, and Deryugina et al. (2020) who use matching estimators to estimate one, three, and ten year elasticities of –0.16, –0.27, and –0.30 to –0.35, respectively, when prices decrease due to utility reform. The rich experimental literature avoids these identification issues, but has to-date focused on the imple-

a Corresponding authors. Department of Economics, The University of Melbourne, 111 Barry Street, Parkville, VIC 3010, Australia. E-mails: byrned@unimelb.edu.au and leslie.martin@unimelb.edu.au.

b Department of Economics, The University of Pittsburgh, Wesley Posvar Hall, Pittsburgh, PA 15260.

mentation of *hourly* changes in prices, like time-of-use, critical peak pricing, or real-time pricing interventions.

In this paper we describe a unique experiment with an electricity retailer partner who was willing to randomize prices for two months and let us publish the results. We provided different levels of discounts, some very large, which allows us to document highly inelastic demand at the monthly level across the entire range of the demand curve. We believe that the most likely explanation for our findings is that an intermediate, month-long time horizon is both too long to allow for significant inter-temporal smoothing of electricity use, and too short to change habits or household appliance stock. Over this time horizon, adjustment costs in behavior and technology can rationally justify non-response.

We exploit the richness of our dataset to address several potential alternative explanations for these findings. Working with our partner, we ensured that our discounts were credible. All correspondence with households in the experiment was sent directly by the retailer via email, in the same format as all previous on-going correspondence with the retailer. All other retailers in the country similarly use email for corresponding with customers. Furthermore, our experimental sample is constructed from a set of households who showed willingness to engage with the retailer by replying to a baseline survey before the experiment started.

We also provide evidence that non-response is not due to households being unaware of the discounts. We obtained server logs that confirm that relevant emails were indeed sent to all of the treated households and not to control households. Using an email tracking platform, we confirmed that the vast majority of treated households opened the treatment emails. Restricting our sample to households we are sure opened treatment emails, we still find next-to-no demand response to our experimental price shocks.

Our price discounts were promised and applied for a month, and then refreshed at different rates for a second consecutive month. Households were told that they would receive the discounts on their first quarterly bill after the experimental period was over. Some households received a non-discounted bill during the experimental period. To address concerns that our results could be driven by household confusion, we check and confirm that all of our findings are robust to restricting our sample to experimental household-days before any intermediate bills were received.

Households were notified about electricity bill savings-to-date after the first month of the experiment. We confirm that households did not change their behavior directly after that communication. We also confirm that they did not change electricity usage levels after the discount was applied to their post-experiment bill.

To what extent is the responsiveness of our customers representative of Victorian customers? On one hand, our partner retailer is a mid–sized retailer that entered the market during retail deregulation. Customers of our retailer are more likely to be price–aware than the two–thirds of Victorian customers still with the large incumbent firms from the pre-deregulation era. That said, we restrict our sample to customers who are not on pay-on-time discounts, who could be less price-responsive. But it is unclear as to how pay-on-time discounts translate to price sensitivity: the discounted price looks attractive, but the Australian Competition and Consumer Commission (ACCC) documents that households on these plans end up paying the much higher non-discounted prices 44% of the time (ACCC, 2018). Our experimental sample pays rates that are representative of actual prices paid by the retailer's average customer and slightly lower rates than the average customer in the state (Table 2).

Did prices disproportionately lack salience for the households that we study? This would have been the case if our households were disproportionately enrolled in automatic payment plans, and hence less likely to pay attention to their bills (Sexton, 2015). However, only 35% percent of the households in our sample have automatic payment plans. When we restrict our sample to households without automatic payment, we still find no evidence of a demand response to our experimental price shocks.

Households' appliance stocks are also unlikely to explain our findings as most households have appliances that allow for varied electricity use. Specifically, 87% of households have air conditioning, and many have electric heating (47%), clothes dryers (27%), and/or electric hot water (16%). We find statistical-zero treatment effects when we restrict our sample to these households that we expected to be relatively more price elastic.

Our results also cannot be explained by household absenteeism: over 92% of the household-days in our sample involve levels of use that exceed 4 kWh/day. We further find no evidence of a demand response when we restrict our analysis to hours of the day when we can infer from smart meter data that someone is likely to be home.

One possible explanation for the observed lack of response is experimenter-goodwill. If households are thankful to receive discounts they could be reluctant to increase electricity use in an experimental setting. We discount this explanation for a few reasons. First, the discounts were framed as a recognition for having answered a baseline survey, so they were *earned*, not given. Second, all emails were sent directly from the retailer, not the University. Finally, even if a goodwill effect did exist, such an effect would have had to grow exactly in proportion to discount size to explain our results.

It is possible that our result differs from findings in the literature due to asymmetry in price responsiveness. With the notable recent exception of Deryugina et al. (2020), many prior studies describe how consumers react to *increases* rather than decreases in electricity prices. We unfortunately did not design the experiment in a way that would allow us to confirm this hypothesis.

Related literature

Our findings contribute most directly to the literature of residential electricity demand that describes how households respond to monthly changes in prices. This large literature typically uses variation in prices induced by geographic and increasing block discontinuities (Ito, 2014), changes in regulatory policies (Deryugina et al., 2020), electricity crises (Bushnell and Mansur, 2005; Reiss and White, 2008; Costa and Gerard, 2015; Alberini et al., 2019), and some combination of the above with structural models (Reiss and White, 2005), and panel and time-series methods (Alberini and Filippini, 2011; Cuddington and Dagher, 2015; Ros, 2017; Burke et al., 2018). Most of this literature relies on data that is aggregated across households, for instance at the state or municipal-level.

Our paper differs from the recent body of experimental evidence primarily in our choice of time horizon. We consider prices shocks that are much longer in duration than is common in the experimental literature. Experimental studies typically examine the effects of price variation or shocks that last for several hours, with a focus on prices that are set ahead of time (time-of-use) as well as a few examples that fluctuate with market conditions (critical-peak-pricing or real-time-pricing). See, for example, Fowlie et al. (2017), Wolak (2011), Allcott (2011a), and Ito et al. (2018) as well as experiments in Australia summarized in Arup et al. (2014). Faruqui and Sergici (2010) and Faruqui and Sergici (2011) provide a survey of this literature that includes experiments involving private companies and utilities.

While many randomized control trials find responses at the hourly-level, recent experiments find that households can also be unresponsive to changes in price. Although Fowlie et al.

(2017) finds price elasticities of –0.075 and –0.31 to critical peak pricing and time-of-use prices, respectively, among *opt-in* participants, the authors find much smaller effects among the much larger set of customers who only engage if defaulted into program participation: –0.028 for critical peak pricing and –0.033 for time of use (authors' own calculations).

Jessoe and Rapson (2014) only find significant household response to hourly price spikes when coupled with in-home displays that increase price salience. Finally, Gillan (2017) finds that responses to hourly critical peak prices depend on whether events take place, regardless of the actual magnitude of the price change. In a policy environment where informational nudges produce small reductions in energy use (Allcott, 2011b; Ferraro and Price, 2013; Costa and Kahn, 2013; Byrne et al., 2018) and where cost-reflective prices are promoted as a politically-costly yet important way to lower system cost and achieve efficiency gains, it is important to understand the power of the price signal.

In the context of the literature, the unique time frame of our experiment and the magnitude of the price shocks that we study is policy relevant. Monthly price variation reflects changes in wholesale costs due to commodity price volatility (Bushnell and Mansur, 2005) and seasonality in both demand and supply. Holland and Mansur (2006) demonstrate that there are significant potential efficiency gains from passing through this variation to households, but the magnitude of the gains depends critically on the price elasticity of demand. Furthermore, monthly price elasticities have implications for the incidence of a price on carbon. In light of our null result, if households are inelastic over the medium-short run as we find, they will bear the full cost of carbon pricing over this period. Our findings directly point to this possibility, and any related distributional and political consequences.

This paper is structured as follows. Section 2 describes the experiment and data. Section 3 presents our empirical results, and we offer concluding remarks in Section 4.

2. THE EXPERIMENT

Our experiment was conducted in Victoria, Australia, the second most populous state in the country. The retail electricity market in Victoria is responsible for end-user electricity pricing and billing. It is a competitive market, with 18 companies vying for households. Retailers in the market frequently offer packages of discounts and bonuses to attract and retain households. We partner with a medium-sized retailer to survey and provide discounted electricity to a random subset of their existing customer base.

We selected our experimental group from a set of approximately 10,500 households without solar panels that answered an online baseline survey from a group of 40,000 households who were emailed the survey. We restricted the experimental sample to survey respondents that had at least one year of smart meter data with our retail partner before the experiment began.[1] We also excluded from the experiment all households on pricing plans that included pay-on-time discounts (89% of those excluded), varied by time-of-use (7%), or involved hardship rates (1%).[2] These restrictions left us with an experimental sample of 988 households. We randomly allocated 300 of these households to one of six treatment groups.

1. Smart meters were rolled out to all Victorian households between 2014 and 2016.

2. Many retailers in Victoria also offer rates that include discounts when customers pay their bills within 2 or 3 weeks of issue (pay-on-time). The discount can apply to the entire bill or a subset of the charges. We excluded customers on these plans to avoid the potential confound in customer-anticipated price.

We stratified the randomization based on whether a household had enabled automatic bill payment, and the ratio of fixed to variable costs in the household's bill from the same period the previous year. Six households from our original treatment group and twenty-two households from our original control group became ineligible during the three-month delay between design and implementation due to either switching retailers, switching to an ineligible price plan, or adopting solar panels. One household also opted-out of treatment, without providing an explanation as to why. Our final sample therefore has 293 treated households and 666 control households.

2.1 Design

We designed our experiment to vary separately, within household, both marginal and average electricity price. In the market that we study, electricity bills include both fixed daily supply charges and variable charges based on the amount of electricity used. Households pay on average \$1/per day in fixed charges and \$0.24/kWh in per kWh variable charges. For a household with an average monthly use of 380 kWh/month, variable charges make up 77% of each bill. We define average price as the ratio between total charges and kWh consumed, using the last bill received before the experiment began to benchmark consumption.[3]

We allocated two types of experimental discount, sequentially, to every treatment household. One treatment gave households a 50%, 25% or 10% discount on their total monthly electricity bill, i.e. on both their fixed and variable charges. The other treatment gave households larger discounts on just the variable charges. The variable-only discount was set to keep the income effect constant across treatments within household.[4]

We also randomized which group received which treatment first. Treatment groups A, B, and C received a 50%, 25% and 10% discount off their total bills, respectively, in the first month. In the second month, these groups received the income-equivalent reduction in their variable charges. Treatment groups D, E, and F were respectively offered the same discounts as groups A, B and C, with the order of treatments reversed. Table 1 summarizes the six treatment conditions.

Figure 1 shows the average change in fixed and variable charges over the two rounds of the experiment for the six treatment groups. For example, in treatment group A both variable and fixed charges drop 50% in the first round, and variable charges drop on average 77% in the second round while fixed charges return to their original levels. Figure 2 shows the distribution of charges across households before and after the first round of experimental variation. Panels (a) and (c) show that the distributions were balanced across the treatment and control groups before the experiment began. Panels (b) and (d) illustrate the large exogenous variation in charges induced by the experiment. For some treatment households, the experiment drops the cost of using a kWh of electricity from \$0.24 to less than \$0.01. For context, this discount reduces the cost using an air conditioner for an hour each day for four weeks from approximately \$10 to less than 50 cents.

3. Unlike Ito (2014) who studies variation in average price created by increasing block marginal rates, our pre-experiment variation in average price is derived from a large daily connection fee accompanying flat marginal rates. So whereas in Ito (2014) average price is lower than marginal price and increases with quantity consumed, in the population we study average price is higher than marginal price and decreases with quantity consumed.

4. Specifically, the discount on the variable charge is the product of the total bill discount and 1 plus the historical ratio of fixed to variable charges. The ratio of fixed to variable charges is calculated based on average daily kWh in October and November 2014, a year before the experiment took place. Therefore, households with lower levels of historical electricity use received larger discounts in the variable-only treatment.

Figure 1: Time variation in variable and fixed charges

(a) Variable charge ($/kWh) for groups A, B, & C

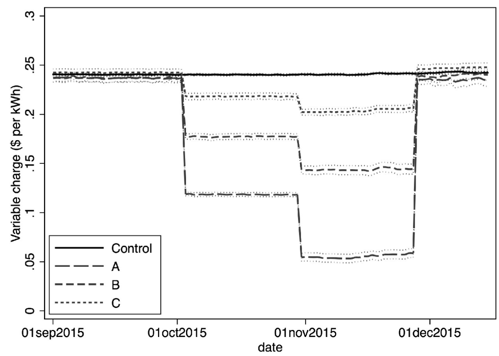

(b) Fixed charge ($/day) for groups A, B, & C

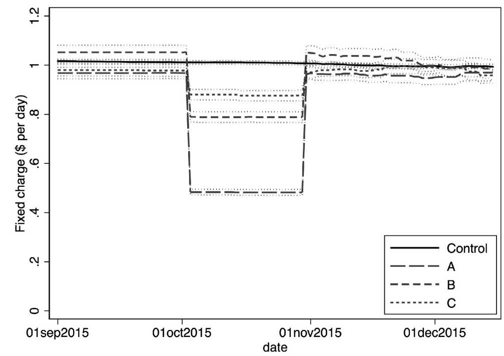

(c) Variable charge ($/kWh) for groups D, E, & F

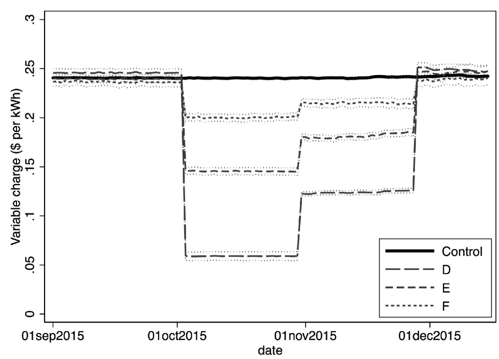

(d) Fixed charge ($/day) for groups D, E, & F

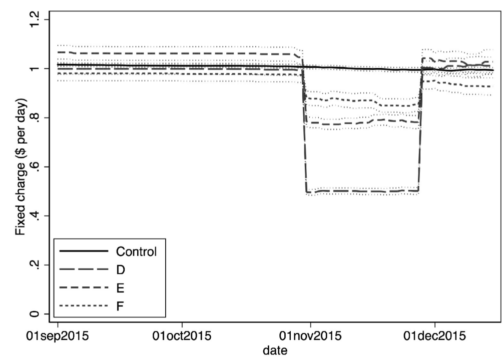

2.2 Implementation

We offered the discounts to treatment households for 8 weeks in October and November 2015. Households were notified via email on October 2 about the price changes two days before the discounts went into effect at 11:59pm on Saturday, October 3. In this first email they were promised four weeks of discount, without any mention of a second four-week discount period.[5]

Using the analytics platform Kissmetrics, our partner retailer confirmed that the vast majority of households opened the treatment emails.[6] Specifically, at least 78% of households opened

5. Figure A.1 in the Appendix provides example emails sent to treatment households. We drafted the text in straightforward language to minimize customer confusion. For example, for group A: "We're giving you 50% off your electricity bill next month" followed by text that clarifies "we're giving you an additional 50% discount off next month's electricity usage and supply charges." That same household later received an email from the retailer with subject heading "we're giving you 86% off your electricity usage charges next month". The body of the email reminded them of what they had received the previous month, how much they had saved. It clarified: "As a future thank you, we would like to give you another discount—this time a 86% discount off next months' electricity usage charges." Customers in Victoria are relatively familiar with the terms "usage charges" and "supply charges" due to widespread marketing by competing retailers to attract customers.

6. Kissmetrics uses the downloading of embedded images to report which households opened which emails. To protect household privacy, some email providers do not display external content automatically in their default settings. In that case, households may have opened the email, but receipt will not be flagged. Furthermore, because the discount was announced in the subject header, it is possible that households that did not open the emails were aware of the discounts.

Figure 2: Distribution of variable and fixed charges

(a) Variable charge, pre-treatment

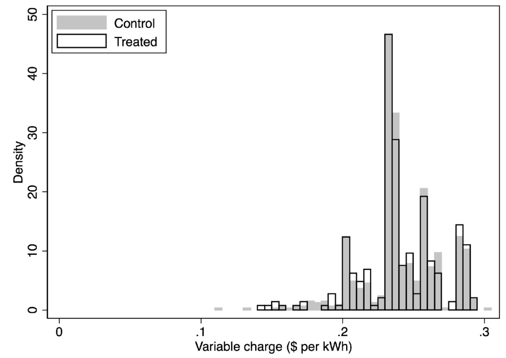

(b) Variable charge, round 1

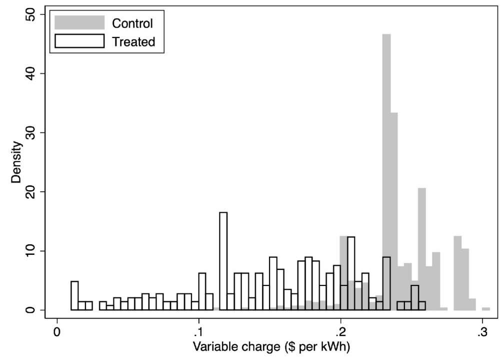

(c) Fixed charge, pre-treatment

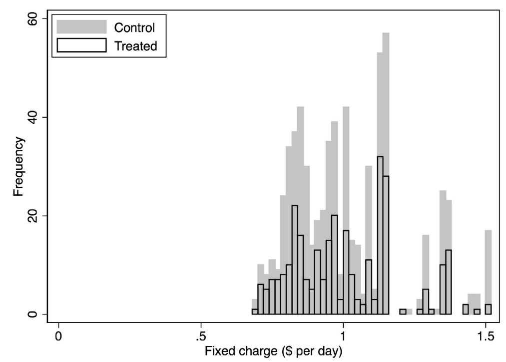

(d) Fixed charge, round 1

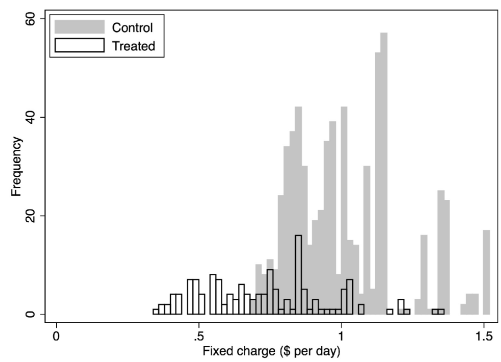

the email announcing the first month of treatment discounts, and at least 63% opened the email announcing the second month of discounts. 84% of households opened at least one of the discount announcement emails.

On October 30, 2015, treatment households received an email announcing another four weeks of discounts. Households that previously received discounts on their total bill received matched discounts on their variable charge, and vice-versa, as per Table 1. Households were also told how much they had saved in the first round of discounts, holding fixed their electricity consumption from the previous month. On average, households in treated groups A and D saved $60, households in treated groups A and D saved $30, and households in treated groups A and D saved $12. The second set of discounts were implemented on October 31 at 11:59pm, and were again promised for four weeks.

On Friday, November 27, 2015, households were sent a final email informing them that the discount period ended. The email also summarized how much they had saved thanks to the second discount. We were unable to run a follow-up survey at the end of the experiment due to unexpected logistical constraints with our partner retailer.[7]

7. A follow-up survey would have been ideal. We would have liked to ask households to recall the fixed and variable electricity charges they faced in the previous two months to confirm that they fully understood the experimental price shocks. The logistical constraint that we ran into was our partner retailer was not willing to allow us to further interact via email with

Table 1: Treatments

Treatment Group	N	Oct 2015	Nov 2015
A	49	50% off $totalbill_{it}$	$59 < X_{it} < 95\%$ off $variablecharge_{it}$
B	49	25% off $totalbill_{it}$	$29 < X_{it} < 70\%$ off $variablecharge_{it}$
C	48	10% off $totalbill_{it}$	$12 < X_{it} < 26\%$ off $variablecharge_{it}$
D	50	$59 < X_{it} < 95\%$ off $variablecharge_{it}$	50% off $totalbill_{it}$
E	48	$29 < X_{it} < 70\%$ off $variablecharge_{it}$	25% off $totalbill_{it}$
F	49	$29 < X_{it} < 70\%$ off $variablecharge_{it}$	10% off $totalbill_{it}$

Notes: X_{it} is set to make the expected total bill savings across months equal for household i in Oct 2015 and Nov 2015 based on use in Oct 2014 and Nov 2014.

2.3 Data and summary statistics

We obtained half-hourly smart electricity meter data along with billing data from the electricity retailer for a full year before treatment and several months afterwards, specifically from October 2014–April 2016. We combine these data with half-hourly temperature observations, household survey responses, and census district characteristics from the Australian Bureau of Statistics. Census data are collected at the "Statistical Area Level 1" (SA1), a census block that typically contains 250 homes.

Table 2 shows average electricity prices, monthly consumption, and bills for a few selected regions. Monthly consumption in our experimental sample is slightly higher than the Victorian average but lower than the Australian average. Similarly, monthly consumption is lower than in California, but higher than the European Union. Average prices are slightly higher than average for Victoria, and are exactly on par with prices in the European Union.

Table 2: Electricity prices and consumption by region

Region	Average price (USD/kWh)	Monthly consumption (kWh)	Monthly bill (USD)
United States	12.65	901	114.03
California	16.99	557	94.59
Australia	18.74	437	83.52
Victoria	21.48	336	89.96
Experimental sample	**23.33**	**387**	**90.29**
Euro area	23.50	310	72.85
Germany	31.49	266	83.76

Notes: Table shows average price of electricity in 2015 USD per kWh, monthly consumption in kWh, and average monthly bill in 2015 USD for selected regions in 2014/15. Euro area are the 19 European Member States who have adopted the euro as their currency. Monthly bills for Victoria and Australia consist of fixed and variable charges, monthly bills in other jurisdictions consist only of variable charges. Sources: AEMC (2015), AER (2014), EIA (2015), Eurostat (2015a), Eurostat (2015b), IRS (2017).

Appendix Table B.2 further compares the experimental sample to the full set of households serviced by our partner retailer, and to the subset of households who answered the online survey. Although our sample pays slightly higher variable costs, it pays slightly lower fixed costs, and neither difference is economically significant. In terms of household characteristics, our sample is more likely to have air-conditioning, less likely to have electric hot water, less likely to be a renter, has a

the customers in the experiment. This is because the retailer perceived an already significant degree of customer interaction (beyond the typical degree of interaction with the company) through the email–based sample recruitment process and experiment implementation.

higher average number of rooms and people and lives in neighborhoods with higher median income. In each case, however, the economic magnitude of the difference is small.

Table 3 presents summary statistics for treatment and control. There are no statistically-significant differences in pre-treatment characteristics across the groups. Appendix Table B.1 further demonstrates balance between each treatment subgroup and the control group.

Table 3: Pre-treatment summary statistics

	(1) Control	(2) Treatment	(3) Difference (SE)
Panel (a): Billing and meter data			
Average daily consumption (kWh)	12.69	12.53	0.16
			(0.46)
Variable charge (AUD $ /kWh)	0.24	0.24	0.00
			(0.00)
Fixed charge (AUD $/day)	1.10	1.10	–0.00
			(0.01)
Has automatic bill payment	0.35	0.35	0.00
			(0.03)
Voluntary green power purchaser	0.20	0.24	–0.03
			(0.03)
Panel (b): Survey data			
Has air conditioning	0.88	0.84	0.03
			(0.02)
Has electric heating	0.49	0.44	0.04
			(0.04)
Has electric hot water	0.15	0.17	–0.01
			(0.03)
Has clothes dryer	0.27	0.27	0.01
			(0.03)
Is a renter	0.20	0.17	0.04
			(0.03)
Has freestanding house	0.19	0.19	0.00
			(0.03)
Number of bedrooms	3.22	3.30	–0.08
			(0.06)
Number of residents	2.98	3.02	–0.04
			(0.09)
Panel (c): Census data			
Median weekly household income	1348.54	1347.35	1.18
			(23.87)
Median age	37.11	37.53	–0.42
			(0.36)
Full time employment rate	0.39	0.39	0.00
			(0.00)
Households	666	293	

Notes: Table shows means and differences in means for control and treatment. Standard error of difference in parentheses. All variables are measured prior to the experiment. Average daily consumption measured from 01 October 2014 through 30 September 2015, inclusive. *** p<0.01, ** p<0.05, * p<0.1

Table 4 presents average daily electricity use, in kWh/day, for each treatment period for the control group, households in any treatment group, and households in each of the treatment groups. Most of the differences represent a slight decrease in use, not an increase, and all are very small, 0.3 kWh/day or less, compared to an average daily use of close to 12 kWh/day. None are statistically-different from zero.

Table 4: Mean differences by treatment group

		Mean Daily kWh					Difference in Mean Daily kWh Relative to Control		
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		Control	Treated	AD	BE	CF	AD	BE	CF
Round 1	Mean	11.9	11.8	11.8	11.7	11.9	–0.19	–0.26	–0.061
	SD	[7.95]	[7.90]	[7.85]	[8.29]	[7.56]	(0.16)	(0.16)	(0.16)
	Obs	18633	8096	2716	2692	2688			
Round 2	Mean	11.8	11.7	11.6	11.5	11.9	–0.18	–0.21	0.15
	SD	[7.85]	[7.79]	[8.05]	[7.86]	[7.42]	(0.16)	(0.16)	(0.16)
	Obs	18,601	8,075	2,715	2,673	2,687			

Notes: First five columns show average daily electricity use, in kWh/day, for each treatment period for the control group, households in any treatment group, and households in treatment groups A and D (50% discount), B and E (25% discount), and C and F (10% discount). Columns (6)–(8) show differences with respect to the control group. Standard deviations in brackets, standard errors of the differences in means in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

3. TREATMENT EFFECTS AND DEMAND ELASTICITIES

Using fixed effects regressions, we estimate treatment effects at the monthly level. More precisely, treatment effects are by 4-week period because each round of treatment lasted 4 weeks. Table 5 presents the results. We first regress total kWh/month for each household on treatment status, by size of discount. This regression only includes observations during the treatment period, so effectively presents simple differences in means, controlling for round of treatment. Standard errors are clustered at the household level. The second column of the table presents the results of the same regression, now also including household baseline electricity use as a control. None of the estimated treatment effects are economically large nor statistically significant.

The third column of Table 5 presents the results of our preferred treatment effect estimates, which are based on the following regression:

$$q_{it} = \sum_j \beta_j T_{ijt} + \mu_i + \tau_t + \varepsilon_{it} \tag{1}$$

where q_{it} is total monthly electricity consumption for household i in month t. The average treatment effect for discount group j is β_j, where the three groups are A and D (50% discount), B and E (25%), and C and F (10%). The treatment groups indicators are denoted T_{ijt}, μ_i are household fixed effects, τ_t are month fixed effects, and ε_{it} is the residual error term. This regression is a difference-in-differences comparison of treated households relative to control households in the treated and baseline periods. The corresponding results in Table 5 are again economically-small and statistically insignificant.

3.1 Time-varying treatment effects

We also estimate treatment effects by calendar date. Panel (a) of Figure 3 plots the raw data at the daily level: average daily consumption of treatment and control groups from September to December 2015. There is no distinguishable difference between the electricity consumption of treatment and control households on any day in the pre-experiment period, or the first or second rounds of treatment that follow. Of note, there is no response on or shortly after October 31, when households were told how much they had saved in the first round. Figure A.2 in the Appendix presents a set of similar figures, further broken down by discount group.

Panel (b) of Figure 3 shows the estimated treatment effect for each date of the treatment period. The point estimates are derived from regressions of daily electricity use on baseline use and

Table 5: Monthly treatment effects

	(1) kWh/month	(2) kWh/month	(3) kWh/month
Treated 50% discount	–5.2792	–4.6775	–1.7201
	(21.251)	(21.261)	(8.092)
Treated 25% discount	–10.0901	–10.4895	8.6642
	(21.541)	(21.462)	(9.164)
Treated 10% discount	1.5493	2.0635	5.1130
	(20.452)	(20.547)	(7.388)
Observations	1,915	1,915	14,339
R^2	0.000471	0.00141	0.242
Control for baseline use	—	Yes	—
Household FE	—	—	Yes
Month FE	Yes	Yes	Yes
Treated households	293	293	293
Control households	666	666	666
Mean baseline use	383.1	383.1	383.1

Notes: Dependent variable is total kWh/month for each household. Because each round of treatment lasted 4 weeks, "month" does not refer to calendar month but to a 4-week period, counting forwards and backwards from the first date of treatment. Columns (1) and (2) only include observations during the treatment period. Column (2) includes household average monthly consumption for the one year prior to treatment as a control variable. Column (3) includes pre-treatment observations. Standard errors (in parentheses) are clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

relevant treatment dates, controlling for household and date fixed effects. The dashed line represents the 95% confidence interval around the treatment effect. Again, there are no statistically-significant observed changes in use.

Figure A.3 in the Appendix further breaks down pre-treatment and treatment consumption by half-hourly interval. In the raw means there appears to be slight shifting of use away from mid-day and into evening hours among treated households. But Appendix Figure A.4 that controls for household-interval fixed effects shows no evidence of a treatment effect, regardless of time of day or level of discount.

3.2 Price elasticities

To compare our results to previous literature, we estimate price elasticities and associated confidence intervals. We obtain elasticities by estimating the following equation:

$$\log(q_{it}) = \eta \log(p_{it}) + \mu_i + \tau_t + \varepsilon_{it} \tag{2}$$

where η is the price elasticity of demand, μ_i are household-register fixed effects, τ_t are date fixed effects, and ε_{it} is the residual error term. Here, i represents the household-register level because 74 out of 959 households have multiple registers with different variable and fixed charges.[8]

We estimate the elasticity with respect to both marginal and variable price. Marginal price is the per kWh variable charge, net of any experimentally-induced discounts. Average price is total charges divided by total use. Equivalently, it is the per kWh variable charge plus the ratio of the daily fixed charge to average daily use, net of any discounts induced by the experiment. We calculate average price based on the average daily use level of the most recent bill that the household received

8. Households with "controlled load" for electric hot water or slab heating have more than one register. Discounts were applied to all prices on all registers. Our results are robust to only using the primary register for each household.

Figure 3: Daily use and treatment effect

(a) Daily use by date

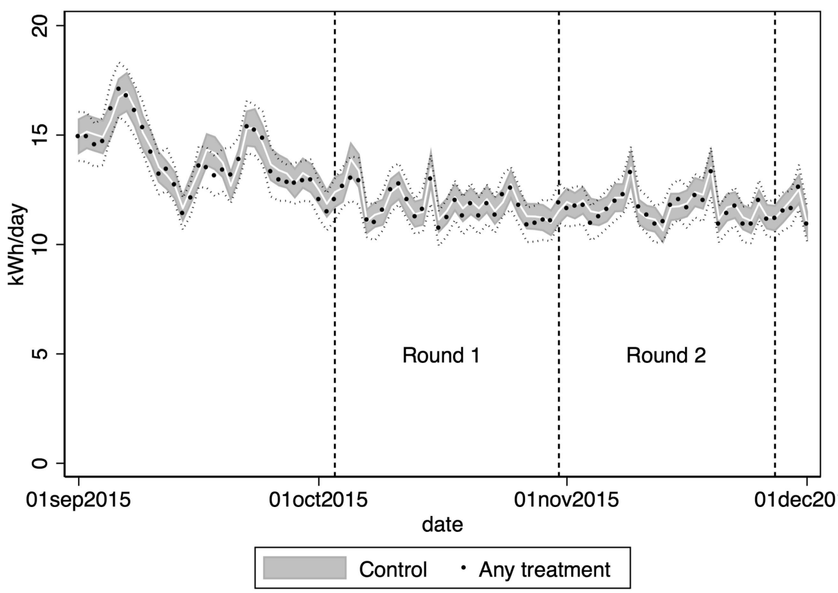

(b) Treatment effect by date

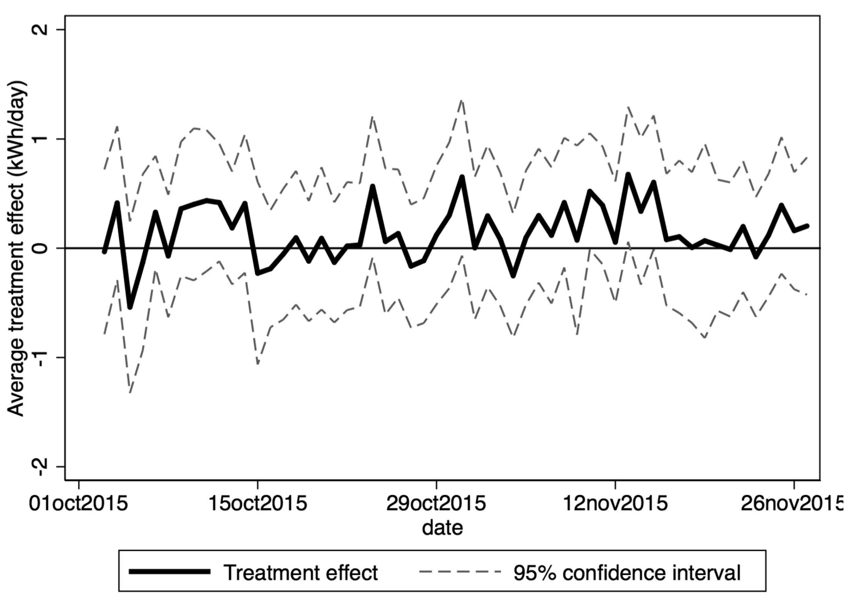

Notes: Panel (a) plots average daily consumption of electricity, in kWh/day, before, during, and after the experiment for treatment group (black dots) and control group (white line on blue). The black dotted line represents a 95% confidence interval around the mean for the treatment group, while the shaded blue area represents a 95% confidence interval around the mean for the control group. Vertical lines show the beginning of the first and second months of treatment and the end of the experimental period. Panel (b) plots the average daily treatment effect, in kWh/day, for each date during the treatment period. Each daily point estimate, and associated confidence interval, is derived from a regression of daily electricity use on treatment status, with household and date fixed effects, over the baseline period and given treatment date. Standard errors (in parentheses) are clustered at the household level.

before the experiment started. Our results are robust to using a longer or more recent historical period.

Our experiment creates random variation in the fixed and variable charge for one of the two rounds of treatment. But in order to hold income effects constant across treatments within household, the discount on the variable charge for the round with no reduction in fixed charge is based on household use one year prior to treatment, with households with higher levels of use receiving smaller discounts. Average price also relies on historical use levels. Because time-varying household characteristics could be correlated with historical household use, we instrument for price using the randomly-allocated fixed charge discount (50%, 25%, 10%) interacted with an indicator variable for the treatment period.

Table 6 presents our 2SLS demand elasticity estimates. The first stage F statistic shows that the randomly-provided discounts are a strong instrument for price. We find similar point estimates for the elasticity of electricity demand with respect to both marginal and average price: –0.003. Both estimates are not statistically or economically different from zero. With 95% confidence, we can rule out that the true elasticity of demand with respect to marginal price is more negative than –0.04. For average price, we can rule out true elasticities more negative than –0.055.

Table 6: Price elasticity estimates

VARIABLES	(1) log(kWh/day)	(2) log(kWh/day)
log Marginal price	–0.0031	
	(0.018)	
log Average price		–0.0034
		(0.026)
Observations	433,590	431,147
R^2	0.128	0.132
Number of hhregid	1,032	1,025
Household-register FE	Yes	Yes
Day FE	Yes	Yes
% confidence interval (low)	–0.0393	–0.0551
% confidence interval (high)	0.0330	0.0483
CDW F-test	1,115	4,700

Notes: Dependent variable is daily register-level log(kWh). A household may have more than one register. Marginal price is the per kWh charge. Average price is the variable charge plus daily fixed charge divided by average historical daily use. Use is based on electricity use on the most recent bill the household received before the experiment. Price is instrumented using treatment status (10%, 25%, 50% discount) × treatment period. Standard errors (in parentheses) are clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1

Figure 4 presents the results of estimating equation (2) separately for each hour of the day for marginal price. We find no evidence that households are responsive to our experimental price shock at any time of day.

3.3 Potential sources of demand inelasticity

To understand the source of demand inelasticity, we look for effects among the households most likely to respond. Table 7 presents results for various selected subsamples. Column (1) restricts the sample to households that opened the treatment email as identified by Kissmetrics. Column (2) drops household-days that may be subject to confusion or credibility concerns because they follow

Figure 4: Estimated price elasticity by hour-of-day

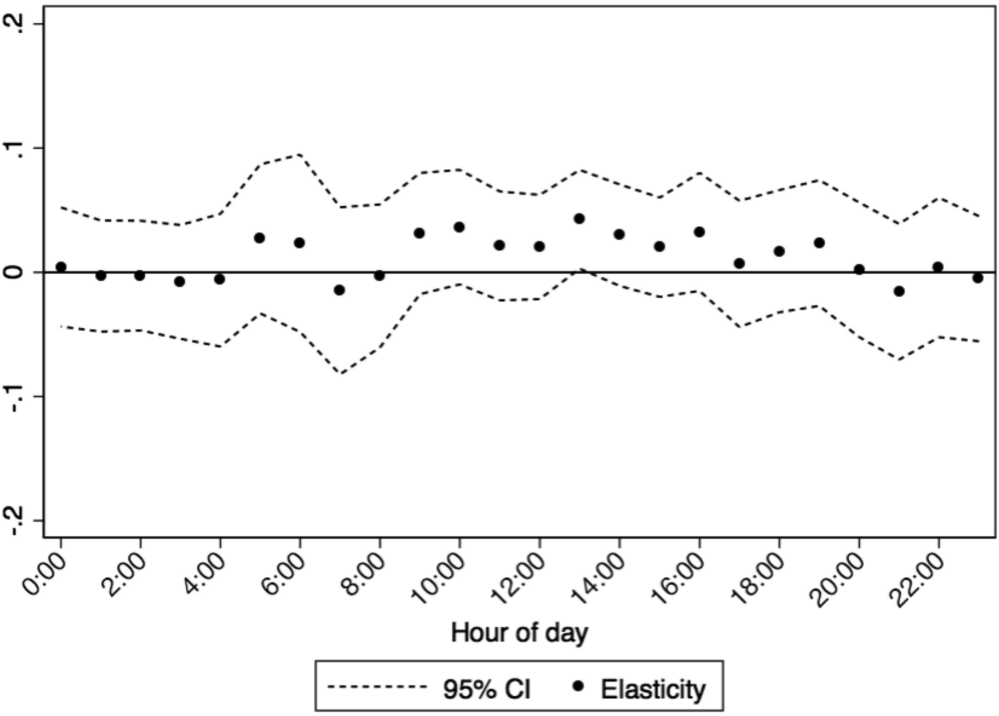

Notes: Figure plots point estimates and 95% confidence intervals of elasticity estimates (marginal price). Each elasticity is estimated separately using household-hour and hour-date fixed effects.

an intermediate bill during our 8-week experimental period that did not yet reflect the experimental discounts.

Bill salience could also explain responsiveness. Column (3) of the table restricts the sample to households who pay each bill as it arrives, that is, households without automatic bill payment. This group opts for a payment method that increases bill salience. Price signals can also be dampened by failures to communicate within-household. Column (4) therefore restricts the sample to households where the person receiving the email is likely to also make many of the use decisions, that is, households with two or fewer residents.

Households who opt to pay more for a renewable energy plan may be more likely to internalize the social costs of their use decisions. They also, on average, pay higher rates, so are receiving larger level discounts. Column (5) therefore presents results restricting the sample to households that have not opted to pay more for a renewable energy plan.

Some households may also have an easier time adjusting their electricity use, due to being at home or their appliance stock. Column (6) restricts the sample to household-days where we are confident based on smart meter data that someone is home, i.e. for which daily consumption levels exceeding 4 kWh/day. Column (7) restricts the sample to households that have central or unit air-conditioning.

The overall message from these auxiliary results is that demand elasticity with respect to marginal or average price is not statistically or economically significant. In some cases we can reject elasticities as small as –0.02 with 95% confidence. The largest confidence interval is around households with only 1 or 2 residents, where we lose two thirds of our sample. Even in that case, we can reject marginal price elasticities as small as –0.05 with 95% confidence.

Table 7: Additional price elasticity estimates

	(1) Opened email	(2) Before bill	(3) Not auto debit	(4) Few residents	(5) Not green plan	(6) At home	(7) Has AC
log Marginal price	0.0014	0.0014	0.0157	0.0064	–0.0037	0.0051	0.0074
	(0.021)	(0.020)	(0.020)	(0.029)	(0.022)	(0.016)	(0.020)
% confidence interval (low)	–0.0393	–0.0379	–0.0234	–0.0501	–0.0470	–0.0272	–0.0317
% confidence interval (high)	0.0420	0.0407	0.0548	0.0629	0.0396	0.0374	0.0465
First stage F	818.9	773.4	588.5	446.5	750.9	1083	832.3
R^2	0.128	0.128	0.131	0.117	0.132	0.171	0.135
Observations	429,007	427,942	281,571	181,394	342,429	401,957	373,106
Registers	1,032	1,032	672	433	815	1,031	887
Household-register FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable is daily register-level log(kWh). A household may have more than one register. Marginal price is the per kWh charge. Price is instrumented using treatment status (10%, 25%, 50% discount) × treatment period. Regressions are restricted to the following subsamples. Column (1): households that definitely opened the treatment email. Column (2): household-days before any intermediate bills were received. Column (3): households that are not signed up for automatic bill payment. Column (4): households with two or fewer residents. Column (5): households that have not opted for a green power plan. Column (6): household-days with consumption levels exceeding 4 kWh/day. Column (7): households that have central or unit air-conditioning. Standard errors (in parentheses) are clustered at the household level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Additional Results

The Appendix presents further subsample results that support our experimental finding of inelastic electricity demand. Table B.4 shows that electricity demand remains unchanged even for households with appliances that variably use relatively large levels of electricity. In Table B.5 we divide the sample into quartiles based on levels of use over the course of the year preceding the experiment. The largest point estimate is –0.036, for households with the highest level of baseline use, which is still not statistically-different from zero. Only for this group are we unable to reject a price elasticity of –0.10 as the 95% confidence interval reaches –0.117.

Figure A.5 presents the elasticities for all subgroups by hour-of-day. The largest, statistically-insignificant point estimates are from households with the highest level of baseline use during early evening hours. Otherwise, households consistently maintain their previous patterns of electricity use despite the experimental price discounts.

Finally, we check whether our lack of response could be due to a run of mild-temperature days, which could have limited discretionary electricity use such as air conditioning. Appendix Figure A.6 Panel (c) shows the distribution of temperatures realized within our sample period. The experimental period in fact covers a range of days that are likely to require energy for heating (those below 65F) and days that are likely to require energy for cooling (those above 65F), with a few very hot days.

Despite this variation, we find electricity use does vary with temperature over our sample period. Panel (a) of Figure A.6 shows average daily consumption of electricity as a function of daily maximum temperature for treatment and control groups. Panel (b) shows the same for the sample of households with unit or central air conditioning. As expected, consumption is highest on hot days when households use electricity for cooling. But there are no differences across treatment and control households in the relationship between electricity consumption and temperature, even among households with air-conditioning.

Finally, as mentioned above, in Table B.2 of the Appendix, we show that households in our experimental sample have similar pre-treatment daily energy usage, fixed charges, variable charges, and demographics to our entire working sample of more than 220,000 households from our partner retailer. All differences in these observable characteristics are very small. This further supports our conclusions regarding demand inelasticity in response to our large experimental monthly price shocks; it is not likely due to sample selection.

4. CONCLUSION

In this paper we document substantial demand inelasticity in the presence of transitory monthly price shocks that are larger than what has been previously considered in observational studies. We also extend the domain of experimental designs in energy pricing experiments to consider monthly price changes as opposed to high-frequency (e.g., hourly) experimental price shocks.

What do we learn from our price experiment? Implementing cost-reflective pricing in the residential sector could expend significant political capital for limited demand-driven efficiency gain in the short-to-intermediate run. We find it surprising that price reductions of 50% or higher yield minimal changes in electricity demand over this time horizon. This result is in line with large behavioral and technological adjustment costs in residential electricity use, and households choosing to ignore long-but-not-too-long price shocks. The behavior is not inconsistent with rational inattention.[9]

This paper documents the behavioral response to one type of price variation. There are of course other types of price variation that we cannot speak to. Time-of-day or critical peak pricing are interesting alternatives for load shifting that could potentially mitigate cost-increasing inefficient grid investments. Potential asymmetric demand responses are another area that we believe should be explored. We are optimistic that the emergence of price experiments at various time horizons and with varying price shocks in terms of their direction and magnitude will, collectively, help paint a complete picture of energy demand with which to inform public policy.

The demand inelasticity that we find at the monthly level to substantial price shocks underscores the challenge policymakers face in getting consumers to shift energy consumption behavior over this horizon. Our study highlights the need to understand how effective prices can be at inducing behavioral change at a range of different time horizons.

ACKNOWLEDGMENTS

We thank Yann Burden, Tatyana Deryugina, Koichiro Ito, Erica Myers, David Rapson, and workshop participants at UC Davis, University of Melbourne, Monash, and the Midwest Energy Fest. We acknowledge financial support of the Australian Research Council Linkage Grant Award LP140100099, a University of Melbourne Faculty of Business and Economics Faculty Research Grant, and from the University of Melbourne's Centre for Market Design. All errors, omissions, and views are our own. The Online Appendix for the paper can be downloaded from https://sites.google.com/view/dprbyrne/researchhttps://sites.google.com/view/dprbyrne/research.

9. Our experiment may have induced income effects in other areas of consumption that have gone unmeasured. However, previous research has abstracted from such income effects arguing they are likely very small. Indeed, this assumption underpins the use of quasi-linear utility in quasi-experimental empirical studies of electricity and water demand including Borenstein (2009), Ito (2014), and Wichman (2014). Given the importance of this assumption, testing for income effects arising from energy price shocks is an important area of future research.

REFERENCES

ACCC (2018). "Restoring Electricity Affordability and Australia's Competitive Advantage" Commonwealth of Australia Technical report.

Alberini, A. and M. Filippini (2011). "Response of Residential Electricity Demand to Price: The Effect of Measurement Error." *Energy Economics* 33(5): 889–895. https://doi.org/10.1016/j.eneco.2011.03.009.

Alberini, A., O. Khymych, and M. Ščasný (2019). "Response to Extreme Energy Price Changes: Evidence from Ukraine" *The Energy Journal* 40: 189–212. https://doi.org/10.5547/01956574.40.1.aalb.

Allcott, H. (2011a). "Rethinking Real-time Electricity Pricing." *Resource and Energy Economics* 33(4): 820–842. https://doi.org/10.1016/j.reseneeco.2011.06.003.

Allcott, H. (2011b). "Social Norms and Energy Conservation." *Journal of Public Economics* 95(9): 1082–1095. https://doi.org/10.1016/j.jpubeco.2011.03.003.

Arup, Energeia, F. Economics, I. for Sustainable Futures, and U. of Technology Sydney (2014). "Smart Grid Smart City: Shaping Australia's Energy Future." Executive Report.

Borenstein, S. (2009). "To What Electricity Price do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing." *UC Berkeley Haas School of Business Working Paper*.

Burke, P.J., A. Abayasekara, et al. (2018). "The Price Elasticity of Electricity Demand in the United States: A Three-Dimensional Analysis." *The Energy Journal* 39(2): 123–145. https://doi.org/10.5547/01956574.39.2.pbur.

Bushnell, J.B. and E.T. Mansur (2005). "Consumption Under Noisy Price Signals: a Study of Electricity Retail Rate Deregulation in San Diego." *Journal of Industrial Economics* 53(4): 493–513. https://doi.org/10.1111/j.1467-6451.2005.00267.x.

Byrne, D.P., A. La Nauze, and L.A. Martin (2018). "Tell Me Something I Don't Already Know: Informedness and the Impact of Information Programs." *Review of Economics and Statistics* 100(3): 510–527. https://doi.org/10.1162/rest_a_00695.

Costa, D.L. and M.E. Kahn (2013). "Energy Conservation "Nudges" and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment." *Journal of the European Economic Association* 3(11): 680–702. https://doi.org/10.1111/jeea.12011.

Costa, F. and F. Gerard (2015). "Hysteresis and the Social Cost of Corrective Policies: Evidence from a Temporary Energy Saving Program." *NBER Working Paper* 24608. https://doi.org/10.2139/ssrn.2666704.

Cuddington, J.T. and L. Dagher (2015). "Estimating Short and Long-Run Demand Elasticities: A Primer with Energy-Sector Applications." *The Energy Journal* 36(1): 185–209. https://doi.org/10.5547/01956574.36.1.7.

Deryugina, T., A. MacKay, and J. Reif (2020). "The Long-Run Elasticity of Electricity Demand: Evidence from Municipal Electric Aggregation." *American Economic Journal: Applied Economics* 12(1):86–114. https://doi.org/10.1257/app.20180256.

Faruqui, A. and S. Sergici (2010). "Household Response to Dynamic Pricing of Electricity: A Survey of 15 Experiments." *Journal of Regulatory Economics* 38(2): 193–225. https://doi.org/10.1007/s11149-010-9127-y.

Faruqui, A. and S. Sergici (2011). "Dynamic Pricing of Electricity in the Mid-Atlantic Region: Econometric Results from the Baltimore Gas and Electric Company Experiment." *Journal of regulatory economics* 40(1): 82–109. https://doi.org/10.1007/s11149-011-9152-5.

Ferraro, P.J. and M.K. Price (2013). "Using Non-Pecuniary Strategies to Influence Behavior: Evidence from a Large Scale Field Experiment." *Review of Economics and Statistics* 95(1): 64–73. https://doi.org/10.1162/REST_a_00344.

Fowlie, M., C. Wolfram, C.A. Spurlock, A. Todd, P. Baylis, and P. Cappers (2017). "Default Effects and Follow-on Behavior: Evidence from an Electricity Pricing Program." *NBER Working Paper* 23553. https://doi.org/10.3386/w23553.

Gillan, J. (2017). "Dynamic Pricing, Attention, and Automation: Evidence from a Field Experiment in Electricity Consumption." *Energy Institute at Haas Working Paper* 284.

Holland, S.P. and E.T. Mansur (2006). "The Short-Run Effects of Time-Varying Prices in Competitive Electricity Markets." *The Energy Journal* : 127–155. https://doi.org/10.5547/ISSN0195-6574-EJ-Vol27-No4-6.

Ito, K. (2014). "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review* 104(2): 537–563. https://doi.org/10.1257/aer.104.2.537.

Ito, K., T. Ida, and M. Tanaka (2018). "The Persistence of Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand." *American Economic Journal: Economic Policy* 10: 240–67. https://doi.org/10.1257/pol.20160093.

Jessoe, K. and D. Rapson (2014). "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review* 104(4): 1417–1438. https://doi.org/10.1257/aer.104.4.1417.

Reiss, P.C. and M.W. White (2005). "Household Electricity Demand, Revisited." *Review of Economic Studies* 72(3): 853–883. https://doi.org/10.1111/0034-6527.00354.

Reiss, P.C. and M.W. White (2008). "What Changes Energy Consumption? Prices and Public Pressures." *RAND Journal of Economics* 39(3): 636–663. https://doi.org/10.1111/j.1756-2171.2008.00032.x.

Ros, A.J. (2017). "An Econometric Assessment of Electricity Demand in the United States Using Utility-specific Panel Data and the Impact of Retail Competition on Prices." *The Energy Journal* 38(4): 73–99. https://doi.org/10.5547/01956574.38.4.aros.

Sexton, S. (2015). "Automatic Bill Payment and Salience Effects: Evidence from Electricity Consumption." *Review of Economics and Statistics* 97(2): 229–241. https://doi.org/10.1162/REST_a_00465.

Wichman, C. (2014). "Perceived Price in Residential Water Demand: Evidence from a Natural Experiment." *Journal of Economic Behavior & Organization* 107(2): 308–323. https://doi.org/10.1016/j.jebo.2014.02.017.

Wolak, F.A. (2011). "Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment." *American Economic Review Papers and Proceedings* 101(3): 83–87. https://doi.org/10.1257/aer.101.3.83.