Title
Essays in Energy and Environmental Economics

Author
Buchsbaum, Jesse Fleischer

Publication Date
2022

Peer reviewed|Thesis/dissertation

Essays in Energy and Environmental Economics

by

Jesse Buchsbaum

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Meredith Fowlie, Chair
Professor James Sallee
Professor Severin Borenstein

Spring 2022

Essays in Energy and Environmental Economics

Abstract

Essays in Energy and Environmental Economics

by

Jesse Buchsbaum

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Meredith Fowlie, Chair

The energy industry is undergoing rapid transformation in the United States. Climate change continues to advance, leading to policies aimed at reducing emissions of greenhouse gases. Meanwhile, the costs of clean energy technologies are declining, leading to increased adoption by firms and consumers. These fundamental changes have led both firms and consumers to contend with a suite of new challenges. For consumers, emissions reduction policies and changes in generation costs have impacted retail prices. Firms, meanwhile, have needed to adapt to shifting grid policies and input costs. In this dissertation, I strive to understand how consumers and firms have responded to these changes in the electricity industry by asking two critical questions of economics: first, how do consumers respond to prices; and second, what are the spillover effects of a policy change? The first two chapters of this dissertation are devoted to the former question with particular attention to responses in the long run, while the final chapter examines the latter with focus on wholesale electricity markets.

In the first chapter of this dissertation, I study how electricity consumers respond to electricity prices in the short and medium run and evaluate heterogeneity in those responses along the income dimension. While the existing literature focuses primarily on the short run, understanding the dynamics of consumer demand over time is critical, as habit formation and durable good investment play an important role. I leverage a novel source of exogenous spatial price variation in combination with dynamic changes in price to evaluate how the responses of electricity consumers vary over time. Consumers are somewhat responsive in the short run, with a price elasticity of -0.36. Responses diminish over time but display some persistence, with an elasticity of -0.12 with respect to a three-year lagged prices. In addition, I evaluate the role that income plays in consumer response, finding that low-income consumers are less responsive to changes in price in both the short and medium run. These findings demonstrate the importance of accounting for consumption dynamics, especially in a setting where habit formation and durable goods play significant roles.

In the second chapter of this dissertation, I study how electricity consumers respond to

electricity prices in the long run. Long-run elasticities are difficult to empirically estimate, and credible quasi-experimental estimates of long-run elasticities are rare, especially in the energy economics literature. However, long-run elasticities are crucial for calculating welfare, forecasting demand, and evaluating policy. Here, I leverage a novel source of plausibly exogenous long-lasting price variation for one of the first quasi-experimental estimates of the long-run price elasticity of demand for residential electricity consumers. I find that consumers are much more responsive to prices in the long run than the short run, with a long-run elasticity estimate of -2.4. Furthermore, I explore some of the mechanisms driving this price response, and find that residential adoption of rooftop solar alone can explain 26% of the observed response in consumption. My findings highlight the impact of price-based policies, and suggest that these types of policies may be more effective than previously thought in inducing energy transitions to cleaner technologies.

In the third chapter of this dissertation, in collaboration with Catherine Hausman, Johanna L Mathieu, and Jing Peng, I explore the spillover effects of policy changes in wholesale electricity markets. In electricity markets, generators are rewarded both for providing energy and for enabling grid reliability. The two functions are compensated with two separate payments: energy market payments and ancillary services market payments. We provide evidence of changes in the generation mix in the *energy* market that are driven by exogenous changes in an *ancillary services* market. We provide a theoretical framework and quasi-experimental evidence for understanding the mechanism, showing that it results from the multi-product nature of conventional power plants combined with discontinuities in costs. While research in economics typically focuses solely on the energy market, our results suggest that spillovers between markets are important as well. Furthermore, policy changes relating to grid operations, grid reliability, or climate change could have unintended effects.

Contents

List of Figures

List of Tables

Acknowledgments

My experience in the Agricultural and Resource Economics Department at the University of California Berkeley has been incredibly formative and life-changing, and there is no place I would have rather completed my PhD.

Professor Meredith Fowlie, my committee chair has inspired me as a researcher and as a person with her intellectual creativity and rigor, her teaching and framing of complex issues, and her passion for research that impacts the world. I am consistently in awe of her creative problem solving and feel incredibly lucky to call myself her student. Professors James Sallee and Severin Borenstien have constantly provided invaluable guidance, ideas, and feedback and I wouldn't be where I am today without their mentorship. Professor Catie Hausman has had a tremendous influence on the way that I approach the research and paper-writing process with her intellectual honesty and rigor, and I am lucky have her as a collaborator and role model. I am eternally grateful to these four professors, without whom I would not be here.

The Agricultural and Resource Economics Department and the Energy Institue at Haas are collaborative atmospheres that have allowed me to be benefit from numerous professors – in particular, David Zilberman, Lucas Davis, Maximilian Auffhammer, Sofia Villas-Boas, and Catherine Wolfram have helped shape me into the researcher that I am today.

I am so grateful to Matthew Tarduno, Jenya Kahn-Lang, and Marshall Blundell, for their constant friendship, the many, many conversations that helped hone ideas into research questions, and for always being willing to lend an ear and a hand. I owe huge thanks to Alex, Jaecheol, Livia, Matthew, Sebastien, and Wenjun – I can't imagine a better group of classmates.

My family has provided unwavering love and support throughout my studies. To my mom, Cathy, who inspired me to become a researcher, you push me to think about research in ways that I have never considered. To my dad, Andy, who never shies from hours-long conversations about nuanced energy topics, thank you for fostering my interest in climate change and love for the environment. To my brother, Seth, who has been one of my role models since I was born, I have you to thank for my love of diving headfirst into the weeds of complex issues.

Finally, to Michelle, thank you for being my rock constantly and without fail. You have been there to pick me up at my lowest points and celebrate with me in my victories. You have believed in me even when I didn't believe in myself and I wouldn't be who I am without you.

Chapter 1

How do residential electricity consumers respond to price? Dynamics and heterogeneity[1]

1.1 Introduction

One fundamental questions of economics asks how consumers respond to prices. Price elasticities of demand are crucial in a variety of settings, to construct demand curves and forecast future demand, to calculate welfare, and to design policy. These elasticities take on even more importance in electricity markets, where the planning decisions of electricity generators, distribution utilities, and grid planners, as well as policy decisions on issues of energy and climate change, all rely on accurate predictions future demand of future. However, the existing literature primarily focuses on average elasticities the short-run, estimating how consumers respond to electricity prices over very short time horizons and often assuming a single elasticity across the full population.

While the short-run is important, planning decisions made by electricity industry stakeholders have much longer time horizons often spanning multiple decades. For instance, solar arrays have a typical lifespan of 25 years, transmission lines last for 35-40 years, and coal-fired power plants often last up to 50-60 years. Understanding how consumers respond to prices over a longer time horizon is critical to inform those planning decisions, as well as the policy that influences investments that will be locked in for decades.

Furthermore, understanding potential heterogeneity is critical for both planning purposes and for understanding welfare. To the extent that different types of consumers respond to

[1] I thank Meredith Fowlie, James Sallee, Catie Hausman, and Severin Borenstein for invaluable advice, feedback, and mentorship throughout this project. I am also grateful to Marshall Blundell, Ellen Bruno, Fiona Burlig, Lucas Davis, Jenya Kahn-Lang, Louis Preonas, Matthew Tarduno, David Zilberman, and participants at various seminars for helpful comments. Finally, thank you to Robert Lucadello for invaluable assistance accessing and understanding the utility data used for this project. I do not have any financial relationships that relate to this research.

prices differently, different types of transmission and distribution investments are necessary in different geographic areas. And the incidence of price-based policies differs dramatically if there are substantial heterogeneities in how consumers respond to prices.

In this paper, I leverage a novel source of persistent spatial price variation to estimate short- and medium-run price elasticities for residential electricity customers in California, and to estimate heterogeneity across income. This price variation is driven by a subtle feature of California's pricing regime. In the increasing block pricing rate structure used throughout California, marginal prices increase when electricity usage exceeds a certain threshold. Because of differences in heating and cooling needs for households across the different climates of California, utilities set these thresholds to different levels depending on where a consumer lives. Within Pacific Gas & Electric's (PGE's) service territory, there are ten different baseline territories[2], with the boundaries for these territories often determined according to discontinuities in a household's elevation. These boundaries have led to long-lasting persistent price variation since they were established in 1982, with one side of the border consistently facing higher prices than the other. I leverage these price discontinuities to estimate elasticities across different time intervals.

In the short run, I follow the methods of Ito (2014), using a simulated instrument to isolate exogenous variation in the price schedule over time. I find that electricity consumption is relatively inelastic, with an elasticity of -0.36 in my preferred specification.

I expand this approach to estimate medium-run price elasticities. Specifically, I estimate how current period consumption is impacted by past changes in prices. Consumers are most responsive to prices in the present period, but do still respond to lagged prices, with estimated elasticities of -0.18, -0.19, -0.13, and -0.12 with respect to prices in the contemporaneous period, one year prior, two years prior, and three years prior respectively.

However, substantial heterogeneity exists in price responsiveness according to a consumer's income level. There are overlapping explanations that might explain this conflict – electricity bills may be more salient to low-income households as they have less discretionary income than higher income households. However, higher income households typically have more appliances, leading to more margins for response to price changes. Furthermore, durable goods that reduce consumption often have high capital costs, leading to potentially greater adoption among higher income households (Borenstein, 2017). I estimate elasticities for households with different income levels, finding that low-income consumers are less responsive to changes in prices in both the short and medium run.

This paper contributes to two distinct literatures. First, there is a large literature that estimates short- and medium-run elasticities for residential electricity customers. This paper builds primarily on works by Ito (2014), Shaffer (2020), and Brolinson (2019) that estimate short-run elasticities with respect to both marginal and average prices in settings with increasing block pricing. I build on this literature in two ways: first, I leverage a novel source of within-utility cross-sectional price variation, which improves on the current literature by

[2]These baseline territories were finalized in 1982 and have largely stayed the same since 1982 (PG&E, 2020).

reducing the potential for confounding non-price effects[3]. Additionally, I expand on existing methods to estimate dynamics in the medium-run, showing that consumers continue to respond to prices lagged up to four years but that responsiveness diminishes over time. There are numerous other papers that estimate short-run elasticities for residential electricity customers. A 2018 meta-analysis (Zhu et al., 2018) of papers estimating price elasticities of demand for residential electricity customers estimates a mean short-run elasticity of -0.23. In this setting, I estimate a short-run elasticity of -0.36, grounding this analysis squarely within the existing literature.

Second, this paper contributes to the literature exploring how household income impacts price responsiveness among residential electricity customers. Evidence in this literature is somewhat conflicting – Alberini, Gans and Velez-Lopez (2011) and Reiss and White (2005) find that price elasticities of demand are highest among the poorest households and monotonically decrease as income grows, while Brolinson (2019) and Schulte and Heindl (2017) find that wealthier households are more responsive to prices. This paper is first to separately quasi-experimentally estimate short- and medium-run price elasticities by income. I show that higher income households are much more responsive to prices in the short- and medium-run, suggesting that adoption of costly durable goods may be a primary driver of income heterogeneity in short-run price responsiveness. However, more research is necessary to capture the mechanisms driving this income heterogeneity, as well as whether the heterogeneity persists in the long run. The second chapter of this dissertation begins to speak to these issues.

The paper proceeds as follows: Section 2 discusses background on the setting and measures of heterogeneity used; Section 3 presents the data and empirical strategy; Section 4 presents estimates of short- and medium-run price elasticities; and Section 5 concludes.

1.2 Background

Increasing block pricing and baseline territories

The setting for this paper is Pacific Gas & Electric (PG&E), a large investor owned utility company in Northern California. PG&E uses a non-linear price schedule called increasing block pricing to set prices for electricity. This pricing mechanism is similar to a graduated income tax, where higher levels of usage face a higher marginal price. As an illustrative example, suppose Customer A uses 1000 kWh in a month. In one region of PG&E's service territory (Baseline Territory Q), she would pay 18 cents per kWh for the first 888 kWh (Tier 1) she uses and 24 cents per kWh for the next 112 kWh (Tier 2) she uses, leading to a total bill of $186.72. In this two-tier example, 888 kWh is the monthly *baseline allowance* –

[3] Brolinson (2019) leverages a similar source of source of cross-sectional variation, but is limited by data on a much sparser set of households. Here, I use a rich set of households, allowing me to directly compare households on either side of the the border in order to credibly estimate elasticities.

after reaching the baseline, all further consumption is in the second tier and faces the higher marginal price.

There is a great deal of variation in climate even within utility service territories. PG&E's service territory includes both Fresno, with an average June high temperature of 92 degrees, and San Francisco, with average June high temperatures of 60 degrees. Because of this wide gap, electricity demand to meet basic heating and cooling needs across a utility's service territory is not equal. As such, customers are divided into climate territories that determine the baseline – in other words, the level of electricity that can be used before the higher marginal price takes effect. Furthermore, baselines are different in summer and winter, as well as for customers with electric versus gas heat. Continuing with our illustrative example, suppose that Customer B has identical usage, but lives in a territory (Baseline Territory T) where the baseline allowance is 447 kWh per month. She pays the same price, 18 cents per kWh, for the first 447 kWh, but then pays 24 cents for the next 553 kWh, leading to a total bill of $213.18 – about $26 higher than Customer A for the exact same level usage.

PG&E divides its service area into ten different baseline territories[4], as shown in Figure 1.1. These baseline territories were established in 1982[5] by the California legislature, and adopted by by the California Public Utility Commission in 1983. Between 1983 and 1990, the CPUC continued to make small changes to where the baseline territory boundaries lay. From 1990 to 2019, the baseline territory map stayed the same, meaning that the baseline territory map is constant over the sample period of this study. More generally, baselines are determined based on two potential factors: geopolitical demarcations (e.g. zip code/city/county boundaries, roads) and elevation discontinuities. For example, Santa Barbara County is divided into Territories R, T, and X according to geopolitical demarcations. However, Trinity Country is divided into Territories X, Y, and Z, where residents of Trinity County below 2,000 feet of altitude are in Territory X, residents between 2,001 feet and 4,500 feet are in Territory Y, and residents above 4,500 feet are in Territory Z. A full list of baseline territory boundaries defined by elevation is provided in Appendix A.2[6].

To determine the level of each baseline, PG&E set quantities so that 50 to 60 percent of expected residential electricity consumption in each climate zone is set as "baseline" consumption (equivalently, so that 50 to 60 percent of consumption is in Tier 1)[7]. Figure 1.2 shows the daily baselines from 2008 to 2020 for each baseline territory within PG&E. There is a great deal of variation both across baseline territories and even within territories from

[4]Note that PG&E's baseline territories are different boundaries than the California Electricity Commission's (CEC's) "climate zones," which are used to determine building codes. Baseline territory boundaries are nearly universally separate from CEC climate zone boundaries, with a very small number of exceptions.

[5]A precursor to baseline territories was established in 1976, called "climate bands," though there were only four climate bands based purely on heating degree-days.

[6]A full list of baseline territory definitions including those defined by non-elevation definitions can be found at https://www.pge.com/tariffs/assets/pdf/tariffbook/ELEC_PRELIM_A.pdf.

[7]One might be concerned that this could lead to endogeneity, where the actions of a household impact the baseline allowance in future periods. I assume that individual households do not exhibit market power, an assumption supported by the fact that each baseline territory contains at least 6,000 households.

Figure 1.1: PGE baseline territories (PGE, 2020)

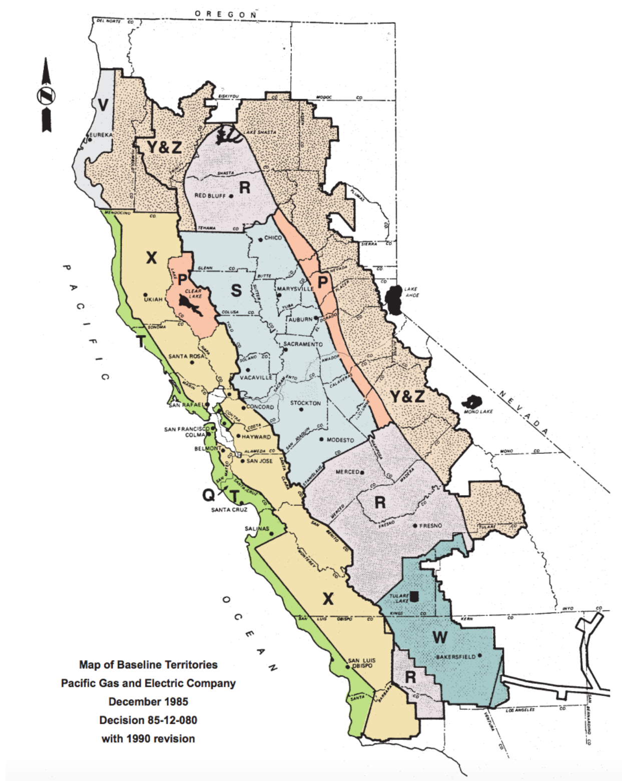

summer to winter and between customers with electric versus gas heat. Over the course the sample for this study (2008 to 2020), baseline quantities change four times.

Because baseline territories are used to determine baselines and therefore marginal and average prices, and customers who live very close to one another might be assigned to different baseline territories, there is variation in the prices faced by customers close to the baseline territory borders. Returning once more to the illustrative example, recall that Customer B (in Territory T) face a monthly bill \$26 higher than Customer A (in Territory Q) for the identical level of usage. Because Territories T and Q are divided by an elevation discontinuity, these two customers might live in the same neighborhood, face the same climatic conditions, and still face substantially different monthly bills. Figure 1.3 exhibits the full price schedule faced by customers in Territory Q compared with T during the winter months of 2017.

Note that the existing literature, including Ito (2014) and Shaffer (2020), often use utility service territory boundaries as a source of exogenous spatial price variation. Utility service territory boundaries, however, are vulnerable to confounding non-price factors along the utility border, such as utility-specific programs and potential household selection effects. Because baseline territories boundaries are within a single utility's service territory, they are not subject to the same confounding effects to prices. Furthermore, utility service boundaries

Figure 1.2: Baselines over time

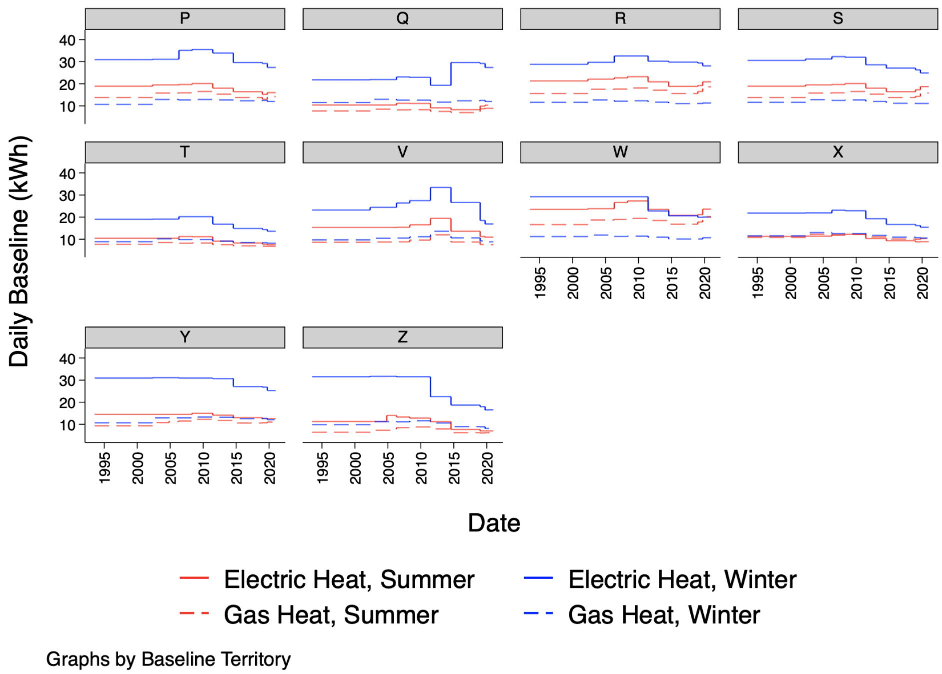

Note: This figure shows the level of baselines over time for each baseline territory and for households with electric heat in summer and winter. Data are shown from 1995 to 2020, as this is the period for which public data is available.

are often limited to only a narrow geographic range. PG&E's baseline territories cover a much broader spatial area, allowing for a more diverse set of households that may be more representative of the broader population.

Not only do marginal prices vary spatially across baseline territory borders, but there is a great deal of price variation over time as well. Figure 1.4 shows the evolution of each price over time for the standard residential tariff, E-1. Note that not only is there price variation within each tier, but there is a compressing of tiers that occurs in December 2016, when rate E-1 moves from four tiers to three. This provides useful identifying variation over time, that can be used in combination with the spatial variation resulting from climate zone boundaries.

With variation in baselines across both space and time, it's important to consider exactly what variation in baselines comes from each source of variation. In Table 1.1, I decompose the variation in baseline territories according to space and time, by using baseline territory and month of sample fixed effects, along with controls for the other determinants of baselines – electric versus gas heating and summer versus winter. Column 1 shows how much of the

Figure 1.3: Price variation in Territory Q versus Territory T

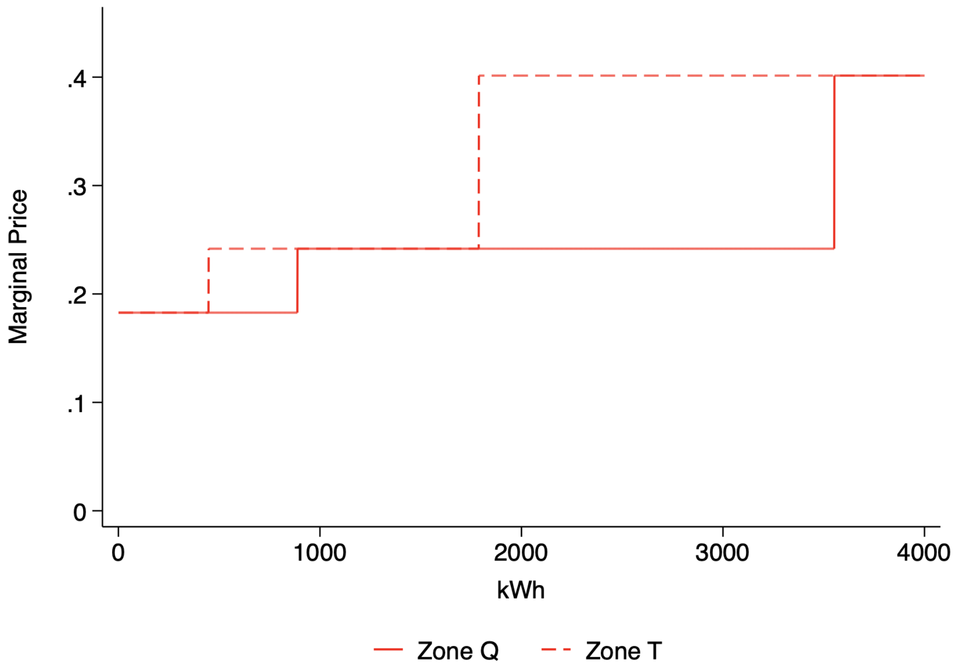

Note: This figure shows the marginal price path for baseline territory Q compared with baseline territory T in January 2017 for customers with electric heat. Territories Q and T are directly adjacent to one another.

variation in baselines can be explained by controls alone, while Columns 2 through 4 add spatial and time series fixed effects sequentially to demonstrate the extent to which each type of fixed effect explains variation in the baseline. The vast majority of variation in baselines not accounted for by the controls can be explained by spatial fixed effects, with a small amount explained by temporal fixed effects, suggesting that cross-sectional variation plays a major role in creating price differences. It is also worth noting that, as expected, almost all variation in baselines (99%) can be explained by spatial and temporal fixed effects in combination with controls for the other determinants of baselines.

Figure 1.4: PG&E price evolution over time

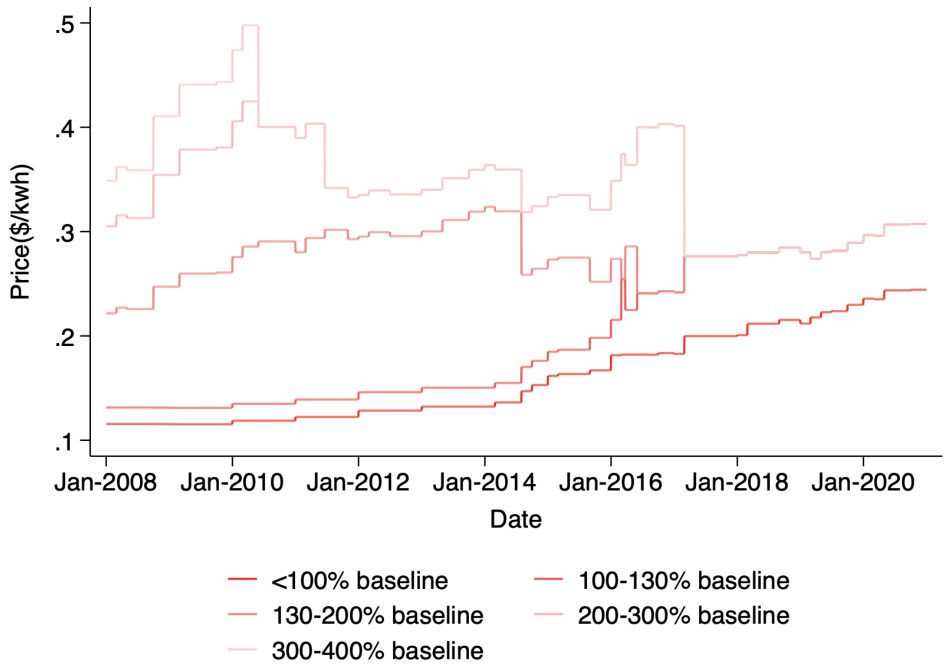

Note: This figure shows the price schedule over time for PG&E's standard default non-time varying tariff (Tariff E-1) from 2008 to 2020. The darkest line shows marginal prices for the lowest level of usage over time (under 100% of the baseline), while lighter shades show marginal prices for higher levels of usage.

Measures of heterogeneity

When estimating how consumers respond to prices, one critical component is to understand who is responding. Many papers, including Shaffer (2020) and Alberini, Gans and Velez-Lopez (2011) show there are significant heterogeneities in how customers respond to prices in their energy choices, driven by factors including information, salience, access to capital, and more. Different responses across customer groups induces heterogeneity in welfare changes. While in theory, transfers could be used to equitably redistribute any gains (or losses) from a policy, work by Sallee (2019) emphasizes the challenge that targeting presents, especially in the context of energy policy. In a context with limited transfers, understanding these mechanisms and heterogeneities is highly important for designing and evaluating policy, especially when equity is a policy objective.

In this setting, the primary demographic variable of interest is income. Because adoption of durable goods requires access to capital, we might expect that higher income customers are more likely to invest in durable goods that impact long-run price responsiveness. On

Table 1.1: Baseline variation decomposition

baseline_{it}	(1)	(2)	(3)	(4)
R^2	0.70	0.93	0.75	0.99
Electric x Summer	Yes	No	No	No
Electric x Summer x BT	No	Yes	No	No
Electric x Summer x MofS	No	No	Yes	No
Electric x Summer x BT x MofS	No	No	No	Yes

*Notes: This table shows the results of four regressions, all with the length of baseline as the dependent variable. Unit of observtion is a customer account by month. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

the other hand, past work (Alberini, Gans and Velez-Lopez, 2011; Reiss and White, 2005) seems to indicate that low-income consumers tend to be more aware of their bills and may therefore may be more responsive to price fluctuations, especially in the short-run.

While I do not directly observe income at a customer level[8], there are two primary ways that I explore demographic heterogeneity. First, I use CBG-level data on income from the American Communities Survey to compare high-income CBGs with lower-income CBGs. Census data has numerous measures of income; my preferred measure in this work is the number of households under the federal poverty line (FPL).

Second, I use a proxy for income that is observed at the account-level: participation in the California Alternative Rates for Energy (CARE) program, following Auffhammer and Rubin (2018) among others. CARE is a program that is available to all energy customers in the state of California with incomes below 200% of the FPL. Customers enroll directly through PG&E, who conducts random income verification checks to ensure that customers are compliant with the income requirements. PG&E estimates that 95% of eligible customers are enrolled in CARE. While there is some endogeneity in which customers are enrolled in CARE that may be correlated with information and bill attention, the high participation rate of CARE implies that it is a good proxy for income.

In Figure 1.5, I compare how CARE participation correlates with CBG-level FPL deciles. While there is strong correlation between CARE enrollment and the number of households within a CBG below the federal poverty line, there is substantial heterogeneity in income levels within each CBG. There are numerous CARE enrollees across all CBGs, including those with the lowest number of households below the federal poverty line. Throughout the paper, CARE will be used as the primary proxy for income, while heterogeneity across

[8]There are two reasons that I don't observe income: (1) high-quality income data at a consumer level are extremely difficult to access; and (2) utility data is anonymized, so that I couldn't match my data with an external income dataset, even could access it.

Figure 1.5: CARE enrollment versus CBG-level FPL decile

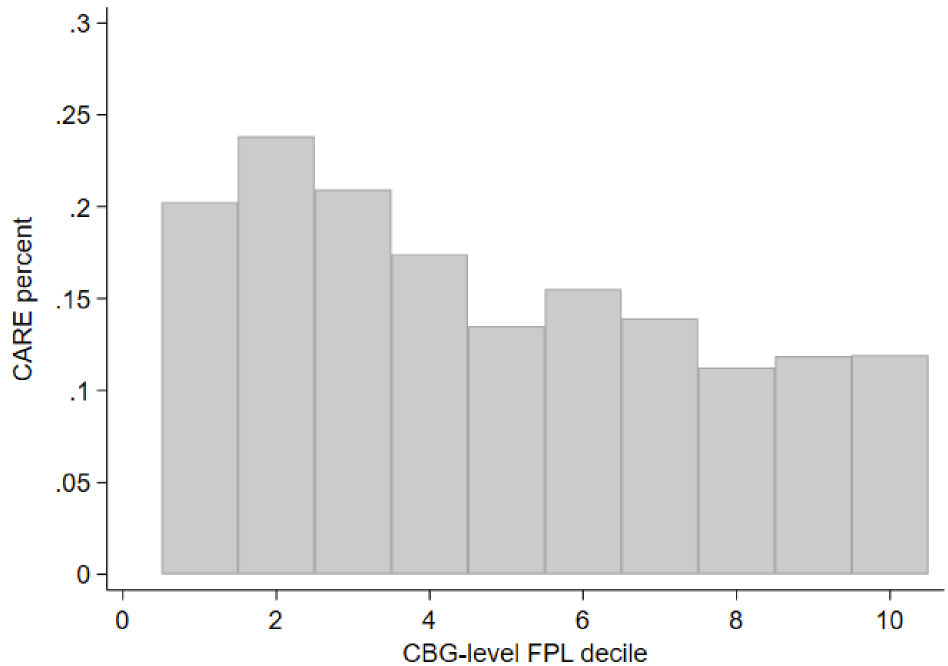

Note: In this figure, CBG-level Federal Poverty Line (FPL) is defined as the proportion of households below the FPL within a CBG. The x-axis shows the decile of this measure of income, while the y-axis represents the proportion of households enrolled in CARE within a CBG.

federal poverty line deciles will be shown in the Appendix.

1.3 Research Design and Data

Data

For this study, I use account-level billing data for all PG&E electricity customers from 2008 to 2020[9]. In each monthly record, I observe data at the monthly level on electricity usage[10], billing, and adoption of durable goods (e.g. solar panels, electric heat, energy efficiency), as well as some limited demographic information. While I do not observe customer names

[9]PG&E has granted me access to this data under a confidentiality agreement

[10]The electricity usage data that I use throughout this paper is *net* monthly electricity consumption. For solar customers who both generate and consume electricity, their net consumption is the difference between their gross monthly consumption and their gross monthly generation.

or addresses, I do observe a customer's Census Block Group (CBG), and I merge PG&E's data with census data from the 2017 5-Year American Community Survey (ACS) to obtain demographic information. In my sample, there are an average of 588 households in each CBG.

In my empirical analysis, I restrict my sample in two ways: first, I omit households with non-standard baselines such as medical baselines; and second, I limit the sample to CBGs where price variation exists. To ensure sufficient overlap within census block groups, I restrict my sample to CBGs with at least 50 customers in each of multiple different baseline territories. This sampling restriction ensures that I compare only customers who face similar weather, climate, and even neighborhood effects, but differ in the electricity baseline and prices that they face. This sampling restrictions result in 235,097 different customer accounts, located in 132 different CBGs.

In Table 1.2, I present summary statistics. I show the means and standard errors for my restricted sample, as well as the broader PG&E sample of customers. The restricted sample has higher usage and baseline territories on average, driven by higher adoption of electric heat than in the full PG&E sample. In addition, while solar adoption in the restricted sample is 5%, it is just 3% in the rest of PG&E's service territory.

Empirical strategy

In order to estimate how customers respond to prices in the short, medium, and long run, I rely on three primary sources of identifying variation: (1) spatial discontinuities in the baseline (and potentially price) that a customer faces; (2) temporal variation in prices; and (3) temporal variation in baselines[11]. In combination, these three sources of variation lead to prices that vary both in time and across space.

I leverage spatial discontinuities in the baseline using sampling restrictions and spatial fixed effects. As mentioned in the previous subsection, I restrict my sample to CBGs with at least 50 customers in multiple different baseline territories. In all specifications estimating price elasticities of demand (including the short, medium, and long run), I include a CBG fixed effect. CBGs are relatively small, and are explicitly drawn to be homogeneous along demographic traits Census (1994). For this reason, my identifying assumption is that customers within the same CBG would be similar in their electricity consumption and durable good adoption but for the differences in baselines (and therefore prices) that they face.

Under this identifying assumption, I pursue several different empirical approaches. First, to anchor this work within the existing literature, I estimate short-run price elasticities using all three sources of identifying variation, following the methods of Ito (2014). Next, I estimate price elasticities in the medium-run, again using all three sources of variation. Within each time scale, I explore heterogeneity across different types of customers in order to acquire a more complete picture of the factors driving the population-wide effects.

[11]Because month-of-sample fixed effects are included in all specifications, temporal variation in baselines is limited to policy changes and does not include seasonal variation.

Table 1.2: Summary statistics

	In Sample			***All PG&E***		
	Mean	Std.Dev.	Accounts	Mean	Std.Dev.	Accounts
Monthly usage (kWh)	569.53	626.72	235,097	395.34	644.39	21,411,577
Monthly baseline (kWh)	472.37	183.04	235,097	349.19	427.85	21,390,754
Average price (kWh)	0.23	0.49	234,138	0.20	0.50	21,265,123
Percent electric heat	0.45	0.50	235,097	0.21	0.40	21,445,946
Percent solar	0.05	0.21	235,097	0.03	0.18	21,445,946
Percent CARE	0.23	0.42	235,097	0.25	0.43	21,445,946

1.4 Short-, and medium-run responses to prices

Short-run responses to prices

I follow the methodology presented by Ito (2014) in order to estimate short-run responses to both average and marginal prices. In that 2014 paper, Ito leverages similar variation in prices over time and along a spatial discontinuity to estimate short-run elasticities.

Let c_{it} denote consumption for customer i in month t and MP_{it} denote the marginal price that customer i faces in month t. For expositional purposes, I assume that all customers respond to marginal price, though I will relax this assumption later in this section. Typically, one could consider the following first differences estimating equation:

$$\Delta\ln(c_{it}) = \beta_1 \Delta\ln(MP_{it}) + \gamma_{ct} + \eta_{it} \tag{1.1}$$

where $\Delta\ln(c_{it}) = \ln(c_{it}) - \ln(c_{i,t-12})$ is the difference between log consumption today and the same month one year prior, $\Delta\ln(MP_{it}) = \ln(MP_{it}) - \ln(MP_{i,t-12})$ is the difference between log marginal price today and the same month one year prior, γ_{ct} denotes CBG-by-time fixed effects, and $\eta_{it} = \epsilon_{it} - \epsilon_{i,t-12}$ is an idiosyncratic error term. Using this first differences estimator removes household-by-month-of-year variation. However, the structure of electricity rates in California raises issues for this estimation.

As described in the background section, electricity providers in California employ increasing-block pricing. Hence, as customers use more electricity, the marginal price of electricity increases. The marginal price of electricity is therefore correlated with consumption, meaning that in the Equation (1), the marginal price is correlated with the unobserved error term η_{it}.

To solve this issue, I follow Ito (2014). Ito instruments for price using the policy-induced price change. The instrument, called a simulated instrument in the tax literature, is

$$\Delta\ln(MP_{it})^I = \ln(MP_t(c_{i,t-6})) - \ln(MP_{t-12}(c_{i,t-6})) \tag{1.2}$$

This instrument isolates the change in price induced by exogenous policy change at a specific consumption level. For it to be valid, $c_{i,t-6}$ must be uncorrelated with the unobserved error η_{it}. Some past studies have used the base year consumption, $c_{i,t-12}$, here. However, as Ito points out, mean reversion presents a challenge in this setting, as transitory shocks to consumption in month $t-12$ will cause mean reversion in consumption that will be correlated $\epsilon_{i,t-12}$ and therefore η_{it}. Blomquist and Selin (2010) and Saez, Slemrod and Giertz (2012) suggest that in an income tax setting, using consumption in a period midway between t and $t-12$ can be used to address this mean reversion problem.

This instrument might still be correlated with η_{it} if specific types of electricity users (e.g. high- and low-usage customers) have different consumption paths over time. This is where I make use of the border discontinuity that results from baseline territories. Ito uses the border discontinuity between utility regions, which was an effective border for understanding how customers respond differently to marginal and average prices. However, it would be less suitable for this project, as long-run price elasticities in residential electricity are likely driven by investment in solar, energy efficiency, and other durable goods. In different utility regions, there are different incentives and marketing strategies for these types of durable goods that go beyond the price that customers face. In order to measure the long-run elasticity of electricity consumption, the baseline territory boundary is better suited to create long-run price variation. Furthermore, baseline territory borders are not limited to one concentrated geographic area as utility borders are, leading to a more representative sample.

Therefore, I restrict my sample to census block groups that have at least 50 service accounts in multiple different climate zones. The resulting identifying assumption is that customers in the same census block groups on either side of the climate zone boundary would consume the same amount of energy absent the price variation that results from the climate zones. With this instrument, I estimate a two-stage least squares regression of consumption on marginal price, instrumenting for marginal price with the simulated instrument described above:

$$\text{First stage: } \Delta\ln(MP_{it}) = \alpha_1\ \Delta\ln(MP_{it})^I + f_t(c_{i,t-6}) + \gamma_{ct} + \eta_{it} \tag{1.3}$$

$$\text{Second stage: } \Delta\ln(c_{it}) = \beta_1\ \widehat{\Delta\ln(MP_{it})} + f_t(c_{i,t-6}) + \gamma_{ct} + \eta_{it} \tag{1.4}$$

where $f_t(c_{i,t-6})$ is a set of dummy variables determined by the decile of consumption in period $t-6$. Formally, for percentile j, $f_{j,t} = 1\{c_{j,t-6} < c_{i,t-6} \leq c_{j+1,t-6}\}$.

In this specification, β_1 represents a short-run elasticity to marginal price – it estimates how any exogenous price change over the previous year leads to a difference in consumption within that period.

As Ito (2014) finds, customers might respond to average prices instead of marginal prices. As such, I include two additional short-run specifications: one in which average prices replace marginal prices as the primary covariates of interest, and one which includes both average and marginal prices as covariates. This final specification is called an encompassing test, and measures whether one pricing model "encompasses" the other. In his work, Ito (2014)

Table 1.3: Short-run price elasticity

	MP (1)	AP (2)	Encompassing (3)
$\Delta \ln(MP_{it})$	-0.18***		-0.052**
	(0.019)		(0.025)
$\Delta \ln(AP_{it})$		-0.36***	-0.28***
		(0.032)	(0.045)
Observations	5692238	5663639	5663639
F	480.9	989.7	99.8

*Note: Across all columns, the dependent variable is $\Delta ln(c_{it})$. Fixed effects include CBG-by-month and 6-month-lagged consumption deciles. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

finds that the average price model encompasses the marginal price model, implying that customers primarily respond to average prices rather than marginal prices.

As shown in Table 1.3, I find results that are consistent with the existing literature (Zhu et al., 2018). Elasticities are approximately -0.18 and -0.36 for marginal and average prices respectively. While I don't find that the average price model encompasses the marginal price model, customers seem to generally be more responsive to average prices than marginal prices.[12] As such, throughout the rest of the paper, my preferred specifications will use average prices as the primary covariate of interest with specifications showing marginal prices in the Appendix. These results primarily serve to anchor my results within the existing literature. Much of the literature on price elasticities in the residential electricity sector have focused on the short-run, and has typically found similar results to those that I present here – customers are more responsive to average price than to marginal price, but short-run responses to prices are relatively inelastic.

To understand how different types of consumers respond differently to prices, I use CARE as a proxy for income, as shown in Table 1.4. I estimate the same specification separately for CARE and non-CARE customers, finding that non-CARE (and therefore higher income) customers tend to be more responsive to prices than CARE customers. While there are some papers that find similar results (Brolinson, 2019; Schulte and Heindl, 2017), this result is in contrast with the majority of the literature, which finds that price elasticities of demand are higher among the poorer households (Alberini, Gans and Velez-Lopez, 2011; Reiss and White, 2005). Appendix Table 4.2 shows similar regressions by federal poverty line decile,

[12]Note that Shaffer (2020) finds a similar result, where customers are heterogeneous in how they respond to prices. I don't take a stand here on whether customers respond to marginal or average prices.

Table 1.4: Short-run price elasticity by CARE

	CARE (1)	nonCARE (2)
$\Delta\ln(AP_{it})$	-0.20***	-0.43***
	(0.035)	(0.037)
Observations	1002596	4660937
F	605.7	1162.3

*Note: Across all columns, the dependent variable is $\Delta ln(c_{it})$. Fixed effects include CBG-by-month and 6-month-lagged consumption deciles. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

finding the same conclusion that low-income consumers are less responsive to prices in the short-run. These results suggest that investment in capital-intensive durable goods may play a significant role price responsiveness.

Medium-run responses to prices

Next, I turn my attention to the medium-run. In the primary medium-run specification in this paper, I refer to the medium run as a time horizon of four years. Currently, there is little existing work in the literature on medium-run elasticities, especially in a quasi-experimental setting. Deryugina, MacKay and Reif (2019) find that customers are more responsive in the medium-run than in the short-run. This is is consistent with a number of studies using aggregated state-level data that similarly find that consumption responses to price build over time. In this section, I estimate how four-year consumption differences can be attributed to price changes that occur within that four-year period.

There are several different channels through which customers might respond to prices. After observing a change in price, consumers might respond in the short run by reducing their consumption of certain appliances – for example, a consumer might turn off their lights more frequently. If this short-run behavior becomes a habit for the consumer, we might continue to see this response carry through to the medium-run. However, customers may also respond by changing their investment of durable goods, such as energy efficient appliances, electric heat, or solar panels. We should expect that durable good adoption will impact consumption in both the short-run and the medium-run. These two channels – habit formation of conservation behaviors and durable good adoption – are the primary channels through which past prices can impact current consumption.

In order to estimate medium-run elasticities in this setting, I extend Ito's approach using time lags. Now, the dependent variable is the difference in consumption between the contemporaneous period and four years prior for a given household. I include the full price path as right-hand side variables, with a series of annual price differences within a four-year window as the primary covariates of interest. I used two-stage least squares, with four endogenous variables and four instruments:

$$\text{First stage: } \Delta\ln(MP_{i,t,l}) = \ln(MP_{i,t-12l}) - \ln(MP_{i,t-12(l+1)}) \text{ for each } l \in 0,1,2,3 \tag{1.5}$$

$$\text{Second stage: } \Delta\ln(c_{i,t,t-48}) = \sum_{l=0}^{3} \beta_l \, \Delta\ln(\widehat{MP}_{i,t,l}) + f_t(c_{i,t-60}) + \gamma_{ct} + \eta_{it} \tag{1.6}$$

where $\Delta\ln(c_{i,t,t-48}) = \ln(c_{i,t}) - \ln(c_{i,t-48})$ and $f_t(c_{i,t-60})$ is a set of dummy variables determined by the percentile of consumption in period $t-60$.

As in the short-run specifications, each price difference is endogenous to consumption due to the nature of increasing block pricing. Again, I use simulated instruments to solve this issue. For each endogenous price covariate, an associated simulated instrument is included.

In the short-run specifications, consumption levels from period $t-6$ were used in the instrument. Here, however, consumption in period $t-6$ is endogenous to the price differences included as covariates. Instead, in the medium-run specifications, consumption levels from period $t-60$ (one year prior to the first included price period) are used. This ensures that the instrument isolates exogenous changes in the price schedule and eliminates all endogenous price variation driven by consumption changes. Note that this specification includes only utility accounts continuously present in the sample over the course of five years (months ranging from t to $t-60$). Any customers who move over the course of this period are dropped from the sample. Hence, the external validity of these medium-run estimates is limited to consumers who are fairly stable and live in a single location for an extended period of time.

It's important to note that the empirical setting in this paper is quite different than in past work, including Deryugina, MacKay and Reif (2019). Deryugina et al. leveraged on a one-time change in prices and followed customers' demand levels over time. Here, price schedule fluctuations frequently occur and impact customers differently depending on their baseline territories and underlying consumption levels. As such, the interpretation of estimates is different in this setting: while elasticity estimates in Deryugina et al. should be interpreted as a consumption response to a single permanent change in price, the estimates in this paper tell us how to attribute changes in consumption to changes in price over the relevant period. When consumption changes over a four year period, how much of that change should be attributed to price changes in each year? Examination of each coefficient in the regression demonstrates how elasticities evolve over time.

Results of these medium-run regressions over a four-year period are shown in Table 1.5. These results demonstrate that responses to prices last over the course of several years, indicating that habits and/or durable good adoption play a vital role. When including past price periods, customers are similarly responsive to short-run fluctuations in price, with a price elasticity of -0.18. This elasticity stays close to -0.2 through two years, before fading towards -0.1 by the fourth year.

Table 1.5: Dynamic medium-run average price elasticities

	kWh
	(1)
$\Delta\ln(AP\_it)$	-0.18***
	(0.044)
$\Delta\ln(AP\_i,t,1)$	-0.19***
	(0.042)
$\Delta\ln(AP\_i,t,2)$	-0.13***
	(0.030)
$\Delta\ln(AP\_i,t,3)$	-0.12***
	(0.034)
Observations	2606624
F	194.9

*Note: Fixed effects include CBG-by-month and consumption deciles from the period twelve months prior to the initial price period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

In the Appendix Table 4.2, I also estimate medium-run elasticities over an eight-year period. Note that this sample is even more highly selected to include only customers who do not move over a nine year period within my twelve year sample. Again, the external validity of these estimates is restricted only to consumers who live in a single location for an extended period of time – in this case, nine years.

These results are consistent with a combination of a short-run transient behavioral responses and significant durable good investment. After price fluctuations, consumers respond by changing their consumption. However, customers may also respond to price changes by investing in durable goods, which last for the duration of the sample. As a result, they still demonstrate responsiveness to price changes that occurred in more distant periods – in this case, three to four years prior to the contemporaneous period.

In addition, I estimate heterogeneity in medium-run elasticities across income, again using CARE enrollment as a proxy for income, as shown in Table 1.6. Consistent with the short-run results, non-CARE (higher-income) consumers are more responsive to changes in their electricity prices in all periods. Once again, this suggests that investment in durable goods may play a substantial role in how consumers respond to energy prices.

Furthermore, these results suggest that consumers' responses to price changes may accumulate over time as consumers continue to respond to prices from four years prior. However, the type of dynamic two-way fixed effects panel regressions shown to this point only allow for evaluation up to the length of the observed sample, and may miss important mechanism. As such, the next chapter of this dissertation is devoted to an empirical approach that leverages

Table 1.6: Dynamic medium-run average price elasticities by CARE

	CARE (1)	nonCARE (2)
$\Delta\ln(AP\_it)$	-0.11**	-0.21***
	(0.053)	(0.050)
$\Delta\ln(AP\_i,t,1)$	-0.12**	-0.23***
	(0.052)	(0.049)
$\Delta\ln(AP\_i,t,2)$	0.0038	-0.18***
	(0.043)	(0.034)
$\Delta\ln(AP\_i,t,3)$	-0.056	-0.13***
	(0.052)	(0.037)
Observations	413612	2192841
F	166.4	200.7

*Note: Fixed effects include CBG-by-month and consumption deciles from the period halfway between the present period and the lagged consumption period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

cross-sectional variation to estimate elasticities over a much longer period of time.

1.5 Conclusions

In this paper, I leverage a novel source of cross-sectional price variation to estimate how residential electricity consumers respond to electricity prices in the short and medium run. I find that consumers are somewhat responsive to electricity prices in the short run, and that the consumption response persists up to at least four years. In the short-run, I follow the existing literature to estimate price elasticities of demand. In the medium run, I expand on this methodology with a novel approach to leverage cross-sectional and dynamic differences in prices to estimate how consumption responds to past changes in prices, up to four years prior.

Furthermore, I estimate heterogeneity in consumption responsiveness based on household income. Low income consumers are about half as responsive as higher-income households in both the short- and medium-run. This result might be explained by the fact that higher income households often have more potential margins of response to changes in electricity prices, due partly to higher electricity usage. One of these potential margins is investment in durable goods, which are frequently capital-intensive and may only be options for higher income households. However, more research is necessary to explore the specific mechanisms driving the differences across income levels. In the second chapter of this dissertation, I begin to explore some of those mechanisms.

The policy implications of these results are clear. Short- and medium-run elasticities are important for electricity demand forecasts, which are used by numerous stakeholders including distribution utilities, generators, grid planners, and many more. Furthermore, policymakers across the country are considering price based policies to mitigate climate change. Understanding how consumers respond to prices and the mechanisms driving these responses is vital when considering the impacts of those policies, especially across different income levels. This paper provides estimates that can inform those policy conversation and promote sound policy that can improve welfare across all consumers.

Chapter 2

Long-run price elasticities and mechanisms: Empirical evidence from residential electricity consumers [1]

2.1 Introduction

Long-run price elasticities of demand are among the most important parameters in the field of economics. To calculate welfare, forecast future demand, and design policy, long-run elasticities are vital components. Still, across many economic fields, there are few experimental or quasi-experimental empirical estimates of long-run price elasticities.

The dearth of experimental and quasi-experimental long-run elasticity estimates is a product of challenging empirical conditions – to empirically estimate a long-run price elasticity, one either needs a plausibly exogenous source of long-lasting variation in prices, or to make strong structural assumptions. Sources of persistent exogenous price variation are rare, however, and many estimates therefore rely heavily on these assumptions (Kamerschen and Porter, 2004; Dergiades and Tsoulfidis, 2008; Alberini and Filippini, 2011).

In this paper, I leverage a novel source of persistent spatial price variation to estimate long-run price elasticities for residential electricity customers in California. This price variation is driven by a subtle feature of California's pricing regime. In the increasing block pricing rate structure used throughout California, marginal prices increase when electricity usage exceeds a certain threshold. Because of differences in heating and cooling needs for

[1] I thank Meredith Fowlie, James Sallee, Catie Hausman, and Severin Borenstein for invaluable advice, feedback, and mentorship throughout this project. I am also grateful to Marshall Blundell, Ellen Bruno, Fiona Burlig, Lucas Davis, Jenya Kahn-Lang, Louis Preonas, Matthew Tarduno, David Zilberman, and participants at various seminars for helpful comments. Finally, thank you to Robert Lucadello for invaluable assistance accessing and understanding the utility data used for this project. I do not have any financial relationships that relate to this research.

households across the different climates of California, utilities set these thresholds to different levels depending on where a consumer lives. Within Pacific Gas & Electric's (PGE's) service territory, there are ten different baseline territories[2], with the boundaries for these territories often determined according to discontinuities in a household's elevation. These boundaries have led to long-lasting persistent price variation since they were established in 1982, with one side of the border consistently facing higher prices than the other. I leverage these price discontinuities to estimate elasticities across different time intervals.

Estimating a long-run elasticity is empirically challenging, as the panel methods commonly used in the literature can miss important margins of response. Typical panel methods compare consumption before and after a price change, which relies on having counterfactual data on a consumer both before and after the change in price. Notably, they miss cross-sectional consumption differences created when homes are built or when new tenants move in (often a time in which home renovations occur) under different price regimes across space.

To estimate long-run elasticities, I rely primarily on cross-sectional price variation driven by the levels of the baselines across baseline territory boundaries. By leveraging this cross-sectional variation, I capture a more comprehensive measure of demand response. I estimate a long-run elasticity of -2.4, indicating that consumers are much more responsive to permanent price changes in the long run than to short-run price fluctuations.

However, substantial heterogeneity exists in price responsiveness according to a consumer's income level. There are overlapping explanations that might explain this conflict – electricity bills may be more salient to low-income households as they have less discretionary income than higher income households. However, higher income households likely have more appliances, leading to more margins for response to price changes. Furthermore, durable goods that reduce consumption often have high capital costs, leading to potentially greater adoption among higher income households (Borenstein, 2017). I estimate elasticities for households with different income levels, finding that low-income consumers are less responsive to changes in prices in the short- and medium-run, but that in the long-run, this gap dissipates and low-income consumers are similarly responsive as higher-income consumers.

To better understand these elasticities, I explore possible mechanisms that would drive such large price responses in the long run. I find that consumers are highly responsive to long-run prices in their adoption of rooftop solar and somewhat responsive in their adoption of energy efficiency[3]. I estimate how adoption of solar and energy efficiency impact household electricity consumption, finding that on average, adoption of solar decreases monthly household consumption by 617 kWh per month. A back-of-the-envelope calculation reveals that this mechanism alone can account for 26% of the observed difference in consumption.

The long-run elasticity and mechanisms estimated here are specific to the geography and climate in this northern and central California sample, and there are reasons to expect

[2]These baseline territories were finalized in 1982 and have largely stayed the same since 1982 PG&E (2020).

[3]In this paper, I observe the adoption of energy efficient appliances when a consumer enrolls in a utility energy efficiency program. This is only a portion of energy efficient appliance sales, and the energy efficiency adoption estimates presented here can be thought of as lower bounds.

a relatively high long-run elasticity in this setting. Electricity prices are particularly high in California, leading consumers to be more aware of their electricity bills. Rooftop solar adoption and potential are also higher in California than many other settings in the United States, creating better conditions for that margin of response. However, one should expect elasticities to be much larger in the long run than the short run across all geographies, as consumers have more time to make adjustments to their behaviors, adopt durable goods, and choose housing characteristics that impact electricity consumption in the long run.

This paper contributes to three distinct literatures. First, this paper contributes to the literature estimating long-run responses to prices. Within the field of energy economics, most papers estimating long-run elasticities use aggregated data and structured dynamic panel models. These papers, including Alberini and Filippini (2011), Kamerschen and Porter (2004), and Dergiades and Tsoulfidis (2008) rely on strong assumptions about the form of serial correlation and typically estimate long-run elasticities in the range of -0.3 to -1.1. There are few papers that use quasi-experimental methods to estimate long-run elasticities, most notably by Deryugina, MacKay and Reif (2019) and Feehan (2018). Deryugina et al. estimates price elasticities spanning a time horizon of up to three years[4], finding a price elasticity of -0.09 in the first six months and -0.28 after 30 months. Feehan (2018) is perhaps more closely related to this work, where the author leverages a natural experiment in Canada to estimate 20-year elasticities for residential electricity customers, finding a long-run elasticity of -1.2. This paper builds on the conclusions of these studies by directly estimating mechanisms of response and by exploring heterogeneity in both the short and long run.

Second, this paper contributes to the literature on durable goods investment, especially in response to input prices. Here, I estimate how adoption of solar and energy efficiency measures varies in response to electricity prices. Work by Chesser et al. (2018) and Crago and Chernyakhovskiy (2017) explore the impact of electricity prices on investment in rooftop solar, using aggregated data to show that electricity prices are an important drivers of residential solar adoption, with higher electricity prices leading to greater solar adoption. This paper builds on that work by using administrative customer-level data to estimate how individual customers respond to within-utility price differences. Furthermore, this paper is the first to directly attribute long-run price responsiveness to durable goods mechanisms.

Finally, this paper contributes to the literature exploring how household income impacts price responsiveness among residential electricity customers. Evidence in this literature is somewhat conflicting – Alberini, Gans and Velez-Lopez (2011) and Reiss and White (2005) find that price elasticities of demand are highest among the poorest households and monotonically decrease as income grows, while Brolinson (2019) and Schulte and Heindl (2017) find that wealthier households are more responsive to prices. This paper is first to separately quasi-experimentally estimate long-run price elasticities by income. I show that while higher income households are much more responsive to prices in the short- and medium-run, as shown in the first chapter of this dissertation, this gap nearly disappears in the long run,

[4]The authors forecast a long-run elasticity spanning up to ten years, but are unable to test this estimate with quasi-experimental data.

suggesting that adoption of costly durable goods may not be the primary driver of income heterogeneity in short-run price responsiveness.

The paper proceeds as follows: Section 2 discusses background on the setting, along with the data and empirical strategy; Section 3 presents estimates of long-run price elasticities; Section 4 explores the mechanisms driving these responses to price; and Section 5 concludes.

2.2 Background, Data, and Research Design

This paper closely follows the first chapter of the dissertation in the research setting and data. As in the first chapter, the setting for this paper is Pacific Gas & Electric (PG&E). PG&E uses increasing block pricing and baseline territories to determine a customers price schedule. Once again, changes to the price schedule over time and across space provide useful identifying variation to estimate the impacts of prices.

This paper also uses the same measures of income heterogeneity as in the first chapter, with the primary measure being enrollment in California Alternative Rates for Energy (CARE) and a secondary measure as the Census-Block-Group (CBG) portion of households below the Federal Poverty Line (FPL).

Finally, the data and sampling restrictions in this paper are identical to those of the first chapter of the dissertation. I use account-level billing data for all PG&E electricity customers from 2008 to 2020, restricting my sample to CBGs baseline territory boundaries lead to price variation. In this paper, I leverage these cross-sectional baseline territory boundaries to estimate long-run price elasticities. I explore heterogeneity across different types of customers and the mechanisms driving those price responses in order to acquire a more complete picture of the factors driving the population-wide effects.

2.3 Estimation

Long-run responses to prices

In the first chapter of this dissertation, I estimate short- and medium-run elasticities by leveraging price variation over time and estimating how consumers react to dynamic changes in the price schedule. However, this type of analysis only captures customers who are continuously present in the sample over a long period of time. Furthermore, as recent work by Davis (2020) suggests, important durable good decisions such as whether a home is heated by gas or electricity may often be decided when a home is built, with substantial switching costs that lead to low incidence of switching behavior. Alternatively, investment decisions may be made when a utility account switches due to a new owner or tenant moving in. The typical dynamic methods used in the literature to estimate elasticities often fails to capture this variation – in fact, any approach that relies on price changes over time will fail to capture this critical margin of response.

This creates a challenge for the empirical researcher – how can one estimate differences in adoption of durable goods (and differences in consumption) without a "pre-period?" Instead of leveraging price changes over time, I leverage cross-sectional differences in prices due to baseline territory divisions to observe long-run differences in durable good adoption and consumption.

Due to the discontinuity of baseline territories, within each CBG in this sample, there is a "high price" region and a "low price" region. Importantly, the ordering of baseline territories tends to be preserved over time and has been relatively consistent since baseline territories were finalized in 1982. To estimate a long-run elasticity, I leverage this cross-sectional price variation. Once again, I use an instrumental variables approach, instrumenting for price with the length of the baseline:

$$\text{First stage: } ln(p_{it}) = \alpha_0 + \alpha_1 \text{baseline}_{it} + \gamma_{ct} + \eta_e + \epsilon \tag{2.1}$$

$$\text{Second stage: } ln(c_{it}) = \beta_0 + \beta_1 \widehat{ln(p_{it})} + \gamma_{ct} + \eta_e + \epsilon^5 \tag{2.2}$$

where c identifies CBG, e is a dummy variable indicating if a customer is all-electric, s denotes whether the present period is summer, and baseline$_{it}$ denotes the length of the monthly baseline for customer i in month t. Recall that a customer's baseline depends on their baseline territory, whether they use electric heat, and whether it is summer or winter. I include fixed effects for electric heat and the month-of-sample fixed effect controls for summer as well. Therefore, the only identifying variation left in baseline lengths is driven by the baseline territory definitions and changes to the baselines over time. Once again, CBG-by-month-of-sample fixed effects are included so that customers within the same CBG but facing different baselines are directly compared. The identifying assumption under this regression remains similar as in the short- and medium-run regressions: customers within the same CBG and with the same type of heating systems would consume similar amounts of electricity absent the differences in prices driven by baseline territories.

Because a binary variable for electric heat is included in the fixed effects, this specification compares customers within the same heating type, because the length of a baseline is partly determined by heating type. This is a necessary fixed effect to prevent an endogenous heat type choice to bias the estimates. However, the inclusion of this fixed effect eliminates heating choice as a potential mechanism to impact consumption. Therefore, the long-run elasticity estimated here is a lower bound, as a theoretical specification that allowed a heating type margin to impact consumption would only increase the estimated elasticity.

Additionally, note that the variation in this regression is not purely cross-sectional, as there is some variation in baselines over time. However, the primary identifying variation is cross sectional and in combination with the sampling restrictions, the regression results are mainly driven by differences across baseline territories. In Appendix A.5, I include similar specifications replacing the instrument with an indicator for whether a consumer is on the "high price" or "low price" side of the border discontinuity. With this specification, I find similar results, demonstrating that spatial variation in baselines in the primary driver of these long-run results. However, the preferred specification shown in Equations 7 and 8 uses

Table 2.1: First stage long-run regression

	All (1)
Monthly baseline (100 kWh)	-0.0041*** (0.00018)
Observations	9399944

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Table 2.2: Reduced form regression

	All (1)
Monthly baseline (100 kWh)	42.8*** (2.05)
Observations	9660076

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

the length of the baseline as an instrument, as the magnitude of baselines reflect the degree of price differences that one would expect to see across baseline territories.

First, results of the first stage and reduced form are shown in Tables 2.1 and 2.2 respectively. Because customers are more responsive to average prices than marginal prices, I show results using average prices here. In the Appendix, I show results of the same specifications using marginal prices.

The first stage results show that for in increase in monthly baselines of 100 kWh, contemporaneous average prices are lower by an average of 0.41 cents per kWh. At median levels of electricity usage, this difference in price would imply a bill difference of about $3 per month. Meanwhile, the reduced form estimates show that for the same increase in monthly baseline,

Table 2.3: Long-run IV estimate of elasticity (average price)

	All
	(1)
Logged average price	-2.35***
	(0.33)
Observations	9331612
F	317.2

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

consumers respond by increasing their consumption by over 40 kWh per month, or about 7.4% of the mean monthly usage.

As shown in Table 2.3, the IV regression results in an elasticity of -2.6, implying that customers are much more responsive to price changes in the long run than in the short or medium run. While this estimate is substantially larger than the existing literature, there are several reasons that one should expect a larger estimate in this setting and under this methodology: first, the existing literature tends to use panel methods that compares consumption for a customer before and after a price change. This type of estimation misses important margins of response – specifically investment decisions at the time that a home is built or when a utility account changes due to a new owner or tenant moving in. By leveraging a persistent source of cross-sectional price variation, the specification here captures the investment margin, including in new and recently transacted homes, both of which are often missed by studies that rely primarily on price variation over time, as opposed to across space.

Second, there are very few existing quasi-experimental estimates of long-run elasticities in the literature. Most estimates rely on strong structural assumptions made by researchers. One of the only quasi-experimental long-run elasticity estimate to date, Deryugina, MacKay and Reif (2019), looks only at a time horizon up to three years, and estimates elasticities using the panel methods described above, which are likely to miss important margins of response. The estimates in that paper should be compared to the medium-run results shown previously in this paper, not these long-run estimates, because of the parallels in both time horizon and in the identifying variation. The other quasi-experimental long-run elasticity estimate to date, Feehan (2018), finds a long-run elasticity of -1.2 in Newfoundland and

Labrador, Canada. Critically, the setting for this paper is in a different climate. Consumers in Newfoundland and Labrador face lower temperatures than California year-round, leading to less flexibility in decisions around heating and cooling. Furthermore, solar irradiance is substantially lower, diminishing the value of one of the most important margins of long-run response observed in California. With additional margins of response and more flexible heating and cooling loads, one would expect consumers to be much more responsive to prices.

One might be concerned about endogeneity in this specification – when a customer adopts a durable good such as solar, their net electricity usage decreases dramatically, often putting them into a different pricing tier and decreasing both their marginal and average prices. Because I use the contemporaneous average price as the variable of interest, there is a concern that endogenous adoption of durable goods may decrease the price difference on either side of the border, thereby biasing upwards the estimated of elasticity. In Appendix A.3, I test alternative definitions of price, where prices are determined by consumption levels in a baseline year, finding similar long-run elasticity estimates.

One may also be concerned that this result is driven by the presence of outliers. To rule out this possibility, I separately estimate coefficients for every census block group in the sample. As shown in Figure 2.1, I find that 57% of CBGs exhibit elasticities between -1 and -5, and that the few outliers that do exist are not the primary factor driving the results. In Figure 2.2, I explore the distribution of elasticity estimates across space, finding no demographic trends that are predictive of elasticity magnitude.

Once again, I estimate in Table 2.4 heterogeneity in elasticities according to whether customers have ever been enrolled in CARE. While in the short- and medium-run specification, non-CARE (higher-income) consumers were much more responsive, in the long run, CARE and non-CARE consumers are similarly responsive to one another. This is a surprising result – higher income consumers have more access to capital with which they can invest in durable goods that impact their consumption. Similarities in long-run price responsiveness indicates that there may be less capital intensive margins of response that low-income households are able to leverage.

The comparison of short- and long-run elasticities suggests the likely channels through which consumers might respond. Because long-run responses are much larger than short-run responses, it is unlikely that these results are primarily driven by intensive margin changes, such as appliance use behaviors. The pattern of response is more consistent with adjustment along an extensive margin: the adoption of durable goods such as solar, energy efficiency, and electric heating. In the next subsection, I empirically estimate how customers respond in their adoption of durable goods to better understand the specific mechanisms driving the observed long-run response.

2.4 Mechanisms

To this point, I have shown that residential electricity customers are much more responsive in the long run than the short run to electricity prices. This begs the question of which

Figure 2.1: Histogram of long-run elasticity distribution across Census Block Groups

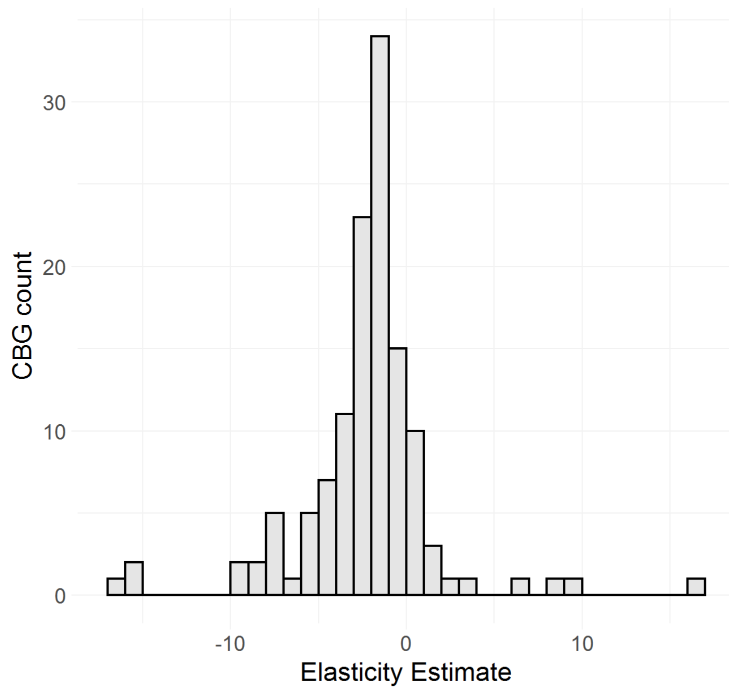

Note: This figure shows the results of estimating long-run elasticities within each Census Block Group. The histogram shows the distribution of CBG-specific elasticity estimates.

mechanisms are driving this response. Are customers making *intensive* margin adjustments to their electricity consumption, by changing their behaviors surrounding heating or appliances? Or are they making *extensive* margin changes, by adopting durable goods such as solar, energy efficiency, or electric heat?

To better understand the primary drivers of these results, it's crucial to understand the specific mechanisms through which customers respond to prices. The observed differences between short-, medium-, and long-run elasticities suggest that investment in durable goods plays a significant role. Here, I directly observe the adoption of two different types of durable goods: residential solar PV and the energy efficient appliances that are supported by utility energy efficiency programs.

To understand how customers respond to prices with durable good adoption in the long run, I estimate a simple cross-sectional specification:

$$\text{Adoption}_i = \beta_0 + \beta_1 hi_i + \gamma_c + \epsilon_i$$

where Adoption_i is a binary variable indicating whether a customer ever adopts the durable good over the course of the sample; hi_i is a binary variable indicating whether a customer

Figure 2.2: Map of long-run elasticities by Census Block Group

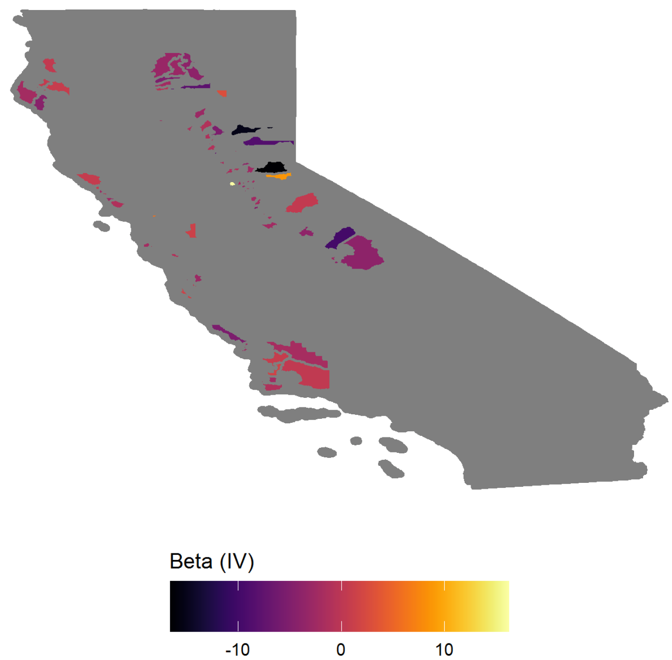

Note: This map shows the results of estimating long-run elasticities within each Census Block Group. The color of each CBG on the map indicates the long-run elasticity estimate for that CBG. One CBG, with an elasticity estimate over 20, has been omitted for scaling purposes.

lives in the "high price" baseline territory within a CBG, and c denotes CBG.

This specification relies solely on cross-sectional variation in prices driven by the baseline territory discontinuity. Once again, I compare customers within the same CBG, leading to the identification assumption that customers within the same CBG would have adopted the same durable goods, absent the difference in price driven by baseline territories.

In Table 2.5, I show the results for the two durable goods of interest: residential solar adoption and utility energy efficiency programs. I find that consumers facing persistently higher prices over time tend to adopt solar more frequently, by 2.0 percentage points (average solar adoption throughout the sample is about 5%). Consumers tend to adopt utility energy efficiency programs more frequently as well, by 0.3 percentage points (mean EE adoption is 2%).

While low-income consumers were similarly responsive to higher-income consumers in their long-run consumption responses, I find that high-income consumers tend to be more responsive in their adoption of observed durable goods. Table 2.6 shows estimates of adoption responses by CARE. While enrollment in energy efficiency programs is fairly balanced across

Table 2.4: Long-run IV estimate of elasticity by CARE

	CARE (1)	nonCARE (2)
Logged average price	-1.71*** (0.24)	-2.24*** (0.34)
Observations	2437275	6894296

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Table 2.5: Mechanisms

	Solar (1)	EE (2)
hi	0.020** (0.0099)	0.0023** (0.00091)
Observations	288386	288386

*Note: Fixed effects include CBG. Standard errors are clustered by CBG-baseline territory. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

income, solar adoption is substantially higher among non-CARE customers.

To summarize, I find that low-income consumers respond similarly to higher income consumers in the long run, yet higher-income consumers seem to be much more responsive across observable mechanisms. More work is needed to understand the departure between the heterogeneity in these mechanisms regressions and in the estimation of long-run elasticities. This result may be driven by unobserved mechanisms, and future research should dig deeper into other mechanisms, such as home characteristics, to better understand the margins of response for low-income consumers.

To fully understand the impact of these mechanisms, it's important to understand the extent to which my estimates of long-run elasticities are driven by mechanisms. Here, I at-

Table 2.6: Estimates of mechanisms by CARE

(a) CARE

	Solar (1)	EE (2)
hi	0.012*	0.0029*
	(0.0073)	(0.0016)
Observations	66095	66095

(b) non-CARE

	Solar (1)	EE (2)
hi	0.021*	0.0022**
	(0.011)	(0.00093)
Observations	222291	222291

*Note: Fixed effects include CBG. Standard errors are clustered by CBG-baseline territory. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

tribute the observed differences in consumption to each durable good of interest. Once again, I use a two-way fixed effects framework, regressing consumption in month i for household t on a binary variables for each durable good:

$$c_{it} = \beta_0 + \beta_1 Solar_{it} + \beta_2 EE_{it} + \beta_3 Electric_{it} + \gamma_i + \lambda_t + \epsilon_{it}$$

where $Solar_{it}$ is a binary variable denoting whether household i has adopted solar in period t and EE_{it} and $Electric_{it}$ follow the same structure for energy efficiency and electric heat respectively. Fixed effects for the household and month of sample are also included. Once again, net monthly consumption (gross consumption minus monthly solar generation) is the dependent variable.

As shown in Table 2.7, I find that adoption of solar leads to a massive reduction in net consumption, by over 600 kWh per month, over a sample-wide baseline mean of about 550 kWh per month[6]. A back-of-the-envelope calculation reveals that 26% of the observed difference in long-run consumption can be attributed to solar adoption alone[7].

The conclusions surrounding energy efficiency are less clear, as I find that enrollment in PG&E's energy efficiency programs results in an increase in consumption of about 31 kWh

[6]This comparison can explained by the fact that solar users tend to have higher-than average consumption before adopting solar. Among customers who eventually adopt solar, mean pre-solar consumption is 1,255

Table 2.7: Consumption impact of durable goods

	(1)
Solar	-616.8***
	(6.66)
Energy Efficiency	21.5***
	(5.68)
Observations	9660065

*Note: Fixed effects include household and month-of-samples. Standard errors are clustered by household. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

per month. There are a number of factors that could explain this result: first, the observed measure of energy efficiency only includes PG&E-specific programs, capturing only a portion of the energy efficiency behaviors. 44% of observed energy efficiency activities are home energy audits, which many past papers have shown often don't deliver on expected savings. Second, numerous papers (Aydin, Kok and Brounen, 2017; Jin, 2007) have documented the "rebound effect" where households consume a greater quantity of energy services after enrolling in an energy efficiency program, though it is exceedingly rare for this increase in energy services to lead to "backfire", where consumption actually increases (Gillingham, Rapson and Wagner, 2020). Finally, enrollment in energy efficiency programs may coincide with other events, such as moving to a new rate schedule, new tenants moving in, or adoption of other durable goods. Regardless, while this observation about energy efficiency is not the focus of this paper, future work can further explore the relationship between energy consumption and utility energy efficiency programs.

2.5 Conclusions

In this paper, I leverage a novel source of cross-sectional price variation estimate how residential electricity consumers respond to electricity prices in the and long run. While the first chapter of this dissertation shows that consumers are somewhat responsive in the short

kWh per month.

[7] In this calculation, I multiply the percentage point difference in solar adoption across the baseline territory boundary (2.0 pp) by the consumption impact of solar adoption (-616.8 kWh), then divide by the reduced-form long-run consumption difference across the baseline territory boundary (-47.3).

run, here I find that consumers are much more price-responsive in the long run. This magnitude of this estimated long-run elasticity is considerably larger than the existing literature, potentially due to methodological differences that allow me to capture additional margins of response. Typical quasi-experimental methods rely on tracking the same consumers before and after a price change, missing investment choices made at the time a home is built or when new tenants move in. In this paper, estimation of long-run elasticities relies primarily on cross-sectional price variation, allowing for comparison of similar households facing different price regimes.

The difference in magnitudes between short- and long-run elasticities suggests that the adoption of durable goods plays an important role in household electricity consumption. I directly observe adoption behaviors of two durable goods that might be used in response to price changes – rooftop solar and energy efficiency programs. Consumers are highly responsive to prices in their adoption of solar and somewhat responsive in their enrollment in energy efficiency programs. Solar adoption alone can explain about 26% of the observed difference in long-run consumption. There are numerous additional margins of response – both durable goods and home characteristics – that may impact household electricity consumption, and further research is needed to understand the impact of each of these potential margins.

In addition, I explore the impact that income has on price responsiveness. While in the first chapter of this dissertation, I find that low-income consumers are much less responsive to price changes in the short- and medium-run, low-income consumers are similarly responsive to higher-income consumers in the long run. These findings highlight that higher income consumers may have more margins to adjust their usage in the short- and medium-run (e.g. more appliances that they are able to turn off in response to price changes), but that across the income spectrum, consumers respond similarly to prices over a longer period of time.

The results presented here have enormous policy implications. Not only are long-run elasticities vital for forecasting electricity demand for a number of applications and stakeholders, but these results have additional importance for climate change and price-based policies. In contrast to past research, I find that electricity consumption is highly responsive to prices in the long run, demonstrating that electricity prices can provide strong incentives for consumers to undertake emissions-saving behaviors. This highlights the role that price-based policies, such as carbon taxes, can play in decarbonization efforts. When consumers respond to prices by adopting consumption-reducing durable goods and thereby reducing their emissions, price-based policies are an appealing option to internalize emissions externalities. It also, however, emphasizes the important of getting prices right. Electricity prices above the social marginal cost may drive too much solar adoption and too little adoption in technologies like electric heat or electric vehicles, potentially leading to losses in social welfare.

Chapter 3

Spillovers from Ancillary Services to Wholesale Energy Markets [1]

In electricity markets, generators are rewarded both for providing energy and for enabling grid reliability. The two functions are compensated separately – energy provision is compensated in the energy market, while grid reliability is compensated in the ancillary services market.[2] To date, the economics literature has largely focused on energy provision, even as other academic literatures, policymakers, and grid regulators have more carefully considered ancillary services. Ancillary service markets are interesting and important in their own right: they procure services that prevent brownouts and blackouts and ensure power quality. Moreover, changes in ancillary services markets can impact the behavior of generators in the much-larger energy market. While these market interactions have been extensively studied by engineers using optimization models and simulations, quasi-experimental evidence is largely lacking. In this paper, we show that exogenous policy changes implemented in the ancillary services portion of a large East Coast electricity market have changed the behavior of coal and natural gas generators in the energy market.

We begin with a stylized model to demonstrate how energy provision and ancillary services provision interact. Key features of the model include (1) power plants are multi-product suppliers; (2) power plants tend to operate within a somewhat narrow operational range defined by non-zero minimum and maximum constraints; (3) policy and market changes can cause power plants to operate at minimum load rather than being off. Importantly, the minimum load of many generators (a physical constraint) is non-negligible – while the related literature frequently ignores this constraint, it may be as high as 50 percent or more

[1] This chapter is coauthored with Catie Hausman, Johanna L Mathieu, and Jing Peng. We are grateful to Jim Archsmith, Severin Borenstein, Karl Dunkle Werner, Ken Gillingham, Akshaya Jha, Jenya Kahn-Lang, Justin Kirkpatrick, Ömer Karaduman, Gordon Leslie, Matt Tarduno, Liz Wachs, and various conference and seminar attendees for helpful comments. This research was funded by the Alfred P. Sloan Foundation. The authors do not have any other financial relationships that relate to this research.

[2] Throughout, we refer to the "energy market" and the "ancillary services market" following the language used in the electricity industry. We note, however, that one could instead think of this as a single electricity market for two related but distinct products: energy provision and grid reliability.

of capacity. To summarize, changes along the extensive margin – which units are turned on – can lead ancillary services markets to have outsized impacts on energy markets.

We next turn to empirical analysis of an ancillary services market: the frequency regulation market in PJM. PJM is the largest wholesale electricity market in the US, serving major population centers on the East Coast and dispatching nearly one fifth of all generation capacity in the lower 48 states. Frequency regulation refers to the short-timescale balancing of supply and demand by grid operators; we describe it in depth below. We leverage policy-induced quasi-experimental variation in the amount of frequency regulation required by grid operators. As a result, we identify how changes in the provision of frequency regulation by power plants impact the electricity market as a whole. We show that increases in the required frequency regulation capacity induce changes in the composition of power plants providing energy generation. For a 100 MW increase in the frequency regulation requirement, we estimate an additional 360 MWh from combined cycle plants over the course of one hour in the energy market and a corresponding decrease from boiler units.[3] These results are qualitatively similarly across a broad suite of robustness checks, with estimated increases of combined cycle units of around 300 to 460 MWh.

Because combined cycle plants primarily use natural gas whereas boiler units use coal, we also see a change in the fuel used to provide energy. Specifically, we find that increases in the regulation requirement lead to an increase in natural gas usage and a decrease in coal and oil. As a result, for every 100 MW of increased frequency regulation, CO_2 emissions fall by 240 metric tons per hour in our sample, with robustness checks showing a range of 200 to 370 tons. For the time period we study, the PJM market implemented several changes to *decrease* the regulation requirement. While doing so can reduce the system-wide private cost of electricity provision, in this case it may have inadvertently led in the short term to higher CO_2 emissions. More generally, changes to ancillary services compensation are being considered as the electrical grid evolves in response to climate policy – as such, unintended CO_2 emissions changes in this context are particularly relevant.

At first glance, the magnitude of the generation mix change is surprising. However, the stylized model points towards potential mechanisms: changes along extensive margins, combined with minimum constraints. As such we next provide empirical evidence supporting this mechanism. We show that increasing the regulation requirement causes boilers to be dispatched at lower *levels* of generation in our sample (i.e., a change along the intensive margin). However, combined cycle plants are dispatched more *frequently* (i.e., a change along the extensive margin). These effects are consistent both with the overall generation impacts and with the stylized model.

These results make several contributions to the energy and environmental economics and the industrial organization literatures. First, we contribute to a still-small empirical liter-

[3]We use megawatts (MW) throughout to measure capacity, and megawatt-hours (MWhs) to measure generation. A 1 MW unit operating at its full capacity for one hour generates 1 MWh of electricity. The regulation requirement is a market-wide capacity or power requirement and is thus measured in MWs. Electricity generation is sold as a function of energy provided over the course of some time frame and is thus measured in MWhs.

ature in energy economics that analyzes electricity markets other than the energy market. One strand of the economics and engineering literature explores ancillary services markets with optimization models or simulation approaches (Hirst and Kirby, 1998; Just and Weber, 2008; Yu and Foggo, 2017). However, empirical analysis of ancillary services markets, particularly in how they interact with energy markets, is limited,[4] despite there being a large empirical literature on wholesale electricity markets.[5] We show that ignoring the ways that energy markets interact with these other markets can lead to incorrect conclusions about the impacts of policy changes.[6]

Second, we contribute to a strand of the electricity literature that emphasizes the importance of understanding how technical constraints impact power plant behavior, especially when considering generators as multi-product firms. Perhaps most closely related is Mansur (2008), which shows that ignoring intertemporal constraints gives an inflated estimate of the welfare impact that restructuring electricity markets had in the late 1990s. Most previous empirical papers had focused on merit-order dispatch without considering intertemporal constraints such as minimum load, startup costs, and ramping. Also related are Wolak (2007); Reguant (2014), and Jha and Leslie (2021), which investigate the role of ramping costs and startup costs in generator behavior and market outcomes. Finally, Gowrisankaran, Reynolds and Samano (2016) shows how the intermittency of renewables interacts with the need for reserve markets, ultimately impacting the system-wide costs of generation. We show that minimum load constraints can have a large impact on plant behavior, a point related to the role of technical constraints investigated in these other papers. We additionally show that the existence of multiple related markets can have significant impacts on plant behavior. We argue that power plants should be considered multi-product firms, and that limiting attention to just one of the markets might lead to incorrect or incomplete conclusions about plant behavior.

The behavior of multi-product firms has been the focus of a growing body of work in industrial organization and in international trade. Researchers have shown that examining how firms optimize *across* different products plays a key role in understanding productivity differences across firms; the behavior of firms with market power; and the impacts of trade policy, exchange rate movements, demand shocks, and more (Johnson and Myatt, 2006; Eckel and Neary, 2010; De Loecker, 2011; Chatterjee, Dix-Carneiro and Vichyanond, 2013).

[4]The primary exceptions are: Doraszelski, Lewis and Pakes (2018), which estimates models of firm learning and convergence to equilibrium in a newly deregulated frequency response market in the UK; Knittel and Metaxoglou (2008), which examines ancillary services in the context of the California electricity crisis; and Schwenen (2015), which examines New York's capacity market. Jha and Wolak (2020) examines the impact of "explicit virtual bidding" on the cost of electricity provision, focusing on fuel costs but also incorporating the costs of ancillary services provision.

[5]See Borenstein, Bushnell and Wolak (2002); Fabrizio, Rose and Wolfram (2007); Borensten and Bushnell (2015); Cicala (2015); Davis and Hausman (2016); Holland et al. (2016*a*); Cullen and Mansur (2017); Leslie (2018); Hortacsu et al. (2019), among many others.

[6]Our results are also somewhat related to the literature on spillover effects within the energy market, e.g. from solar to fossil power plants (Bushnell and Novan, 2021) and from hydro to fossil power plants (Archsmith, 2020).

Nonetheless, studies of the electricity market have tended to treat firms as providers of a single good – electricity generation – rather than multiproduct firms, with the exception of an older literature on the optimal regulation of natural monopolies that provide multiple goods (for instance, Mayo, 1984). Our ability to model the production of electricity with an engineering-based model makes clear how the degree of complementarity versus substitutability in production matters for outcomes in multi-product firms. Given the large number of industries with multi-product features (refining, airlines, freight transportation, manufacturing, to name a few), these mechanisms are of widespread relevance.

Third, this paper contributes to policy discussions about several ongoing developments in electricity markets: changes in the way frequency regulation is procured and compensated, the introduction of utility-scale batteries, and the increasing deployment of renewable electricity (MIT Energy Initiative, 2011; Department of Energy, 2013, 2016; Hledik et al., 2017). Across the country, frequency regulation markets have seen multiple changes in recent years. In 2011, the Federal Energy Regulatory Commission (FERC) issued Order 755 (Federal Energy Regulatory Commission, 2011), which required grid operators to change their frequency regulation compensation mechanisms; we give details below. Various electricity markets across the US have responded with differing changes to their frequency regulation markets (Department of Energy, 2013; Tabari and Shaffer, 2020). These compensation mechanisms favor some resource types more than others, so supply in the regulation market (and therefore dispatch in the energy market) is likely to be affected. The extent to which these changes in compensation mechanisms have impacted electricity markets has not been thoroughly analyzed in the energy economics literature.

Moreover, there has been a growing interest in energy storage devices, such as utility-scale batteries, to provide frequency regulation and other grid support services. Within the PJM market we study, batteries have largely been deployed for frequency regulation, rather than for the intertemporal arbitrage[7] potential explored in energy economics papers (Carson and Novan, 2013; Holladay and LaRiviere, 2018; Antweiler, 2021; Kirkpatrick, 2018; Ambec and Crampes, 2019; Linn and Shih, 2019; Castro, 2020; Butters, Dorsey and Gowrisankaran, 2021; Karaduman, 2021).[8] Worldwide, a primary use of battery storage is for frequency regulation (International Renewable Energy Agency, 2017; Deloitte, 2018). The behavior of batteries providing ancillary services may differ tremendously from that of batteries providing arbitrage – both because of the way the relevant markets are designed and because of the timescale of use (e.g., seconds versus hours), and our results point towards the need for careful study of the effects of this kind of battery deployment.

Finally, the rise of intermittent renewables such as wind and solar generation can increase the amount of frequency regulation required in electricity markets (Kirby, 2004; MIT Energy

[7]For batteries, arbitrage involves charging from the grid when the price is low and selling electricity to the grid when the price is high. As we discuss later, this is not the primary use of most grid-connected batteries in PJM.

[8]There is also a related literature on the use of hydroelectric facilities for storage, again focusing on storage for arbitrage purposes; see e.g. Liski and Vehvilainen (2020).

Initiative, 2011).[9] As generating technologies continue to evolve and batteries and renewable resources play a larger role in both markets, the interactions between energy and ancillary services markets are likely to continue to be important.

3.1 Background

Ancillary Services and Frequency Regulation

Electricity markets are actually made up of numerous interrelated markets: energy, capacity, and several types of ancillary services markets.[10] The energy market is the most studied and best understood by economists – this is where firms are compensated for generating electricity to be used by residential, commercial, and industrial customers. In ancillary service markets, generators are compensated for providing services that enable grid reliability – for example, frequency regulation and other types of reserves. Because the same generators are suppliers of both energy and ancillary services, the structure of ancillary service markets may have important spillovers in the energy provision market. Despite this, there has been very little empirical research into ancillary service markets. In this paper, we focus on frequency regulation, motivated by ongoing policy changes in this market.

Electricity markets are unique in that demand must constantly equal supply, a responsibility that falls on grid operators. However, there are frequent fluctuations in demand and supply, creating small mismatches between the two. When supply exceeds demand, frequency (the number of cycles per second of the alternating current) rises above the nominal frequency (i.e., 60 Hz in North America); when demand exceeds supply, frequency falls below the nominal frequency. If the grid frequency departs enough from the nominal level, it can cause damaged equipment or brownouts and blackouts for customers (Federal Energy Regulatory Commission, 2011).

To prevent this from happening, system operators have created markets to regulate the frequency of the grid, a service called "frequency regulation," "regulation reserves," "load frequency control, "secondary frequency control," or simply "regulation."[11] In the typical market, the system operator sets a frequency regulation requirement – this is the total capacity (in MW) that must be set aside for provision of frequency regulation. It is sometimes time-invariant, sometimes a function of forecasted demand, or possibly also a function of

[9] Ovaere and Gillingham (2019) examines the empirical impact of renewables on the cost of ancillary services provision.

[10] By energy market, we refer to the sale of electricity, in MWh, in a wholesale market. In PJM, this is called simply the "energy market." In this paper, we sometimes refer to "energy provision" to distinguish it from energy markets more generally such as natural gas and oil markets. Since the electricity economics literature frequently does not cover ancillary services markets, it frequently uses terms like "electricity market," "power market," and "wholesale market" without specifying whether the market is for providing energy or capacity or ancillary services or some combination thereof.

[11] Background on frequency regulation and other ancillary services is provided in Hirst and Kirby (1997); Kirby (2004); Hummon et al. (2013); Tacka (2016); Zhou, Levin and Conzelmann (2016).

Figure 3.1: Example Regulation Signals

(a) RegA Signal

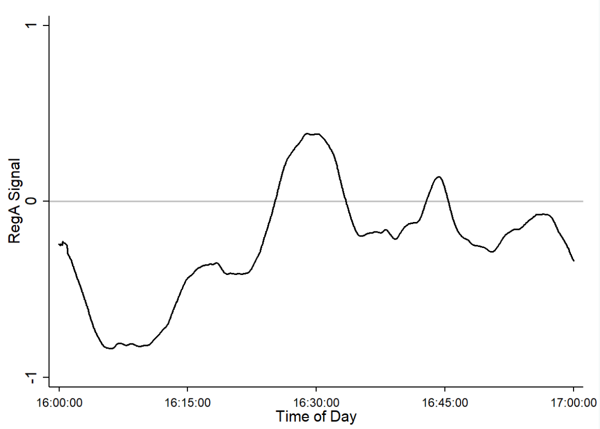

(b) RegD Signal

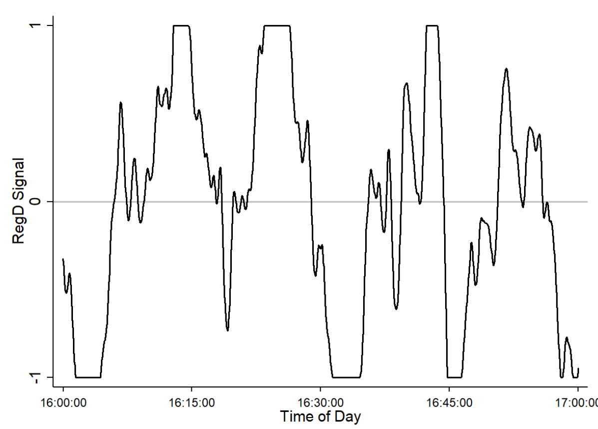

Note: This figure shows the RegA and RegD signals in PJM from 4 pm to 5 pm on July 19, 2019. Data are from PJM.

renewables forecasts. Generators can then bid a portion of their capacity to be available to grid operators to either increase or decrease generation (relative to their set point) at any time (within the frequency regulation contract duration), depending on the needs of the grid. The independent system operator[12] sends out a signal (automatic generation control, or AGC) to participating units, to which they automatically make small adjustments in their generation to balance supply and demand.[13] Typically, these small adjustments are made within seconds (Zhou, Levin and Conzelmann, 2016). An example signal is shown in Figure 3.1.

Some system operators use separate signals for "regulation up" versus "regulation down." The PJM market that we study does not use separate signals for up and down movements. For the time period we analyze, this signal is energy-neutral within a short time-frame (15 minutes), so that units always return to their initial set point (Monitoring Analytics, 2018).

To participate in the market, a power plant must have the technical capability to follow the operator's regulation signal. It must also be dispatched at a non-zero level of generation, with headroom and footroom to follow the signal. That is, it cannot be operating at its minimum or maximum constraints because it must be able to move up and down in response to the operator's signal (Kirby, 2004). The resource mix contributing to frequency regulation varies across regions. In PJM, it is a mix of coal, natural gas, hydro, and battery storage (Monitoring Analytics, 2018).

[12]An independent system operator is non-profit entity that operates the wholesale electricity market.

[13]Specifically, FERC Order 755 defines frequency regulation as "the capability to inject or withdraw real power by resources capable of responding appropriately to a system operator's automatic generation control signal in order to correct for actual or expected Area Control Error needs" ((Federal Energy Regulatory Commission, 2011), p 67266).

Generators take several cost considerations into account when deciding whether and how to bid in a regulation market. Small fluctuations around the generator's set point impact the plant's heat rate, thus changing fuel costs. Providing regulation also imposes wear and tear on the plant (Hirst and Kirby, 1997; Hummon et al., 2013) and can change SO_2 and NO_x emissions, which in some markets are priced.[14]

Over time, system operators have moved towards rewarding regulation providers for both the *capacity* committed to regulation and the *quality* of regulation services provided.[15] Specifically, some units, such as coal-fired boilers, have significant physical inertia that prevents them from responding quickly to regulation signals. In contrast, units such as batteries, hydro generators, and some natural gas generators are able to respond very quickly to the signal, which gives the system operator greater flexibility and speed in restoring the system-wide frequency to its desired level. To incorporate this difference across suppliers, PJM uses a more complicated set of payments.[16] First, PJM sends out two separate regulation signals, one termed "RegA" (for slower-responding units) and one termed "RegD" (for faster-responding units). Moreover, units receive payments both for their capability (i.e., the quantity of MWs offered) and for their performance (the accuracy with which the unit responds to the operator's signal). In Section 3.2, we discuss how the compensation of both capability and performance might impact our results. We discuss the entry of utility-scale batteries for frequency regulation in Section 3.6.

Related Literature

One strand of the economics and engineering literature explores ancillary services markets with optimization models and simulations (Hirst and Kirby, 1998; Just and Weber, 2008; Yu and Foggo, 2017). In particular, Hirst and Kirby (1997) notes that the minimum and maximum constraints of individual generators, when combined with the need to provide headroom and footroom for regulation provision, can lead to a dispatch of units across the system that would not appear least-cost if one only considered the marginal cost of energy provision, and it can also lead to significant complexity in which units are dispatched. Simulations showing changes in dispatch to meet regulation requirements are also given in Hummon et al. (2013). This informs our stylized model in Section 3.2.

A small number of ex-post empirical papers examine how additional operational constraints can lead to out of merit dispatch (e.g., Mansur (2008); Reguant (2014)); these papers focus on dynamic constraints related to startup costs and ramping. However, empirical

[14]The exact cost associated with additional wear and tear is not generally known for individual generators; simulations typically assume a cost that varies across fuel types (see, e.g., Hummon et al. (2013)).

[15]This has been spurred by FERC Order 755 (FERC 2011), which requires that independent system operators change their frequency regulation compensation mechanisms to "pay for performance" systems that recognize the differential speed and accuracy with which different resources respond to the regulation signal. Each independent system operator has designed its pay for performance compensation mechanisms differently (Department of Energy, 2013; Tabari and Shaffer, 2020).

[16]This payment mechanism in PJM was established in October 2012, following FERC Order 755.

analysis of ancillary services markets, particularly in terms of how they interact with energy markets, is limited, despite there being a large literature on wholesale electricity markets.

Several papers use optimization models and/or small-scale simulations to show how frequency regulation markets and other ancillary services markets interact with energy markets. Notably, one report (Atanacio et al., 2012) simulates the impact of storage providing regulation in the PJM system.[17] Specifically, it looks at the emissions changes expected to result across the PJM system when fast-acting storage devices provide 10 percent, 25 percent, or 50 percent of frequency regulation services. It is not clear if the study's proprietary simulation model includes the unit commitment and/or economic dispatch algorithms used by PJM. So, it is not clear if storage providing frequency regulation can impact commitment in the model, nor how exactly it impacts dispatch. The simulation results show small emissions reductions from storage entry, because conventional plants operate less efficiently when providing regulation services. However, the report noted that the interaction between the energy market and the frequency regulation market are complicated and could limit the emissions benefits of storage, and the study's analysis of California's system showed the potential for increased CO_2 emissions (Atanacio et al., 2012).

Two additional papers that examine emissions impacts of using storage for frequency regulation are Lin, Johnson and Mathieu (2019) and Ryan et al. (2018). Ryan et al. (2018) uses a small test system to show that changes in the frequency regulation market can lead to changes in the fuel mix and therefore emissions changes. Indeed, the authors find that "[c]hanges in generator commitment and dispatch caused by the addition of energy storage were the most significant contributors to the energy storage system's environmental impact" (p 10172). Hummon et al. (2013) also explores the interaction between reserve markets and energy markets using a simulation approach, although the paper does not examine emissions outcomes. Cho and Kleit (2015) explores the optimal bidding strategy for a storage device that can provide ancillary services. And finally, in Yu and Foggo (2017), "[s]imulation results with a realistic battery storage system reveal that the majority of the market revenues comes from frequency regulation services" (p 177).

While several of these papers examine ancillary service markets, and even consider spillovers into the energy provision market, this paper is the first to our knowledge to show quasi-experimental evidence of these spillover effects, which lead to fuel use changes and therefore changes in emissions.

3.2 Stylized Model of the Electricity Market

The System Operator's Optimization Problem

In this section, we develop a stylized model of the regulation and energy markets, following Kirschen and Strbac (2004). The goal is not to fully represent all the complexities of the

[17]Papers from the empirical economics literature exploring other aspects of the PJM market include Mansur (2008); Mansur and White (2012); Abito et al. (Forthcoming).

electrical grid, but to show some of the mechanisms by which an increase in the regulation requirement will change generation decisions in the energy market. Specifically, we elucidate how a system operator optimally procures energy and regulation services, especially in the presence of constraints over minimum generation. We use a stylized version of the system operator's problem, which is essentially a single-period unit commitment problem.[18] While we focus on regulation provision, the model has features that apply to reserves markets in general.

We assume that there are multiple heterogeneous thermoelectric generating units participating in an energy market and a regulation market. We evaluate how the plants operate before and after the regulation requirement increases. We use the following notation:

- x_i - generation for energy market for unit i[19]
- y_i - one-sided capacity committed to the regulation market for unit i[20]
- p_x - energy market price per MWh
- p_y - regulation market price per MW
- m_i - marginal cost of generation for unit i, a function of fuel use and wear and tear
- n_i - marginal cost of regulation for unit i, a function of wear and tear
- M_i - minimum generation for unit i, a physical constraint
- C_i - maximum generation for unit i, a physical constraint
- r_i - maximum regulation for unit i, a physical constraint[21]

Additionally, we make the following assumptions:

- Firms are sufficiently small that their actions do not influence the market price in either the regulation or the energy market (i.e. no market power).

[18]Note we do not model the bidding behavior of individual plants. We are essentially assuming that there is no market power, and so plants bid their marginal costs. This is realistic if there are many firms and/or if regulators are able to observe marginal cost and thus punish anti-competitive bidding. The latter is especially likely to be true in our context. Annual market monitoring reports state that the exercise of market power has not generally been observed in PJM's frequency regulation and energy markets (for instance, `https://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2014/2014-som-pjm-volume2-sec1.pdf`, and other years' reports).

[19]This is the set point around which regulation will be provided, if the generator offers frequency regulation.

[20]By one-sided, we mean the capacity available in either direction. The generator must be available to deviate up or down from its set point x_i by the amount y_i.

[21]A typical power plant can commit at most 10 to 20 percent of its capacity to regulation (Makarov et al., 2008; Atanacio et al., 2012).

Figure 3.2: Short-Run Marginal Cost Curves

(a) Energy Market

(b) Regulation Market

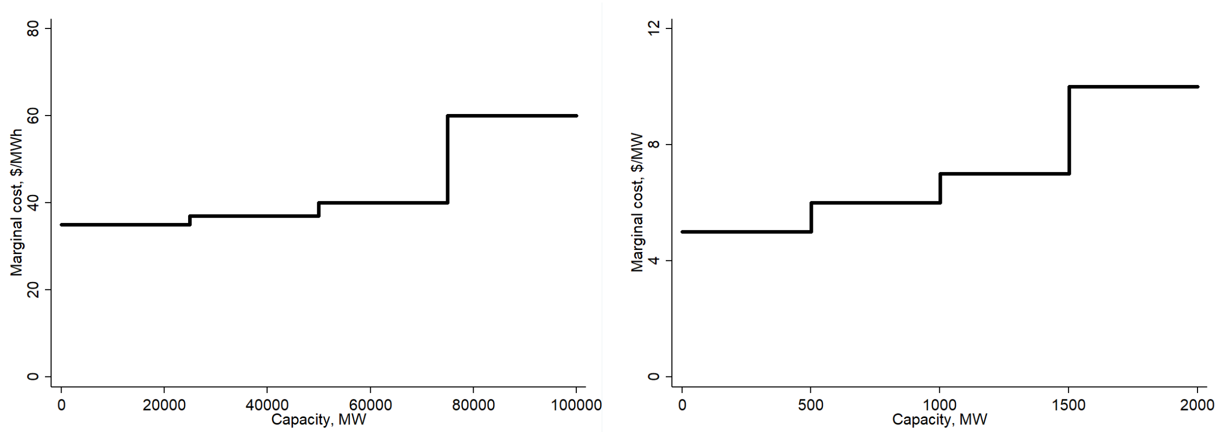

Note: These figures show stylized marginal cost curves for a hypothetical energy market and regulation market, with four types of participating generating units.

- Constant marginal costs of generation and of regulation for each unit, with no fixed costs. These are heterogeneous across plants, a function of the technology installed (fuel choice, prime mover type, etc).[22]

- $r_i < C_i - M_i$ for each unit i.[23]

Focusing on the energy market and leaving aside (for now) minimum constraints, this yields a short-run marginal cost curve that is a step function, as shown in the left-hand panel of Figure 3.2. The height of each step is the marginal cost m_i for each fuel and technology combination, and the width is the maximum generation C_i for each fuel and technology combination. Typically natural gas combined cycle units and coal-fired units are the cheapest to operate, and natural gas combustion turbines operate at much higher cost and with smaller capacities.

Many papers in the economics literature focus on the energy market and elide minimum constraints, yielding a supply curve much like the left-hand panel in Figure 3.2 (Borenstein, Bushnell and Wolak, 2002; Davis and Hausman, 2016). Dispatch order would then follow a least-cost framework, in which the lowest cost units are dispatched, up until demand (exogenously determined) has been fulfilled. In such a supply curve, when demand exogenously

[22]A fully realistic model of the electric grid would allow for non-linear costs within each unit, as in Hirst and Kirby (1997). However, constant marginal costs are frequently assumed in the electricity economics literature (Borenstein, Bushnell and Wolak, 2002; Mansur, 2008; Davis and Hausman, 2016) and are sufficient to illustrate the mechanisms at play in our model.

[23]This is reasonable to assume in our empirical setting. As noted above, a power plant can typically commit 10 to 20 percent of capacity to regulation. Typical minimum constraints are around 30 to 50 percent of maximum capacity. For further discussion of this constraint, see Kirschen and Strbac (2004).

changes, one can examine whether a different unit is on the margin to examine price impacts as well as emissions impacts. This kind of framework is also used in the literature to examine what happens when fuel price changes lead to a re-ordering of the dispatch, i.e. a change in which units are least-cost.

Focusing on the regulation market, our framework implies a similar short-run marginal cost curve, again a step function (right-hand panel of Figure 3.2). Here the height of each step is the marginal cost n_i for each fuel and technology combination, and the width is the maximum regulation r_i for each fuel and technology combination.

Solving for the competitive equilibrium in the energy market when there is no regulation market and there are no minimum constraints is simple, as described above. However, when the system operator is minimizing the cost across these two markets and when minimum constraints are incorporated, we have a more complicated mixed integer linear programming problem:

$$\min_{x_i,y_i} \left(\sum_{i\in(1,2,\ldots I)} m_i x_i + n_i y_i \right) \quad s.t. \sum x_i = \text{demand};$$

$$\begin{aligned} &\sum y_i = \text{regulation requirement}; \\ &x_i + y_i \le C_i \quad \forall i; \\ &x_i - y_i \ge M_i \quad \text{or} \quad x_i = y_i = 0 \quad \forall i; \\ &0 \le y_i \le r_i \quad \forall i; \end{aligned}$$

The operator minimizes the total cost of generation and regulation provision, subject to a number of constraints. The total demand and regulation requirements must be satisfied.[24] Units cannot commit more than their total capacity across the two different markets; they must have sufficient headroom if they offer regulation. The system operator can choose not to dispatch any particular unit, but if the unit operates, it must be at least at its minimum constraint. As a result, if it commits non-zero capacity to the regulation market, it must be at its minimum constraint *plus* its regulation provision in the energy market, i.e. $x_i - y_i \ge M_i$ (i.e. have footroom). These minimum constraints are frequently elided in empirical papers, but they are nontrivial. The typical unit in our data has a minimum constraint of 30 to 50 percent of its maximum capacity.[25] The constraint is related to technical restrictions – operating a plant below minimum load can damage plant equipment.[26]

[24]The magnitude of the regulation requirement refers to the one-directional capacity needed across all units.

[25]The median PJM combustion turbine in our data has a minimum constraint at 50 percent of its maximum; the median non-CT in our PJM data has a minimum constraint at 30 percent of its maximum capacity.

[26]Minimum load constraints are sometimes determined by environmental compliance, if emissions rates are very high at low levels of generation. Cost considerations can also be a factor, if fuel efficiency is very low at low levels of generation.

Table 3.1: Some Potential Effects of an Increase in Regulation Requirement

Pre-period x_i	Post-period x_i	Change in x_i	Pre-period y_i	Post-period y_i	Change in y_i
$x_i = C_i$	$x_i = C_i - r_i$	$-r_i$	$y_i = 0$	$y_i = r_i$	r_i
$x_i = 0$	$x_i = M_i + r_i$	$M_i + r_i$	$y_i = 0$	$y_i = r_i$	r_i

Note: This table shows two potential effects of an increase in the regulation requirement on a generating unit participating in the energy market. In the first row, the generating unit backs down from maximum capacity C to have enough headroom to provide regulation. In contrast, in the second row, the generating unit enters the energy market to have enough footroom to offer regulation. In both cases, the generating unit increases its regulation provision (from zero to r). However the cases show effects in opposite directions and of differing magnitudes in the energy market. Other outcomes are possible as well – for instance if the generating unit had previously been offering some quantity between zero and C in the energy market, or because of market-wide re-dispatch in the energy market.

In this model, corner solutions are possible for many individual generating units. For instance, a unit might operate at maximum capacity in the energy market ($x_i = C_i$) and not participate in the regulation market ($y_i = 0$). It might instead operate just below maximum capacity to fully participate in the regulation market, i.e. with $x_i = C_i - r_i$ and with $y_i = r_i$. It might similarly operate just above minimum capacity to fully participate in the regulation market, with $x_i = M_i + r_i$ and with $y_i = r_i$. There may also be marginal units with generation levels between minimum and maximum capacity, and/or with regulation commitments between 0 and r_i.

Now suppose that the regulation requirement is exogenously increased, and that the change is large enough that the marginal unit cannot provide the additional regulation. The system operator will procure regulation from an additional unit. Suppose this does not require a change in which plants are dispatched in the energy market (an unrealistic assumption to which we return momentarily). Then the system operator might change the energy and regulation procurement from an individual generating unit in multiple ways, as shown in Table 3.1.

To be able to provide regulation services, it is possible that a unit that had previously been operating at maximum capacity C_i would need to back down from its maximum capacity. This could occur if the regulation price change is large enough to outweigh the lost revenues from participating less in the energy market. This outcome is demonstrated in the first row of Table 3.1. It is also possible, however, that a unit that had not been participating in either market could be induced to enter *both* markets. In this scenario, a unit would move from zero generation in the energy market to at least a bit above its minimum operational constraint, operating at $M_i + r_i$ (or more) to be able to sell r_i regulation services. This could occur if its marginal energy cost m_i is above the market clearing energy price p_x, but the additional revenues in the regulation market make up for losses in the energy market. This outcome is demonstrated in the second row of Table 3.1. In short, an increase in the regulation requirement could lead an individual unit to either increase or decrease the energy it sells in the energy market. Moreover, it is possible for the change in the energy market to be *larger* than the change in the regulation market, if a unit is induced to move from not generating at all to generating above its minimum constraint.

We must also consider the follow-on changes for *other* units in the energy market. Since energy demand is exogenous and inelastic, any changes induced by one unit, as described in Table 3.1, must be offset by an equal amount across all other units (neglecting changes in losses due to changes in power flows, which we have not modeled). That is, the change in the regulation requirement could induce not only changes in p_y but also changes in p_x and therefore a different set of plants committed, and different dispatch levels for those plants, in both the regulation and energy markets. The system as a whole could move to a different equilibrium with different inframarginal units, and with some units changing by more than the regulation requirement change. How the system changes will depend on the ways short-run operating profits in one market (e.g., $p_x - m_i$) compare to short-run operating profits in the other market (e.g., $p_y - n_i$) across the entire set of generators.

Simulated Market

We construct a four-unit model, showing that a regulation requirement change can have a wide range of impacts in the energy market. We then solve for the equilibrium, exogenously changing the regulation requirement to show how changes occur along the extensive margin for various plants. The four units represent three baseload units, with differing marginal costs of energy and regulation, and one peaker unit with higher marginal cost for both services (details in Appendix). All four units have minimum operational constraints.[27] All four units are capable of following the regulation signal, which is an energy-neutral signal (units return to their initial set point within a specified time frame). The four units combined must meet an exogenous perfectly inelastic demand requirement as well as a perfectly inelastic regulation requirement.

Figure 3.3 shows how each unit changes its generation (top row) and regulation provision (bottom row) as the regulation requirement is scaled up (results are also shown in table form in the Appendix, Table 4.10). In the top row, note the axes are scaled differently across units.

If there were no minimum constraints and no regulation requirement, Units A and B would provide energy at their maximum (25,000 MWh each) as they have the lowest marginal

[27]The minimum is the same for the three low-cost plants. The high-cost unit has a smaller minimum operational constraint, representing the fact that the peaking portion of the electricity market is made up of many small peaker units that can each be dispatched at quite small levels of generation.

Figure 3.3: Simulated Four-Unit Model

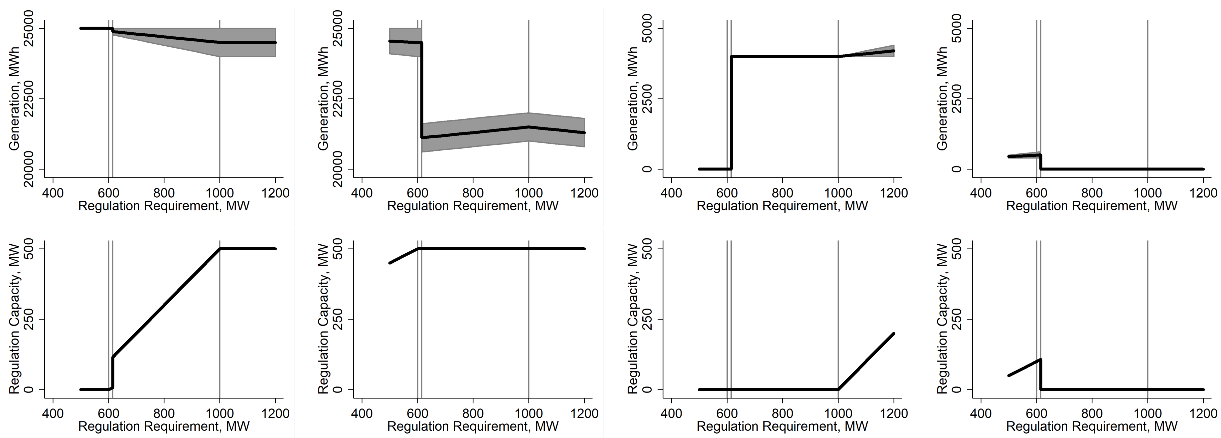

Note: This figure displays the equilibrium results for a four-unit model with energy and regulation output. The top row shows generation outcomes with a black line, surrounded by the regulation band in grey. Generation is in MWh provided (over one hour). The bottom row shows the capacity committed to regulation. Units are ordered left to right from least to most expensive, with the ranking the same across the generation and regulation markets (Unit A has the lowest marginal cost for both services; Unit D the highest). All four units face minimum and maximum constraints. Energy demand is held constant, while the regulation requirement varies exogenously across the x-axis. Discontinuities and kinks are shown with vertical grey lines, at a regulation requirement of 600, 615, and 1000 MW. Quantities are given in table form, along with cost and constraint details, in Appendix Table 4.10.

costs of energy provision. However, the minimum constraints and the regulation requirements change the equilibrium in qualitatively important ways.

Consider first the setting where the regulation requirement is set at 500 MW. Unit A produces the maximum possible energy (25,000 MWh), but Unit B provides only 24,550. It then uses its remaining capacity (450 MW) to provide regulation. Unit C is not dispatched to fill in the remaining 450 MWh to satisfy energy demand. Instead, the most expensive unit, Unit D, provides 450 MWh of energy and 50 MW of regulation services. This is because Unit D, a peaker, has a lower minimum generation requirement than does Unit C – in fact, it operates just at its minimum generation requirement. This illustrates how the minimum generation constraints can alter the dispatch order. (We caution that the results are particularly "lumpy" in that the minimum generation has an especially pronounced impact because there are only four generators; results would be less lumpy in a large market.)

Next, consider what happens as the regulation requirement increases from 500 to 550 or 600 MW. Unit B decreases its energy a bit, to be able to provide additional regulation – it must decrease energy provision to do so, since it had been operating at its maximum constraint when combining both energy and regulation. Unit D also provides a bit more regulation, but to do so, it must *increase* its generation, to maintain status above its minimum constraint. This illustrates how a unit wishing to provide additional regulation could conceivably increase *or* decrease its energy provision, depending on which (if any) of its constraints are binding.

Once the regulation requirement is increased to 601 MW, Unit A begins to decrease its energy, to be able to provide regulation. This is because Unit B is providing as much regulation as possible (500 MW) and cannot provide additional regulation. As the regulation requirement continues to increase, up to 615 MW, additional 1-unit changes in the regulation requirement lead to Unit A decreasing its energy provision and increasing its regulation provision.

The most interesting change occurs when the regulation requirement increases to 616 MW. At this point, the least-cost solution involves dispatching Unit C at its minimum generation. Thus Unit D exits both markets. Previously, this had not been least-cost because of Unit C's minimum constraint. However, it is now least-cost for Unit C to operate at its minimum constraint, rather than to be turned off. So a change in the regulation requirement from 615 MW to 616 MW leads to a very different equilibrium across all units. Unit A decreases its energy provision by around 100 MWh and increases its regulation provision by around 100 MW. Unit B decreases its energy provision by around 3000 MWh and maintains the same amount of regulation provision. Unit C increases from 0 to 4000 MWh of energy, but does not provide any regulation. Unit D drops from over 500 MWh to 0 MWh of energy, and also exits the regulation market. This illustrates how a small change in the regulation requirement can lead to oversized changes in the energy market if it changes whether a unit switches between zero generation and operating at its minimum constraints.

Finally, the last change to occur in this figure is when the regulation requirement goes from 1000 to 1001 MW. Unit C provides the marginal regulation. To do so, it must increase its energy provision since it had been operating at its minimum constraint. Recalling that total generation is fixed, we see that Unit B decreases its energy provision as a result.

We make several caveats regarding the generalizability of our model. We have deliberately presented a stylized version of the energy and regulation markets, to show some of the ways that minimum constraints combine with the multi-product nature of power plants. In practice, there are two different regulation prices in the PJM market: units are rewarded separately for the quantity of MWs offered and for the accuracy with which they respond to the regulation signal. This will impact the mix of, for instance, natural gas combined cycle units versus coal units in the regulation market, since they have differing levels of accuracy. The mix of coal versus natural gas units in the regulation market will also be impacted by secular changes in fuel prices and thus the generation mix of the broader electrical grid.

Also, we have not modeled other features of the market that could interact with the minimum constraints. For instance, our model is static, and dynamic constraints such as minimum up and down times could matter. How exactly they would interact with frequency regulation provision would depend on demand profiles across the day, among other things. Moreover, in reality, system operators run multiple optimization algorithms because there are markets on different timescales: the system operator must make decisions at the day-ahead, hour-ahead, and real-time levels. How regulation interacts with unit commitment will be determined by what algorithms the system operator uses across these different timescales. Also, transmission congestion can matter; as Ryan et al. (2018) write, "we cannot make generalizations about the effects of congestion because, in practice, results would be strongly dependent on grid topology and parameters, generator sizing and location, and so on" (p 10172). Finally, our model consists of only four units; in practice electricity markets are of course much larger. More generating units will mean that supply is less "lumpy" and so one might expect less discontinuous changes than what is observed in Figure 3.3; however in a real-world market, transmission congestion could shrink the number of units that are able to respond.

Overall, there are four primary takeaways from this stylized model: (1) there are potential nonlinearities in the impacts of a regulation requirement change; (2) high marginal cost units can be dispatched over cheaper units because of minimum constraints; (3) an increase in the regulation requirement can either increase or decrease generation at a given unit; and (4) changes in generation at an individual unit can be bigger than the change in the regulation requirement. Thus we see that minimum constraints can be quite important here in that they create lumpiness in how the market responds to exogenous changes, although the economics literature tends to elide them for simplicity. Moreover, the two output markets can interact in surprising ways, with outsized impacts of the regulation market on the energy market. Finally, the specific changes that will be observed following a regulation market change will depend on a suite of parameters. To understand how frequency regulation provision has spillover effects on energy provision, we next turn to an empirical exercise for a large electricity market in the U.S.

3.3 Data

We collect data from the Environmental Protection Agency's Continuous Emissions Monitoring System (CEMS) on hourly generation (MWh) and CO_2 emissions (metric tons)[28] at the generator level, for fossil-fuel-fired units.[29] From CEMS, we also observe the primary fuel source used by a generator each year (coal, natural gas, or oil) and the technology type of the generator (boiler, combined cycle, or combustion turbine). From EIA-860 data, we observe whether units are located in PJM. We also observe in EIA-860 whether units are operated by electric utilities, independent power producers, or as part of a commercial or industrial operation. We drop all commercial and industrial units, as they are unlikely to

[28]The CEMS-reported CO_2 emissions are missing for approximately 9% of observations with non-zero heat input data, representing 2% of generation. In place of these missing values, we assume an emissions rate (per mmBtu of fuel used) equal to the median rate at the unit; see Appendix for details.

[29]CEMS reports *gross* generation rather than *net*, i.e. not accounting for the generation used by the plant itself (or instance, to run pollution control equipment). Net generation is the variable of interest, since that is what is sold in the electricity market. Following Cicala (2022), we scale each unit's generation down from gross to net using monthly generation data from the Energy Information Administration's (EIA) form 923 dataset. This approach also resolves incomplete reporting for some combined cycle units. See Appendix for details.

sell into the electricity market. We report our primary results for four technology categories of PJM units: boilers, combined cycle (CC) units, combustion turbine (CT) units, and all other PJM CEMS units aggregated. We also report results across four fuel types: coal, natural gas, oil, and others aggregated.[30]

The CEMS data do not provide information on non-fossil units (e.g., nuclear, hydro, wind, municipal solid waste), nor does CEMS cover fossil-fuel fired units with capacity less than 25 MW. (For context, units smaller than 25 MW are quite small; the average capacity in EIA-860 data is over 350 MW for steam units in PJM.) To observe the behavior of these units, we calculate a residual category of generation (in MWh), equal to the difference between total demand reported by PJM and total generation reported in CEMS. This residual variable thus covers PJM units not in CEMS as well as net imports into PJM from other regions.

Ideally, we would also observe individual participation in the regulation market. However, this is considered sensitive market information and is not published by PJM. Instead, we can use the CEMS data to observe how generation behavior in the energy market changes as a function of the regulation requirement, which we do observe. Specifically, we collect hourly data, from PJM, on regulation market activity. Our primary explanatory variable is the hourly regulation requirement, in MW. The regulation requirement refers to a pre-determined quantity of regulation capacity that the independent system operator announces it will purchase. As we describe below, this amount varies over time because of several policy changes. In some periods, the regulation requirement is tied directly to the forecasted peak and valley demand for each day, so we also assemble data on these forecasts from PJM.

Summary statistics are provided in Appendix Tables 4.11, 4.12, and 4.14. A time-series of generation by fuel types and technology types is provided in Appendix Figure 4.5. The sample is characterized by a reliance on coal generation and on natural gas generation. The mix of coal versus gas is fairly constant over the time period we study. Other fuel types are present (oil-fired generations; etc.), but are a very small portion of the energy market.

3.4 Empirical Evidence on Generation and Emissions

Identifying Variation

We are interested in how exogenous changes in the frequency regulation market spill over into the behavior of generators in the energy market. As such, we leverage policy changes in PJM's regulation market over the period October 1, 2012 to December 31, 2014. We leverage quasi-experimental variation in the total regulation requirement set by the independent system operator, which changes several times. The benefit of focusing on this period is twofold: battery capacity for frequency regulation is limited and conventional generator entry and exit is small. This allows us to focus on the total regulation requirement and its impact on conventional generators.

[30]For details on "others aggregated," see Appendix.

As of October 1, 2012, the regulation requirement for peak hours was set at 0.78% of forecasted peak load; for off-peak hours it was set at 0.78% of forecasted valley load. On November 22, 2012, this ratio was reduced to 0.74%, then to 0.70% on December 18, 2012. More important than these minor changes is the change on December 1, 2013: the regulation requirement was changed from a percentage of the forecasted peak and valley load to a requirement of 700 MW of effective regulation during peak hours and 525 effective MW during off-peak hours.[31] This variation in the regulation requirement can be seen in Figure 3.4. Also apparent in Figure 3.4 is the considerable variation in the regulation requirement over this time period, from less than 500 to around 1000 MW, which provides us with substantial identifying variation.

A PJM report from 2011 describes the motivation behind these changes: "[d]ecreasing regulation requirements reduces regulation payments" and "[f]ewer resources providing regulation means more resources available for the energy market."[32]

In Section 3.1, we discuss how the regulation *signal* at any point in time is a function of mismatch between supply and demand – thus the signal itself is endogenous to market activity. However, the regulation *requirement* is a pre-determined capacity procurement – set in advance by policy – and thus is not endogenous to day-to-day or hour-to-hour market activity.

While exogenously determined, the requirement could still be correlated with other determinants of generator behavior. For instance, the peak and valley forecasted load may impact generator behavior directly, in addition to being determinants of the regulation requirement – necessitating the inclusion of some controls in our regression.

Our regression takes the form:

$$G_{i,t} = \alpha_i + \beta_i R_t + X_t \Theta_i + \varepsilon_{i,t}, \tag{3.1}$$

where $G_{i,t}$ is generation at unit type i in hour t, R is the regulation requirement, and X_t is a vector of controls. Note this is estimated as a single time-series, but we index G with unit type i, since we can separately estimate the regression for multiple unit types, allowing the parameters to vary by unit type.[33] Standard errors are clustered by sample week.[34]

In this regression, we must control for forecasted peak and valley load,[35] as they directly determine the regulation requirement in the first part of the sample and may also directly

[31]Sources: `http://www.monitoringanalytics.com/reports/pjm_state_of_the_market/2013/2013-som-pjm-volume2.pdf` and `https://www.pjm.com/~/media/documents/manuals/m12.ashx`. "Effective" regulation is a measurement that takes into account the performance of the units providing regulation and the substitutability across RegA and RegD units; see `https://pjm.com/~/media/documents/manuals/m11.ashx`.

[32]`https://www.pjm.com/-/media/committees-groups/committees/mrc/20110915/20110915-item-13-rpstf-update-presentation.ashx`

[33]For expositional purposes, we focus on generation aggregated up to a prime mover type (e.g., boiler) or a fuel type (e.g., coal), but the regression analysis could also be done at the individual unit level.

[34]Clustering by week accounts for serial correlation in unobservables, such as an exogenous multi-day outage at a plant.

[35]Specifically, we use the day-ahead forecasted peak load in hours for which the peak regulation require-

Figure 3.4: Regulation Requirement in PJM

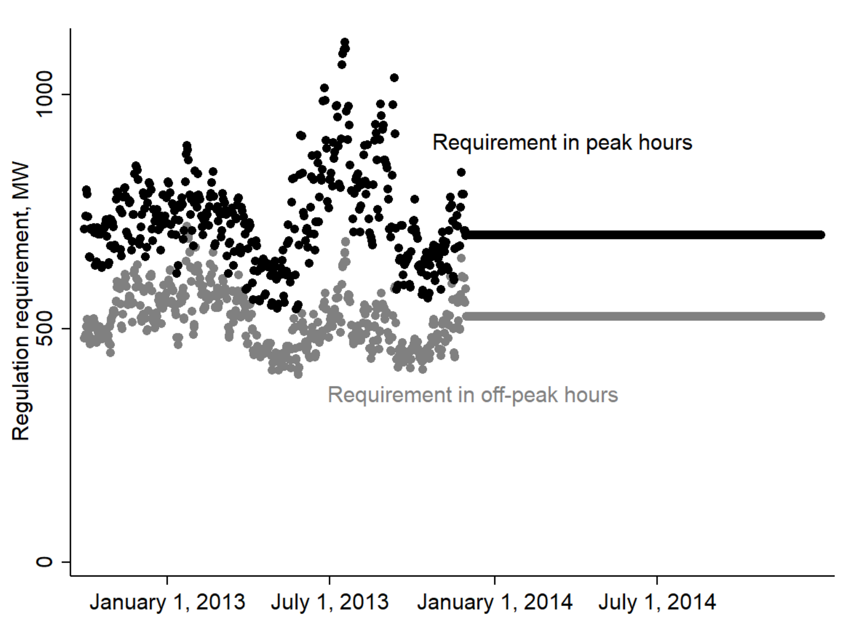

Note: The regulation requirement changes across hours within a day, hence the two different levels plotted on each day. Peak hours are defined 4 a.m. to midnight, and off-peak hours as midnight to 4 a.m. In the raw data, one hour (on 4/2/2013) is listed as having a regulation requirement of zero; this hour is dropped from the regressions. Data source is PJM.

impact generator behavior (i.e., not including these controls X_t would mean that $\varepsilon_{i,t}$ would be correlated with R_t).[36] The policy change thus allows us to separately identify the impact of forecasted peak and valley load from the impact of the regulation requirement. Specifically, variation in forecasted peak and valley loads during the second half of the sample identifies the effect Θ of these variables on generator behavior. Assuming that Θ does not change between the first and second halves of the sample, the first time period can then be used to estimate the impact of the regulation requirement on generator behavior conditional on the peak and valley forecasts (i.e., we can assume that $\varepsilon_{i,t}|X_t$ is uncorrelated with R_t). In the Appendix, we provide a simulation of this identification strategy, where we directly control the data-generating process.

One might also worry that other variables are correlated with both the regulation requirement and with generator behavior. Following the literature on generator behavior (Holland and Mansur, 2008; Wolak, 2011; Cullen, 2013; Cullen and Mansur, 2017; Fell and Kaffine,

ment applies (4 a.m. to midnight) and set the variable equal to zero in off-peak hours. Similarly, we use the day-ahead forecasted valley load in hours for which the off-peak regulation requirement applies (midnight to 4 a.m.) and set the variable equal to zero in peak hours.

[36]We also include a peak versus off-peak hour dummy variable.

2018; Leslie, 2018; Bushnell and Novan, 2021), we include additional control variables. We control for total demand in the PJM system; for fuel prices; and for weather.[37] Failing to include these could introduce bias if they are correlated with the regulation requirement.

We also follow the literature in including time period dummies (month of sample, hour of day, and day of week dummies) to control flexibly for other exogenous changes in the PJM market over this time period (e.g. seasonality, macroeconomic shocks, and secular trends in entry and exit). Finally, we also control for the total generation by PJM units appearing in the CEMS data. Combined with the control for total demand, this is equivalent to controlling for total demand net of nuclear, solar, wind, biomass, and other renewable or non-CEMS generation.

Regression Results

We first show regression results aggregated by technology type, in Table 3.2 (coefficients on control variables are shown in the Appendix, Table 4.15). We see that when more regulation is required, the unit types providing services in the energy market change, with decreased generation from boiler units and increased generation from combined cycle units. Specifically, for each 100 MW of additional regulation capacity required by the system operator, boiler units decrease their generation by 390 MWh and combined cycle units increase their generation by 360 MWh (both statistically different from zero at the five percent level). The difference between these two comes from a small number of combustion turbines and units with other technology types (right-most columns).

Because different technology types use different primary fuel sources, the results in Table 3.2 will have implications for fuel use and therefore environmental impacts. We next show regression results aggregated by fuel type, in Table 3.3. Most PJM boiler units burn coal, although some burn natural gas or oil. Similarly, most combined cycle units and combustion turbines use natural gas, but some use oil. Consistent with the boiler results in Table 3.2, we see a decrease in coal-fired generation and oil-fired generation in Table 3.3. Similarly, consistent with the combined cycle results, we see an increase in natural gas generation.

To understand the emissions impact of this change in the generation mix, we can estimate a similar regression with CO_2 emissions as the dependent variable. We now aggregate CO_2 emissions across all units in CEMS. We see in the right-most column of Table 3.3 that CO_2 emissions fall when the regulation requirement is raised: for every 100 MW increase in the regulation requirement, we estimate 240 fewer tons of CO_2 are emitted by units participating in the energy market, statistically significant at the one percent level. This is in line with the fuel use changes described above. Specifically, the PJM-wide emissions rate for natural gas units (combining CC and CT units) in our sample is 0.42 tons CO_2 per MWh; the emissions rate for coal and other units combined is 0.95. This difference implies that a shift of 400

[37] Demand data are from PJM. Natural gas, oil, and coal prices are from the Energy Information Administration. Weather is from NOAA. See Appendix for details.

Table 3.2: The Impact of the Regulation Requirement on the Energy Market

	Boiler (MWh)	CC (MWh)	CT (MWh)	Other tech. (MWh)
Regulation requirement, 100 MW	-388.9**	357.1**	16.2	15.5
	(189.7)	(162.4)	(208.6)	(22.1)
Observations	19,693	19,693	19,693	19,693
Within R^2	0.91	0.71	0.49	0.13
Mean of dep. var.	40,185	14,290	848	554

Note: This table shows estimates from four separate time-series regressions. In all columns, the dependent variable is total MWh of electricity generated per hour in the PJM market by units that appear in CEMS data (where each column aggregates across all units of a particular type). Coefficients on control variables are shown in the Appendix, Table 4.15. The unit of analysis is an hour. Standard errors are clustered by sample week.

MWh from coal to gas would lead to a decrease in CO_2 emissions of around 210 tons, close to our estimate of 240 tons in Table 3.3.[38]

Robustness Checks

In the Appendix, we explore a suite of robustness checks and placebo regressions. We show a number of regressions that include additional controls. For instance, we control for the standard deviation of the regulation requirement over the previous 72 hours – one might be concerned that our main results are driven by the change in the variance of the regulation requirement from the first to second half of our sample, visible in Figure 3.4. We also show regressions with alternative functional forms and/or non-parametric specifications for the control variables. Next, we collapse to a daily specification to allow for across-hour effects as in Bushnell and Novan (2021). We also show a specification that limits the sample to units that do not retire during our sample period, to ensure that the main results are driven by these non-retiring units as opposed to units that retired (perhaps for secular reasons) over this time period. Next, we calculate alternative standard errors, specifically Newey-West standard errors with a maximum lag length of 168 hours (one week).[39]

Across all of these specifications, we estimate fuel use shifts (coal to natural gas), with magnitudes and statistical significance comparable to what we display in Tables 3.2 and 3.3.

[38]The difference between this back-of-the-envelope calculation and our estimate may be due to differential heat rates of marginal versus average units, and/or to a heat rate effect of regulation provision that we describe in the Appendix.

[39]If there is cross-day correlation, then a rule of thumb like $N^{1/4}$ (in this case = 12 hours) would be insufficient, so we use a more conservative lag length.

Table 3.3: The Impact of the Regulation Requirement on the Energy Market

	Coal (MWh)	Natural gas (MWh)	Oil (MWh)	Other fuel (MWh)	CO2 (tons)
Regulation requirement, 100 MW	-351.0	432.6**	-97.7	16.1	-242.4***
	(234.3)	(192.9)	(68.2)	(22.0)	(89.0)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.88	0.82	0.46	0.13	0.98
Mean of dep. var.	39,818	14,906	570	581	45,215

Note: This table shows estimates from five separate time-series regressions. For the first four columns, the dependent variable is total MWh of electricity generated per hour in the PJM market by units that appear in CEMS data (where each column aggregates across all units of a particular fuel type). For the right-most column, the dependent variable is CO_2 emissions (tons) per hour for all PJM units in CEMS (i.e. combining across the four unit types from the first four columns). Coefficients on control variables are shown in the Appendix, Table 4.16. The unit of analysis is an hour. Standard errors are clustered by sample week.

We also estimate a negative and statistically significant impact of the regulation requirement on CO_2 emissions. We also show more parsimonious regressions, for which we estimate qualitatively similar results but where, not surprisingly, we lose precision.

Additional regressions in the Appendix show estimated effects at a number of placebo units (Table 4.21). We run our primary regression using generating units in nearby states that do not participate in the PJM wholesale energy market.[40] We also separately use CEMS-reporting units that are classified as part of commercial or industrial operations; this includes generation from facilities such as hospitals and petroleum refineries. We also estimate our main specification with wind generation as the dependent variable. Finally, we consider a "residual" category of generation as the dependent variable. Following Davis and Hausman (2016), we estimate the effects of the regulation requirement on the amount of generation that would be needed to satisfy total demand, after accounting for the generation quantity reported in CEMS. This category accounts for nuclear generation, hydro generation, net imports, and small units not reported in CEMS. Across all but one of these categories, we estimate only small effects of the regulation requirement. For the non-PJM coal units, we estimate a larger effect, but it is noisy and has the *opposite* sign of the coal estimates in Table 3.3.

[40]Specifically, our full dataset contains data on all CEMS units in Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, North Carolina, New Jersey, Ohio, Pennsylvania, Tennessee, Virginia, Washington DC, and West Virginia. These states are covered in part by PJM. However, in states such as Indiana, only a minority of the plants participate in the PJM market, with the rest primarily participating in the MISO wholesale market. Our placebo sample consists therefore of, e.g., MISO-participating units in states such as Indiana.

We also estimate a series of two-way fixed effects regressions, in which the plants in nearby states serve as "control" units, allowing us to include day-of-sample effects. Results are qualitatively similar, with comparable point estimates.[41]

Overall, our empirical estimates show that when more frequency regulation is needed, substantial fuel use changes occur in the energy market. In particular, coal units sell less in the energy market while natural gas units sell more. This leads to an overall decrease in CO_2 emissions, holding generation constant. This magnitude is quite large – why would a 100 MW change in the regulation market lead to a roughly 400 MWh change in the energy market? Our simulated model suggests that this could occur if unit behavior around minimum constraints is changing. To further understand the magnitude and the mechanism behind it, we turn to regressions that incorporate information about minimum and maximum constraints.

3.5 Evidence of Changes Along Extensive Versus Intensive Margins

To further connect our modeling results to our empirical results, we next estimate various intensive versus extensive margin changes. We separate hourly generation into five bins for each unit: *Off*, *Below Minimum Constraint*, *At Minimum Constraint*, *Between Minimum Constraint and Maximum Capacity*, and *At Maximum Capacity*. Then we count the number of units of each fuel type in each bin in each hour, giving us a time series of bin-level counts for each fuel type. Details of the data and variable construction are in the Appendix.

We regress the count of generators falling into each bin on the regulation requirement and a vector of controls separately for each fuel/mover type. The regressions take the form:

$$N_{i,t} = \alpha_i + \beta_i R_t + X_t \Theta_i + \varepsilon_{i,t}, \tag{3.2}$$

where $N_{i,t}$ is a count of units of fuel type i in hour t that have generation levels falling in a particular bin (e.g., the number of coal boiler units at 5 a.m. on November 1, 2012 with capacity factors below their minimum constraint). Again R_t is the regulation requirement and X_t is a vector of controls (the same controls as in the generation regressions above). Standard errors are clustered by sample week.

Table 3.4 shows the results of these regressions for each of the three units types of interest: boiler, combined cycle, and combustion turbine (results for "other" units are shown in the Appendix, Table 4.23). These effects are in line with both the generation changes shown in

[41] This two-way fixed effects specification does not serve as our main specification for a couple of reasons, as described in the Appendix. Even in this specification, our primary source of identification is the switch from a varying regulation requirement in the first half of the sample to a flat requirement in the second half. Thus "control" units in other states are not expected to provide much additional help with identification. The time-series specifications are a more transparent way to implement this identification strategy. The two-way fixed effects can help if there additional regional or national secular trends that are not adequately controlled for, and as such it is reassuring that they yield qualitatively similar estimates.

Table 3.2 and with the stylized model. Panel A shows that boilers are less likely to be at their maximum (Column 5) and more likely to operate within their main operational range (Column 4) when the regulation requirement is higher. Specifically, a 100 MW increase in the regulation requirement causes one fewer units to be at maximum capacity. This is consistent with this unit providing additional regulation, and needing its set point to be below its maximum to have the flexibility for upward movements in response to a regulation signal. It is also consistent with the boiler decreasing its energy provision to accommodate additional energy provision by combined cycle plants, discussed next.

Panel B shows that combined cycle units are more likely to be dispatched when the regulation requirement is higher. For every 100 MW additional regulation requirement, around two combined cycle units are more likely to be dispatched and in their main operational range (Column 4). This could be consistent either with dispatching with positive generation to be able to themselves provide frequency regulation, or with needing to fill a gap left by reduced boiler generation, shown in Panel A.

Panel C shows that combustion turbines are also more likely to be operating within their main operational range, some combination of units being less likely to be off (Column 1) or less likely to be at their maximum capacity (Column 5). This could be because they move to the middle of their range to provide frequency regulation, because they are newly dispatched to fill in for lost boiler generation, or some combination of both.

As shown in the Appendix, results are robust to an alternative construction of the minimum constraint variable. Results are also robust to using ten bins, identically spaced across the capacity of each unit (0 to 10 percent of capacity, 11 to 20 percent of capacity, etc.), rather than a minimum constraint definition.

Overall, the primary effect we see when the regulation requirement is higher is that more units operate within their main operational range, rather than being off or at the maximum constraint. This is consistent with the model in Section 3.2. The magnitudes are also consistent with a back-of-the-envelope calculation of the number of units that would be needed to provide 100 MW of regulation. Recall that the typical unit can commit 10-20 percent of its capacity to regulation. In our sample, the average boiler or combined cycle unit has a capacity of around 230-280 MW, implying that two to four units would be needed to provide 100 MW of regulation. Combustion turbine units in our sample have a capacity of around 80 MW, so more of these plants would be needed to provide the same amount of regulation.

We see empirically that generators behave in intuitive ways along both the intensive and extensive margins when the regulation requirement is exogenously changed. Because they are multi-product firms, they adjust their outputs in multiple markets. Indeed, recall that one of the motivations behind the policy changes to the regulation requirement was to free up capacity for the energy market (Section 3.4). Moreover, minimum and maximum constraints can lead to changes in the energy market that are outsized in comparison with the change in the regulation requirement.

Table 3.4: The Regulation Requirement and Extensive Versus Intensive Margins

Panel A. Boilers	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.29 (1.03)	0.20 (0.16)	0.39 (0.40)	1.18 (0.90)	-1.48*** (0.47)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.71	0.11	0.38	0.43	0.75
Mean of dep. var.	199	5	11	99	29
Panel B. Combined Cycle Plants	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.33*** (0.86)	0.03 (0.25)	0.40* (0.24)	2.31*** (0.70)	-0.40 (0.25)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.22	0.07	0.61	0.32
Mean of dep. var.	56	13	9	55	2
Panel C. Combustion Turbines	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.13 (3.38)	0.01 (0.46)	0.31 (0.33)	3.06 (2.44)	-1.26** (0.59)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.31	0.27	0.45	0.19
Mean of dep. var.	410	3	1	8	1

Note: This table shows estimates from 15 separate regressions. The dependent variable is a variable representing the count of units of each type generating at each level in PJM. The unit of analysis is an hour. Effects for other unit types are shown in the Appendix, Table 4.23. Standard errors are clustered by sample week.

3.6 Discussion

In this section, we discuss several broader implications of our results. First, as described in Section 3.4, we find CO_2 impacts resulting from changes in frequency regulation requirements. Specifically, for every 100 MW increase in the regulation requirement, we estimate a decrease of 240 tons of CO_2 per hour. This represents a decrease of 0.5% of CO_2 emissions, or an annual total of 2 million tons of reduced CO_2. Valued at the IWG Social Cost of

Carbon, the short-term generation changes from an increase in the regulation requirement would reduce climate damages by around \$90M per year.[42] Recent peer-reviewed estimates would place the value even higher (Pindyck, 2017; Moore et al., 2017; Ricke et al., 2018; Bastien-Olvera and Moore, 2020).

A full welfare analysis would compare these emissions impacts to changes in fuel costs, operations and maintenance costs, and the external costs associated with other pollutants (such as particulate matter and sulfur dioxide). We do not attempt this calculation because we do not observe operations and maintenance (O&M) cost data across units or across time. Changes in frequency regulation provision would be expected to impact wear-and-tear on the plant, and therefore O&M costs. Similarly, changes in plant utilization (e.g., capacity factors) would also be expected to impact O&M costs. Moreover, changes to regulation requirements also impact the grid in hard-to-quantify ways relating to resilience and the probability of blackouts in rare events.

However, we note a few things. First, we would expect that an increase in the regulation requirement would lead to an increase in total overall costs, as regulation provision is costly. For our four-unit model in Section 3.2, total costs (combining energy provision and regulation provision costs) increase with the regulation requirement. Whether these private cost increases outweigh the emissions cost savings described above is an empirical question. For the time period we study, we observe a typical coal price of \$2.34 per mmbtu and a natural gas price of \$4.71 per mmbtu.[43] Supposing a coal heat rate of 10.3 MWh per mmbtu and a natural gas heat rate of 7.7 MWh per mmbtu, this would imply that a 400 MWh change from coal to natural gas from an increased regulation requirement would cost around \$5,000 per hour in increased fuel costs (again, note this ignores changes in O&M costs).[44] In contrast, the CO_2 reductions would save over \$10,000 per hour in external climate damages. In short, because coal is so much more CO_2-intensive than natural gas, the private cost changes could easily be outweighed by the external cost changes.

Second, our results on frequency regulation markets have implications for utility-scale battery storage. There is a growing interest in using energy storage, including batteries, flywheels, and loads coordinated to behave like storage, to help operate the electrical grid. The energy economics literature has focused on the use of batteries for arbitrage: charging when demand is low (e.g., at night) and discharging when demand is high (e.g., the late afternoon). However, batteries can also be used to provide ancillary services such as frequency regulation (Department of Energy, 2013; International Renewable Energy Agency, 2017; Deloitte, 2018; Ryan et al., 2018). PJM has been at the forefront of incorporating storage into ancillary service markets; storage providers found RegD particularly lucrative when it was first introduced.[45] Indeed, nearly all storage capacity in PJM is built for the

[42]We use the 2015 SCC from the Interagency Working Group on Social Cost of Greenhouse Gases (2016) and convert to 2019 dollars using the CPI, implying a value of \$44 per metric ton.

[43]Source: EIA's "Average cost of fossil fuels for electricity generation."

[44]If we include oil costs, we would estimate *decreased* fuel costs, as oil prices paid by generators are over \$20 per mmbtu, and we show a negative albeit very noisy estimate on oil-fired generation in Table 3.3.

[45]See Maloney (2017), "Is the bloom off the RegD rose for battery storage in PJM?" in *Utility Dive*, `https:`

provision of frequency regulation.[46]

Consider a battery entering the market in order to participate in the regulation market. Suppose this battery is inframarginal – batteries generally have high fixed costs but low marginal costs – and suppose that the battery does not participate in the energy market (recall from above that batteries in PJM generally do not provide arbitrage in the energy market). If this battery participates in each period, its entry represents a reduction in the residual regulation requirement faced by conventional generators, by a magnitude equal to the capacity of the battery. Based on our analysis of the PJM regulation market, we expect that the entry of batteries in PJM would lead to generation mix changes and emissions changes in the energy market. For the time period we study (2012-2014), we could infer that battery entry (akin to a reduction in the regulation requirement) would lead to *increased* CO_2 emissions, with a gas to coal shift. This provides empirical support for simulation evidence in the engineering literature. Specifically, Ryan et al. (2018) uses a unit commitment and dispatch model of a small power system, not calibrated to PJM, to show the life-cycle environmental impacts of batteries. That research similarly finds fuel switching effects for conventional generators when batteries are introduced to provide frequency regulation.[47]

In sum, our results have qualitatively important implications for climate policy and for battery deployment. The generation mix changes that could result from frequency regulation changes will have impacts on CO_2 emissions. A few brief caveats are worth noting. Recall that in the stylized model, we see that the changes in regulation can have many potential impacts on the energy market. Thus, the mechanisms we describe will be broadly relevant, but our specific regression results are not necessarily externally valid: the results cannot be simply extrapolated to alternative time periods in PJM nor to other system operators. Different fuel costs or the secular retirements of power plants could mean very different impacts of changes to the regulation market.[48,49] In addition, the introduction of batteries could

//www.utilitydive.com/news/is-the-bloom-off-the-regd-rose-for-battery-storage-in-pjm/503793/.

[46]22 of PJM's 27 facilities list "frequency regulation" as a service; only five facilities list "load management" and just one lists "arbitrage" as a service, according to Energy Information Administration data. Source: EIA-860 data for 2018.

[47]Note we also expect two additional impacts of battery entry. Battery entry can impact the heat rates of conventional generators, similar to how frequency regulation impacts the heat rate, a point we discuss in Section 3.1 and in the Appendix. Furthermore, we note that batteries are net users of electricity (they do not have 100% round-trip efficiency), and so their entry impacts the amount of conventional power plant generation required (Department of Energy, 2016). This mechanism is not captured with the regression approach we have taken, which conditions on total quantity demanded across the system.

[48]For related mechanisms in the context of wind energy, see Callaway, Fowlie and McCormick (2018); in the electric vehicle context, see Holland et al. (2016*a*,*b*).

[49]The fact that we see non-linearities in the model in Section 3.2 suggests that even *within* a given geographic region or broad time period, one might expect to see heterogeneity in the impact of a change to the regulation requirement across different levels of demand or different ex-ante levels of the regulation requirement itself. Unfortunately, we have insufficient power for exploring this heterogeneity ourselves. But, future work could look at this in depth with either additional sources of identifying variation or with additional simulations calibrated to various markets.

enable additional renewables entry,[50] which could lower CO_2 emissions. The introduction of carbon pricing would also be expected to change the relationship between ancillary services and energy markets. Future empirical work in these areas will be important as the electrical grid continues to evolve.

3.7 Conclusion

Overall, we see that changes in the structure and makeup of the frequency regulation market impact generators that participate in the energy market. We present a model that allows generators to be multi-product suppliers, additionally making the realistic assumption that units are constrained by non-zero minimum and maximum constraints. With these assumptions, we show that policy and market changes can cause power plants to move from fully off to operating at non-negligible minimum load, and vice-versa. Moreover, our model and results suggest that generating units should be thought of as multi-product firms participating in multiple markets, and that the concept of a "marginal" firm is much more complex after considering both minimum constraints and multiple markets.

Turning to empirical estimation for the PJM market, we show that for every additional 100 MW increase in frequency regulation required of plants in PJM, there is an approximately 360 MWh increased use of combined cycle power plants in the energy market, and a corresponding decrease in the use of boiler units. The results are directionally robust to considering alternative controls and various alternative specifications. Results also pass a series of placebo tests using generators outside of the PJM market. However, the magnitudes and even the direction of the effects are specific to the time period and market we study – as our theoretical model shows, the results depend on the composition of generators operating in the market (and therefore on fuel prices, etc).

Most importantly, this paper demonstrates that ancillary services markets and energy markets are far more intertwined than economics researchers might have previously thought. The structure and policy details of ancillary service markets have important impacts for generators and for the energy market, and more careful research across the country is necessary to better understand these complexities in different settings. Future economics research could use structural methods to estimate power plant costs in both markets and evaluate market power after accounting for the existence of both minimum constraints and multiple markets. Future work could also incorporate dynamic constraints (e.g. minimum up and down times) and transmission constraints. Finally, we have focused on short-term impacts. If regulation market changes impact profitability and therefore retirement decisions, long-term effects could differ.

We also find that regulation market changes can impact the fuel mix in the energy market and therefore CO_2 emissions. Our results are specific to a second-best world in which CO_2 emissions are not priced. This has policy importance as the grid continues to evolve to

[50]For research on the interaction of renewables and storage, see Gowrisankaran, Reynolds and Samano (2016); Ovaere and Gillingham (2019), and Karaduman (2021).

mitigate climate change (e.g., by incorporating renewables). How batteries, renewables, and conventional power plants will interact in markets for both generation and grid reliability will continue to be an important question, pointing towards the complexities inherent in designing second-best greenhouse gas abatement policy.

Bibliography

Abito, Jose-Miguel, Christopher R. Knittel, Konstantinos Metaxoglou, and Andre Trindade. Forthcoming. "The Role of Output Reallocation and Investment in Coordinating Externality Markets." *International Journal of Industrial Organization*.

Alberini, Anna, and Massimo Filippini. 2011. "Response of residential electricity demand to price: The effect of measurement error." *Energy Economics*, 33(5): 889–895.

Alberini, Anna, Will Gans, and Daniel Velez-Lopez. 2011. "Residential consumption of gas and electricity in the US: The role of prices and income." *Energy Economics*, 33(5): 870–881.

Ambec, Stefan, and Claude Crampes. 2019. "Decarbonizing Electricity Generation with Intermittent Sources of Energy." *Journal of the Association of Environmental and Resource Economists*, 6(6): 1105–1134.

Antweiler, Werner. 2021. "Microeconomic models of electricity storage: Price forecasting, arbitrage limits, curtailment insurance, and transmission line utilization." *Energy Economics*, 101.

Archsmith, James. 2020. "Dam Spillovers: The direct and indirect costs from environmental constraints on hydroelectric generation." *Working paper*.

Atanacio, Manuel, Richard Fioravanti, Warren Katzenstein, and Khoi Vu. 2012. "Emission II Study of Advanced Storage used for Frequency Regulation." *Department of Energy – Sandia National Labs, Project Number: Sandia Negotiation 96362*.

Auffhammer, Maximilian, and Edward Rubin. 2018. "Natural gas price elasticities and optimal cost recovery under consumer heterogeneity: Evidence from 300 million natural gas bills." *National Bureau of Economic Research*, , (w24295).

Aydin, Erdal, Nils Kok, and Dirk Brounen. 2017. "Energy efficiency and household behavior: the rebound effect in the residential sector." *The RAND Journal of Economics*, 48(3): 749–782.

Bastien-Olvera, Bernardo A., and Frances C. Moore. 2020. "Use and non-use value of nature and the social cost of carbon." *Nature Sustainability*.

Blomquist, Sören, and Håkan Selin. 2010. "Hourly wage rate and taxable labor income responsiveness to changes in marginal tax rates." *Journal of Public Economics*, 94(11-12): 878–889.

Borenstein, Severin. 2017. "Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates." *Journal of Public Economics*, 4(S1): S85–S122.

Borenstein, Severin, James B. Bushnell, and Frank A. Wolak. 2002. "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review*, 92(5): 1376–1405.

Borensten, Severin, and James Bushnell. 2015. "The US Electricity Industry After 20 Years of Restructuring." *Annual Review of Economics*, 7: 437–464.

Brolinson, Becka. 2019. "Does Increasing Block Pricing Decrease Energy Use?: Evidence from the Residential Electricity Market." *Georgetown University, Department of Economics.*

Bushnell, James, and Kevin Novan. 2021. "Setting with the Sun: The Impacts of Renewable Energy on Conventional Generation." *Journal of the Association of Environmental and Resource Economists*, 8(4): 759–796.

Butters, R. Andrew, Jackson Dorsey, and Gautam Gowrisankaran. 2021. "Soaking Up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium." *NBER Working Paper No. 29133.*

Callaway, Duncan S., Meredith Fowlie, and Gavin McCormick. 2018. "Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources." *Journal of the Association of Environmental and Resource Economists*, 5(1): 39–75.

Carson, Richard, and Kevin Novan. 2013. "The Private and Social Economics of Bulk Electricity Storage." *Journal of Environmental Economics and Management*, 66: 404–423.

Castro, Miguel. 2020. "Intermittent Renewable Energy, Hydropower Dynamics and the Profitability of Storage Arbitrage." *IDB Working Paper 1107.*

Census, US. 1994. "Census Tracts and Block Numbering Areas." *https://www2.census.gov/geo/pdfs/reference/GARM/Ch10GARM.pdf.*

Chatterjee, Arpita, Rafael Dix-Carneiro, and Jade Vichyanond. 2013. "Multi-Product Firms and Exchange Rate Fluctuations." *American Economic Journal: Economic Policy*, 5(2): 77–110.

Chesser, Michael, Jim Hanly, Damien Cassells, and Nicholas Apergis. 2018. "The positive feedback cycle in the electricity market: Residential solar PV adoption, electricity demand and prices." *Energy Policy*, 122: 36–44.

Cho, Joohyun, and Andrew N. Kleit. 2015. "Energy Storage Systems in Energy and Ancillary Markets: A Backwards Induction Approach." *Applied Energy*, 147: 176–183. doi: 110.1016/j.apenergy.2015.01.114.

Cicala, Steve. 2015. "When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation." *American Economic Review*, 105(1): 411–444.

Cicala, Steve. 2022. "Imperfect Markets versus Imperfect Regulation in U.S. Electricity Generation." *American Economic Review*, 112(2): 409–441.

Crago, Christine Lasco, and Ilya Chernyakhovskiy. 2017. "Are policy incentives for solar power effective? Evidence from residential installations in the Northeast." *Journal of Environmental Economics and Management*, 81: 132–151.

Cullen, Joseph. 2013. "Measuring the Environmental Benefits of Wind-Generated Electricity." *American Economic Journal: Economic Policy*, 5(4): 107–133.

Cullen, Joseph A., and Erin T. Mansur. 2017. "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution." *American Economic Journal: Economic Policy*, 9(3): 106–133.

Davis, Lucas. 2020. "What Matters for Electrification? Evidence from 70 Years of US Home Heating Choices." *National Bureau of Economic Research*, , (w28324).

Davis, Lucas, and Catherine Hausman. 2016. "Market Impacts of a Nuclear Power Plant Closure." *American Economic Journal: Applied Economics*, 8(2): 92–122.

De Loecker, Jan. 2011. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica*, 79(5): 1407–1451.

Deloitte. 2018. "Supercharged: Challenges and opportunities in global battery storage markets."

Department of Energy. 2013. "Grid Energy Storage." *https://www.energy.gov/sites/prod/files/2014/09/f18/Grid%20Energy%20Storage%20 December%202013.pdf.*

Department of Energy. 2016. "Demand Response and Energy Storage Integration Study."

Dergiades, Theologos, and Lefteris Tsoulfidis. 2008. "Estimating residential demand for electricity in the United States, 1965–2006." *Energy Economics*, 30(5): 2722–2730.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2019. "The long-run dynamics of electricity demand: Evidence from municipal aggregation." *American Economic Journal: Applied Economics*, 12(1): 86–114.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes. 2018. "Just Starting Out: Learning and Equilibrium in a New Market." *American Economic Review*, 108(3): 565–615. https://doi.org/10.1257/aer.20160177.

Eckel, Carsten, and J. Peter Neary. 2010. "Multi-Product Firms and Flexible Manufacturing in the Global Economy." *The Review of Economic Studies*, 77: 188–217.

Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram. 2007. "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency." *American Economic Review*, 97(4): 1250–1277.

Federal Energy Regulatory Commission. 2011. "Order No. 755, Frequency Regulation Competition in Organized Wholesale Power Markets." *Federal Register, Vol. 76, No. 210.*

Feehan, James. 2018. "The long-run price elasticity of residential demand for electricity: Results from a natural experiment." *Utility Policy*, 51: 12–17.

Fell, Harrison, and Daniel T. Kaffine. 2018. "The Fall of Coal: Joint Impacts of Fuel Pices and Renewables on Generation and Emissions." *American Economic Journal: Economic Policy*, 10(2): 90–116.

Gillingham, Kenneth, David Rapson, and Gernot Wagner. 2020. "The rebound effect and energy efficiency policy." *Review of Environmental Economics and Policy.*

Gowrisankaran, Gautam, Stanley S. Reynolds, and Mario Samano. 2016. "Intermittency and the Value of Renewable Energy." *Journal of Political Economy*, 124(4): 1187–1234.

Hirst, Eric, and Brendan Kirby. 1997. "Creating Competitive Markets for Ancillary Services." *Oak Ridge National Laboratory ORNL/CON-448.*

Hirst, Eric, and Brendan Kirby. 1998. "Simulating the Operation of Markets for Bulk-Power Ancillary Services." *The Energy Journal*, 19(3): 49–68.

Hledik, Ryan, Roger Lueken, Colin McIntyre, and Heidi Bishop. 2017. "Stacked Benefits: Comprehensively Valuing Battery Storage in California." *Brattle Group.*

Holladay, J. Scott, and Jacob LaRiviere. 2018. "How Does Welfare from Load Shifting Electricity Policy Vary with Market Prices? Evidence from Bulk Storage and Electricity Generation." *The Energy Journal*, 39(6).

Holland, Stephen P., and Erin T. Mansur. 2008. "Is Real-Time Pricing Green? The Environmental Impacts of Electricity Demand Variance." *The Review of Economics and Statistics*, 90(3): 550–561.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016*a*. "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors." *American Economic Review*, 106(12): 3700–3729.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016*b*. "Distributional Effects of Air Pollution from Electric Vehicle Adoption." *Journal of the Association of Environmental and Resource Economists*, 6(S1): S65–S94.

Hortacsu, Ali, Fernando Luco, Steven L. Puller, and Dongni Zhu. 2019. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets." *American Economic Review*, 109(12): 4302–4342.

Hummon, Marissa, Paul Denholm, Jennie Jorenson, David Palchak, Brendan Kirby, and Ookie Ma. 2013. "Fundamental Drivers of the Cost and Price of Operating Reserves." *National Renewable Energy Laboratory Technical Report NREL/TP-6A20-58491.*

Interagency Working Group on Social Cost of Greenhouse Gases, United States Government. 2016. "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866."

International Renewable Energy Agency. 2017. "Electricity Storage and Renewables."

Ito, Koichiro. 2014. "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing." *American Economic Review*, 104(2): 537–63.

Jha, Akshaya, and Frank Wolak. 2020. "Can Forward Commodity Markets Improve Spot Market Performance? Evidence from Wholesale Electricity." *Working Paper.*

Jha, Akshaya, and Gordon Leslie. 2021. "Start-up Costs and Market Power: Lessons from the Renewable Energy Transition." *Working Paper.*

Jin, Sang-Hyeon. 2007. "The effectiveness of energy efficiency improvement in a developing country: Rebound effect of residential electricity use in South Korea." *Energy Policy*, 35(11): 5622–5629.

Johnson, Justin P, and David P Myatt. 2006. "Multiproduct Cournot Oligopoly." *RAND Journal of Economics*, 37(3).

Just, Sebastian, and Christoph Weber. 2008. "Pricing of reserves: Valuing system reserve capacity against spot prices in electricity markets." *Energy Economics*, 30: 3198–3221.

Kamerschen, David, and David Porter. 2004. "The demand for residential, industrial and total electricity, 1973–1998." *Energy Economics*, 26(1): 87–100.

Karaduman, Ömer. 2021. "Economics of Grid-Scale Energy Storage in Wholesale Electricity Markets." *Working Paper.*

Kirby, Brendan J. 2004. "Frequency Regulation Basics and Trends." *Oak Ridge National Laboratory ORNL/TM-2004/291.*

Kirkpatrick, Justin. 2018. "Estimating Congestion Benefits of Batteries for Unobserved Networks: A Machine Learning Approach." *Working paper.*

Kirschen, Daniel, and Goran Strbac. 2004. *Fundamentals of Power System Economics.* John Wiley & Sons, Ltd.

Knittel, Christopher R., and Konstantinos Metaxoglou. 2008. "Diagnosing Unilateral Market Power in Electricity Reserves Market." *The Journal of Energy Markets.*

Leslie, Gordon. 2018. "Tax induced emissions? Estimating short-run emission impacts from carbon taxation under different market structures." *Journal of Public Economics*, 167: 220–239.

Linn, Joshua, and Jhih-Shyang Shih. 2019. "Do lower electricity storage costs reduce greenhouse gas emissions?" *Journal of Environmental Economics and Management*, 96: 130–158.

Lin, Yashen, Jeremiah X. Johnson, and Johanna L. Mathieu. 2019. "Emissions Impacts of Using Energy Storage for Power System Reserves." *Applied Energy*, 168: 444–456.

Liski, Matti, and Iivo Vehvilainen. 2020. "Gone with the Wind? An Empirical Analysis of the Equilibrium Impact of Renewable Energy." *Journal of the Association of Environmental and Resource Economists*, 7(5): 873–900.

Makarov, Y. V., S. Lu, J. Ma, and T. B. Nguyen. 2008. "Assessing the Value of Regulation Resources Based on Their Time Response Characteristics." *Pacific Northwest National Laboratory, Prepared for California Energy Commission.*

Mansur, Erin T. 2008. "Measuring Welfare in Restructured Electricity Markets." *The Review of Economics and Statistics*, 90(2): 369–386.

Mansur, Erin T, and Matthew W White. 2012. "Market Organization and Efficiency in Electricity Markets." *Working paper.*

Mayo, John W. 1984. "Multiproduct Monopoly, Regulation, and Firm Costs." *Southern Economic Journal*, 51(1): 208–218.

MIT Energy Initiative. 2011. "Managing Large-Scale Penetration of Intermittent Renewables."

Monitoring Analytics. 2018. "State of the Market Report for PJM."

Moore, Frances C., Uris Baldos, Thomas Hertel, and Delavane Diaz. 2017. "New Science of Climate Change Impacts on Agriculture Implies Higher Social Cost of Carbon." *Nature Communications*, 8(1): 1607.

Ovaere, Marten, and Kenneth Gillingham. 2019. "The Heterogeneous Value of Solar and Wind Energy: Empirical Evidence from the United States and Europe." *Working paper.*

PG&E. 2020. "Pacific Gas and Electric Company 2020 General Rate Case II: Baseline Territory Study."

Pindyck, Robert S. 2017. "Coase Lecture–Taxes, Targets and the Social Cost of Carbon." *Economics*, 84(335): 345–364.

Reguant, Mar. 2014. "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets." *Review of Economic Studies*, 81: 1708–1742.

Reiss, Peter, and Matthew White. 2005. "Household electricity demand, revisited." *The Review of Economic Studies*, 72(3): 853–883.

Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni. 2018. "Country-Level Social Cost of Carbon." *Nature Climate Change.*

Ryan, Nicole A., Yashen Lin, Noah Mitchell-Ward, Johanna L. Mathieu, and Jeremiah X. Johnson. 2018. "Use-Phase Drives Lithium-Ion Battery Life Cycle Environmental Impacts When Used for Frequency Regulation." *Environ. Sci. Technol.*, 52: 10163–10174.

Saez, Emmanuel, Joel Slemrod, and Seth Giertz. 2012. "The elasticity of taxable income with respect to marginal tax rates: A critical review." *Journal of Economic Literature*, 50(1): 3–50.

Sallee, Jim. 2019. "Pigou creates losers: On the implausibility of achieving pareto improvements from efficiency-enhancing policies." *National Bureau of Economic Research*, , (w25831).

Schulte, Isabella, and Peter Heindl. 2017. "Price and income elasticities of residential energy demand in Germany." *Energy Policy*, 102: 512–528.

Schwenen, Sebastian. 2015. "Strategic bidding in multi-unit auctions with capacity constrained bidders: the New York capacity market." *RAND Journal of Economics*, 46(4): 730–750.

Shaffer, Blake. 2020. "Misunderstanding nonlinear prices: Evidence from a natural experiment on residential electricity demand." *American Economic Journal: Economic Policy*, 12(3): 433–61.

Tabari, Mokhtar, and Blake Shaffer. 2020. "Paying for performance: The role of policy in energy storage deployment." *Energy Economics*, 92.

Tacka, Natalie. 2016. "RTO/ISO Regulation Market Comparison." *Regulation Market Issues Senior Task Force, April 13.*

Wolak, Frank A. 2007. "Quantifying the Supply-Side Benefits from Forward Contracting in Wholesale Electricity Markets." *Journal of Applied Econometrics*, 22: 1179–1209.

Wolak, Frank A. 2011. "Measuring the Benefits of Greater Saptial Granularity in Short-Term Pricing in Wholesale Electricity Markets." *American Economic Review: Papers and Proceedings*, 101(3): 247–252.

Yu, Nanpeng, and Brandon Foggo. 2017. "Stochastic valuation of energy storage in wholesale power markets." *Energy Economics*, 64: 177–185.

Zhou, Zhi, Todd Levin, and Guenter Conzelmann. 2016. "Survey of U.S. Ancillary Services Markets." *Argonne National Laboratory ANL/ESD-16/1.*

Zhu, Xing, Lanlan Li, Kaile Zhou, Xiaoling Zhang, and Shanlin Yang. 2018. "A meta-analysis on the price elasticity and income elasticity of residential electricity demand." *Journal of Cleaner Production*, 201: 169–177.

Chapter 4

Appendix

4.1 Appendix A: How do residential electricity consumers respond to price? Dynamics and heterogeneity

Marginal price regressions

This section of the Appendix shows marginal price elasticity estimates (in contrast with the average price elasticity estimates shown in the main body of this paper). While short-run marginal price elasticity estimates are shown in the main body, this Appendix section shows medium- and long-run marginal price elasticity estimates.

Heterogeneity by Federal Poverty Line

This section of the Appendix shows heterogeneity estimates, using CBG-level decile of percentage of households below the Federal Poverty Line (FPL) as a proxy for income. I estimate short- and long-run elasticities by FPL decile.

Table 4.1: Dynamic medium-run price elasticities - all marginal prices (4 year stable sample)

	kWh
	(1)
$\Delta\ln(MP\_it)$	-0.11***
	(0.023)
$\Delta\ln(MP\_i,t,1)$	-0.13***
	(0.021)
$\Delta\ln(MP\_i,t,2)$	-0.085***
	(0.014)
$\Delta\ln(MP\_i,t,3)$	-0.072***
	(0.015)
Observations	2626159
F	54.3

*Note: Fixed effects include CBG-by-month and consumption deciles from the period twelve months prior to the initial price period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Figure 4.1: Short-run estimates by Federal Poverty Level Decile

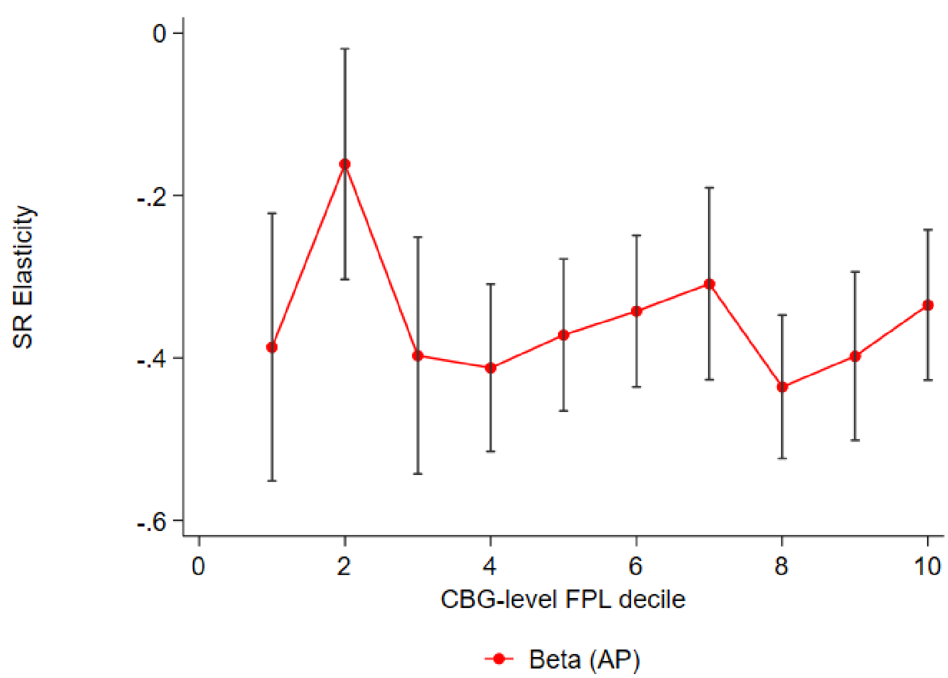

Figure 4.2: Short-run estimates by Federal Poverty Level Decile

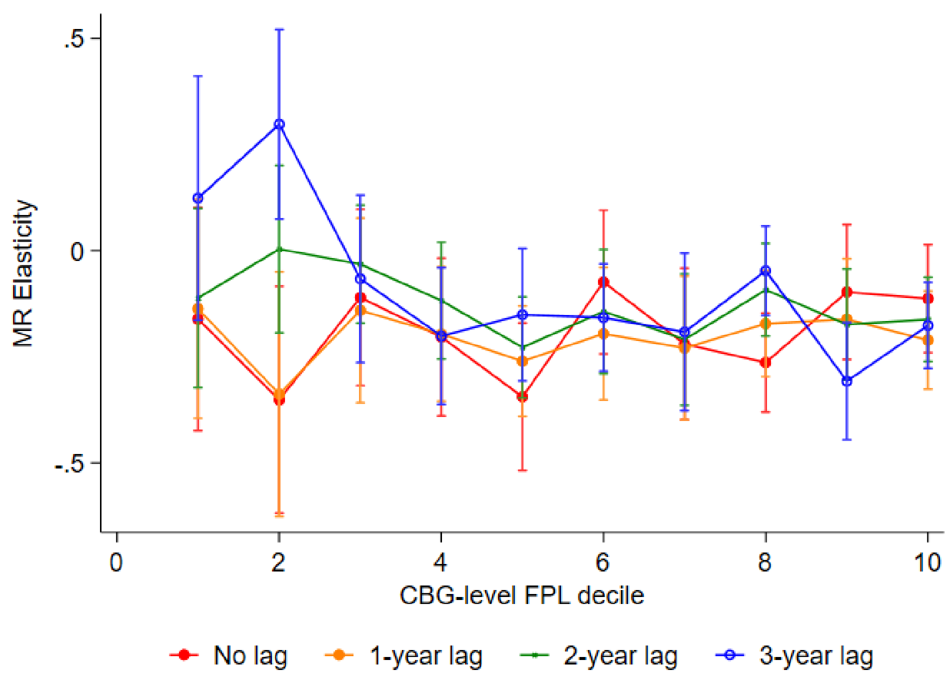

Table 4.2: Dynamic medium-run price elasticities - all average prices (8 year stable sample)

	12
	(1)
$\Delta\ln(AP\_it)$	-0.072
	(0.059)
$\Delta\ln(AP\_i, t, 1)$	-0.14**
	(0.058)
$\Delta\ln(AP\_i, t, 2)$	-0.030
	(0.053)
$\Delta\ln(AP\_i, t, 3)$	0.048
	(0.056)
$\Delta\ln(AP\_i, t, 4)$	0.027
	(0.068)
$\Delta\ln(AP\_i, t, 5)$	-0.0088
	(0.056)
$\Delta\ln(AP\_i, t, 6)$	-0.058
	(0.061)
$\Delta\ln(AP\_i, t, 7)$	-0.21**
	(0.088)
Observations	955837
F	62.6

*Note: Fixed effects include CBG-by-month and consumption deciles from the period twelve months prior to the initial price period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Definition of baseline territories

Pacific Gas and Electric Company®
U 39 San Francisco, California

	Revised	Cal. P.U.C. Sheet No.	34601-E
Cancelling	Revised	Cal. P.U.C. Sheet No.	12081-E

ELECTRIC PRELIMINARY STATEMENT PART A Sheet 1
DESCRIPTION OF SERVICE AREA & GENERAL REQUIREMENTS

A. DESCRIPTION OF SERVICE AREA AND GENERAL REQUIREMENTS

1. TERRITORY SERVED BY PG&E

a. The Pacific Gas and Electric Company (PG&E) supplies electric service in all or portions of 47 counties in the northern and central part of the State of California. A map of counties and associated zip codes that PG&E provides service to can be found on PG&E's website at http://www.pge.com/tariffs/ under electric maps. (N) I (N)

b. The baseline territories used in the residential rate schedules are shown below for each county. Baseline territories correspond with elevation lines, unless specific boundaries were drawn to avoid dividing communities or neighborhoods as described in Section A.1.c. (T) (T)

County	Locations, Elevation Range or Description at c. Below	Baseline Territory Code
ALAMEDA	c.(1)(S)	S
	c.(1)(T)	T
	All Other	X
ALPINE*	All	Z
AMADOR	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-6,000'	Y
	Over 6,000'	Z
BUTTE	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-4,800'	Y
	Over 4,800'	Z
CALAVERAS	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-6,000'	Y
	Over 6,000'	Z
COLUSA	All	S
CONTRA COSTA	c.(2)(S)	S
	c.(2)(T)	T
	All Other	X
EL DORADO*	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-6,000'	Y
	Over 6,000'	Z
FRESNO*	Under 3,500'	R
	3,501'-6,500'	Y
	Over 6,500'	Z
GLENN	Under 3,000'	S
	Over 3,000'	Y
HUMBOLDT	c.(3)(V)	V
	All Other	Y
KERN*	Under 1,000'	W
	Over 1,000'	R
KINGS*	All	W
LAKE*	All	P
LASSEN*	Under 4,800'	Y
	Over 4,800'	Z

*Pertains to PG&E electric service area only.

(Continued)

Advice	4535-E-A	Issued by	Date Filed	December 15, 2014
Decision		**Steven Malnight**	Effective	December 17, 2014
		Senior Vice President	Resolution	
		Regulatory Affairs		

ELECTRIC PRELIMINARY STATEMENT PART A Sheet 2
DESCRIPTION OF SERVICE AREA & GENERAL REQUIREMENTS

A. DESCRIPTION OF SERVICE AREA AND GENERAL REQUIREMENTS (Cont'd.)

1. TERRITORY SERVED BY PG&E (Cont'd.)

County	Locations, Elevation Range or Description at c. Below	Baseline Territory Code	
MADERA*	Under 4,000'	R	
	4,001'-6,500'	Y	
	Over 6,500'	Z	
MARIN	c.(4)(T)	T	
	All Other	X	
MARIPOSA	Under 3,500'	R	
	3,501'-6,000'	Y	
	Over 6,000'	Z	
MENDOCINO	c.(5)(T)	T	
	All Other	X	
MERCED	All	R	
MONTEREY	c.(6)(T)	T	
	All Other	X	
NAPA	All	X	
NEVADA	Under 1,500'	S	
	1,500'-3,000'	P	
	3,001'-5,500'	Y	
	Over 5,500'	Z	
PLACER*	Under 1,500'	S	
	1,500'-3,000'	P	
	3,001'-5,500'	Y	
	Over 5,500'	Z	
PLUMAS*	Under 4,800'	Y	
	Over 4,800'	Z	
SACRAMENTO	All	S	
SAN BENITO	c.(7)(T)	T	
	All Other	X	
SAN FRANCISCO	All	T	
SAN JOAQUIN	All	S	
SAN LUIS OBISPO	c.(8)(R)	R	
	c.(8)(T)	T	
	All Other	X	
SAN MATEO	c.(9)(T)	T	
	c.(9)(Q)	Q	
	All Other	X	
SANTA BARBARA*	c.(10)(R)	R	
	c.(10)(T)	T	
	All Other	X	
SANTA CLARA	c.(11)(Q)	Q	
	All Other	X	
SANTA CRUZ	Under 1,500'	T	
	1,500' & Over	Q**	(T)

* Pertains to PG&E electric service area only.
** Territory Q also includes customers in the following locations (zip codes) within Santa Cruz County at elevations less than 1,500 feet: Ben Lomond (95005), Boulder Creek (95006), Brookdale (95007), Felton (95018), Mount Hermon (95041) and unincorporated areas (95033). (N)

(Continued)

ELECTRIC PRELIMINARY STATEMENT PART A Sheet 3
DESCRIPTION OF SERVICE AREA & GENERAL REQUIREMENTS

A. DESCRIPTION OF SERVICE AREA AND GENERAL REQUIREMENTS (Cont'd.) (T)

1. TERRITORY SERVED BY PG&E (Cont'd.) (T)

County	Locations, Elevation Range or Description at c. Below	Baseline Territory Code	
SHASTA	Under 2,000'	R	(L)
	2,001'-4,500'	Y	|
	Over 4,500'	Z	|
SIERRA	Under 5,500'	Y	|
	Over 5,500'	Z	|
SISKIYOU*	Under 4,500'	Y	|
	Over 4,500'	Z	|
SOLANO	c.(12)(X)	X	|
	All Other	S	|
SONOMA	c.(13)(T)	T	|
	All Other	X	|
STANISLAUS	All	S	|
SUTTER	All	S	|
TEHAMA	Under 2,500'	R	|
	2,501'-4,800'	Y	|
	Over 4,800'	Z	|
TRINITY	Under 2,000'	X	|
	2,001'-4,500'	Y	|
	Over 4,500'	Z	|
TULARE*	Under 1,000'	W	|
	1,001'-3,500'	R	|
	3,501'-6,500'	Y	|
	Over 6,500'	Z	|
TUOLUMNE*	Under 1,500'	S	|
	1,500'-3,500'	P	|
	3,501'-6,000'	Y	|
	Over 6,000'	Z	|
YOLO	All	S	|
YUBA	Under 1,500'	S	|
	1,500' & Over	P	(L)

* Pertains to PG&E electric service area only. (D)

(Continued)

4.2 Appendix B: Long-run price elasticities and mechanisms: Empirical evidence from residential electricity consumers

Regression discontinuity

While my preferred specifications rely on a two-way fixed effects approach, an alternative approach would use regression discontinuity (RD) methods. Before discussing why I prefer the approach presented in the main body of the paper, it's important to note that any RD in this research setting would need to be a fuzzy RD. When PG&E defined baseline territory boundaries in 1982, there was substantial measurement error in their observations of elevation. This measurement error has resulted in incorrect baseline territory assignment for 17% of households in the final sample, where the observed household elevation does not match the assigned baseline territory. While this measurement error excludes a sharp RD as a possible method, it is still possible to use a fuzzy RD, where the primary independent variable is a binary variable indicating whether a household is on the "high price" side of the border, and I instrument for this variable with observed elevation:

$$\text{First stage: } hi_i = \alpha_0 + \alpha_1 elevation_i + \gamma_{ct} + \epsilon \tag{4.1}$$

$$\text{Second stage: } ln(c_{it}) = \beta_0 + \beta_1 \widehat{hi_i} + \gamma_{ct} + \epsilon \tag{4.2}$$

This specification gives the consumption impact of being assigned to a higher-price baseline territory. I run a similar regression with $ln(AP)$ as the second-stage dependent variable to estimate the average price impact of baseline territory assignment. Dividing the consumption coefficient by the price coefficient reveals a point estimate of long-run elasticity. Results of these regressions are shown in Tables 4.3 and 4.4, for households within 300 feet of elevation from the border discontinuity and 200 feet of elevation from the border discontinuity.

Table 4.3: Fuzzy Regression Discontinuity - Log Consumption

	300 (1)	200 (2)
hi	-0.10*** (0.013)	-0.079*** (0.016)
Observations	4324573	3048760
F		

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Table 4.4: Fuzzy Regression Discontinuity - Log Average Price

	300 (1)	200 (2)
hi	0.0040 (0.0055)	0.016** (0.0070)
Observations	4325205	3049265
F		

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

These results imply a point elasticity estimate of approximately -1, which is substantially smaller than the two-way fixed effects approach in the main body of text. However, I believe that the two-way fixed effects approach presented in the main body of the paper is a more sound approach. The primary reason for this preference is the heterogeneity in baseline variation across different borders. Across some borders, the difference in baselines is almost negligible, while in others, the difference in baselines is immense. The two-way fixed effects specification accommodates using the length of the baseline itself as an instrument, which allows the variation in baselines to be used for identification. A fuzzy RD approach requires a more blunt instrument – a binary variable indicating whether a consumer is in the "high" or "low" price regime. This approach throws away a large amount of this baseline variation.

A similar approach can be used to test specific mechanisms driving these price responses.

Table 4.5: Fuzzy Regression Discontinuity - Solar

	300 (1)	200 (2)
hi	0.018*** (0.0020)	0.019*** (0.0024)
Observations	133163	93897
F		

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Instead of using panel data, I collapse the data to a customer level to be purely cross-sectional. I then use a similar fuzzy regression discontinuity approach:

$$\text{First stage: } hi_i = \alpha_0 + \alpha_1 elevation_i + \gamma_c + \epsilon \tag{4.3}$$

$$\text{Second stage: } Solar_i = \beta_0 + \beta_1 \widehat{hi_i} + \gamma_c + \epsilon \tag{4.4}$$

where $Solar_i$ denotes an indicator for whether a household ever adopts solar. Results of this regression are shown in Table 4.5 for households within 300 feet of elevation from the border discontinuity and 200 feet of elevation from the border discontinuity.

Similar to the main text, I find that households on the high price side of the border are more likely to adopt solar, by 2-3 percentage points.

Alternative price definitions

Here, I explore how long-run elasticity estimates vary under different definitions of price. In the main body specifications, one might be concerned about endogeneity, where the adoption of an energy-saving technology might cause a dramatic change in usage and push a household into a lower pricing tier. I test this concern with two alternative definitions of price: first, I calculate what monthly average price would have been under monthly consumption levels from 2008 (the first year of data in my sample) and under the present-period price schedule[1]. Second, to confirm that prices aren't dependent on that single year of data, I repeat the same exercise with 2009 consumption levels. As shown in Tables 4.6 and 4.7, I find similar estimates, demonstrating that this potential endogeneity is not driving the observed results.

[1]Specifically, I match according to the month. For instance, average price in February 2011 would be determined from consumption levels in February 2008 and the price schedule in February 2011.

Table 4.6: Price under 2008 consumption

	All (1)
Logged average price (premise 2008)	-2.17*** (0.29)
Observations	9199687
F	357.1

*Note: Fixed effects include CBG-by-month and a binary variable indicating electric heat. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Table 4.7: Price under 2009 consumption

	All (1)
Logged average price (premise 2009)	-2.10*** (0.28)
Observations	9010362
F	378.2

*Note: Fixed effects include CBG-by-month and a binary variable indicating electric heat. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Marginal price regressions

This section of the Appendix shows marginal price elasticity estimates (in contrast with the average price elasticity estimates shown in the main body of this paper).

Table 4.8: Long-run IV estimate of elasticity (marginal price)

	All
	(1)
Logged marginal price	-1.19***
	(0.17)
Observations	9331612
F	332.7

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Alternative instrument

This section of the Appendix shows long-run elasticity estimates using an alternative instrument – a binary variable indicating whether a household is in the high price portion of a Census Block Group. This instrument is identical to the instrument used in the mechanism regressions in Section 5 of the main text.

Table 4.9: Long-run IV estimate of elasticity (marginal price)

	All (1)
Logged average price	-5.67*** (2.08)
Observations	9331612
F	24.7

*Note: Fixed effects include CBG-by-month. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Heterogeneity by Federal Poverty Line

This section of the Appendix shows heterogeneity estimates, using CBG-level decile of percentage of households below the Federal Poverty Line (FPL) as a proxy for income. I estimate short- and long-run elasticities by FPL decile.

Figure 4.3: Long-run estimates by Federal Poverty Level Decile

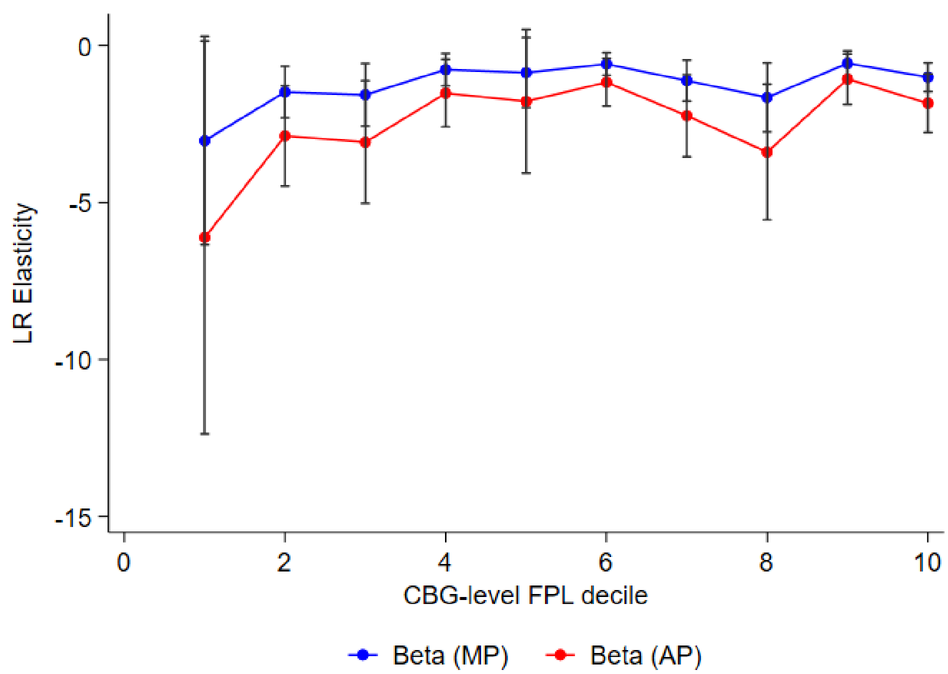

4.3 Appendix C: Spillovers from Ancillary Services to Wholesale Power Markets

This Appendix provides additional details on the stylized four-unit model, on the data used, and on the regression coefficients. It also provides additional robustness checks for the regression results.

Table 4.10: Four-Unit Model

	Generation in Equilibrium				Regulation in Equilibrium			
Regulation Requirement	Unit A, \$35/MWh	Unit B, \$37/MWh	Unit C, \$40/MWh	Unit D, \$60/MWh	Unit A, \$5/MW	Unit B, \$6/MW	Unit C, \$7/MW	Unit D, \$10/MW
500	25000	24550	0	450	0	450	0	50
550	25000	24525	0	475	0	475	0	75
599	25000	24500.5	0	499.5	0	499.5	0	99.5
600	25000	24500	0	500	0	500	0	100
601	24999.5	24500	0	500.5	0.5	500	0	100.5
602	24999	24500	0	501	1	500	0	101
610	24995	24500	0	505	5	500	0	105
615	24992.5	24500	0	507.5	7.5	500	0	107.5
616	24884	21116	4000	0	116	500	0	0
620	24880	21120	4000	0	120	500	0	0
624	24876	21124	4000	0	124	500	0	0
625	24875	21125	4000	0	125	500	0	0
650	24850	21150	4000	0	150	500	0	0
675	24825	21175	4000	0	175	500	0	0
700	24800	21200	4000	0	200	500	0	0
750	24750	21250	4000	0	250	500	0	0
800	24700	21300	4000	0	300	500	0	0
850	24650	21350	4000	0	350	500	0	0
900	24600	21400	4000	0	400	500	0	0
950	24550	21450	4000	0	450	500	0	0
999	24501	21499	4000	0	499	500	0	0
1000	24500	21500	4000	0	500	500	0	0
1001	24500	21499	4001	0	500	500	1	0
1050	24500	21450	4050	0	500	500	50	0
1100	24500	21400	4100	0	500	500	100	0
1150	24500	21350	4150	0	500	500	150	0
1200	24500	21300	4200	0	500	500	200	0

Note: Table 4.10 lists the equilibrium results for a four-unit model with energy and regulation output. Units A, B, and C face a maximum capacity of 25,000 MW each. They also face a minimum constraint, when generating, of 4,000 MW. Unit D faces a maximum capacity of 25,000 MW and a minimum when generating of 400 MW. (This lower minimum operational constraint is meant to represent the fact that the peaking portion of the electricity market is made up of many small peaker units that can each be dispatched at quite small levels of generation.) Marginal costs of energy and regulation provision are listed in the table. For both services, Unit A is cheapest, Unit B next cheapest, etc. Energy demand is held constant at 50,000 MWh, while the regulation requirement varies exogenously across rows. The equilibrium is found using the online tool `https://online-optimizer.appspot.com/`. We check whether results are global, not just local, solutions by forcing individual units on or off, finding that alternative solutions do not achieve a lower system-wide cost. We further explore whether the solutions are unique (as opposed to, e.g., having a flat objective function) by imposing additional constraints forcing an individual unit's generation or regulation to be $\varepsilon = 0.001$ higher than the optimal solution in an effort to find other equal-cost solutions – however, for all cases we explored, doing so yields a higher total system cost (or no feasible solution) indicating that the reported solutions are likely unique.

Data Appendix

Data Sources for Control Variables

From PJM, we observe total hourly electricity demand, in MWh. We also collect PJM data on the forecasted peak and valley demand for each day.

From the Energy Information Administration , we observe the daily price of natural gas (measured at Henry Hub), the daily price of oil (West Texas Intermediate), and the monthly price of coal paid by power plants. In regressions with month-of-sample effects, the coal price drops out. Also, we observe daily average temperature at the Philadelphia airport (degrees Fahrenheit, from NOAA), a relatively central location within PJM, which we use to calculate cooling and heating degree days. In some regressions we add the daily average temperature in Chicago, also from NOAA.

For some robustness checks, we include the hourly requirement (in MW) for other types of ancillary services: synchronized and non-synchronized reserves. PJM sets requirements both for the territory as a whole and for the Mid-Atlantic Dominion area; we control for both sets of requirements. None of these variables were directly tied to policy changes on December 1, 2013, and they are not generally correlated with the frequency regulation requirement (the correlation coefficient between each of these variables and the regulation requirement is less than 0.1), so these controls are not expected to be necessary for identification.

In one placebo specification, we use hourly data on wind generation (MWh, from PJM). For other placebo tests, we collect hourly generation and CO_2 emissions data at the generator level for two other types of units: CEMS fossil-fuel-fired units not in PJM but in nearby states; and units in CEMS data that that are not categorized by CEMS as electrical generating units (e.g., refineries).

CO2 Data

The CEMS-reported CO_2 emissions are missing for approximately 9% of observations with non-zero heat input data, representing 2% of generation. In place of these missing values, we assume an emissions rate (per mmBtu of fuel used) equal to the median rate at the unit, typically around 0.093 metric tons per mmBtu for coal-fired units and 0.054 for natural gas fired units. Below, we show alternative results using CEMS-reported CO_2 emissions.

Gross to Net Conversion

As described in the main text, we must re-scale the CEMS-reported hourly generation to account for both in-house load and incomplete reporting of combined cycle units. Specifically, we do as follows.

In the EIA-923 dataset, we observe annual generation by plant. While EIA-923 reports monthly generation, it is imputed for some units. Thus we focus on the annual generation variable, which is not imputed. EIA-923 reports generation at a somewhat finer scale: prime mover by fuel type within a plant (e.g., aggregating across all coal boilers within a plant).

However, we are most confident in the matching at the plant level as opposed to the prime mover by fuel type level, since there may be some differences in the reporting of technology between EIA and CEMS.

We merge annual CEMS data with annual EIA data at the plant level. For each plant-year, we calculate the ratio of net to gross generation. At plant-year combinations with small generation quantities, this may lead to outliers, so we take the median across years for each plant. We also winsorize the upper and lower 2% to deal with outliers – the 2nd percentile is 0.4 and the 98th percentile is 2.3. Across all electrical generating units in PJM, the median is 0.95, fairly consistent with (Cicala, 2022). The median for boilers is 0.92. The median for combustion turbines is 0.98. The distribution for combined cycles is bimodal, with one mass at around 0.97 (consistent with reporting both cycles) and one mass at around 1.5 (consistent with reporting only one cycle).

This approach also solves a problem we see with some combined cycle units: they do not report the full value of their electrical output in the CEMS data. Finally, some units report only steam load in CEMS, but report non-zero net generation in EIA-923. We similarly scale from steam load to net generation for these units; they account for 4% of our final net generation variable.

Minimum Constraints

First, we estimate the minimum constraint for each generator, using EIA-860 data on minimum operational constraints. We observe reported minimum operational constraints for the years 2013-2014; they are not reported in the 2012 EIA-860. Unfortunately, a comprehensive merge between EIA-860 and CEMS at the unit level does not exist. However, merging at the plant level, or even at the plant by prime mover by fuel type level, is straightforward. Accordingly, we bring in minimum operational data as follows. For around half of generator-year combinations at electrical generating units in PJM, the minimum operational constraint is the same across all units within a plant (when expressed as a percentage of maximum capacity), so merging at the plant level is appropriate. For the remaining generator-year combinations, we use the median operational constraint within the plant at the prime mover by fuel type level. Some units (representing 3% of generation) do not appear in the minimum constraints data in EIA-860, and for these units we use the median constraint by prime mover and fuel type across all PJM plants.

Example plots of hourly capacity factors show that these minimum constraints are visible in hourly data (Figure 4.6). Here we show nine histograms – one unit at three large plants for each of our three main technology types. A vertical black line depicts the minimum operational load in EIA-860 data. For most of these units, the vertical line is close to a discontinuity in the hourly histogram.

However, in a robustness check, we construct an alternative minimum operational load using the unit-level observed behavior, as follows. We calculate the portion of hours a plant is generating at a capacity factor of 0, a capacity factor between 0 and 10 percent, between 10 and 20 percent, etc. We then use as the minimum operational load whatever is the smallest bin in which at least 5 percent of non-zero generating hours fall. This is a proxy for the discontinuities observed visually in the histograms. We generally calculate minimum operational loads of around 40 to 60 percent for the boilers and CC plants, although we also observe units with a very small minimum constraint (0-10% of capacity), especially for the CT units. (Regression results using this alternative minimum constraint measure are shown in Table 4.24.)

Once we have a measure of minimum constraints for each unit, we proceed as follows. We calculate the capacity factor of each unit in each hour, defined as net generation divided by maximum observed generation. We then place each unit-hour observation into one of five bins: *Off* (capacity factor of zero), *Below Minimum Constraint* (capacity factor between 0% and less than 5% of the minimum constraint to maximum capacity ratio), *At Minimum Constraint* (capacity factor within 5% of the minimum to maximum ratio), *Between Minimum Constraint and Maximum Capacity*, and *At Maximum Capacity* (capacity factor between 95% and 100%).

Fuel Types and Unit Types

From CEMS, we observe fuel types and unit types. The raw CEMS data lists 37 unique primary fuel types. The most common are coal, pipeline natural gas, and diesel oil. Less common categories include, e.g., "residual oil" "process gas," "wood," etc., as well as combinations of these fuels, e.g., "coal, natural gas." We generate four categories: "coal" (which aggregates across coal as well as a small number of units using "coal refuse" or "petroleum coke"), "pipeline gas" + "natural gas," "oil" (diesel, residual, or other oil), and "other," where "other" aggregates across, e.g., wood, units listing combinations of fuels, and units for which we do not have a fuel type.

The raw CEMS data similarly lists 22 different technology types, with the most common being "combustion turbine", "dry bottom wall-fired boiler," and "combined cycle." We generate four categories: "boiler" (an aggregation of all boilers, stokers, and tangentially-fired units), "combined cycle," "combustion turbine," and "other." The latter includes a small number of other technology types, a small number with unreported technology type, and some units that changed technology over this 2012-2014 sample.

For our 2012-2014 sample, total gross generation by category is shown in Table 4.11.

CEMS Versus EIA Generation Data

In addition to the net-versus-gross distinction described above, the CEMS and EIA data differ in their coverage across plants. EIA data include hydro, nuclear, solar, wind units, etc. EIA data also include small coal, gas, and oil units not in CEMS. Total generation by fuel type can be compared in Tables 4.11 and 4.12. The difference between CEMS and EIA data is accounted for by the "residual" generation variable we construct, equal to the difference between total demand reported by PJM and total generation reported in CEMS.

Table 4.11: Total Annual Generation by Unit Type, CEMS Data, 2012-2014

Unit Type	Generation, TWh
Coal, Boiler	347
NG, CC	127
NG, CT	8
Switch	4
Oil, Boiler	3
Oil, CC	2
NG, Boiler	<1
Oil, CT	<1
Other, Boiler	<1

Note: This table shows annual generation over 2012-2014 for the aggregations of fuel by technology type that we have used. Data coverage is all CEMS-reporting electrical generating units in PJM. Data source is CEMS for generation, fuel type, technology type; and EIA for electrical generating unit designation and PJM designation.

This residual variable thus captures the behavior of nuclear, etc. units; as well as in-house load and imports and exports between PJM and other ISO/RTOs.

Hour Naming Conventions

PJM data are reported in both Coordinated Universal Time (UTC) and Eastern Prevailing Time (EPT). CEMS data, in contrast, are reported in local, standard time (Central or Eastern, depending on the plant's location). We convert all PJM data to Eastern Standard Time (EST). For CEMS units in Illinois and parts of Indiana, Kentucky, Michigan, and Tennessee, we convert from Central Standard Time (CST) to Eastern Standard Time. Thus all regressions use variables in Eastern Standard Time. Regression results are similar if one uses the raw data, mixing EPT, CST, and EST across variables and plants.

Other

We drop one hour (5 a.m. on April 2, 2013) when the regulation requirement is listed as zero. This represents less than 0.01 percent of our sample (19,693 hours).

Table 4.12: Total Annual Generation by Fuel Type, EIA Data, 2012-2014

Unit Type	PJM Generation, TWh
Coal	336
Nuclear	276
Natural Gas	143
Wind	16
Waste Coal	9
Hydroelectric Conventional	7
Biogenic Municipal Solid Waste and Landfill Gas	5
Distillate Petroleum	2
Other (including nonbiogenic MSW)	1.8
Petroleum Coke	1
Wood and Wood Waste	0.9
Other Gases	0.8
Solar PV and thermal	0.6
Residual Petroleum	0.3
Waste Oil	0.2
Other Renewables	<0.01
Hydroelectric Pumped Storage	-2

Note: This table shows annual generation over 2012-2014. Data coverage is all PJM units in EIA-923 data operating as independent power producers or electric utilities. Data source is EIA for generation, fuel type, sector, and PJM designation.

Simulated Dataset to Illustrate the Identification Strategy

As is shown in Figure 3.4 in the main text, and discussed in Section 3.4, the bulk of our identifying variation comes from a policy change mid-way through our sample, in which the regulation requirement changes from being a function of forecasted peak and valley load (which change daily) to a flat requirement (albeit with separate levels in peak versus off-peak hours). In addition, as a secondary source of variation, we leverage two policy changes that modified the *multiplier* used to convert from forecasted peak and valley load to the regulation requirement

Thus, the second half of our sample, during which the regulation requirement does not vary across days, allows us to identify the effects of control variables (including the forecasted peak and valley load, as well as other things that may be correlated with these forecasts) separately from the effects of the regulation requirement, our variable of interest.

To illustrate how this works, we conduct a simulation in which we directly control the data-generating process. We set our sample size to 20,000, roughly equal to the sample size in Tables 3.2 and 3.3. We construct a peak forecast variable, normally distributed with mean zero and standard deviation equal to one.[2] We then construct a treatment variable, equal to the peak forecast variable in the first half of the sample and equal to zero in the second half, as shown in Figure 4.4. The outcome variable is a function of a constant, the peak forecast,

[2]For simplicity, we use only a peak forecast variable and not separate peak and valley forecasts across hours.

the treatment variable, and random noise (also normally distributed with mean zero and standard deviation equal to one):

$$y_t \equiv 1 + 1 \cdot \text{peak}_t + 1 \cdot \text{treatment}_t + \varepsilon_t.$$

A successful identification strategy will thus recover a coefficient on the peak variable equal to one, and a coefficient on the treatment variable equal to one. Table 4.13 illustrates a series of regressions. In Column 1, the regression is correctly specified, including the entire sample and controlling for the peak forecast variable. As expected, all three coefficients are estimated to be 1.0, with a high degree of precision. Column 2 shows that adding a post dummy does not affect our ability to estimate the treatment and peak effects; this is relevant as some of our specifications include time effects.

Columns 3 through 6 illustrate how identification is achieved by displaying specifications that are *not* identified. In Column 3, only the first half of the sample is included. Thus the effects of the peak forecast and treatment variable cannot be estimated. The software has dropped the coefficient on the treatment variable because of perfect collinearity, and the effects of both variables have been rolled into the coefficient on "Peak," which is now biased upwards.

In Column 4, only the second half of the sample is included. The software has again dropped the coefficient on the treatment variable because of perfect collinearity. The effect of the peak forecast variable can be correctly estimated, but the treatment effect of interest cannot be recovered.

In Column 5, the entire sample is included, but the crucial "Peak" control has been left out by the researcher. Again, the effects of both variables have been rolled into the coefficient on "Peak," which is now biased upwards.

Columns 3 through 5 thus show how having the policy change as well as the peak control variable are the crucial components for identification. The second half of the sample allows the researcher to estimate the "Peak" effect, which can then be controlled for to allow the researcher to estimate the "Treatment" effect. This is comparable to our main specification in Tables 3.2 and 3.3, for which we observe a policy change mid-way through the sample, and where we know (both based on policy documentation and what we observe in the data itself) that the treatment variable is a direct multiplier of the peak variable.

Finally, Column 6 illustrates that identification is not achieved via just a simple pre/post comparison. In this example, the mean level of the treatment variable has not changed from the pre-period to the post-period, and indeed a simple regression on a post-period dummy would not uncover the coefficient of interest.

Figure 4.4: Treatment Variable for Simulation of Identification Strategy

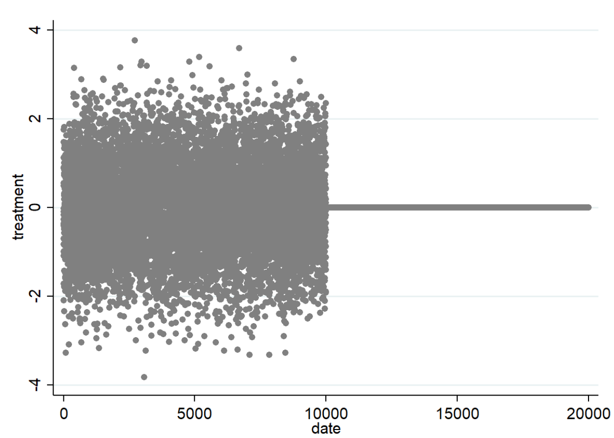

Note: This figure shows the treatment variable constructed for the simulation exercise. It is normally distributed with mean zero and standard deviation equal to one for the first part of the sample, and it is equal to zero for the second part of the sample.

Table 4.13: Simulated Dataset to Illustrate the Identification Strategy

	Correctly specified	Alternative	First half	Second half	Dropping peak control	Only dummy
Treatment	0.99***	0.99***			1.99***	
	(0.01)	(0.01)			(0.01)	
Peak	1.01***	1.01***	1.99***	1.01***		
	(0.01)	(0.01)	(0.01)	(0.01)		
Post dummy		-0.01				-0.01
		(0.01)				(0.03)
Constant	1.01***	1.02***	1.02***	1.01***	1.01***	1.01***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Observations	20,000	20,000	10,000	10,000	20,000	20,000
R^2	0.72	0.72	0.80	0.50	0.57	0.00

Note: This table shows five regressions using a simulated dataset constructed by the researchers. The true data-generating process is $y_t \equiv 1 + 1 \cdot \text{peak}_t + 1 \cdot \text{treatment}_t + \varepsilon_t$, where "peak" and ε are each normally distributed with mean zero and standard deviation equal to one. The "treatment" variable is equal to "peak" in the first half of the sample and equal to zero in the second half. The first column includes the entire sample and both the Treatment variable of interest and the Peak control variable; it correctly uncovers all three coefficients. This column mimics the identification strategy used in the main text for the impact of the regulation requirement on the generation mix. The second column shows that including a post-period dummy does not affect our ability to recover the treatment and peak effects. However, dropping either half of the data or not including the "Peak" control leads to mis-specification, as shown in the third through fifth columns. Estimation using a only a post-period dummy is not possible, as the mean value of the Treatment variable is constant throughout the sample, as shown in the last column. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Robustness Checks

Tables 4.17, 4.18, and 4.19 show additional robustness checks, mentioned briefly in the main text. Here we describe the rationale for examining these alternative specifications.

First, we examine more parsimonious specifications, dropping various control variables. In the main text, we argue that peak and valley forecast controls are necessary, as they directly impact the regulation requirement in some hours and as they may also directly impact generator behavior. We also argue that additional controls, common in the literature, may similarly be correlated with both the regulation requirement and generator behavior (for instance, weather). However, might worry about "oversaturation" in our regressions (for instance, that controlling for so many things leaves only measurement error, as in Fisher et al. 2012). In these parsimonious regression, we continue to control for peak and valley forecasts (and a peak versus valley hour dummy), following the logic of Section 4.3, but we drop all other control variables. These results provide reassurance that the effects we estimate are not driven by the inclusion of too many controls. The estimates from the parsimonious regression are generally similar to the main reported results, albeit with (not surprisingly) much less precision. An exception is the "other technology" and "other fuel" results, for which we estimate somewhat different results. However, this does not change any of our main conclusions about the changes to boilers versus combined cycle units, nor about the changes to coal versus natural gas.

Table 4.19 next shows additional controls; more flexible non-parametric conrols; etc.[3] See table notes for details. Overall, across our robustness checks, we estimate qualitatively similar fuel use shifts (increased generation by natural gas units) and CO_2 reductions.

We also show CO_2 results using reported rather than constructed emissions (Table 4.20). With this variable, we again estimate statistically significant emissions reductions.

Heat Rate Effect of Regulation Provision

We briefly note that our results incorporate an additional effect on CO_2 emissions. When a power plant supplies frequency regulation, its heat rate is impacted – the amount of fuel it must use per unit of electricity sold. This is for two reasons. First, the heat rate at an individual generator depends on its generation level; it is non-linear (and frequently modeled as quadratic). Thus because generators are operating at new set points (the point around which they move in response to the regulation signal), their heat rate could change. Second, the generator must move up and down around its operating set point, rather than holding steady at a given level of output. This will worsen the heat rate, i.e. require greater heat input (and therefore more CO_2 emissions) per unit of electricity sold (Hirst and Kirby, 1997; Hummon et al., 2013).

Our regressions implicitly incorporate these two effects. Because our CO_2 emissions rate is time-varying, our left-hand side variable in Table 3.3, Column 5 will vary as the heat

[3] In previous versions of this paper, we additionally controlled for some things like power plant retirements; these are subsumed in this version by month of sample effects.

rate changes. These two effects do not appear to be the main drivers of our results, given the magnitude of the generation mix changes we observe and how closely our back-of-the-envelope CO_2 calculations line up, in the main text.

References

Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker. 2012. “The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment.” American Economic Review, 102(7): 3749–3760.

Additional Tables and Figures

This section contains additional tables and figures referenced in the text, including summary statistics, robustness checks, etc.

Figure 4.5: Monthly Generation

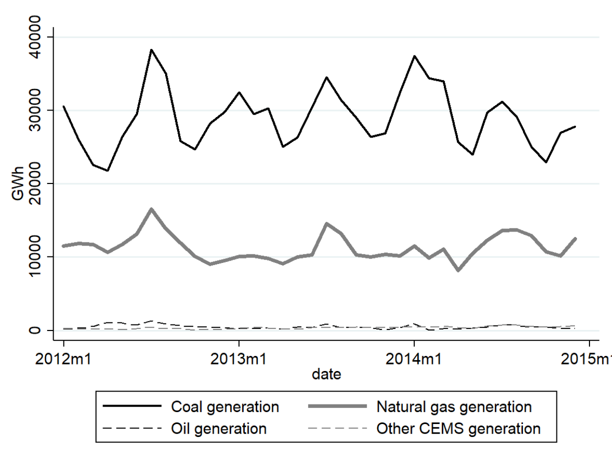

Note: This figure shows monthly generation by unit type for PJM units that appear in CEMS data.

Figure 4.6: Minimum Constraints

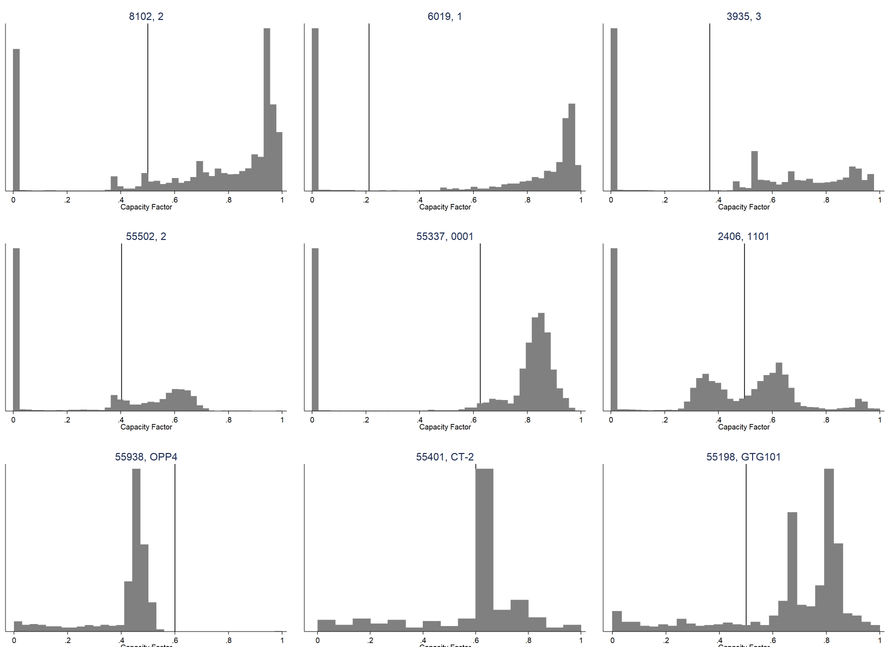

Note: These are nine units at large plants, three for each technology type. The top row shows three coal-fired boilers, with the plant id and unit id given at the top of each histogram. The second row shows natural gas combined cycle plants, and the bottom row shows natural gas combustion turbines. In the bottom row, zeros are not displayed – because CTs operate infrequently, displaying zeros makes it difficult to visualize the non-zero portion of the histogram. Vertical lines are placed at the minimum operating constraint constructed from EIA data (which in some cases is a plant-level proxy, rather than measured at the individual unit level - that may be why some panels appear to show measurement error). See Appendix text for details.

Table 4.14: Summary Statistics

	Mean	Std. Dev.	N
Regulation requirement, 100 MW	6.78	1.06	19693
Generation, by technology:			
Boilers, MWh	40184.7	8233.9	19728
Combined cycle, MWh	14289.3	3635.1	19728
Combustion turbine, MWh	847.9	1923.4	19728
Other unit types in CEMS, MWh	553.4	294.7	19728
Generation, by fuel type:			
Coal generation, MWh	39818.1	7934.4	19728
Natural gas generation, MWh	14906.0	4836.5	19728
Oil generation, MWh	570.1	699.0	19728
Other fuel types in CEMS, MWh	581.0	295.1	19728
CEMS CO2 emissions, tons	45215.2	9378.2	19728
Generation not in CEMS, MWh	34044.3	4568.1	19728
PJM load, MWh	89919.5	15730.8	19728
Peak forecast, in peak hours, MWh	104021.8	15040.0	15618
Valley forecast, in off-peak hours, MWh	74386.2	10460.2	4106
Henry hub natural gas price, dollars per mmbtu	3.97	0.64	19728
WTI oil price, dollars per barrel	94.7	10.3	19728
Coal price, dollars per mmbtu	2.34	0.040	19728
Cooling degree days in Philadelphia	3.15	5.29	19728
Heating degree days in Philadelphia	13.1	13.5	19728

Note: Data cover the period October 1, 2012 through December 31, 2014. Unit of observation is one hour. Data sources: PJM, EPA, and EIA. Peak and valley forecasts apply only in the peak (4 a.m. to midnight) and valley (midnight to 4 am) hours, respectively. A small number of observations (<1%) are missing for the regulation requirement and peak/valley forecast variables.

Table 4.15: Displaying Control Coefficients: The Impact of the Regulation Requirement on the Energy Market

	Boiler	CC	CT	Other
Regulation requirement, 100 MW	-388.86**	357.11**	16.24	15.50
	(189.65)	(162.37)	(208.63)	(22.13)
PJM load, MWh	-0.15***	-0.01	0.16***	-0.00
	(0.02)	(0.02)	(0.02)	(0.00)
CEMS units generation, MWh	0.75***	0.27***	-0.02	0.01
	(0.02)	(0.02)	(0.01)	(0.00)
Peak forecast, in peak hours, MWh	0.02**	-0.01*	-0.01	0.00
	(0.01)	(0.01)	(0.01)	(0.00)
Valley forecast, in off-peak hours, MWh	0.08***	-0.01	-0.07***	0.00
	(0.01)	(0.01)	(0.02)	(0.00)
Henry hub price	556.01	-722.09**	165.99	0.09
	(351.02)	(302.18)	(148.48)	(22.15)
WTI price	5.31	-16.78	9.15	2.31
	(28.85)	(25.06)	(13.81)	(4.93)
Cooling degree days in Philadelphia	21.94	-29.40	7.78	-0.31
	(24.13)	(20.82)	(16.28)	(2.76)
Heating degree days in Philadelphia	20.11**	6.56	-28.23***	1.57
	(9.82)	(10.25)	(8.30)	(1.94)
Observations	19,693	19,693	19,693	19,693
Within R^2	0.91	0.71	0.49	0.13

Note: This table shows coefficients on the control variables for the regression results shown in the main text in Table 3.2. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.16: Displaying Control Coefficients: The Impact of the Regulation Requirement on the Energy Market

	Coal	NG	Oil	Other	CO2
Regulation requirement, 100 MW	-351.04	432.57**	-97.67	16.14	-242.37***
	(234.32)	(192.86)	(68.25)	(22.00)	(89.00)
PJM load, MWh	-0.19***	0.15***	0.04***	-0.00	-0.05***
	(0.02)	(0.02)	(0.01)	(0.00)	(0.01)
CEMS units generation, MWh	0.76***	0.24***	-0.01	0.01	0.78***
	(0.03)	(0.02)	(0.01)	(0.00)	(0.01)
Peak forecast, in peak hours, MWh	0.01	-0.02**	0.01***	0.00	0.03***
	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)
Valley forecast, in off-peak hours, MWh	0.07***	-0.07***	0.00	0.00	0.04***
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)
Henry hub price	561.19	-609.41*	45.82	2.40	580.30***
	(345.19)	(344.25)	(65.92)	(22.02)	(200.33)
WTI price	15.45	-16.05	-1.81	2.41	26.91*
	(29.96)	(28.68)	(6.00)	(4.97)	(15.56)
Cooling degree days in Philadelphia	-7.62	-11.45	19.27***	-0.21	31.40**
	(26.23)	(24.27)	(6.73)	(2.75)	(13.16)
Heating degree days in Philadelphia	28.43**	-24.14**	-5.84*	1.55	5.63
	(10.87)	(10.41)	(3.42)	(1.94)	(6.39)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.88	0.82	0.46	0.13	0.98

Note: This table shows coefficients on the control variables for the regression results shown in the main text in Table 3.3. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.17: Parsimonious Robustness Check: The Impact of the Regulation Requirement on the Energy Market

	Boiler (MWh)	CC (MWh)	CT (MWh)	Other tech. (MWh)
Regulation requirement, 100 MW	-637.8	280.3	101.6	-136.6***
	(410.2)	(204.0)	(239.0)	(25.3)
Observations	19,694	19,694	19,694	19,694
Within R^2	0.66	0.43	0.25	0.27

Note: This table shows a specification similar to that shown in the main text in Table 3.2, but controlling only for the peak and valley load forecasts. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.18: Parsimonious Robustness Check: The Impact of the Regulation Requirement on the Energy Market

	Coal (MWh)	Natural gas (MWh)	Oil (MWh)	Other fuel (MWh)	CO2 (tons)
Regulation requirement, 100 MW	-659.5	403.2	0.3	-136.5***	-460.8
	(442.3)	(275.0)	(84.1)	(25.4)	(333.9)
Observations	19,694	19,694	19,694	19,694	19,694
Within R^2	0.64	0.47	0.30	0.27	0.72

Note: This table shows a specification similar to that shown in the main text in Table 3.3, but controlling only for the peak and valley load forecasts. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.19: Robustness to Alternative Specifications: The Impact of the Regulation Requirement on the Energy Market

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Panel A. Boiler																
Reg. req.	-408**	-381**	-389**	-348*	-425**	-407**	-340	-512***	-389*	-385**	-166	-358*	-403**	-406**	-337*	-376*
	(174)	(163)	(189)	(176)	(185)	(162)	(234)	(182)	(211)	(190)	(162)	(186)	(192)	(193)	(196)	(192)
Panel B. Combined Cycle																
Reg. req.	435**	414**	360**	375***	381**	363***	345*	357**	357**	356**	461***	369**	350**	368**	303*	348**
	(175)	(168)	(162)	(139)	(157)	(136)	(181)	(162)	(160)	(162)	(146)	(162)	(166)	(162)	(157)	(161)
Panel C. Combustion Turbines																
Reg. req.	-37	-61	17	-42	33	27	-20	20	16	13	-308**	-30	35	22	22	17
	(211)	(200)	(208)	(121)	(209)	(130)	(211)	(204)	(221)	(208)	(121)	(196)	(209)	(213)	(216)	(209)
Panel D. Other Units																
Reg. req.	10	29	13	16	11	16	15	18	16	17	13	18	17	16	12*	11*
	(23)	(23)	(21)	(23)	(21)	(21)	(29)	(23)	(21)	(22)	(23)	(22)	(22)	(22)	(7)	(6)
Panel E. Coal																
Reg. req.	-317	-280	-353	-318	-399*	-377*	-318	-467**	-351	-346	-66	-304	-374	-368	-296	-338
	(217)	(199)	(234)	(205)	(228)	(191)	(288)	(224)	(256)	(234)	(192)	(229)	(236)	(239)	(247)	(236)
Panel F. Natural Gas																
Reg. req.	445**	402**	435**	404**	474**	450***	387	429**	433**	427**	226	393**	449**	448**	376*	424**
	(177)	(167)	(192)	(183)	(187)	(167)	(238)	(190)	(211)	(193)	(168)	(189)	(195)	(196)	(193)	(189)
Panel G. Oil																
Reg. req.	-139**	-152**	-96	-102*	-87	-90	-85	-97	-98	-98	-173***	-107	-93	-97	-92	-98
	(69)	(64)	(69)	(54)	(66)	(56)	(81)	(67)	(70)	(68)	(59)	(67)	(68)	(69)	(75)	(68)
Panel H. Other Units																
Reg. req.	11	29	13	16	12	17	16	18	16	17	13	19	18	17	12*	11*
	(23)	(23)	(21)	(23)	(20)	(21)	(29)	(22)	(21)	(22)	(22)	(22)	(22)	(22)	(7)	(6)
Panel I. CO2 Emissions																
Reg. req.	-274***	-239***	-243***	-232**	-247***	-242***	-276**	-365***	-242***	-240***	-207**	-247***	-252***	-247***	-201**	-266**
	(87)	(83)	(89)	(91)	(86)	(86)	(110)	(80)	(93)	(89)	(90)	(89)	(90)	(89)	(95)	(104)

Note: This table shows alternative specifications for the regressions displayed in the main text in Tables 3.2 and 3.3, using various alternative controls, variable definitions, and subsamples. Column 1 controls for the standard deviation of the regulation requirement over the previous 72 hours. Column 2 controls for the standard deviation of each of: the regulation requirement, PJM-wide load, CEMS generation, and the peak and valley forecasted load over the previous 72 hours. Column 3 adds a linear time trend. Column 4 uses a spline with five knots, rather than a linear control, for PJM-wide load. Column 5 uses a spline with three knots, rather than a linear control, for fuel prices. Column 6 adds flexible (binned) controls for PJM-wide load, CEMS generation, and fuel prices. Column 7 collapses to the daily level. Column 8 restricts the sample to units that do not retire. Column 9 calculates Newey-West standard errors with a maximum lag length of 168 hours (one week). Column 10 adds controls for other ancillary services requirements. Column 11 limits the sample to time periods with overlap in the peak/valley forecasts between the pre and post-policy change periods. Column 12 adds weather controls for Chicago. Column 13 uses a constructed regulation requirement variable as an instrument for the reported regulation requirement. Column 14 uses PJM variables in their raw form, i.e., in Eastern Prevailing Time rather than Eastern Standard Time. Column 15 uses CEMS-reported gross generation rather than re-scaled net generation. Column 16 uses only CEMS-reported electrical generation, rather than also incorporating steam load in the net generation scaling. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.20: Alternative CO_2 Measurement: The Impact of the Regulation Requirement on the Energy Market

	Coal	NG	Oil	Other	Total
Panel A. CO_2 Variable in Main Text, Metric Tons					
Regulation requirement, 100 MW	-355*	168	-67	11	-242***
	(214)	(122)	(55)	(8)	(89)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.89	0.79	0.43	0.25	0.98
Panel B. CO_2 as Reported, Metric Tons					
Regulation requirement, 100 MW	-361*	142	-63	8	-274***
	(214)	(122)	(51)	(7)	(92)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.89	0.78	0.44	0.20	0.98
Panel C. Using Unit-Level Emissions Rates, Metric Tons					
Regulation requirement, 100 MW	-355*	197*	-61	6	-213**
	(212)	(118)	(53)	(8)	(87)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.89	0.79	0.44	0.21	0.98

Note: Panel A shows the CO_2 emissions results by fuel type (Columns 1 through 4) and aggregated (Column 5), matching the specifications used in the main text, Table 3.3. Panel B shows analogous specifications, but using CEMS-reported CO_2 emissions (which are occasionally missing) rather than emissions constructed from the heat input variable. Panel C uses the unit-level emissions rate for all hours, not just hours with missing CO_2 data. All panels are reported in metric tons (i.e., in Panel B we convert CEMS-reported short tons into metric tons). ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.21: Placebo and Residual Units: The Impact of the Regulation Requirement on the Energy Market

	Non-PJM Coal	Non-PJM NG	Non-PJM Oil	Non-PJM Other	PJM Comm+Ind	PJM Wind	PJM Residual
Regulation requirement, 100 MW	348.6*	40.2	-22.2	6.9	-3.1	-32.9	-12.1
	(196.8)	(141.8)	(21.3)	(4.4)	(6.7)	(67.1)	(153.0)
Observations	19,693	19,693	19,693	19,693	19,693	19,688	19,693
Within R^2	0.71	0.67	0.12	0.11	0.09	0.17	0.53

Note: This table shows estimates from seven separate regressions, analogous to those presented in the main text, Tables 3.2 and 3.3. The dependent variable in the first four columns is MWh of electricity generated per hour for the electrical generating units that are located in PJM states but are *not* part of PJM; see footnote 40. The dependent variable in the fifth column is MWh of electricity generated by commercial and industrial units in PJM. The dependent variable in the sixth column is MWh of wind generation in PJM. The dependent variable in the seventh column is the difference between PJM-wide demand and the generation reported by electrical generating units in CEMS; this accounts for fuel types not in CEMS (nuclear, wind, solar, etc.), small units not in CEMS, and net imports. The unit of analysis is an hour. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.22: Specifications with Two-Way Fixed Effects

	Boiler	CC	CT	Other tech	Coal	NG	Oil	Other fuel	CO2
Panel A. Time Effects, Location Effects, and Peak/Valley Forecast Controls									
Reg. req., 100 MW, in PJM	-498	335*	97	-123***	-496	389*	27	-110***	-263
	(373)	(191)	(142)	(27)	(377)	(226)	(77)	(27)	(371)
Observations	39,388	39,388	39,388	39,388	39,388	39,388	39,388	39,388	39,388
Within R^2	0.22	0.26	0.13	0.24	0.16	0.32	0.26	0.19	0.33
Panel B. Plus Additional Controls									
	Boiler	CC	CT	Other tech	Coal	NG	Oil	Other fuel	CO2
Reg. req., 100 MW, in PJM	-494***	413***	68	13	-419**	483***	-73	9	-303***
	(149)	(146)	(133)	(22)	(180)	(155)	(61)	(23)	(90)
Observations	39,386	39,386	39,386	39,386	39,386	39,386	39,386	39,386	39,386
Within R^2	0.75	0.40	0.25	0.11	0.71	0.55	0.40	0.07	0.92

Note: Panel A shows specifications in which plants in nearby states (see footnote 40) serve as controls. The unit of observation is an hour in a region (PJM, or nearby states grouped together). The variable of interest takes on the value of the regulation requirement in PJM, and a value of zero in nearby states, as the regulation requirement does not directly affect them. Controls are: hour-of-sample effects, region fixed effects, and two interaction variables. These latter controls are (1) forecasted peak load interacted with a PJM dummy, and forecasted valley load interacted with a PJM dummy; these may be important for avoiding omitted variables bias as the regulation requirement is a direct function of these forecasts in the first half of the sample. Panel B includes the same controls, but also adds all the controls from the main specification in Tables 3.2 and 3.3 interacted with a PJM dummy for additional precision. We also include a control for the total CEMS generation within each region. As shown when moving from Panel A to Panel B, these additional controls aid with precision even in the two-way fixed effects specification. This is intuitive if the response to control variables such as the natural gas price varies between regions, in which case hour of sample effects will not fully account for the natural gas price effect across the two regions. This limitation of the two-way fixed effects specification – it is still aided by additional control variables – combined with the fact that identification is less transparent than it is in the time-series regression, is why our primary specification in the main text is a time-series regression. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.23: Showing Other CEMS Units: The Regulation Requirement and Intensive/Extensive Margins

Panel A. Boilers	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.29	0.20	0.39	1.18	-1.48***
	(1.03)	(0.16)	(0.40)	(0.90)	(0.47)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.71	0.11	0.38	0.43	0.75
Panel B. Combined Cycle Plants	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.33***	0.03	0.40*	2.31***	-0.40
	(0.86)	(0.25)	(0.24)	(0.70)	(0.25)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.22	0.07	0.61	0.32
Panel C. Combustion Turbines	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.13	0.01	0.31	3.06	-1.26**
	(3.38)	(0.46)	(0.33)	(2.44)	(0.59)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.31	0.27	0.45	0.19
Panel D. Other Units	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.03	-0.09	-0.10***	0.14	0.07
	(0.17)	(0.06)	(0.03)	(0.15)	(0.04)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.19	0.03	0.03	0.14	0.04

Note: This table expands on Table 3.4 by showing results at other units. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.24: Alternative Minimum Constraints Data: The Regulation Requirement and Intensive/Extensive Margins

Panel A. Boiler	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.29	0.04	0.48	1.26	-1.48***
	(1.03)	(0.53)	(0.66)	(0.96)	(0.47)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.71	0.36	0.51	0.65	0.75
Panel B. Combined Cycle	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.33***	0.55	0.73**	1.46*	-0.40
	(0.86)	(0.34)	(0.36)	(0.80)	(0.25)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.22	0.15	0.59	0.32
Panel C. Combustion Turbine	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.13	0.01	0.20	3.18	-1.26**
	(3.38)	(0.20)	(0.32)	(2.72)	(0.59)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.20	0.43	0.44	0.19
Panel D. Other Tech	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.03	-0.18***	0.15***	-0.02	0.07
	(0.17)	(0.06)	(0.06)	(0.15)	(0.04)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.19	0.03	0.05	0.18	0.04

Note: This table is analogous to Table 3.4, but uses an alternative variable to construct the minimum constraint. Rather than EIA-reported minimum constraints, it uses the smallest bin with at least 5 percent of non-zero generating hours. Note this alternative definition does not impact the "off" or "at max" counts. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 4.25: Bins: The Regulation Requirement and Intensive/Extensive Margins

Panel A. Boiler

	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-0.25	0.12	0.18	0.04	-0.51	0.11	0.60	0.79***	0.34	-0.13	-1.26*
	(1.04)	(0.11)	(0.15)	(0.14)	(0.34)	(0.36)	(0.37)	(0.26)	(0.26)	(0.42)	(0.73)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.70	0.05	0.06	0.20	0.34	0.31	0.22	0.11	0.07	0.36	0.82

Panel B. Combined Cycle

	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-2.33***	-0.14***	-0.07***	0.03	-0.02	-0.06	0.33	1.65***	-0.14	1.23***	-0.46
	(0.86)	(0.04)	(0.02)	(0.03)	(0.07)	(0.13)	(0.27)	(0.37)	(0.40)	(0.46)	(0.46)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.05	0.04	0.04	0.02	0.07	0.17	0.14	0.08	0.54	0.46

Panel C. Combustion Turbine

	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-2.15	-0.08	0.02	0.06	0.07	0.27	0.57*	0.73	1.99**	0.14	-1.63**
	(3.39)	(0.14)	(0.09)	(0.08)	(0.12)	(0.23)	(0.30)	(0.44)	(0.96)	(0.84)	(0.78)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.14	0.14	0.14	0.14	0.22	0.26	0.29	0.39	0.42	0.24

Panel D. Other Tech

	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-0.03	-0.02	-0.02	-0.08***	-0.01	0.22***	-0.03	-0.10	-0.23**	0.01	0.28***
	(0.17)	(0.03)	(0.02)	(0.03)	(0.06)	(0.08)	(0.06)	(0.06)	(0.09)	(0.06)	(0.09)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.19	0.01	0.02	0.03	0.02	0.06	0.01	0.03	0.05	0.06	0.13

Note: This table is analogous to Table 3.4, but rather than using data on minimum constraints, it simply counts the number of units generating at 0 percent of capacity, 0 to 10 percent of capacity, etc. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.