NBER WORKING PAPER SERIES

THE LONG-RUN DYNAMICS OF ELECTRICITY DEMAND: EVIDENCE FROM MUNICIPAL AGGREGATION

Tatyana Deryugina
Alexander MacKay
Julian Reif

Working Paper 23483
http://www.nber.org/papers/w23483

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2017

We thank ComEd for generously sharing electricity usage data with us, and are particularly grateful to Renardo Wilson for many helpful discussions. We also thank Torsten Clausen and the Illinois Commerce Commission. We thank Severin Borenstein, Mike Florio, Don Fullerton, Katrina Jessoe, Nolan Miller, Erica Myers, Mar Reguant, and participants at the University of Illinois IGPA seminar, 2017 ASSA meetings, the 2017 CeMENT workshop, the EPIC lunch series, the POWER Conference on Energy Research and Policy, the Midwest Economics Association meetings, and the University of Pittsburgh seminar for excellent comments and suggestions. Noah Baird, Dylan Hoyer, and Chitra Jogani provided excellent research assistance. Views reflected here are solely those of the authors. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation
Tatyana Deryugina, Alexander MacKay, and Julian Reif
NBER Working Paper No. 23483
June 2017
JEL No. D12,Q41,Q48

ABSTRACT

Understanding the response of consumers to electricity prices is essential for crafting efficient energy market regulations, evaluating climate change policy, and investing optimally in infrastructure. We study the dynamics of residential electricity demand by exploiting price variation arising from a natural experiment: the introduction of an Illinois policy that enabled communities to select electricity suppliers on behalf of their residents. Participating communities experienced average price decreases in excess of 10 percent in the two years following adoption. Using a flexible difference-in-differences matching approach, we estimate a one-year price elasticity of -0.14 and three-year elasticity of -0.29. We also present evidence that consumers increased usage in anticipation of the price changes. Finally, we estimate a forward-looking demand model and project that the price elasticity converges to a value between -0.30 and -0.35 after ten years. Our findings demonstrate the importance of accounting for long-run dynamics in this context.

Tatyana Deryugina
Department of Finance
University of Illinois at Urbana-Champaign
515 East Gregory Drive, MC-520
Champaign, IL 61820
and NBER
deryugin@illinois.edu

Alexander MacKay
Strategy Unit
Harvard Business School
Soldiers Field
Boston, MA 02163
amackay@hbs.edu

Julian Reif
University of Illinois at Urbana-Champaign
College of Business
515 East Gregory Drive
Champaign, IL 61820
jreif@illinois.edu

1 Introduction

Economic theory suggests that demand is typically more elastic in the long run relative to the short run. When consumption depends on goods that are durable or habit-forming, consumers may take years to respond fully to a price change (Topel and Rosen, 1988; Becker et al., 1994). These demand-side dynamics are important considerations in the electricity sector, where suppliers, market regulators, and policymakers make decisions with long-run ramifications.

Generators and distributors require forecasts of the long-run demand response to price changes to invest optimally in capacity and infrastructure. Likewise, market regulators require these forecasts to design efficient allocation mechanisms and renewable energy subsidies. Policies in related areas, such as regulations to reduce greenhouse gas emissions, may significantly raise electricity prices by taxing coal and other inputs with low private costs (Karnitschnig, 2014). Evaluating the effects and incidence of these policies depends crucially on the price elasticity of electricity demand. For example, the realized emissions reduction from a carbon tax is a function of this parameter. Likewise, the economic costs of an emissions cap also depend on this price elasticity. More generally, estimates of electricity demand are important inputs to most partial and general equilibrium models that include the electricity sector.

Yet, there is little consensus about the magnitude of the price elasticity of electricity demand, especially in the long run. The absence of consensus is in part due to a lack of quasi-experimental studies. The few that exist are limited, analyzing short-run responses that last no more than four months. It is challenging to find both exogenous variation in prices and a suitable control group in electricity markets. Much of the observed price variation, such as an increase in the electricity price due to an unusually hot summer, is demand-driven and affects a large geographic area. Short-run supply shocks are often not passed through to consumers due to regulation, and those that are passed through often result in small, temporary price changes that are not salient. In addition, adjustment costs, such as the psychic costs of changing habits and the replacement costs of durable goods, may cause the long-run response to greatly exceed the short-run response. Measuring the full effect of a price change in such settings requires a long time series or a model of dynamics.

We provide the first quasi-experimental estimate of the long-run price elasticity of residential electricity demand, a category that accounts for 36 percent of U.S. electricity consumption (U.S. Energy Information Administration, 2017b). We find that the long-run elasticity is between -0.30 and -0.35, over twice the magnitude of the elasticity in the first year (-0.14). This difference in magnitude matters when predicting the effects of taxes, estimating policy incidence, and calculating returns to investment. Therefore, these estimates have important implications for climate policy, supply-side decisions, and the design of electricity markets. We discuss these implications in Section 7.

To construct our estimates, we exploit large, long-lasting, and plausibly exogenous variation in residential electricity prices arising from Illinois' Municipal Aggregation program (hereafter referred to as "aggregation").[1] This setting provides a clean natural experiment to study the effect of price on usage. Enacted by law in 2009, the aggregation program allowed communities in Illinois to choose an electricity supplier on behalf of their residents. The decision to implement aggregation depended on a community-wide referendum, and, as we show, was plausibly independent from expected changes in electricity usage. Communities that passed the referendum then selected an alternative supplier with a corresponding new rate. Customers in those communities were automatically enrolled in aggregation, unless they chose to opt out. Aggregation customers continued to receive their electricity bill from the utility in the same format, so the price variation in our analysis is not confounded with changes in how consumers are billed. Moreover, Illinois employs a linear price schedule for residential electricity, consisting of a modest fixed fee and a constant marginal price. With rare and short-lived exceptions, aggregation affected only the marginal price of electricity, greatly simplifying our analysis.

We employ monthly community-level usage data obtained from ComEd, a utility that services 3.8 million people in 885 Illinois communities, including the city of Chicago (Exelon, 2017). In our sample of 768 communities, 289 implemented aggregation. These communities obtained lower electricity prices than the prevailing ComEd rate, and the vast majority of consumers switched to the community-chosen supplier. We compare usage and price changes in communities that implemented aggregation to those that did not in order to estimate the price elasticity of consumer demand.

Our usage data span the years 2007-2014, which allows us to estimate trends in electricity usage over long periods of time both before and after most communities' implementation of aggregation. The relatively large number of ComEd communities that did not pass a referendum on aggregation (479 in our sample of 768) in combination with a lengthy

[1] In other settings, these programs are sometimes called "community choice aggregation." Opt-out aggregation is also available in California, Massachusetts, New Jersey, New York, Ohio, and Rhode Island.

pre-period provides an excellent empirical setting for a matching estimator. Specifically, we combine a difference-in-differences methodology with the matching estimator developed by Abadie and Imbens (2006, 2011). Matching estimators are particularly well-suited to our study because electricity usage is highly seasonal, and these seasonal patterns vary substantially across different communities. We demonstrate that our matching estimator obtains more precise estimates than a traditional difference-in-differences estimator. We view our application as a useful demonstration of matching for applied researchers, in the vein of Fowlie et al. (2012). In contrast to their paper, we employ subsampling to calculate confidence intervals for a richer space of estimates that do not have pre-existing formulas for standard errors.

We match each aggregation community to five "nearest neighbors" that did not pass a referendum on aggregation. We construct the matching criteria using monthly electricity usage profiles from 2008 and 2009. This matching period long precedes our natural experiment; more than ninety percent of the referenda in our sample are held after February 2012, over two years later. Our identifying assumption is that the average observed differences in usage between aggregation communities and their matched controls in the post-period are attributable to aggregation. We support this assumption by showing that usage trends among aggregation communities and their matched controls are parallel between the end of the matching period and the beginning of aggregation.

We estimate that prices fell by 22 percent (0.25 log points) and usage increased by 5.1 percent in the 7-12 months following an aggregation referendum, relative to control communities that did not pass a referendum on aggregation. These estimates imply an average price elasticity of -0.16 for that time period. In the second and third years, the relative price differences shrank to 13 percent and 10 percent, respectively. The narrowing of the price difference is due primarily to the expiration of a long-term ComEd contract in June of 2013, which led to lower prices among control communities. Correspondingly, we find that the usage difference shrinks during this period. Although our data have both a large initial price decrease due to aggregation and a modest subsequent (relative) price increase attributable to ComEd's long-term contract expiration, the price elasticity smoothly declines from -0.14 in the first year, to -0.27 in the second year, and to -0.29 in the third year, illustrating the importance of long-run dynamics in this setting. We find little evidence that the elasticity depends on economic characteristics, but we do find some variation in the elasticity across social demographics.

Following the passage of the aggregation referendum and choice of a new supplier, all residents were notified by mail of the exact month and size of the upcoming price decrease. Consistent with an economic model of forward-looking consumers who have adjustment costs, we find that usage increased shortly after passage of the referendum, but before the actual price decrease several months later. This result rejects the myopic "partial adjustment" model frequently estimated in the energy literature (Hughes et al., 2008; Alberini and Filippini, 2011; Blázquez et al., 2013). Although we do not identify the mechanism driving this behavior, our result provides compelling evidence that consumers begin adjusting their electricity usage prior to realized price changes.

Finally, we develop and estimate a forward-looking dynamic model of demand in which usage is a flexible function of past, current, and future prices. As before, our identifying variation comes only from price changes brought about by aggregation. We then use the estimates to calculate the long-run price elasticity beyond our sample period. Because the ComEd and average aggregation supply prices converged around 2015, how consumers form expectations may play an important role: if consumers expected the convergence, they may have been less responsive than if they thought the price decreases were permanent. We estimate the model under two different assumptions about consumer expectations. If we assume consumers have perfect foresight, then we estimate a long-run elasticity of -0.35. If, by contrast, we assume that consumers employ a status quo ("no change") forecast when forming expectations about prices, then we estimate a long-run elasticity of -0.31, which is similar to our three-year quasi-experimental estimate. In both cases, we estimate that the elasticity converges to its long-run value after approximately ten years.

The primary contribution of our study is a credible, quasi-experimental estimate of the long-run price elasticity of residential electricity demand. Existing short-run and long-run estimates vary widely: from nearly zero to about -0.9 in the short run (one year or less) and from -0.3 to about -1.1 in the long run.[2] None of the long-run estimates is based on quasi-experimental variation, and many rely on state-level data and dynamic panel models, which include lagged consumption as an independent variable.[3] Consistency generally

[2]Also, a growing literature investigates the impact of real-time pricing (e.g., Wolak, 2011; Allcott, 2011; Jessoe and Rapson, 2014). The elasticity we identify here is fundamentally different from the elasticity estimated in the real-time pricing literature, which reflects intra-day substitution patterns as well as any overall reductions in electricity.

[3]See Alberini and Filippini (2011) for a brief review. Some have argued that a state's average price of electricity can be considered exogenous because it is regulated (Paul et al., 2009) or because the unregulated component is driven by national trends (Bernstein and Griffin, 2005). However, electricity rates may be set based on the anticipated cost of electricity to suppliers and that cost, in turn, may be based on anticipated demand. Therefore, it is not possible to separate supply-side variation from demand-side variation in national changes in fuel prices without explicitly constructing instruments.

requires strong assumptions about the form of serial correlation, and Alberini and Filippini (2011) show that these models are particularly sensitive to the exact specification used. Conversely, our main approach makes relatively few assumptions, and our forward-looking dynamic model is substantially more flexible than what has previously been estimated.

Papers outside of the dynamic panel literature typically focus on short-run rather than long-run price elasticities. Using a structural model and exploiting the non-linearity of the electricity price schedule in California, Reiss and White (2005) estimate an average annual elasticity of -0.39. Ito (2014) uses quasi-experimental variation by comparing households located near a boundary between two California utilities, which vary in when and by how much they change prices. He estimates an average price elasticity of -0.07 to -0.09 in the 1-4 months following a price change. Neither study estimates the price elasticity over a longer time period. Our findings demonstrate that residential electricity consumers are more than twice as responsive in the long run relative to the short run, which, as we discuss in more detail below, has important policy implications.

We also contribute to the large literature on forward-looking demand. It is challenging to provide persuasive evidence that consumers react to future prices; indeed, pre-period changes in outcome variables are often interpreted as model misspecification rather than forward-looking behavior (Malani and Reif, 2015). However, similar to Gruber and Köszegi (2001), we argue that the pre-price-change responses we find in our setting likely reflect anticipation effects, because they appear only following the passage of the aggregation referendum. Additionally, our rich data allow us to estimate a forward-looking model under fewer assumptions than is typically required. To avoid being too demanding of the data, prior studies have generally estimated these models in the form of an Euler equation (Becker et al., 1994; Malani and Reif, 2015), which requires assuming exponential discounting and finding appropriate instrumental variables. By contrast, the long length of our panel allows us to estimate a flexible model of demand that imposes minimal structure on the form of discounting and does not require instruments.

The rest of this paper is organized as follows. Section 2 discusses electricity market regulation and Municipal Aggregation in Illinois and compares the state to the U.S. as a whole. Sections 3 and 4 describe our data and empirical approach, respectively. Section 5 presents the quasi-experimental results. Section 6 develops a simple dynamic framework to motivate the construction of long-run projections, and then estimates these projections. Section 7 discusses the implications of our main results, and Section 8 concludes.

2 The Illinois Electricity Market

The provision of electricity to residential customers consists of two components: supply and distribution. Suppliers generate or purchase electricity and sell it to customers, and distributors provide the infrastructure to deliver the electricity and often handle billing. Illinois has two regulated electricity distributors: Commonwealth Edison Co. ("ComEd") and Ameren Illinois Utilities ("Ameren"). Prior to 1997, they owned generating units as well as the distribution network. In 1997, they were given ten years to transition to a competitive electricity procurement process due to widespread agreement that, unlike distribution, electricity generation is not a natural monopoly (Illinois General Assembly, 1997). The price for electricity supplied by Ameren or ComEd is, by law, equal to their procurement cost and does not vary geographically.[4] The procurement cost is determined by an auction, and unanticipated changes to this cost are passed through to customers.

Customers are assigned their distributor on the basis of their geographic location. ComEd serves northern Illinois, and Ameren serves central and southern Illinois. While customers have no choice in distributors, in 2002 residential and small commercial customers gained the ability to choose an alternative retail electric supplier (ARES) who would be responsible for supplying (but not delivering) their electricity.[5] However, the residential ARES market was practically nonexistent between 2002 and 2005. This was blamed on barriers to competition in the residential market and a rate freeze that kept the default utility rate low. In 2006, the state removed some of these barriers and instituted a discount program for switchers. These changes increased the savings of switching to an ARES, but still had little effect on behavior. By 2009, only 234 residential customers had switched electricity providers. By contrast, 71,000 small commercial, large commercial, and industrial customers had switched (Spark Energy, 2011). Individual residential customers faced high switching costs relative to their usage, rendering alternative suppliers less attractive.

In 2009, the Illinois Power Agency Act was amended to allow for Municipal Aggregation, whereby municipalities and counties could negotiate the purchase of electricity on behalf of their residential and small commercial customers. Another form of local government, townships, gained this ability in 2012. The 2009 amendment, which went into effect on January 1, 2010, was motivated by the observation that few consumers switched away from the incumbent supplier on their own, even when the potential savings were

[4] Their profits stem from delivery fees, which are set by the Illinois Commerce Commission (ICC) (DeVirgilio, 2006).

[5] Large commercial and industrial customers gained this ability at the end of 1999.

large. To ensure that individual consumers retained the ability to choose their supplier, the amendment requires municipalities to allow individuals to opt out of aggregation.

To implement an opt-out aggregation program, municipalities must first educate their communities about aggregation using local media and community meetings, register the proposed aggregation program with the state, and hold a referendum. The wording of the referendum question is specified in the Illinois Power Agency Act and given in the Appendix. If the referendum is approved, the municipality must develop a plan, hold public hearings, and then have the plan approved by the local city council. Usually, the government hires a consultant to solicit and negotiate terms with a number of suppliers. In most cases, multiple suppliers submit (private) bids for predetermined contract lengths (e.g., one-, two-, and three-year contracts). In other cases, the municipality negotiates directly with a supplier. When determining the bid or negotiating directly, each supplier obtains aggregate community-level usage data from the distributor. These usage data, along with electricity futures, are the main factors in each offered price. Importantly, our analysis controls for these same community-level usage data, which reduces the likelihood that price changes are affected by confounding factors that are unobservable to us.[6]

The two main ways in which suppliers differentiate themselves are price and the fraction of generation derived from renewable sources. Nearly all communities select the supplier with the lowest price, although environmental preferences occasionally induce communities to select a more expensive one. Once a supplier is chosen, the price is guaranteed for the length of the contract. Because many of these initial contracts were in effect through the end of our usage data, the price variation employed by our study derives mainly from the first set of aggregation contracts signed by Illinois communities.

Importantly, instead of "opting in" to an ARES, customers in a community that has passed an aggregation referendum are automatically switched to the newly chosen electricity supplier unless they opt out by mailing in a card, calling, or filling out a form online.[7] The few residential customers who had already opted into an ARES or into real-time pricing prior to aggregation are not switched over to the chosen supplier. aggregation officially begins at the conclusion of the opt-out process. From the consumer's point of view, the

[6]Suppliers may also base their bids on the number of electric space heat customers, which we do not observe.

[7]While we do not have an exact number, ComEd and several energy suppliers have told us that the opt-out rate is low. Community-specific opt-out rates mentioned in newspapers range from 3 to 10 percent (e.g., Lotus, 2011; Wade, 2012; Ford, 2013). The number of non-aggregated customers does, however, grow slowly over time because new residents who move to an aggregation community are not defaulted into the aggregation program.

only thing that changes is the supply price of electricity on her bill. The bill is still issued by the incumbent distributor (Ameren or ComEd). The other items on the the bill, such as charges for distribution and capacity, remain the same for all customers in the distributor's territory regardless of aggregation. Conveniently, this means that the price effects of aggregation will not be confounded by billing confusion. We discuss the different prices in more detail in the Data section. The appendix to our paper includes a sample ComEd bill, a sample letter notifying households of the aggregation program, and a sample opt-out card.

In our setting, the realized savings from Municipal Aggregation come largely from the timing of the program. During our sample period, alternative suppliers were able to offer lower rates due to the unexpected boom in shale gas, while ComEd was locked into a long-term high-price procurement contract. In addition to a market-timing advantage, the availability of alternative suppliers allows communities with favorable load profiles to secure rates at discounts relative to the aggregated pool. Indeed, Municipal Aggregation is very popular in Illinois: as of March 2016, 741 (out of about 2,100) Illinois communities had voted to implement aggregation.

It is worth mentioning that Illinois is similar to the U.S. as a whole along several dimensions related to electricity consumption. Using daily station-level weather data from NOAA, we calculate that the average number of population-weighted heating degree days between 1965–2010 in Illinois (U.S.) is 5,997 (5,115), while the average number of population-weighted cooling degree days is 881 (1,086). Considering that the state-level standard deviations for these two measures are approximately 2,100 and 746, Illinois provides a close proxy for the average U.S. temperature patterns. Likewise, from 2000 to 2015, the difference between the average annual electricity price in Illinois and the U.S. as a whole was -0.4 cents per kWh, which is only about four percent of the U.S. mean (U.S. Energy Information Administration, 2017a).

From 2011 to 2015, Illinois had an average per capita income of $30,494 (in 2015 dollars), which is very close to the U.S. average of $28,930. Similarly, the Illinois employment rate was 65.6 percent, while the U.S. employment rate was 63.6 percent (United States Census Bureau, 2016). The 2010 demographic characteristics of Illinois are also comparable to the U.S. average: 14.5 percent of the Illinois population is black (versus 12.6 percent for the U.S.), 12.5 percent is over the age of 65 (versus 13 percent for the U.S.), and 24.4 percent is under the age of 18 (versus 24 percent in the U.S.). To understand how the generalizability of our results is affected by the particular characteristics of our sample, we explore how our elasticity estimates vary by socioeconomic characteristics in Section 5.3.

Table 1: Count of MEA Communities in Sample

Referendum Date	Implemented	Passed, Not Implemented	Voted, Not Passed
November 2010	1	0	0
April 2011	18	0	0
March 2012	164	0	28
November 2012	57	5	2
April 2013	38	3	6
March 2014	8	1	0
November 2014	3	2	0
Total	289	11	36

3 Data

We obtain electricity usage data directly from ComEd, which serves the vast majority of communities in northern Illinois, including the city of Chicago. The data contain monthly residential electricity usage at the municipality level for ComEd's 885 service territories from February 2007 through June 2014. We drop 106 communities from our analysis that are missing data or that experience changes in their coverage territory during our sample period (see the Appendix for further details). For our main analysis, we drop an additional 11 communities that passed a referendum approving aggregation but never implemented the program. We estimate our model using a balanced panel of monthly usage for the remaining 768 ComEd communities, of which 289 implemented aggregation.

We constructed the time series of ComEd electricity rates using ComEd ratebooks, which we obtained from the Illinois Commerce Commission. Prior to June of 2013, customers with electric space heating faced a lower rate than those with non-electric space heating. Because electric space heating is relatively rare in Illinois, we assume that the incumbent rate is equal to the non-electric space heating rate, which will be true for the majority of non-aggregation customers. Data on aggregation referenda dates, aggregation supply prices, and aggregation implementation dates were obtained from a variety of sources, including PlugInIllinois, websites of electricity suppliers, and municipal officials. The median length of time between passage of the aggregation referendum and commencement of the aggregation program is 4 months.

As shown in Table 1, 300 communities in the ComEd territory passed a referendum on aggregation during our sample period, and 289 of those communities eventually imple-

Figure 1: Spatial Distribution of Communities in Sample

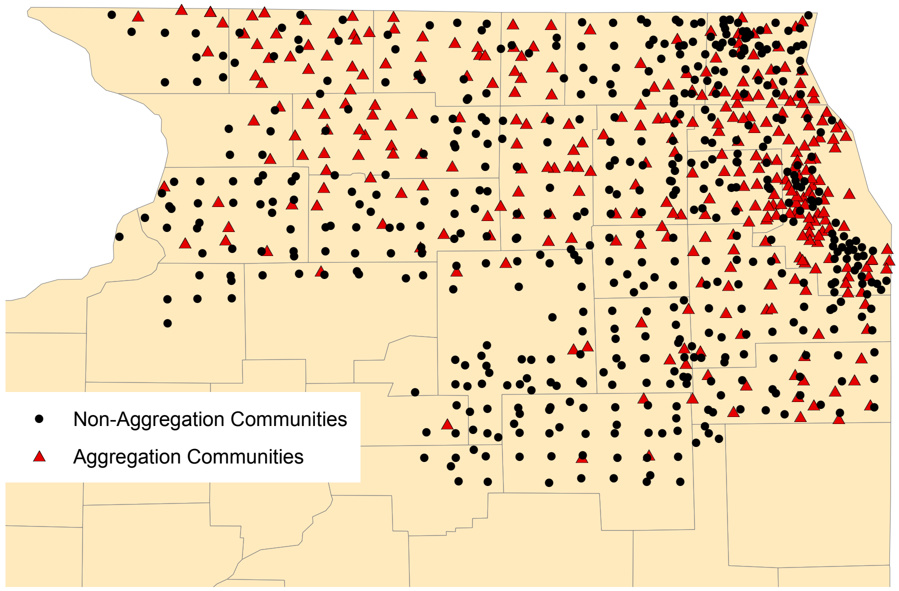

Notes: Figure displays the locations of communities in our sample. Red triangles indicate communities that implemented aggregation. Black dots indicated communities that did not pass aggregation.

mented an aggregation program.[8] In addition, 36 communities voted on but did not pass aggregation. Anecdotally, the reasons why some communities voted against aggregation include: (1) lack of trust in the local government to secure savings relative to the incumbent; (2) loyalty to the utility; (3) concern about the environmental impact of the resulting electricity use increase; (4) a misunderstanding about the opt-out provision; and (5) the belief that choosing an electricity provider for residents was not an appropriate government function.[9] The geographic locations of communities in our sample are displayed in Figure 1. Aggregation communities are well-dispersed throughout the ComEd territory but are more prevalent in the Chicago metropolitan area.

Many states employ a "block pricing" schedule where the marginal price of electricity increases with quantity purchased. Illinois, by contrast, employs a constant marginal price and a moderate fixed fee, which is appealing for us because it reduces confusion over the "price" to which consumers might be responding. This constant marginal price can be broken down into several components. Implementing aggregation entails a community signing

[8]Five of these communities passed a referendum in November of 2014, five months after the end of our usage data.

[9]These are based on authors' attendance of a public hearing in Champaign, notes from town hall meetings, discussions with ComEd, and discussions with the Illinois Commerce Commission, which regulates Illinois electricity providers and distributors.

a contract for a particular supply rate (the largest component of the marginal price) with an electricity supplier; non-aggregation communities pay the default ComEd supply rate. Thus, aggregation only affects the marginal price of electricity. All fixed fees and remaining usage rates are nearly identical across the aggregation and non-aggregation communities in our sample. The average fixed fee for customers residing in ComEd service territories during our sample period is $12.52 per month. We ignore the fixed fees in our analysis, because they do not vary across communities and because our independent variable is the difference in the (log) marginal price between aggregation and non-aggregation communities.[10] Municipal tax rates do vary across communities, but the variation is small. For our analysis, we use the median tax rate across ComEd communities (0.557 cents/kWh).

The thick dotted red and thin green lines in Figure 2a display ComEd's monthly supply rate and the total of all other usage rates, respectively, during and after our sample period. ComEd's supply rate dropped significantly in 2013, when its last remaining high-priced power contracts expired.[11] The blue line in Figure 2a shows the average monthly supply rate for communities that implemented aggregation beginning in June 2011, when the first community implemented aggregation. During our sample period, the average aggregation supply rate is always lower than the default ComEd supply rate. But starting around June 2015, the two rates became very similar and the number of communities choosing to continue aggregation decreased. The price variable in our estimating equations is equal to the aggregation supply rate plus all other usage rates if the community has implemented aggregation; otherwise, it is equal to the ComEd supply rate plus all other usage rates.

Figure 2b plots the mean and the 10th and 90th percentile of the log difference between aggregation and ComEd supply rates as a function of the number of months since the aggregation referendum. Aggregation supply rates are quite heterogeneous, and the average difference between them and the ComEd supply rate changes over time. There is also variation in how long implementation takes: at least 10 percent of aggregation communities switched suppliers within 3 months of the referendum, whereas 10 percent had not done so 6 months afterward. In most specifications, we construct estimates relative to the referendum date to capture any usage responses that occur prior to the actual price change.

[10] In rare instances the aggregation supplier charged an additional fixed fee, but these scenarios were short-lived.

[11] This drop was not a surprise. See, e.g., http://citizensutilityboard.org/pdfs/CUBVoice/SummerCUBVoice12.pdf.

Figure 2: Price Differences Between Aggregation and Non-Aggregation Communities

(a) Monthly Electricity Rates, 2008-2016

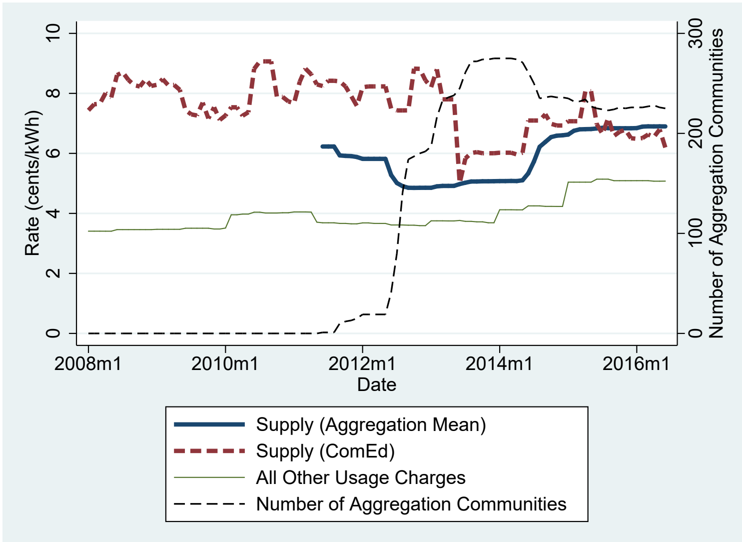

(b) Heterogeneity in Price Changes Realized by Aggregation Communities

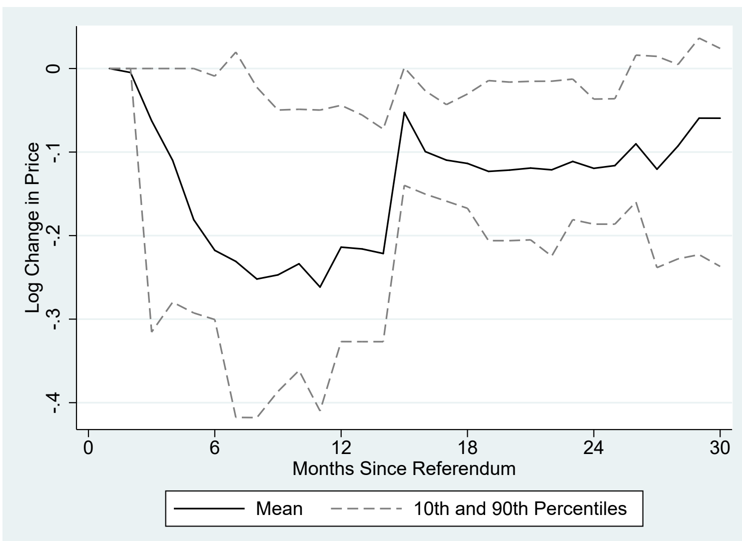

Notes: The thick blue line in panel (a) displays the average supply rate among all communities that adopted aggregation. The first community adopted aggregation in June of 2011. Non-aggregation communities pay the supply rate charged by ComEd, indicated by the thick, dashed red line in panel (a). The green line in panel (a) displays the total of all other electricity rates on a consumer's residential bill, which do not depend on whether a community has adopted aggregation. These displayed rates correspond to those for a single family residence with non-electric heating. The thin dashed line in panel (a) indicates the cumulative number of communities that have implemented aggregation.

Because ComEd's supply price does not vary geographically, we can calculate the savings aggregation communities obtained from switching suppliers. Specifically, we multiply aggregation communities' observed electricity usage by the price difference each month and aggregate over our sample period. Overall, residential aggregation consumers in our sample saved $566 million through June 2014.

4 Empirical Strategy

4.1 Difference-in-Differences Matching Framework

We construct counterfactual electricity usage by matching communities that implemented aggregation (the "treated" group) to communities that did not (the "control" group) based on their pre-aggregation electricity usage. Our setting is an ideal application for a matching estimator. The large, diverse set of control communities makes it likely that a nearest-neighbor matching approach will successfully find a suitable comparison group. Additionally, we have enough data in the period before aggregation for internal validation of the approach. As we show below, treated and control communities that are matched based on their 2008-2009 usage also have very similar usage patterns in 2010. The usage patterns diverge only after communities begin implementing aggregation in June of 2011.

Specifically, we apply a difference-in-differences adjustment to the bias-corrected matching estimator developed by Abadie and Imbens (2006, 2011). For each of the 289 treated communities, we use 2008-2009 electricity usage data to identify the five nearest neighbors from the 479 control communities available in our sample. We average annual log usage and monthly log deviations from annual usage across 2008 and 2009 to construct 13 match variables. We standardize the variables and use an equal-weight least squares metric to calculate distance. That distance is then used to select (with replacement) the five nearest neighbors for each treated community. We use these nearest neighbors to construct counterfactual usage, and we employ standard difference-in-differences techniques to adjust for pre-period differences. The identification assumption is that, conditional on 2008-2009 usage, the passage of aggregation and subsequent price changes are unrelated to anticipated electricity use. We provide evidence that this assumption is reasonable by showing that trends in usage for the control and treated groups remain parallel after the matching period but before the passage of aggregation.

A key advantage of the nearest-neighbor approach is that it eliminates comparison com-

Figure 3: Monthly Electricity Usage in Illinois ComEd Service Territories, 2007-2014

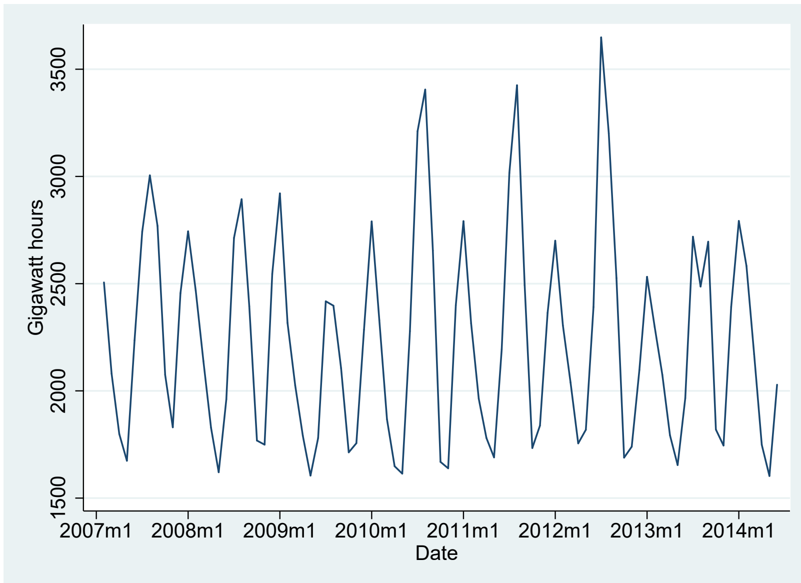

Notes: Figure displays total electricity usage across the ComEd service territories in our sample.

munities that are not observationally similar to treated communities and whose inclusion would add noise (and possibly bias) to the estimation. Electricity usage is highly seasonal, with peaks in winter and summer and troughs in spring and fall. These patterns are shown in aggregate in Figure 3. More importantly, the degree of seasonality varies widely across the different communities in our sample. Identifying control communities with usage profiles similar to aggregation communities can therefore greatly increase precision.

Figure 4 provides a demonstration of this benefit. Panel (a) displays electricity usage adjusted for community-level monthly seasonal patterns both for treated (aggregation) communities and communities that never passed aggregation. Even after accounting for community-specific seasonality, usage varies greatly within and across years: the largest peak occurs in July 2012, which corresponds to a record heat wave. By contrast, summer peaks are much less pronounced in 2009 and 2013, when the summers were mild. The difference between these two time series, which corresponds to an event study regression with community-specific month-of-year fixed effects, is displayed in panel (c). The increase in the difference is visible beginning in late-2011, which can be attributed to the implementation of aggregation, but this difference is quite noisy. The heterogeneity in seasonal patterns poses a challenge for a standard regression that compares treated communities to all control communities in the sample: it is difficult to estimate an effect when the baseline month-to-month divergence in usage is of the same order of magnitude as the effect.

Figure 4: Comparing Regression to Nearest-Neighbor Matching

(a) Regression (b) Nearest-Neighbor Matching

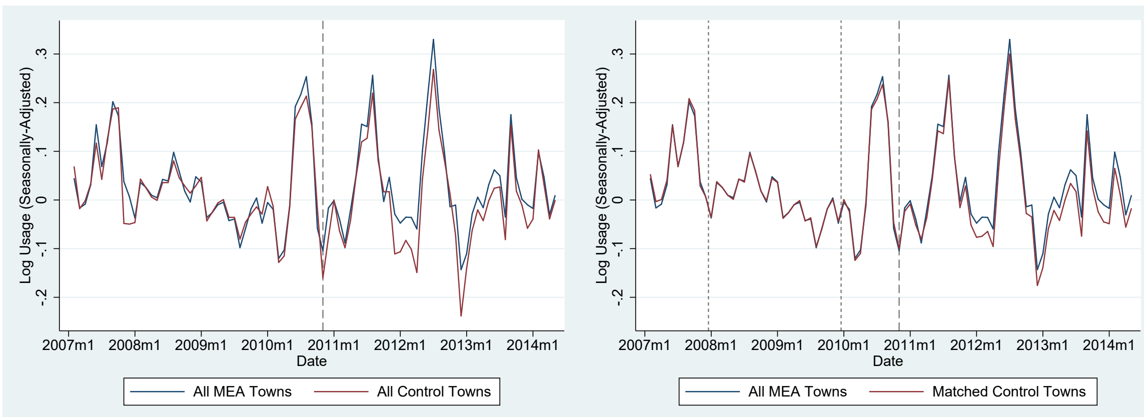

(c) Treatment and All Controls: Difference (d) Treatment and Matched Controls: Difference

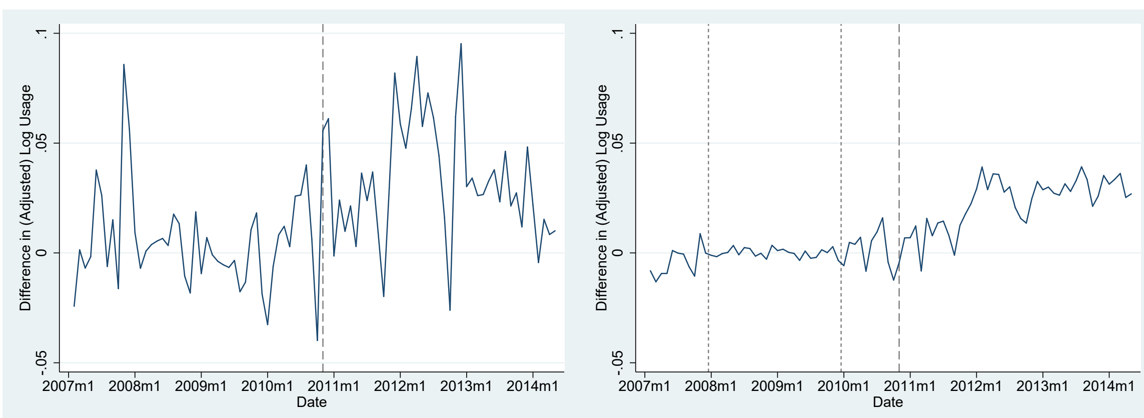

Notes: Panel (a) displays seasonally-adjusted usage for all aggregation and non-aggregation communities. The red line corresponds to the control group in a typical regression, with community-specific month-of-year fixed effects. Panel (b) employs the nearest-neighbor matching procedure, in which five communities are selected for each aggregation community, and the control line is weighted by how often each control community is selected. Panels (c) and (d) plot the differences between the treatment and control lines in Panels (a) and (b), respectively. The vertical dashed lines indicate the first referendum date. The vertical dotted lines in Panels (b) and (d) indicate the window used to match based on usage.

Panels (b) and (d) of Figure 4 show analogous plots for the nearest-neighbor matching approach discussed above. Panel (d) shows again that the difference in log usage between treatment and (matched) control communities increases beginning in late-2011. The differ-

Table 2: Characteristics of Aggregation, Non-Aggregation, and Matched Control Communities

	(1)	(2)	(3)	(4)	(5)
	Implemented Aggregation	Did Not Implement Aggregation		Matched Controls	
	Mean	Mean	p-value of Difference from (1)	Mean	p-value of Difference from (1)
Per Capita Electricity Usage in 2010, kWh	4,893	5,078	0.790	4,862	0.964
Total Population (Log)	8.63	7.20	<0.001	8.43	0.135
Percent Black	4.92	5.41	0.663	8.26	0.038
Percent White	86.54	89.06	0.055	83.49	0.087
Median Income	71,848	68,371	0.119	71,437	0.876
Median Age	38.63	40.80	<0.001	38.90	0.625
Total Housing Units (log)	7.69	6.27	<0.001	7.45	0.083
Median Year Built	1,969	1,965	0.006	1,972	0.023
Median Housing Value	264,723	222,617	0.001	250,355	0.310
Percent with High School Education	29.80	36.29	<0.001	32.75	0.005
Percent with Some College Education	29.73	31.39	0.008	30.53	0.227
Percent with Bachelor Degree	18.32	14.31	<0.001	16.71	0.087
Percent with Graduate Degree	11.22	7.43	<0.001	9.01	0.007
Latitude	41.91	41.67	<0.001	41.80	0.005
Longitude	-88.41	-88.53	0.025	-88.20	<0.001
Number of Unique Communities	286	385		271	

Electricity usage data come from ComEd. All other characteristics are from the 2005-2009 American Community Survey. Number of observations in column (1) is smaller for median year built (285). Number of observations in column (2) is smaller for median housing value (383). Estimates in columns (4) and (5) are weighted by the number of times the control community is a match for a treated community.

ence in panel (d) exhibits far less noise than the difference displayed in panel (c), because the matching estimator selects only those control communities that are similar to treated communities. This method of selection allows the matching estimator to generate more precise estimates than the standard difference-in-differences estimator.

To see whether matching also helps improve the similarity between treated and control communities on other dimensions, we matched the names of the communities in our ComEd sample to data obtained from the 2005-2009 American Community Survey (ACS) Summary File. We obtained ACS matches for 286 out of 289 aggregation communities, and 385 out of 479 non-aggregation communities.[12] Table 2 reports the mean character-

[12] Aggregation is generally implemented at the township, village, or city level. These different levels of municipal governments frequently overlap and have similar names. In order to minimize incorrect matches,

istics of the communities with a successful ACS match. We display results separately for communities that implemented aggregation (column 1), all communities that did not pass aggregation (column 2), and matched control communities (column 4). Columns 3 and 5 report the p-values for the null hypothesis that the difference between aggregation and all non-aggregation or between aggregation and matched controls is zero, respectively.

Compared to non-aggregation communities, communities that implemented aggregation are significantly larger, younger, and more educated. They are also less white and have more expensive and slightly newer housing. However, the average per capita electricity use in 2010 is very similar for aggregation and non-aggregation communities. After matching on electricity usage, the weighted pool of matched controls has more similar characteristics to the aggregation communities. With two exceptions (percent black and longitude), the p-values for differences are substantially larger, which indicates that matching on usage also selects control communities with more similar socioeconomic characteristics.

4.2 Estimating the Effect of Aggregation on Usage

We estimate how the price elasticity of electricity evolves over time using a two-stage approach. First, we estimate the effect of the policy change on usage for each treated (aggregation) community, using our matching estimator. Second, we use the observed change in price and the estimated change in usage to estimate demand elasticities.

Let Y_{it} denote log usage for community i in period t, where $t = 0$ corresponds to the referendum date for each treated community. For control communities, $t = 0$ corresponds to the referendum date of the treated community to which they have been matched. Let the indicator variable D_i be equal to 1 if a community ever implements aggregation and 0 otherwise. Y_{it} is a function of D_i, so that $Y_{it}(1)$ indicates usage when treated, and $Y_{it}(0)$ indicates usage when not treated. To calculate the effect of aggregation on electricity usage, we construct an estimate of untreated usage for aggregation communities, $\hat{Y}_{it}(0)$, which we describe below. Finally, let N denote the total number of communities in the sample, and $N_1 < N$ denote the number of treated (aggregation) communities in our sample. The average treatment effect on the treated in each period t is then estimated by

$$\hat{\tau}_t = \frac{1}{N_1} \sum_{i=1}^{N} D_i \left(Y_{it}(1) - \hat{Y}_{it}(0) \right).$$

we avoided ambiguous matches.

We observe the outcome $Y_{it}(1)$ for the treated communities in our data. The counterfactual outcome, $\hat{Y}_{it}(0)$, is unobserved and is calculated as follows. For each treated community i, we select $M = 5$ nearest neighbors using the procedure previously discussed. Let $\mathcal{J}_M(i)$ denote the set of control communities for community i. The counterfactual outcome, $\hat{Y}_{it}(0)$, is then equal to

$$\begin{aligned}\hat{Y}_{it}(0) &= \hat{\mu}_i^{m(t)} + \frac{1}{M}\sum_{j\in\mathcal{J}_M(i)}\left(Y_{jt}(0) - \hat{\mu}_j^{m(t)}\right)\\ &= \frac{1}{M}\sum_{j\in\mathcal{J}_M(i)} Y_{jt}(0) + \left(\hat{\mu}_i^{m(t)} - \frac{1}{M}\sum_{j\in\mathcal{J}_M(i)}\hat{\mu}_j^{m(t)}\right),\end{aligned}$$

where

$$\hat{\mu}_i^{m(t)} = \frac{1}{2}\left(Y_{i,m(t)}^{2008} + Y_{i,m(t)}^{2009}\right)$$

represents the average log usage in the calendar month corresponding to t for the years 2008 and 2009. The parameter $\hat{\mu}_i^{m(t)}$ is a standard bias correction that accounts for the average month-by-month usage patterns of each community. The variables $Y_{i,m(t)}^{2008}$ and $Y_{i,m(t)}^{2009}$ represent observed usage for community i in calendar month $m(t)$ in 2008 and 2009, respectively. For example, if $t = 25$ corresponds to January 2014, then $\hat{\mu}_i^{m(25)} = \frac{1}{2}\left(Y_{i,m(25)}^{2008} + Y_{i,m(25)}^{2009}\right) = \frac{1}{2}\left(Y_{i,January}^{2008} + Y_{i,January}^{2009}\right)$ is equal to the average log usage in January 2008 and January 2009. Thus, the estimated counterfactual $\hat{Y}_{it}(0)$ is equal to the average usage for a treated community's nearest neighbors plus the difference in usage between that community and its neighbors averaged across the 2008-2009 calendar months corresponding to t.

Finally, the "difference-in-differences" matching estimator that we employ subtracts the difference corresponding to the year prior to treatment from the average treatment effect on the treated, $\hat{\tau}_t$. It is defined as

$$\hat{\tau}_t^{DID} = \hat{\tau}_t - \frac{1}{N_s}\sum_{s=1}^{N_s}\hat{\tau}_{-s}, \tag{1}$$

where N_s indicates the number of periods in the year prior to the policy change.[13] Our difference-in-differences estimate thus reflects the change in usage between treated and

[13] $N_s = 12$ for monthly data. For biannual data, $N_s = 2$.

control communities in period t relative to the average difference in the year leading up to the policy change.

Malani and Reif (2015) show that failing to account for anticipation effects in a difference-in-differences framework can lead to biased treatment effect estimates when consumers are forward-looking. Ample evidence from the energy literature suggests this concern is valid. For example, Myers (2016) finds that expected future heating costs appear to be capitalized fully into housing values. While Allcott and Wozny (2014) reject full capitalization of gasoline prices into used vehicle prices, they nonetheless find that vehicle prices are significantly affected by future expected fuel costs. In our setting, the aggregation electricity price is announced months before the actual price change takes place. Thus, consumers may respond to future price changes by, for example, purchasing less energy-efficient appliances, or changing their thermostat program or energy use habits. Because we do not observe exactly when the price change is announced, our main specification estimates effects that are relative to the date on which the referendum was passed, rather than when aggregation was implemented.[14] In a second specification, we explicitly test for anticipation behavior by estimating changes in electricity use after the referendum is passed but before a community switches to a new supplier.

4.3 Estimating Elasticities

To estimate demand elasticities, we regress community-specific estimates of the change in usage on the observed community-specific price changes. The community-specific estimate $\hat{\tau}_{it}^{DID}$ is the single-community analog of Equation (1). It serves as the outcome variable in the following regression:

$$\hat{\tau}_{it}^{DID} = \beta_g \cdot \Delta \ln p_{it} + \eta_{it}, \tag{2}$$

where $\Delta \ln p_{it}$ is the difference in the logs of the marginal price of electricity between an aggregation community and its matched controls. Because electricity rates do not vary in the cross-section for non-aggregation communities, this difference will be exactly zero prior to a community's implementation of aggregation and will only reflect differences in marginal prices after aggregation has been implemented.

[14]It is also possible for consumers to make changes even prior to the passage of the referendum, in anticipation that it will pass and that electricity prices will fall. As we discuss later, we find no evidence of such behavior.

Equation (2) allows us to estimate period-specific elasticities that show how the response changes over time. In our main results, we run separate regressions with g corresponding to six-month intervals. The parameter of interest, β_g, corresponds to the average elasticity over the time interval g.

4.4 Inference

Because matching estimators do not meet the regularity conditions required for bootstrapping (Abadie and Imbens, 2008), we employ a subsampling procedure to construct confidence intervals for our matching estimates.[15] Subsampling, like bootstrapping, obtains a distribution of parameter estimates by sampling from the observed data.

For each of $N_b = 500$ subsamples, we select without replacement $B_1 = R \cdot \sqrt{N_1}$ treated communities and $B_0 = R \cdot \frac{N_0}{\sqrt{N_1}}$ control communities, where R is a tuning parameter (Politis and Romano, 1994) and N_0 is the number of control communities. As before, N_1 is the number of treated (aggregation) communities. For each subsample, we calculate $\hat{\tau}_b$. The matching estimator of the average treatment effect on the treated converges at rate $\sqrt{N_1}$ (Abadie and Imbens, 2006, 2011). The estimated CDF of $\hat{\tau}$ is given by:

$$\hat{F}(x) = \frac{1}{N_b} \sum_{b=1}^{N_b} \mathbf{1} \left\{ \frac{\sqrt{B_1}}{\sqrt{N_1}} (\hat{\tau}_b - \hat{\tau}) + \hat{\tau} < x \right\}$$

The lower and upper bounds of the confidence intervals can then be estimated as $\hat{F}^{-1}(0.025)$ and $\hat{F}^{-1}(0.975)$.

Subsampling requires a large number of effective observations (i.e., treated units) in each subsample, but it also requires that this number be small relative to the total number of effective observations in the full sample. We employ $R = 3$ ($B_1 = 51$) for the confidence intervals and standard errors reported in our main tables, which prioritizes the large-sample properties within each subsample. All other standard errors and confidence intervals are calculated similarly to those of $\hat{\tau}$. For example, we calculate elasticity estimates of β_g for each subsample to generate confidence intervals for our elasticities. Table A.2 in the appendix compares standard errors for different values of the tuning parameter, and shows they are robust to different values of R.

[15] Abadie and Imbens (2006, 2011) provide a formula for the standard errors for bias-corrected matching estimators of average treatment effects. Our panel data structure and the use of match-specific indexing for the control communities relative to the treated communities preclude a simple implementation of this formula. Further, the majority of our estimates are not simple average treatment effects.

5 Results

5.1 Event Study Results

We begin by showing that electricity prices fell substantially and persistently following the passage of aggregation referenda. Figure 5a displays the average change in log prices for aggregation communities, relative to their matched controls. The price change is exactly equal to zero in the pre-period because the treated communities face the same ComEd supply prices as their matched control communities during that time period. Within 12 months of passing the referendum, prices in aggregation communities decrease by nearly 0.3 log points relative to control communities, although they rebound significantly a few months into the second year. The rebound is attributable to a sharp decrease in ComEd's supply price in June of 2013 (see Figure 2a). Nonetheless, aggregation prices stay at least 10 percent lower than the control communities for most of the remaining estimation period.

Figure 5b displays the corresponding estimates for electricity usage. Prior to the referendum, the difference in usage between aggregation and control communities is relatively constant and never significantly different from zero. We emphasize that this result is not mechanical, as our 2008-2009 matching period predates the vast majority of aggregation instances by at least two years (see Table 1).[16] Following the referendum, usage in aggregation communities increases and eventually stabilizes at approximately 0.04 log points higher. The large increase in usage in the first year followed by a modest decrease mirrors the price patterns illustrated in Figure 5a. This zig-zag effect, which demonstrates that customers respond to both relative price decreases and increases, provides persuasive evidence that we are capturing the causal effect of price changes on electricity consumption.

Figure 6a displays our estimates of the price elasticity of demand, a key parameter of interest. To reduce noise, we estimate elasticities at the biannual, rather than monthly level.[17] The estimates increase in magnitude from about -0.09 in the first 6 months following the referendum up to -0.27 after two years, indicating that consumers are much more elastic in the long run than the short run. We also estimate a specification that models the price elasticity as a quadratic function of the number of months since the referendum. The results, displayed in Figure 6b, are very similar.

[16] Specifically, 270 of the 289 aggregation communities have virtually no overlap between the matching period and the pre-period estimates in Figure 5b.

[17] Monthly elasticity estimates are shown in Figure A.1 and Table A.1.

Figure 5: Effect of Implementing Aggregation on Electricity Prices and Usage

(a) Effect of Aggregation on Log Price

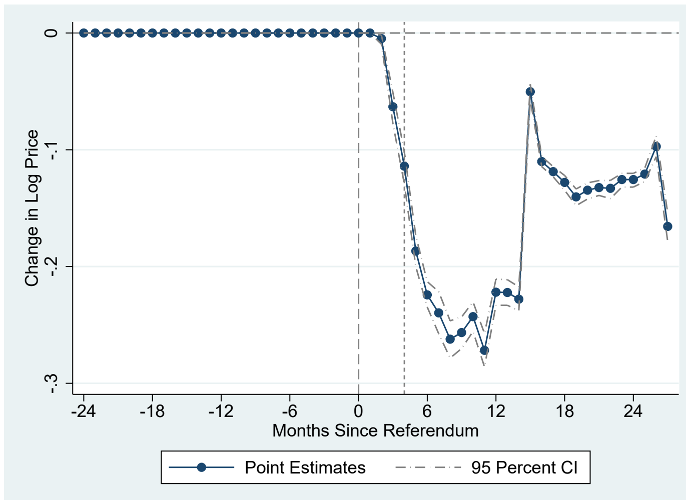

(b) Effect of Aggregation on Log Usage

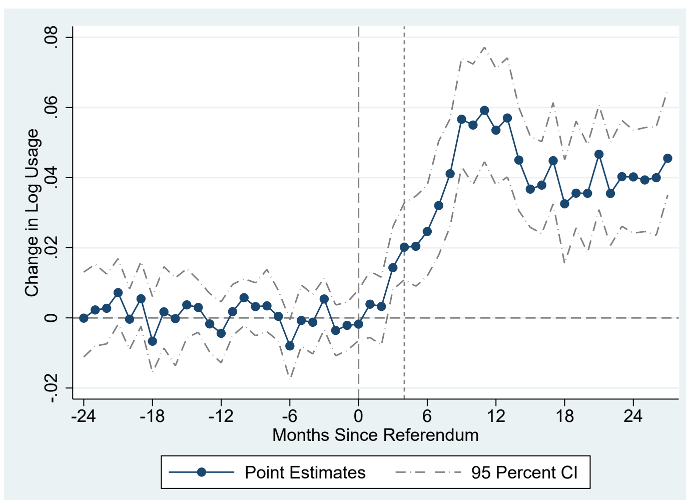

Notes: Panels (a) and (b) displays estimates of the mean price and usage effect, respectively, of implementing aggregation in a community. Prices differences are calculated using the natural log of the marginal electricity supply rates. The pre-period price difference is exactly zero. Usage is normalized so that the average usage difference in the year prior to the referendum is zero. After this normalization, the difference in the month prior to the referendum is -0.002. The short dashed line indicates the median implementation date relative to when the referendum was passed. Confidence intervals are constructed via subsampling.

Figure 6: Estimated Price Elasticities, Biannual

(a) Biannual

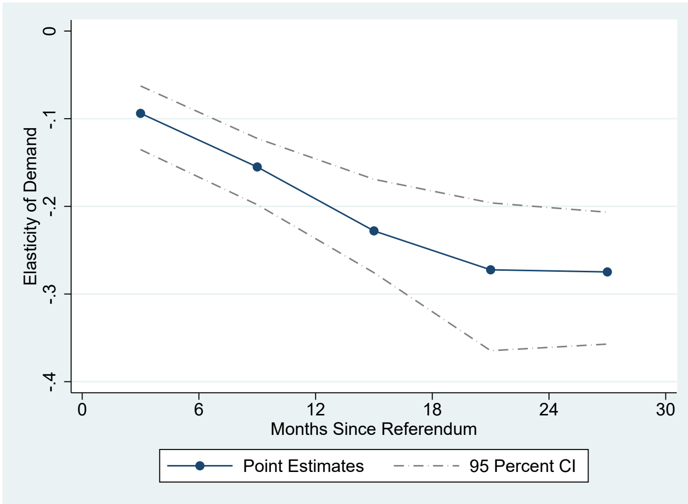

(b) Quadratic Fit

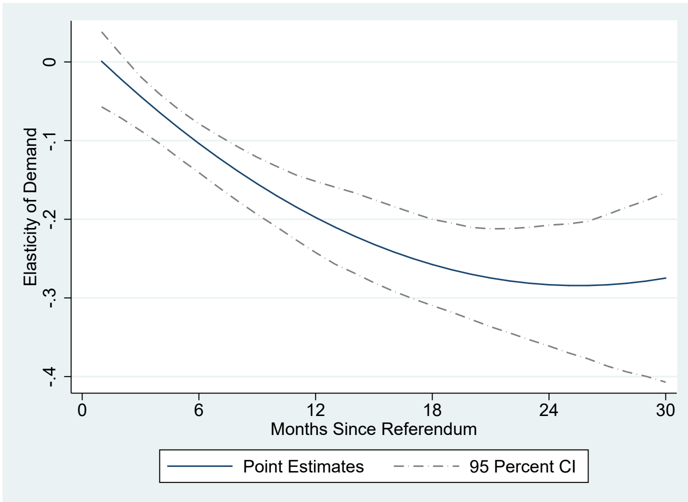

Notes: Elasticities in panel (a) are calculated for each six-month period by regressing community-month changes in log usage on the observed change in log price. The corresponding counts of observations for each six-month group are: 1685, 1656, 1504, 1144, and 589. In panel (b), the time-dependent elasticity is estimated using a quadratic specification. Community-month changes in log usage are regressed on changes in log price, where the log price changes are also interacted with months since referendum and the square of months since referendum. These three parameters are used to construct the estimated elasticity response curve as a function of time. Confidence intervals for both panels are constructed via subsampling.

Our main results from the difference-in-differences matching approach are summarized in Table 3. Table 4 reports the corresponding yearly estimates. All specifications report an

Table 3: Matching Estimates of the Effect of Aggregation on Usage and Prices

	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
1-6 Months Post-Referendum	0.014***	-0.098***	-0.094***	1692	1692
	(0.003)	(0.003)	(0.019)		
7-12 Months Post-Referendum	0.050***	-0.249***	-0.155***	1668	1668
	(0.007)	(0.007)	(0.020)		
13-18 Months Post-Referendum	0.043***	-0.147***	-0.228***	1516	1515
	(0.005)	(0.002)	(0.027)		
19-24 Months Post-Referendum	0.039***	-0.132***	-0.272***	1155	1155
	(0.006)	(0.003)	(0.043)		
25-30 Months Post-Referendum	0.043***	-0.120***	-0.275***	606	604
	(0.007)	(0.004)	(0.039)		

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each MEA town is matched to the five non-MEA towns with the most similar usage in 2008 and 2009. The number of price observations corresponds to the number of observations for each elasticity estimate, as we always observe usage where we observe a price change. Standard errors are in parentheses. Significance is determined by subsampling to construct confidence intervals.

increase in the magnitude of the price elasticity from about -0.1 in the first six months to nearly -0.3 in the third year. The finding that consumers are more elastic in the long run than in the short run is consistent with several mechanisms, including habit formation, learning, and appliance replacement over time. Because we do not have data on consumer behavior beyond electricity consumption, we do not attempt to distinguish among specific adjustment channels.

Our estimates suggest that the price elasticity converges to approximately -0.3 around 24 months after the policy change, although the limited length of our panel precludes us from confidently ruling out alternative possibilities. Because the price decrease following the aggregation referendum is not constant throughout the estimation period, the estimated elasticity at any given time captures the response to the longer-run price decrease due to aggregation, to the drop in ComEd prices in June 2013, as well as to monthly variations in the ComEd rate. Because the elasticity is growing smoothly throughout the estimation period, it appears that the long-run effects dominate the shorter-run responses to changes in the ComEd rate. However, in Section 6, we develop a model that – with a few assumptions – allows us to explicitly isolate the effects of price changes lasting across periods and to estimate the long-run elasticity.

Table 4: Matching Estimates of the Effect of Aggregation on Usage and Prices, Yearly

	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
1-12 Months Post-Referendum	0.032***	-0.173***	-0.140***	3360	3360
	(0.005)	(0.004)	(0.018)		
13-24 Months Post-Referendum	0.041***	-0.141***	-0.243***	2671	2670
	(0.005)	(0.002)	(0.028)		
25-36 Months Post-Referendum	0.046***	-0.108***	-0.285***	720	718
	(0.008)	(0.006)	(0.041)		

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each MEA town is matched to the five non-MEA towns with the most similar usage in 2008 and 2009. The number of price observations corresponds to the number of observations for each elasticity estimate, as we always observe usage where we observe a price change. Standard errors are in parentheses. Significance is determined by subsampling to construct confidence intervals.

5.2 Anticipation Effects

Forward-looking individuals should respond to policies prior to their implementation if those policies can be anticipated and if there is a benefit of responding in advance. For example, prior studies have documented that expectations of future policies and prices matter when purchasing durables such as cars or houses or when making human capital investments (e.g., Poterba, 1984; Ryoo and Rosen, 2004; Allcott and Wozny, 2014; Myers, 2016). The effect of aggregation on electricity usage is an ideal setting for detecting anticipation effects: the implementation of aggregation was widely announced months ahead of time, and electricity usage depends on durable goods like air conditioners, water heaters, and dishwashers, as well as on consumer habits and knowledge.

To demonstrate the presence of a pre-price-change usage response, we re-estimate the usage effects of aggregation, setting $t = 0$ to be the date that aggregation was *implemented* (which corresponds to the date of the price change) rather than the date of the aggregation *referendum*. The results are shown in panel (a) of Figure 7. The price difference between aggregation communities and their matched controls is exactly zero prior to implementation. However, usage begins increasing three months prior to the price change. Consulting Figure 5b reveals that this increase did not occur prior to the referendum. Together, these results suggest that the referendum and/or the subsequent price change announcement alerted consumers to the impending price decrease and caused them to increase their usage in anticipation, perhaps by, relative to their control counterparts, purchasing less energy efficient appliances or changing their electricity usage habits. Because mailers were sent to all residents informing them of the exact month of the price change, it is unlikely that customers were confused about the timing of the price change, although we cannot definitively rule out that possibility.

Figure 7: Anticipation Effects of Implementing Aggregation on Log Usage

(a) Defining $\tau = 0$ as the Implementation Date

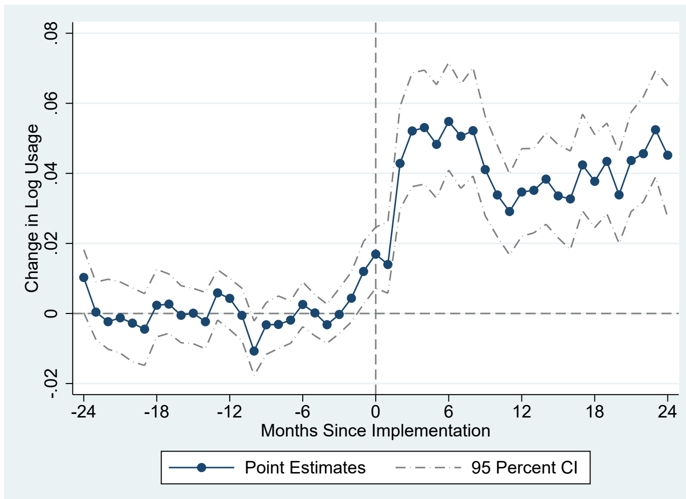

(b) Using Pre-Implementation Observations Only

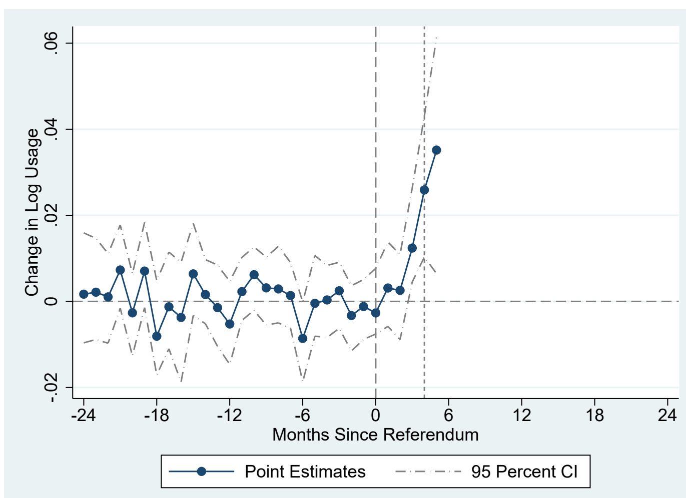

Notes: The figure displays estimates of the anticipation effects of implementing aggregation. In panel (a), the vertical dashed line corresponds to the month when aggregation was implemented. In panel (b), the vertical dashed line corresponds to the month when aggregation was passed, and the short-dashed line indicates the median implementation period. Confidence intervals are constructed via subsampling.

To explicitly isolate the anticipation effect, panel (b) of Figure 7 displays estimates of changes in electricity usage relative to the date of the aggregation *referendum* using only pre-implementation data.[18] Electricity usage increases steadily and significantly 3-5 months after the referendum despite the fact that prices have not yet changed for the observations in this sample. Specifically, usage is 0.012 log points higher 3 months after the referendum and 0.035 log points higher 5 months after the referendum, confirming the existence of substantial anticipation effects in our sample.

Because we do not calculate how large the pre-price-change usage response *should* be for a perfectly informed forward-looking individual facing adjustment costs, these estimates should not be interpreted as necessarily reflecting such behavior. Confusion and other behavioral mechanisms that affect our estimate should translate to other settings. Regardless of alternative explanations, what is important is that the anticipation effect we estimate is policy relevant; other policies that affect the price of electricity are likely to be implemented in a similar manner and generate pre-implementation effects similar to those we find here.

5.3 Elasticities by Demographic Characteristics

This section investigates how much variation in the price elasticity can be explained by demographic characteristics in order to understand both the distributional effects of policies that affect electricity prices and the generalizability of our results. For this exercise, we regress log usage on log price change and add interactions between the log price change and dummy variables indicating whether or not a community is in the top half of the distribution for x^j, a characteristic of interest:

$$\hat{\tau}_{it}^{DID} = \beta_g \cdot \Delta \ln p_{it} + \sum_{j=1}^{J} (\beta_j \cdot \Delta \ln p_{it} \cdot \mathbf{1}[x_i^j > median(x^j))] + \eta_{it}.$$

The indicator function $\mathbf{1}[x_i^j > median(x^j)]$ is equal to 1 if the value of the (time-invariant) variable x^j for community i is above the median of the distribution and 0 otherwise. We estimate this regression using $J = 8$ different characteristics obtained from the ACS, and report our results in Figures 8 and 9. Because the estimation is done jointly, the displayed elasticities for any given characteristic control for the other characteristics.

[18]Figure A.2 displays results for communities that passed a referendum but never implemented aggregation. Although the estimates are noisy, they suggest that there was no increase in usage due to the referendum itself in those communities.

Figure 8: Elasticities by Housing Demographics

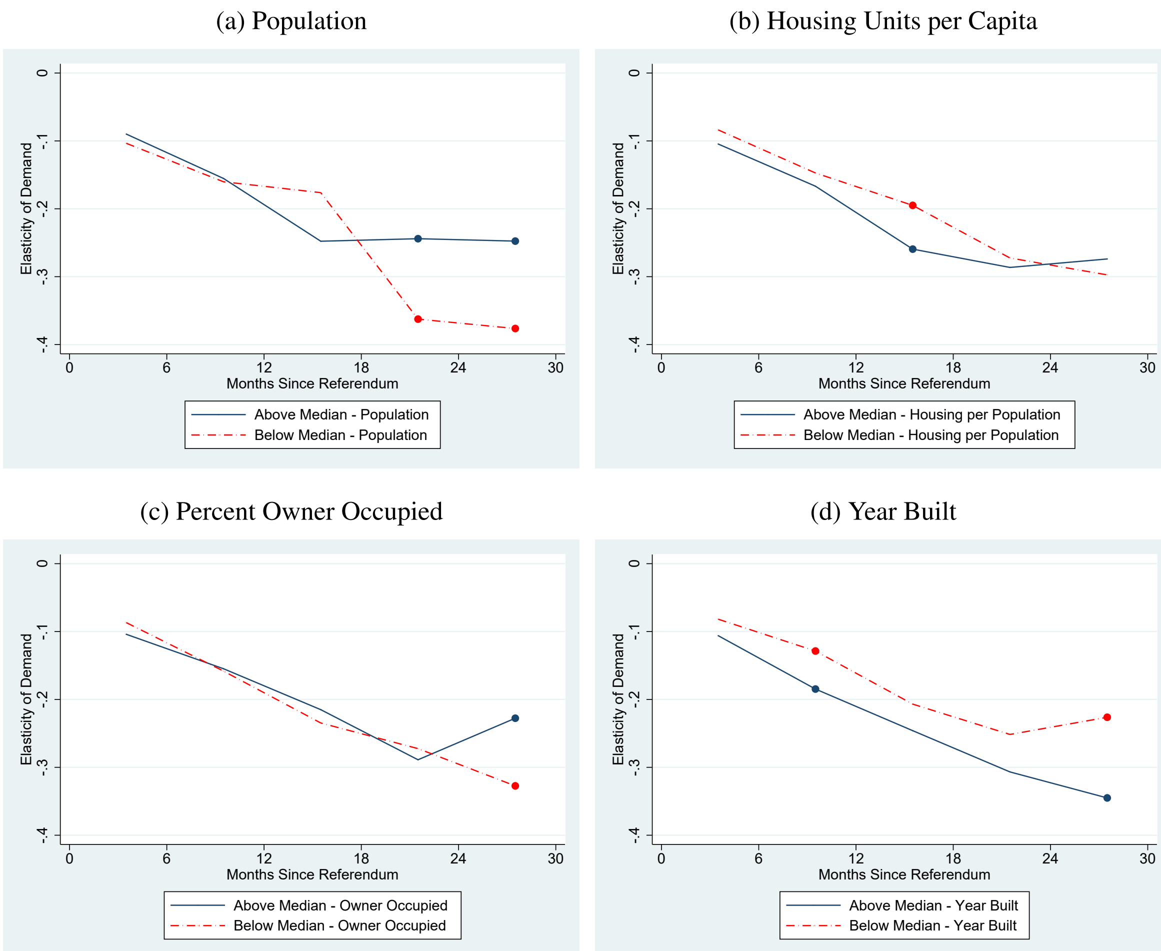

Notes: These panels display elasticity estimates for the upper half and lower half of demographic variables. The estimates are calculated by regressions of log usage on log price, where the price change is interacted with a dummy indicating whether or not the two is in the upper half of the distribution. The regressions control for eight interactions simultaneously - total population, housing units per capita, percent owner occupied, median year built, median income, percent with bachelor's degree, median age, and percent white. Significant coefficients ($\alpha = 0.10$) are indicated by the presence of a marker.

Figure 9: Elasticities by Socioeconomic Demographics

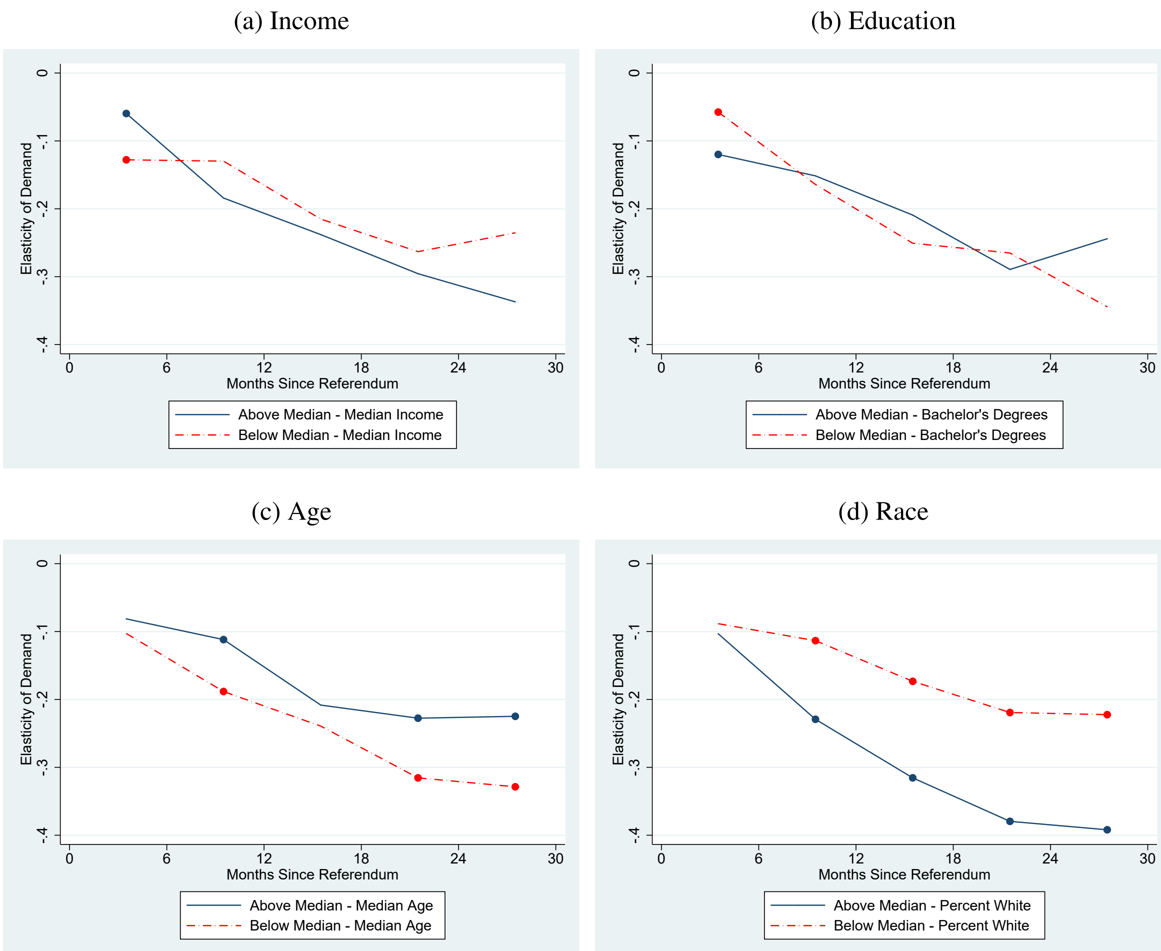

Notes: These panels display elasticity estimates for the upper half and lower half of demographic variables. The estimates are calculated by regressions of log usage on log price, where the price change is interacted with a dummy indicating whether or not the two is in the upper half of the distribution. The regressions control for eight interactions simultaneously - total population, housing units per capita, percent owner occupied, median year built, median income, percent with bachelor's degree, median age, and percent white. Significant coefficients ($\alpha = 0.10$) are indicated by the presence of a marker.

Figure 8 reports heterogeneity results for variables related to the housing stock. Estimates that are statistically different from each other (at the 10 percent level) are indicated with a marker. Communities with newer homes (as measured by "year built") have a more elastic demand response, conditional on the other characteristics. This difference could arise because newer homes are more likely to have technology such as programmable thermostats, which make it easier for consumers to control electricity consumption.

Figure 9 reports heterogeneity results for socioeconomic characteristics. Surprisingly, age and race appear to matter for the elasticity of demand more than economic variables such as income and education. Younger communities have a more elastic response, as do communities with a greater percentage of white people. By contrast, our elasticity estimates are relatively stable across economic characteristics.

5.4 Robustness Checks

One concern raised by our empirical approach is whether the magnitude of the price change that a community experiences is correlated with its expected demand elasticity. For example, suppliers might offer lower rates to more inelastic customers, as more elastic customers would demand more at the same price and drive up supply costs. We think this is unlikely because most individual communities are small relative to the total market. However, we empirically check for this possibility by splitting our treated communities into seven groups based on the price change in the first two years after the referendum. We then calculate elasticities separately for each group. Figure A.3 plots these estimates. As expected, we find no evidence of a relationship between the price change and the *ex post* estimated elasticity.

We chose January 2008 through December 2009 as our matching period because it allowed us to match treated communities to their controls in a high quality manner while also leaving us with a fairly long post-matching period that we used to test for pre-trends. We have also estimated our model using an alternative matching period of February 2007 through January 2009. Those results, available upon request, are very similar to the results we present in the paper.

Finally, we have also estimated the effect of aggregation on electricity usage using a standard difference-in-differences event study. These results, discussed in detail in the Appendix, are qualitatively similar to the results presented in the main text. In particular, we again find no evidence of pre-trends, supporting the identifying assumption that passage of aggregation was not prompted by growth in electricity usage.

6 Long-Run Projections

6.1 Conceptual Framework for Dynamics

As our estimates above demonstrate, electricity usage does not adjust to price changes instantaneously. It takes time to change one's habit of turning off the lights or air conditioning when away from home. Usage also depends on the energy efficiency and utilization of durables such as dishwashers, dryers, and air conditioners, which are typically purchased only once every 10 or 20 years. Finally, some consumers may not immediately realize that the electricity price has changed, especially if the benefit of tracking price changes is small relative to the cost of paying attention. Adjustment costs like these, both tangible and psychic, mediate the consumer response to electricity prices. In particular, these adjustment costs suggest that the long-run response to a price change will exceed the short-run response. Moreover, if consumers are forward-looking, then the presence of adjustment costs may cause them to respond in anticipation of future prior changes. In this section we present a simple framework to demonstrate these issues formally and to motivate the empirical model in the following section.

A simple way to model consumption dynamics is to employ the habit model of Becker et al. (1994) and allow current utility in each period to depend on y_t, the consumption of electricity in that period, and on y_{t-1}, the consumption of electricity in the previous period.[19] In this case, the consumer's problem is:

$$\max_{y_t, x_t} \sum_{t=1}^{\infty} \beta^{t-1} U\left(y_t, y_{t-1}, x_t\right),$$

where y_0 is given, $\beta < 1$ is the consumer's discount factor, and x_t represents consumption of a composite good that is taken as numeraire. The consumer's budget constraint is

$$W_0 = \sum_{t=1}^{\infty} (1/\beta)^{-(t-1)} \left(x_t + p_t y_t\right),$$

where W_0 is the present value of wealth, and p_t denotes the price of the electricity.

We assume that consumers are forward-looking, and for simplicity we assume they have perfect foresight. Finally, we will assume utility is quadratic. This functional form allows us to illustrate the different types of dynamics that can arise in our setting while

[19] Similarly, one could allow utility in each period to depend on a "stock" of appliances (Filippini et al., 2015). The resulting model will exhibit dynamics similar to what we present here.

still allowing us to derive analytical solutions to the consumer's problem.[20] Under this assumption, the demand equation is:

$$y_t = \alpha_1 y_{t-1} + \alpha_2 y_{t+1} + \alpha_3 p_t, \tag{3}$$

where the coefficients in (3) depend on the parameters of the quadratic utility function. The "adjustment cost" model frequently estimated in the energy demand literature corresponds to the special case where consumers are myopic, in which case the demand equation simplifies to:

$$y_t = \theta_1 y_{t-1} + \theta_2 p_t. \tag{4}$$

In the forward-looking model (3), consumers adjust their consumption in anticipation of future price changes. This does not occur in the myopic adjustment model (4). Prior studies have noted that one can therefore test the myopic model by testing whether consumers respond to future prices (Becker et al., 1994; Gruber and Köszegi, 2001).

The effect of a price change on consumption will depend on whether the change was anticipated and on how long consumers expect the price change to last. Researchers are typically interested in estimating the long-run effect of a permanent change in price, which, in terms of the parameters from equation (3), is equal to:

$$\frac{dy}{dp^*} = \frac{\alpha_3}{1 - \alpha_1 - \alpha_2}. \tag{5}$$

While equation (3) is parsimonious and can be estimated using short panels, consistently estimating the parameters α_1, α_2, and α_3 requires at least two valid instruments for the endogenous variables y_{t-1} and y_{t+1}. Alternatively, one can estimate consumption as a function of all past and future prices:

$$y_t = \sum_{s=0}^{t-1} \beta_{t-s} p_{t-s} + \sum_{s=1}^{\infty} \beta_{t+s} p_{t+s}. \tag{6}$$

In this case, the long-run effect of a permanent change in price is equal to

$$dy/dp^* = \sum_{s=0}^{t-1} \hat{\beta}_{t-s} + \sum_{s=1}^{\infty} \hat{\beta}_{t+s} \tag{7}$$

[20]Alternatively, one could allow utility to be general and take a linear approximation to the first-order conditions to analyze dynamics near a steady state, which would yield the same equations we present here.

Because aggregation provides us with plausibly exogenous variation in prices, this is the approach we take here. The appendix provides a derivation that demonstrates the equivalence of equations (5) and (7) in the case where consumers are exponential discounters and have perfect foresight.

6.2 Empirical Estimates

As with other studies estimating a long-run elasticity, we face the challenge that our dataset spans a limited time window: the results reported in Section 5 estimate the monthly response with reasonable accuracy only up to 27 months beyond the referendum date. Moreover, Figure 2a shows a convergence in average prices for aggregation and non-aggregation communities starting in June of 2015. If this convergence was anticipated by consumers, then it might have affected the usage response in earlier years. Investigating how the long-run response to a permanent reduction in price corresponds to our quasi-experimental estimate requires accounting for these future price changes, which, in turn, requires imposing some assumptions regarding consumer expectations.

Motivated by the adjustment cost model presented above, we estimate the long-run price elasticity of demand by regressing our matching estimates of log usage changes on lags and leads of log price changes:

$$\hat{\tau}_{it}^{DID} = \sum_{s=-L_1}^{L_2} \beta_s \cdot \Delta \ln p_{i(t-s)} + \eta_{it}. \tag{8}$$

The number of leads in the regression is equal to L_1, and the number of lags is L_2.[21] Although we do not observe usage past June 2014, we observe prices through 2016, allowing us to estimate equation (8) without losing observations. To reduce sensitivity to outliers, we estimate this model using median (quantile) regression instead of least squares.[22]

We first present results assuming that consumers have perfect foresight, as implicitly assumed by the formulation of equation (8). By contrast, Anderson et al. (2013) present evidence that consumers have status quo ("no change") expectations, at least with respect to gasoline prices. While the anticipation results presented earlier reject that model of belief

[21] We follow our earlier specification and estimate our model in logs. Our results are very similar if we instead estimate the model in levels, as given by equation (6).

[22] Point estimates from least squares regressions are similar, but the subsampling routine for least squares does not converge for a meaningful number of subsamples. This reflects the sensitivity of least squares to outliers and the relatively small size of each subsample.

formation, it is possible that this is due to the unusual salience of aggregation implementation during the time period we analyze. As a robustness check, we also present results assuming that consumers have status quo expectations with respect to aggregation savings after the end of our usage data.[23] Specifically, we assume that aggregation communities expect savings from future aggregation contracts to be equal to the average price difference in the twelve months leading up to the expiration date of the contract that is in place as of June 2014. Thus, differences between the perfect foresight and the status quo estimates will be driven by changes in price occurring after June 2014 (see Figure 2a).

The markers in Figure 10a correspond to estimates assuming that consumers have perfect foresight. We use four sets of leads and lags, starting with 18-month leads and ending with a specification with lags only. The sum of the number of leads and lags in each specification is equal to 30. To account for compositional effects, we also estimate versions of equation (8) with community fixed effects (red markers). Figure 10b presents corresponding estimates when we assume status quo expectations rather than perfect foresight.

To calculate elasticities beyond the time span covered by the leads and lags, we estimate a parametric specification that restricts the relationships among the set $\{\beta_s\}$ in equation (8) to fit a four-parameter adjustment curve. Based on the non-parametric estimates presented in Figure 10, we assume an exponential form for the lags and a linear form for the leads:

$$\beta_s = (\gamma_1 - \frac{\gamma_1}{\gamma_2} s \cdot \mathbf{1}[s \leq 0]) \cdot \mathbf{1}[s \geq \gamma_2] + \gamma_3(1 - \exp(\gamma_4 s)) \cdot \mathbf{1}[s >= 0] \tag{9}$$

The elasticity corresponding to a contemporaneous price shock is equal to γ_1, and the long-run elasticity is equal to $\gamma_1 + \gamma_3$. The parameters γ_2 and γ_4 govern the speed of adjustment. The fitted models are plotted as solid lines in Figure 10 and line up closely with the unrestricted estimates from equation (8).

[23]One could also employ other models of belief formation such as rational or adaptive expectations. See Malani and Reif (2015) for discussion of different approaches to modeling forward-looking regressions.

Figure 10: Long-Run Elasticity Projections

(a) Perfect Foresight

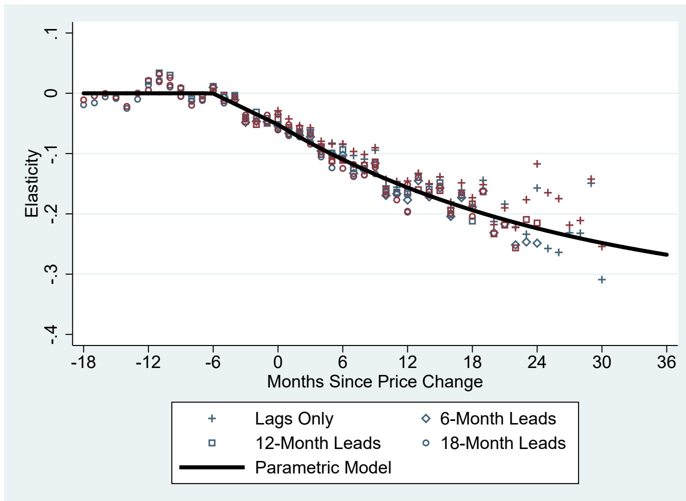

(b) Status Quo Expectations

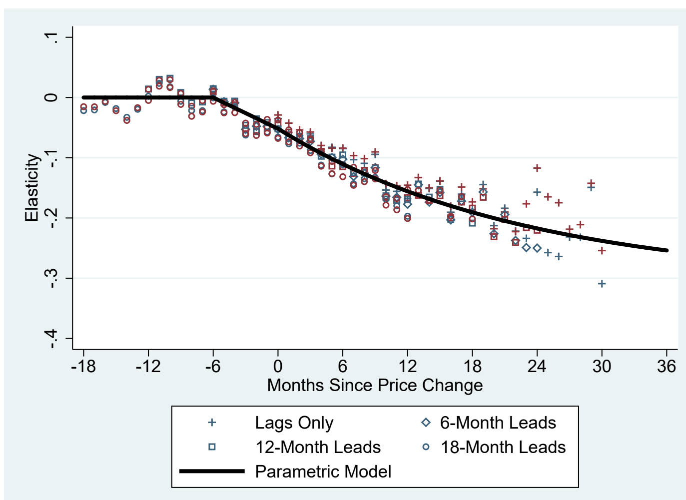

Notes: These graphs compare non-parametric and parametric estimates of the dynamic elasticity curve. The plotted points display non-parametric estimates derived from equation (8), where each set of points is calculated using 30 lags and leads, beginning with the lead displayed in the legend. Each point displays the estimated elasticity as a function of the number of months since a price change. The point corresponding to the cumulative effect at lag l is calculated as $\sum_{s=-L_1}^{l} \widehat{\beta}_s$, where L_1 is the number of leads included in the specification. The red points correspond to non-parametric models estimated with community fixed effects; the blue points are from specifications with no fixed effects. The solid lines display the estimates derived from the parametric model constrained by equation (9).

Table 5: Parametric Estimates of the Dynamic Elasticity Curve

Time Period	Perfect Foresight	Status Quo Expectations
Month 7-12 Leads	0.000 [-0.004, 0.000]	0.000 [-0.005, 0.000]
Month 1-6 Leads	-0.022*** [-0.038, -0.014]	-0.022*** [-0.039, -0.014]
Contemporaneous	-0.052*** [-0.075, -0.033]	-0.052*** [-0.076, -0.034]
Month 1-6 Lags	-0.087*** [-0.114, -0.063]	-0.087*** [-0.113, -0.065]
Month 7-12 Lags	-0.138*** [-0.174, -0.103]	-0.138*** [-0.173, -0.106]
Month 13-18 Lags	-0.179*** [-0.227, -0.134]	-0.177*** [-0.223, -0.134]
Month 19-24 Lags	-0.212*** [-0.271, -0.158]	-0.207*** [-0.264, -0.155]
Month 25-30 Lags	-0.239*** [-0.310, -0.178]	-0.230*** [-0.300, -0.171]
Month 31-36 Lags	-0.260*** [-0.345, -0.190]	-0.248*** [-0.332, -0.182]
Year 5 Lag	-0.315*** [-0.462, -0.225]	-0.288*** [-0.442, -0.210]
Year 10 Lag	-0.345*** [-0.606, -0.243]	-0.306*** [-0.601, -0.221]
Year 25 Lag	-0.348*** [-0.841, -0.245]	-0.307*** [-0.825, -0.221]
Long Run	-0.348*** [-1.007, -0.245]	-0.307*** [-1.079, -0.221]

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a regression of log usage changes on leads and lags of log price changes. The coefficients are constrained to match a four-parameter model. Changes in log usage and log price are estimated using a nearest-neighbor matching approach where each MEA town is matched to the five non-MEA towns with the most similar usage in 2008 and 2009. 95 percent confidence intervals are displayed in brackets and are constructed via subsampling.

Table 5 summarizes the estimates of our dynamic elasticity curve. The first column displays the results under the perfect foresight assumption, and the second column displays results assuming status quo expectations for contracts that begin after June 2014. The reported coefficients are quite similar for lags up to two years after the price change. The estimated elasticities, which reach -0.21 in months 19-24, are slightly smaller than what we find in our model-free estimates from Section 5. We should not be surprised to find some differences, as the model-free estimates capture impacts arising from past, present, and future price changes, while the estimates in this section explicitly isolate effects from a price change s months ago. However, the fact that the two sets of estimates are very similar suggests that long-run dynamics are dominant in our setting.

For longer-run effects, the perfect foresight assumption results in slightly larger elasticities than the status quo assumption. Specifically, the perfect foresight model implies a long-run price elasticity of -0.35. In the status-quo model, our long-run elasticity estimate is -0.31. This is sensible: if consumers correctly anticipated that prices would converge after June of 2014, then this expectation would dampen the usage increase in earlier years. We would therefore attribute more of the observed response to past prices and estimate a larger price elasticity.

Our model estimates that roughly 80 percent of the long-run consumer response occurs within three years after the price change. We forecast that it takes about ten years for consumers to fully adjust, at which point our estimate is nearly identical to the long-run response. Taking into account the 95 percent confidence intervals of the two formulations of expectations, we conclude that that the ten-year elasticity is bounded by -0.60 and -0.22. Both specifications find a small but significant anticipation response in the six months prior to a price change.

7 Implications of the Results

Emissions Pricing Our empirical estimates show that the price elasticity of residential electricity demand is substantially larger in magnitude in the long run than in the short run. This result has two important implications for regulation that targets electricity generation and subsequently affects the price of electricity (e.g., emissions pricing). First, the long-run change in electricity consumption following such regulation is likely to be significantly larger than estimates using the short-run elasticity would imply. Second, the share of the overall regulatory burden borne by consumers is likely to be smaller than expected, at least in partial equilibrium.

Specific quantitative conclusions are beyond the scope of this paper, however, because two other parameters that are important for understanding the likely effect and incidence of policy are not well-known: the price elasticity of electricity supply and the extent of pass-through to consumers in regulated electricity markets. Below, we explain how these different parameters matter for incidence and for policy.

Policymakers who wish to use a tax to target a specific level of emissions must correctly quantify how equilibrium electricity consumption responds to changes in taxes. In perfectly competitive markets, it is well known that the fall in equilibrium quantity following a small tax increase depends closely on both the demand and supply elasticities (Salanie, 2011):

$$\frac{\partial Q}{\partial t} = -\frac{\epsilon_s \epsilon_d}{\epsilon_s + \epsilon_d}\frac{Q}{P} = -\frac{1}{1/\epsilon_d + 1/\epsilon_s}\frac{Q}{P},$$

where Q is quantity, t is a per-unit tax, P is price, and ϵ_s and ϵ_d are the absolute values of supply and demand elasticities, respectively. It is clear from this expression that more elastic demand corresponds to a greater fall in equilibrium electricity consumption. Underestimating the demand elasticity by using its short-run value will thus lead to an understatement of emissions reduction from a small carbon tax.[24] The degree of understatement increases with the supply elasticity. For example, if the supply elasticity is $\epsilon_s = 0.5$, then the change in quantity when $\epsilon_d = 0.3$ is 2.25 times larger compared to $\epsilon_d = 0.1$. In other words, quantity is projected to fall by more than twice as much when the larger (e.g., long-run) demand elasticity is used. If supply is very elastic, say, $\epsilon_s = 5$, the change in quantity is 2.89 larger when $\epsilon_d = 0.3$ compared to $\epsilon_d = 0.1$.

It is also well known that the partial equilibrium incidence of a tax depends on the relative elasticity of supply and demand (Salanie, 2011). In a perfectly competitive market, the change in the tax-inclusive price (or "pass-through rate") for a small change in the tax rate is approximately

$$\rho \equiv \frac{\partial P}{\partial t} = \frac{\epsilon_s}{\epsilon_s + \epsilon_d} = \frac{1}{1 + \frac{\epsilon_d}{\epsilon_s}}.$$

The relative incidence of a tax, defined as the ratio of the tax burden borne by the consumers to that borne by the producers, can be shown to equal $\frac{\rho}{1-\rho}$, or simply the ratio of supply and demand elasticities, ϵ_s/ϵ_d. Weyl and Fabinger (2013) extend the analysis of incidence to imperfectly competitive markets and show that, while the incidence formula becomes slightly more complicated, it is still a function of the pass-through rate ρ and thus of the demand elasticity. Thus, underestimating the demand elasticity by using its short-run value leads to an overestimate of the share of the overall tax burden borne by consumers.

[24]Because carbon emission rates depend on an electricity generator's type (coal, natural gas, wind, etc.), translating a carbon tax into changes in electricity consumption is more complicated in practice. But the influence of the demand elasticity is similar.

One concern that arises when discussing incidence in the electricity sector is that residential electricity rates are often controlled by utility regulators. In such cases, any changes in generator costs (including a carbon tax or cap-and-trade) are passed onto consumers gradually, making it difficult to estimate the long-run pass-through rate. As of 2016, however, fifteen states had retail choice for electricity supply (Morey and Kirsch, 2016). Some of these states, as well as others that do not allow retail choice, impose profit constraints on incumbent suppliers (e.g., the zero-profit condition in Illinois). Thus, the supply component of price is more tightly linked to changes in the costs of electricity generation. In such cases, the incidence framework above remains relevant.

Carbon Leakage The price elasticity of electricity demand is also relevant for simulations of carbon leakage (e.g., Wing and Kolodziej, 2009; Elliott and Fullerton, 2014; Baylis et al., 2014; McKibbin et al., 2014). A typical leakage model specifies two or more goods, at least one of which is subject to a carbon tax. If there is at least one good that generates carbon emissions but is not taxed (e.g., because it is produced in a location with no carbon tax or by a sector not subject to the tax), carbon leakage, whereby consumption of the non-taxed carbon-emitting goods changes, can result. Elliott and Fullerton (2014) argue that taxing only electricity sector emissions is a likely scenario for carbon policy, at least initially, even though other sectors also generate carbon emissions. In this case, the price elasticity of demand for electricity is a key determinant of the extent of carbon leakage. More generally, this price elasticity is important for any general equilibrium model that includes the electricity sector.

Supply-Side Dynamics Conceptually, the long-run price elasticity of demand is also important for any study of long-term dynamics in electricity markets, whether such efforts are aimed at understanding the effects of deregulation (Olsina et al., 2006), generator entry decisions (Takashima et al., 2008), the diffusion of renewable energy (Kumbaroğlu et al., 2008), or the effects of environmental regulation (Cullen, 2014; Bushnell et al., 2017). For example, if generator entry raises electricity prices, then a larger long-run price elasticity of demand dampens entry incentives by reducing the size of the market. Conversely, if entry lower prices, then a larger consumer response will increase the size of the market and thus the profit from entering, all else equal. The assumptions made about the price elasticity of demand vary widely in the literature, and, with the exception of Bushnell et al. (2017), the sensitivity of results with respect to this parameter is often not discussed. Our study provides a well-identified estimate that can be employed in future research in this area.

Market Design Finally, our findings matter for the organization of electricity markets. One example is recent industry discussions regarding whether to use a downward-sloping or vertical demand curve in forward capacity auctions for the Midcontinent Independent System Operator (MISO) electricity area (Cook, 2016b,a). Because the relevant time period for these auctions spans multiple years, our estimates lend more weight to the case for using a downward-sloping demand curve in such settings.

8 Conclusion

An accurate estimate of the long-run price elasticity of electricity demand is valuable to electricity generators, distributors, and regulators. The price elasticity is also vital for projecting the effects of emissions policies such as a carbon tax. Policies that address climate change or other issues associated with electricity generation are likely to alter the price of electricity, which in turn will affect demand and any accompanying emissions. However, few reliable estimates of the price elasticity exist, as price changes in this market are often endogenous, short-lived, small, or unnoticed. In addition, models that use short-run price changes to infer the long-run elasticity of demand require strong assumptions.

Our study provides the first quasi-experimental estimate of the long-run price elasticity of residential electricity demand, and finds that it is about double the short-run elasticity. Our results underscore the importance of identifying settings that accurately capture long-run elasticities, as short-run data may vastly understate total effects. We also demonstrate that consumers began adjusting their electricity usage in the months leading up to the price change. Thus, in settings where changes are known ahead of time, accounting for anticipation effects is crucial for obtaining a correct estimate of the consumer response.

Finally, we note that the natural experiment created by aggregation decreased electricity prices, whereas price-based climate policies would increase prices to reduce total carbon emissions. It is therefore important to know whether the demand response is symmetric for price increases and price decreases. For example, appliance replacement rates might respond more quickly to a price increase than a price decrease, although this difference may matter less in the long-run. Future research in this area would be valuable.

References

Abadie, A. and G. W. Imbens (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica 74*(1), 235–267.

Abadie, A. and G. W. Imbens (2008). On the failure of the bootstrap for matching estimators. *Econometrica 76*(6), 1537–1557.

Abadie, A. and G. W. Imbens (2011). Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics 29*(1), 1–11.

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

Allcott, H. (2011). Rethinking real-time electricity pricing. *Resource and Energy Economics 33*(4), 820–842.

Allcott, H. and N. Wozny (2014). Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics 96*(5), 779–795.

Anderson, S. T., R. Kellogg, and J. M. Sallee (2013). What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management 66*(3), 383–403.

Baylis, K., D. Fullerton, and D. H. Karney (2014). Negative leakage. *Journal of the Association of Environmental and Resource Economists 1*(1/2), 51–73.

Becker, G. S., M. Grossman, and K. M. Murphy (1994). An empirical analysis of cigarette addiction. *American Economic Review*, 396–418.

Bernstein, M. A. and J. Griffin (2005). Regional differences in the price-elasticity of demand for energy. *RAND Corporation Technical Report*.

Blázquez, L., N. Boogen, and M. Filippini (2013). Residential electricity demand in Spain: new empirical evidence using aggregate data. *Energy Economics 36*, 648–657.

Bushnell, J. B., S. P. Holland, J. E. Hughes, and C. R. Knittel (2017). Strategic policy choice in state-level regulation: The epa's clean power plan. *American Economic Journal: Economic Policy*. forthcoming.

Cook, A. D. (2016a). Board orders negotiation in auction disagreement. *RTO Insider*. https://www.rtoinsider.com/miso-board-capacity-auction-imm-26870/, Accessed March 13, 2017.

Cook, A. D. (2016b). MISO considering changes to proposed auction design. *RTO Insider*. https://www.rtoinsider.com/miso-auction-design-imm-26062/, Accessed March 13, 2017.

Cullen, J. A. (2014). Dynamic response to environmental regulation in the electricity industry. mimeo.

DeVirgilio, B. (2006). A discussion of the deregulation of the energy industry in Illinois and its effects on consumers. *Loyola Consumer Law Review 19*, 256.

Elliott, J. and D. Fullerton (2014). Can a unilateral carbon tax reduce emissions elsewhere? *Resource and Energy Economics 36*(1), 6–21.

Exelon (2017). Our companies: ComEd. http://www.exeloncorp.com/companies/comed, Accessed: 2017-02-16.

Filippini, M., B. Hirl, and G. Masiero (2015). Rational habits in residential electricity demand. *Working Paper*.

Ford, M. A. (2013). Five percent of Normal residents opted out of electric aggregation. *Pantagraph*. http://www.pantagraph.com/news/local/five-percent-of-normal-residents-opted-out-of-electric-aggregation/article_48bcb79e-64f2-11e2-acb9-0019bb2963f4.html, Accessed March 13, 2017.

Fowlie, M., S. P. Holland, and E. T. Mansur (2012). What do emissions markets deliver and to whom? Evidence from Southern California's NOx trading program. *American Economic Review 102*(2), 965–993.

Gruber, J. and B. Köszegi (2001). Is addiction "rational"? Theory and evidence. *The Quarterly Journal of Economics 116*(4), 1261–1303.

Hughes, J. E., C. R. Knittell, and D. Sperling (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal 29*(1), 113–134.

Illinois General Assembly (1997). Electric service customer choice and rate relief law of 1997 (220 ilcs 5/article xvi).

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

Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review 104*(4), 1417–1438.

Karnitschnig, M. (August 26, 2014). Germany's expensive gamble on renewable energy. *The Wall Street Journal*.

Kumbaroğlu, G., R. Madlener, and M. Demirel (2008). A real options evaluation model for the diffusion prospects of new renewable power generation technologies. *Energy Economics 30*(4), 1882–1908.

Lotus, J. (2011). Oak Park seals green deal on electric. *OakPark.com*. http://www.oakpark.com/News/Articles/10-18-2011/Oak-Park-seals-green-deal-on-electric/, Accessed March 13, 2017.

Malani, A. and J. Reif (2015). Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform. *Journal of Public Economics 124*, 1–17.

McKibbin, W. J., A. C. Morris, and P. J. Wilcoxen (2014). Pricing carbon in the US: A model-based analysis of power-sector-only approaches. *Resource and Energy Economics 36*(1), 130–150.

Morey, M. J. and L. D. Kirsch (2016). Retail choice in electricity: What have we learned in 20 years? Technical report. Christensen Associates Energy Consulting, LLC.

Myers, E. (2016). Are home buyers myopic? Evidence from housing sales. E2e Working Paper 024.

Olsina, F., F. Garcés, and H.-J. Haubrich (2006). Modeling long-term dynamics of electricity markets. *Energy Policy 34*(12), 1411–1433.

Paul, A. C., E. C. Myers, and K. L. Palmer (2009). A partial adjustment model of us electricity demand by region, season, and sector. RFF Discussion Paper 08-50.

Politis, D. N. and J. P. Romano (1994). Large sample confidence regions based on subsamples under minimal assumptions. *The Annals of Statistics*, 2031–2050.

Poterba, J. M. (1984). Tax subsidies to owner-occupied housing: An asset-market approach. *The Quarterly Journal of Economics 99*(4), 729–52.

Reiss, P. C. and M. W. White (2005). Household electricity demand, revisited. *The Review of Economic Studies 72*(3), 853–883.

Ryoo, J. and S. Rosen (2004). The engineering labor market. *Journal of Political Economy 112*(S1), S110–S140.

Salanie, B. (2011). *The Economics of Taxation*. MIT press.

Spark Energy (2011, July). The history of electricity deregulation in Illinois. http://www.sparkenergy.com/blog/illinois-electricity-deregulation-history/.

Takashima, R., M. Goto, H. Kimura, and H. Madarame (2008). Entry into the electricity market: Uncertainty, competition, and mothballing options. *Energy Economics 30*(4), 1809–1830.

Topel, R. and S. Rosen (1988). Housing investment in the united states. *Journal of Political Economy 96*(4), 718–740.

United States Census Bureau (2016). Quickfacts. https://www.census.gov/quickfacts/.

U.S. Energy Information Administration (2017a). Average retail price of electricity to ultimate customers. https://www.eia.gov/electricity/data.cfm, Accessed February 2, 2017.

U.S. Energy Information Administration (2017b). February 2017 monthly energy review. http://www.eia.gov/totalenergy/data/monthly/#electricity, Accessed March 3, 2017.

Wade, P. (2012). Urbana residents to get opt-out letters on electric rates soon. *The News-Gazette*. http://www.news-gazette.com/news/local/2012-06-15/urbana-residents-get-opt-out-letters-electric-rates-soon.html.

Weyl, E. G. and M. Fabinger (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy 121*(3), 528–583.

Wing, I. S. and M. Kolodziej (2009). The regional greenhouse gas initiative: Emission leakage and the effectiveness of interstate border adjustments. Technical report. Available at SSRN: https://ssrn.com/abstract=1448748.

Wolak, F. A. (2011). Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment. *American Economic Review: Papers and Proceedings 101*(3), 83–87.

A Appendix (For Online Publication Only)

A.1 Conceptual Framework Derivations

As shown in Becker et al. (1994), the effect of a price change on consumption at a particular point in time depends on whether or not the change was anticipated; when the change occurred; and whether the change is temporary or permanent. This can be shown by solving the second-order difference equation (3):

$$y_t = K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} h(t+s) + K_2 \sum_{s=0}^{t-1} (\lambda_2)^s h(t-s) + (\lambda_2)^t \left(y_0 - K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} h(s) \right) \tag{10}$$

where $h(t) = \alpha_3 p_{t-1}$ and

$$\begin{aligned} K_1 &= \frac{\lambda_1}{\alpha_2 (\lambda_1 - \lambda_2)} \\ K_2 &= \frac{\lambda_2}{\alpha_2 (\lambda_1 - \lambda_2)} \end{aligned}$$

with roots

$$\begin{aligned} \lambda_1 &= \frac{2\alpha_1}{1 - \sqrt{1 - 4\alpha_1\alpha_2}} > 1 \\ \lambda_2 &= \frac{2\alpha_1}{1 + \sqrt{1 - 4\alpha_1\alpha_2}} < 1 \end{aligned}$$

We assume $4\alpha_1\alpha_2 < 1$, so that our solutions are real-valued.

Equation (10) shows that consumption in period t is a function of all future prices, all past prices, and the initial condition y_0. In long-run equilibrium ($t \to \infty$), the third term in equation (10) becomes zero, so that consumption no longer depends on the initial condition, y_0. It is straightforward to show that the solution to the first-order difference equation (4), the myopic "adjustment cost" model, depends only on past prices, and not on future prices.

A.2 Data Processing Details

In the usage data provided by ComEd, several communities change definitions over time, moving customers from one community to another or creating a new community. This appears as large, discrete changes in our community-level aggregate usage data. To eliminate this noise, we apply two filters to search for large structural breaks. For each community, we run 89 separate regressions of log usage on month dummies and a structural break indicator, where we start the structural break indicator at each month in the sample. We then compare the maximum R-squared to the minimum R-squared among a community's set of regressions. If this difference exceeds 0.5, then it is dropped from the sample.

For the second filter, we run a series of similar regressions with the addition of a linear time trend. For this filter, we drop any communities for which the explanatory power of the break increases the R-squared by more than 0.2.

One concern with this filter is that we may eliminate actual structural breaks arising from our policy of interest. The communities that are removed in this fashion are primarily small communities that did not implement aggregation. Further, the coefficient on the structural break indicator implies an unrealistic response to the price change.

A.3 Event Study Difference-in-Differences Estimates

In this section, we describe how we estimate the effect of implementing aggregation on electricity prices and usage using a standard difference-in-differences model:

$$Y_{cmy} = \sum_{\tau=-24,\tau\neq-1}^{24} \beta_\tau A_{c\tau} + \beta_{25} A_{c,25} + \beta_{-25} A_{c,-25} + \alpha_{cm} + \alpha_{my} + \varepsilon_{cmy}, \tag{11}$$

where Y_{cmy} is either the natural logarithm of the monthly price or the natural logarithm of total monthly electricity use in community c in calendar month m and year y. The main parameter of interest is β_τ. The variable $A_{c\tau}$ is an indicator equal to 1 if, as of month m and year y, community c implemented aggregation τ months ago. The month before aggregation implementation ($\tau = -1$) is the omitted category. To ensure that our estimated coefficients are relative to this category, we include indicators for aggregation having been implemented 25 or more months ago ($A_{c,25}$) and for aggregation being implemented 25 or more months in the future ($A_{c,-25}$). We include a full set of month-by-year (α_{my}) and community-by-month (α_{cm}) fixed effects and cluster standard errors at the community level. We discuss the robustness of our estimates to different sets of fixed effects in Section 5.

We also estimate a second, more parametric specification that assesses the effect by six-month periods and uses the entire two years prior to aggregation as the reference period:

$$\begin{aligned} Y_{cmy} = {} & \gamma_1 A_{c,0\text{ to }6} + \gamma_2 A_{c,7\text{ to }12} + \gamma_3 A_{c,13\text{ to }18} + \gamma_4 A_{c,19\text{ to }24} \\ & + \beta_{25} A_{c,25} + \beta_{-25} A_{c,-25} + \alpha_{cm} + \alpha_{my} + \varepsilon_{cmy}. \end{aligned} \tag{12}$$

In this specification, $A_{c,0\text{ to }6}$ is an indicator variable equal to 1 if the community implemented aggregation in the past 6 months and 0 otherwise. Similarly, $A_{c,7\text{ to }12}$ is an indicator equal to 1 if the community implemented aggregation between 7 and 12 months ago, and so on. The other variables are defined as in equation (11).

One could use this framework to estimate the effect of implementing aggregation by comparing communities that implemented aggregation to those that did not implement aggregation. However, this raises the concern that communities that did not adopt aggregation may not serve as adequate counterfactual for communities that did adopt aggregation. That is, the decision to adopt aggregation may be correlated with future energy usage. We therefore restrict our estimation sample to communities that implemented aggregation. Our main identifying assumption for these estimates is that, conditional on a host of fixed effects, the timing of aggregation adoption is exogenous with respect to electricity use.

Figure A.4 presents the change in electricity prices following aggregation, in logs, as estimated by equation (11). Similar to our matching results, prices do not drop immediately following the referendum because it takes time for communities to switch to a new supplier. Unlike the matching estimator, the pre-period change is not exactly equal to zero in the event-study difference-in-difference. Although treatment and control communities face identical prices in the pre-period in *calendar time*, they do not face identical prices in *event-study time* because ComEd's prices fluctuate month-to-month. This distinction does not matter for the matching estimator, which creates counterfactuals separately for each treated community. The second vertical dashed line in Figure A.4 shows the point at which half of all communities have implemented aggregation (4 months after passing the referendum). Prices continue to drop as more communities switch and then eventually stabilize. Within 8 months of passing the referendum, the average electricity price has decreased by more than 0.3 log points (26 percent) in aggregation communities relative to the control group. There is an increase in the relative aggregation price 28 months after passing aggregation, which is due to the fact that electricity prices fell sharply for ComEd customers in June of 2013 (see Figure 3), the middle of our sample period. Despite this increase, prices in aggregation communities remain significantly lower than those in the control group for the entire sample period.

Figure A.5 shows the corresponding estimates for electricity usage. Prior the referendum, the difference in usage between aggregation and the control communities is statistically indistinguishable from zero. Usage in aggregation communities then begins to increase following the referendum. By the end of the first year, usage in aggregation communities is about 0.1 log points (9.5 percent) higher relative to the counterfactual.

Table A.3 shows the estimated impact of aggregation on the log of the electricity price in these communities 0-6, 7-12, 13-18, and 19-24 months after implementation, as estimated by equation (12). Overall, the results consistently show large and significant price drops. Our preferred specification is presented in Column 4 and includes community-by-month and month-by-year fixed effects. This specification estimates that electricity prices fell by 0.1 log points in the first six months, and eventually stabilizes at around 0.3 log points by the end of the first year. These estimates are robust to including different fixed effects.

Table A.4 shows the estimated change in usage as estimated by equation (12) for the sample of communities that implemented aggregation. Our preferred specification, presented in Column 4, estimates that electricity usage is 0.048 log points higher in the first 6 months following the referendum, and this increases to 0.114 log points within one year.

Finally, Figure A.6 shows the elasticities implied by the two preceding tables. Specifically, we show the ratio of coefficients from Tables A.4 and A.3, which estimate the aggregation-induced change in electricity quantities and prices, respectively. Because the outcomes are in logs, their ratio will be approximately equal to the elasticity. The implied elasticity ranges from -0.33 7-12 months after passage of aggregation to -0.45 two and a half years after passage.

A.4 Municipal Aggregation Materials

After the proposed aggregation program has been registered with the state, the municipality must hold a referendum. The wording of the referendum question is specified in the Illinois Power Agency Act:[25]

> The election authority must submit the question in substantially the following form:
>
> Shall the (municipality, township, or county in which the question is being voted upon) have the authority to arrange for the supply of electricity for its residential and small commercial retail customers who have not opted out of such program?
>
> The election authority must record the votes as "Yes" or "No".

Figure A.7 displays an example of a letter sent to residents of a community following the passage of an aggregation referendum and selection of a new aggregation supplier. The letter informs residents about their new supply price for electricity, and lets them know that they will have an opportunity to opt out of aggregation. Figure A.8 displays an example of the opt-out card that a customer must fill out and mail if they wish to retain their current electricity supplier.

Figures A.9 and A.10 display the front and back page of a typical electricity bill for a customer residing in ComEd's service territory. If a customer switches suppliers, e.g., her community adopts aggregation and she does not opt out, then the Electricity Supply Charge rate (see Figure A.10) will change. Otherwise her bill will remain the same.

[25]From 20 ILCS 3855/1-92, Text of Section from P.A. 98-404. Available from http://www.ilga.gov/legislation/ilcs/fulltext.asp?DocName=002038550K1-92.

Appendix Tables

Table A.1: Matching Estimates of the Effect of Aggregation on Usage and Prices, Monthly

	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
Month 3	0.014***	-0.063***	-0.061**	286	286
	(0.005)	(0.007)	(0.037)		
Month 4	0.020***	-0.114***	-0.081***	278	278
	(0.006)	(0.007)	(0.032)		
Month 5	0.020***	-0.187***	-0.095***	278	278
	(0.006)	(0.007)	(0.028)		
Month 6	0.025***	-0.224***	-0.107***	278	278
	(0.007)	(0.005)	(0.027)		
Month 7	0.032***	-0.240***	-0.094***	278	278
	(0.008)	(0.010)	(0.025)		
Month 8	0.041***	-0.262***	-0.114***	278	278
	(0.008)	(0.008)	(0.020)		
Month 9	0.057***	-0.257***	-0.175***	278	278
	(0.008)	(0.007)	(0.024)		
Month 10	0.055***	-0.243***	-0.182***	278	278
	(0.009)	(0.007)	(0.028)		
Month 11	0.059***	-0.272***	-0.170***	278	278
	(0.008)	(0.008)	(0.023)		
Month 12	0.054***	-0.222***	-0.227***	278	278
	(0.009)	(0.006)	(0.032)		
Month 13	0.057***	-0.222***	-0.236***	278	278
	(0.009)	(0.006)	(0.033)		
Month 14	0.045***	-0.228***	-0.161***	278	277
	(0.008)	(0.005)	(0.026)		
Month 15	0.037***	-0.050***	-0.418***	240	240
	(0.007)	(0.003)	(0.097)		
Month 16	0.038***	-0.110***	-0.321***	240	240
	(0.007)	(0.002)	(0.061)		
Month 17	0.045***	-0.119***	-0.361***	240	240
	(0.007)	(0.002)	(0.058)		
Month 18	0.033***	-0.128***	-0.220***	240	240
	(0.008)	(0.003)	(0.058)		
Month 19	0.036***	-0.140***	-0.232***	240	240
	(0.008)	(0.004)	(0.053)		
Month 20	0.036***	-0.135***	-0.248***	183	183
	(0.008)	(0.004)	(0.058)		
Month 21	0.047***	-0.132***	-0.325***	183	183
	(0.008)	(0.003)	(0.055)		
Month 22	0.035***	-0.133***	-0.246***	183	183
	(0.008)	(0.004)	(0.055)		
Month 23	0.040***	-0.125***	-0.309***	183	183
	(0.007)	(0.003)	(0.057)		
Month 24	0.040***	-0.125***	-0.308***	183	183
	(0.007)	(0.003)	(0.056)		
Month 25	0.039***	-0.121***	-0.327***	183	182
	(0.007)	(0.003)	(0.058)		
Month 26	0.040***	-0.097***	-0.290***	183	182
	(0.008)	(0.005)	(0.062)		
Month 27	0.046***	-0.166***	-0.236***	183	183
	(0.008)	(0.006)	(0.038)		

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each MEA town is matched to the five non-MEA towns with the most similar usage in 2008 and 2009. The number of price observations corresponds to the number of observations for each elasticity estimate, as we always observe usage where we observe a price change. Standard errors are in parentheses. Significance is determined by subsampling to construct confidence intervals.

Table A.2: Comparison of Tuning Parameters for Subsampling

R	B_1	Type	Months 1-6	Months 7-12	Months 13-18	Months 19-24	Months 25-30
		Point Estimate	-0.0939	-0.1550	-0.2280	-0.2723	-0.2748
1	17	Standard Error	0.0208	0.0221	0.0283	0.0476	0.0444
2	34	Standard Error	0.0197	0.0204	0.0275	0.0471	0.0430
3	51	Standard Error	0.0190	0.0199	0.0265	0.0430	0.0386
5	85	Standard Error	0.0176	0.0185	0.0242	0.0388	0.0352
7	119	Standard Error	0.0169	0.0158	0.0217	0.0364	0.0334

Results from our bi-annual elasticity estimates are reported above. The first row reports the point estimates. The remaining rows report the standard errors calculated via subsampling with different values of the tuning parameter, R, and the corresponding subsample size in terms of treated communities, B_1. Confidence intervals throughout the paper are calculated with $R = 3$.

Table A.3: Effect of Aggregation on Electricity Prices, Communities that Passed Aggregation

	(1)	(2)	(3)	(4)
0-6 Months Post-Aggregation	-0.119***	-0.100***	-0.123***	-0.101***
	(0.005)	(0.005)	(0.005)	(0.005)
7-12 Months Post-Aggregation	-0.307***	-0.313***	-0.312***	-0.320***
	(0.007)	(0.007)	(0.007)	(0.007)
13-18 Months Post-Aggregation	-0.297***	-0.265***	-0.303***	-0.267***
	(0.008)	(0.009)	(0.008)	(0.010)
19-24 Months Post-Aggregation	-0.283***	-0.285***	-0.285***	-0.287***
	(0.010)	(0.013)	(0.010)	(0.013)
25-30 Months Post-Aggregation	-0.281***	-0.264***	-0.296***	-0.279***
	(0.013)	(0.017)	(0.014)	(0.018)
Community Fixed Effects	X	X		
Month and Year Fixed Effects	X		X	
Month-by-Year Fixed Effects		X		X
Community-by-Month Fixed Effects			X	X
Dep. Var. Mean	2.202	2.202	2.202	2.202
Observations	25,716	25,716	25,716	25,716
Adjusted R-squared	0.793	0.898	0.802	0.907

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Standard errors (in parentheses) clustered by community. Outcome variable is the log of the per-kWh electricity price.

Table A.4: Effect of Aggregation on Electricity Usage, Communities that Passed Aggregation

	(1)	(2)	(3)	(4)
0-6 Months Post-Aggregation	0.073***	0.059***	0.066***	0.048***
	(0.008)	(0.009)	(0.005)	(0.006)
7-12 Months Post-Aggregation	0.054***	0.095***	0.065***	0.114***
	(0.012)	(0.016)	(0.012)	(0.016)
13-18 Months Post-Aggregation	0.107***	0.140***	0.088***	0.114***
	(0.015)	(0.019)	(0.014)	(0.017)
19-24 Months Post-Aggregation	0.084***	0.073***	0.109***	0.114***
	(0.016)	(0.023)	(0.015)	(0.021)
25-30 Months Post-Aggregation	0.067***	0.139***	0.067***	0.133***
	(0.020)	(0.025)	(0.020)	(0.024)
Community Fixed Effects	X	X		
Month and Year Fixed Effects	X		X	
Month-by-Year Fixed Effects		X		X
Community-by-Month Fixed Effects			X	X
Dep. Var. Mean	14.371	14.371	14.371	14.371
Observations	25,716	25,716	25,716	25,716
Adjusted R-squared	0.991	0.993	0.996	0.998

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Standard errors (in parentheses) clustered by community. Outcome variable is the log of total electricity usage.

Appendix Figures

Figure A.1: Estimated Price Elasticities, Monthly

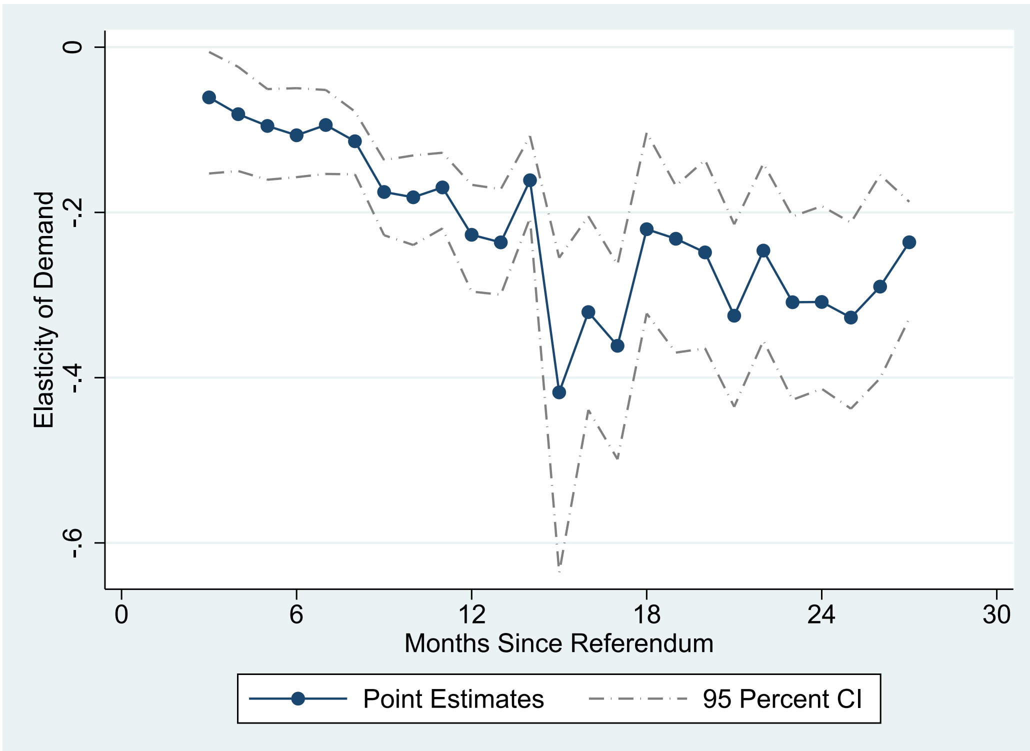

Notes: Elasticities are calculated for each month by regressing community-month changes in log usage on the observed change in log price. Confidence intervals are constructed via subsampling.

Figure A.2: Effect on Log Usage: Communities that Passed but Did Not Implement Aggregation

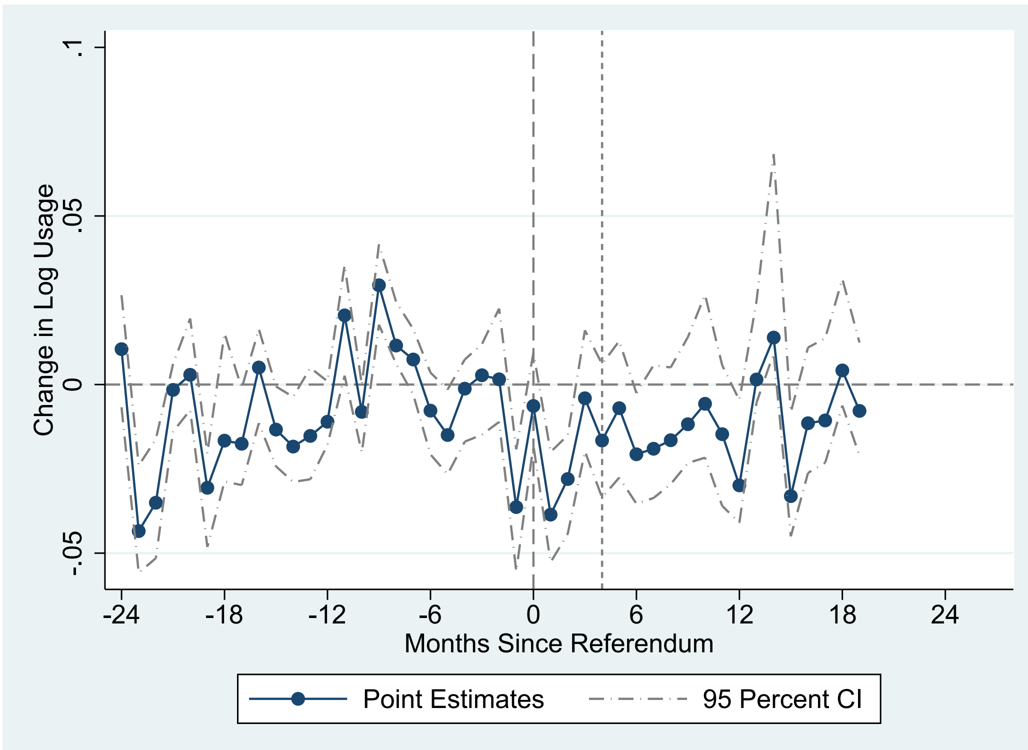

Notes: The figure displays estimates of the mean usage effect for the eleven communities that pass aggregation but never implement it. The effect is estimated relative to that community's five nearest-neighbors, as defined by the difference-in-differences matching procedure outlined in the main text. The short dashed line indicates the median implementation date relative to when the referendum was passed. Confidence intervals are constructed via subsampling.

Figure A.3: Estimated Elasticities and Mean Log Price Change

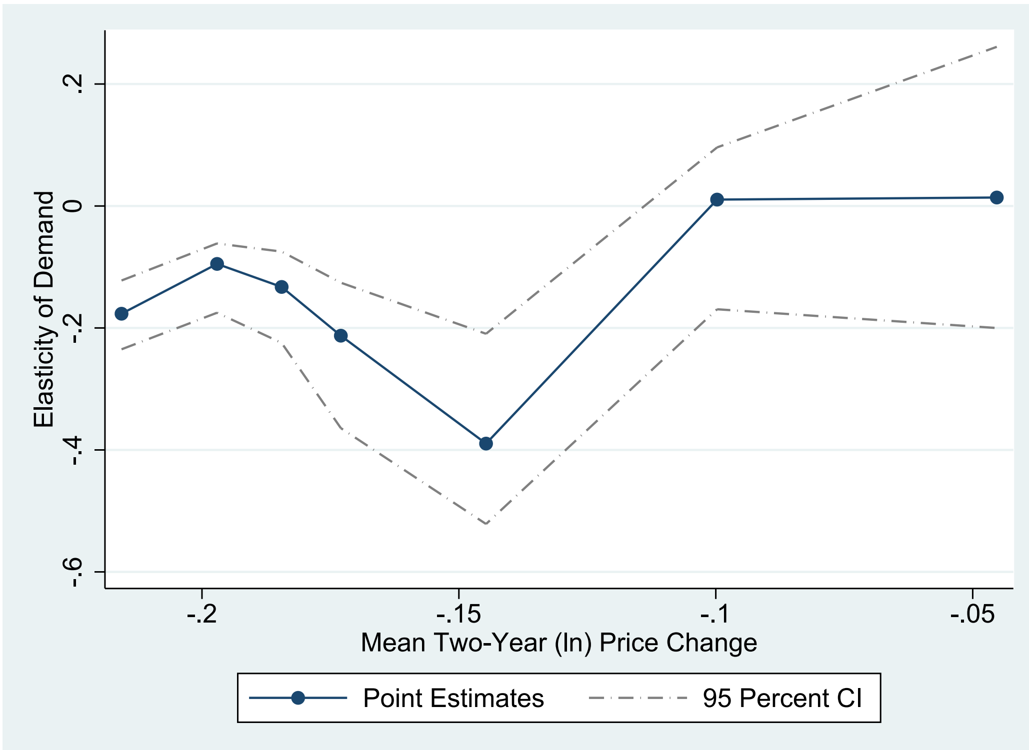

Notes: Communities are split into seven groups based on the average two-year price change. Elasticities are calculated separately for each group. The graph shows no relationship between the estimated group elasticity and the price change, mitigating some concerns about endogeneity. Confidence intervals are constructed via subsampling.

Figure A.4: Regression Estimates of the Effect of Aggregation on Electricity Prices, Communities that Passed Aggregation

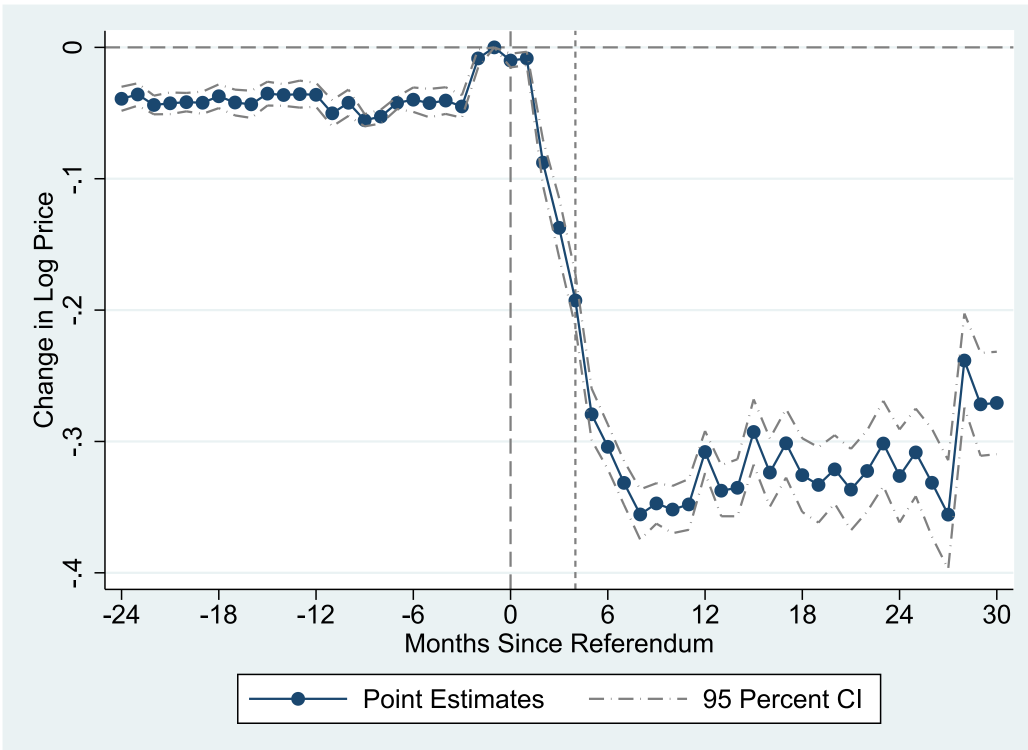

Notes: Outcome is the natural log of the electricity price. The first vertical dashed line indicates the date of the aggregation referendum. The second dashed line indicates the date of aggregation implementation. Regressions include month-by-year and community-by-month fixed effects. Standard errors are clustered by community. Sample includes only communities that passed aggregation at some point during our sample.

Figure A.5: Regression Estimates of the Effect of Aggregation on Electricity Usage, Communities that Passed Aggregation

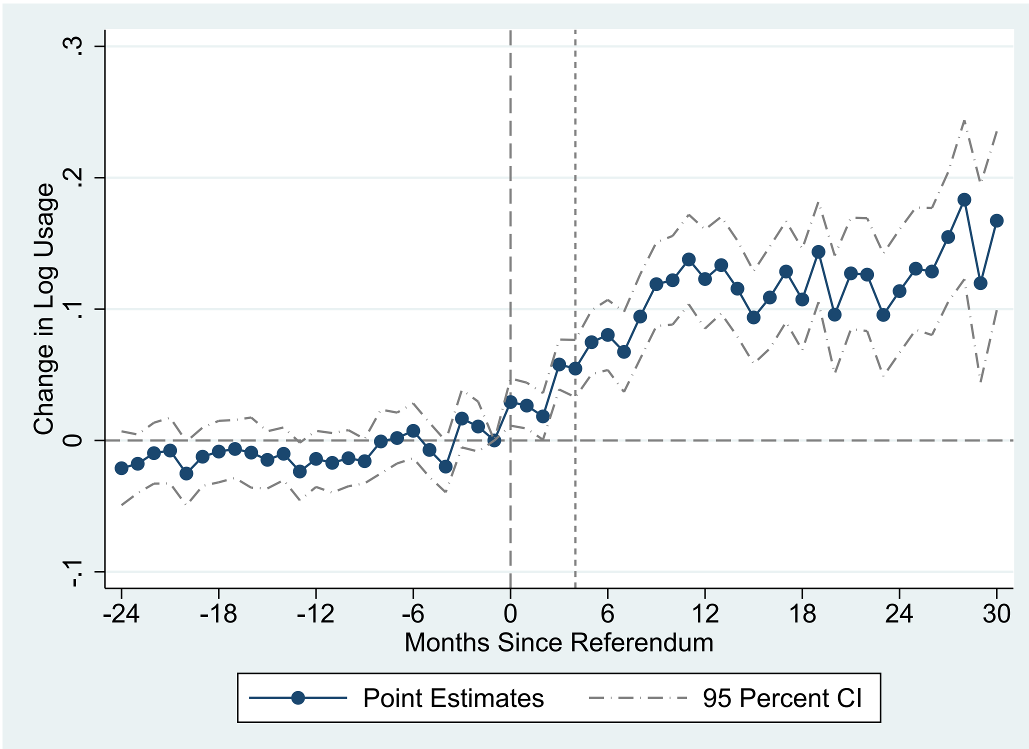

Notes: Outcome is the natural log of total electricity use. The first vertical dashed line indicates the date of the aggregation referendum. The second dashed line indicates the date of aggregation implementation. Regressions include month-by-year and community-by-month fixed effects. Standard errors are clustered by community. Sample includes only communities that passed aggregation at some point during our sample.

Figure A.6: Estimated Price Elasticities, Communities that Passed Aggregation

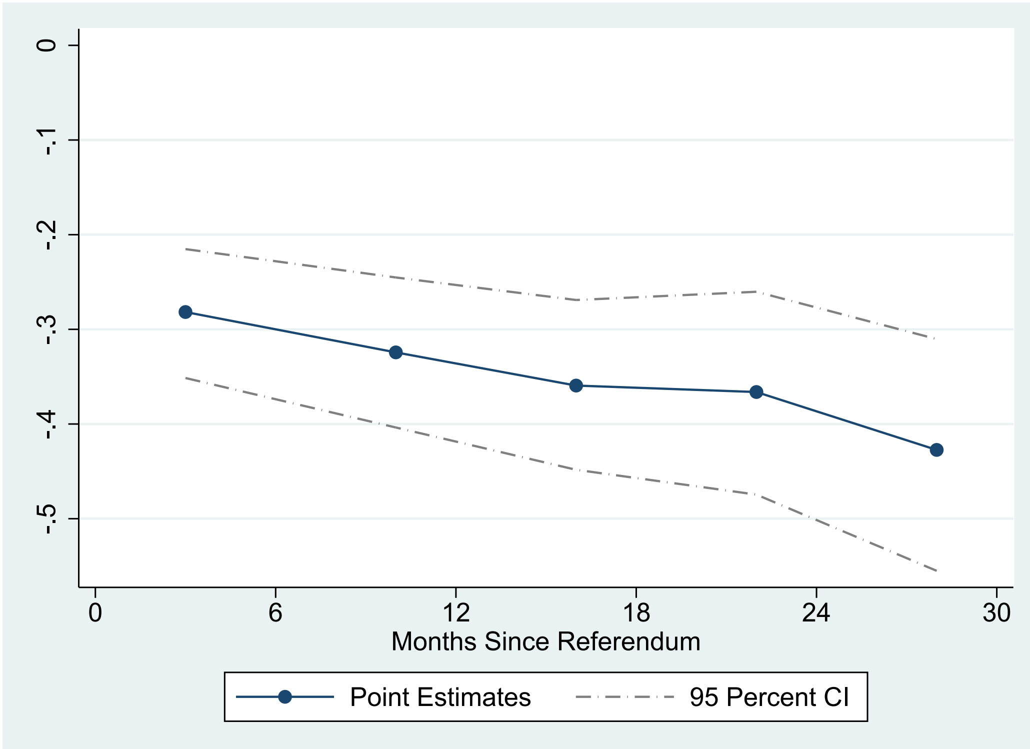

Notes: Sample includes only communities that passed aggregation at some point. Elasticities are calculated for each six-month period by regressing community-month changes in log usage on the observed change in log price. Confidence intervals are constructed by bootstrap.

Figure A.7: Example of an Aggregation Mailing

Kane County
C/O Dynegy Energy Services
1500 Eastport Plaza Dr.
Collinsville, IL 62234

John A. Smith
123 Main St
Anytown, IL 65432

Kane County is pleased to announce that Dynegy Energy Services, LLC ("DES") has been selected as the Supplier for its Municipal Aggregation program. This includes a 24-month program with a fixed price of $**0.06533 per kilowatt hour (kWh)** for the first 12 months (August 2015 to August 2016) and steps down to **$0.06065 per kWh** for the last 12 months (August 2016 to August 2017). DES is an independent seller of power and energy service and is certified as an Alternative Retail Electricity Supplier by the Illinois Commerce Commission (ICC Docket No. 14-0336).

As an eligible residential or small business customer located in unincorporated portions of Kane County, you will be automatically enrolled unless you opt out.

HOW TO OPT-OUT
You need do nothing to receive this new fixed rate. However, if you choose not to participate, simply return the enclosed Opt-Out Card **or call DES at 844-351-7691** by **July 10, 2015.** For more information, visit www.DynegyEnergyServices.com or contact DES Customer Care at 866-694-1262 from 8:00am to 7:00pm Mon- Fri or via email at DESCustCare@Dynegy.com.

There is no enrollment fee, no switching fee, and no early termination fee. This is a firm, fixed all-inclusive rate guaranteed until **August 2017**. This program offers automatic enrollment in Traditionally-sourced Power, but you have an option of purchasing Renewable Power at a rate of $**0.06766 per kWh** for the first 12 months (August 2015 to August 2016) which steps down to $**0.06327 per kWh** for the last 12 months (August 2016 to August 2017).

ENROLLMENT PROCESS
Once your account is enrolled, you will receive a confirmation letter from ComEd confirming your switch to DES. A sample ComEd notice is attached. Approximately 30 to 45 days after enrollment you will receive your first bill with your new DES price. Please review the enclosed Terms and Conditions for additional information.

Please be advised you also have the option to purchase electricity supply from a Retail Electric Supplier (RES) or from ComEd pursuant to Section 16-103 of the Public Utilities Act. Information about your options can be found at the Illinois Commerce Commission website: www.pluginillinois.org and www.ComEd.com. You may request a list of all supply options available to you from the Illinois Power Agency.

See Reverse for Frequently Asked Questions…

Sincerely,

Christopher J. Lauzen
Board Chairman
Kane County

Kurt R. Kojzarek
Development Committee Chairman
Kane County

Figure A.8: Example of an Opt-Out Card

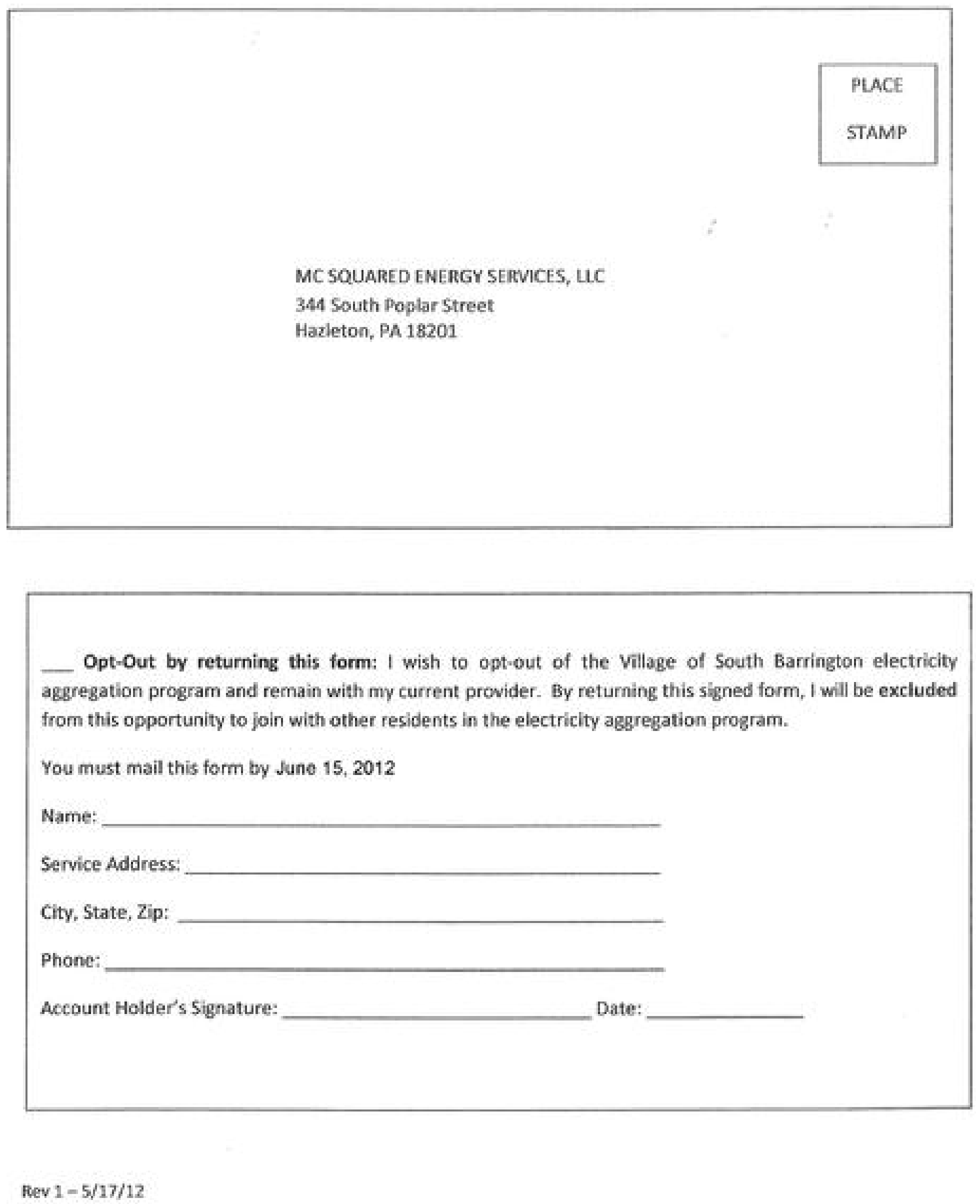

PLACE

STAMP

MC SQUARED ENERGY SERVICES, LLC
344 South Poplar Street
Hazleton, PA 18201

___ **Opt-Out by returning this form:** I wish to opt-out of the Village of South Barrington electricity aggregation program and remain with my current provider. By returning this signed form, I will be **excluded** from this opportunity to join with other residents in the electricity aggregation program.

You must mail this form by June 15, 2012

Name: ____________________

Service Address: ____________________

City, State, Zip: ____________________

Phone: ____________________

Account Holder's Signature: ____________________ Date: __________

Rev 1 – 5/17/12

Figure A.9: Example of a ComEd Bill (page 1 of 2)

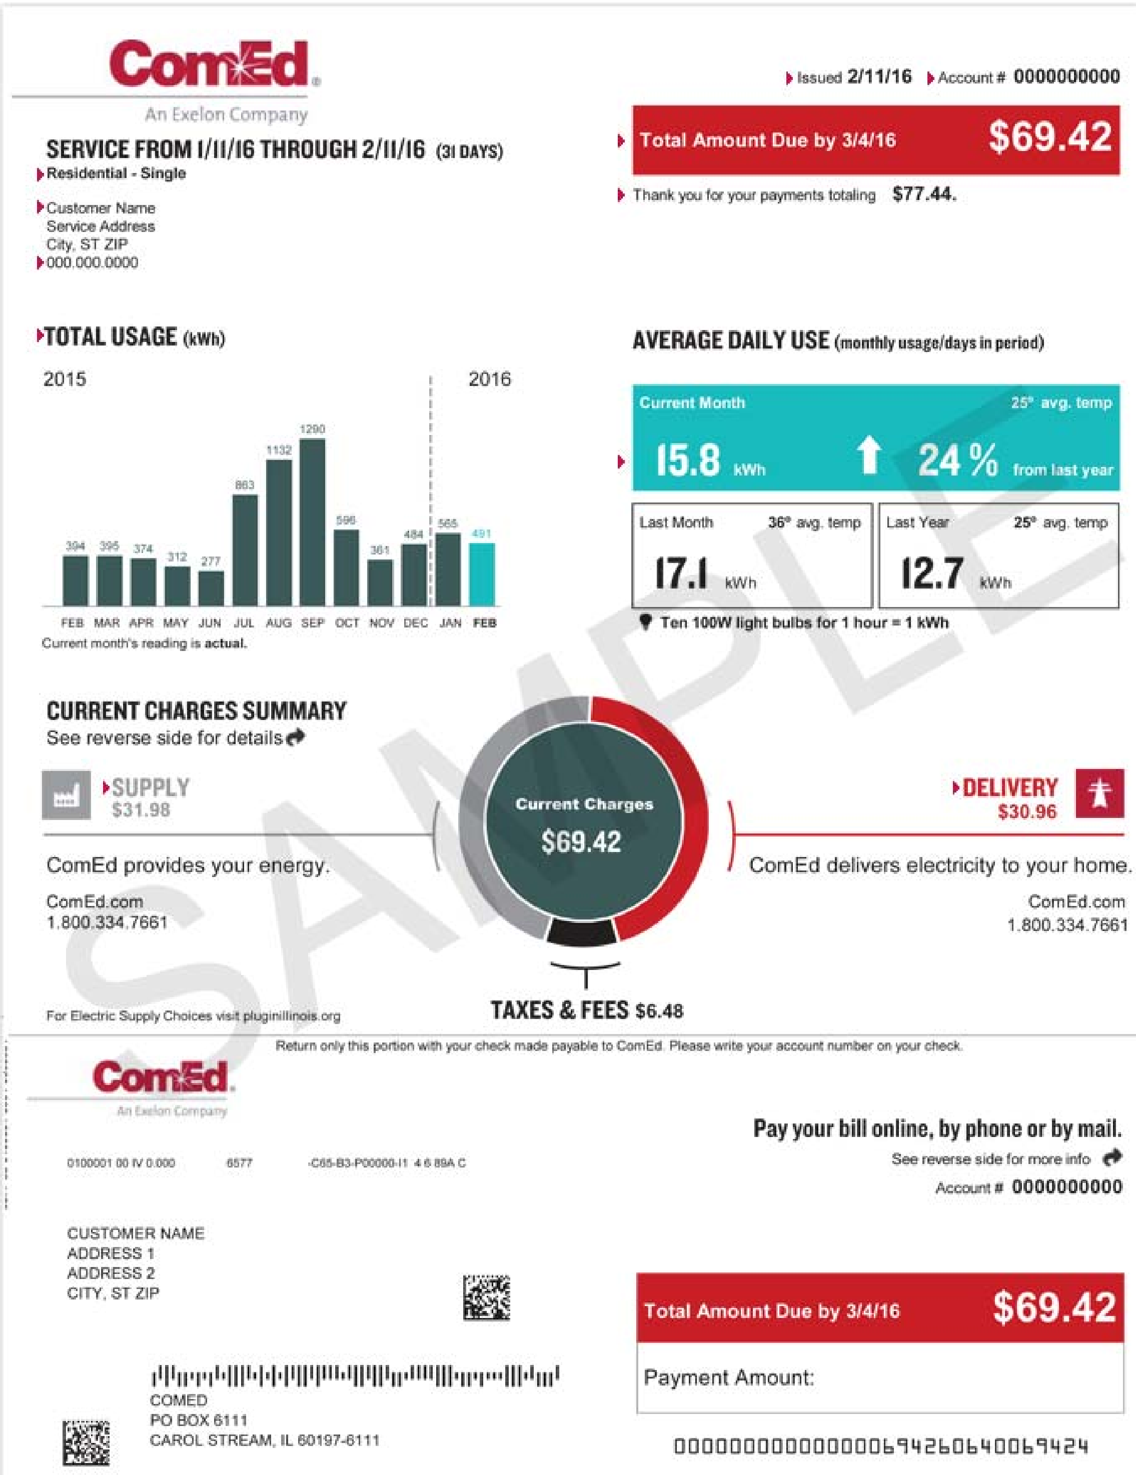
ComEd
An Exelon Company

Issued 2/11/16 Account # 0000000000

SERVICE FROM 1/11/16 THROUGH 2/11/16 (31 DAYS)
Residential - Single

Customer Name
Service Address
City, ST ZIP
000.000.0000

Total Amount Due by 3/4/16 $69.42

Thank you for your payments totaling $77.44.

TOTAL USAGE (kWh)

2015 2016

FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	FEB
394	395	374	312	277	863	1132	1290	596	361	484	565	491

Current month's reading is actual.

AVERAGE DAILY USE (monthly usage/days in period)

Current Month 25° avg. temp
15.8 kWh ↑ 24 % from last year

Last Month 36° avg. temp
17.1 kWh

Last Year 25° avg. temp
12.7 kWh

Ten 100W light bulbs for 1 hour = 1 kWh

CURRENT CHARGES SUMMARY
See reverse side for details

SUPPLY
$31.98
ComEd provides your energy.
ComEd.com
1.800.334.7661

Current Charges
$69.42

DELIVERY
$30.96
ComEd delivers electricity to your home.
ComEd.com
1.800.334.7661

For Electric Supply Choices visit pluginillinois.org

TAXES & FEES $6.48

SAMPLE

Return only this portion with your check made payable to ComEd. Please write your account number on your check.

ComEd
An Exelon Company

0100001 00 IV 0.000 6577 -C65-B3-P00000-I1 4 6 89A C

Pay your bill online, by phone or by mail.
See reverse side for more info
Account # 0000000000

CUSTOMER NAME
ADDRESS 1
ADDRESS 2
CITY, ST ZIP

Total Amount Due by 3/4/16 $69.42

Payment Amount:

COMED
PO BOX 6111
CAROL STREAM, IL 60197-6111

0000000000000000694260640069424

Figure A.10: Example of a ComEd Bill (page 2 of 2)

For Questions, Support, and Outages visit ComEd.com

English	1.800.EDISON1 (1.800.334.7661)
Español	1.800.95.LUCES (1.800.955.8237)
Hearing/Speech Impaired	1.800.572.5789 (TTY)
Federal Video Relay Services (VRS)	Fedvrs.us/session/new

Issued **2/11/16** Account # **0000000000**

Total Amount Due by 3/4/16 $69.42

METER INFORMATION

Read Dates	Meter Number	Load Type	Reading Type	Previous		Present		Difference	Multiplier	Usage
1/11-2/11	000000000	General Service	Total kWh	94278	Actual	94769	Actual	491	x 1	491

CHARGE DETAILS

Residential - Single 1/11/16 - 2/11/16 (31 Days)

SUPPLY		**$31.98**
Electricity Supply Charge	491 kWh X 0.05865	$28.80
Transmission Services Charge	491 kWh X 0.01122	$5.51
Purchased Electricity Adjustment		-$2.33
DELIVERY - ComEd		**$30.96**
Customer Charge		$10.53
Standard Metering Charge		$4.36
Distribution Facilities Charge	491 kWh X 0.03156	$15.50
IL Electricity Distribution Charge	491 kWh X 0.00116	$0.57
TAXES & FEES		**$6.48**
Environmental Cost Recovery Adj	491 kWh X 0.00038	$0.19
Energy Efficiency Programs	491 kWh X 0.00345	$1.69
Franchise Cost	$30.39 X 2.36300%	$0.72
State Tax		$1.62
Municipal Tax		$2.26
Service Period Total		**$69.42**
Thank you for your payment of $77.44 on January 29, 2016		
Total Amount Due		**$69.42**

UPDATES

ComEd

- WAYS TO PAY: ComEd offers a variety of ways to Pay. Learn more at ComEd.com/pay
- WE ARE HERE FOR YOU: As of Feb. 1st, we'll be open 1 hour later. Call us from 7 am to 7pm Monday through Friday. 1-800-EDISON-1 ComEd.com/ContactUs
- APPLIANCE REBATES: Get Appliance Rebates from the ComEd Energy Efficiency Program to upgrade your appliances. ComEd.com/Rebates
- SCAM ALERT: ComEd will never call you to request cash or ask you to buy a prepaid credit card to pay a bill. ComEd.com/ScamAlert
- PART 280: View a copy of the ICC Commission 83 Ill. Adm. Code 280 rules at ComEd.com/Part280
- YOUR COMED BILL: Need help understanding your bill line item definitions? Please visit us at ComEd.com/UnderstandBill or call us at 1-800-334-7661
- ENVIRONMENTAL DISCLOSURE STATEMENT: ComEd's Environmental Disclosure Statement can now be found online at ComEd.com/EnvironmentalDisclosure
- Past due balances are subject to late charges.

6577-65-0100001-0001-0000001

SAMPLE

OTHER WAYS TO PAY YOUR BILL

Visit ComEd.com/PAY for more information including applicable fees for some transactions.

Online

Set up an automatic payment, enroll in paperless billing, or make a convenience payment at ComEd.com/Pay.

Mobile App

Download the ComEd mobile app on your Apple® or Android™ device to view and pay your bill, or manage your account.

Phone

Call us to make a convenience payment with a credit card, ATM card, or your bank account: 1.800.588.9477. (Fee Applies)

In-Person

Pay your bill in-person at many ComEd authorized agents located throughout the region. Visit ComEd.com/Pay for details.

When you provide a check as payment, you authorize us to use information from your check either to make a one-time electronic fund transfer from your account or to process the payment as a check transaction.

10% total recycled fiber