The Long(er)-Run Elasticity of Electricity Demand: Evidence from Municipal Electric Aggregation*

Tatyana Deryugina
University of Illinois
and NBER

Alexander MacKay
Harvard University

Julian Reif
University of Illinois

PRELIMINARY AND INCOMPLETE. NOT FOR GENERAL CIRCULATION.

Abstract

Understanding how consumers respond to electricity prices is essential for predicting the effects of climate change policy and other policies that target electricity markets. To date, studies of long-run electricity demand have relied on price changes that are transient or endogenous, and none have utilized experimental or quasi-experimental variation. We study the dynamics of residential electricity demand by exploiting price variation arising from a natural experiment: an Illinois community aggregation policy that enabled many communities to select new electricity suppliers on behalf of their residents. Participating communities experienced average price decreases in excess of 10 percent in the three years following adoption of a new supplier. Employing a flexible matching approach, we estimate a one-year average 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, which is consistent with a dynamic model of demand. Our results demonstrate the importance of accounting for long-run adjustments when predicting policy effects.

JEL codes: Q41, Q48, D12

*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 Erica Myers, and participants at the 2017 ASSA meetings, the 2017 CeMENT workshop, and the EPIC lunch series 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. Emails: `deryugin@illinois.edu`, `amackay@hbs.edu`, and `jreif@illinois.edu`.

1 Introduction

The threat of climate change has prompted many governments to work toward reducing greenhouse gas emissions. The electricity industry is a natural target because it accounts for a large share of emissions and has relatively few emitters. Electricity prices in some countries, such as Germany, have already increased substantially as a result of climate change policies (Karnitschnig, 2014). Regulation of other pollutants emitted by electricity generators, such as sulfur dioxide, nitrogen dioxide, or mercury, can similarly raise electricity prices. Evaluating the effects and incidence of these policies depends crucially on the price elasticity of electricity demand. For example, the magnitude of an emissions decrease resulting from a carbon tax and the costs of an emissions cap are both functions of this parameter. It is particularly important to obtain an estimate of the long-run price elasticity of demand because most policies targeting the electricity sector are intended to be long-lasting. In the United States, the residential sector accounts for almost 40% of retail electricity sales (U.S. Energy Information Administration, 2017), making its demand elasticity particularly important.

Yet, there is little consensus on the magnitude of the short-run or long-run price elasticity of electricity demand. The absence of consensus is in part due to a lack of quasi-experimental studies. Isolating a source of exogenous variation in prices and finding a suitable comparison group have been significant challenges in this setting. Some components of the price variation are demand-driven and affect a geographic area uniformly, such as an increase in the electricity price due to an unusually hot summer. Further, price changes that are temporary or small, such as those due to short-run changes in fuel costs, may not be salient to consumers. Moreover, in the face of adjustment costs, either through the replacement of durable goods or other psychic costs, the long-run response of consumers may differ significantly from the short-run response.

We provide the first quasi-experimental estimate of the long-run price elasticity of electricity demand by exploiting large, long-lasting, and plausibly exogenous variation in electricity prices arising from Illinois' Municipal Electric Aggregation (MEA) program. In 2009, Illinois passed a law allowing municipalities to choose an electricity supplier on behalf of their residents. Interested communities were required to first pass a local referendum approving the policy. Residents whose electricity was supplied by the incumbent would then be automatically switched to a newly chosen supplier unless they opted out.[1] The first referendum took place in November of 2010, and hundreds of others followed in subsequent years. As of February 2016, 741 Illinois communities had approved aggregation programs, and the vast majority of consumers in these communities switched suppliers and paid significantly lower electricity prices as a result (Plug In Illinois, 2016).

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

Our analysis employs monthly community-level usage data from ComEd, one of the two electricity distributors in Illinois. ComEd services 3.8 million people in 885 Illinois communities, including the city of Chicago (Exelon, 2017). We compare usage and price changes in communities that implemented MEA to those that did not pass MEA to evaluate the consumer response to price changes. Importantly for our analysis, MEA communities only saw a change in the marginal price of electricity. Residents retained the incumbent distributor for billing, keeping the appearance of the bill and distribution charges the same. Illinois also has a simple linear price schedule for residential electricity (i.e., no block pricing), which greatly simplifies our analysis.

Our data span the years 2007-2014, which allows us to estimate trends in electricity over large periods of time both before and after most communities' passage of MEA. The large number of ComEd communities that did not pass MEA (479) in combination with a lengthy pre-period makes our setting an ideal application 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 desirable for electricity markets because usage is highly seasonal, and the seasonal patterns vary substantially across different communities. High levels of idiosyncratic baseline variation (i.e., noise) can present a challenge for the traditional difference-in-differences estimator. A key advantage of our matching estimator is that it can significantly increase precision, as we demonstrate in the main text.

We match each MEA community to five "nearest neighbors" that did not pass MEA on the basis of their electricity usage in 2008 and 2009. The end of our matching window precedes the first implementation of MEA by a year and a half. Our identifying assumption is that average observed differences in usage between MEA towns and their matched controls are due to MEA. Specifically, we assume that MEA implementation and usage in the absence of MEA are uncorrelated, conditional on 2008 and 2009 usage. One implication of this assumption is that communities do not select into the policy based on expected future usage, controlling for past usage. We show that treated and control communities have very similar usage patterns in 2010, after the matching period, but before any community implements MEA. These usage patterns diverge only after communities begin to implement MEA, which supports the plausibility of our identifying assumption. Additionally, we show that usage trends in MEA and non-MEA communities are very similar in the two years prior to the MEA referenda.[2]

We estimate that prices in MEA communities fell by 22 percent (0.25 log points) and that usage increased by 5.1 percent in months 7 through 12 following the MEA referendum, relative to control communities that did not pass MEA. The usage response over this period implies an average price elasticity of -0.16. In the second and third year, the relative price differences shrank to 13 percent

[2]This result is not mechanical: the vast majority of MEA referenda in our sample are held after February of 2012, more than three years after the end of the matching period.

and 10 percent, respectively. This partial price convergence was caused by the expiration of a long-term ComEd contract in June of 2013, which lowered prices among control communities. We find a corresponding (relative) decrease in usage among MEA communities during this time period. Despite employing price variation from two different sources–a large initial price decrease following implementation of MEA, and a modest subsequent (relative) price increase following the long-term contract expiration–our estimates of the price elasticity smoothly fall from -0.14 in the first year to -0.27 in the second year and -0.29 in the third year, demonstrating the importance of long-run dynamics in this setting. Our results show that consumers are significantly more elastic in the long run than the short run, which has important policy implications. For example, our results demonstrate that employing short-run estimates to predict long-run effects of climate policies like carbon taxes will significantly underestimate the degree of consumer response.

We also find that usage in MEA communities began increasing shortly after passage of the referendum, but before the actual price decrease, which occurred four months after the referendum in the median MEA community. We interpret this as evidence consistent with an economic model of forward-looking consumers who have adjustment costs, whether through the purchase of durable, energy-intensive goods or consumption habits.[3] Our results demonstrate the importance of properly accounting for anticipation when estimating responses to price changes and anticipated policies in general.

Then, we consider whether our estimates appear to be converging to a long-run elasticity by explicitly modeling the demand elasticity as a function of past, current, and future prices. We estimate a parametric model with exponential decay and find a long-run elasticity of -0.29, exactly in line with our three-year quasi-experimental estimate. Our parametric model results indicate that consumers approximately reach the long-run response within three years. We demonstrate that expectations matter for such projections. Our baseline analysis assumes that customers expect MEA savings to continue at the average level for contracts that start after June 2014. Alternatively, if we assume customers perfectly predict the realized prices, we estimate a long-run elasticity of -0.39.

Finally, to evaluate the generalizability of our results, we explore how much the response varies across demographic characteristics of the communities in our sample. We find little evidence that the elasticity depends on economic characteristics, but we do find support for some variation in the elasticity along social demographics. This finding suggests that researchers should account for social demographics when making out-of-sample elasticity projections.

The main contribution of our study is a credible, quasi-experimental estimate of the long-run price elasticity of residential electricity demand. A number of papers have attempted to estimate

[3]It is also possible that some consumers were confused about when the price change would occur.

the short-run and long-run price elasticity of residential electricity demand.[4] The estimated elasticities vary widely: from close to zero and insignificant to about -0.9 in the short run (one year or less) and from -0.3 to about -1.1 in the long run. Much of this literature relies on state-level data and dynamic panel models, which include current prices, lagged consumption, and lagged prices as independent variables (see Alberini and Filippini (2011) for a brief review). The consistency of estimates from such models requires strong assumptions about the form serial correlation takes, and Alberini and Filippini (2011) show that the models are particularly sensitive to the exact specification used. Conversely, our approach relies on quasi-experimental variation and makes relatively few assumptions.

In addition, the vast majority of existing literature does not explicitly isolate an exogenous source of variation in electricity prices. 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, these facts do not eliminate the possibility of endogeneity. For example, 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 constructing instruments. Alberini and Filippini (2011) also point out that there also appears to be non-trivial measurement error in state-level average prices, causing attenuation bias in the demand elasticity estimates. By contrast, we exploit quasi-experimental within-state price variation, and provide strong evidence that this variation is not driven by demand-side factors. Importantly, we demonstrate that trends in electricity use in the years leading up to MEA implementation are similar between MEA and non-MEA communities.

Papers outside of the dynamic panel literature typically focus on short-run rather than longer-run elasticities. Using a structural model and exploiting the non-linearity of the electricity price schedule in California, Reiss and White (2005) estimate the average annual elasticity to be about -0.39. Ito (2014) uses quasi-experimental variation to estimate average price elasticities 1-4 months following a price change, finding that they range from -0.07 to -0.09. Neither study estimates the price elasticity over a longer time period.

We also contribute to the literature on residential energy choice. To our knowledge, our study is the first to formally evaluate the effects of community choice aggregation.[5] In Illinois, MEA was introduced following the overwhelming failure of electric supply choice: only 234 residential consumers had switched suppliers after seven years of that program. Research has shown that failing to exploit available savings appears to be a general pattern of consumer behavior in this

[4]There is also a growing literature on 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.

[5]Littlechild (2008) provides a descriptive overview and some basic statistics about municipal aggregation in Ohio.

area, due to a combination of large switching costs, search costs, or lack of information (Giulietti et al., 2014; Hortaçsu et al., 2017). Allowing communities to choose electricity providers on behalf of their residents while preserving consumers' ability to remain with the incumbent essentially turns an opt-in program into an opt-out program, with correspondingly higher participation rates. While determining the welfare benefits of MEA is beyond the scope of this paper, we estimate that MEA saved residential customers in our sample over $566 million through June 2014.

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

2 The Illinois Electricity Market

The provision of generated 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. Illinois has two regulated electricity distributors: Commonwealth Edison Co. ("ComEd") and Ameren Illinois Utilities ("Ameren"). Since 1997, they have been prohibited from owning or profiting from the generation of electricity due to concerns about market power and widespread agreement that, unlike distribution, electricity generation is not a natural monopoly. The price for electricity supplied by Ameren or ComEd is, by law, equal to their procurement cost and does not vary geographically.[6] 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 suppliers (ARES) who would be responsible for supplying (but not delivering) their electricity.[7] 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 made switching unprofitable. In 2006, the state removed some of these barriers and instituted a discount program for switchers. These changes made savings from switching to ARES much larger, but still had little effect on behavior. By 2009, only 234 residential customers had switched electricity providers. By contrast, there were 71,000 small commercial, large commercial,

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

[7]Large commercial and industrial customers have had this ability since the end of 1999.

and industrial customers who had switched (Spark Energy, 2011).

In 2009, the Power Agency Act was amended to allow for Municipal Electric Aggregation (MEA), 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 large. To ensure that individual consumers retained the ability to choose their provider, the amendment requires municipalities to allow individuals to opt out of aggregation.

To implement an opt-out MEA program, municipalities are required to first educate their communities about MEA using local media and community outreach meetings. After the proposed MEA program has been registered with the state, the municipality must hold a referendum.[8] 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 the majority of 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. This usage data, along with electricity futures, are the main factors in each offered price. Importantly, our analysis employs the same community-level usage data, which reduces the likelihood that price changes are affected by unobservable (to us) confounding factors.[9]

The two main ways in which suppliers differentiate themselves are price and environmental friendliness (as measured by the fraction of generation derived from renewable sources). Nearly all communities select the supplier with the lowest price, although environmental preferences occasionally cause 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 the 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 MEA contracts signed by Illinois communities.

Importantly, instead of "opting in" to an ARES, customers in a community that has passed an MEA referendum are automatically switched to the electricity supplier chosen by their community unless they opt out by mailing in a card, calling, or filling out a form online.[10] The few residen-

[8]The wording of the referendum question is specified in the Act and given in the Appendix.

[9]Suppliers may also base their bids on the number of electric space heat customers, which we do not have access to.

[10]While we do not have specific numbers, ComEd and several energy suppliers have told us that the opt-out rate is low. The number of non-MEA customers does, however, grow slowly over time because new residents who move to an MEA community are not defaulted into the MEA program.

tial customers who had already opted into an ARES or into real-time pricing prior to MEA are not switched over to the chosen supplier. MEA officially begins at the conclusion of the opt-out process. From the consumer's point of view, the 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 MEA. Conveniently, this means that the price effects of MEA 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 MEA program, and a sample opt-out card.

Savings from Municipal Electric Aggregation come largely from the timing of switching behavior. Because ARES purchase electricity from the same competitive generation market as ComEd, they cannot offer a permanently lower supply price unless the community has a favorable load profile. However, the availability of ARES allows MEA communities to exploit falls in electricity costs and lock in those savings for a period of time. Indeed, Municipal Electric Aggregation is very popular in Illinois: as of March 2016, 741 Illinois communities had voted to implement MEA.

3 Data

We obtain electricity usage data directly from ComEd, one of the two electricity distributors in Illinois. ComEd 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 for the time period February 2007 to 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 pass a referendum approving MEA but never implement the program. We estimate our model using a balanced panel of monthly usage for the remaining 768 ComEd communities, of which 289 implemented MEA.

As shown in Table 1, 300 communities passed a referendum on MEA during our sample period, and 289 of those communities eventually implemented an MEA program.[11] In addition, 36 communities in the ComEd territory voted on but did not pass MEA. 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

[11]This includes 5 communities that passed a referendum in November of 2014, five months after the end of our usage data.

provision; and (5) the belief that choosing an electricity provider for residents was not an appropriate government function.[12] The geographic locations of communities in our sample are displayed in Figure 1. MEA communities are well-dispersed throughout the ComEd territory, although they are disproportionately represented in the suburbs of the Chicago area.

We constructed the time series of ComEd electricity rates using ComEd ratebooks, which were made available to us by 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 and because we do not observe household-level usage, we assume that the incumbent rate is equal to the non-electric space heating rate, which will be true for the majority of non-MEA customers. Data on MEA referenda dates, MEA supply prices, and MEA 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 MEA referendum and commencement of the MEA program is 4 months.

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. This is appealing for our research study because it eliminates confusion over which "price" consumers might be responding to. This constant marginal price can be broken down into different components. Implementing MEA requires a community to sign a contract with an electricity supplier that specifies a particular supply rate, which is the largest component of the marginal price; non-MEA communities pay instead the default ComEd supply rate. Thus, MEA only affects the marginal price of electricity. All fixed fees and remaining usage rates are nearly identical across the MEA and non-MEA communities in our sample. The average fixed fee for customers residing in ComEd service territories during our sample period is \$12.52. This fee does not vary across communities, and we ignore it in our analysis as we are concerned with the marginal price.[13] Municipal taxes vary across communities, but the variance is small relative to the total rate. For our analysis, we use the median tax across ComEd communities (0.557 cents/kWh).

The thick red and green lines in Figure 2a display ComEd's monthly supply rate and the total of all remaining 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.[14] The blue line in Figure 2a shows the average monthly supply rate for communities that implemented MEA. This line begins June 2011, when the first community implemented MEA. During

[12]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.

[13]There are a few instances where the MEA supplier charges an additional fixed fee, but these are rare and short-lived.

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

our sample period, the average MEA supply rate is always lower than the default ComEd supply rate. However, starting around June 2015, the two rates became very similar and the number of communities choosing to continue MEA decreased. The price variable in our estimating equations is equal to the MEA supply rate plus all other usage rates if the community has implemented MEA; otherwise it is equal to the ComEd supply rate plus all other usage rates.

Figure 2b plots the mean, the 90th percentile, and the 10th percentile of the log difference between the MEA and ComEd supply rates as a function of the number of months since the MEA referendum. This figure shows heterogeneity in the MEA supply rates and that the average difference between these rates and the ComEd supply rate changes over time. The graph also indicates that at least 10 percent of MEA communities implemented the change within three months of the referendum, whereas 10 percent had not implemented a price change six months afterward. We use the referendum date as our base period to conservatively capture anticipation effects that might occur prior to the actual price change.

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

4 Empirical Strategy

4.1 Difference-in-Differences Matching Framework

We select control communities by matching on pre-MEA electricity usage. Our setting is an ideal application for a matching estimator. The large, diverse set of communities available in the control group makes it likely that a nearest-neighbor matching approach will successfully find suitable comparison groups. Additionally, we have enough data in the period before MEA to internally validate 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 MEA 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 treatment communities (i.e., those that implemented MEA), we find the five nearest neighbors by matching on 2008 and 2009 usage from the pool of 479 control communities available in our sample. We use these nearest neighbors to construct counterfactual usage, and we employ the standard difference-in-differences technique to adjust for pre-period differences. The identification assumption is that, conditional on 2008-2009 electricity usage, the passage of MEA is 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 MEA.

A key advantage of the nearest-neighbor approach is that it eliminates comparison communities 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. Filtering out less relevant control communities can therefore greatly increase precision.

Figure 4 provides a demonstration of this benefit. Panel (a) displays seasonally-adjusted usage for treated (MEA) communities and all communities that never passed MEA. Even after accounting for community-specific monthly seasonal patterns, 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). There is a visible increase in the difference beginning in late-2011, which can be attributed to the implementation of MEA, but this difference is quite noisy. The heterogeneity in seasonal patterns poses a challenge for a standard regression that compares treatment 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.

Panels (b) and (d) of Figure 4 show analogous plots for the nearest-neighbor matching estimator that we employ. Panel (d) shows again that the difference in log usage between treatment and (matched) control communities increases beginning in late-2011. This difference 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 (MEA) communities. This allows the matching estimator to generate more precise estimates than the standard difference-in-differences estimator.

To select the five nearest neighbors, we match on both mean usage and seasonal (monthly) fluctuations from 2008 and 2009. We use annual log usage and monthly log deviations from annual usage to construct 13 match variables, taking the average for each measure across 2008 and 2009. 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.

To see whether matching also helps improve the similarity between treatment 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 MEA communities, and 385 out of 479 non-MEA communities.[15] Table 2 reports the mean characteristics of the communities with a successful ACS match. We display results separately for communities that implemented MEA (column 1), all communities that did not pass MEA (column 2), and matched control communities (column 4). Columns 3 and 5 report the p-values for the null hypothesis that the difference between MEA and all non-MEA or between MEA and matched controls is zero, respectively.

Compared to non-MEA communities, communities that implemented MEA 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 MEA and non-MEA communities. After matching on electricity usage, the weighted pool of matched controls has more similar characteristics to the MEA communities. In cases where significant differences persist, the p-values are substantially larger, indicating that matching on usage is also selecting control communities with much more similar socioeconomic characteristics.

4.2 Estimating the Effect on Usage

Our ultimate goal is to estimate how the price elasticity of electricity evolves over time. We do so using a two-stage approach. First, we estimate the effect of the policy change on usage for each MEA community using our matching estimator. Second, we use the observed change in price and the estimated change in usage to estimate elasticities.

Let Y_{it} denote log usage for community i in period t, where $t = 0$ corresponds to the referendum date for treatment communities. 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 MEA 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. $Y_{it}(1)$ is observed for MEA communities and $Y_{it}(0)$ is observed for non-MEA communities. To calculate the effect of MEA, we construct an estimate of untreated usage for MEA 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 MEA (treated) communities in our sample.

We estimate the average treatment effect on the treated, calculated for each period t, where $t = 0$ indicates the date the referendum was passed:

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

[15] MEA 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, 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 treatment community i, we select $M = 5$ nearest neighbors using the procedure discussed above. 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

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

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. For example, if $t = 25$ corresponds to January 2014, then $\hat{\mu}_i^{m(25)} = \frac{1}{2}\left(Y_{i,January}^{2008} + Y_{i,January}^{2009}\right)$ 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 community's nearest neighbors plus the average (seasonal) difference in usage between that community and its neighbors during 2008-2009.

Finally, the difference-in-differences matching estimator identifies the effect relative to the immediate pre-period. 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 (e.g., $N_s = 12$ for monthly data). Our difference-in-differences estimate thus reflects the change in usage between treated and 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 differences-in-differences framework can lead to biased treatment effect estimates when consumers are forward-looking. There is ample evidence from the energy literature that this is a potential concern in our setting. For example, Myers (2016) finds that expected future heating costs appear to be fully capitalized 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 MEA electricity price is announced months before the actual price change takes place. This raises the possibility that consumers may respond to future price changes by, for example, purchasing less energy efficient appliances, 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 MEA was implemented.[16] 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 construct elasticities, we regress community-specific estimates of the change in usage on the observed community-specific price changes. The community-specific measure $\hat{\tau}_{it}^{DID}$ 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 observed electricity price difference between an MEA community and its matched controls. Because electricity rates do not vary cross-sectionally for non-MEA communities, this difference will be exactly zero prior to a community's implementation of MEA and will only reflect differences in marginal prices after MEA has been implemented.

Equation (2) allows us to flexibly construct period-specific estimates for elasticities to 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. Subsampling, like bootstrapping, obtains a distribution of parameter estimates by sampling from the observed data. The procedure is described below.

For each of $N_b = 500$ subsamples, we select without replacement $B_1 = R \cdot \sqrt{N_1}$ treatment 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 corresponds to the pool of control communities. For each subsample, we

[16]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.

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 $\hat{F}^{-1}(0.975)$. We employ $R = 3$ ($B_1 = 51$) for the confidence intervals reported in our main tables, but our confidence intervals are robust to different values of R (results available upon request). Similarly, we calculate elasticity estimates of β_g for each subsample to generate confidence intervals for our elasticities.

5 Results

5.1 Main Quasi-Experimental Results

We begin by showing that electricity prices fell substantially and persistently following the passage of MEA referenda. Figure 5a displays the average change in log prices for the treatment communities in our sample, relative to their matched control communities. The price change is exactly equal to zero in the pre-period because the treatment 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 MEA 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, MEA prices stay at least 10% 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 treatment 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 MEA instances by at least two years (see Table 1).[17] Following the referendum, usage in MEA communities increases and eventually stabilizes at around 0.04 log points. 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 price increases, provides persuasive evidence that we are capturing the causal effect of price changes on demand for electricity.

[17] Specifically, there is virtually no overlap between the matching period and the pre-period estimates in Figure 5b for 270 of the 289 MEA communities.

Figure 6a displays our estimates of the price elasticity of demand, a key parameter of interest for policymakers. To reduce noise, we estimate elasticities at the biannual, rather than monthly level.[18] 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.

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 increase in the magnitude of the price elasticity from about -0.1 in the short run to nearly -0.3 in the long run. The finding that consumers are more elastic in the longer run than in the short run is consistent with several mechanisms, including habit formation, learning, and appliance replacement over time. Our quasi-experimental estimates, including the quadratic and biannual forms, also suggest convergence around 24 months after the policy change, although data limitations preclude us from drawing a definitive conclusion based on these estimates alone. While we do not attempt to distinguish among specific adjustment channels, we do test whether the estimated elasticity is continuing to grow at the end of our sample period or seems to have converged with a more formal model developed in Section 6.

5.2 Anticipation Effects

Economic theory suggests that forward-looking individuals will respond to policies prior to their implementation if those policies can be anticipated and if there is a benefit of responding before the policy is implemented. For example, prior studies have documented that expectations of future policies matter when purchasing durables such as cars or houses or when making human capital investments (e.g., Poterba, 1984; Ryoo and Rosen, 2004). The effect of MEA on electricity usage is an ideal setting for detecting anticipation effects: the implementation of MEA 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.

Figure 7 displays changes in electricity prices and usage relative to the date of MEA *implementation* rather than the date of the MEA *referendum*. The price difference between MEA towns and their matched controls is exactly zero prior to implementation. However, usage begins increasing in the 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 or the subsequent price change announcement served as a signal to consumers that prices would soon

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

decrease, and that at least some consumers were forward looking and increased their usage immediately, perhaps by reducing their purchases of energy efficient appliances or changing their electricity usage habits relative to communities that did not pass MEA. It is also possible that some consumers were confused and thought their electricity prices had already decreased immediately following the referendum, although we should note that all residents were notified by mail of the exact month of the price change.

In order to explicitly isolate the anticipation effect, Figure 8 displays estimates of changes in electricity usage relative to the date of the MEA *referendum* using only pre-implementation data.[19] 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 over 0.01 log points higher 3 months after the referendum and over 0.035 log points higher 5 months after the referendum, confirming the existence of substantial anticipation effects in our sample.

5.3 Elasticities by Demographic Characteristics

We now consider how much variation in the elasticity of response can be explained by demographic characteristics, and we show how the demographics in our sample relate to the demographics of the U.S. as a whole. This type of analysis matters for understanding the distributional effects of policies that affect electricity prices and for 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 that variable:

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

The indicator function $\mathbf{1}[x_i^j > median(x^j)]$ is equal to 1 if the the value of the (time-invariant) variable x^j for town i is above the median of the distribution and 0 otherwise. We estimate this regression using $n = 8$ different characteristics obtained from the ACS, and report our results in Figures 9 and 10. Because the estimation is done jointly, the displayed elasticities for any given characteristic control for the other characteristics.

Figure 9 reports heterogeneity results for variables related to the housing stock. Statistically significant (at the 10 percent level) estimates 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 could be because newer homes are more likely to have technology such as

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

programmable thermostats, which make it easier for consumers to control electricity consumption.

Figure 10 reports heterogeneity results for socioeconomic characteristics. Younger communities have a more elastic response, as do communities with a greater percentage of white people. Perhaps surprisingly, age and race appear to matter for the elasticity of demand more than economic variables such as income and education.

Our elasticity estimates are relatively stable with respect to economic demographics. These results provide some guidance on how elasticity estimates may be extrapolated to populations within the U.S. that have different demographic 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 with market power might offer lower rates to more inelastic customers, as more elastic customers would demand more at the same price and drive up supply costs. Because communities always have ComEd supply as the "outside option" and the ComEd price is equal to the procurement cost, we think this is unlikely. However, we empirically check for this possibility by splitting our treatment 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 11 plots these estimates. As expected, we find no evidence of a relationship between the price change and the *ex post* estimated elasticity.

Our choice to match on January 2008 through December 2009 electricity usage is driven by the tradeoff between allowing for a long enough post-matching period to evaluate pre-trends and matching on recent enough electricity usage to maximize match quality. However, our results are very similar if we match on electricity usage from February 2007 through January 2009 (available upon request).

Finally, we have also estimated the impact of MEA using a standard differences-in-differences event study and the sample of communities that passed MEA at some point. These results, discussed in detail in the Online Appendix, are qualitatively similar. In particular, we also see no evidence of pre-trends, supporting the identification assumption that passage of MEA was not prompted by growth in electricity usage.

6 Long-Run Projections

6.1 Conceptual Framework for Dynamics

Electricity usage does not adjust to price changes instanteously. It takes time to change one's habit of keeping the lights or air conditioning on when away from home. Usage depends on the energy efficiency of durables such as dishwashers, dryers, and air conditioners, which are typically purchased only once every 10 or 20 years. 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 likely exceed the short-run response. Moroever, if consumers are forward-looking, then the presence of adjustment costs may cause them to respond in anticipation of future prior changes.

A simple way to model these 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.[20] 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 rate, and x_t represents consumption of a composite good that is taken as numeraire. The consumer's budget constraint is

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

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 still allowing us to derive analytical solutions to the consumer's problem. 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 literature corresponds to the special case

[20] 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.

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 or not 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:

$$\frac{dy}{dp^*} = \frac{\hat{\alpha}_3}{1 - \hat{\alpha}_1 - \hat{\alpha}_2}.$$

Alternatively, one can estimate consumption as a function of all past and future prices:

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

In this case, $dy/dp^* = \sum\limits_{s=0}^{t-1} \hat{\beta}_{t-s} + \sum\limits_{s=1}^{\infty} \hat{\beta}_{t+s}$. The appendix provides a derivation that demonstrates the equivalence of these two formulas for dy/dp^*.

The framework above demonstrates the importance of accounting for anticipation effects and for modeling consumer expectations appropriately.

6.2 Empirical estimates

One of the limitations of our study (and other existing studies that attempt to estimate long-run demand elasticities) is that our data include only a limited window for which we measure the consumer response. Specifically, in our non-parametric framework, we can only estimate the monthly response with reasonable accuracy up to 27 months beyond implementation. Without additional data or assumptions, it is impossible to evaluate whether the long-run response to price changes is more elastic than what we measure at the end of our sample.

To evaluate whether our quasi-experimental estimates approximately obtain the long-run response, we regress our matching estimates of log usage changes on lags and leads of log price changes. We consider this a flexible, reduced-form approach to constructing a long-run projection. However, this form is compatible with the adjustment cost model developed above. The model we estimate is:

$$\hat{\tau}_{it}^{DID} = \sum_{s=-L_1}^{L_2} \beta_s \cdot \Delta \ln p_{i(t-s)} + \eta_{it}. \tag{5}$$

The number of leads in the regression is equal to L_1, and the number of lags is equal to L_2. We employ a parametric approach in which we constrain the parameter estimates of the regression to fit a four-parameter adjustment curve. We assume an exponential decay for the impact of lags and a linear response to leads. Estimating a parametric response to lags with an exponential decays allows us to estimate the long-run effect. The curve has the following form:

$$\beta_s = (\alpha_1 - \frac{\alpha_1}{\alpha_2} s \cdot \mathbf{1}[s \leq 0]) \cdot \mathbf{1}[s \geq \alpha_2] + \alpha_3(1 - \exp(\alpha_4 s)) \cdot \mathbf{1}[s >= 0] \tag{6}$$

In this form, α_1 is the elasticity corresponding to a contemporaneous, anticipated price shock, and $\alpha_1 + \alpha_3$ is the long-run elasticity. α_2 and α_4 are parameters governing the speed of adjustment.

Figure 12a plots the estimated adjustment curve. To verify the fit, we also plot the point estimates for non-parametric versions of (5). For the non-parametric regressions, we can only estimate 24 parameters accurately. We include the results for four different sets of leads and lags, starting with 18-month leads and ending with a specification with lags only. To account for compositional effects in each point estimate, we run specifications with community fixed effects and include those as well. The parametric model and the non-parametric estimates produce similar results.

Table 5 displays the estimates of our dynamic elasticity curve. Our preferred specification, which applies status quo expectations to contracts starting after June 2014, is the first column. The resulting long-run estimate is -0.287, which is close to our three-year quasi-experimental estimate of -0.285. For 25-30 months beyond the price change, the adjustment curve predicts an elasticity of -0.278, compared to -0.275 for our quasi-experimental estimate. On the other hand, the parametric estimate for 7-12 months after a price change is -0.225, which is greater in magnitude than our quasi-experimental estimate of -0.165. Properly accounting for expectations matters, and expectations may account for differences between these estimates. In our quasi-experimental setting, we exploit the fact customers do not anticipate any price change in advance of the referendum date. For the adjustment curve, we measure the impact for customers that have a high degree of sophistication in terms of expectations about prices several months ahead. Thus, the methodology in this section does not account for the element of surprise that came with the referendum. The anticipation impact four months prior, which is the median time between the referendum and implementation, is -0.067. If we net out this anticipation impact, the parametric prediction for 7-12 months becomes -0.157, which is closer to our quasi-experimental finding. If the parametric model is correct for expected shocks, and if consumers are able to speed up adjustment in response to unexpected shocks, we would expect our quasi-experimental estimator to find a coefficient between -0.225 and -0.157. Our quasi-experimental estimate of -0.165 lies in this range.

Expectations As mentioned above, the adjustment curve depends on properly accounting for expectations about future prices. For our projections, we assume that consumer expectations align with realized prices for through June 2014, which is the end of our usage panel. For future prices, we assume that customers accurately forecast the price difference between the MEA rate and the ComEd rate for the duration of the existing fixed-price contract. For the subsequent contract, we assume that customers expect savings equal to the average price difference for the twelve months leading up to that contract. Variations on this ”status-quo” assumption, where savings remain as they were toward the end of the sample, generate similar estimates. However, if we use the actual prices as expectations for future prices, our long-run estimates change, increasing in magnitude to -0.39. The dynamic elasticity curve under these alternative perfect foresight assumptions is displayed in Figure 12b. The difference in the estimates is driven by price patterns occuring after the end of our usage panel. As shown in Figure 2a, the average MEA price began to rise in the latter half of 2014 and coverged with the ComEd rate around June 2015. Based on anecdotal evidence, customers did not anticipate this convergence.[21]

7 Incidence and Policy Implications

It is 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 for a small change in the tax rate is approximately

$$\frac{\partial P}{\partial t} = \frac{\epsilon_s}{\epsilon_s + \epsilon_d}, \tag{7}$$

where P is price, t is the tax, and ϵ_s and ϵ_d are the absolute values of supply and demand elasticities, respectively. The share of the tax that is passed through to consumers, $\frac{\partial P}{\partial t}$, is sometimes called the “pass-through” rate. If it is close to 1, the majority of the tax is passed through to consumers, causing them to bear the incidence of the tax. Conversely, if $\frac{\epsilon_s}{\epsilon_s+\epsilon_d}$ is close to zero, then the tax is passed through to producers rather than consumers, and consumers do not bear the burden. Equation (7) demonstrates that, all else equal, a larger demand elasticity (in absolute value) implies a lower incidence of tax on consumers, and can be used to evaluate the significance of our findings.

[21] We estimated an alternative parameteric form that allowed for exponential decay in anticipation, which we found to be overly sensitive to the prices occuring after our usage panel. Our linear anticipation model reduces this sensitivity to anticipation to better pin down the long-run response.

For example, suppose that the elasticity of supply is equal to 1.5. Using the one-year estimate of -0.1 for the elasticity of demand implies a pass-through rate of 94 percent for consumers. Using the three-year estimate of -0.3, by contrast, implies a pass-through rate of 83 percent. The corresponding pass-through rates for suppliers are 6 percent and 17 percent. These differences become more stark for smaller values of supply elasticity.

An alternative to estimating the demand elasticity is estimating the pass-through rate directly. However, this is difficult to do in the electricity sector because consumer prices are typically controlled by utility regulators. Thus, any changes in generator costs (including a carbon tax or cap-and-trade) are passed onto consumers gradually, making it difficult to directly estimate the long-run pass-through rate.

Finally, we note that policies which employ a tax to target a specific level of emissions need to correctly quantify how equilibrium electricity consumption responds to changes in taxes.[22] For a small tax increase, the fall in equilibrium quantity is equal to:

$$\frac{\partial Q}{\partial t} = -\frac{\epsilon_s \epsilon_d}{\epsilon_s + \epsilon_d}\frac{Q}{P} = -\frac{1}{1/\epsilon_s + 1/\epsilon_d}\frac{Q}{P}. \tag{8}$$

Thus, more elastic demand corresponds to a greater fall in equilibrium electricity consumption. Underestimating the demand elasticity will thus lead to an underestimate of how much a carbon tax would reduce emissions, while overestimating it would do the opposite.

In an analysis of the Clean Power Plan Rule, Bushnell et al. (2015) set the median demand elasticity to -0.05. Additionally, in a study of the Spanish wholesale electricity market, Fabra and Reguant (2014) estimate the pass-through of carbon prices to be nearly full, also implying very inelastic demand. Our empirical results suggest that the long-run elasticities are substantially greater in magnitude and would be more appropriate for long-run predictions.

8 Conclusion

An accurate estimate of the price elasticity of electricity demand is essential for evaluating the effects of energy policies such as a carbon tax. Policies that address climate change can be expected to permanently affect the price of electricity, which in turn will affect emissions. However, few reliable estimates of the price elasticity exist, as price changes in this market are often endogenous, short-lived, small, or unnoticed.

We overcome these challenges by exploiting a policy change in Illinois that allowed munic-

[22] 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.

ipalities to select electricity suppliers on behalf of their residents. We show that communities implementing Municipal Electric Aggregation experienced large and lasting price changes relative to communities that did not implement MEA. These price drops, in turn, led to increased electricity usage. Overall, we estimate that the price elasticity of electricity is twice as large in the long-run as in the short-run.

Our finding underscores the importance of identifying settings that accurately capture long-run elasticities, as short-run data may grossly understate total effects. We demonstrate that consumers need at least two years to full adjust to a new policy. Although our data suggest that the consumer response to elasticity prices stabilizes after a two-year period, it is possible that further adjustments may occur over a longer time window. Estimates of elasticities that exploit data spanning more periods is a good avenue for future research.

Finally, we note that the natural experiment created by MEA decreased electricity prices, whereas price-based climate policies would increase prices to reduce total carbon emissions. One consideration for policy-makers is whether the demand response is symmetric for price increases and price decreases. In our data, we find some evidence that consumers respond similarly to the relative price increase occurring over a year after MEA implementation. More direct evidence of the demand response to price increases is welcome.

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, January). 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.

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

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

Bushnell, J. B., S. P. Holland, J. E. Hughes, and C. R. Knittel (2015). Strategic policy choice in state-level regulation: The epa's clean power plan. NBER Working Paper 21259.

DeVirgilio, B. (2006). A discussion of the deregulation of the energy industry in illinois and its effects on consumers. *Loy. Consumer L. Rev. 19*, 256.

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

Fabra, N. and M. Reguant (2014). Pass-through of emissions costs in electricity markets. *American Economic Review 104*(9), 2872–2899.

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

Giulietti, M., M. Waterson, and M. Wildenbeest (2014). Estimation of search frictions in the b ritish electricity market. *The Journal of Industrial Economics 4*(62), 555–590.

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

Hortaçsu, A., S. A. Madanizadeh, and S. L. Puller (2017). Power to choose? an analysis of consumer inertia in the residential electricity market. *American Economic Journal: Economic Policy forthcoming*.

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. *The American Economic Review 104*(4), 1417–1438.

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

Littlechild, S. (2008). Municipal aggregation and retail competition in the ohio energy sector. *Journal of Regulatory Economics 34*(2), 164–194.

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.

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

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.

Plug In Illinois (2016). List of Communities Pursuing an Opt-Out Municipal Aggregation Program. `https://www.pluginillinois.org/MunicipalAggregationList.aspx`. Accessed: 2016-02-11.

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.

U.S. Energy Information Administration (2017). February 2017 monthly energy review. Accessed March 3, 2017.

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.

Tables and Figures

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

Figure 1: Spatial distribution of MEA and non-MEA communities in our sample

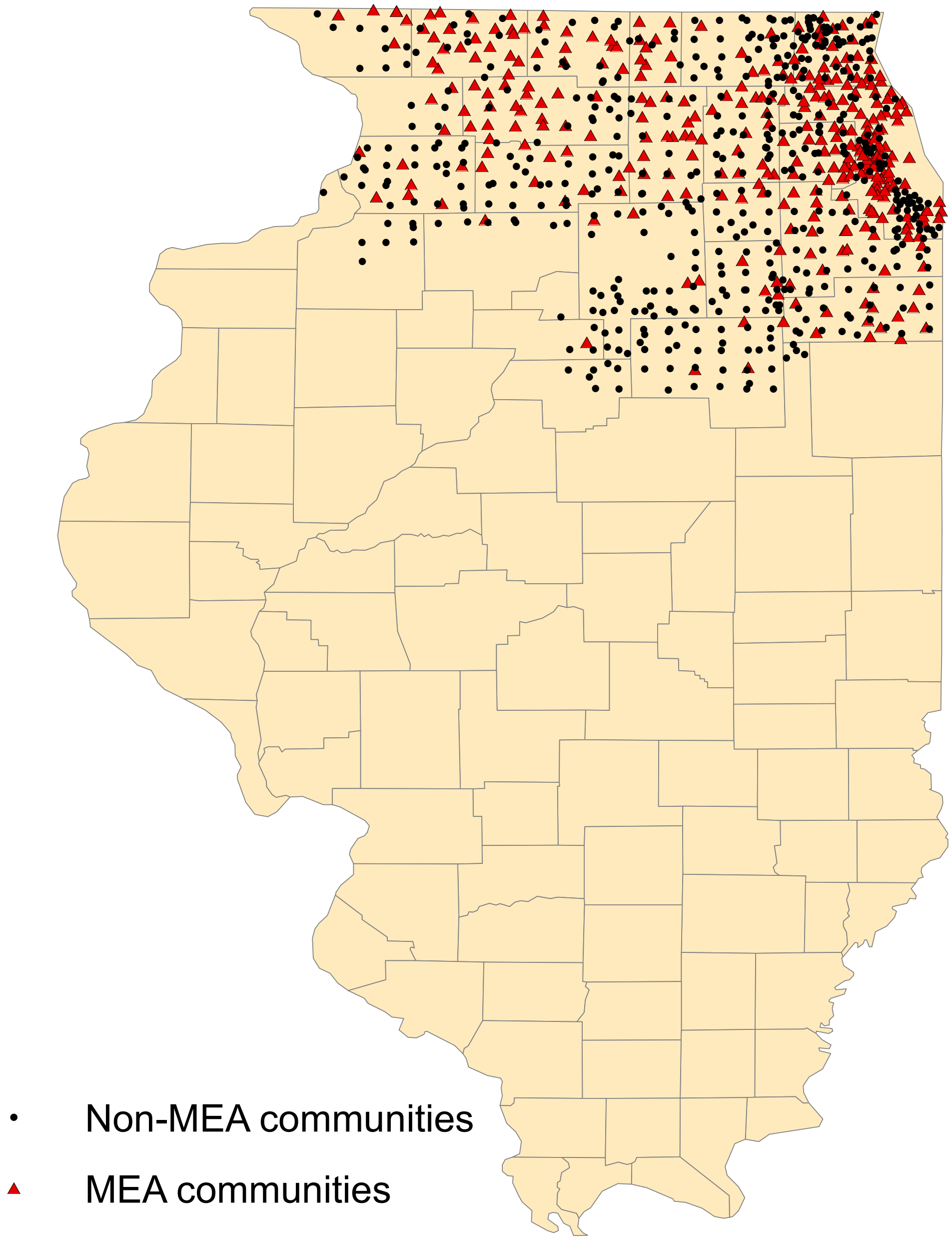

Notes: Figure displays the locations of communities in our sample. Red triangles indicate communities that implemented MEA. Black dots indicated communities that did not implement MEA.

Figure 2: Price differences between MEA and non-MEA communities

(a) Monthly electricity rates, 2008-2016

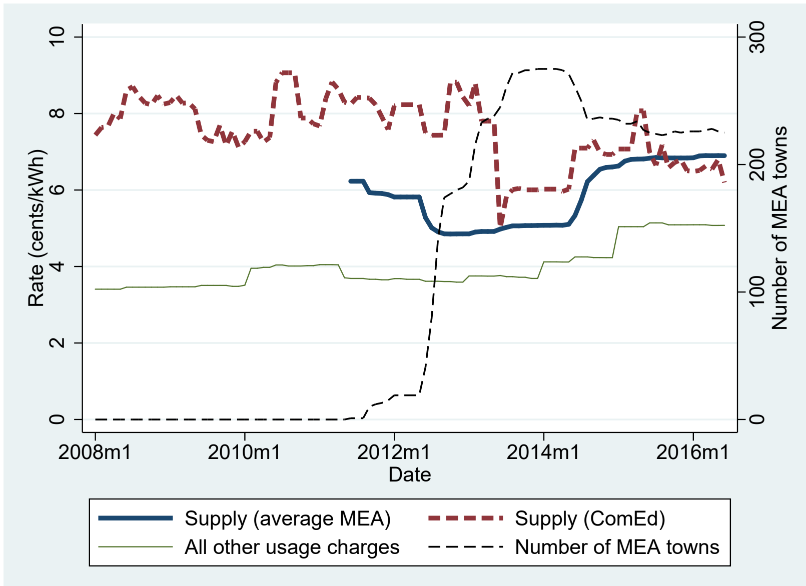

(b) Heterogeneity in price changes realized by MEA communities

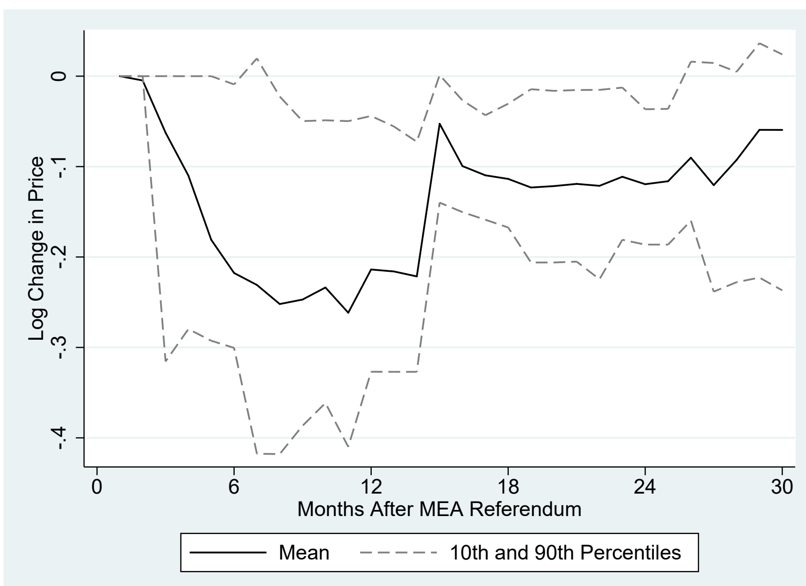

Notes: The thick blue line in panel (a) displays the average supply rate among all communities that adopted municipal electric aggregation (MEA). The first community adopted MEA in June of 2011. Non-MEA communities pay the supply rate charged by ComEd, indicated by the thick, dashed red line in panel (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 MEA. 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 MEA.

Figure 3: Monthly electricity usage in Illinois ComEd service territories, 2007-2014

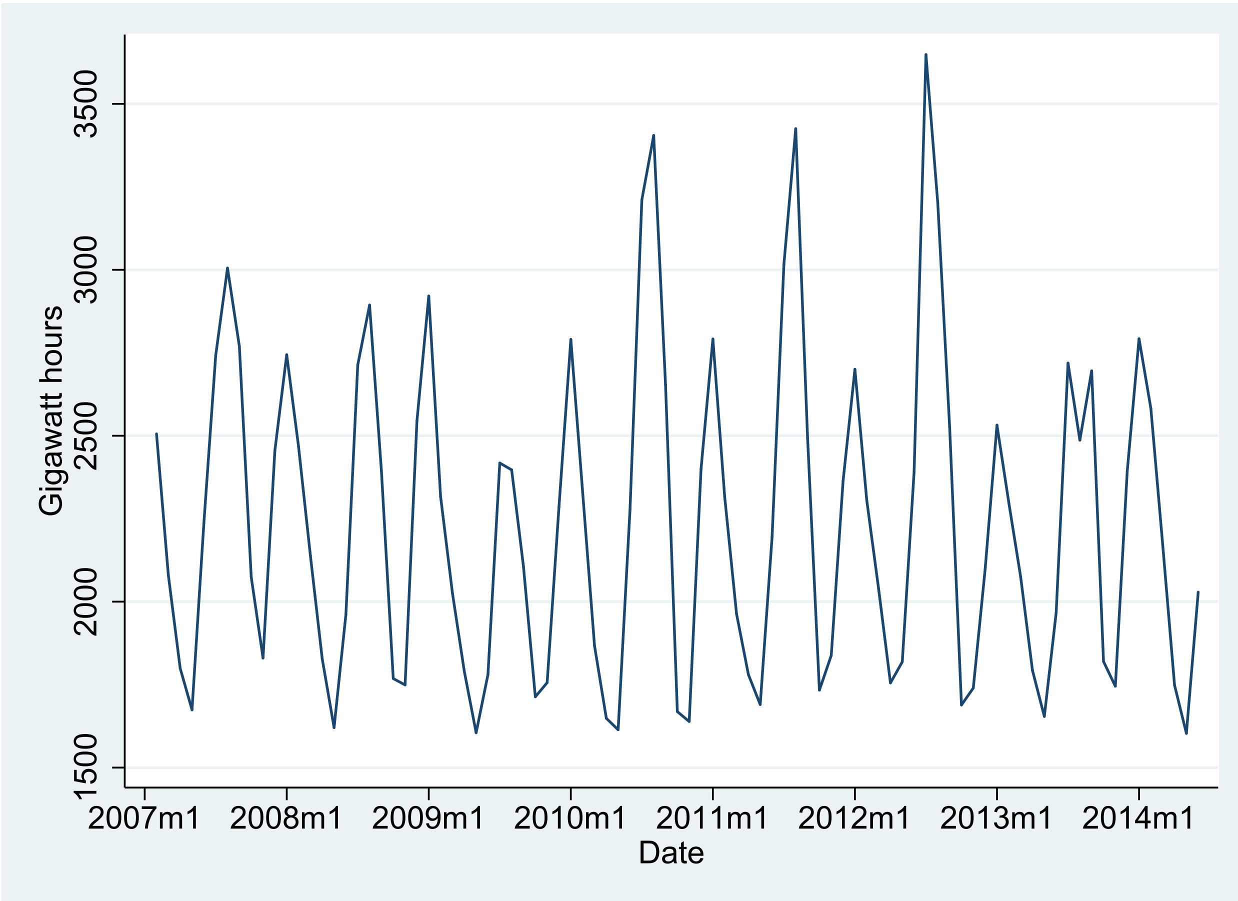

Notes: Figure displays total electricity usage across the 768 ComEd service territories in our estimation sample.

Figure 4: Comparing regression to nearest-neighbor matching

(a) Regression

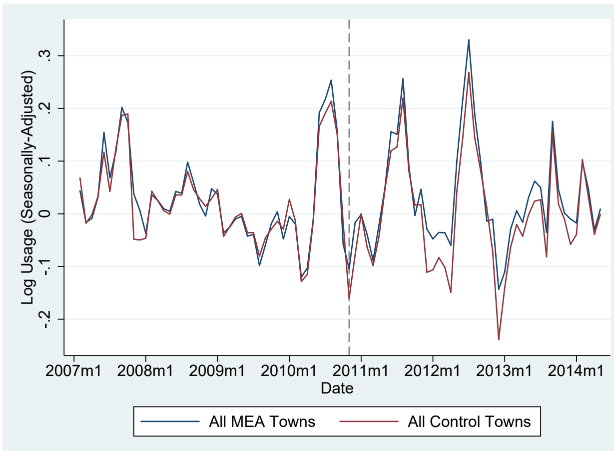

(b) Nearest-neighbor matching

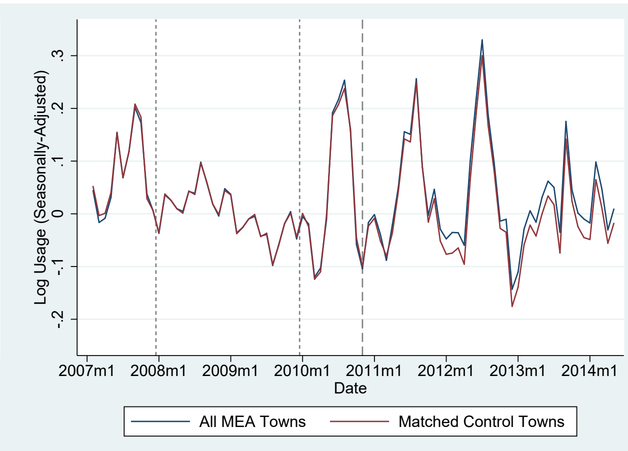

(c) Treatment and All Controls: Difference

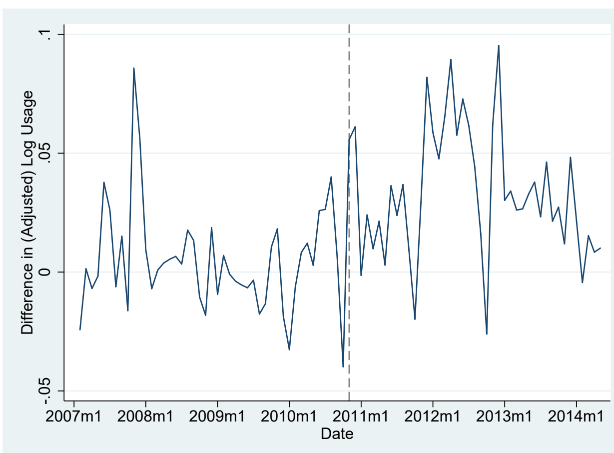

(d) Treatment and Matched Controls: Difference

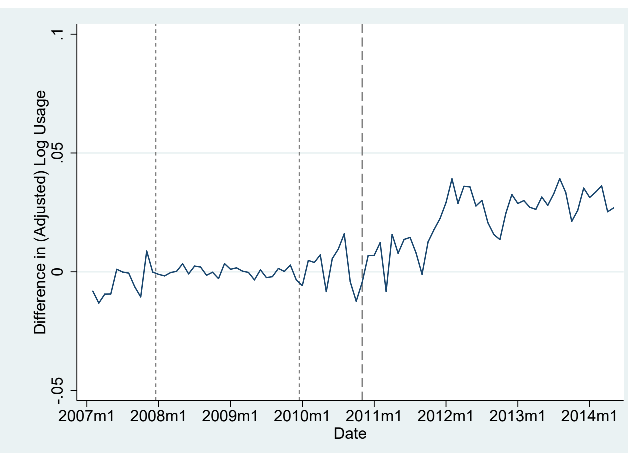

Notes: Panel (a) displays seasonally-adjusted usage for all MEA and non-MEA communities. The red line corresponds to the control group in a typical regression, in which community-specific month-of-year fixed effects are used as controls. Panel (b) employs the nearest-neighbor matching procedure, in which five communities are selected for each MEA 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. Panel (d) demonstrates that the matching procedure greatly reduces noise compared to a standard regression. The pre-period fit is much better, even for 2010 usage, which was not used in the matching procedure. 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.

Table 2: Characteristics of MEA, non-MEA, and matched control communities

	(1) Implemented MEA	(2) Did not pass MEA	(3)	(4) Matched controls	(5)
	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.

Figure 5: Effect of implementing MEA on electricity prices and usage

(a) Effect of MEA on log price

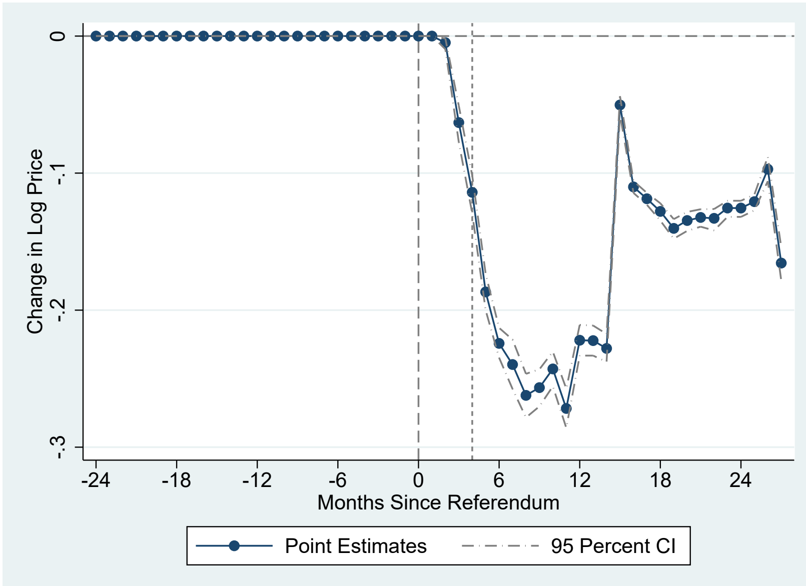

(b) Effect of MEA on log usage

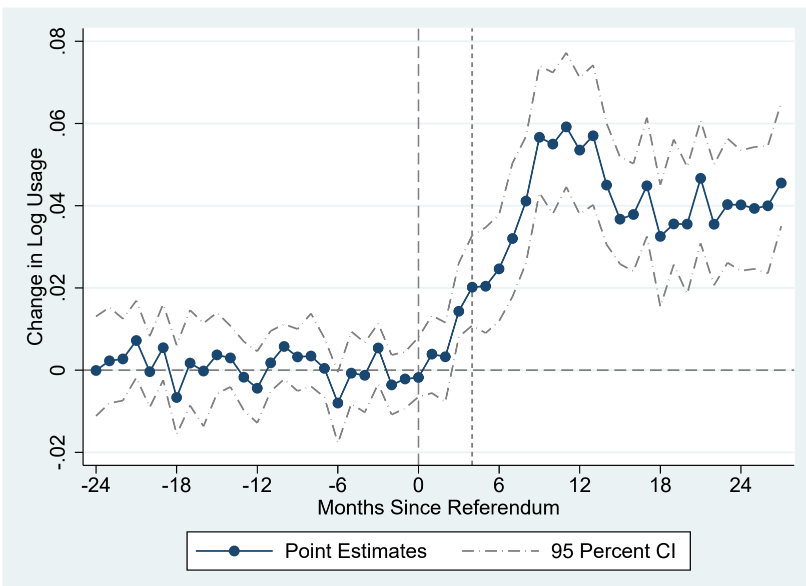

Notes: Panels (a) and (b) displays estimates of the mean price and usage effect, respectively, of implementing MEA in a community relative to that community's five nearest-neighbors, as defined by the difference-in-differences matching procedure outlined in the main text. Prices differences are calculated using the natural log of the marginal electricity supply rates. The pre-period price difference is exactly zero because all communities faced the ComEd price during that period. 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

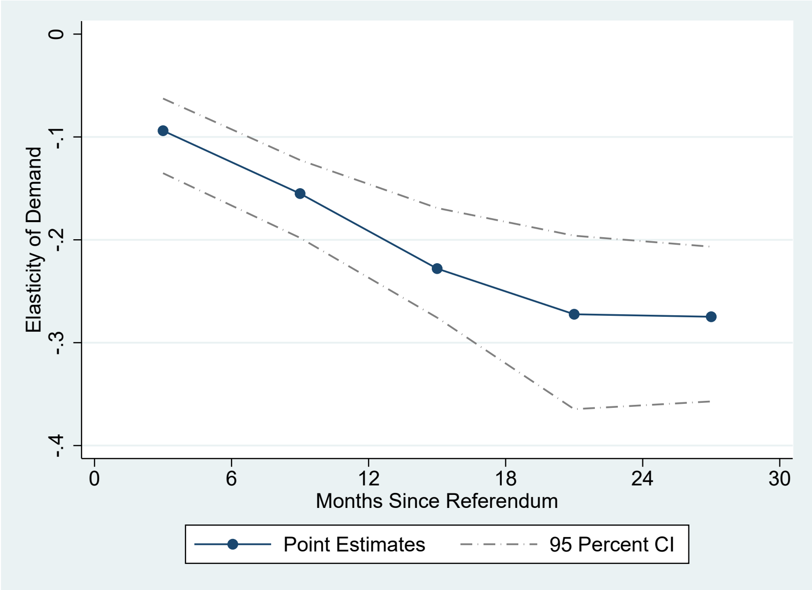

(b) Quadratic fit

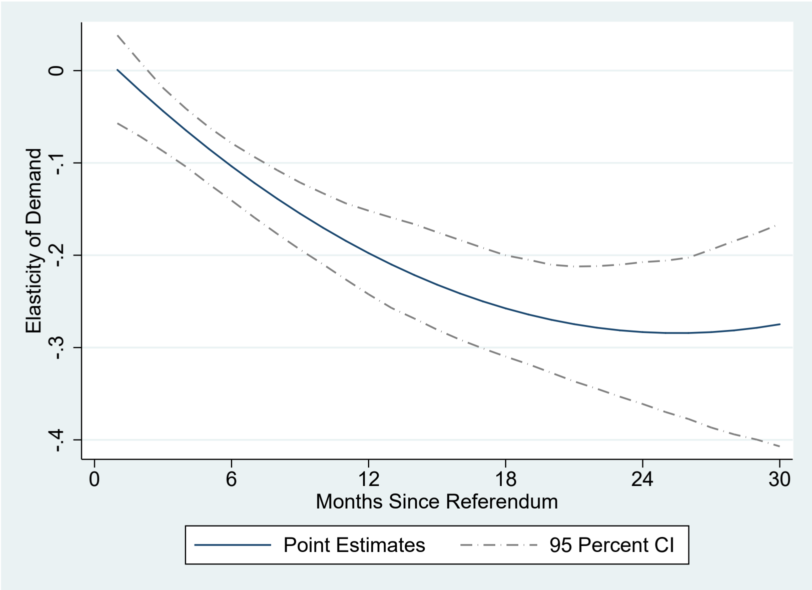

Notes: Elasticities in panel (a) are calculated for each six-month period by regressing community-month changes in log usage on the observed change in log price. The corresponding counts of observations for each six-month group are: 1685, 1656, 1504, 1144, and 589. In panel (b), the time-dependent elasticity 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.

Table 3: Matching estimates of the effect of MEA 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.

Table 4: Matching estimates of the effect of MEA 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.

Figure 7: Anticipation effects of implementing MEA on log usage

(a) Log price

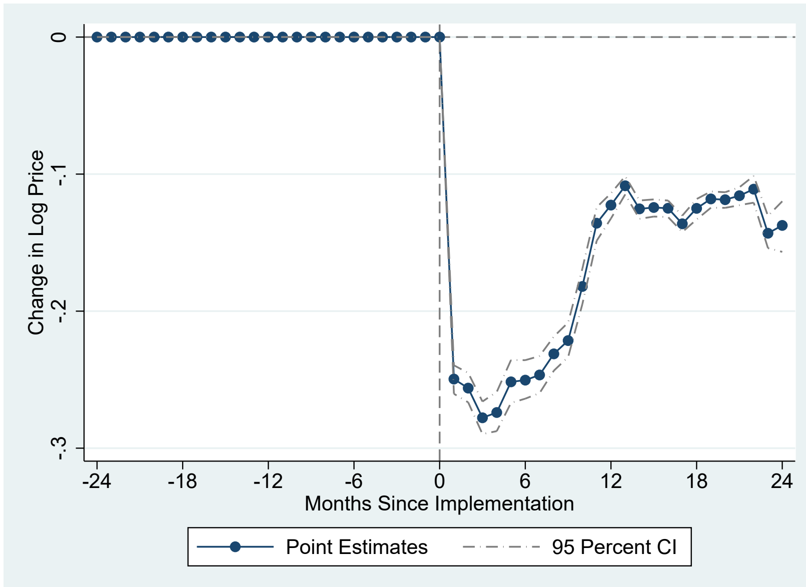

(b) Log usage

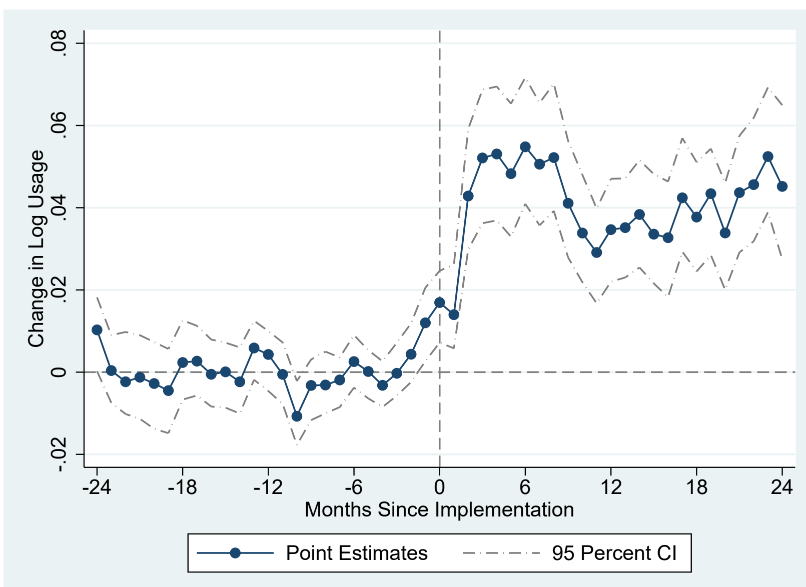

Notes: The figure displays estimates of the anticipation effect of implementing MEA in a community relative to that community's five nearest-neighbors, as defined by the difference-in-differences matching procedure outlined in the main text. The vertical dashed line corresponds to the month before MEA was implemented (i.e., the month before the price change). Confidence intervals are constructed via subsampling. The usage is normalized so that the average usage difference in the year prior to the referendum is zero.

Figure 8: Usage response to MEA: Anticipation only

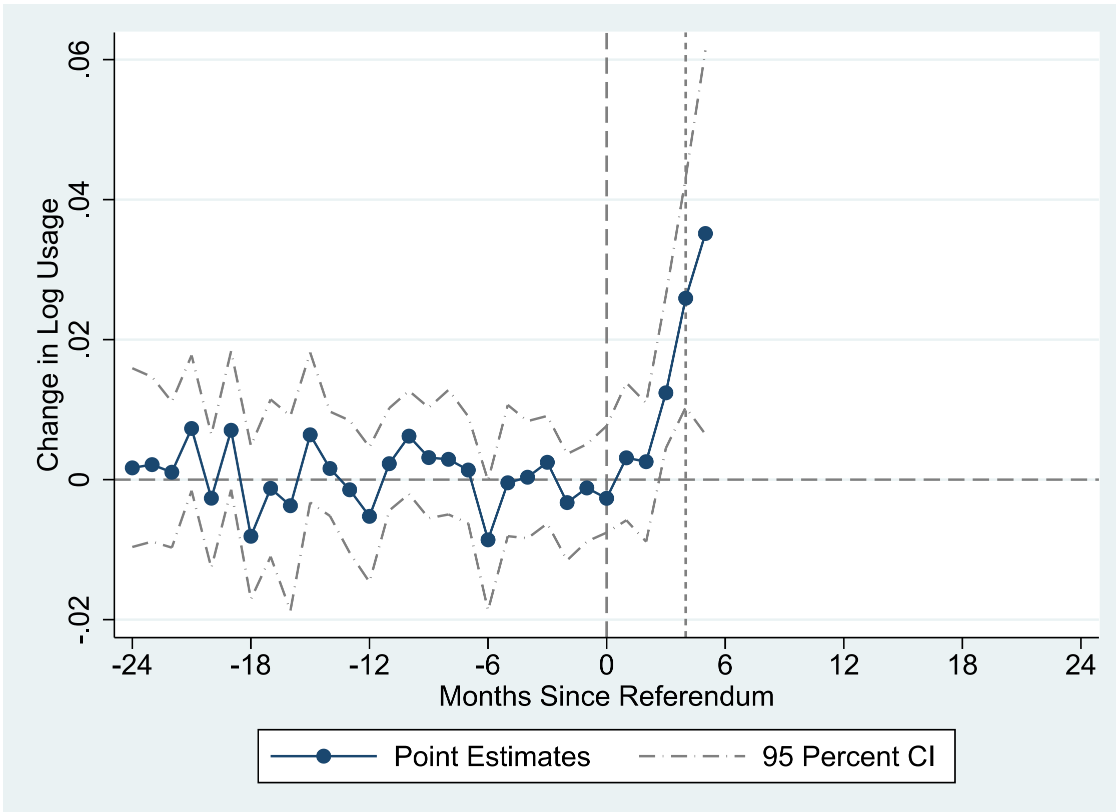

Notes: The figure displays estimates of the mean anticipation effect for communities that approved MEA prior to the implementation of a price change. The effect is estimated relative to that community's five nearest-neighbors, as defined by the difference-in-differences matching procedure outlined in the main text. The horizontal axis indicates months since the referendum was passed. The short dashed line indicates the median implementation period. Confidence intervals are constructed via subsampling.

Figure 9: Elasticities by Housing Demographics

(a) Population

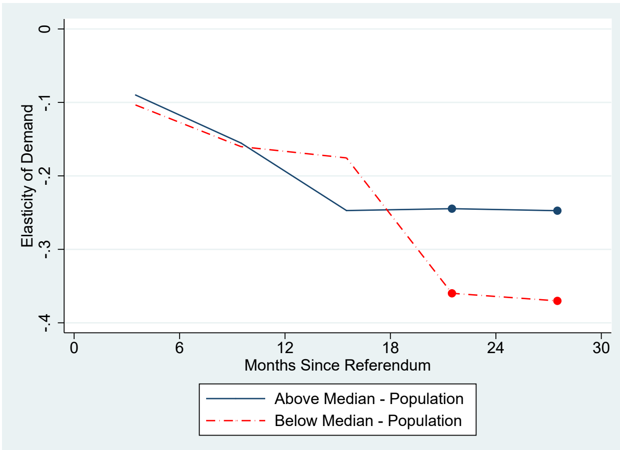

(b) Housing Units per Capita

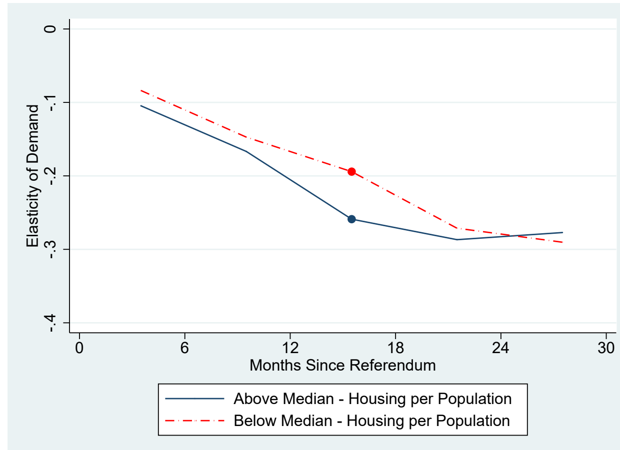

(c) Percent Owner Occupied

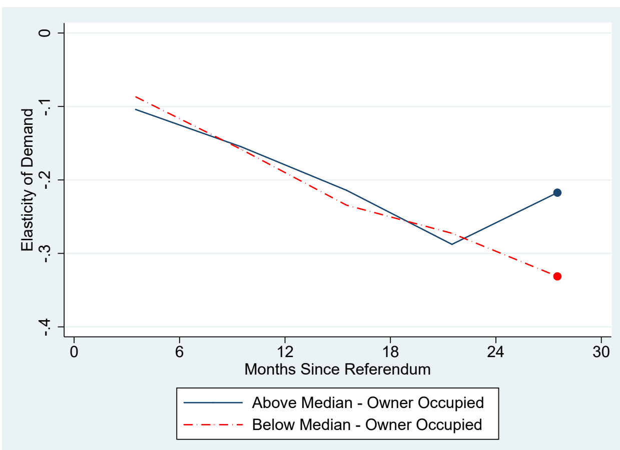

(d) Year Built

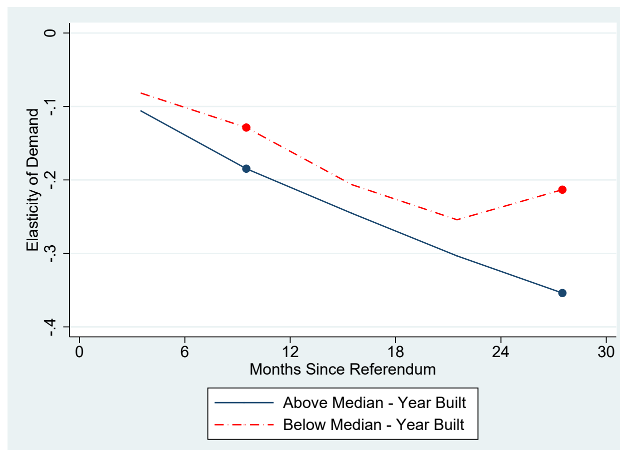

Notes: These panels display elasticity estimates for the upper half and lower half of demographic variables. The estimates are calculated by regressions of log usage on log price, where the price change is interacted with a dummy indicating whether or not the two is in the upper half of the distribution. 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 10: Elasticities by Socioeconomic Demographics

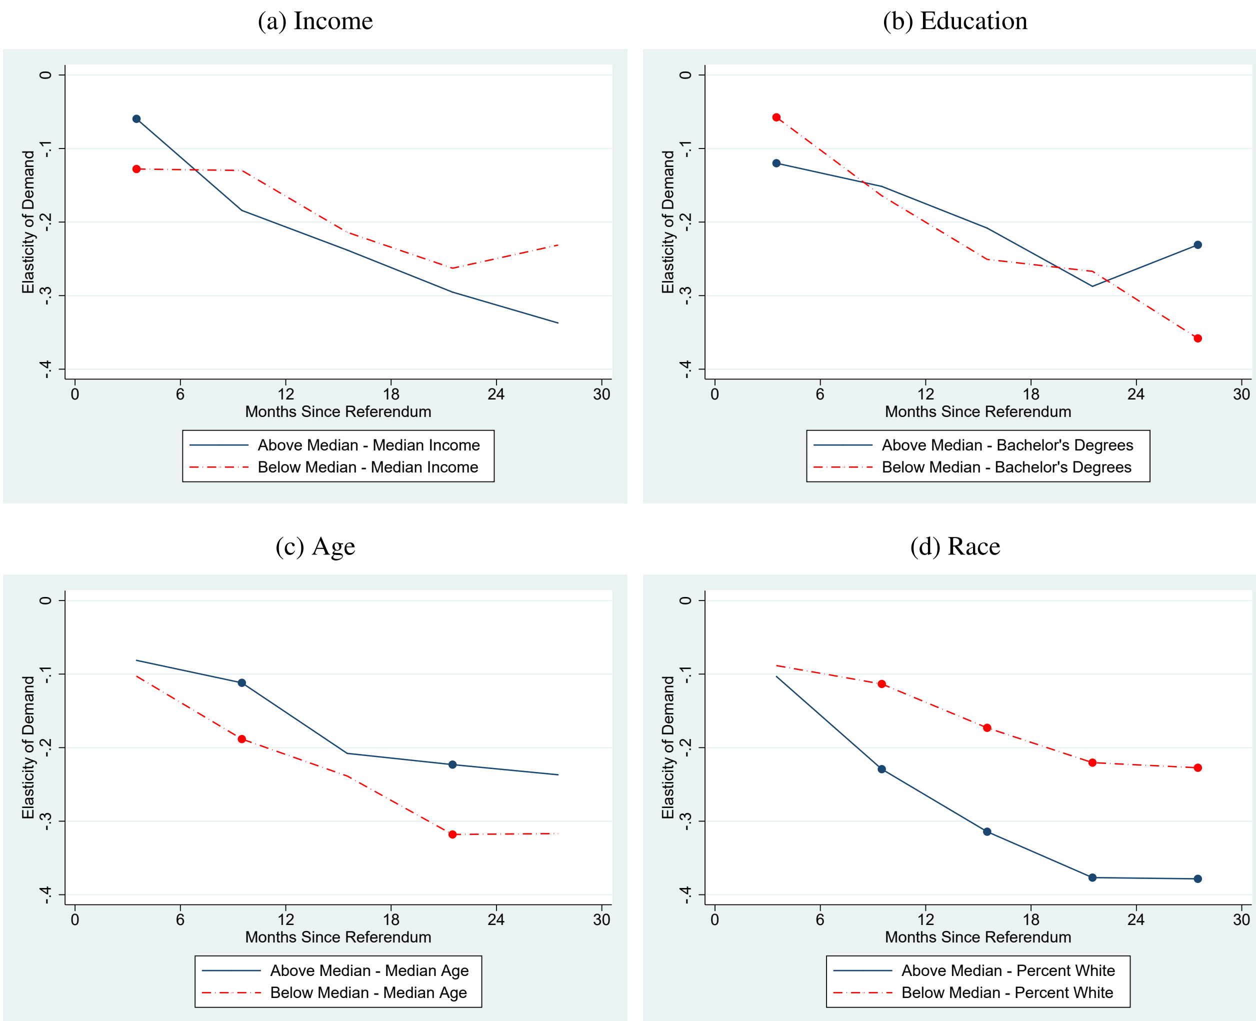

Notes: These panels display elasticity estimates for the upper half and lower half of demographic variables. The estimates are calculated by regressions of log usage on log price, where the price change is interacted with a dummy indicating whether or not the two is in the upper half of the distribution. 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 11: Estimated elasticities versus mean log price change

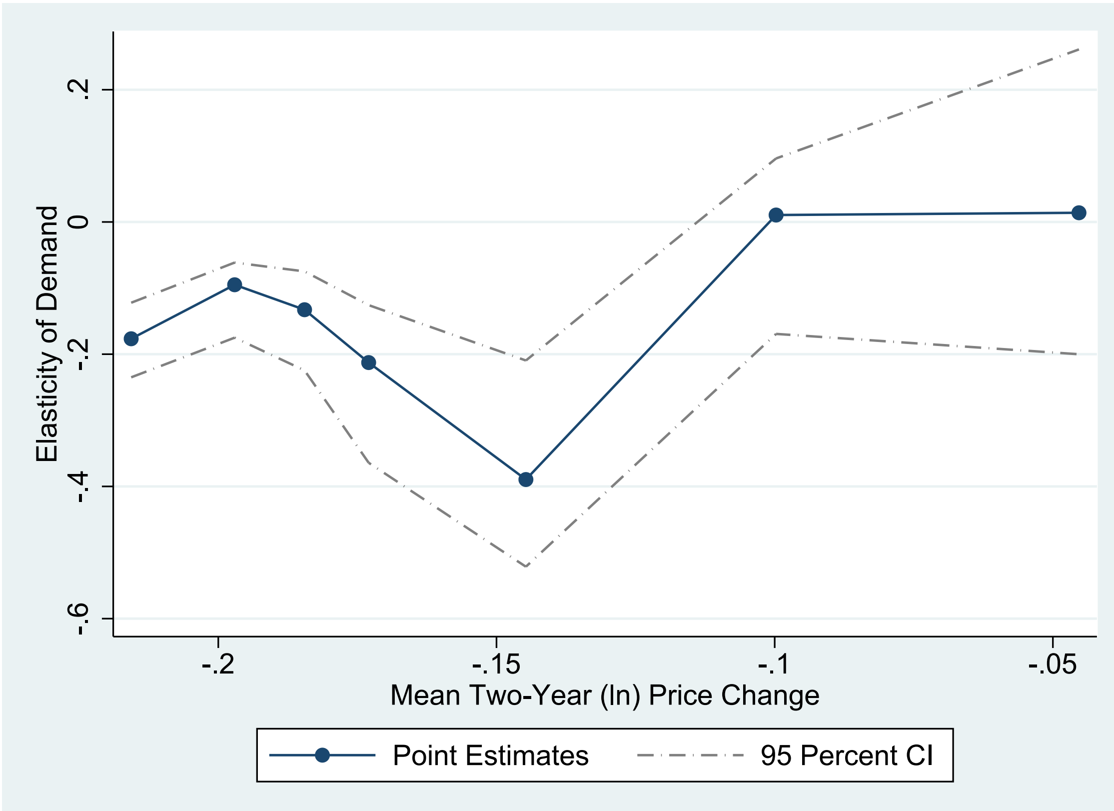

Notes: Communities are split into seven groups based on the average two-year price change. Elasticities are calculated separately for each group. The graph shows no relationship between the estimated group elasticity and the price change, mitigating some concerns about endogeneity. Confidence intervals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

Figure 12: Long-run elasticity projections

(a) Status quo expectations

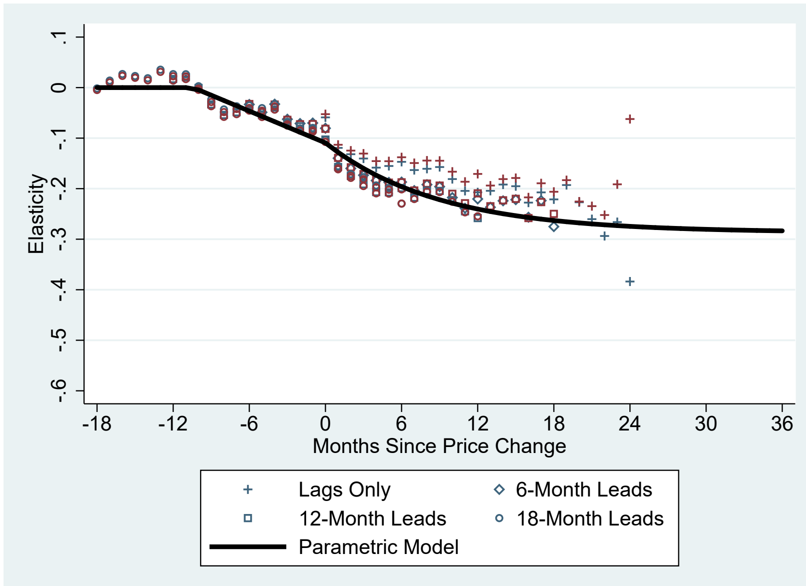

(b) Perfect foresight

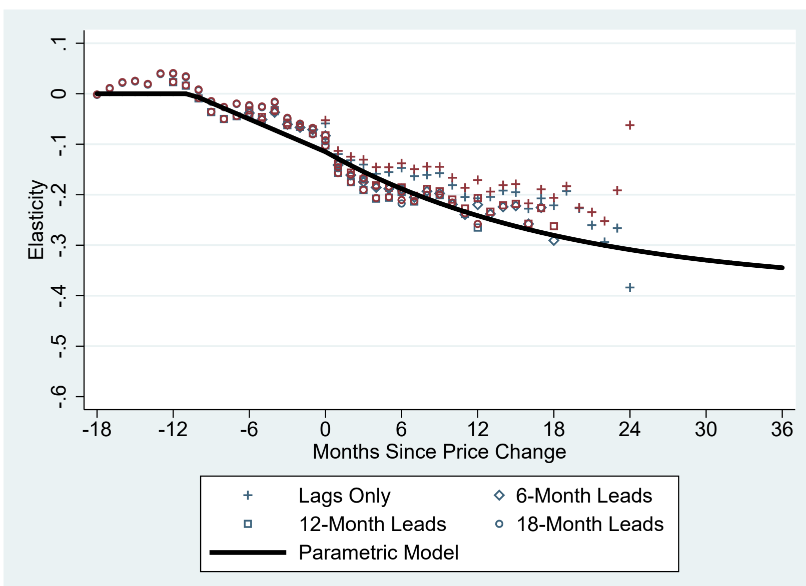

Notes: The solid lines display the estimated elasticity as a function of the months since a price change. The line is estimated by fitting a parametric model to the usage changes and price changes obtained by matching. The plotted points display non-parametric analogs of the line, where each set of points is calculated using 24 combined lags and leads, beginning with the lead displayed in the legend. The red points correspond to non-parametric models estimated with community fixed effects; the blue points are from specifications with no fixed effects.

Table 5: Parametric Estimates of the Dynamic Elasticity Curve

Time Period	Status-Quo Expectations	Perfect Foresight
Month 19-24 Leads	0.000	0.000
	(0.000)	(0.000)
Month 13-18 Leads	0.000	0.000
	(0.000)	(0.000)
Month 7-12 Leads	-0.014**	-0.016**
	(0.007)	(0.006)
Month 1-6 Leads	-0.073***	-0.077***
	(0.012)	(0.013)
Contemporaneous	-0.109***	-0.115***
	(0.014)	(0.017)
Month 1-6 Lags	-0.165***	-0.159***
	(0.023)	(0.019)
Month 7-12 Lags	-0.224***	-0.220***
	(0.025)	(0.024)
Month 13-18 Lags	-0.255***	-0.265***
	(0.025)	(0.026)
Month 19-24 Lags	-0.270***	-0.297***
	(0.029)	(0.035)
Month 25-30 Lags	-0.278***	-0.321***
	(0.035)	(0.049)
Month 31-36 Lags	-0.282***	-0.338***
	(0.039)	(0.064)
Long Run	-0.287***	-0.385***
	(0.047)	(0.141)

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. Standard errors are in parentheses. Significance is determined by the jackknife.

A Appendix

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 (1) whether or not the change was anticipated; (2) when the change ocurred; and (3) whether the change is temporary or permanent.

This can be shown clearly by solving the second-order difference equation (3):

$$y_t = K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} h(t+s) + K_2 \sum_{s=0}^{t-1} (\lambda_2)^s h(t-s) + (\lambda_2)^t \left(y_0 - K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} h(s) \right) \quad (9)$$

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 our solutions are real).

Equation (9) 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 ($t \to \infty$), 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, is a function of past prices, but not of future prices.

The effect of a one-period change in price on current consumption, holding all other prices constant, is equal to

$$\begin{aligned} \frac{dy_t}{dp_t} &= K_1 (\lambda_1)^{-1} \alpha_3 - (\lambda_2)^t K_1 (\lambda_1)^{-(t+1)} \alpha_3 \\ &= K_1 \alpha_3 \left((\lambda_1)^{-1} - (\lambda_2)^t (\lambda_1)^{-t} (\lambda_1)^{-1} \right) \\ &= K_1 \frac{\alpha_3}{\lambda_1} \left(1 - \left(\frac{\lambda_2}{\lambda_1} \right)^t \right) \end{aligned}$$

The fully anticipated one-period price effect corresponds to letting $t \to \infty$:

$$\frac{dy_\infty}{dp_\infty} = K_1 \frac{\alpha_3}{\lambda_1} = \frac{\alpha_3}{\alpha_2(\lambda_1 - \lambda_2)}$$

To obtain the completely unanticipated price effect, let $t = 1$:

$$\frac{dy_1}{dp_1} = K_1 \frac{\alpha_3}{\lambda_1}\left(1 - \frac{\lambda_2}{\lambda_1}\right) = \frac{\alpha_3}{\alpha_2 \lambda_1}$$

We are generally interested in estimating the impact on consumption of a permanent change in price. The short-run effect of this is equal to

$$\frac{dy_t}{dp_t^*} = \sum_{\tau=t}^{\infty} \frac{dy_t}{dp_\tau} = \frac{\alpha_3 \lambda_1}{\alpha_2 \left(\lambda_1 - \lambda_2\right)\left(\lambda_1 - 1\right)} \left(1 - \left(\frac{\lambda_2}{\lambda_1}\right)^t\right) \tag{10}$$

As before, we can obtain the fully anticipated or completely unanticipated effects by letting $t \to \infty$ or setting $t = 1$, respectively. The long-run effect of a permanent change in price is equal to

$$\frac{dy_\infty}{dp^*} = \lim_{t \to \infty} \sum_{\tau=-\infty}^{t-1} \frac{dy_t}{dp_\tau} = \frac{\alpha_3}{\alpha_2 \left(1 - \lambda_2\right)\left(\lambda_1 - 1\right)} \tag{11}$$

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 MEA. 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 MEA on electricity prices and usage using a standard difference-in-differences model:

$$Y_{cmy} = \sum_{\tau=-24,\tau\neq-1}^{24} \beta_\tau MEA_{c\tau} + \beta_{25} MEA_{c,25} + \beta_{-25} MEA_{c,-25} + \alpha_{cm} + \alpha_{my} + \varepsilon_{cmy}, \quad (12)$$

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 $MEA_{c\tau}$ is an indicator equal to 1 if, as of month m and year y, community c implemented MEA τ months ago. The month before MEA implementation ($\tau = -1$) is the omitted category. To ensure that our estimated coefficients are relative to this category, we include indicators for MEA having been implemented 25 or more months ago ($MEA_{c,25}$) and for MEA being implemented 25 or more months in the future ($MEA_{c,-25}$). 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 MEA as the reference period:

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

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

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

Figure A.3 presents the change in electricity prices following MEA, in logs, as estimated by equation (12). 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 treatment community. The second vertical dashed line in Figure A.3 shows the point at which half of all communities have implemented MEA (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 MEA communities relative to the control group. There is an increase in the relative MEA price 28 months after passing MEA, 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 MEA communities remain significantly lower than those in the control group for the entire sample period.

Figure A.4 shows the corresponding estimates for electricity usage. Prior the referendum, the difference in usage between MEA and the control communities is statistically indistinguishable from zero. Usage in MEA communities then begins to increase following the referendum. By the end of the first year, usage in MEA communities is about 0.1 log points (9.5 percent) higher relative to the counterfactual.

Table A.2 shows the estimated impact of MEA 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 (13). 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.3 shows the estimated change in usage as estimated by equation (13) for the sample of communities that implemented MEA. Our preferred specification, presented in Column 4, estimates that electricity usage is 4.8 log points higher in the first 6 months following the referendum, and this increases to 11.4 log points within one year.

Finally, Figure A.5 shows the elasticities implied by the two preceding tables. Specifically, we show the ratio of coefficients from Tables A.3 and A.2, which estimate the MEA-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 MEA to -0.45 in the first six months after passage.

A.4 MEA Materials

Referendum wording

Excerpt from Sec. 1-92. A[23]
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”.

[23]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.

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

Electric Aggregation Program
Frequently Asked Questions

Overview of Municipal Aggregation

What is Municipal Aggregation?
Illinois law allows municipalities and counties to negotiate the purchase price of electricity on behalf of residential and small business utility customers living within their borders. While these governmental entities choosing community aggregation would be responsible for negotiating the price of power from a supplier other than the traditional utility, your utility would still be responsible for delivering that power to your home, and billing you for it.

How can I get more information about the municipality or county's aggregation program?
Contact your municipality or county for information related to the referendum and the aggregation program. Additional resources can be found at:
http://www.dynegyenergyservices.com/residential/municipal-aggregation/communities-we-serve.php

Eligibility and Enrollment

Who is eligible to participate?
Residential or small business customers located in the participating governmental entity boundaries may participate. Customers enrolled in real time pricing, Power Smart Pricing, space electric heat rate, or served by an alternative retail supplier may not be eligible.

How do I enroll?
It's simple. It's automatic. Unless you "opt-out" of the program, all eligible ComEd customer accounts within the boundaries will be enrolled in the program. You will receive a "switch" letter from your utility, ComEd, confirming your enrollment.

Do I have to participate in the municipal or county aggregation plan?
All eligible ComEd utility customers within the municipal or county boundaries will receive an opt-out notification letter via U.S. mail. You may "opt-out" by returning the Opt-Out card by the deadline date identified in your notification. If you choose to opt-out, your account remains with ComEd at the current utility rate.

What if I decide to opt-out after the opt-out deadlines have passed?
You may opt out at any time by calling our toll free number or sending us an email.

Rate and Term Information

What are the Rates and Terms for my Municipality or County?
A listing of communities served by DES can be found at www.DynegyEnergyServices.com. Select your municipality or county to find the applicable rates, contract length, and the terms and conditions for your particular governmental entity. You can expect to receive your first bill with the new DES rate in September 2015.

What if ComEd rates decrease?
If at any time during the term of this Agreement ComEd rates fall lower than the DES price, you will have the option to return to the utility without penalty.

Why does the price change in the second and third year?
DES is committed to offering the lowest possible price to participants in municipal aggregation programs. Cost factors in the power market will change during the second & third year. Specifically, the price DES is charged for capacity changes in year two and three and is reflected in the price.

What happens at the end of the Agreement term?
At the end of the Agreement term, as defined in the Terms and Conditions you have the option of staying with a new Municipal Aggregation program, returning to the utility, or signing with a new supplier independent of the Municipal Aggregation program.

Billing and Service Information

Who will bill me for electricity? Will I get two bills?
You will continue to receive one monthly bill from ComEd. The bill will include the charges for electricity supplied by us, as well as the delivery service charges from ComEd.

Can I still have my payment automatically deducted from my checking account?
Yes, how you pay your bill will not change.

Can I stay on budget billing?
Yes, your budget billing will not be affected by your participation in this program.

Who is responsible for the delivery of power to my home or business?
ComEd will continue to deliver your electricity and will be responsible for maintaining the system that delivers power into your home. As your energy delivery company, they will continue to respond around-the-clock to outages, service calls and emergencies regardless of your electric supplier.

Who do I call to report a power outage or problems with my electric service?
You will continue to call ComEd for power outages, problems with your service or questions regarding your monthly bill.

ComEd Residential Customers: 800.334.7661
ComEd Business Customers: 800.334.7661

Who do I call if I have questions regarding the Municipal or County Opt-Out Electricity Aggregation Program?
Questions should be referred to a member of our DES Customer Care team.

DES Customer Care: 844.351.7691
DESCustCare@Dynegy.com

A complete list of Frequently Asked Questions can be found at
http://www.dynegyenergyservices.com/residential/municipal-aggregation/faq-residential-municipal.php
or by calling DES at 844.351.7691

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

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)

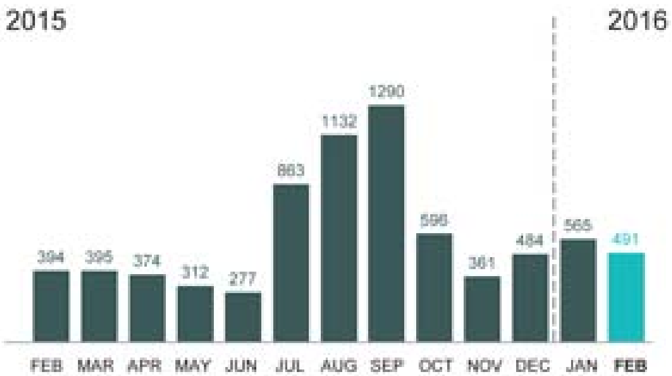

Current month's reading is **actual.**

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	Last Year 25° avg. temp
17.1 kWh	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

DELIVERY
$30.96

ComEd delivers electricity to your home.

ComEd.com
1.800.334.7661

TAXES & FEES $6.48

For Electric Supply Choices visit pluginillinois.org

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

ComEd
An Exelon Company

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

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

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

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

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.

A.5 Appendix Tables

Table A.1: Matching estimates of the effect of MEA 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: Effect of MEA adoption on electricity prices, communities that passed MEA

	(1)	(2)	(3)	(4)
0-6 months post-MEA	-0.119***	-0.100***	-0.123***	-0.101***
	(0.005)	(0.005)	(0.005)	(0.005)
7-12 months post-MEA	-0.307***	-0.313***	-0.312***	-0.320***
	(0.007)	(0.007)	(0.007)	(0.007)
13-18 months post-MEA	-0.298***	-0.266***	-0.303***	-0.267***
	(0.008)	(0.009)	(0.008)	(0.010)
19-24 months post-MEA	-0.283***	-0.285***	-0.285***	-0.287***
	(0.010)	(0.013)	(0.010)	(0.013)
25-30 months post-MEA	-0.283***	-0.264***	-0.298***	-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	25,710	25,710	25,710	25,710
Adjusted R-squared	0.794	0.899	0.804	0.908

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.3: Effect of MEA adoption on electricity usage, communities that passed MEA

	(1)	(2)	(3)	(4)
0-6 months post-MEA	0.073***	0.059***	0.066***	0.048***
	(0.008)	(0.009)	(0.005)	(0.006)
7-12 months post-MEA	0.054***	0.095***	0.065***	0.114***
	(0.012)	(0.016)	(0.012)	(0.016)
13-18 months post-MEA	0.107***	0.140***	0.088***	0.114***
	(0.015)	(0.019)	(0.014)	(0.017)
19-24 months post-MEA	0.084***	0.073***	0.109***	0.113***
	(0.016)	(0.023)	(0.015)	(0.021)
25-30 months post-MEA	0.067***	0.139***	0.068***	0.134***
	(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	25,710	25,710	25,710	25,710
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.

A.6 Appendix Figures

Figure A.1: Estimated price elasticities, monthly

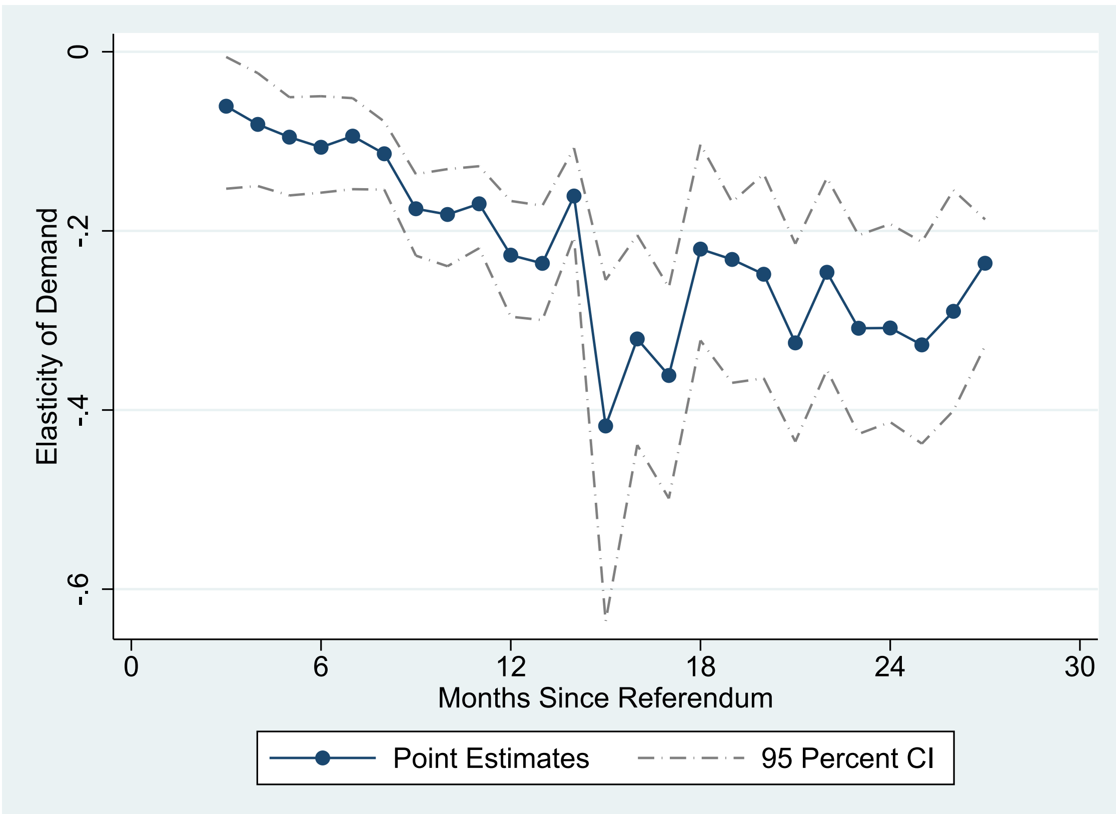

Notes: Elasticities are calculated for each month by regressing community-month changes in log usage on the observed change in log price. Confidence intervals are constructed via subsampling.

Figure A.2: Effect on log usage for those that passed but did not implement MEA

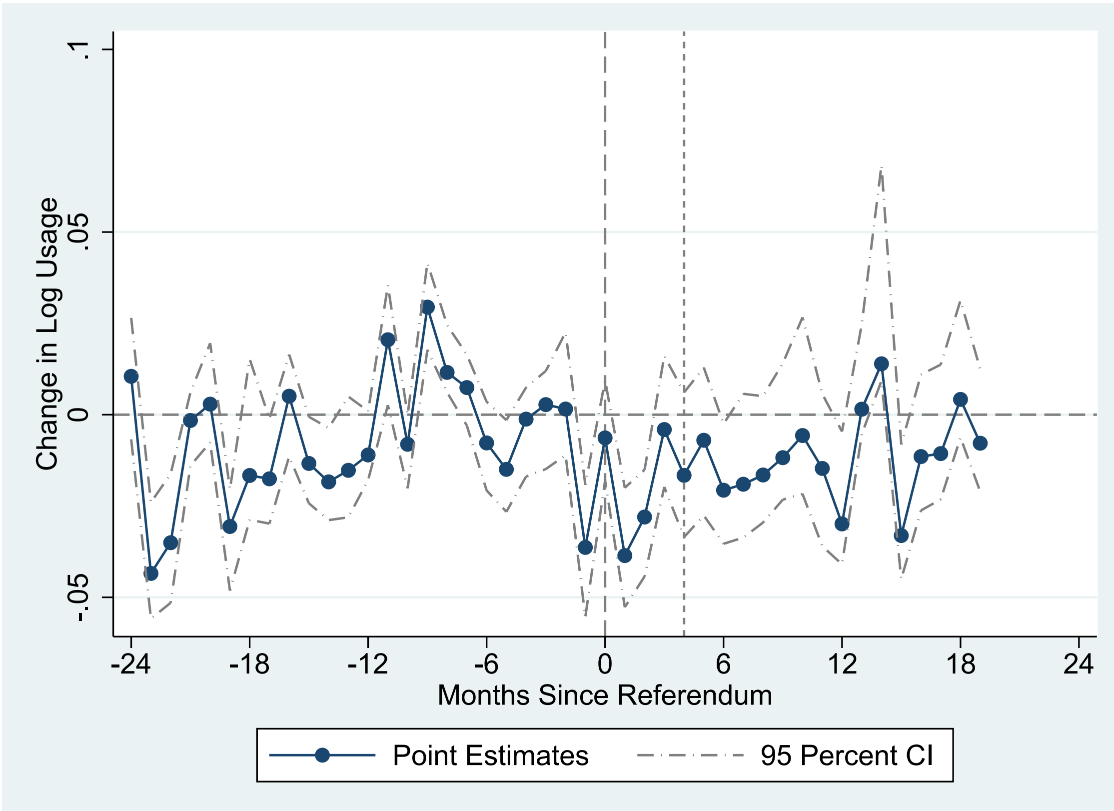

Notes: The figure displays estimates of the mean usage effect for the eleven communities that pass MEA but never implement it. The effect is estimated relative to that community's five nearest-neighbors, as defined by the difference-in-differences matching procedure outlined in the main text. The short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

Figure A.3: Regression estimates of the effect of MEA on electricity prices, communities that passed MEA

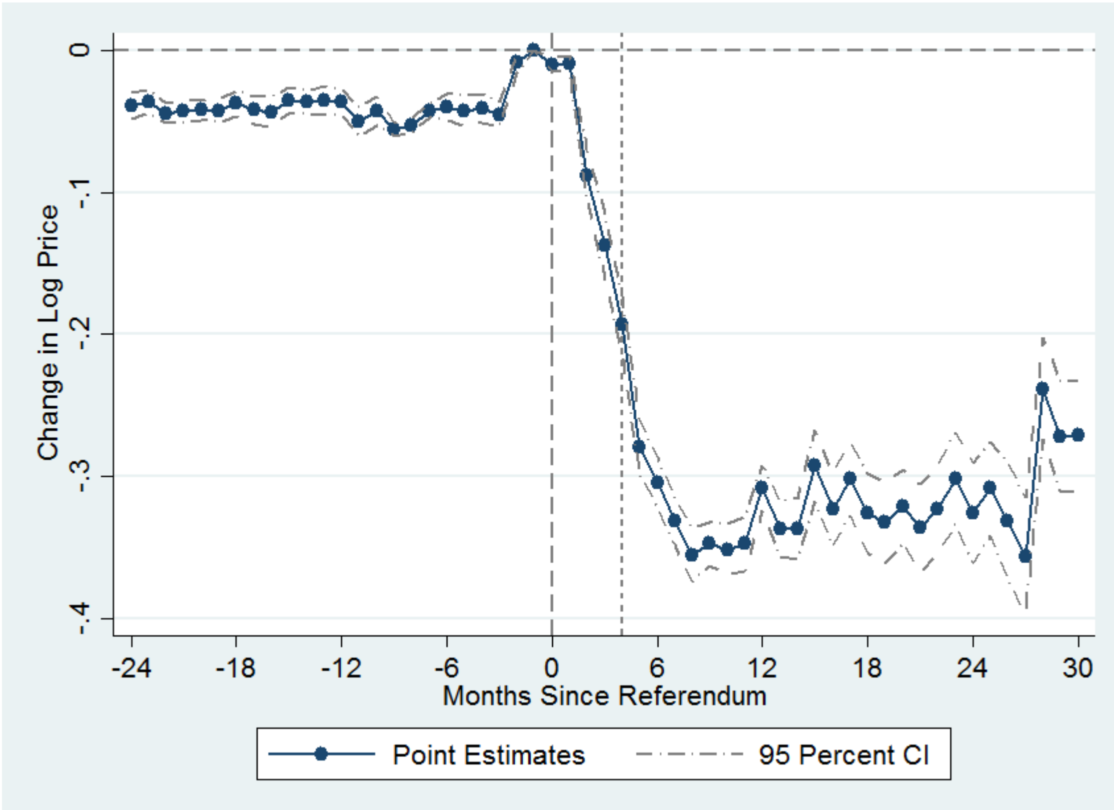

Notes: Outcome is the natural log of the electricity price. The first vertical dashed line indicates the date of the MEA referendum. The second dashed line indicates the date of MEA implementation. Regressions include month-by-year and community-by-month fixed effects. Standard errors are clustered by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

Figure A.4: Regression estimates of the effect of MEA on electricity usage, communities that passed MEA

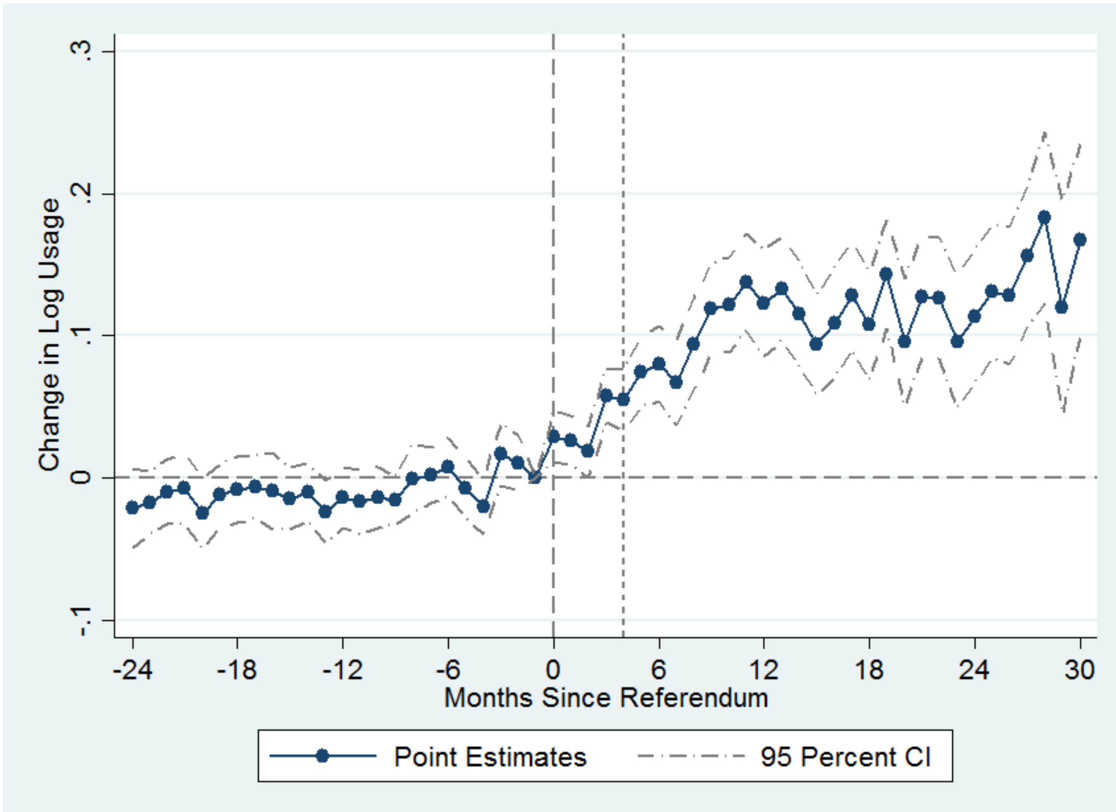

Notes: Outcome is the natural log of total electricity use. The first vertical dashed line indicates the date of the MEA referendum. The second dashed line indicates the date of MEA implementation. Regressions include month-by-year and community-by-month fixed effects. Standard errors are clustered by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

Figure A.5: Estimated price elasticities, communities that passed MEA

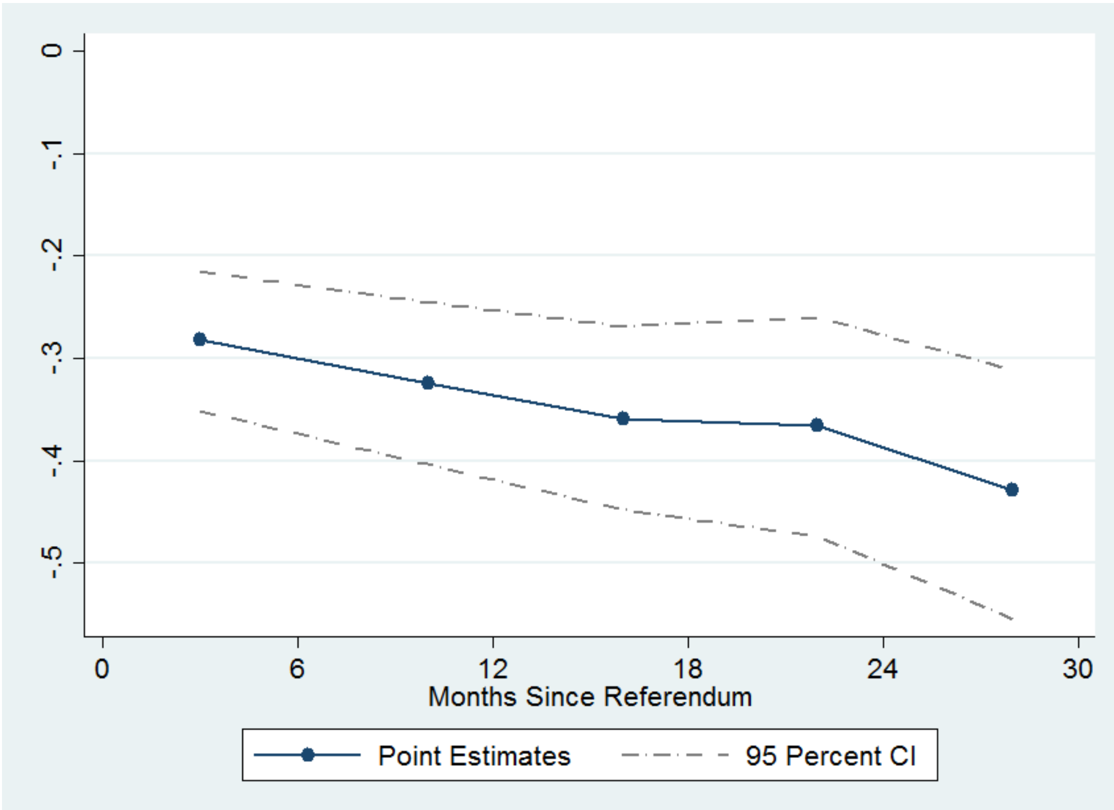

Notes: Sample includes only communities that passed MEA at some point. Elasticities are calculated for each six-month period by regressing town-month changes in log usage on the observed change in log price. Confidence intervals are constructed by bootstrap.