

NBER WORKING PAPER SERIES

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

Tatyana Deryugina
Alexander MacKay
Julian Reif

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

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

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

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Tatyana Deryugina, Alexander MacKay, and Julian Reif. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

ABSTRACT

Economic theory suggests that demand is more elastic in the long run relative to the short run, but evidence on the empirical relevance of this phenomenon is scarce. We study the dynamics of residential electricity demand by exploiting price variation arising from a natural experiment: the introduction of an Illinois policy that enabled communities to select electricity suppliers on behalf of their residents. Using a flexible difference-in-differences matching approach, we estimate a one-year price elasticity of -0.16 and three-year elasticity of -0.27. We also present evidence that consumers increased usage ahead of these announced price changes. Finally, we project that the price elasticity converges to a value between -0.30 and -0.35 after ten years. Our findings highlight the importance of accounting for consumption dynamics when evaluating energy policy.

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

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

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

1 Introduction

Economic theory suggests that demand is typically more elastic in the long run relative to the short run. When consumption depends on goods that are durable or habit-forming, consumers may take years to respond fully to a price change (Topel and Rosen, 1988; Becker et al., 1994). The presence of dynamics in household consumption poses fundamental challenges to demand estimation. In any period, consumption may depend upon both current and past prices, and, when consumers are forward looking, also upon future prices. When prices fluctuate, the demand response will reflect a mix of both short-term and long-term changes in consumption. Finally, unless one imposes restrictive functional form assumptions, unbiased estimation requires a source of exogenous price variation large enough to produce detectable effects in longer-run behavior.

In this paper, we develop and apply a flexible econometric model that decomposes changes in residential electricity consumption into effects arising from current and past price changes, as well as expected future price changes. A central output of our model is an estimate of how consumption changes over time with respect to a (hypothetical) permanent price change. The demand-side dynamics captured by our model are important to the electricity sector, where suppliers, market regulators, and policymakers make decisions with long-run ramifications. Failing to account for such dynamics will lead to biased estimates of policy changes, such as the long-run effect of a carbon tax on emissions reductions. Indeed, we estimate that household electricity consumption is twice as elastic in the long run, compared to one year after a permanent price change.

Our unique empirical setting allows us to estimate our dynamic model using large, long-lasting, and plausibly exogenous variation in residential electricity prices. In doing so, we provide the first quasi-experimental estimate of the long-run price elasticity of residential electricity demand, a category that accounts for 36 percent of U.S. electricity consumption (U.S. Energy Information Administration, 2017b). We find that the long-run elasticity is between -0.30 and -0.35, twice the magnitude of the one-year elasticity of -0.16. This difference in magnitude matters when predicting the effects of taxes, estimating policy incidence, and calculating returns to investment.

We estimate the model using monthly consumption and price data from 2007 to 2014 from the largest utility in Illinois, ComEd, whose territory encompasses approximately 70 percent of residential consumers in the state. During this period, Illinois implemented a municipal aggregation program,¹ which allowed individual communities to select new electricity suppliers (and correspondingly, new rates) on behalf of their residents. For communities that implemented aggregation, the results of this program were large, long-lasting price changes which, as we discuss below, were plausibly exogenous with respect to expected future usage. Our setting provides a clean natural experiment: aggregation customers continue to receive their electricity bill from the

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

utility in the same format as before, so the price variation in our analysis is not confounded with other billing changes. Moreover, Illinois employs a linear price schedule for residential electricity, consisting of a modest fixed fee and a constant marginal price. With rare and short-lived exceptions, aggregation affected only the marginal price of electricity, greatly simplifying our analysis.

Our empirical approach combines a difference-in-differences methodology with the matching estimator developed by Abadie and Imbens (2006, 2011). Matching estimators are particularly well-suited to our study because electricity usage is highly seasonal, and these seasonal patterns vary substantially across different communities and over time. The relatively large number of ComEd communities that did not pass a referendum on aggregation (479 in our sample of 768) in combination with a lengthy pre-period provides an excellent empirical setting for a matching estimator. We find that our matching estimator obtains more precise estimates than a traditional non-matching estimator, and we view our application as a useful demonstration of matching for applied researchers, in the vein of Fowlie et al. (2012). In contrast to their paper, we employ subsampling to calculate confidence intervals for a richer space of estimates that do not have pre-existing formulas for standard errors.

We match each aggregation community to five “nearest neighbors” that did not pass a referendum on aggregation.² We construct the matching criteria using communities’ monthly electricity usage profiles from 2008 and 2009. This matching period long precedes our natural experiment: more than ninety percent of the referenda in our sample are held after February 2012, over two years later. Our identifying assumption is that the average observed differences in usage between aggregation communities and their matched controls in the post-period are caused only by aggregation. In support of this assumption, we document that usage patterns between aggregation communities and matched controls are parallel after the matching period but prior to the referenda. Thus, systematic differences in usage arise only after the referenda. Because we observe usage at a high (monthly) frequency and community referenda occur only during infrequent statewide elections,³ our finding of no effect in the months immediately preceding the referenda suggests that communities did not select into aggregation based on expected usage changes. Whether or not a community pursued aggregation was likely influenced by social and political factors, including loyalty to the utility and trust in the local government. To the best of our knowledge, expected future usage was not discussed when considering aggregation.⁴ In light of these factors and the absence of pre-trends, aggregation therefore provides plausibly exogenous price variation in our sample.

²As a robustness check, we also estimate a traditional difference-in-differences model without matching, using only communities that implemented aggregation. These results, presented in the Appendix, are similar.

³Referenda could be passed in March and November in even years and in April in odd years.

⁴Our understanding of the motivations behind aggregation is guided by conversations with industry participants and local government meeting minutes.

For transparency and simplicity, we first estimate a reduced-form relationship between post-referendum price and usage changes. We estimate that prices fell by 22 percent (0.25 log points) and usage increased by 5.1 percent in the 7-12 months following an aggregation referendum, relative to control communities that did not pass a referendum on aggregation. After the first year of implementation, our data have a second large price change, arising from the expiration of a long-term ComEd contract in June 2013. This event leads to a price decrease for control communities, which we also see reflected in usage. Despite these large level shifts in prices, the price elasticity smoothly declines from -0.14 in the first year, to -0.27 in the second year, and to -0.29 in the third year, illustrating the importance of long-run dynamics in this setting.

As discussed above, the reduced-form results reflect a mix of both short-run and longer-run responses. Moreover, the reduced-form model cannot estimate long-run responses beyond the sample window. To address these shortcomings, we estimate a forward-looking dynamic model of demand in which usage is a flexible function of past, current, and future prices. As before, our identifying variation comes only from price differences caused by aggregation.

Because our model is forward looking, we must make an assumption about how expectations are formed. This is particularly important in our setting because the ComEd and average aggregation supply prices converged in June 2015, one year after our usage panel ends. If consumers expected this price convergence, then anticipatory behavior could have mitigated their response to the price decrease that municipal aggregation initially caused. We therefore estimate the dynamic model under two different assumptions about consumer expectations. If we assume consumers have perfect foresight, then we estimate a long-run elasticity of -0.35. If, by contrast, we assume that consumers employ a status quo (“no change”) forecast when forming expectations about the longer-run savings from aggregation, then we estimate a long-run elasticity of -0.31. In both cases, we estimate that the elasticity converges to its long-run value after approximately ten years.

Our setting also allows us to test for anticipation. Following the passage of the aggregation referendum and choice of a new supplier, all residents were notified by mail of the new price and the exact month it took effect. We find that consumers responded in anticipation of the price change: usage increased shortly after passage of the referendum, but before the actual price decrease several months later. Unfortunately, we lack data to identify the primary mechanism driving this behavior, which could be rational or due to confusion. Regardless of the mechanism, this result provides compelling evidence that consumers begin adjusting their electricity usage prior to realized price changes, rejecting the myopic “partial adjustment” model typically estimated in the energy literature (Hughes et al., 2008; Alberini and Filippini, 2011; Blázquez et al., 2013).

The primary contribution of our study is a credible, quasi-experimental estimate of the long-run price elasticity of residential electricity demand. Generators and distributors require forecasts of the long-run demand response to price changes to invest optimally in capacity and infrastructure.

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

Existing short-run and long-run estimates of the price elasticity of electricity demand vary widely: from nearly zero to about -0.9 in the short run (one year or less) and from -0.3 to about -1.1 in the long run.⁵ None of the long-run estimates is based on quasi-experimental variation, and many rely on state-level data and dynamic panel models, which include lagged consumption as an independent variable.⁶ Consistency in those models generally requires strong assumptions about the form of serial correlation, and Alberini and Filippini (2011) show that they are particularly sensitive to the exact specification used. Conversely, our main approach makes relatively few assumptions, and our forward-looking dynamic model is substantially more flexible than what has previously been estimated.

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

We also contribute to the large literature on forward-looking demand. It is challenging to provide persuasive evidence that consumers react to future prices; indeed, pre-period changes in outcome variables are often interpreted as model misspecification rather than forward-looking be-

⁵Also, a growing literature investigates the impact of real-time pricing (e.g., Wolak, 2011; Allcott, 2011; Jessoe and Rapson, 2014). The elasticity we identify here is fundamentally different from the elasticity estimated in the real-time pricing literature, which reflects intra-day substitution patterns as well as any overall reductions in electricity.

⁶See Alberini and Filippini (2011) for a review. Some have argued that a state's average price of electricity is exogenous because it is regulated (Paul et al., 2009) or because the unregulated component is driven by national trends (Bernstein and Griffin, 2005). However, electricity rates may be set based on the anticipated cost of electricity to suppliers and that cost, in turn, may be based on anticipated demand. Therefore, it is not possible to separate supply-side variation stemming from national changes in fuel prices from demand-side variation without explicitly constructing instruments, and we are not aware of any papers in this literature that employ instruments.

havior (Malani and Reif, 2015). However, similar to Gruber and Köszegi (2001), we argue that anticipatory responses in our setting represent a consumer reaction to aggregation, and not an unrelated trend, because they appear only following the passage of the aggregation referendum. Additionally, our rich data allow us to estimate a forward-looking model under fewer assumptions than is typically required. To avoid being too demanding of the data, prior studies have generally estimated these models in the form of an Euler equation (Becker et al., 1994; Malani and Reif, 2015), which requires assuming exponential discounting. By contrast, the long length of our panel allows us to estimate a flexible model of demand that imposes no structure on the form of discounting.

The rest of this paper is organized as follows. Section 2 presents a conceptual framework that captures the dynamics of electricity demand. Section 3 discusses electricity market regulation and municipal aggregation in Illinois. Sections 4 and 5 describe our data and empirical approach, respectively. Section 6 presents results. Section 7 discusses the implications of our main results, and Section 8 concludes.

2 Framework for Consumption Dynamics

Residential electricity usage is unlikely to adjust immediately to price changes. It takes time to change habits, such as turning off the lights or turning down the air conditioning when away from home. Usage also depends on the energy efficiency of durables such as dishwashers, dryers, and air conditioners, which are purchased infrequently and are continuously replaced within a population. Finally, some consumers may need time to learn that the electricity price has changed, especially if the benefit of tracking price changes is small relative to the cost of paying attention.

The presence of such adjustment costs suggests that the long-run response to a price change will exceed the short-run response. Moreover, if consumers are forward-looking, they may respond in anticipation of future price changes. Our preferred empirical specification therefore models consumption as a non-parametric function of past and future prices. Here, we present a simple framework demonstrating that our empirical approach is consistent with a standard rational model of a forward-looking consumer. We emphasize, however, that the empirical specification we employ for our estimates does not require rationality.

To illustrate the dynamics inherent in electricity usage more formally, we employ the habit model of Becker et al. (1994). For a derivation of the results presented below, the interested reader should consult the Appendix. Utility in each period depends on y_t , the consumption of electricity in that period, and on y_{t-1} , the consumption of electricity in the previous period.⁷ The consumer's problem is:

⁷Similarly, 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.

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

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

$$W_0 = \sum_{t=1}^{\infty} R^{t-1} (x_t + p_t y_t),$$

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

We assume that consumers are forward-looking, and, for expositional purposes, we assume they have perfect foresight.⁸ Finally, to illustrate the dynamics that can arise in our setting, we assume utility is quadratic.⁹ Under these assumptions, the demand equation is:

$$y_t = \alpha_1 y_{t-1} + \alpha_2 y_{t+1} + \alpha_3 p_t, \quad (1)$$

where the coefficients in (1) depend on the parameters of the quadratic utility function. The “adjustment cost” model frequently estimated in the energy demand literature corresponds to the special case where consumers are myopic, in which case the demand equation simplifies to:

$$y_t = \theta_1 y_{t-1} + \theta_2 p_t. \quad (2)$$

In the forward-looking model (1), consumers adjust their consumption in anticipation of future price changes, which are implicitly captured by y_{t+1} . They do not do so in the myopic adjustment model (2). Prior studies have noted that one can therefore test the myopic model by testing whether demand responds to future prices (Becker et al., 1994; Gruber and Köszegi, 2001).

One can equivalently express demand as a function of all past and future prices by solving the second-order difference equation (1):

$$y_t = \sum_{s=0}^{t-1} \delta_{t-s} p_{t-s} + \sum_{s=1}^{\infty} \delta_{t+s} p_{t+s}. \quad (3)$$

The effect of a price change on consumption will depend on whether the change was anticipated and on how long consumers expect the price change to last. Given estimates of equation (3), the long-run effect of a permanent change in price, p^* , on consumption can then be calculated as

$$\frac{dy}{dp^*} = \sum_{s=0}^{t-1} \hat{\delta}_{t-s} + \sum_{s=1}^{\infty} \hat{\delta}_{t+s}. \quad (4)$$

An appealing feature of equation (3) is that it is more flexible than equation (1). For example, it will generate a valid estimate of dy/dp^* even if consumers do not discount the future exponentially.

⁸We consider alternative assumptions in our empirical implementation.

⁹Alternatively, to analyze dynamics near a steady state, one could allow utility to be general and take a linear approximation to the first-order conditions, which would yield identical equations.

This non-parametric specification is rarely estimated, however, because most studies do not have data panels of sufficient length. We do not face this constraint, so we base our empirical approach on equation (3) rather than equation (1).

3 The Illinois Electricity Market

The provision of electricity to residential customers consists of two components: supply and distribution. Suppliers generate or purchase electricity and sell it to customers, and distributors provide the infrastructure to deliver the electricity and often handle billing. Illinois has two regulated electricity distributors: Commonwealth Edison Co. (“ComEd”) and Ameren Illinois Utilities (“Ameren”). Prior to 1997, they owned generating units as well as the distribution network. In 1997, the passage of the Consumer Choice Act allowed for competitive supply in the market, due to widespread agreement that, unlike distribution, electricity generation is not a natural monopoly (Illinois General Assembly, 1997). As part of the deregulation measures, the two utilities were encouraged to divest their generation assets. These policies led to the entry of several alternative suppliers into the market.

Customers are assigned their distributor on the basis of geographic location. ComEd, the distributor for whom we have usage data, serves northern Illinois and approximately 70 percent of the customers in the state. The price for electricity supplied by Ameren or ComEd is, by law, equal to their procurement cost and does not vary geographically.¹⁰ While customers have no choice in distributors, in 2002 residential and small commercial customers gained the ability to choose an alternative retail electric supplier (ARES) who would be responsible for supplying (but not delivering) their electricity.¹¹ However, the residential ARES market was practically nonexistent between 2002 and 2005. This was blamed on barriers to competition and a rate freeze that kept the default utility rate low. In 2006, the state removed some of these barriers and instituted a discount program for switchers, but this still had little effect on behavior. By 2009, only 234 residential customers had switched suppliers. By contrast, 71,000 small commercial, large commercial, and industrial customers had switched (Spark Energy, 2011).

Motivated by these patterns, the Illinois Power Agency Act was amended in 2009 to allow for municipal aggregation, whereby municipalities and counties could negotiate the purchase of electricity on behalf of their residential and small commercial customers. Townships gained this ability in 2012. To ensure that individual consumers retained the ability to choose their supplier, the amendment requires municipalities to allow individuals to opt out of aggregation.

To implement an opt-out aggregation program, municipalities must first educate their communities about aggregation using local media and community meetings, register the proposed aggre-

¹⁰Profits stem from delivery fees set by the Illinois Commerce Commission (DeVirgilio, 2006).

¹¹Large commercial and industrial customers gained this ability at the end of 1999.

gation program with the state, and hold a referendum. The wording of the referendum question is specified in the Illinois Power Agency Act and given in the Appendix. In most cases where the referendum is approved, multiple suppliers then 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. The two main ways in which suppliers differentiate themselves are price and the share of generation from renewable sources. Nearly all communities select the supplier with the lowest price, although environmental preferences occasionally induce communities to select a more expensive one.

When determining the bid or negotiating directly, each supplier obtains aggregate community-level usage data from the distributor. These usage data, along with electricity futures, are the main factors in each offered price.¹² Importantly, our analysis utilizes the same usage data, which reduces the likelihood that price changes are affected by confounding factors that are unobservable to us. Because many of communities' first contracts were in effect through the end of our usage data, the price variation we employ comes mainly from the first set of aggregation contracts in Illinois. Overall, municipal aggregation is very popular in Illinois: as of March 2016, 741 (out of about 2,100) communities had voted to implement it.

In our setting, the realized savings from aggregation come largely from the timing of the program. During our sample period, alternative suppliers were able to offer lower rates due to the unexpected boom in shale gas, while ComEd was locked into a long-term high-price procurement contract.

Customers in a community that has passed an aggregation referendum are automatically switched to the newly chosen electricity supplier unless they opt out by mailing in a card, calling, or filling out a form online.¹³ Aggregation officially begins at the conclusion of the opt-out process. From the consumer's point of view, the only visible change is the supply price of electricity on her bill, which is still issued by the incumbent distributor (Ameren or ComEd). The other items on the bill remain the same for all customers in the distributor's territory regardless of aggregation. Conveniently, this means that the price effects of aggregation will not be confounded by billing confusion. The Appendix includes a sample ComEd bill, a sample letter notifying households of aggregation, and a sample opt-out card.

¹²Suppliers may also base their bids on the number of electric space heat customers, which we do not observe. In Illinois, only about 10 percent of households heat their homes with electricity (U.S. Energy Information Administration, 2009).

¹³While we do not have an exact number, ComEd and several energy suppliers have told us that the opt-out rate is low. Community-specific opt-out rates mentioned in newspapers range from 3 to 10 percent (e.g., Lotus, 2011; Wade, 2012; Ford, 2013). The number of non-aggregated customers does, however, grow slowly over time because new residents who move to an aggregation community are not defaulted into the aggregation program. The few residential customers who had already opted into an ARES or into real-time pricing prior to aggregation are not switched over to the chosen supplier.

4 Data

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

We constructed the time series of ComEd electricity rates using ComEd ratebooks, which we obtained from the Illinois Commerce Commission. Prior to June of 2013, customers with electric space heating faced a lower rate than those with non-electric space heating. Because only about 10 percent of households in Illinois heat their homes with electricity (U.S. Energy Information Administration, 2009), we assume that the incumbent rate is equal to the non-electric space heating rate, which will be true for the majority of non-aggregation customers. Data on aggregation referenda dates, aggregation supply prices, and aggregation implementation dates were obtained from a variety of sources, including PlugInIllinois, websites of electricity suppliers, and municipal officials.

Table 1: Count of Aggregation Communities in Sample

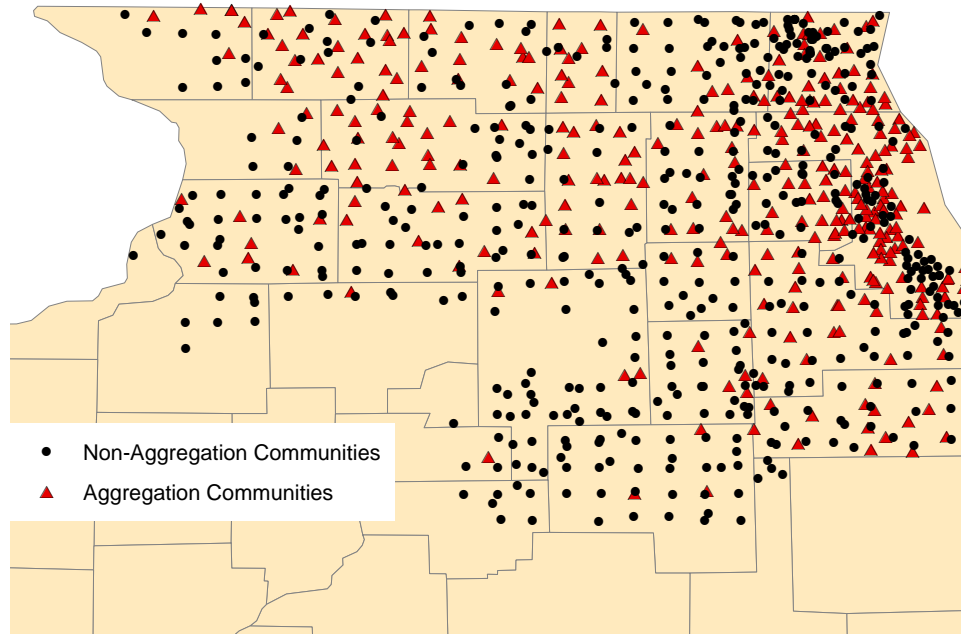
Referendum Date	Implemented	Passed, Not Implemented	Voted, Not Passed
November 2010	1	0	0
April 2011	18	0	0
March 2012	164	0	28
November 2012	57	5	2
April 2013	38	3	6
March 2014	8	1	0
November 2014	3	2	0
Total	289	11	36

As shown in Table 1, 300 communities in the ComEd territory passed a referendum on aggregation during our sample period, and 289 of those communities eventually implemented an aggregation program.¹⁴ In addition, 36 communities voted on, but did not pass, the referendum. The geographic locations of the 768 communities in our sample are displayed in Figure 1. Aggregation communities are well-dispersed throughout the ComEd territory but are slightly more prevalent in the greater Chicago area.

Our discussions with industry participants suggest that social and political factors, such as

¹⁴Five of these communities passed a referendum in November of 2014, five months after the end of our usage data.

Figure 1: Spatial Distribution of Communities in Sample



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

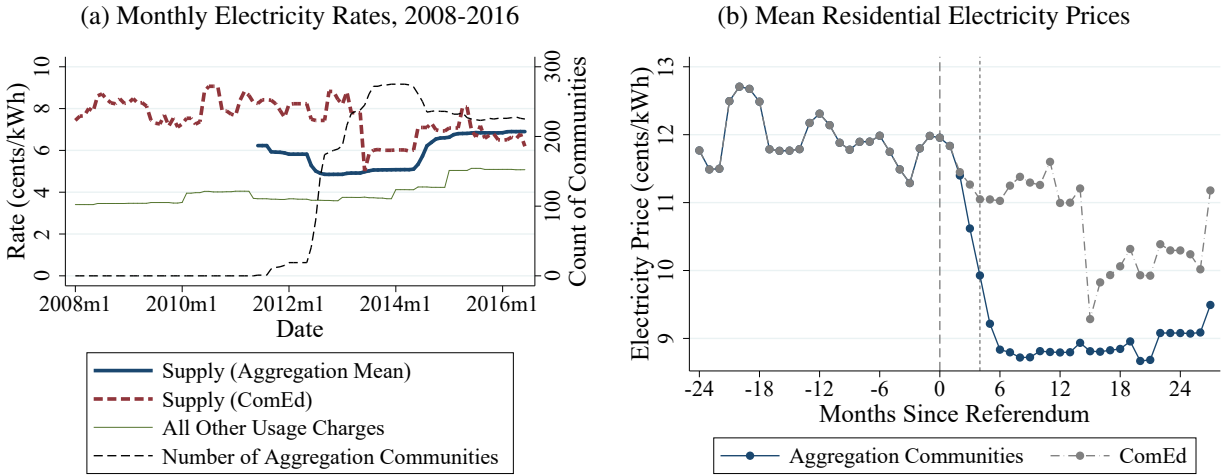
whether the local government should involve itself in negotiating electricity prices, played a role in a community's decision to hold a referendum.¹⁵ Anecdotally, the reasons why some communities voted against aggregation include: (1) lack of trust in the local government to secure savings relative to the incumbent; (2) loyalty to the utility; (3) concern about the environmental impact of the resulting electricity use increase; (4) a misunderstanding about the opt-out provision; and (5) the belief that choosing an electricity provider for residents was not an appropriate government function.

There is some variation in how long communities take to implement price changes after approving aggregation. The median length of time between passage of the aggregation referendum and commencement of the aggregation program is 4 months. At least 10 percent of aggregation communities switched suppliers within 3 months of the referendum, whereas 10 percent had not done so 6 months afterward. In most specifications, we construct estimates relative to the referendum date to capture any usage responses that occur prior to the actual price change.

Many states employ a "block pricing" schedule where the marginal price of electricity increases with quantity purchased. Illinois, by contrast, employs a constant marginal price and a moderate fixed fee, which simplifies our analysis because it reduces confusion over the "price" to which consumers might be responding. This constant marginal price can be broken down into three main

¹⁵These observations are also informed by the authors' attendance of a public hearing in Champaign, in-person discussions with ComEd and the Illinois Commerce Commission (Illinois' electricity regulator), and local government meeting minutes.

Figure 2: Prices for Aggregation and Non-Aggregation (ComEd) Communities



Notes: The thick blue line in Panel (a) displays the average supply rate among all communities that adopted aggregation. The first community adopted aggregation in June of 2011. Non-aggregation communities pay ComEd’s supply rate (thick, dashed red line in Panel (a)). The green line in Panel (a) displays the total of all other electricity rates on a consumer’s residential bill, which do not depend on whether a community has adopted aggregation. All displayed rates are for a single family residence with non-electric heating. The thin dashed line in Panel (a) indicates the cumulative number of communities that have implemented aggregation. Panel (b) displays the mean (total) electricity price for aggregation communities as a function of the time since referendum and compares it to the corresponding ComEd price. The short dashed line indicates the median implementation date relative to when the referendum was passed.

components: supply, delivery, and taxes/fees. Implementing aggregation entails a community signing a contract for a particular supply rate (the largest component of the marginal price) with an electricity supplier; non-aggregation communities pay the default ComEd supply rate. Thus, aggregation only affects the marginal price of electricity. Nearly all suppliers offered contracts with a constant supply rate, with most terms ranging from from 9 months to 3 years.¹⁶

Figure 2 demonstrates the price variation in our sample. The thick dotted red and thin green lines in Figure 2a display ComEd’s monthly supply rate and the total of all other usage rates, respectively, during and after our sample period. ComEd’s supply rate decreased significantly in 2013, when its long-term high-priced power contract expired. This price decrease was not a surprise, and we discuss its implications for the interpretation of our estimates in Section 6. In response to this price drop, several communities switched back to ComEd when their aggregation contracts expired in June of 2014. This reversal is visible in the black dashed line in Figure 2a, which displays the count of aggregation communities over time.

The blue line in Figure 2a shows the average monthly supply rate for aggregation communities, starting from when the first community implemented aggregation (June 2011). During our sample period (2008-2014), the average aggregation supply rate is always lower than the default ComEd supply rate, although the two rates converge in mid-2015. We also note that the month-to-month

¹⁶There was no significant variation among communities in other charges during this period. See the Appendix for more details.

variation in the ComEd rate is small relative to the level shocks realized by aggregation and the June 2013 ComEd contract expiration.

Another way to visualize the price variation in our sample is to plot the mean residential electricity prices for aggregation communities *relative to the referendum date* (Figure 2b). The displayed price combines the supply rate with the other usage charges from Figure 2a. For comparison, we also show the mean contemporaneous ComEd rate. The figure illustrates the two large price reductions during the time period spanned by our sample. Averaging the price changes along the x-axis, we observe a 25 percent reduction for communities upon implementation, and a 10 percent reduction for non-aggregation communities about 15 months after the average date of aggregation adoption, reflecting the June 2013 drop in the ComEd supply price. There was also significant cross-sectional variation in the aggregation price shocks (see Appendix Figure A.1), but for nearly all communities in nearly every month, aggregation reduced prices.

Variation in aggregation prices is due to differences in (i) the timing of the referenda, (ii) community procurement strategies, and (iii) the load profiles of the communities.¹⁷ The last source of variation raises endogeneity concerns in our setting: Do communities with greater expected usage responses receive different prices? By matching treatment communities to their nearest neighbors in terms of usage patterns, we mitigate some of this concern as long as the responses are proportional. Empirically, we use the same data (monthly community-level usage) that is used by suppliers to formulate their bids, which accounts for the relevant heterogeneity in load profiles. As a check for price endogeneity, we also estimate elasticities separately by percentiles of the realized price change, and we find no relationship. We discuss this further in Section 6.4.¹⁸

Both the aggregation-driven price changes and the 2013 drop in the ComEd price were likely more salient to customers than typical month-to-month changes in electricity rates. Aggregation is publicized in advance of the referendum, and each household receives a mailer informing them of the new aggregation price. Likewise, the drop in ComEd prices received statewide attention in the press.¹⁹ Other potential energy policies (e.g., carbon tax) that affect electricity prices are also likely to be anticipated and publicized. Aggregation therefore provides a potentially policy-relevant setting for assessing the residential usage response to such policies.

Finally, because ComEd's supply price does not vary geographically, we can use our raw data to

¹⁷For a discussion of these first two points, see the description of aggregation in Section 3.

¹⁸An additional concern is that the implementation of aggregation by some communities might influence the supply price paid by remaining ComEd customers through a reverse causality channel: as communities with favorable load profiles implement aggregation, the communities that remain with ComEd could be more costly to serve, causing ComEd to increase its price. However, ComEd still controlled a massive pool of customers throughout our sample period, and there is no large rival purchaser to challenge their buyer power. Indeed, we observe declining ComEd rates throughout our sample, suggesting this is not an important concern.

¹⁹"Some ComEd customers to see lower prices," *Chicago Tribune*, April 1, 2013. "What's happening to your electric bill June 1?" Citizens Utility Board, May 30, 2013.

calculate the savings aggregation communities obtained from switching suppliers. Specifically, we multiply aggregation communities' observed electricity usage by the price difference each month and aggregate over our sample period. We estimate that the residential aggregation consumers in our sample saved \$566 million through June 2014.

5 Empirical Strategy

5.1 Overview

Measuring price differences brought about by aggregation is straightforward. Each treatment community has a constant marginal electricity supply rate that is set within a few months of passing a referendum. The counterfactual supply price for each treatment community is the ComEd price, as non-aggregation supply prices do not vary within ComEd's service area. Similarly, there is no meaningful geographic variation in the other components of marginal price (see Appendix for additional details). Thus, the primary challenge is estimating counterfactual electricity usage for aggregation communities.

We estimate the community-level usage response to changes in the electricity price by matching communities that implemented aggregation (the "treated" group) to communities that did not (the "control" group) based on their pre-aggregation electricity usage. We exploit the fact that the implementation of aggregation depended on a referendum, which was plausibly exogenous with respect to expected changes in electricity usage. We provide some empirical evidence for this identifying assumption in subsequent sections.

We estimate two empirical models. Our "reduced-form" model regresses usage changes, as calculated by our matching estimator, on corresponding aggregation-driven price shocks to measure the contemporaneous response of usage to electricity prices. However, because the post-aggregation price shocks are not constant, this model does not directly account for dynamics that might arise from adjustment costs, learning, or expectations. In other words, the elasticity estimated by this model is accurate only to the extent that the post-period price changes in our sample can be reasonably approximated as a one-time, permanent change in price. Nevertheless, the reduced-form model allows us to test whether the price elasticity is constant over time, i.e., whether it is necessary to estimate a dynamic model. The estimates can also be used to detect anticipation effects and quantify heterogeneity.

To disentangle the short-run response to price changes from the long-run response, we estimate a second model that is motivated by the dynamic model presented in Section 2. This specification allows us to quantify explicitly how the price changes in each future and past time period affect contemporaneous usage, thereby allowing us to estimate long-run elasticities.

As a robustness check, we present results from a standard (non-matching) difference-in-differences model in the Appendix. That analysis exploits variation in policy timing among aggregation com-

munities rather than variation in policy adoption across all communities in the territory. We find similar patterns to our main results.

5.2 Difference-in-Differences Matching Estimator

We estimate counterfactual electricity usage by matching communities that implemented aggregation (the “treated” group) to communities that did not (the “control” group), based on their pre-aggregation usage. Our setting is an ideal application for a matching estimator. The large, diverse set of control communities makes it likely that a nearest-neighbor matching approach will successfully find a suitable comparison group. Additionally, we have enough pre-aggregation usage data for internal validation of the approach.

Specifically, we apply a difference-in-differences adjustment to the bias-corrected matching estimator developed by Abadie and Imbens (2006, 2011). For each of the 289 treated communities, we use 2008–2009 electricity usage data to identify the five nearest neighbors from the 479 control communities available in our sample. We average annual log usage and monthly log deviations from annual usage across 2008 and 2009 to construct 13 match variables. We standardize the variables and use an equal-weight least squares metric to calculate distance. That distance is then used to select (with replacement) the five nearest neighbors for each treated community.²⁰ We use these nearest neighbors to construct counterfactual usage, and we employ standard difference-in-differences techniques to adjust for pre-period differences. The identifying assumption is that, conditional on 2008–2009 usage, the passage of aggregation and subsequent price changes are unrelated to anticipated electricity use. We provide evidence that this assumption is reasonable by showing that trends in usage for the control and treated groups remain parallel after the matching period but before the passage of aggregation.

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

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

²⁰The results are very similar if we instead select the single nearest neighbor or the ten nearest neighbors.

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

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

where

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

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

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

$$\hat{\tau}_t^{DID} = \hat{\tau}_t - \frac{1}{N_s} \sum_{s=1}^{N_s} \hat{\tau}_{-s}, \quad (5)$$

where N_s indicates the number of periods in the year prior to the policy change.²¹ 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. As the policy change is at the community level, and we only include communities that implement aggregation, our estimate is the effect of the treatment on the treated.²²

Failing to account for anticipation effects in a difference-in-differences framework can lead to

²¹ $N_s = 12$ for monthly data. For biannual data, $N_s = 2$.

²²If we were to interpret our results as a household-level estimate, we would consider it an intent-to-treat effect (conditional on implementing aggregation), as a small proportion of households opted out.

biased treatment effect estimates when consumers are forward-looking (Malani and Reif, 2015). Ample evidence from the energy literature suggests this concern is valid. For example, Myers (2016) finds that expected future heating costs appear to be capitalized fully into housing values. While Allcott and Wozny (2014) reject full capitalization of gasoline prices into used vehicle prices, they nonetheless find that vehicle prices are significantly affected by future expected fuel costs. In our setting, the aggregation electricity price is announced months before the actual price change takes place. Thus, consumers may respond to future price changes by, for example, placing less weight on energy efficiency when replacing old appliances, changing their thermostat program, or changing 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 referendum date, rather than when aggregation was implemented.²³ 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.

5.3 Advantages of Matching Estimators with Electricity Data

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 (see Appendix Figure A.2), and the degree of seasonality varies widely across the different communities in our sample. Identifying control communities with usage profiles similar to aggregation communities can therefore greatly increase precision.

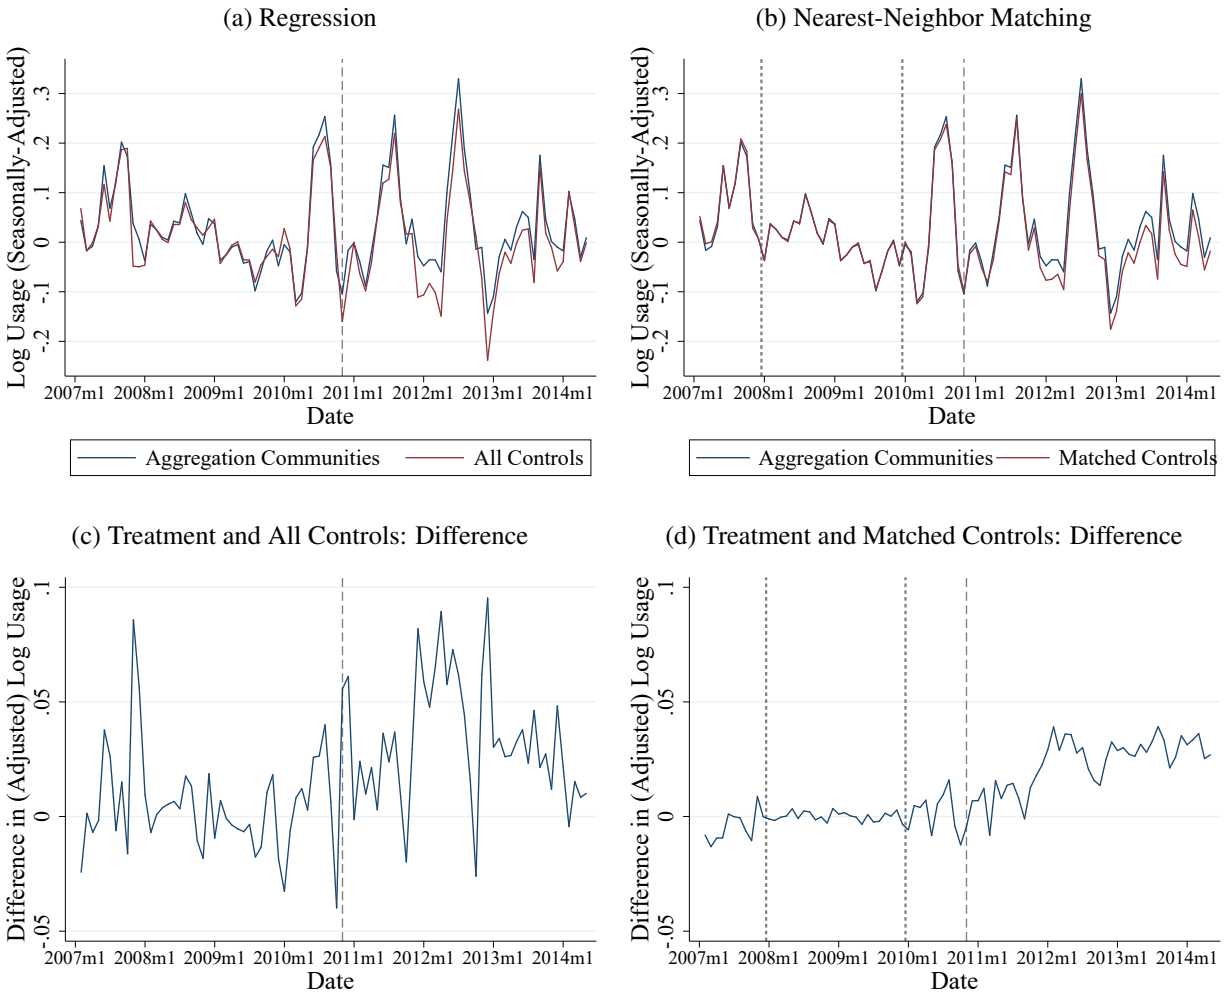
Figure 3 provides a demonstration of this benefit. Panel (a) displays electricity usage adjusted for community-level monthly seasonal patterns both for treated (aggregation) communities and communities that never passed aggregation. Even after accounting for community-specific seasonality, usage varies greatly within and across years: the largest peak occurs in July 2012, which corresponds to a record heat wave. By contrast, summer peaks are much less pronounced in 2009 and 2013, when the summers were mild. The difference between these two time series, which corresponds to an event study regression with community-specific month-of-year fixed effects, is displayed in Panel (c).²⁴ The increase in the difference is visible beginning in late-2011, which can be attributed to the implementation of aggregation, but this difference is quite noisy. The heterogeneity in seasonal patterns poses a challenge for a standard regression that compares treated communities to all control communities in the sample: it is difficult to estimate an effect when the baseline month-to-month divergence in usage is of the same order of magnitude as the effect.

Panels (b) and (d) of Figure 3 show analogous plots for the nearest-neighbor matching ap-

²³It 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.

²⁴The fixed effects are calculated using monthly data from 2008 and 2009.

Figure 3: Comparing Regression to Nearest-Neighbor Matching



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

proach discussed above. Panel (d) shows again that the difference in log usage between treatment and (matched) control communities increases beginning in late-2011. The difference in Panel (d) exhibits far less noise than the difference displayed in Panel (c), because the matching estimator selects only those control communities that are similar to treated communities. This method of selection allows the matching estimator to generate more precise estimates than the standard difference-in-differences estimator.

To see whether matching also helps improve the similarity between treated and control communities on other dimensions, we matched the names of the communities in our ComEd sample to data

obtained from the 2005-2009 American Community Survey (ACS) Summary File. We obtained ACS matches for 286 out of 289 aggregation communities, and 385 out of 479 non-aggregation communities.²⁵ Compared to non-aggregation communities, communities that implemented aggregation are significantly larger, younger, and more educated (see Appendix Table A.1 for more details). They are also less white, and have more expensive and slightly newer housing. However, the 2010 electricity usage per capita is very similar for aggregation and non-aggregation communities. After matching on monthly and annual electricity usage, the weighted pool of matched controls is more similar to the aggregation communities. With two exceptions (percent black and longitude), the differences between non-aggregation and aggregation communities are less likely to be statistically significant, which indicates that matching on usage also selects control communities with more similar socioeconomic characteristics.

5.4 Estimating Elasticities

Reduced Form

In our first specification, we regress community-specific estimates of the change in usage on the observed community-specific price changes. The community-specific estimate $\hat{\tau}_{it}^{DID}$ is the single-community analog of equation (5). It serves as the outcome variable in the following regression:

$$\hat{\tau}_{it}^{DID} = \beta_g \cdot \Delta \ln p_{it} + \eta_{it}, \quad (6)$$

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

Equation (6) allows us to estimate period-specific elasticities that show how the response changes over time with respect to the referendum date. In our main results, we run separate, pooled regressions with g corresponding to six-month intervals. The parameter of interest, β_g , corresponds to the average elasticity over interval g .

Accounting for Dynamics

As discussed earlier, electricity usage is unlikely to adjust immediately to contemporaneous price changes. Motivated by the conceptual framework in Section 2, we estimate the dynamics of adjustment by regressing our matching estimates of log usage changes on lags and leads of log price

²⁵Aggregation 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.

changes:

$$\hat{\tau}_{it}^{DID} = \sum_{s=-L_1}^{L_2} \delta_s \cdot \Delta \ln p_{i(t-s)} + \eta_{it}. \quad (7)$$

The number of leads in the regression is equal to L_1 , and the number of lags is L_2 . For consistency with the reduced-form specification, we estimate our model in logs, but our results are very similar if we instead estimate the model in levels, as given by equation (3). Although we do not observe usage past June 2014, we observe prices through 2016, allowing us to estimate equation (7) without losing observations. To reduce sensitivity to outliers, we estimate this model using median (quantile) regressions instead of least squares.²⁶

A central challenge in the estimation of forward-looking models is how to account for expectations. Equation (7) implicitly assumes that consumers have perfect foresight with respect to future prices. This is a strong assumption, however, and can lead to biased estimates if it is inaccurate. We therefore also estimate an alternative specification inspired by Anderson et al. (2013), who present evidence that consumers have status quo (“no change”) expectations, at least with respect to gasoline prices.²⁷ Specifically, we assume that consumers expect that aggregation contracts starting after June 2014 will provide savings equal to the average price difference in the twelve months leading up to the expiration date of the previous contract. Thus, differences between the perfect foresight and the status quo estimates are driven by price changes occurring after June 2014 (see Figure 2a).²⁸

5.5 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.²⁹ Subsampling, like bootstrapping, obtains a distribution of parameter estimates by sampling from the observed data.

For each of $N_b = 500$ subsamples, we select without replacement $B_1 = R \cdot \sqrt{N_1}$ treated

²⁶Point estimates from least squares regressions are similar, but the subsampling routine we describe below does not converge for some subsamples in the least squares specification. This reflects the sensitivity of least squares to outliers and the relatively small size of each subsample.

²⁷One could also employ other models of belief formation such as rational or adaptive expectations. See Malani and Reif (2015) for discussion of different modeling approaches.

²⁸We do not assume status quo expectations for changes in price prior to June 2014 because, as we show later, consumers clearly reacted to the pending implementation of aggregation, which is at odds with a status quo model of expectations. Because the expiration of aggregation contracts was not as widely advertised, however, the corresponding change in prices may not have been anticipated by consumers.

²⁹Abadie and Imbens (2006, 2011) provide a formula for the standard errors of bias-corrected matching estimators of average treatment effects. Our panel data structure and the use of match-specific indexing for the control communities relative to the treated communities preclude a simple implementation of this formula. Further, our main estimates are not simple average treatment effects.

communities and $B_0 = R \cdot \frac{N_0}{\sqrt{N_1}}$ control communities, where R is a tuning parameter (Politis and Romano, 1994) and N_0 is the number of control communities. As before, N_1 is the number of treated (aggregation) communities. For each subsample, we calculate $\hat{\tau}_b$. The matching estimator of the average treatment effect on the treated converges at rate $\sqrt{N_1}$ (Abadie and Imbens, 2006, 2011). The estimated CDF of $\hat{\tau}$ is given by:

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

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

Subsampling requires a large number of effective observations (i.e., treated units) in each subsample, but it also requires that this number be small relative to the total number of effective observations in the full sample. We employ $R = 3$ ($B_1 = 51$) for the confidence intervals and standard errors reported in our main tables, which prioritizes the large-sample properties within each subsample. All other standard errors and confidence intervals are calculated similarly to those of $\hat{\tau}$. For example, we calculate elasticity estimates of β_g for each subsample to generate confidence intervals for our elasticities. Table A.4 in the Appendix compares standard errors for different values of the tuning parameter, and shows they are robust to different values of R .

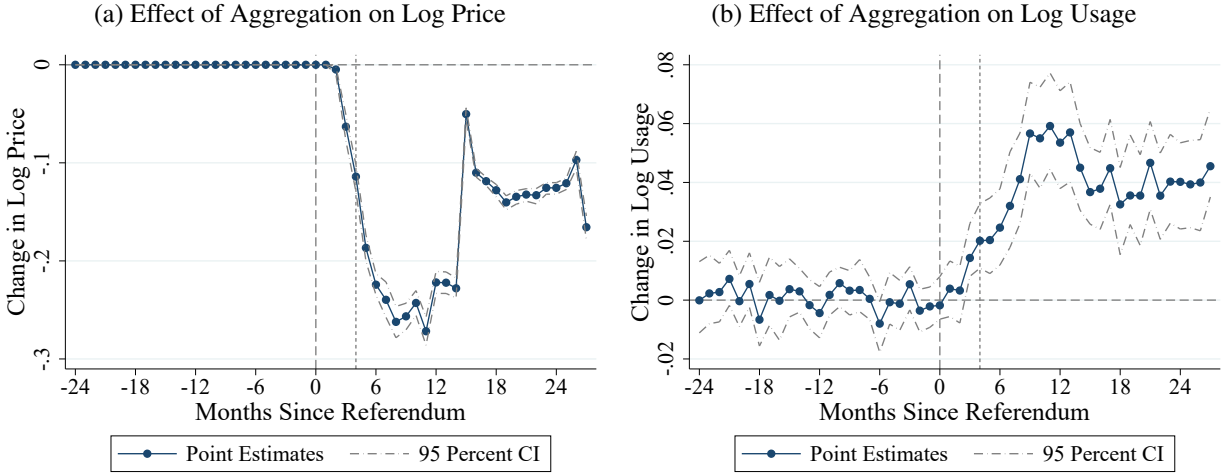
6 Results

6.1 Reduced-Form Model

We first show that electricity prices fell substantially and persistently following the passage of aggregation referenda. Figure 4a displays the average change in log prices for aggregation communities, relative to their matched controls. The price change is exactly equal to zero in the pre-period because the treated communities face the same ComEd supply prices as their matched control communities during that time period. Within 12 months of passing the referendum, prices in aggregation communities decrease by nearly 0.3 log points relative to control communities, although they rebound significantly a few months into the second year. The rebound is attributable to a sharp decrease in ComEd’s supply price in June of 2013 (see Figure 2a). Nonetheless, aggregation prices stay at least 10 percent lower than the control communities for most of the remaining estimation period.

Figure 4b displays the corresponding estimates for electricity usage. Prior to the referendum, the difference in usage between aggregation and control communities is relatively constant and never significantly different from zero. We emphasize that this result is not mechanical, as our 2008-2009 matching period predates the vast majority of aggregation instances by at least two

Figure 4: Reduced-Form Effects of Aggregation on Electricity Prices and Usage



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

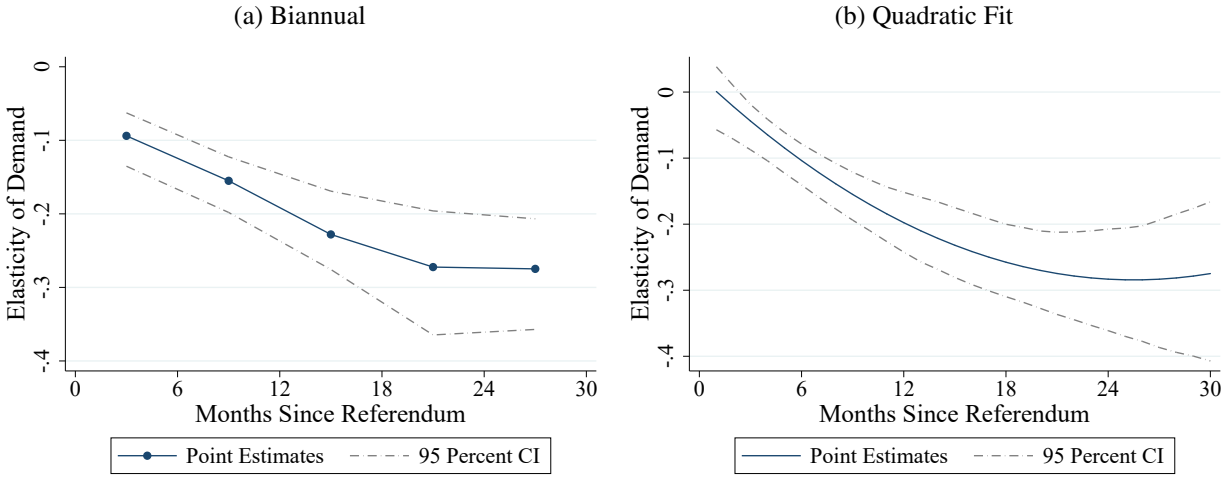
years (see Table 1).³⁰ Following the referendum, usage in aggregation communities increases by as much as 0.06 log points in the first year after the referendum. When the ComEd price decreases (about fifteen months after most referenda), the difference in usage between the treatment and control group shrinks to about 0.04 log points and then stabilizes. Because the aggregation prices are stable around this time period (see Figure 2b), this pattern provides persuasive evidence that both aggregation and non-aggregation communities are responding to their respective price changes. Thus, our results are not driven merely by the salience of the price change brought about by aggregation.

Figure 5a displays reduced-form estimates of the price elasticity of demand. To reduce noise, we estimate elasticities at the biannual, rather than monthly level. Monthly elasticity estimates can be found in the Appendix (Figure A.3 and Table A.2). 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 and ruling out a constant price elasticity. In other words, usage does not appear to respond fully to price changes in the near term. 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 5b, are very similar.

The time-varying elasticity shown in Figure 5 is not simply due to a delay between referenda and when prices actually take effect. While it is true that usage patterns in Figure 4b reflect the

³⁰Specifically, 270 of the 289 aggregation communities have virtually no overlap between the matching period and the pre-period estimates in Figure 4b.

Figure 5: Reduced-Form Price Elasticities, Biannual



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.

lag between referenda and implementation, so do the price patterns in Figure 4a. As the usage changes are scaled by the price changes, the implementation lag does not matter when calculating elasticities. For example, if the price elasticity were in fact constant over time, we would estimate it as such, even in the presence of implementation delays.

Our main results from the reduced-form matching approach are summarized in Table 2.³¹ Because the price decrease following the aggregation referendum is not constant throughout the estimation period, the estimated elasticity at any given time captures the response to the longer-run price decrease due to aggregation, to the drop in ComEd prices in June 2013, as well as to monthly variations in the ComEd rate. At the same time, because the elasticity is growing smoothly throughout the estimation period, it appears that the long-run effects of aggregation dominate the shorter-run responses to changes in the ComEd rate. In the following section, we estimate a model that explicitly accounts for the dynamic effects of these price changes, thereby allowing us to estimate the long-run elasticity.

6.2 Dynamic Model

To account for the dynamics of electricity demand, we model consumption as a function of past and future prices. Therefore, how consumers form expectations about future prices may matter for our estimates. The precise nature of these expectations may be relevant because – as shown

³¹Appendix Table A.3 reports the corresponding yearly estimates.

Table 2: Matching Estimates of the Effect of Aggregation on Usage and Prices

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

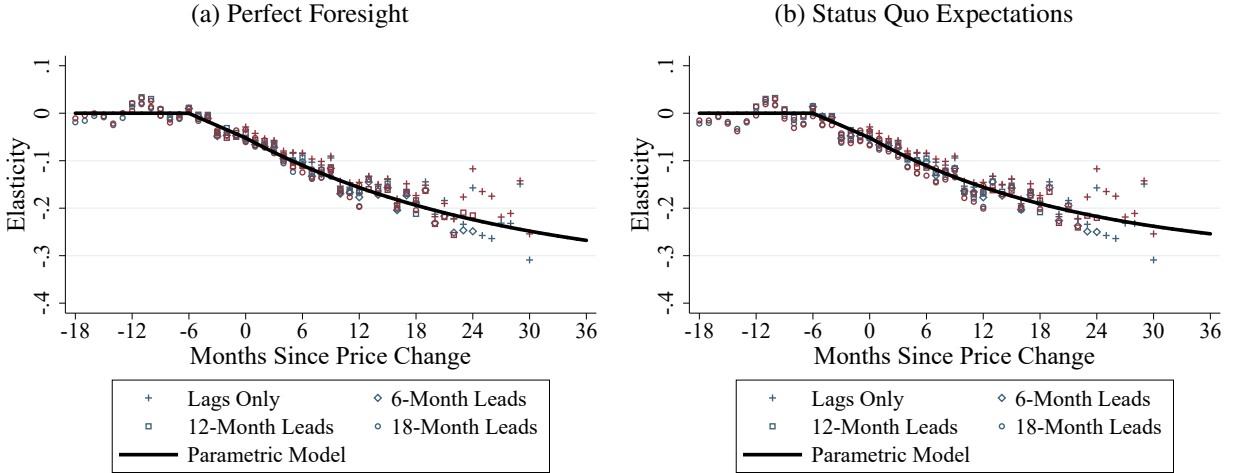
Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each aggregation community is matched to the five non-aggregation communities with the most similar usage in 2008 and 2009. The number of price observations corresponds 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.

in Figure 2a – average prices for aggregation and non-aggregation communities converged in June 2015. If this convergence was anticipated by consumers, then it might have affected the usage response in earlier years.

We therefore estimate our dynamic model under two different sets of assumptions regarding expectations of future prices. In the first specification, given by equation (7), we assume that consumers have perfect foresight. In the second specification, we assume that consumers have perfect foresight for the initial contract and for renewals through June 2014, which is the end of our usage sample, but that after June 2014 consumers form status quo expectations for aggregation prices relative to the ComEd price. In other words, consumers expect that price differences for new contracts will equal to the average difference over the twelve months leading up the prior contract's expiration. These assumptions are plausible because aggregation prices and contract lengths are widely publicized, and many of the initial contracts last until June 2014, while prices following the end of a contract are uncertain. While our findings below indicate that consumers respond ahead of price changes, thus rejecting a status quo model of belief formation, it is possible that anticipation with respect to the first aggregation contract is due to the unusual salience of aggregation during the time period we analyze. As both sets of assumptions have perfect foresight through the end of our usage sample, differences between these two sets of estimates will be driven by price changes occurring after June 2014.

The markers in Figure 6a correspond to non-parametric monthly estimates under the assumption that consumers have perfect foresight. We use four different sets of leads and lags, starting with 18-month leads and ending with a specification with lags only. The sum of the number of leads and lags in each specification is equal to 30. To account for compositional effects, we also estimate versions of equation (7) with community fixed effects (red markers). Figure 6b presents

Figure 6: Estimated Price Elasticities: Dynamic Model



Notes: These graphs compare non-parametric and parametric estimates of the dynamic elasticity curve. The plotted points display non-parametric estimates derived from equation (7), where each set of points is calculated using 30 lags and leads, beginning with the lead displayed in the legend. Each point displays the estimated elasticity as a function of the number of months since a price change. The point corresponding to the cumulative effect at lag l is calculated as $\sum_{s=-L_1}^l \hat{\delta}_s$, where L_1 is the number of leads included in the specification. The red points correspond to non-parametric models estimated with community fixed effects; the blue points are from specifications with no fixed effects. The solid lines display the estimates derived from the parametric model constrained by equation (8).

corresponding estimates when we assume status quo expectations rather than perfect foresight. Each marker in these two sub-figures represents the cumulative electricity usage change (in percent) in response to an anticipated 1-percent price change that begins in month 0 and persists. Note that period $t = 0$ in Figure 6b corresponds to the date of the price change, but in Figures 4 and 5 period $t = 0$ corresponds to the date of the referendum.³² Thus, it is helpful to recall that the median time between the referendum and implementation is about four months.

Overall, the estimates in the two plots are quite similar to each other. They indicate that consumers begin responding to price changes several months before such changes occur, which we will later confirm also happens in our reduced-form specification. The pre-price-change elasticities are small, around -0.05. Following the price change, the usage response grows over time, with an implied elasticity of about -0.1 six months post-implementation (ten months after the referendum for the median community) and an elasticity of about -0.2 eighteen months post-implementation.

As with other studies of long-run elasticities, we face the challenge that our dataset spans a limited time window: the non-parametric results estimate the monthly response with reasonable accuracy only up to 24 months beyond the date of a price change. To calculate elasticities beyond the time span covered by the leads and lags, we estimate a parametric specification that restricts the relationships among the set $\{\delta_s\}$ in equation (7). We fit the cumulative elasticities, which

³²Reduced-form results where $t = 0$ corresponds to implementation of aggregation are presented in Section 6.3.

correspond to a permanent price change, to a four-parameter adjustment curve. Based on the non-parametric estimates presented in Figure 6, we assume an exponential form for the lags and a linear form for the leads:

$$\beta_s = \sum_{\tau=-\infty}^s \delta_\tau = \left(\gamma_1 - \frac{\gamma_1}{\gamma_2}s \cdot \mathbf{1}[s \leq 0]\right) \cdot \mathbf{1}[s \geq \gamma_2] + \gamma_3(1 - \exp(\gamma_4 s)) \cdot \mathbf{1}[s \geq 0] \quad (8)$$

The elasticity corresponding to an anticipated, contemporaneous price shock is equal to γ_1 , and the long-run elasticity is equal to $\gamma_1 + \gamma_3$. The parameters γ_2 and γ_4 govern the speed of adjustment. The point estimates for the parameters $(\hat{\gamma}_1, \hat{\gamma}_2, \hat{\gamma}_3, \hat{\gamma}_4)$ are $(-0.052, -6.034, -0.296, 0.036)$ under perfect foresight expectations and $(-0.052, -6.000, -0.255, 0.044)$ under status quo expectations. The fitted models are plotted as solid lines in Figure 6 and line up closely with the unrestricted estimates from equation (7).

Table 3 reports estimates from our dynamic specification (7), when we impose the parametric restriction specified in (8). Again, these estimates correspond to the effect of a permanent, anticipated increase in price beginning at time $t = 0$ on usage in month t .³³ The first column displays the results under the perfect foresight assumption, and the second column displays results assuming status quo expectations for contracts that begin after June 2014. The reported coefficients are quite similar for up to two years after the price change. The estimated elasticities, which reach -0.16 after one year and -0.22 at the end of two years, are slightly smaller than what we find in our reduced-form estimates from Section 6, which capture the net post-referendum usage response relative to the net price change at each point in time. We should not be surprised to find some differences, as the reduced-form estimates capture a mix of effects arising from past, present, and future price changes, while the dynamic estimates in this section explicitly isolate effects from a price change s months ago. That the two sets of estimates are similar is partly due to the fact that the aggregation price shocks are fairly constant and persistent. The similarity also suggests that long-run dynamics are dominant in our setting.

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

Our model estimates that roughly 80 percent of the long-run consumer response occurs within three years after the price change. We forecast that it takes about ten years for consumers to fully adjust, at which point our estimate is nearly identical to the long-run response. Based on the 95

³³For annual averages, which correspond more closely to the reduced-form estimates, see Table A.8 in the Appendix.

Table 3: Parametric Estimates of the Dynamic Elasticity Curve

Period After Price Change	Perfect Foresight	Status Quo Expectations
Month -3	-0.026*** [-0.043, -0.015]	-0.026*** [-0.043, -0.015]
Contemporaneous	-0.052*** [-0.075, -0.033]	-0.052*** [-0.076, -0.034]
Month 6	-0.110*** [-0.140, -0.082]	-0.111*** [-0.140, -0.084]
Month 12	-0.156*** [-0.196, -0.117]	-0.156*** [-0.196, -0.118]
Month 18	-0.194*** [-0.247, -0.146]	-0.191*** [-0.240, -0.144]
Month 24	-0.224*** [-0.288, -0.168]	-0.217*** [-0.280, -0.164]
Month 30	-0.248*** [-0.326, -0.184]	-0.238*** [-0.314, -0.176]
Month 36	-0.268*** [-0.359, -0.195]	-0.254*** [-0.344, -0.187]
Year 5	-0.315*** [-0.462, -0.225]	-0.288*** [-0.442, -0.210]
Year 10	-0.345*** [-0.606, -0.243]	-0.306*** [-0.601, -0.221]
Year 25	-0.348*** [-0.841, -0.245]	-0.307*** [-0.825, -0.221]
Long Run	-0.348*** [-1.007, -0.245]	-0.307*** [-1.079, -0.221]

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. The elasticities are interpreted as the cumulative effect of a permanent one-percent price change on current usage. Estimates are constructed from 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 aggregation community is matched to the five non-aggregation communities with the most similar usage in 2008 and 2009. 95 percent confidence intervals are displayed in brackets and are constructed via subsampling.

percent confidence intervals generated by the two sets of assumptions, we conclude that the ten-year elasticity is bounded by -0.60 and -0.22.

The finding that consumers are more elastic in the long run than in the short run is consistent with several mechanisms, including habit formation, learning, and continuous appliance replacement within a community. Because we do not have data on consumer behavior beyond electricity consumption, we do not attempt to distinguish among these different adjustment channels.

6.3 Anticipation Effects

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

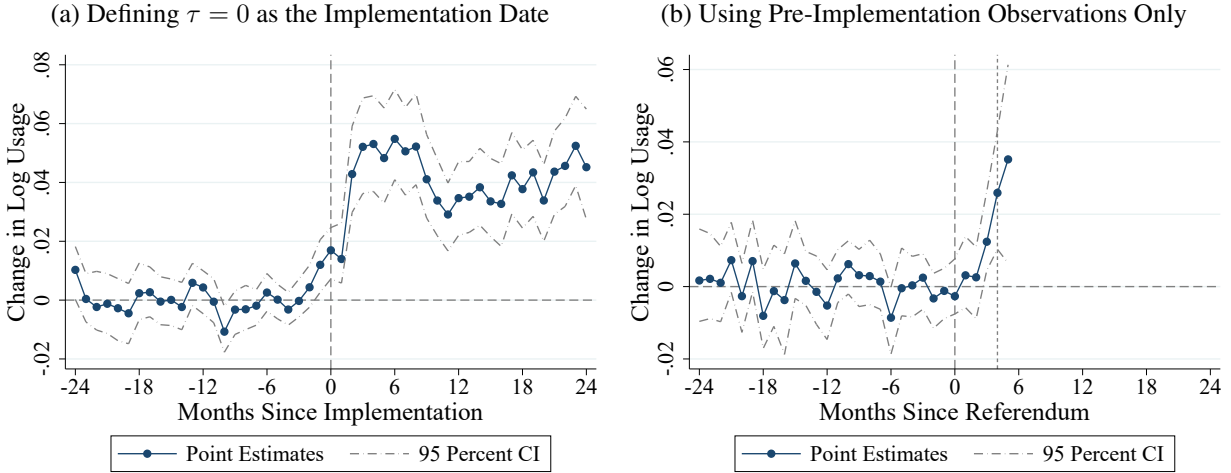
Both dynamic specifications reported in Table 3 find a small but significant anticipation response in the six months prior to a price change. To further validate this result, we re-estimate the reduced-form specification (6), setting $t = 0$ to be the date that aggregation was *implemented* (which corresponds to the date of the price change) rather than the date of the aggregation *referendum*. The results are shown in Panel (a) of Figure 7. The price difference between aggregation communities and their matched controls is exactly zero prior to implementation. However, usage begins increasing three months prior to the price change.

To explicitly isolate the anticipation effect, Panel (b) of Figure 7 displays estimates of changes in electricity usage relative to the *referendum* date, using only pre-implementation data.³⁴ Electricity usage increases steadily and significantly 3-5 months after the referendum despite the fact that prices have not yet changed for the observations in this sample. Specifically, usage is 0.012 log points higher 3 months after the referendum and 0.035 log points higher 5 months after the referendum, confirming the existence of non-trivial anticipation effects in our sample. We do not observe any significant change in usage prior the referendum.

Together, these results suggest that the referendum and/or the subsequent price change an-

³⁴Figure A.4 displays results for communities that passed a referendum but never implemented aggregation. Although the estimates are noisy, they suggest that there was no increase in usage due to the referendum itself in those communities.

Figure 7: Anticipation Effects of Implementing Aggregation on Log Usage



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

nouncement alerted consumers to the impending price decrease and caused them to increase their usage, perhaps by placing less weight on energy efficiency when replacing old appliances or by changing their electricity usage habits (relative to their control counterparts). Because mailers were sent to all residents informing them of the exact month of the price change, it is unlikely that customers were confused about the timing of the price change, although we cannot definitively rule out that possibility. More generally, while our estimates are consistent with a standard forward-looking model of rational consumers, they are also consistent with confusion and other behavioral mechanisms. Regardless of the precise mechanism, what is important is that the anticipation effect we estimate is policy relevant; other policies that affect the price of electricity are likely to be implemented in a similar manner and generate pre-implementation effects similar to those we find here.

6.4 Robustness Checks

One endogeneity concern raised by our empirical approach is whether the magnitude of the price change that a community experiences is correlated with its expected demand elasticity. For example, suppliers might offer lower rates to more inelastic customers, as more elastic customers would demand more at the same price and drive up supply costs. We think this is unlikely because most individual communities are small relative to the total market. However, we empirically check for this possibility by splitting our treated communities into seven equal groups based on the price change they experienced in the first two years following their referenda. We then calculate elasticities separately for each group. Figure A.5 plots these estimates. As expected, we find no evidence

of a relationship between the price change and the *ex post* estimated elasticity.

We chose January 2008 through December 2009 as our matching period because it allowed us to more closely match controls to treated communities while also leaving us with a fairly long post-matching period to test for pre-trends. We have also estimated our model using an alternative matching period of February 2007 through January 2009. These results are very similar to those in the paper. Additionally, our results are robust to the number of neighbors chosen for the matching procedure. We obtain similar results if we match to the single nearest neighbor or the ten nearest neighbors, instead of the five we use as our baseline. For these alternative number of matches, the yearly reduced-form results and selected coefficients from the dynamic elasticity curve are reported in Tables A.7 and A.8 in the Appendix.

We have also estimated the effect of aggregation on electricity usage using a difference-in-differences reduced form approach without matching. In this analysis, which we discuss in detail in the Appendix, we exploit the variation of the timing of implementation among aggregation communities. The results are qualitatively similar to the results presented in the main text. In particular, we again find a price elasticity that grows post-referendum but no evidence of pre-trends, supporting the identifying assumption that passage of aggregation was not prompted by growth in electricity usage.

Finally, it is worth noting that Illinois is similar to the U.S. as a whole along several dimensions related to electricity consumption. We describe the relationship between key demographics in Section A.3 in the Appendix. To understand how our results generalize to geographies with different demographics, the Appendix also explores how our elasticity estimates vary by socioeconomic characteristics.

7 Implications of the Results

Emissions Pricing Our empirical estimates show that the price elasticity of residential electricity demand is substantially larger in magnitude in the long run than in the short run. This has two important implications for regulation that targets electricity generation and subsequently affects the price of electricity (e.g., emissions pricing). First, the long-run change in electricity consumption following such regulation is likely to be significantly larger than estimates using the short-run elasticity would imply. Second, the share of the overall regulatory burden borne by consumers is likely to be smaller than expected, at least in partial equilibrium. Though specific quantitative conclusions are beyond the scope of this paper, we explain in broad strokes below how these different parameters matter for policy evaluation.³⁵

Policymakers who wish to use a tax to target a specific level of emissions must correctly quan-

³⁵Specific impacts depend on two other highly uncertain parameters: the price elasticity of electricity supply and the extent of pass-through to consumers in regulated electricity markets.

tify how equilibrium electricity consumption responds to changes in taxes. In perfectly competitive markets, it is well known that the fall in equilibrium quantity following a small tax increase depends closely on both the demand and supply elasticities (Salanie, 2011):

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

where Q is quantity, t is a per-unit tax, P is price, and ϵ_s and ϵ_d are the absolute values of supply and demand elasticities, respectively. It is clear from this expression that more elastic demand corresponds to a greater fall in equilibrium electricity consumption. Underestimating the demand elasticity by using its short-run value will thus lead to an understatement of emissions reduction from a small carbon tax.³⁶ The degree of understatement increases with the supply elasticity. For example, if the supply elasticity is $\epsilon_s = 0.5$, then the change in quantity when $\epsilon_d = 0.3$ is 2.25 times larger compared to $\epsilon_d = 0.1$. In other words, quantity is projected to fall by more than twice as much when the larger (e.g., long-run) demand elasticity is used. If supply is very elastic, say, $\epsilon_s = 5$, the change in quantity is 2.89 times larger when $\epsilon_d = 0.3$ compared to $\epsilon_d = 0.1$.

It is also well known that the partial equilibrium incidence of a tax depends on the relative elasticity of supply and demand (Salanie, 2011). In perfectly competitive markets, the relative incidence of a tax, defined as the ratio of the tax burden borne by consumers to that borne by producers, can be shown to equal the ratio of supply and demand elasticities, ϵ_s/ϵ_d . Weyl and Fabinger (2013) extend the analysis of incidence to imperfectly competitive markets and show that, while the incidence formula becomes more complicated, it is still a function of the demand elasticity. Thus, underestimating the demand elasticity by using its short-run value leads to an overestimate of the share of the overall tax burden borne by consumers.

Though residential electricity prices are often regulated, which affects how taxes are passed through to customers and the incidence of these taxes, competitive electricity supply is increasingly common. As of 2016, fifteen states had retail choice for electricity supply (Morey and Kirsch, 2016).

Carbon Leakage The price elasticity of electricity demand is also relevant for general-equilibrium simulations of carbon leakage (e.g., Wing and Kolodziej, 2009; Elliott and Fullerton, 2014; Baylis et al., 2014; McKibbin et al., 2014). In these models, at least one good is subject to a carbon tax. If there is at least one good that generates carbon emissions but is not taxed (e.g., because it is produced in a location with no carbon tax or by a sector not subject to the tax), consumption of the non-taxed carbon-emitting goods can increase. Elliott and Fullerton (2014) argue that taxing only electricity sector emissions is a likely scenario for carbon policy, at least initially, even though other

³⁶Because 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.

sectors also generate carbon emissions. In this case, the price elasticity of demand for electricity is a key determinant of the extent of carbon leakage. More broadly, this price elasticity is important for any general equilibrium model that includes the electricity sector.

Supply-Side Dynamics Conceptually, the long-run price elasticity of demand is also important for any study of long-term dynamics in electricity markets, whether such efforts are aimed at understanding the effects of deregulation (Olsina et al., 2006), generator entry decisions (Takashima et al., 2008), the diffusion of renewable energy (Kumbaroğlu et al., 2008), or the effects of environmental regulation (Cullen, 2014; Bushnell et al., 2017). For example, a larger long-run price elasticity dampens entry incentives by making it harder to recover any initial investment. The assumptions made about the price elasticity of demand vary widely in the literature, and, with the exception of Bushnell et al. (2017), the sensitivity of results with respect to this parameter is often not discussed. Our study provides a well-identified estimate that can be employed in future research in this area.

Market Design Finally, our findings matter for the organization of electricity markets. For example, there is some debate as to whether to use a downward-sloping or vertical demand curve in forward capacity auctions for the Midcontinent Independent System Operator (MISO) electricity area (Cook, 2016a,b). As the relevant time period for these auctions spans multiple years, our findings lend more weight to using a downward-sloping demand curve in this setting.

8 Conclusion

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

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

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

References

- Abadie, A. and G. W. Imbens (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica* 74(1), 235–267.
- Abadie, A. and G. W. Imbens (2008). On the failure of the bootstrap for matching estimators. *Econometrica* 76(6), 1537–1557.
- Abadie, A. and G. W. Imbens (2011). Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics* 29(1), 1–11.
- Alberini, A. and M. Filippini (2011). Response of residential electricity demand to price: The effect of measurement error. *Energy Economics* 33(5), 889–895.
- Allcott, H. (2011). Rethinking real-time electricity pricing. *Resource and Energy Economics* 33(4), 820–842.
- Allcott, H. and N. Wozny (2014). Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics* 96(5), 779–795.
- Anderson, S. T., R. Kellogg, and J. M. Sallee (2013). What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management* 66(3), 383–403.
- Baylis, K., D. Fullerton, and D. H. Karney (2014). Negative leakage. *Journal of the Association of Environmental and Resource Economists* 1(1/2), 51–73.
- Becker, G. S., M. Grossman, and K. M. Murphy (1994). An empirical analysis of cigarette addiction. *American Economic Review*, 396–418.
- Bernstein, M. A. and J. Griffin (2005). Regional differences in the price-elasticity of demand for energy. *RAND Corporation Technical Report*.
- Blázquez, L., N. Boogen, and M. Filippini (2013). Residential electricity demand in Spain: new empirical evidence using aggregate data. *Energy Economics* 36, 648–657.
- Bushnell, J. B., S. P. Holland, J. E. Hughes, and C. R. Knittel (2017). Strategic policy choice in state-level regulation: The epa’s clean power plan. *American Economic Journal: Economic Policy*. forthcoming.
- Cook, A. D. (2016a). Board orders negotiation in auction disagreement. *RTO Insider*. <https://www.rtoinsider.com/miso-board-capacity-auction-imm-26870/>, Accessed March 13, 2017.
- Cook, A. D. (2016b). MISO considering changes to proposed auction design. *RTO Insider*. <https://www.rtoinsider.com/miso-auction-design-imm-26062/>, Accessed March 13, 2017.
- Cullen, J. A. (2014). Dynamic response to environmental regulation in the electricity industry. mimeo.
- DeVirgilio, B. (2006). A discussion of the deregulation of the energy industry in Illinois and its effects on consumers. *Loyola Consumer Law Review* 19, 256.

- Elliott, J. and D. Fullerton (2014). Can a unilateral carbon tax reduce emissions elsewhere? *Resource and Energy Economics* 36(1), 6–21.
- Filippini, M., B. Hirl, and G. Masiero (2015). Rational habits in residential electricity demand. *Working Paper*.
- Ford, M. A. (2013). Five percent of Normal residents opted out of electric aggregation. *Pantagraph*. http://www.pantagraph.com/news/local/five-percent-of-normal-residents-opted-out-of-electric-aggregation/article_48bcb79e-64f2-11e2-acb9-0019bb2963f4.html, Accessed March 13, 2017.
- Fowlie, M., S. P. Holland, and E. T. Mansur (2012). What do emissions markets deliver and to whom? Evidence from Southern California’s NOx trading program. *American Economic Review* 102(2), 965–993.
- Gruber, J. and B. Köszegi (2001). Is addiction “rational”? Theory and evidence. *The Quarterly Journal of Economics* 116(4), 1261–1303.
- Hughes, J. E., C. R. Knittel, and D. Sperling (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal* 29(1), 113–134.
- Illinois General Assembly (1997). Electric service customer choice and rate relief law of 1997 (220 ilcs 5/article xvi).
- Ito, K. (2014, February). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review* 104(2), 537–63.
- Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review* 104(4), 1417–1438.
- Karnitschnig, M. (August 26, 2014). Germany’s expensive gamble on renewable energy. *The Wall Street Journal*.
- Kumbaroğlu, G., R. Madlener, and M. Demirel (2008). A real options evaluation model for the diffusion prospects of new renewable power generation technologies. *Energy Economics* 30(4), 1882–1908.
- Lotus, J. (2011). Oak Park seals green deal on electric. *OakPark.com*. <http://www.oakpark.com/News/Articles/10-18-2011/Oak-Park-seals-green-deal-on-electric/>, Accessed March 13, 2017.
- Malani, A. and J. Reif (2015). Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform. *Journal of Public Economics* 124, 1–17.
- McKibbin, W. J., A. C. Morris, and P. J. Wilcoxon (2014). Pricing carbon in the US: A model-based analysis of power-sector-only approaches. *Resource and Energy Economics* 36(1), 130–150.
- Morey, M. J. and L. D. Kirsch (2016). Retail choice in electricity: What have we learned in 20 years? Technical report. Christensen Associates Energy Consulting, LLC.
- Myers, E. (2016). Are home buyers myopic? Evidence from housing sales. E2e Working Paper 024.
- Olsina, F., F. Garcés, and H.-J. Haubrich (2006). Modeling long-term dynamics of electricity markets. *Energy Policy* 34(12), 1411–1433.
- Paul, A. C., E. C. Myers, and K. L. Palmer (2009). A partial adjustment model of us electricity demand by region, season, and sector. RFF Discussion Paper 08-50.
- Politis, D. N. and J. P. Romano (1994). Large sample confidence regions based on subsamples under minimal assumptions. *The Annals of Statistics*, 2031–2050.
- Poterba, J. M. (1984). Tax subsidies to owner-occupied housing: An asset-market approach. *The Quarterly Journal of Economics* 99(4), 729–52.

- Reiss, P. C. and M. W. White (2005). Household electricity demand, revisited. *The Review of Economic Studies* 72(3), 853–883.
- Ryoo, J. and S. Rosen (2004). The engineering labor market. *Journal of Political Economy* 112(S1), S110–S140.
- Salanie, B. (2011). *The Economics of Taxation*. MIT press.
- Spark Energy (2011, July). The history of electricity deregulation in Illinois. <http://www.sparkenergy.com/blog/illinois-electricity-deregulation-history/>.
- Takashima, R., M. Goto, H. Kimura, and H. Madarame (2008). Entry into the electricity market: Uncertainty, competition, and mothballing options. *Energy Economics* 30(4), 1809–1830.
- Topel, R. and S. Rosen (1988). Housing investment in the united states. *Journal of Political Economy* 96(4), 718–740.
- United States Census Bureau (2016). Quickfacts. <https://www.census.gov/quickfacts/>.
- U.S. Energy Information Administration (2009). Household energy use in illinois. https://www.eia.gov/consumption/residential/reports/2009/state_briefs/pdf/IL.pdf, Accessed August 30, 2017.
- U.S. Energy Information Administration (2017a). Average retail price of electricity to ultimate customers. <https://www.eia.gov/electricity/data.cfm>, Accessed February 2, 2017.
- U.S. Energy Information Administration (2017b). February 2017 monthly energy review. <http://www.eia.gov/totalenergy/data/monthly/#electricity>, Accessed March 3, 2017.
- Wade, P. (2012). Urbana residents to get opt-out letters on electric rates soon. *The News-Gazette*. <http://www.news-gazette.com/news/local/2012-06-15/urbana-residents-get-opt-out-letters-electric-rates-soon.html>.
- Weyl, E. G. and M. Fabinger (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy* 121(3), 528–583.
- Wing, I. S. and M. Kolodziej (2009). The regional greenhouse gas initiative: Emission leakage and the effectiveness of interstate border adjustments. Technical report. Available at SSRN: <https://ssrn.com/abstract=1448748>.
- Wolak, F. A. (2011). Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment. *American Economic Review: Papers and Proceedings* 101(3), 83–87.

A Appendix (For Online Publication)

A.1 Conceptual Framework Derivations

As shown in Becker et al. (1994), the effect of a price change on consumption at a particular point in time depends on whether or not the change was anticipated; when the change occurred; and whether the change is temporary or permanent. This can be shown by solving the second-order difference equation (1) to obtain:

$$y_t = K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} \alpha_3 p_{t+s-1} + K_2 \sum_{s=0}^{t-1} (\lambda_2)^s \alpha_3 p_{t-s-1} + (\lambda_2)^t \left(y_0 - K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} \alpha_3 p_{s-1} \right) \quad (\text{A.1})$$

where

$$K_1 = \frac{\lambda_1}{\alpha_2 (\lambda_1 - \lambda_2)}$$
$$K_2 = \frac{\lambda_2}{\alpha_2 (\lambda_1 - \lambda_2)}$$

with roots

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

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

Equation (A.1) 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 ($t \rightarrow \infty$), the third term in equation (A.1) becomes zero, so that consumption no longer depends on the initial condition, y_0 . The long-run effect of a permanent change in price in all periods is

$$\frac{dy}{dp^*} = K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} \alpha_3 + K_2 \sum_{s=0}^{\infty} (\lambda_2)^s \alpha_3$$

which corresponds to equation (4) from the main text as $t \rightarrow \infty$.

It is straightforward to show that the solution to the first-order difference equation (2), the myopic “adjustment cost” model, depends only on past prices, and not on future prices.

A.2 Data Appendix

Consistent Usage Data

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. These changes appears

as large, discrete jumps 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 the community 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. However, the communities that are removed in this fashion are primarily small communities that did not implement aggregation.

Other Components of the Electricity Bill

All fixed fees and remaining usage rates are nearly identical across the aggregation and non-aggregation communities in our sample. The average fixed fee for customers residing in ComEd service territories during our sample period is \$12.52 per month. We ignore the fixed fees in our analysis, because they do not vary across communities and because our independent variable is the difference in the (log) marginal price between aggregation and non-aggregation communities. In rare instances the aggregation supplier charged an additional fixed fee, but these scenarios were short-lived. Municipal tax rates do vary across communities, but the variation is small and they rarely change over time. For our analysis, we use the median tax rate across ComEd communities (0.557 cents/kWh). The price variable in our estimating equations is equal to the aggregation supply rate plus all other usage rates if the community has implemented aggregation; otherwise, it is equal to the ComEd supply rate plus all other usage rates.

A.3 Demographic Characteristics

Illinois versus the United States

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

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

Elasticities by Demographic Characteristics

Next, we investigate how much variation in the reduced-form price elasticity can be explained by demographic characteristics. Doing so helps us understand both the distributional effects of policies that affect electricity prices and the generalizability of our results. For this exercise, we regress the log net usage change at time t on the contemporaneous log price change and add interactions between the log price change and indicator variables for whether or not a community is in the top half of the distribution for x^j , a characteristic of interest:

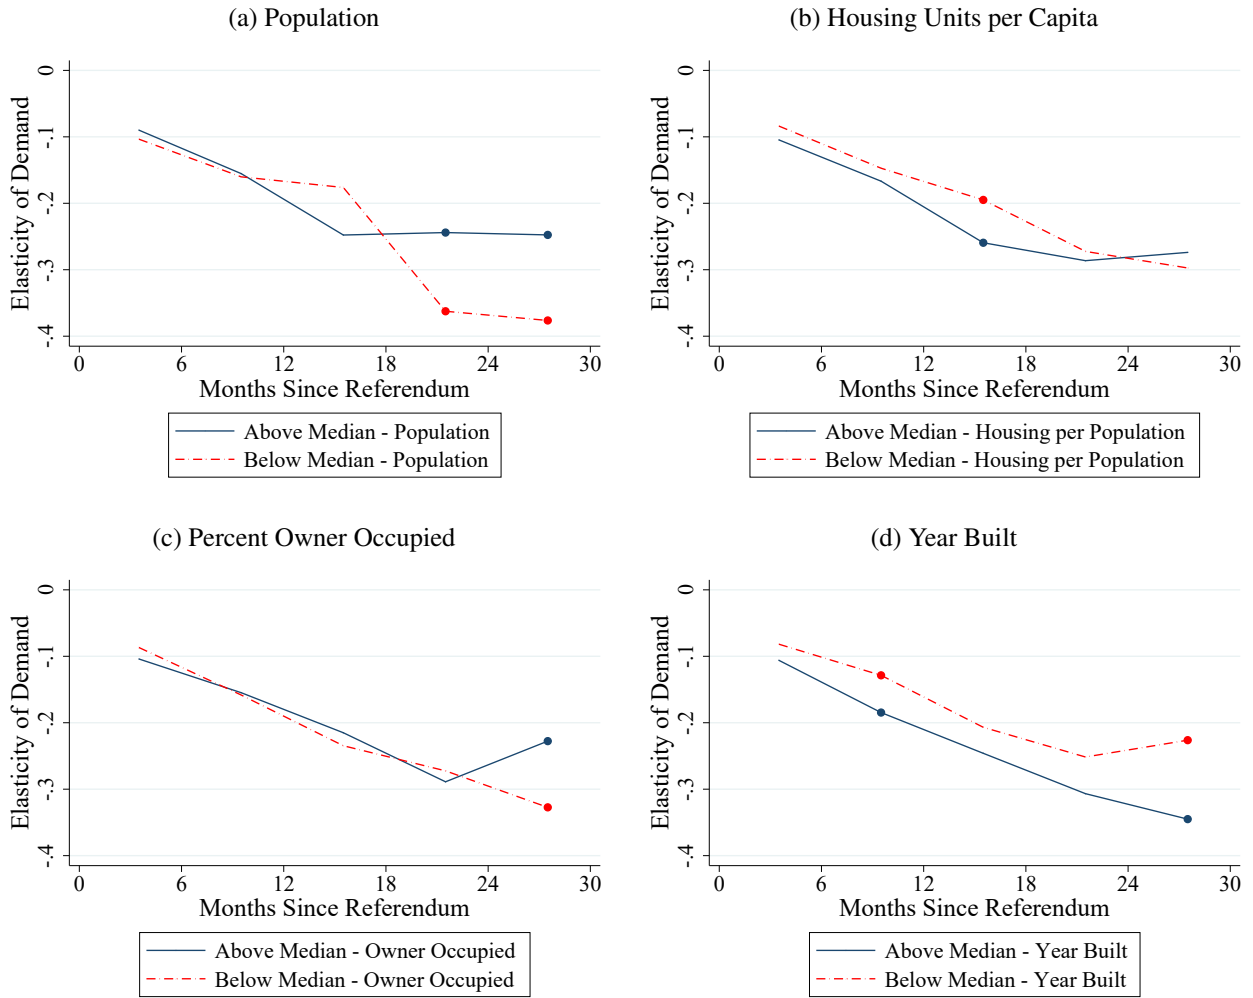
$$\hat{\tau}_{it}^{DID} = \beta_g \cdot \Delta \ln p_{it} + \sum_{j=1}^J (\beta_j \cdot \Delta \ln p_{it} \cdot \mathbf{1}[x_i^j > \text{median}(x^j)]) + \eta_{it}.$$

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

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

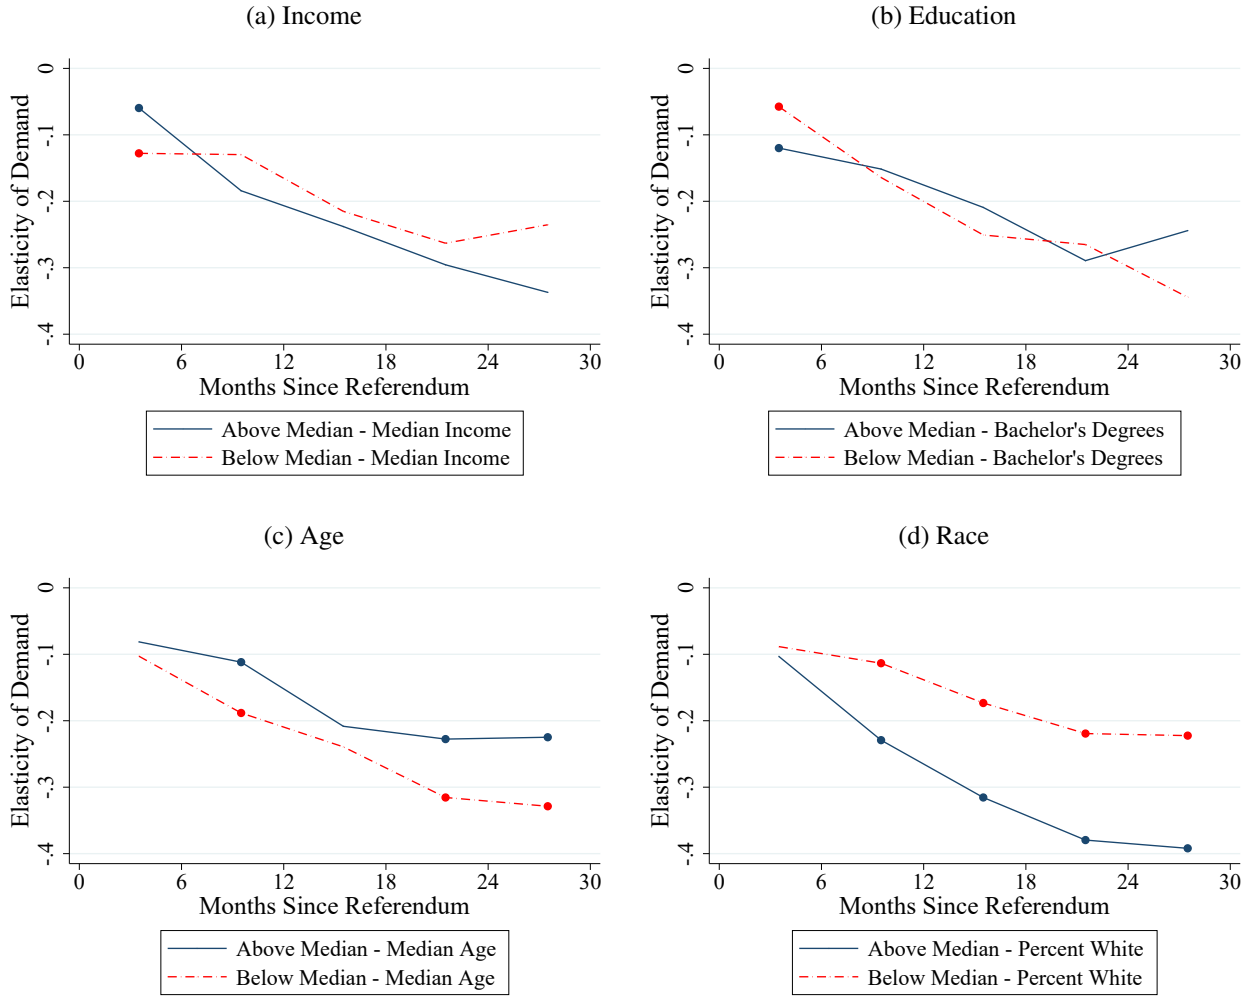
Figure 9 reports heterogeneity results for socioeconomic characteristics. Surprisingly, age and race appear to matter for the elasticity of demand more than economic variables such as income and education. Younger communities have a more elastic response, as do communities with a greater percentage of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

Figure 8: Elasticities by Community Characteristics: Housing and population



Notes: These panels display elasticity estimates for communities that are below and above median for the specified characteristic. The estimates are from a reduced-form specification augmented with an interaction term for whether the community is in the upper half of the distribution. The regressions include 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. Coefficients significant at the 10 percent level are indicated by a marker.

Figure 9: Elasticities by Community Characteristics: Socioeconomics



Notes: These panels display elasticity estimates for communities that are below and above median for the specified characteristic. The estimates are from a reduced-form specification augmented with an interaction term for whether or not the community is in the upper half of the distribution. The regressions include 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. Coefficients significant at the 10 percent level are indicated by a marker.

A.4 Difference-in-Differences Estimates without Matching

In this section, we describe how we estimate the effect of implementing aggregation on electricity prices and usage using a standard difference-in-differences model with no matching. Our estimating equation is given by

$$Y_{cmy} = \sum_{\tau=-24, \tau \neq -1}^{24} \beta_{\tau} A_{c\tau} + \beta_{25} A_{c,25} + \beta_{-25} A_{c,-25} + \alpha_{cm} + \alpha_{my} + \varepsilon_{cmy}, \quad (\text{A.2})$$

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

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

$$Y_{cmy} = \gamma_1 A_{c,0 \text{ to } 6} + \gamma_2 A_{c,7 \text{ to } 12} + \gamma_3 A_{c,13 \text{ to } 18} + \gamma_4 A_{c,19 \text{ to } 24} + \beta_{25} A_{c,25} + \beta_{-25} A_{c,-25} + \alpha_{cm} + \alpha_{my} + \varepsilon_{cmy}. \quad (\text{A.3})$$

In this specification, $A_{c,0 \text{ to } 6}$ is an indicator variable equal to 1 if the community passed aggregation in the past 6 months and 0 otherwise. Similarly, $A_{c,7 \text{ to } 12}$ is an indicator equal to 1 if the community passed aggregation between 7 and 12 months ago, and so on. The other variables are defined as in equation (A.2).

One could use this framework to estimate the effect of implementing aggregation by comparing communities that implemented aggregation to those that did not implement aggregation. However, the latter may not serve as adequate counterfactual for the former because, as we show in the paper, communities that implemented aggregation were on average different from those that did not. We therefore restrict our estimation sample to communities that implemented aggregation.³⁷ Our main identifying assumption here is that, conditional on a host of fixed effects, the timing of aggregation adoption is exogenous with respect to electricity use.

Figure A.6 presents the change in electricity prices following aggregation, in logs, as estimated by equation (A.2). Similar to our matching results, prices do not drop immediately following the referendum because it takes time for communities to switch to a new supplier. Unlike the matching estimator, the pre-period change is not exactly equal to zero in the event-study difference-in-difference. Although treatment and control communities face identical prices in the pre-period in *calendar time*, they do not face identical prices in *event-study time* because ComEd's prices fluctuate month-to-month. This distinction does not matter for the matching estimator, which creates counterfactuals separately for each treated community. The second vertical dashed line in Figure A.6 shows the point at which half of all communities have implemented aggregation (4

³⁷Difference-in-difference estimates that use all non-aggregation communities as controls are available upon request. These estimates are similar to those we present here but the coefficients are much more unstable.

months after passing the referendum). Prices continue to drop as more communities switch and then eventually stabilize. Within 8 months of passing the referendum, the average electricity price has decreased by more than 0.3 log points (26 percent) in aggregation communities relative to the control group. There is an increase in the relative aggregation price 28 months after passing aggregation, which is due to the fact that electricity prices fell sharply for ComEd customers in June of 2013, the middle of our sample period, and that some aggregation communities may have returned to ComEd. Despite this increase, prices in aggregation communities remain significantly lower than those in the control group for the entire sample period.

Figure A.7 shows the corresponding estimates for electricity usage. Prior the referendum, the difference in usage between aggregation and the control communities is statistically indistinguishable from zero. Usage in aggregation communities then begins to increase following the referendum. By the end of the first year, usage in aggregation communities is about 0.1 log points (9.5 percent) higher relative to the counterfactual.

Table A.5 shows the estimated impact of aggregation on the log of the electricity price in these communities 0-6, 7-12, 13-18, and 19-24 months after implementation, as estimated by equation (A.3). 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.6 shows the estimated change in usage as estimated by equation (A.3) for the sample of communities that implemented aggregation. Our preferred specification, presented in Column 4, estimates that electricity usage is 0.048 log points higher in the first 6 months following the referendum, and this increases to 0.114 log points within one year.

Finally, Figure A.8 shows the elasticities implied by the usage and price results. Similarly to the main text, we regress the predicted aggregation-driven log change in a community's electricity usage on the change in log prices. To account for uncertainty in the price and usage estimates, we cluster bootstrap the confidence intervals using 1000 draws. The implied elasticity ranges from -0.33 7-12 months after passage of aggregation to -0.45 two and a half years after passage. The differences in the magnitude between these and our matching estimates may be due to the fact that they reflect a slightly different mix of communities and a different combination of short-run and longer-run responses.

A.5 Municipal Aggregation Materials

After the proposed aggregation program has been registered with the state, the municipality must hold a referendum. The wording of the referendum question is specified in the Illinois Power Agency Act:³⁸

The election authority must submit the question in substantially the following form:

Shall the (municipality, township, or county in which the question is being voted upon) have the authority to arrange for the supply of electricity for its residential and small commercial retail customers who have not opted out of such program?

The election authority must record the votes as “Yes” or “No”.

Figure A.9 displays an example of a letter sent to residents of a community following the passage of an aggregation referendum and selection of a new aggregation supplier. The letter informs residents about their new supply price for electricity, and lets them know that they will have an opportunity to opt out of aggregation. Figure A.10 displays an example of the opt-out card that a customer must fill out and mail if they wish to retain their current electricity supplier.

Figures A.11 and A.12 display the front and back page of a typical electricity bill for a customer residing in ComEd’s service territory. If a customer switches suppliers, e.g., her community adopts aggregation and she does not opt out, then the Electricity Supply Charge rate (see Figure A.12) will change. Otherwise her bill will remain the same.

³⁸From 20 ILCS 3855/1-92, Text of Section from P.A. 98-404. Available from <http://www.ilga.gov/legislation/ilcs/fulltext.asp?DocName=002038550K1-92>.

Appendix Tables

Table A.1: Characteristics of Aggregation, Non-Aggregation, and Matched Control Communities

	(1) Aggregation Mean	(2) Non- Aggregation Mean	(3) <i>p</i> -value of Difference (2)-(1)	(4) Matched Controls Mean	(5) <i>p</i> -value of Difference (4)-(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.

Table A.2: Matching Estimates of the Effect of Aggregation on Usage and Prices, Monthly

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

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each aggregation community is matched to the five non-aggregation communities 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.3: Matching Estimates of the Effect of Aggregation on Usage and Prices, Yearly

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

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each aggregation community is matched to the five non-aggregation communities 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.4: Comparison of Tuning Parameters for Subsampling

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

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

Table A.5: Effect of Aggregation on Electricity Prices, Communities that Passed Aggregation

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

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

Table A.6: Effect of Aggregation on Electricity Usage, Communities that Passed Aggregation

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

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

Table A.7: Number of Nearest Neighbors and Reduced-Form Estimates

(a) 5 Nearest Neighbors

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

(b) 1 Nearest Neighbor

	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
1-12 Months Post-Referendum	0.034*** (0.005)	-0.173*** (0.004)	-0.144*** (0.018)	3360	3360
13-24 Months Post-Referendum	0.046*** (0.006)	-0.141*** (0.002)	-0.265*** (0.032)	2671	2670
25-36 Months Post-Referendum	0.049*** (0.009)	-0.108*** (0.006)	-0.318*** (0.049)	720	718

(c) 10 Nearest Neighbors

	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
1-12 Months Post-Referendum	0.033*** (0.005)	-0.173*** (0.004)	-0.146*** (0.018)	3360	3360
13-24 Months Post-Referendum	0.042*** (0.005)	-0.141*** (0.002)	-0.246*** (0.029)	2671	2670
25-36 Months Post-Referendum	0.047*** (0.009)	-0.108*** (0.006)	-0.295*** (0.043)	720	718

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Estimates are constructed by a nearest-neighbor matching approach where each aggregation community is matched to non-aggregation communities 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.8: Number of Nearest Neighbors and the Dynamic Elasticity Curve

(a) 5 Nearest Neighbors

Period After Price Change	Perfect Foresight	Status Quo Expectations
Contemporaneous	-0.052*** [-0.075, -0.033]	-0.052*** [-0.076, -0.034]
Mean: Months 1-12	-0.112*** [-0.144, -0.083]	-0.113*** [-0.142, -0.086]
Mean: Months 13-24	-0.195*** [-0.249, -0.147]	-0.192*** [-0.242, -0.145]
Mean: Months 25-36	-0.249*** [-0.328, -0.184]	-0.239*** [-0.316, -0.177]
Year 10	-0.345*** [-0.606, -0.243]	-0.306*** [-0.601, -0.221]
Long Run	-0.348*** [-1.007, -0.245]	-0.307*** [-1.079, -0.221]

(b) 1 Nearest Neighbor

Period After Price Change	Perfect Foresight	Status Quo Expectations
Contemporaneous	-0.056*** [-0.081, -0.037]	-0.056*** [-0.081, -0.038]
Mean: Months 1-12	-0.108*** [-0.142, -0.087]	-0.108*** [-0.145, -0.085]
Mean: Months 13-24	-0.190*** [-0.249, -0.145]	-0.188*** [-0.245, -0.143]
Mean: Months 25-36	-0.256*** [-0.331, -0.186]	-0.247*** [-0.321, -0.180]
Year 10	-0.464*** [-0.691, -0.307]	-0.402*** [-0.642, -0.272]
Long Run	-0.509*** [-1.425, -0.333]	-0.422*** [-1.018, -0.284]

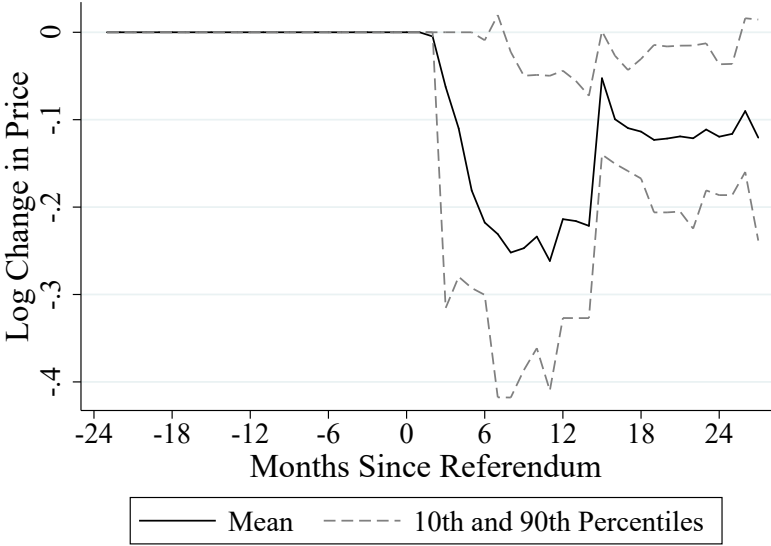
(c) 10 Nearest Neighbors

Period After Price Change	Perfect Foresight	Status Quo Expectations
Contemporaneous	-0.058*** [-0.078, -0.035]	-0.057*** [-0.078, -0.037]
Mean: Months 1-12	-0.118*** [-0.147, -0.088]	-0.119*** [-0.148, -0.089]
Mean: Months 13-24	-0.202*** [-0.254, -0.149]	-0.200*** [-0.250, -0.147]
Mean: Months 25-36	-0.257*** [-0.336, -0.183]	-0.248*** [-0.326, -0.178]
Year 10	-0.355*** [-0.663, -0.245]	-0.321*** [-0.638, -0.222]
Long Run	-0.359*** [-1.124, -0.253]	-0.322*** [-1.041, -0.223]

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. The elasticities are interpreted as the cumulative effect of a permanent one-percent price change on current usage. Estimates are constructed from 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 aggregation community is matched to the five non-aggregation communities with the most similar usage in 2008 and 2009. 95 percent confidence intervals are displayed in brackets and are constructed via subsampling.

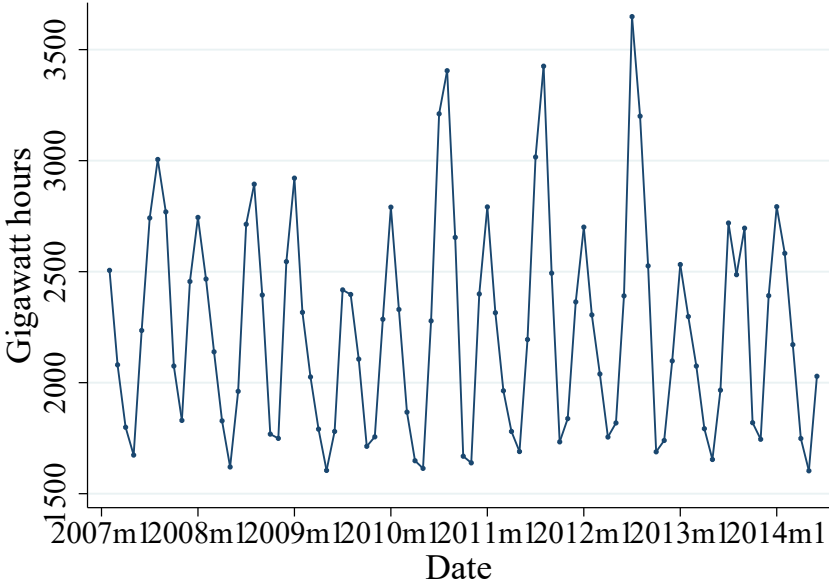
Appendix Figures

Figure A.1: Heterogeneity in Price Changes



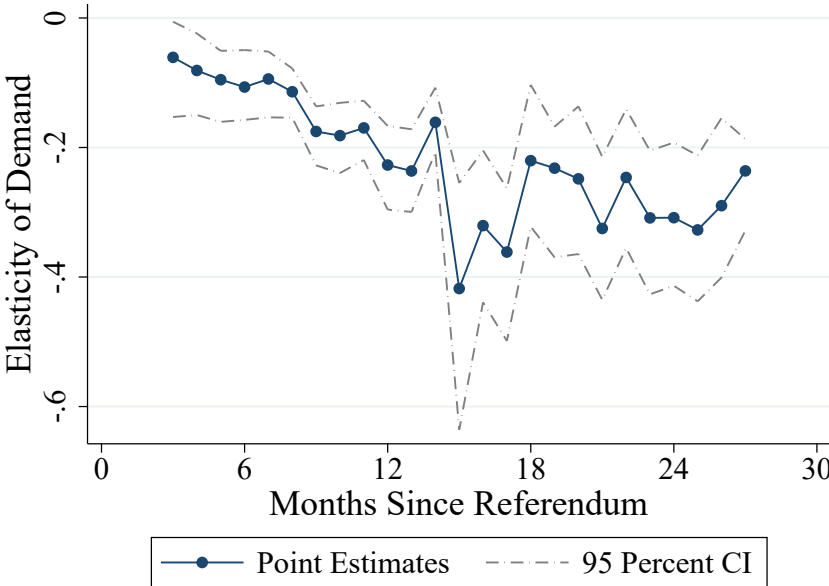
Notes: This figure displays the distribution of (log) price changes for aggregation communities relative to the contemporaneous ComEd price. The solid line displays the mean, and the dashed lines represent the 10th and 90th percentiles.

Figure A.2: Monthly Electricity Usage in ComEd Service Territories, 2007-2014



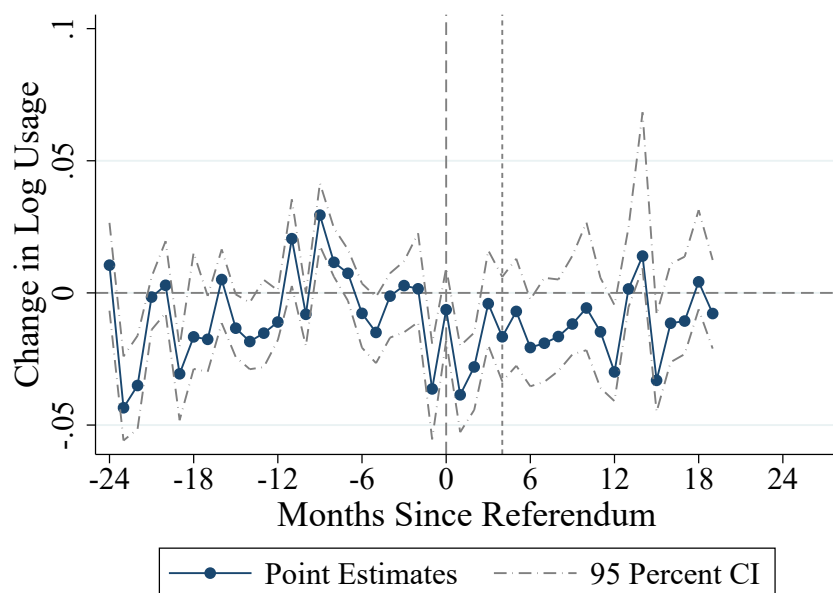
Notes: Figure displays total electricity usage across the ComEd service territories in our sample.

Figure A.3: Estimated Price Elasticities, Monthly



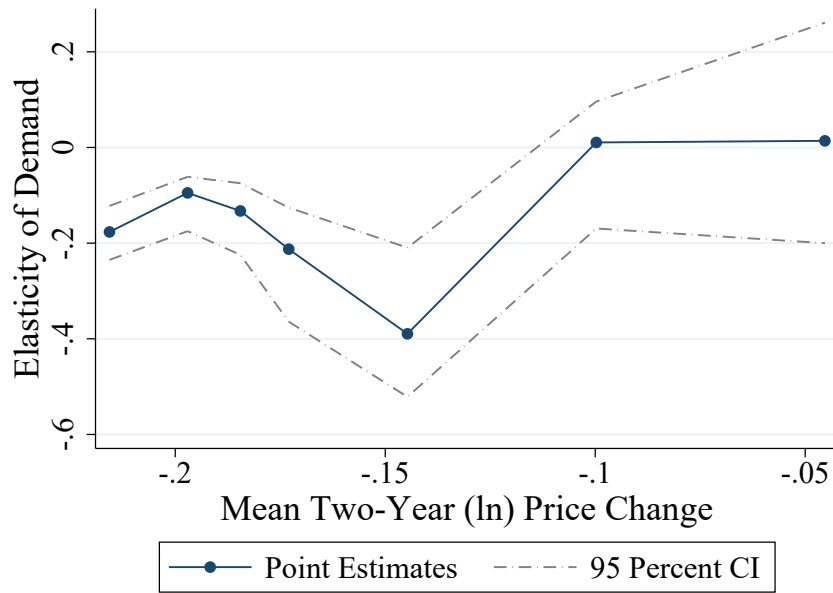
Notes: Elasticities are calculated for each month by regressing community-month changes in log usage on the observed change in log price. Confidence intervals are constructed via subsampling.

Figure A.4: Effect on Log Usage: Communities that Passed but Did Not Implement Aggregation



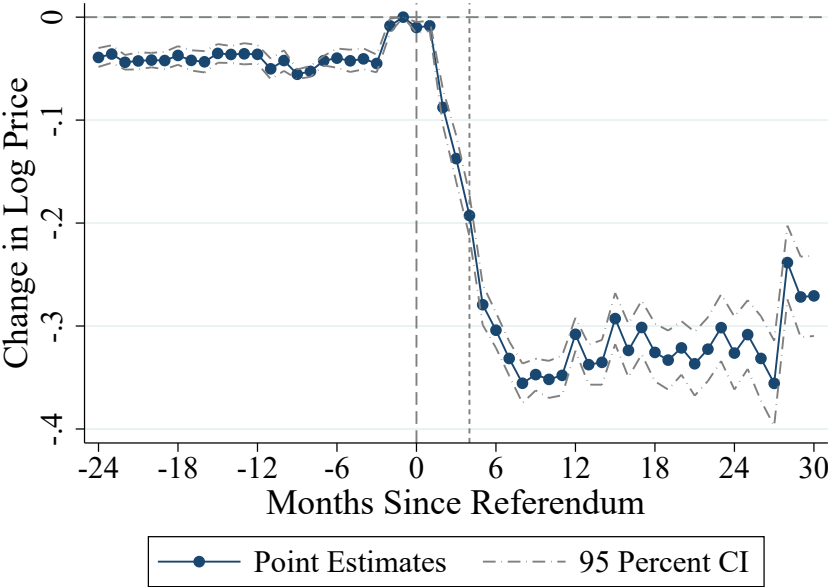
Notes: The figure displays estimates of the mean usage effect for the eleven communities that pass aggregation but never implement it. The effect is estimated relative to that community's five nearest neighbors, as defined by the difference-in-differences matching procedure outlined in the main text.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

Figure A.5: Estimated Elasticities and Mean Log Price Change



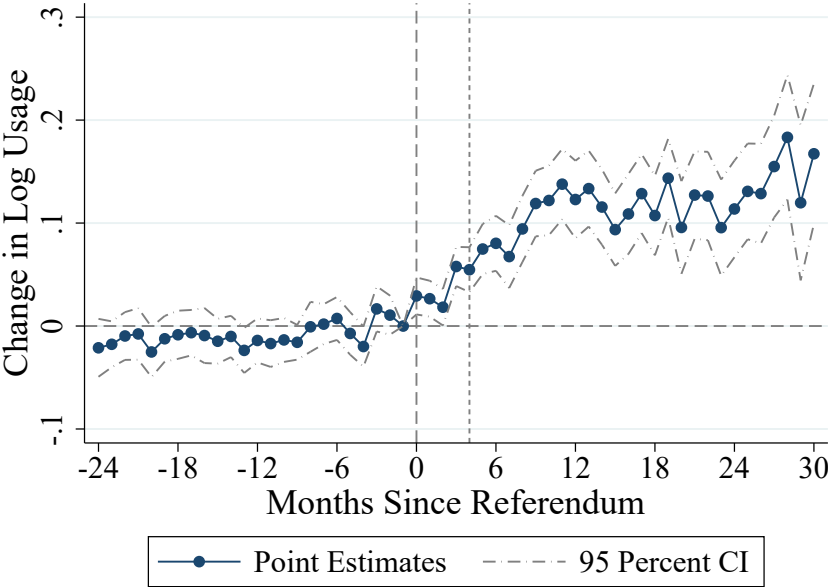
Notes: Communities are split into seven groups based on the average two-year price change they experienced in the first two years following their referenda. Elasticities are calculated separately for each group. The graph shows no relationship between the estimated group elasticity and the price change, mitigating concerns that the price change might be correlated with a community's elasticity of demand. Confidence intervals are constructed via subsampling.

Figure A.6: Regression Estimates of the Effect of Aggregation on Electricity Prices, Communities that Passed Aggregation



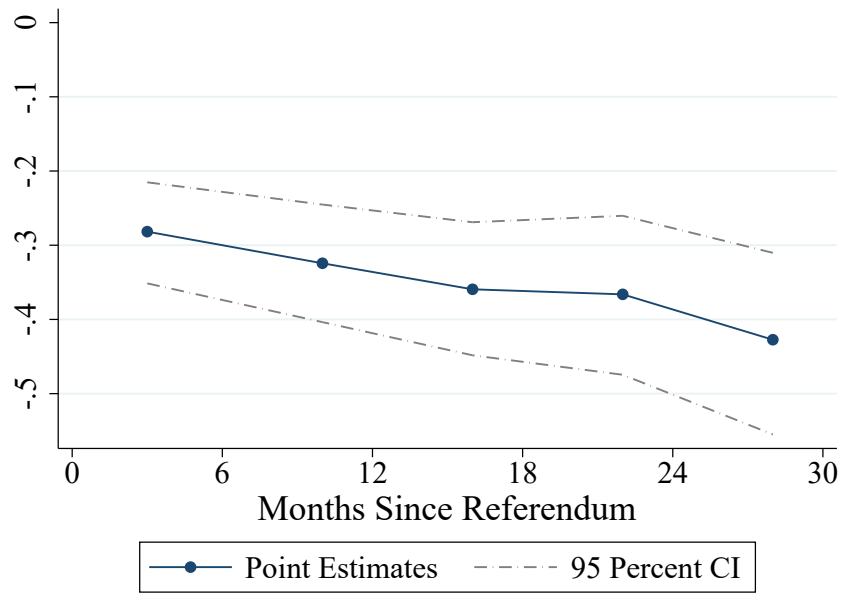
Notes: Outcome is the natural log of the electricity price. The first vertical dashed line indicates the date of the aggregation referendum. The second dashed line indicates the date of aggregation implementation. Regressions include month-by-year and community-by-month fixed effects. Standard errors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

Figure A.7: Regression Estimates of the Effect of Aggregation on Electricity Usage, Communities that Passed Aggregation



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

Figure A.8: Estimated Price Elasticities, Communities that Passed Aggregation



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

Figure A.9: Example of an Aggregation Mailing



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

John A. Smith
123 Main St
Anytown, IL 65432

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

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

HOW TO OPT-OUT

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

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

ENROLLMENT PROCESS

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

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

Sincerely,

See Reverse for Frequently Asked Questions...

Christopher J. Lauzen
Board Chairman
Kane County

Kurt R. Kojzarek
Development Committee Chairman
Kane County

Figure A.10: 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: _____

Figure A.11: Example of a ComEd Bill (page 1 of 2)

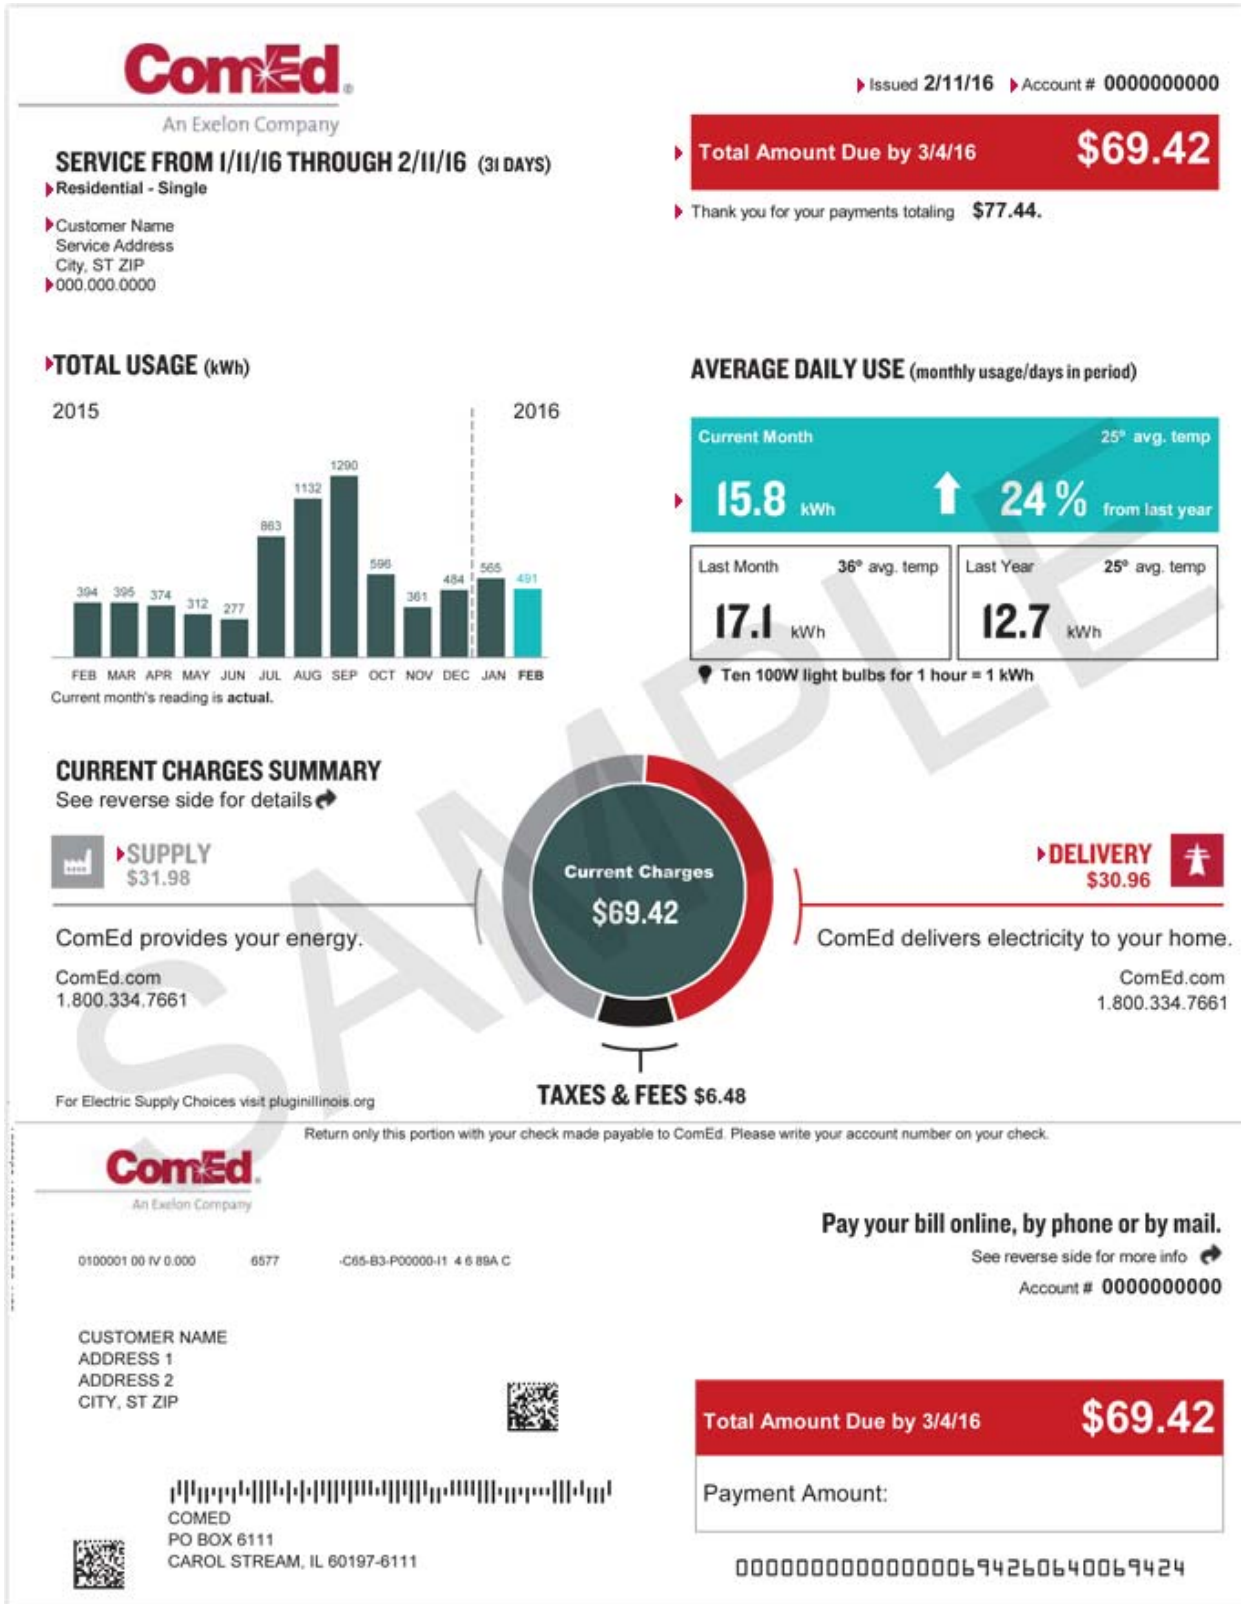


Figure A.12: Example of a ComEd Bill (page 2 of 2)

For Questions, Support, and Outages visit ComEd.com

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

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

▶ Total Amount Due by 3/4/16
\$69.42

METER INFORMATION

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

CHARGE DETAILS

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

SUPPLY **\$31.98**

▶ Electricity Supply Charge ▶ 491 kWh X 0.05865 \$28.80

Transmission Services Charge 491 kWh X 0.01122 \$5.51

Purchased Electricity Adjustment -\$2.33

DELIVERY - ComEd **\$30.96**

▶ Customer Charge \$10.53

▶ Standard Metering Charge \$4.36

▶ Distribution Facilities Charge 491 kWh X 0.03156 \$15.50

▶ IL Electricity Distribution Charge 491 kWh X 0.00118 \$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.

OTHER WAYS TO PAY YOUR BILL

Visit [ComEd.com/PAY](#) for more information including applicable fees for some transactions.

Online

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

Mobile App

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

Phone

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

In-Person

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

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

10% total recycled fiber

6577-65-0100001-0001-0000001