

The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation*

Tatyana Deryugina
University of Illinois
and NBER

Alexander MacKay
Harvard University

Julian Reif
University of Illinois
and NBER

[Final Version]

Abstract

We study the dynamics of residential electricity demand by exploiting a natural experiment that produced large and long-lasting price changes in over 250 Illinois communities. Using a flexible difference-in-differences matching approach, we estimate that the price elasticity of demand grows from -0.09 in the first six months to -0.27 two years later. We find similar results with a dynamic model in which usage is a function of past and future prices. Our findings highlight the importance of accounting for consumption dynamics when evaluating energy policy.

JEL codes: Q41, Q48, D12

*We thank ComEd for generously sharing electricity usage data with us, and we are particularly grateful to Renardo Wilson for many helpful discussions. We also thank Torsten Clausen, the Illinois Commerce Commission, and Dave Hoover 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. Mohammad Ahmadizadeh, Noah Baird, Dylan Hoyer, Chitra Jogani, and Prakrati Thakur provided excellent research assistance. Views reflected here are solely those of the authors. Emails: deryugin@illinois.edu, amackay@hbs.edu, and jreif@illinois.edu.

Economic theory suggests that demand is typically more elastic in the long run relative to the short run. When consumption depends on goods that are durable or habit-forming, consumers may take years to respond fully to a price change (Topel and Rosen, 1988; Becker et al., 1994). These dynamics pose several 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, as they typically do, the demand response will reflect a mix of both short-term and long-term changes in consumption. Finally, unless one accepts restrictive functional form assumptions, unbiased estimation requires a source of exogenous price variation large enough to produce detectable effects in long-run behavior. Quantifying these demand-side dynamics is important in the electricity sector, where suppliers, market regulators, and policymakers make decisions with long-run ramifications.

We estimate how the price elasticity of electricity demand evolves over time by exploiting an Illinois policy that generated plausibly exogenous shocks to residential electricity prices in over 250 communities. Because these price shocks were large and lasted over two years, we are able to estimate the demand response more flexibly and over a longer period than prior quasi-experimental studies. We show that residential electricity consumers take multiple years to adjust to price changes. Our event study finds an elasticity of -0.27 in the period 25–30 months after the policy change, almost three times the magnitude of the elasticity in the first six months (-0.09).

Our monthly consumption and price data span 2007–2014 and come from the largest utility in Illinois, ComEd, whose territory encompassed approximately 70 percent of residential consumers in the state. During this period, Illinois implemented a municipal aggregation program.¹ The program allowed individual communities to select new electricity suppliers on behalf of their residents with the approval of a local referendum. This policy change resulted in large, long-lasting price changes for communities that implemented aggregation. Our setting provides a clean natural experiment: aggregation customers continued to receive their electricity bill from the utility in the same format as before, so the price variation in our analysis is not confounded with other billing changes. Additionally, 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.

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

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 regression. 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 conduct inference using subsampling, which allows for a richer space of estimators whose distributions do not have pre-existing formulas.

Our estimator matches 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. Because we observe usage at a high (monthly) frequency and community referenda occur only during infrequent statewide elections, 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. In light of these factors and the absence of pre-trends, aggregation therefore provides plausibly exogenous price variation in our sample.

We find that prices fell by 24 percent and usage increased by 6.1 percent by the end of the first year following an aggregation referendum, relative to control communities that did not pass a referendum on aggregation. Toward the beginning of the second year, the

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

expiration of a long-term ComEd contract caused a price decrease and an accompanying usage increase in control communities. Nevertheless, during our estimation period, the estimated price elasticity declines smoothly from -0.09 in the first six months to -0.27 two years later, illustrating the importance of long-run dynamics in this setting.

Although they are long-lasting, the relative price decreases in our sample are not constant over time. Thus, the estimates above reflect a mix of both short-run and longer-run responses. To address this shortcoming, we also estimate a 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. Using this model, we estimate an elasticity of -0.08 in the first six months following a price change and an elasticity of -0.21 19–24 months after a price change. Accounting for the fact that the median community implements a price change four months after a referendum, we conclude that the patterns estimated by the more flexible dynamic model are similar to those from the reduced-form approach described above. Finally, we employ a parametric formulation of the dynamic model to forecast that the long-run elasticity converges to -0.35 after approximately 10 years.

Our results have significant implications for energy policy and market participants. 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. Utilizing a short-run elasticity will lead to an underestimate of the longer-run demand response. Using simple back-of-the-envelope calculations, we show that predicted quantity reductions following the imposition of a carbon tax are 2.3–2.9 times larger when using the two-year price elasticity of demand compared to using the six-month elasticity, depending on the elasticity of supply. This difference implies that the carbon tax required to reduce emissions by one percent is 56–65 percent smaller. Although our exact estimates may not directly translate to other settings, the relative magnitudes of the short- and longer-run elasticities suggest that other energy policies can cause a substantially larger usage response in the long run than in the short run.

The existing electricity demand literature relies on state-level data and dynamic panel models to estimate consumption dynamics over a longer period. The long-run elasticity estimates vary widely, from -0.3 to about -1.1 (e.g., Kamerschen and Porter, 2004; Dergiades and Tsoulfidis, 2008; Alberini and Filippini, 2011), and none of these estimates is

based on quasi-experimental variation.³ Consistency in these models generally requires strong assumptions about the form of serial correlation, and the estimates are particularly sensitive to the exact specification used (Alberini and Filippini, 2011). Conversely, our approach makes relatively few assumptions and is substantially more flexible than what has been done previously.

Papers outside of the dynamic panel literature typically estimate only short-run or static price elasticities.⁴ 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 -0.09 in the first 4 months following a price change, which is similar in magnitude to our six-month estimate. Reiss and White (2005) employ cross-sectional data from California and a structural model that exploits the non-linearity of the electricity price schedule. They estimate an average (static) price elasticity of -0.39 , but do not investigate how it evolves following a price change.⁵ Our findings demonstrate the importance of such dynamics with respect to the price of electricity: residential electricity consumers are more than twice as responsive in the longer run relative to the short run.

A few recent papers estimate consumption dynamics in response to non-price interventions or in other contexts. Allcott and Rogers (2014) estimate that the energy reductions following random assignment of a home energy report are larger 7–12 months after the beginning of the program, relative to 1–6 months after. On the other hand, Ito (2015) finds that the one-year and three-year effects of an appliance rebate program are similar, suggesting that not every policy induces dynamics of the kind we find in our study. In the context of borrowing, Karlan and Zinman (2018) find that the elasticity with respect to the interest rate nearly triples over time, from -1.1 in the first year to -2.9 in the third year.

The rest of this paper is organized as follows. Section I discusses the electricity market

³See Alberini and Filippini (2011) for a review of this literature. 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 that employ instruments to estimate a longer-run price elasticity.

⁴A growing literature investigates the impact of real-time pricing (e.g., Wolak, 2011; Allcott, 2011; Jessee 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 consumption.

⁵Reiss and White (2005) estimate a same-month elasticity, which they allow to vary seasonally, and then report the average annual value of these estimates.

and municipal aggregation in Illinois. Sections II and III describe our data and reduced-form empirical approach, respectively. Section IV presents reduced-form results, which Section V then extends to a dynamic framework. Section VI discusses the implications of our main results, and Section VII concludes.

I The Illinois Electricity Market

The provision of electricity to residential customers consists of two components: supply and distribution. Suppliers generate or procure electricity and distributors provide the infrastructure to deliver it 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, where approximately 70 percent of the state’s residents live. The supply price component of electricity delivered by Ameren or ComEd is, by law, equal to their procurement cost and does not vary geographically.⁶ 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.⁷ 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

⁶Profits stem from delivery fees set by the Illinois Commerce Commission (DeVirgilio, 2006).

⁷Large commercial and industrial customers gained this ability at the end of 1999.

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 required municipalities to allow individuals to opt out of aggregation.

To implement an opt-out aggregation program, a municipality had to publicize the proposal, hold a town hall meeting to educate the community, register the proposed aggregation program with the state, and hold a referendum. Referenda dates followed the state electoral calendar: they were held in March or November in even years and in April in odd years. The wording of the referendum question was specified in the Illinois Power Agency Act and is reproduced in Section 1 of the Appendix. In most cases where the referendum was approved, multiple suppliers submitted bids for predetermined contract lengths (e.g., one-, two-, and three-year contracts). In other cases, the municipality negotiated directly with a supplier. The two main ways in which suppliers differentiated themselves were price and the share of generation from renewable sources. Nearly all communities selected the supplier with the lowest price, although environmental preferences occasionally induced communities to select a more expensive one.

When determining the bid or negotiating directly, each supplier obtained community-level usage data from the distributor. These usage data, along with electricity futures, were the main factors in each offered price.⁸ Importantly, our analysis employs 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 was popular in Illinois. Of the roughly 2,100 communities in the ComEd and Ameren service territories, 741 had voted to implement aggregation as of March 2016. In our setting, the realized savings from aggregation came 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 passed an aggregation referendum were automatically switched to the newly chosen electricity supplier unless they opted out by mailing in a

⁸Suppliers may have also based 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).

card, calling, or filling out a form online.⁹ Aggregation officially began at the conclusion of the opt-out process. From the consumer’s point of view, the only visible change was the supply price of electricity on her bill, which was still issued by the incumbent distributor (Ameren or ComEd). The design of the bill, as well as the value of the non-supply prices, remained identical for all customers in the distributor’s territory regardless of aggregation. Conveniently, this means that the price effects of aggregation were not confounded with changes in a bill’s appearance. Figures A.1–A.4 in the Appendix display a sample letter notifying households of aggregation, a sample opt-out card, and a sample ComEd bill.

Our estimation approach assumes that municipal aggregation was not accompanied by other programs that could also affect electricity usage. A careful review of news articles, aggregation-related announcements, and other online materials did not turn up any evidence contradicting this assumption. ComEd runs several energy-related rebate and discount programs, but these were offered uniformly to all communities in its service territory during our sample period (with the exception of small pilot programs, which are typically conducted in communities near ComEd’s headquarters in Chicago). Moreover, ComEd does not have a strong profit incentive to target its rebate and discount programs to aggregation or non-aggregation communities. By law, ComEd makes zero profits from electricity supply. Instead, its profits come from distributing electricity, and ComEd remained the distributor in all the aggregation communities in our sample.

II Data

We obtained electricity usage data directly from ComEd. The data contain monthly residential electricity usage at the municipality level for ComEd’s 887 service territories from February 2007 through June 2014. We drop 108 communities that were missing data, did not appear to have consistent geographic boundaries during our sample period, or could not be assigned an aggregation status with confidence. For our main analysis, we also drop an additional 11 communities that passed a referendum approving aggregation but never implemented the program. The resulting dataset is a balanced panel of monthly electricity

⁹While we do not have an exact number, ComEd and several energy suppliers have told us that the opt-out rate was 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 did, however, grow slowly over time because new residents who moved to an aggregation community were 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 were not switched over to the chosen supplier.

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

usage for 768 ComEd communities.¹⁰

We constructed the time series of ComEd electricity rates using ComEd ratebooks 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 was equal to the non-electric space heating rate, which was 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. We describe these sources in further detail in Section 2 of the Appendix.

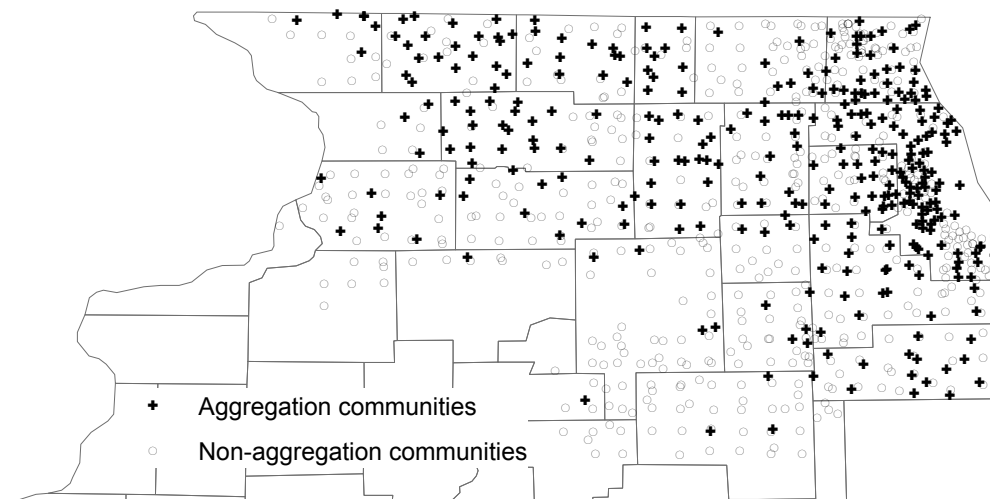
As shown in Table 1, 300 communities in the ComEd territory passed a referendum on aggregation during our sample period, and 289 of those communities eventually implemented an aggregation program. In addition, 36 communities voted on, but did not pass, the referendum.¹¹ The geographic locations of the 768 aggregation and non-aggregation 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 whether the local government should involve itself in negotiating electricity prices,

¹⁰See Section 2 of the Appendix for additional data details. Ameren, the other distributor in Illinois, declined our data request.

¹¹We investigated the possibility of employing a regression discontinuity approach based on the fraction voting “yes” or “no”, but we do not have enough power. There are only 20 communities where the referendum failed by less than five percentage points.

Figure 1: Spatial Distribution of Communities in Sample



Notes: Figure displays the locations of communities in our sample. Plus signs indicate communities that implemented aggregation. Circles indicate communities that did not pass aggregation.

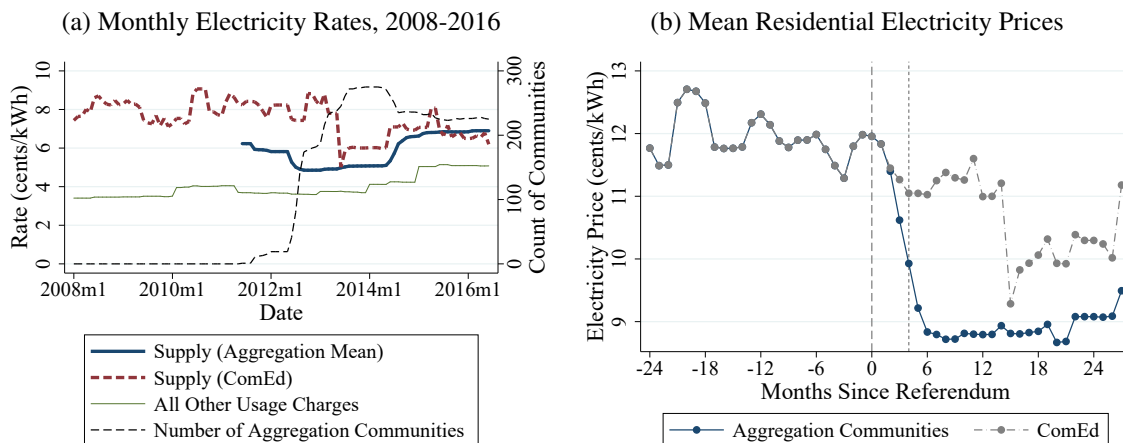
played a role in a community’s decision to hold a referendum. To the best of our knowledge, changes in expected future usage were not a consideration.¹² Anecdotally, the reasons why some communities voted against aggregation include: (1) lack of trust in the local government to secure savings relative 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. For our reduced-form results, we construct estimates relative to the referendum date to capture any usage response that occurred 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

¹²Our understanding of the motivations behind aggregation is guided by local news articles, local government meeting minutes, in-person discussions with ComEd and the Illinois Commerce Commission (Illinois’ electricity regulator), and phone conversations with industry participants.

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 thin 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 vertical line indicates the median implementation date relative to when the referendum was passed.

the “price” to which consumers might be responding. This constant marginal price can be broken down into three main components: supply, delivery, and taxes/fees. Implementing aggregation entails a community signing a contract with a supplier for a particular supply rate (the largest component of the marginal price); delivery charges and other fees remain the same. 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 9 months to 3 years.

Figure 2 demonstrates the price variation in our sample. The thick dashed line and thin green line in panel (a) 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 a long-term high-price power contract expired. We discuss the implications of this shock for the interpretation of our estimates in Section IV. 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, which displays the count of aggregation communities over time.

The thick blue line in panel (a) shows the average monthly supply rate for aggrega-

tion 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 variation in the ComEd rate is small relative to the level shocks brought about 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 (panel (b) of Figure 2). The displayed price combines the supply rate with the other usage charges from panel (a). 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 a successful referendum, reflecting the June 2013 drop in the ComEd supply price. There was also significant cross-sectional variation in the aggregation price shocks (see Figure A.5 in the Appendix), 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. This last source of variation raises the concern that suppliers might have charged different prices to communities with different expected future usage. Our estimation approach mitigates this issue because it matches aggregation communities to their non-aggregation controls based on the same community-level usage data that suppliers used to determine prices. Furthermore, we investigate the potential for price endogeneity by estimating elasticities separately by percentiles of the realized price change, and we find no relationship (see Section IV.B).¹³

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

¹³An additional concern is that the implementation of aggregation could influence the supply price paid by remaining ComEd customers through a reverse causality channel: if communities with favorable load profiles implemented aggregation, then the remaining ComEd communities would be more costly to serve, causing ComEd to increase its price. However, ComEd still supplied a large pool of customers throughout our sample period, and the fact that aggregation prices converged to the ComEd price in 2015–2016 suggests that selection based on the load profile is not a significant concern in our setting.

mailer informing them of the new aggregation price. Likewise, the drop in ComEd prices received statewide attention in the press.¹⁴ Though we cannot say whether other potential energy policies that affect electricity prices (e.g., a carbon tax) will be more or less salient, they are also likely to be anticipated and publicized. Aggregation therefore provides a good setting for assessing the residential usage response to energy policies.

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

III Empirical Strategy

A *Difference-in-Differences Matching Estimator*

To estimate the price elasticity of demand, we first estimate changes in electricity usage brought about by aggregation. We match communities that implemented aggregation (the “treated” group) to communities that did not (the “control” group) based on their pre-aggregation electricity usage. We then apply a difference-in-differences adjustment to the bias-corrected matching estimator developed by Abadie and Imbens (2006, 2011). For each of the 289 communities that implemented aggregation, we use 2008–2009 electricity usage data to identify the five nearest neighbors from the 479 communities in our sample that did not implement aggregation. 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 distances between communities. That distance is then used to select the five nearest neighbors for each treated community.¹⁵

We use these nearest neighbors to construct counterfactual usage and 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 subse-

¹⁴See, for example, “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.

¹⁵We allow non-aggregation communities to be selected as neighbors for multiple aggregation communities; i.e., each aggregation community draws from an identical set of potential controls. Table A.1 in the Appendix shows that the results are very similar if we instead select the single nearest neighbor or the ten nearest neighbors.

quent 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. The outcome Y_{it} is a function of D_i , so that, for aggregation communities, $Y_{it}(1)$ indicates usage when treated and $\widehat{Y}_{it}(0)$ indicates estimated counterfactual usage when not treated. Given $Y_{it}(1)$ and $\widehat{Y}_{it}(0)$, we can obtain a community-specific estimate of the effect of aggregation on usage, $\widehat{\Delta Y}_{it}$:

$$\widehat{\Delta Y}_{it} = Y_{it}(1) - \widehat{Y}_{it}(0). \quad (1)$$

We observe the outcome $Y_{it}(1)$ for the treated communities in our data. The counterfactual outcome, $\widehat{Y}_{it}(0)$, 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, $\widehat{Y}_{it}(0)$, is then equal to

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

where

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

The parameter $\widehat{\mu}_i^{m(t)}$ 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.¹⁶ Thus, the estimated counterfactual $\widehat{Y}_{it}(0)$ is equal to the average usage for a treated community's nearest neighbors plus the difference in usage between that community and

¹⁶For example, if $t = 25$ is January 2014, then $\widehat{\mu}_i^{m(25)} = \frac{1}{2} (Y_{i,m(25)}^{2008} + Y_{i,m(25)}^{2009}) = \frac{1}{2} (Y_{i,January}^{2008} + Y_{i,January}^{2009})$ is equal to the average log usage in January 2008 and January 2009.

its neighbors averaged across the 2008–2009 calendar months corresponding to t .

Finally, we obtain a community-specific “difference-in-differences” estimate of the impact on usage by netting out the average change in usage in the year prior to treatment. This estimate is defined as

$$\widehat{\tau}_{it} = \widehat{\Delta Y}_{it} - \frac{1}{N_s} \sum_{s=1}^{N_s} \widehat{\Delta Y}_{i,-s}, \quad (2)$$

where N_s indicates the number of periods in the year prior to the policy change.¹⁷ The $\widehat{\Delta Y}$ terms on the right-hand side of equation (2) are defined in equation (1). Our difference-in-differences estimate $\widehat{\tau}_{it}$ thus reflects the difference in usage between a treated community and its matched control communities in period t , relative to the average difference in the year leading up to the policy change.

To quantify the average impact of the policy on usage, we take the mean of the community-specific treatment effects:

$$\widehat{\tau}_t = \frac{1}{N_1} \sum_{i=1}^{N_1} \widehat{\tau}_{it}. \quad (3)$$

where N_1 denotes the number of treated communities in our sample. Because the policy change occurs at the community level, and we only include communities that implement aggregation, we interpret this estimate as the effect of the treatment on the treated.¹⁸

In our setting, the aggregation electricity price is announced months before the actual price change occurs. 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. Prior studies have found evidence of forward-looking behavior by energy consumers (Allcott and Wozny, 2014; Myers, 2016). Failing to account for such behavior could lead to biased estimates (Malani and Reif, 2015).

Since 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. Section 4 of the Appendix presents evidence of anticipatory behavior by estimating changes in electricity use after the referendum is passed but before a community switches to a new supplier.¹⁹ Using only pre-implementation data, we find

¹⁷For monthly estimates, $N_s = 12$. For biannual estimates, $N_s = 2$.

¹⁸Because a small proportion of households opted out of aggregation, at the household level our estimate reflects an intent-to-treat effect (conditional on implementing aggregation).

¹⁹It 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. We find no evidence of such behavior.

small but noticeable usage increases of 0.012 to 0.035 log points three to five months after passage of the referendum but before the price changes had occurred.

We estimate price changes the same way we estimate usage changes: by comparing prices in each treated community to its matched controls. Because electricity rates do not vary in the cross-section for non-aggregation communities, the price difference is exactly zero prior to a community’s implementation of aggregation. Therefore, the difference-in-differences estimate of log price changes is simply the observed difference between the community’s aggregation rate and the ComEd rate, $\Delta \ln p_{it}$. The mean effect of aggregation on prices is the average of these differences:

$$\overline{\Delta \ln p_t} = \frac{1}{N_1} \sum_{i=1}^{N_1} \Delta \ln p_{it}. \quad (4)$$

B Estimating Elasticities

To obtain elasticities, we regress community-specific estimates of the change in usage on community-specific price changes. We estimate separate elasticities over time in order to show how the elasticity changes dynamically. For each post-referendum period, g , the corresponding elasticity, β_g , is obtained via the regression

$$\widehat{\tau}_{it} = \beta_g \cdot \Delta \ln p_{it} + \eta_{it} \quad \forall t \in g. \quad (5)$$

We estimate two versions of this specification. First, we estimate equation (5) separately for each month post-referendum (i.e., $\{g\} = \{t\}$). Second, to increase the precision of our estimates, we group the data into six-month intervals and estimate equation (5) for each of these groups. In this second specification, β_g is the average elasticity for each six-month period. We also report estimates of the mean change in log price and in log usage in each six-month period, $\frac{1}{|g|} \sum_{t \in g} \overline{\Delta \ln p_t}$ and $\frac{1}{|g|} \sum_{t \in g} \widehat{\tau}_t$.

The elasticities estimated by equation (5) are accurate to the extent that the post-period price changes in our sample can be reasonably approximated by a one-time, permanent change in price. We account for the empirical variation in prices over time in Section V, where we estimate a more flexible model that disentangles the short-run response to price changes from the longer-run response.

C 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.²⁰ Subsampling, like bootstrapping, obtains a distribution of parameter estimates by sampling from the observed data.

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

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

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

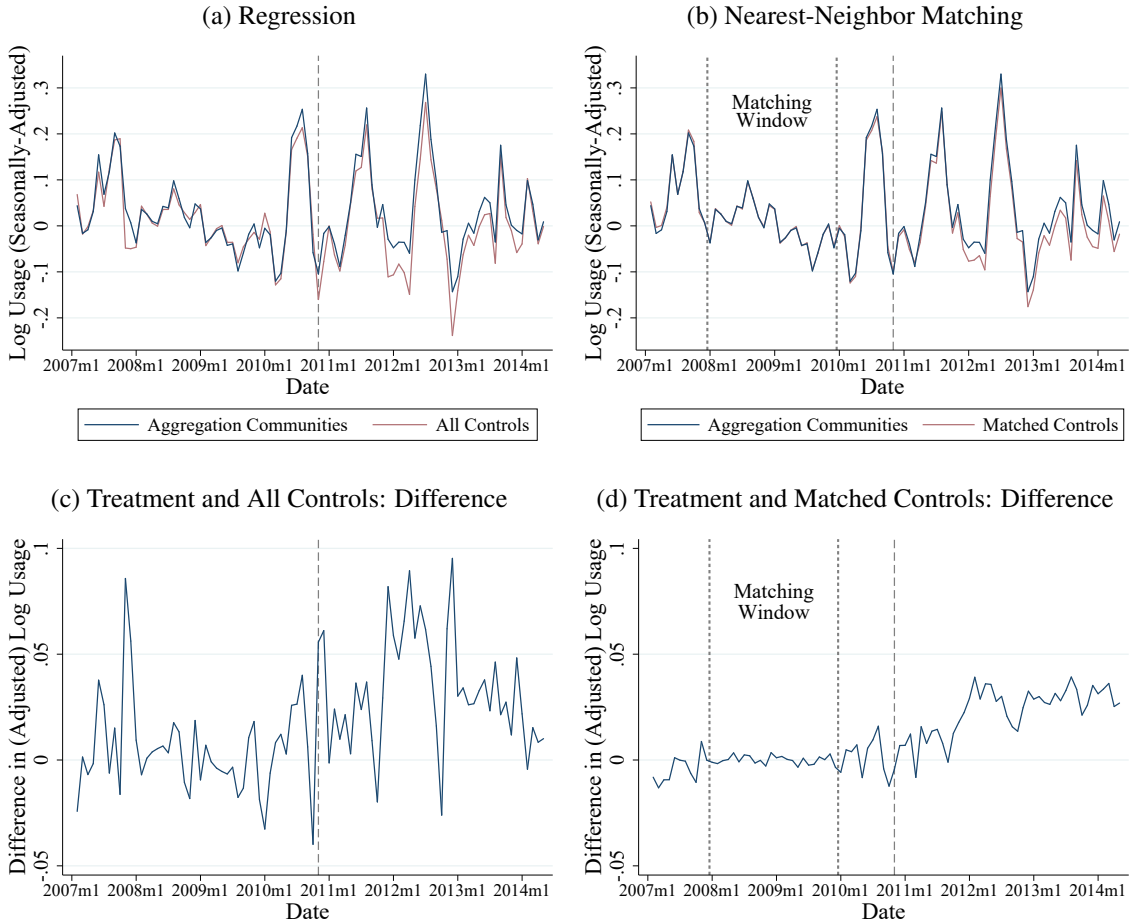
Subsampling requires a large number of effective observations (i.e., treated units) in each subsample, but it also requires that this number be small relative to the total number of effective observations in the full sample. We employ $R = 3$ ($B_1 = 51$) for the confidence intervals and standard errors reported in the paper, which prioritizes the large-sample properties within each subsample. Table A.2 in the Appendix compares standard errors for different values of the tuning parameter R and shows they are robust.

D Advantages of Matching Estimators with Electricity Data

A key advantage of the nearest-neighbor approach is that it eliminates control communities that are not observationally similar to treated communities and whose inclusion would thus 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 Figure A.6 in the Appendix), and the degree of seasonality varies widely across the different communities in our sample.

²⁰Abadie 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.

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. We match on usage during the 24 months in 2008 and 2009. The matching window is indicated by the vertical dotted lines in panels (b) and (d).

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 ($Y_{it} - \hat{\mu}_i^{m(t)}$) for aggregation (treated) and non-aggregation (control) communities. Even after accounting for community-specific seasonality using monthly data from 2008 and 2009, usage varies greatly within and across years: the largest peak occurs in July 2012, which corresponds to a record heat

wave. By contrast, summer peaks are much less pronounced in 2009 and 2013, when the summers were mild. The difference between these two time series, which corresponds to an event study regression with community-specific month-of-year fixed effects, is displayed in panel (c). The increase in the difference is visible beginning in late-2011, which can be attributed to the implementation of aggregation, but this difference is quite noisy. The heterogeneity in seasonal patterns poses a challenge for a standard regression that compares treated communities to all control communities in the sample: it is difficult to estimate an effect when the baseline month-to-month divergence in usage is of the same order of magnitude as the effect.

Panels (b) and (d) of Figure 3 show analogous plots for the nearest-neighbor matching approach discussed above. Vertical dotted lines indicate the matching window of 2008 through 2009. 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), however, 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.²¹


IV Results

A Main Results

We first show that electricity prices fell substantially and persistently following the passage of aggregation referenda. Panel (a) in Figure 4 displays the average change in log prices for aggregation communities, relative to their matched controls ($\overline{\Delta \ln p_t}$). 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 0.27 log points (about 24 percent) 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 2). Nonetheless, ag-

²¹In terms of demographic characteristics, communities that implemented aggregation are significantly larger, younger, and more educated than those that did not. However, after matching on electricity usage, the matched controls are more similar to the aggregation communities along these and other dimensions. See Table A.3 in the Appendix for a comparison. Note that matching directly on these additional variables would select controls that are less similar in terms of pre-aggregation usage and would add noise to our estimates.

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




Notes: Panel (a) shows the effect of aggregation on the log of the electricity price, as calculated by equation (4). Because treatment and control communities faced the same price prior to aggregation, the pre-period difference is exactly zero. Panel (b) shows the effect of aggregation on average log electricity usage, as calculated by equation (3). These estimates are normalized so that the average usage difference in the year prior to the referendum is zero. The short dashed line displayed in both panels indicates the median implementation date relative to when the referendum was passed (4 months). Confidence intervals are constructed via subsampling.

gregation prices stay at least 10 percent lower than the control communities for most of the remaining estimation period.

Panel (b) in Figure 4 displays the corresponding estimates for electricity usage ($\widehat{\tau}_t$). Prior to the referendum, the difference in usage between aggregation and control communities is nearly 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 successful aggregation referenda by at least two years (see Table 1).²² Following the referendum, relative usage in aggregation communities increases by about 0.06 log points (6.1 percent) by the end of first year after the referendum. This difference shrinks to about 0.04 log points (4.1 percent) at month 15 because of the ComEd price decrease; it then stabilizes. Because the aggregation prices are stable around the date of the ComEd price drop (see Figure 2), 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.

²²Specifically, 270 of the 289 aggregation communities have virtually no overlap between the matching period and the pre-period estimates in Figure 4.

Figure 5: Reduced-Form Price Elasticities



Notes: This figure plots estimates of $\hat{\beta}_g$ from equation (5). Each point estimate comes from a separate regression of estimated community-specific changes in usage on the corresponding price changes for a particular post-referendum time period g . Panel (a) reports estimates from a series of monthly regressions. Panel (b) reports estimates from an analogous specification that groups the data into six-month intervals. The estimates in panel (b) are also reported in Table 2. Confidence intervals are constructed via subsampling.

Figure 5 displays reduced-form estimates of the price elasticity of demand (equation (5)). We present both monthly and biannual elasticities. Due to their greater precision, we consider the biannual estimates our primary specification, and we summarize them in Table 2.²³ The elasticity estimates increase in magnitude from about -0.09 in the first 6 months following the referendum up to -0.27 two years later, 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 respond fully to price changes in the near term.

The time-varying elasticity shown in Figure 5 is not due to the delay between the dates of the referenda and the dates of the actual price changes. While it is true that usage patterns in Figure 4 reflect the lag between referenda and implementation, so do the price patterns. 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.

The sharp decrease in the estimated monthly elasticity at month 15 in Figure 5 corre-

²³Point estimates for monthly elasticity estimates can be found in Table A.4 in the Appendix. Table A.5 reports the corresponding yearly estimates. We have also estimated a specification that models the price elasticity as a quadratic function of the number of months since the referendum. The results, displayed in Figure A.7, are very similar.

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 Log Usage and Log Price columns report average effects as calculated by equations (3) and (4), respectively. The Elasticity column reports regression estimates of $\hat{\beta}_g$ from equation (5). Each elasticity estimate corresponds to a separate regression. The number of price observations corresponds to the number of observations used for each elasticity estimate. Standard errors, given in parentheses, and statistical significance levels are calculated via subsampling.

sponds to the large ComEd price decrease. Because the adjustment process is dynamic, the usage difference does not shrink as quickly as the price difference, which causes the estimated monthly price elasticity to increase in magnitude temporarily. This demonstrates a challenge with interpreting the reduced-form price elasticity: our estimate captures a mix of short-run and longer-run responses to (1) the price decrease due to aggregation; (2) the drop in ComEd prices in June 2013; and (3) the monthly variation in the ComEd rate. In addition, consumers may respond to anticipated price changes. These caveats aside, the growth over time in the magnitude of the estimated elasticity suggests that the long-run effects of aggregation dominate the shorter-run responses to other price changes.

The finding that consumers are more elastic in the long run than in the short run is consistent with several mechanisms, including habit formation, inattention, learning, and slow-moving investments in appliances and home insulation. For example, Brandon et al. (2017) find that 35–55 percent of the energy reductions due to a non-price intervention (a home energy report) persist even after the treated household moves out of the dwelling, suggesting that investments in physical capital matter for energy consumption changes in

their setting. In order to gauge the potential role of appliance replacement in our setting, we perform a simple back-of-the-envelope calculation. Suppose that new appliances bought in non-aggregation communities used 10–20 percent less energy than new appliances bought in aggregation communities.²⁴ To account for the estimated 4 percent energy use difference two years after aggregation, 20–40 percent of appliances would need to be replaced within this time period. This value is higher than the typical two-year replacement rate for major household appliances, which ranges from 5–20 percent.²⁵ Thus, while appliance replacement behavior can explain some of our findings, other mechanisms—such as those mentioned above—likely play a role. Because we do not have data on consumer behavior beyond electricity consumption, we do not attempt to further distinguish among these different adjustment channels.

B Robustness Checks and Extensions

Our main identifying assumption is that the price changes caused by aggregation were uncorrelated with expected future changes in usage, conditional on our control communities. If suppliers offered different prices based on predicted future usage, this *price targeting* might generate bias in our estimates. Our nearest-neighbor controls are selected based on a community’s historical usage data, which are the same data used by suppliers to determine prices. Therefore, we think that systematic correlation between price and unobservable factors that affect post-referendum usage is unlikely. We also perform a simple empirical test for price targeting by suppliers. We split the treated communities into seven equal groups based on the price change they experienced in the first two years following their referenda. We then calculate the average elasticity separately for each group by pooling observations in this two-year period and estimating equation (5). If suppliers were offering prices to communities based on expected changes in demand, then we might find a systematic relationship between estimated elasticities and prices. However, we find no clear relationship between the two. Figure A.8 in the Appendix plots these estimates.

²⁴This difference reflects the gap between ENERGY STAR efficiency and the minimum required efficiency. Houde (2018) shows that appliance manufacturers tend to bunch their product offerings at these two levels.

²⁵Clothes washers, clothes dryers, and water heaters have a lifespan of about 10 years, refrigerators have a lifespan of about 15 years, and air conditioners have a lifespan of 15–20 years. Moreover, some aggregation households are likely to continue purchasing ENERGY STAR appliances, while some non-aggregation households purchase non-ENERGY STAR appliances, implying that an even larger replacement rate would be necessary to generate the effects we observe. For reference, in 2009, 54 percent of clothes washers, 4 percent of water heaters, 40 percent of refrigerators, 34 percent of air conditioners, and 76 percent of dishwashers sold in Illinois were ENERGY STAR certified (Environmental Protection Agency, 2018).

We chose January 2008 through December 2009 as our matching period because it allowed us to closely match controls to treated communities while also providing 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 2009.²⁶ These results are very similar to those in the paper and are available upon request. Additionally, we show in Table A.1 of the Appendix that 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.

As a placebo test, we estimate how usage evolves for communities that passed a referendum but never implemented aggregation (see Figure A.9 in the Appendix). Although the estimates are noisy, they suggest that there was no increase in usage due to the referendum itself in those communities.

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 Section 3 of 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 expected growth in electricity usage.

Our setting also allows us to test for pre-period changes in behavior. 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. In results shown and discussed in Section 4 of the Appendix, we find that consumers responded prior to 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 anticipation or confusion. Regardless of the mechanism, this result provides evidence that consumers begin adjusting their electricity usage prior to realized price changes, suggesting that the myopic “partial adjustment” model typically estimated in the energy literature (Hughes et al., 2008; Alberini and Filipini, 2011; Blázquez et al., 2013) could be systematically biased.

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

²⁶We do not include January 2007 in the alternative matching period because reported usage in this month is inexplicably low for many communities. This shortfall suggests that the usage data in the first month of the sample are incomplete.

U.S. and Illinois demographics in Section 5 in the Appendix. To understand how our results generalize to geographies with different demographics, that section also explores how our elasticity estimates vary by communities' socioeconomic characteristics.

V Accounting for Dynamics

A Framework

The presence of dynamics poses a challenge to interpreting the results in the previous section: because relative prices are changing over time (see Figure 4), contemporaneous usage reflects both a short-run response to recent price changes and a longer-run response to earlier price changes. In general, dynamics in consumer behavior imply that usage is a function of past and (potentially) future prices.²⁷ In this section, we explicitly account for these dynamics by regressing our community-specific matching estimates of log usage changes, $\hat{\tau}_{it}$, on lags and leads of log price changes:

$$\hat{\tau}_{it} = \sum_{r=-L_1}^{L_2} \delta_r \cdot \Delta \ln p_{i(t-r)} + \eta_{it}. \quad (6)$$

The number of leads in the regression is equal to L_1 , and the number of lags is L_2 . We consider four specifications, constructed such that each has a different set of 31 coefficients: $(L_1, L_2) \in \{(18, 12), (12, 18), (6, 24), (0, 30)\}$. To convert these estimates to a price elasticity s months after a price change, we construct the dynamic elasticity parameter

$$\beta_s := \sum_{r=-L_1}^s \delta_r. \quad (7)$$

Although our usage data end in June 2014, we observe prices through 2016, allowing us to estimate equation (6) without losing observations. To reduce sensitivity to outliers, we estimate this model using median (minimum absolute deviation) regressions instead of least squares.²⁸

To reduce noise in the estimates, we also estimate a parametric specification of equation (6). This specification still relies on our quasi-experimental variation, but it restricts the

²⁷See the conceptual framework presented in Section 6 of the Appendix for a simple demonstration.

²⁸Point estimates from least squares regressions are similar, but the subsampling routine described in Section III 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.

relationships among the set $\{\delta_r\}$ in equation (6). Specifically, we assume the cumulative elasticity can be modeled as an exponential function of lags and a linear function of leads:

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

where s is again the number of months since a price change. Thus, negative numbers indicate the price change has not happened yet. According to equation (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 before and after a price change. In addition to delivering more precise estimates, the parametric specification also allows us to calculate elasticities outside our data window. We estimate equation (8) by selecting the parameters that minimize the mean absolute deviation between the observed usage response and the response predicted by the model.

Equations (6) and (8) implicitly assume that consumers have perfect foresight. Following Anderson et al. (2013), we have also estimated the model using status quo (“no change”) expectations. The results are similar to the ones we present below.


B Results

The results from the dynamic model are presented in Figure 6. The markers report dynamic elasticities, as estimated by equations (6) and (7), for the four specifications described earlier. Each marker represents the cumulative electricity usage change (in percent) in response to an anticipated one-percent price change that begins in month 0 and persists. The solid line corresponds to the parametric model specified by equation (8) and lines up closely with the unrestricted estimates.²⁹ Note that the period $t = 0$ in Figure 6 corresponds to the date of the price change, but in Figures 4 and 5, period $t = 0$ corresponds to the date of the referendum, which occurs four months earlier for the median community.

Overall, the estimates follow a similar pattern to the reduced-form results presented in Section IV. Following the price change, the usage response grows over time, with an implied elasticity of about -0.1 six months post-implementation and an elasticity of about -0.2 eighteen months post-implementation. Thus, consumer dynamics play an important role in the first two years after a price change. The dynamic model also estimates that consumers begin to respond in advance of a price change, though the pre-price-change

²⁹The point estimates for the parameters $(\hat{\gamma}_1, \hat{\gamma}_2, \hat{\gamma}_3, \hat{\gamma}_4)$ are $(-0.052, -6.034, -0.296, 0.036)$.

Figure 6: Estimated Price Elasticities: Dynamic Model



Notes: This figure reports non-parametric and parametric estimates of the dynamic elasticity curve. The plotted points are derived from equation (6). Each set of points is estimated using a different set of 30 lags and leads, beginning with the lead displayed in the legend. For example, the “12-Month Leads” specification includes 12 leads and 18 lags. Each point displays the estimated elasticity as a function of the number of months since a price change, calculated as $\sum_{r=-L_1}^s \hat{\delta}_r$, where s is the number of months since a price change and L_1 is the number of leads included in the specification. The solid line displays estimates derived from the parametric model specified by equation (8). Six-month average elasticities for both the non-parametric and parametric estimates are reported in Table 3.

elasticities are small, averaging -0.02 in the six months prior to a price change.

The first two columns of Table 3 summarize the non-parametric and parametric estimates of the dynamic elasticity curve, respectively. The reported coefficients are the mean of the point estimates in Figure 6 within each interval indicated by each row. Again, these estimates correspond to the effect of a permanent, anticipated increase in price beginning at time $t = 0$ on usage s months after. These estimates indicate an elasticity of approximately -0.085 in months 1–6 and an elasticity of -0.21 in months 19–24.³⁰

For comparison, we also provide analogous reduced-form estimates in column (3) of Table 3. To create them, we replicate the reduced-form elasticity estimates presented in Section IV, but re-index the estimates to the date of the price change, rather than the date of the referendum.³¹ To provide the closest possible comparison for the dynamic model,

³⁰Table A.8 in the Appendix shows that these results are not sensitive to the number of neighbors used.

³¹Because the price change prior to the implementation date is zero, we cannot estimate an elasticity for the months prior to the policy change with the reduced-form model. Additionally, this reduced-form estimator

Table 3: Estimates of the Dynamic Elasticity Curve

Period After Price Change	Non-Parametric	Parametric	Reduced Form
1-6 Months Prior	-0.023** (0.010)	-0.022*** (0.006)	N/A
Contemporaneous	-0.050*** (0.015)	-0.052*** (0.010)	-0.034** (0.021)
1-6 Months	-0.083*** (0.013)	-0.087*** (0.012)	-0.069*** (0.011)
7-12 Months	-0.145*** (0.018)	-0.138*** (0.017)	-0.122*** (0.017)
13-18 Months	-0.168*** (0.022)	-0.179*** (0.023)	-0.218*** (0.035)
19-24 Months	-0.209*** (0.031)	-0.212*** (0.028)	-0.260*** (0.038)

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. The dynamic estimates are constructed from a regression of log usage changes on leads and lags of log price changes. The non-parametric elasticities are calculated according to equation (7) and have been averaged across the four different specifications of equation (6) discussed in the main text. The parametric elasticities are calculated according to equation (8). The reported estimates represent the average cumulative effect of a permanent one-percent price change on log usage over the six-month interval shown in the first column. The reduced-form estimates are constructed by regressing log usage changes on log price changes in the same period. These reduced-form estimates differ from those presented in Table 2 because the reference period here is time since price change, rather than time since referendum. Standard errors, given in parentheses, and statistical significance levels are constructed via subsampling.

we estimate this specification using median (minimum absolute deviation) regressions. The resulting reduced-form estimates are similar to those reported in Table 2 but shifted by the number of elapsed months between the referendum and the price change.

The estimated elasticities in Table 3 grow steadily over time, though the dynamic model estimates are somewhat smaller than the reduced-form estimates in later periods. Given that municipal aggregation did not lead to a time-invariant price difference between aggregation and non-aggregation communities, it is not surprising that there are some differences in these estimates. The decrease in ComEd prices occurring about 11 months after implementation causes the reduced-form estimates to overstate the elasticity after this point, as contemporaneous usage still partially reflects the earlier (larger) price difference. Nevada does not account for anticipation, but this has little impact on the resulting estimates.

Table 4: Long-Run Elasticity Forecasts from Dynamic Model

Period After Price Change	Point Estimate
Year 3	-0.268*** [-0.358, -0.195]
Year 5	-0.315*** [-0.462, -0.222]
Year 10	-0.345*** [-0.616, -0.242]
Long Run	-0.348*** [-1.007, -0.245]

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. These forecasts are an extrapolation based on the parameters estimated by equation (8). The reported coefficients represent the cumulative effect of a permanent one-percent price change on log usage. 95 percent confidence intervals, given in brackets, and statistical significance levels are constructed via subsampling.

ertheless, the estimates are quite similar overall, which suggests that modeling municipal aggregation as a permanent price shock is reasonable in our setting.

Although the primary benefit of the dynamic model is that it disentangles the effects of contemporaneous versus lagged price changes on current usage, the estimated parameters can also be used to project the usage response beyond the two-and-a-half-year post period in our data. As a calibration exercise, we report the long-run projections arising from our parametric model in Table 4. We estimate a three-year elasticity of -0.27 and a five-year elasticity of -0.32 . Both the ten-year and the long-run elasticity are close to -0.35 . Thus, over 90 percent of the long-run consumer response occurs within five years of the price change, and nearly 100 percent occurs within ten years. These estimates imply that prices from more than ten years ago have little (direct) impact on residential electricity consumption today. These dynamics are similar to those found in Allcott and Rogers (2014), who find that the consumption effects of a home energy report decay by 10–20 percent per year after it is discontinued, implying that consumers fully adjust after 5–10 years.

VI Discussion

The fact that the price elasticity of residential electricity demand grows significantly over time has two important implications. First, price changes will have much larger effects on consumption in the long run than in the short run. Second, any change to the market, such as a tax or increased generation, will have a smaller effect on consumer prices in the long run relative to the short run. Policymakers who do not anticipate these dynamics will both underestimate the long-run effects of regulations that affect electricity prices and overestimate the share of the long-run regulatory burden borne by consumers.

We demonstrate the quantitative magnitudes of these implications with a few simple calculations. Because our model was estimated using Illinois data, our exercises employ Illinois data on electricity consumption, prices, and carbon dioxide (CO₂) emissions. The exact data we use do not affect the primary implications of our analysis, as we focus on relative comparisons (ratios) between the short- and longer-run elasticities.

According to the U.S. Energy Information Administration (2018), in 2016 Illinois generated 187,441,635 MWh of electricity at an average retail price of \$93.8 per MWh. This in turn produced 72,226 thousand metric tons of CO₂. For simplicity, we assume that the residential electricity market is perfectly competitive and that all generated electricity is sold at the average retail price. Because our goal is to highlight the role of the demand elasticity, we abstract away from supply-side issues such as changes in the composition of generation. Finally, because there is little consensus on the magnitude of the supply elasticity, ϵ_s , we perform all calculations under two different assumptions: $\epsilon_s = 0.5$ (low elasticity of supply) and $\epsilon_s = 5$ (high elasticity of supply).

Our first exercise considers the implementation of a tax intended to reduce emissions. Policymakers who wish to target a specific level of emissions must first predict how equilibrium electricity consumption responds to changes in taxes. In perfectly competitive markets, the fall in equilibrium quantity following a small tax increase depends on both the demand and supply elasticities (Salanie, 2011):

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

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 underestimating the demand elasticity will underestimate the demand response following a price change. We demonstrate the significance of this point in Table 5. Employing our six-month

Table 5: Implications of the Demand Elasticity for Carbon Taxes and Incidence

	(1) $\epsilon_D = -0.09$ (1-6 Months)	(2) $\epsilon_D = -0.27$ (19-24 Months)	(3) Ratio of (2) to (1)
(a) Inelastic Supply ($\epsilon_S = 0.5$)			
Change in Quantity for a \$1/MWh tax (MWh)	-152,414	-350,353	2.30
Tax to Reduce Emissions by 1% (\$/ton CO ₂)	31.92	13.88	0.44
Consumer Share of Tax Burden (%)	85	65	0.77
(b) Elastic Supply ($\epsilon_S = 5$)			
Change in Quantity for a \$1/MWh tax (MWh)	-176,668	-511,901	2.90
Tax to Reduce Emissions by 1% (\$/ton CO ₂)	27.53	9.50	0.35
Consumer Share of Tax Burden (%)	98	95	0.97

This table presents calculations of the effects of an electricity tax on electricity output; the size of a carbon dioxide (CO₂) tax required to reduce electricity-related emissions by 1%; and the incidence of these policies. Column (1) presents results employing the estimated short-run (6-month) elasticity from Table 2 (−0.09). Column (2) presents results using the corresponding two-year elasticity (−0.27). Calculations are based on 2016 Illinois statistics obtained from U.S. Energy Information Administration (2018): 187,441,635 MWh of electricity generated; 72,226 thousand metric tons of CO₂ emitted; average retail price of \$93.8 per MWh. All calculations assume perfectly competitive markets.

estimated elasticity and assuming a supply elasticity of $\epsilon_s = 0.5$, we calculate that electricity output falls by 152,000 MWh per dollar increase in tax. By contrast, employing our two-year estimate results in an estimated decrease that is over twice as large, 350,000 MWh per dollar. Panel (b) shows that this difference becomes even larger if we instead assume a supply elasticity of $\epsilon_s = 5$.

Similarly, our second exercise shows that employing smaller demand elasticity estimates will reduce the predicted effectiveness of a carbon tax. Assuming that all emissions reductions must come from reduced electricity generation, a demand elasticity of −0.09 implies that a 1 percent reduction in emissions would require a tax of \$27.5–\$31.9 per ton of CO₂, while a demand elasticity of −0.27 implies that a \$9.5–\$13.9 tax would be sufficient.³² In other words, the two-year elasticity we estimate yields tax rates that are 56–65 percent smaller than our six-month elasticity would imply. It is worth noting that, because

³²To calculate the required CO₂ tax to obtain emissions reductions of $\alpha \in (0, 1)$, we calculate $\frac{\alpha Q}{\eta \frac{\partial Q}{\partial t}}$, where $\eta = \frac{72,226,000}{187,441,635} = 0.385$ is the Illinois emissions intensity. The corresponding carbon tax would be 3.67 times larger than the CO₂ tax, as 3.67 tons of CO₂ contain one ton of carbon. Because carbon emission rates depend on an electricity generator’s type (coal, natural gas, wind, etc.), calculating the required carbon tax in practice is more complicated. Nevertheless, the influence of the demand elasticity is similar.

both the six-month and two-year estimated elasticities are quite inelastic, the taxes required to reduce emissions through reducing the quantity of electricity consumed are large.

Finally, our third exercise considers the partial equilibrium incidence of a tax, which also depends on the elasticity of demand (Salanie, 2011). In perfectly competitive markets, the share of the overall burden borne by consumers is equal to $\frac{\epsilon_s}{\epsilon_s + \epsilon_d}$. We report the short-run and longer-run incidence in Table 5.³³ If we assume a supply elasticity of $\epsilon_s = 0.5$, then we estimate that consumers would bear 85 percent of the burden when the demand elasticity is -0.09 , but only 65 percent of the burden when the demand elasticity is -0.27 . When supply is elastic ($\epsilon_s = 5$), however, the consumer share of the tax burden is 95 percent or larger across our range of demand elasticity estimates.³⁴

VII Conclusion

It is essential for electricity suppliers, market regulators, and policymakers to understand how electricity consumption responds to price changes. For example, accurately predicting the long-run effect of a carbon tax on electricity consumption (and any accompanying emissions) requires a good estimate of the long-run price elasticity of demand. Few reliable estimates of this important parameter exist because electricity price changes are often endogenous, short-lived, small, or unnoticed. Our study provides the first quasi-experimental estimate of the two-year price elasticity of demand for residential electricity, and we find that it is more than double the short-run (six-month) elasticity, although it is still inelastic. In addition, our long-run projections suggest that it takes ten years for consumers to fully respond, which implies that consumer behavior may depend upon price changes that occurred up to ten years ago.

There are several possible reasons why our estimated elasticity grows over time. 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

³³Weyl 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.

³⁴One concern that arises when discussing incidence in the electricity sector is that residential electricity rates are often controlled by utility regulators. As of 2016, however, fifteen states had retail choice for electricity supply (Morey and Kirsch, 2016). Some of these states, as well as others that do not allow retail choice, impose profit constraints on incumbent suppliers (e.g., the zero-profit condition in Illinois). Thus, the supply component of price is tightly linked to changes in the costs of electricity generation. In such cases, the standard incidence framework remains relevant.

continuously replaced within a population. 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. Whatever the underlying mechanism, our results underscore the importance of identifying settings that accurately capture longer-run responses to price changes, as short-run data may cause the researcher to significantly underestimate the long-run response.

Our results also 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 suggest that using a downward-sloping demand curve in this setting is appropriate.

Finally, we note that the natural experiment created by municipal aggregation decreased electricity prices, whereas price-based climate policies would increase prices to reduce total carbon emissions. It is therefore important to know whether price increases and decreases have symmetric effects on demand. In addition, other energy policies such as a carbon tax may be less salient or more salient to consumers than aggregation, which in turn may affect the speed and magnitude of the consumer response. Future research along these dimensions 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 T. Rogers (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review* 104(10), 3003–3037.
- 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.
- 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.
- Brandon, A., P. J. Ferraro, J. A. List, R. D. Metcalfe, M. K. Price, and F. Rundhammer (2017). Do the effects of social nudges persist? theory and evidence from 38 natural field experiments. National Bureau of Economic Research working paper 23277.
- 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.
- Dergiades, T. and L. Tsoulfidis (2008). Estimating residential demand for electricity in the united states, 1965–2006. *Energy Economics* 30(5), 2722–2730.

- 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.
- Environmental Protection Agency (2018). Energy star unit shipment and sales data archives. https://www.energystar.gov/index.cfm?c=partners.unit_shipment_data_archives, Accessed September 17, 2018.
- Filippini, M., B. Hirl, and G. Masiero (2018). Habits and rational behaviour in residential electricity demand. *Resource and Energy Economics* 52, 137–152.
- 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.
- Houde, S. (2018). Bunching with the stars: How firms respond to environmental certification. E2e Working Paper 037.
- 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.
- Ito, K. (2015). Asymmetric incentives in subsidies: Evidence from a large-scale electricity rebate program. *American Economic Journal: Economic Policy* 7(3), 209–37.
- Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review* 104(4), 1417–1438.
- Kamerschen, D. R. and D. V. Porter (2004). The demand for residential, industrial and total electricity, 1973–1998. *Energy Economics* 26(1), 87–100.
- Karlan, D. and J. Zinman (2018). Long-run price elasticities of demand for credit: Evidence from a countrywide field experiment in Mexico. *The Review of Economic Studies*. forthcoming.
- 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.

- 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.
- 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.
- Reif, J. (2018). A model of addiction and social interactions. *Economic Inquiry*. forthcoming.
- 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/>.
- 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 September 6, 2018.
- U.S. Energy Information Administration (2017). Average retail price of electricity to ultimate customers. <https://www.eia.gov/electricity/data.cfm>, Accessed February 2, 2017.
- U.S. Energy Information Administration (2018). State electricity profiles: Illinois electricity profile 2016. <https://www.eia.gov/electricity/state/illinois/index.php>, Accessed August 20, 2018.
- 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.

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.

Appendix (For Online Publication)

1 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:¹

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.1 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.2 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.3 and A.4 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.4) will change. Otherwise her bill will remain the same.

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-of-year indicators 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.

¹From 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>.

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.

Together, these two filters remove 79 communities from our initial sample of 887 service territories. We remove an additional 29 communities from our sample because they were missing data and/or we could not confidently assign them to our treatment or control groups. Finally, we also drop 11 communities that passed a referendum approving aggregation but never implemented the program, leaving us with 768 (= 887 - 79 - 29 - 11) communities in our final sample.

One concern with the two filters that we applied is that we may eliminate actual structural breaks arising from municipal aggregation. However, the communities that were removed in this fashion were 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.

Sources for Aggregation Rates

This section summarizes the data sources used to verify the rates obtained by aggregation communities. The primary data source is Plug In Illinois. If information was unavailable from Plug In Illinois, we tried to retrieve it from websites of the participating communities or from FOIA requests made to the local governments.

Plug In Illinois The website <https://pluginillinois.org/MunicipalAggregationList.aspx> lists the current state of aggregation programs in Illinois. If a community is currently on an aggregation plan, the rate and the end of the contract are listed. Repeated snapshots of the website were used

to track the rates over time. For many of the communities, Plug In Illinois provides a link to access the opt-out letter and a link to the terms and conditions of service.

Opt-Out Letter The opt-out letter is sent to all households following passage of the referendum. The letter confirms the aggregation program and provides them with an opportunity to opt out and remain with ComEd. The letter also includes the aggregation rate and the relevant time period.

Terms and Conditions of Service The terms and conditions are sent to households that do not opt out of the aggregation program. The terms and conditions are an agreement between a household and the alternative retail electric supplier, and they indicate the aggregation rate and the relevant time period.

Master Agreement The master agreement is the contract between the alternative retail electric supplier and the local government. The agreement indicates the term of the agreement and the agreed-upon rate. We obtained these through FOIA requests.

3 *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

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

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 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. The main parameter of interest is β_{τ} .

To ensure that β_{τ} reflects changes relative to the omitted 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 also include 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:

$$Y_{cmy} = \gamma_1 A_{c,0 \text{ to } 6} + \gamma_2 A_{c,7 \text{ to } 12} + \gamma_3 A_{c,13 \text{ to } 18} + \gamma_4 A_{c,19 \text{ to } 24} + \beta_{25} A_{c,25} + \beta_{-25} A_{c,-25} + \alpha_{cm} + \alpha_{my} + \varepsilon_{cmy}. \quad (\text{A.2})$$

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.1).

One could 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 an adequate counterfactual for the former without correcting for systematic differences via matching, as we do in the main text. We therefore restrict our estimation sample here to communities that implemented aggregation.² Our main identifying assumption is that, conditional on the included fixed effects, the timing of aggregation adoption is exogenous with respect to electricity use.

Figure A.10 presents the change in electricity prices following aggregation, in logs, as estimated by equation (A.1). 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 price 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.10 shows the point at which half of all communities have implemented aggregation (4 months after passing the referendum). Prices continue to drop as more communities switch and then eventually stabilize. Within 8 months of passing the referendum, the average electricity price has decreased by more than 0.3 log points (26 percent) in aggregation communities relative to the control group.

Figure A.11 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.6 shows the estimated impact of aggregation on the log of the electricity price in

²We also estimated a specification that uses all non-aggregation communities as controls. Those estimates are similar to the ones we present here, although the coefficients are much less stable.

these communities 0–6, 7–12, 13–18, and 19–24 months after implementation, as estimated by equation (A.2). The results consistently show large and significant price drops, and are robust to including different fixed effects. 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 around 0.3 log points by the end of the first year.

Table A.7 shows the estimated change in usage as estimated by equation (A.2) 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.12 shows the elasticities implied by the usage and price results. Similar 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.

4 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 a good 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 given in equation (5), 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

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.

aggregation communities and their matched controls is 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. 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 announcement 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). Alternatively, it is possible that customers were confused about the timing of the price change, although we note that mailers were sent to all residents informing them of the exact month of the price change. 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, to the extent that other policies are implemented in a similar manner, they are likely to generate pre-implementation effects similar to those we find here.

5 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, 2017).

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.).

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:

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

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 A.13 and A.14. Because the estimation is done jointly, the displayed elasticities for any given characteristic control for the other characteristics.

Figure A.13 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 A.14 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.

6 Framework for Consumption Dynamics

Summary

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. Here, we present a simple framework demonstrating that our empirical approach in Section V is consistent with a standard rational model of a forward-looking consumer.

We employ the habit model of Becker et al. (1994), and provide a full derivation of these results at the end of this 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.³ The consumer’s problem is:

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

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

³Similarly, one could allow utility in each period to depend on a “stock” of appliances (Filippini et al., 2018). The resulting model will exhibit dynamics similar to what we present here.

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

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. Finally, to illustrate the dynamics that can arise in our setting, we assume utility is quadratic.⁴ Under these assumptions, the demand equation is:

$$y_t = \alpha_1 y_{t-1} + \alpha_2 y_{t+1} + \alpha_3 p_t, \quad (\text{A.3})$$

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

$$y_t = \theta_1 y_{t-1} + \theta_2 p_t. \quad (\text{A.4})$$

In the forward-looking model (A.3), 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 (A.4). 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 (A.3):

$$y_t = \sum_{s=0}^{t-1} \delta_{t-s} p_{t-s} + \sum_{s=1}^{\infty} \delta_{t+s} p_{t+s}. \quad (\text{A.5})$$

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 (A.5), the long-run effect of a permanent change in price, p^* , on consumption can then be calculated as

$$\frac{dy}{dp^*} = \sum_{s=0}^{t-1} \hat{\delta}_{t-s} + \sum_{s=1}^{\infty} \hat{\delta}_{t+s}. \quad (\text{A.6})$$

An appealing feature of equation (A.5) is that it is more flexible than equation (A.3). For example, it will generate a valid estimate of dy/dp^* even if consumers do not discount the future exponentially. This non-parametric specification is rarely estimated, however, because most studies do not have data panels of sufficient length.

⁴Alternatively, 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.

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 (A.3) to obtain:

$$y_t = K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} \alpha_3 p_{t+s-1} + K_2 \sum_{s=0}^{t-1} (\lambda_2)^s \alpha_3 p_{t-s-1} + (\lambda_2)^t \left(y_0 - K_1 \sum_{s=1}^{\infty} (\lambda_1)^{-s} \alpha_3 p_{s-1} \right) \quad (\text{A.7})$$

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

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

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

See Reif (2018) for additional details.

Equation (A.7) 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.7) 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 (A.6) as $t \rightarrow \infty$.

It is straightforward to show that the solution to the first-order difference equation (A.4), the myopic “adjustment cost” model, depends only on past prices, and not on future prices.

Appendix Tables

Table A.1: Number of Nearest Neighbors and Reduced-Form Estimates

(a) 5 Nearest Neighbors					
	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
(b) 1 Nearest Neighbor					
	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
1-6 Months	0.016*** (0.004)	-0.098*** (0.003)	-0.098*** (0.022)	1692	1692
7-12 Months	0.051*** (0.007)	-0.249*** (0.007)	-0.159*** (0.020)	1668	1668
13-18 Months	0.047*** (0.006)	-0.147*** (0.002)	-0.248*** (0.030)	1516	1515
19-24 Months	0.044*** (0.007)	-0.132*** (0.003)	-0.301*** (0.049)	1155	1155
25-30 Months	0.047*** (0.008)	-0.120*** (0.004)	-0.306*** (0.046)	606	604
(c) 10 Nearest Neighbors					
	Log Usage	Log Price	Elasticity	Usage Obs.	Price Obs.
1-6 Months	0.015*** (0.003)	-0.098*** (0.003)	-0.099*** (0.019)	1692	1692
7-12 Months	0.051*** (0.007)	-0.249*** (0.007)	-0.162*** (0.020)	1668	1668
13-18 Months	0.044*** (0.006)	-0.147*** (0.002)	-0.228*** (0.027)	1516	1515
19-24 Months	0.040*** (0.006)	-0.132*** (0.003)	-0.283*** (0.044)	1155	1155
25-30 Months	0.044*** (0.007)	-0.120*** (0.004)	-0.283*** (0.040)	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 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. Standard error are determined by subsampling.

Table A.2: Comparison of Tuning Parameters for Subsampling

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

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

Table A.3: Characteristics of Aggregation, Non-Aggregation, and Matched Control Communities

	(1)	(2)	(3)	(4)	(5)
	Aggregation Mean	Non- Aggregation Mean	p -value of Difference (2)-(1)	Matched Controls Mean	p -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	

Mean characteristics for aggregation communities, non-aggregation communities, and matched control (non-aggregation) communities are reported. Electricity usage data come from ComEd. All other characteristics are from the 2005-2009 American Community Survey (ACS). We obtained non-ambiguous ACS matches for 286 out of 289 aggregation communities and 385 out of 479 non-aggregation communities. The number of observations in column (1) is smaller for median year built (285). The 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.4: 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. Months 0–2 are not shown as no aggregation community switches suppliers that quickly following a referenda, resulting in price changes that are identically zero and undefined reduced-form elasticities. 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.5: 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.6: 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.7: 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.8: Number of Nearest Neighbors and the Dynamic Elasticity Curve

(a) 5 Nearest Neighbors

Period After Price Change	Non-Parametric	Parametric	Reduced Form
1-6 Months Prior	-0.023** (0.010)	-0.022*** (0.006)	N/A
Contemporaneous	-0.050*** (0.015)	-0.052*** (0.010)	-0.034** (0.021)
1-6 Months	-0.083*** (0.013)	-0.087*** (0.012)	-0.069*** (0.011)
7-12 Months	-0.145*** (0.018)	-0.138*** (0.017)	-0.122*** (0.017)
13-18 Months	-0.168*** (0.022)	-0.179*** (0.023)	-0.218*** (0.035)
19-24 Months	-0.209*** (0.031)	-0.212*** (0.028)	-0.260*** (0.038)

(b) 1 Nearest Neighbor

Period After Price Change	Non-Parametric	Parametric	Reduced Form
1-6 Months Prior	-0.024** (0.012)	-0.023*** (0.007)	N/A
Contemporaneous	-0.043*** (0.018)	-0.056*** (0.012)	-0.010 (0.025)
1-6 Months	-0.084*** (0.014)	-0.085*** (0.013)	-0.071*** (0.012)
7-12 Months	-0.135*** (0.020)	-0.131*** (0.018)	-0.115*** (0.018)
13-18 Months	-0.162*** (0.023)	-0.172*** (0.024)	-0.226*** (0.040)
19-24 Months	-0.191*** (0.033)	-0.208*** (0.030)	-0.223*** (0.046)

(c) 10 Nearest Neighbors

Period After Price Change	Non-Parametric	Parametric	Reduced Form
1-6 Months Prior	-0.025*** (0.010)	-0.024*** (0.007)	N/A
Contemporaneous	-0.052*** (0.016)	-0.057*** (0.010)	-0.056*** (0.020)
1-6 Months	-0.086*** (0.013)	-0.092*** (0.012)	-0.071*** (0.011)
7-12 Months	-0.146*** (0.019)	-0.144*** (0.018)	-0.119*** (0.016)
13-18 Months	-0.176*** (0.024)	-0.185*** (0.024)	-0.244*** (0.035)
19-24 Months	-0.211*** (0.033)	-0.219*** (0.030)	-0.253*** (0.039)

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. Standard errors are constructed via subsampling.

Appendix Figures

Figure A.1: 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...

Handwritten signature of Christopher J. Lauzen in red ink.

Christopher J. Lauzen
Board Chairman
Kane County

Handwritten signature of Kurt R. Kojzarek in blue ink.

Kurt R. Kojzarek
Development Committee Chairman
Kane County

Figure A.2: Example of an Opt-Out Card

<div style="border: 1px solid black; padding: 5px; display: inline-block;">PLACE STAMP</div>
<p>MC SQUARED ENERGY SERVICES, LLC 344 South Poplar Street Hazleton, PA 18201</p>

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.3: Example of a ComEd Bill (page 1 of 2)

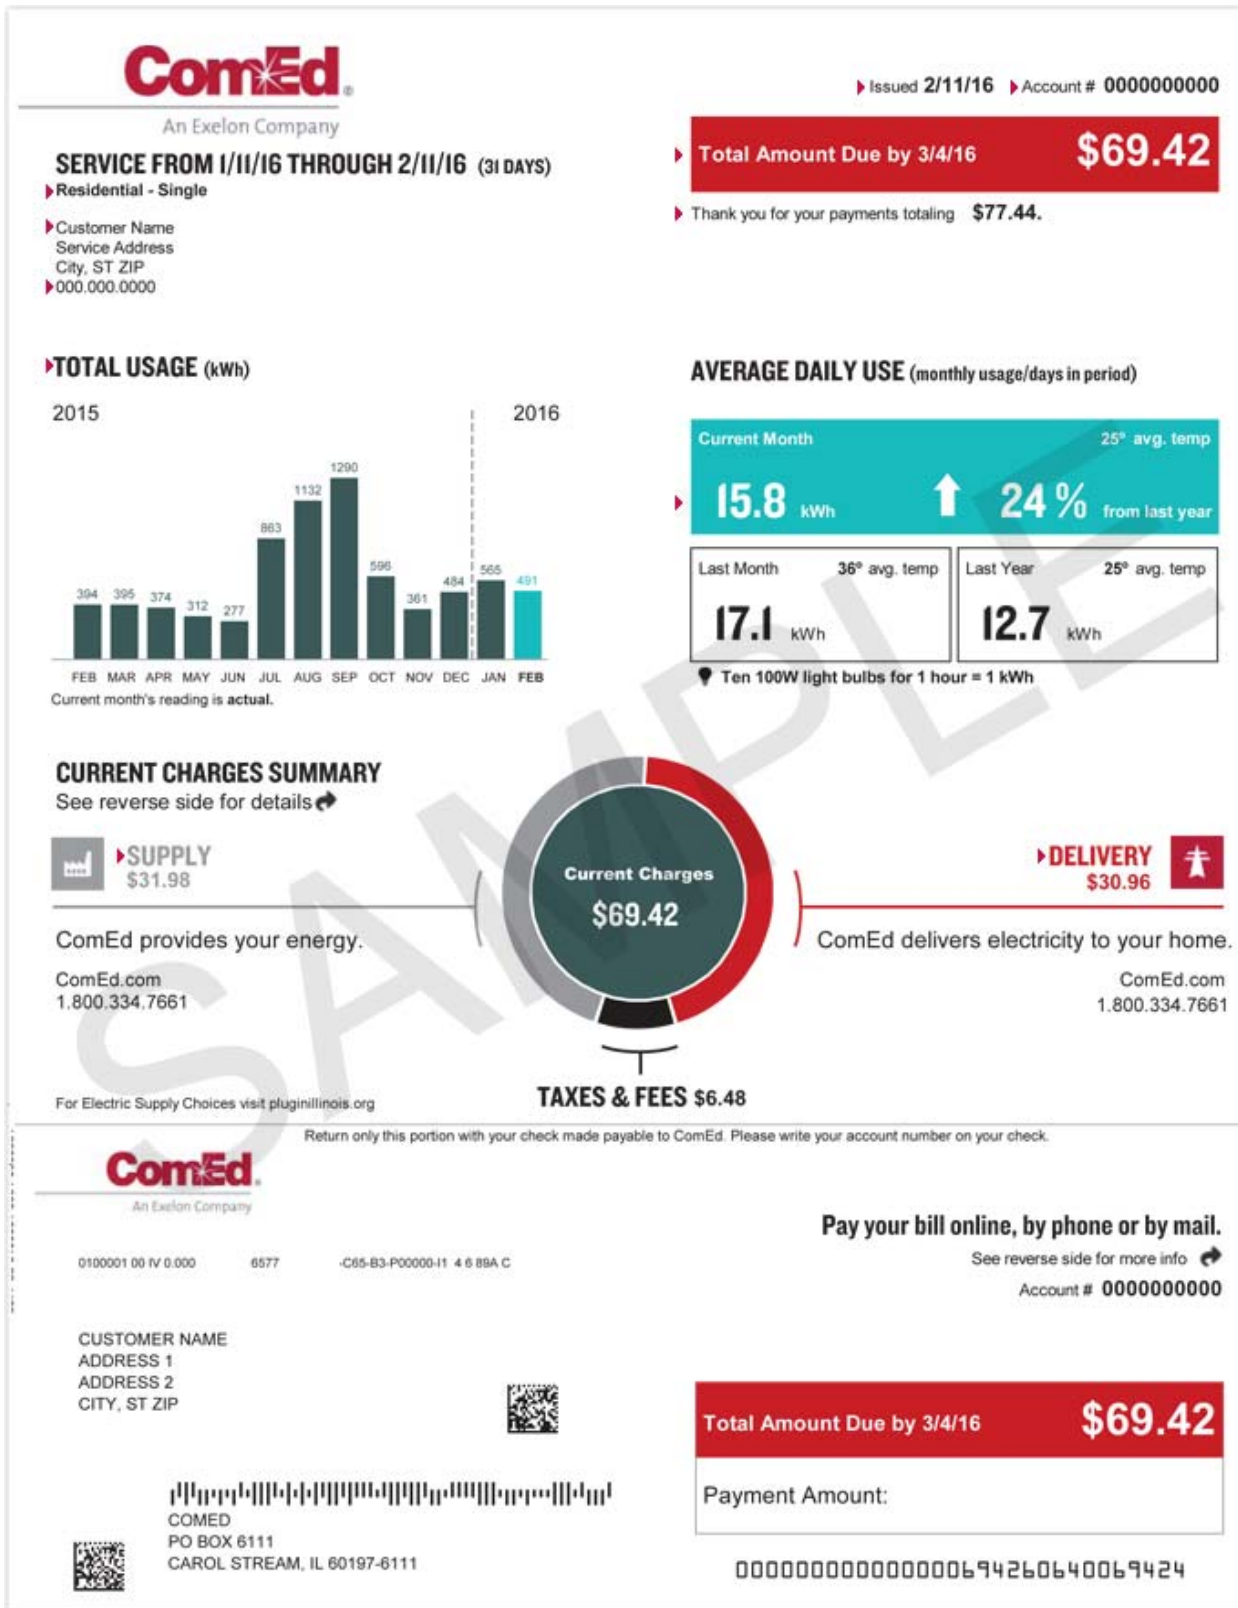


Figure A.4: Example of a ComEd Bill (page 2 of 2)

For Questions, Support, and Outages visit ComEd.com

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

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

▶ Total Amount Due by 3/4/16
\$69.42

METER INFORMATION

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

CHARGE DETAILS

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

SUPPLY **\$31.98**

▶ Electricity Supply Charge ▶ 491 kWh X 0.05865 \$28.80
 Transmission Services Charge 491 kWh X 0.01122 \$5.51
 Purchased Electricity Adjustment -\$2.33

DELIVERY - ComEd **\$30.96**

▶ Customer Charge \$10.53
 ▶ Standard Metering Charge \$4.36
 ▶ Distribution Facilities Charge 491 kWh X 0.03156 \$15.50
 ▶ IL Electricity Distribution Charge 491 kWh X 0.00116 \$0.57

TAXES & FEES **\$6.48**

▶ Environmental Cost Recovery Adj 491 kWh X 0.00038 \$0.19
 ▶ Energy Efficiency Programs 491 kWh X 0.00345 \$1.69
 ▶ Franchise Cost \$30.39 X 2.36300% \$0.72
 ▶ State Tax \$1.62
 ▶ Municipal Tax \$2.26

Service Period Total **\$69.42**

Thank you for your payment of \$77.44 on January 29, 2016

Total Amount Due **\$69.42**

UPDATES

ComEd

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

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

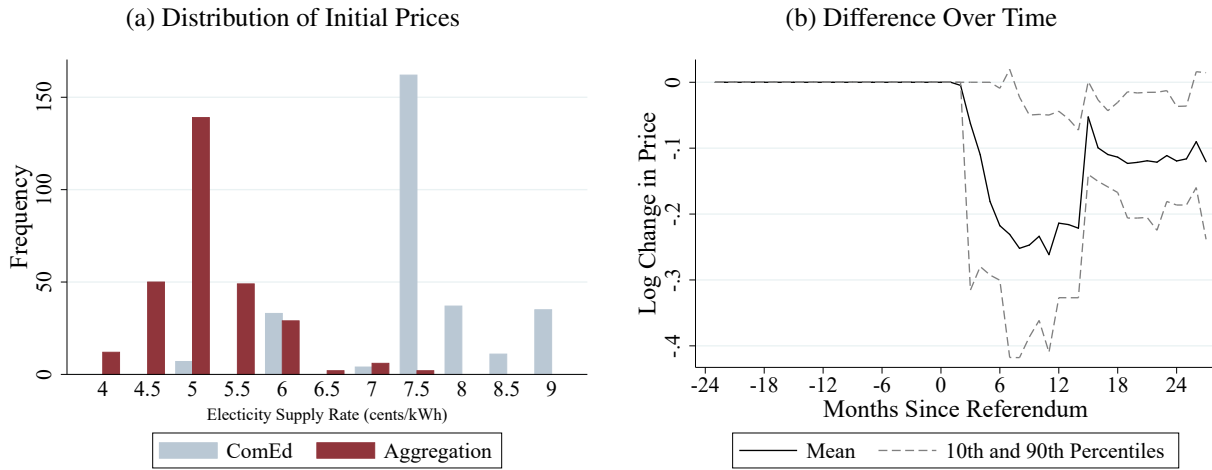
Pay your bill in-person at many ComEd authorized agents located throughout the region. Visit ComEd.com/Pay for details.

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

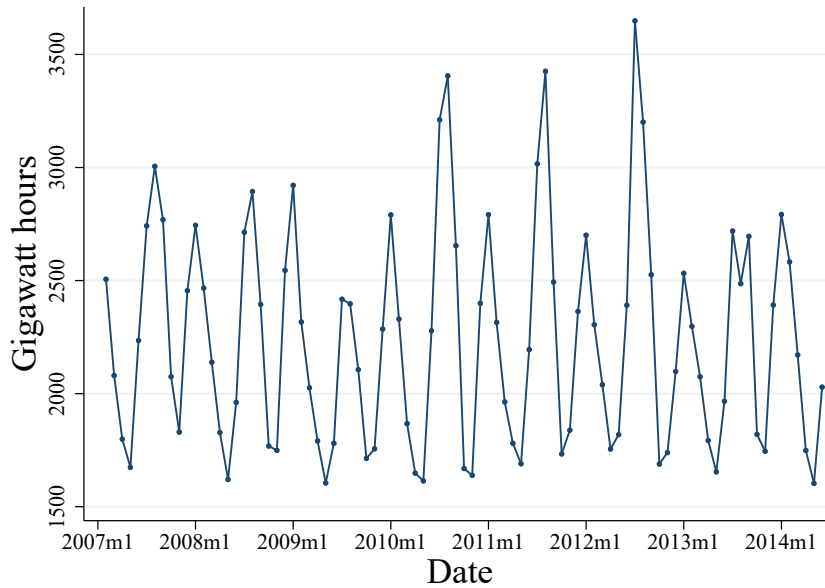
6577-65-0100001-0001-0000001

Figure A.5: Heterogeneity in Price Changes




Notes: This figure displays the distribution of (log) price changes for aggregation communities relative to the contemporaneous ComEd price. Panel (a) provides a histogram of the initial aggregation rate and the corresponding ComEd rate in the first month of the program for each community. Panel (b) plots the the difference between aggregation and ComEd rates over time. The solid line displays the mean, and the dashed lines represent the 10th and 90th percentiles.

Figure A.6: 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.7: Estimated Price Elasticities, Quadratic Fit




Notes: The reduced-form, 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 are constructed via subsampling.

Figure A.8: 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 by pooling observations in this two-year period and estimating equation (5). 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.9: 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.10: 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.11: Regression Estimates of the Effect of Aggregation on Electricity Usage, Communities that Passed Aggregation



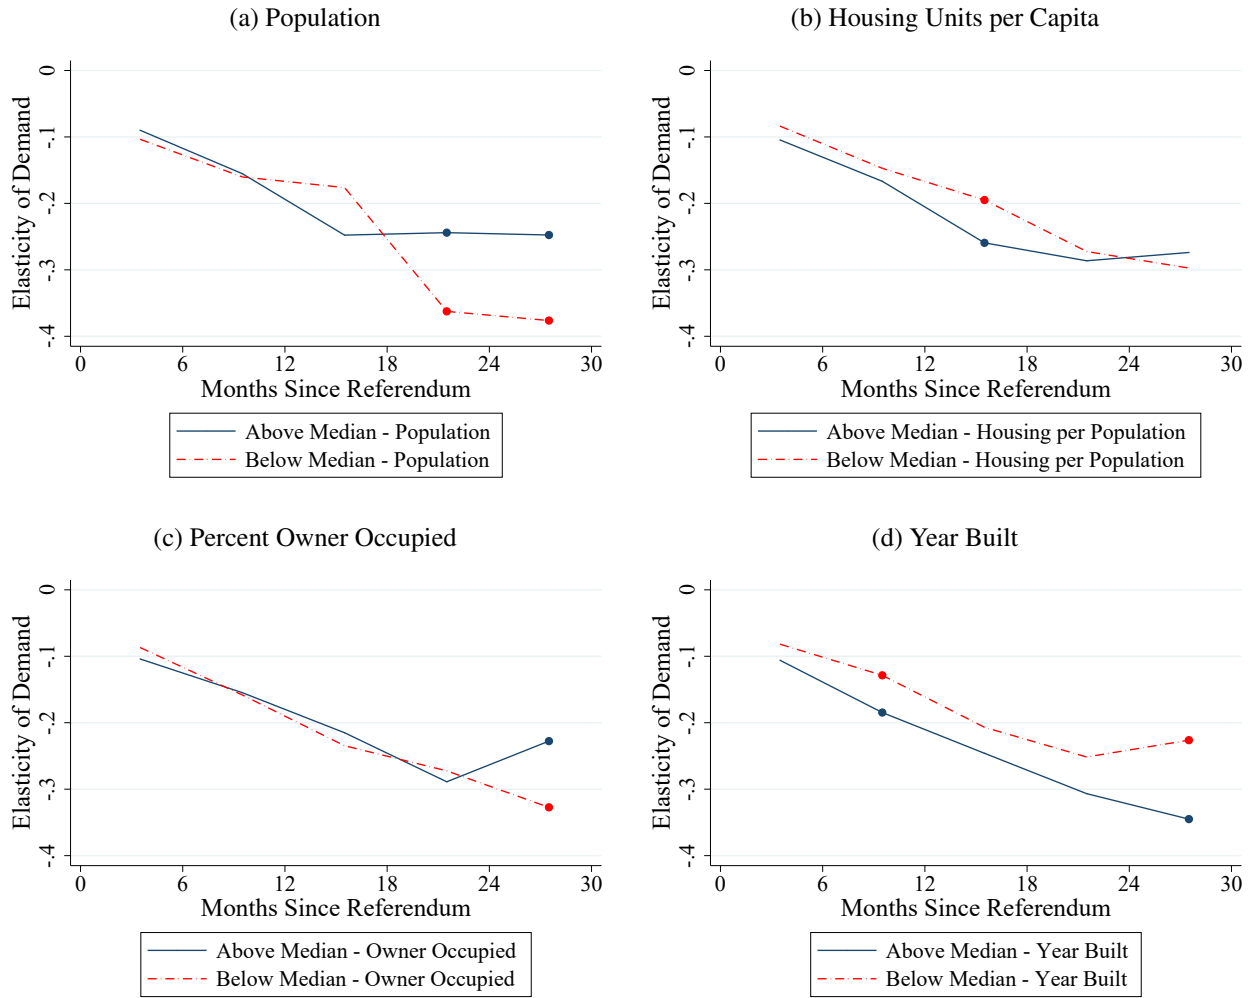
Notes: Outcome is the natural log of total electricity use. The first vertical dashed line indicates the date of the aggregation referendum. The second dashed line indicates the date of aggregation implementation. Regressions include month-by-year and community-by-month fixed effects. Standard errors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

Figure A.12: Estimated Price Elasticities, Communities that Passed Aggregation



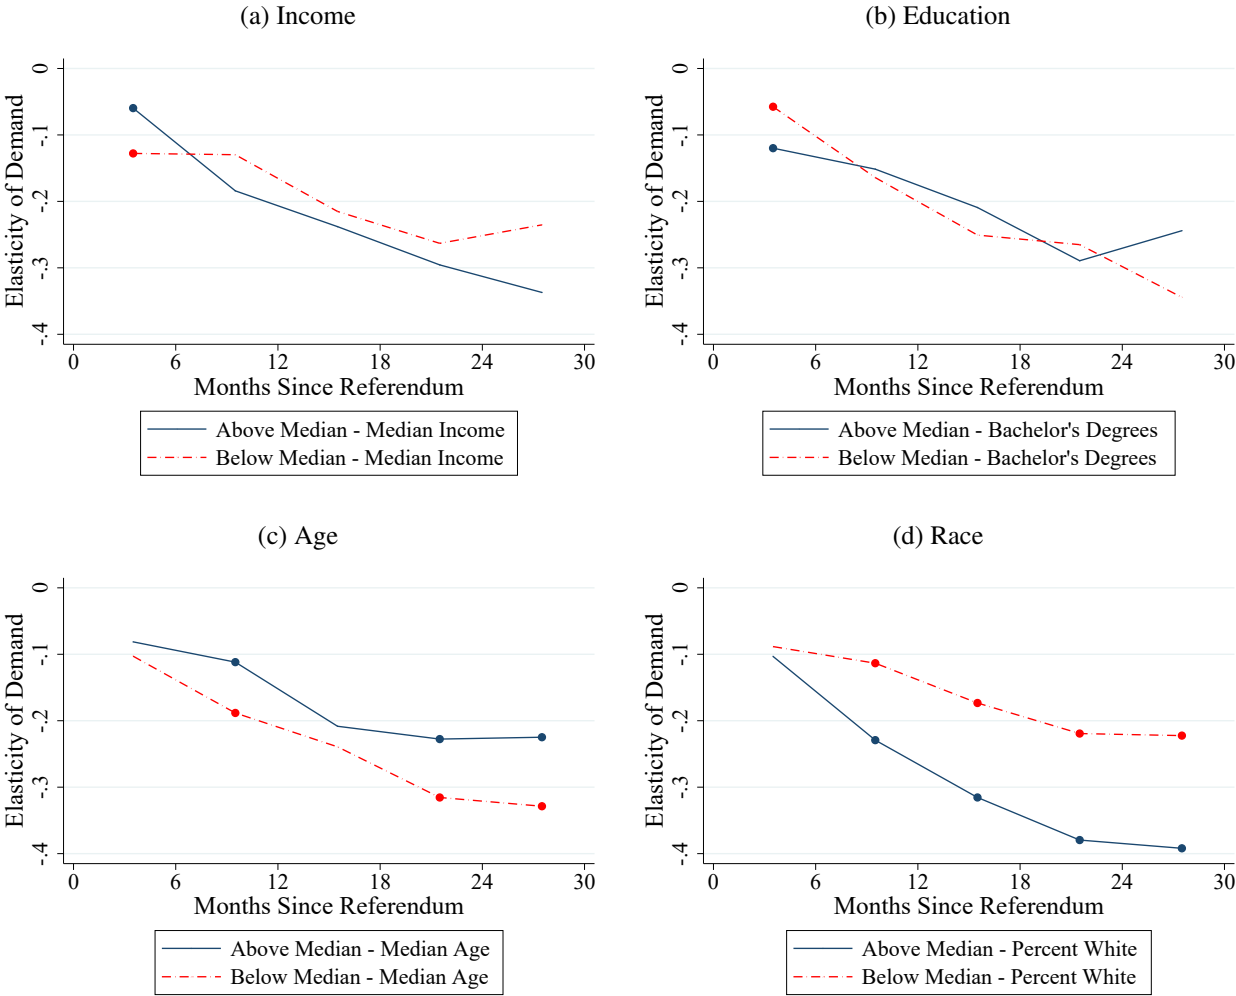
Notes: Sample includes only communities that passed aggregation at some point. Elasticities are calculated for each six-month period by regressing community-month changes in log usage on the observed change in log price. Confidence intervals are constructed by clustered bootstrap.

Figure A.13: 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 A.14: 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.