Energy Institute WP 294R

Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency

Severin Borenstein and James Bushnell

Revised July 2021

Revised version published in
American Economic Journal: Economic Policy
14(4), 80-110, 2022

Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency

By Severin Borenstein and James B. Bushnell*

July 2021

ABSTRACT: Economists favor pricing pollution in part so that consumers face the full social marginal cost (SMC) of goods and services. But even without valuing externalities, retail electricity prices typically exceed private marginal cost, due to a utility's need to cover average costs. Furthermore, due to costly storage, the marginal cost of electricity can fluctuate widely hour-to-hour, while retail prices do not. We show that residential electricity rates exceed average SMC in most of the US, but there is large variation, both geographically and temporally. This finding has important implications for pass-through of pollution costs, as well as for policies to promote dynamic pricing, alternative energy and reduced electricity consumption.

* Borenstein: Haas School of Business, University of California, Berkeley, CA 94720-1900, severinborenstein@berkeley.edu. Bushnell: Department of Economics, University of California, Davis, CA 95616, jbbushnell@ucdavis.edu. For outstanding research assistance and very helpful comments, we thank Stephen Jarvis, Sara Johns, and Jenya Kahn-Lang. For valuable comments and discussions, we thank Hunt Allcott, Joseph Cullen, Lucas Davis, Larry Goulder, Stephen Holland, Paul Joskow, Erin Mansur, Nick Muller, three referees, the editor, and participants in seminars at The Energy Institute at Haas-UC Berkeley, Berkeley-Stanford IOfest, the University of Arizona's Environmental & Energy Economics Workshop, the Asia-Pacific Industrial Organization Conference, the National Bureau of Economic Research program in Environmental & Energy Economics, and the 23rd Annual POWER conference. This research was funded in part by a gift from Google to the Energy Institute at Haas and a grant from the Sloan Foundation. An earlier version of this paper was circulated under the title "Are Residential Electricity Prices Too High or Too Low? Or Both?".

Economically efficient decision making by producers and consumers relies on prices accurately reflecting the short-run social marginal cost of supply. However, in utility industries that have traditionally been viewed as natural monopolies, the theoretical ideal of marginal cost pricing has been elusive in practice. One stream of research dating back to Ramsey (1927) has examined how price discrimination and non-linear tariffs can be used to mitigate deadweight loss while still allowing a utility with declining average cost to recover its total costs. Another research literature, growing out of Pigou's (1920) seminal work, has shown that environmental externalities lead firms to charge prices below social marginal cost. A third literature – starting with Boiteaux (1960) and Steiner (1957) – has emphasized that the highly time-varying cost of delivering electricity, due to its high cost of storage, suggests the need for dynamic pricing in order to reflect the constantly changing cost.

In this paper, we examine the relationship between marginal retail prices and the social marginal cost of supply in the electricity industry from 2014 to 2016. We focus on the most common residential electricity tariffs. In the $174 billion per year residential market, the efficiency implications of a gap between the marginal cost of service and the marginal price paid by consumers are growing more serious with the increasing availability of substitute technologies, such as rooftop solar photovoltaics and small-scale battery storage. These technologies make the demand of end-use consumers more price elastic, and therefore can magnify the deadweight loss from mis-pricing. Utilities around the world have expressed concern about the prospect of a "death spiral," in which reduced consumption leads to higher regulated prices, which in turn leads to more consumption decline (Costello and Hemphill 2014).

Retail pricing in electricity markets suffers from at least three distortions: (a) because neither buyers nor sellers bear the pollution costs of electricity generation, prices will tend to be below their optimal level, (b) because there are significant economies of scale in electricity distribution, and possibly other parts of the value chain, a linear price likely will need to exceed private marginal cost of the utility in order to recover its total costs, and (c) because electricity is not economically storable and demand fluctuates continuously, the private marginal cost changes constantly within a day, yet retail prices do not reflect those fluctuations. Importantly, these distortions do not all work in the same direction and can at times potentially offset one another. Research on the electricity industry and the policies that impact it, however, has tended to focus on each of these distortions in isolation. Since at least Lipsey and Lancaster (1956) and Buchanan (1969) it has been well understood in economics that markets with multiple distortions may not be improved by addressing one of the distortions in isolation.

In this paper, we take a step towards a holistic view by attempting to measure, at high frequency, the departure of residential electricity prices from the economic ideal of short-run social marginal cost (SRSMC). We then decompose the departure from SRSMC into the component caused by charging a price that differs from the average SRSMC and the component caused by charging a constant price that does not vary over short time periods, as SRSMC does. The analysis is primarily an exercise in measurement of various aspects of SRSMC and the marginal prices faced by customers. Some of these measures are available in public data, while some we estimate, because direct measures are not available.

We break the construction of price versus social marginal cost into three components: retail price, private marginal cost, and external marginal cost. Section II presents the residential electricity price data and our calculation of marginal electricity price. Section III discusses private marginal cost, for which we begin with hourly wholesale electricity price data, but then make adjustments to incorporate time-varying costs associated with local distribution. Section IV brings in externalities, estimating marginal externality costs for the marginal consumption of electricity by region. In section V, we bring the three measures together to analyze the deviation of price from SRSMC, then calculate and decompose the implied deadweight loss. In section VI we discuss several potential policy implications of our calculation. We conclude in section VII with a discussion of the broader relevance of our findings.

I. Related Literature

This paper relates to three strands of literature that have examined electricity pricing from different perspectives. The first concerns itself with the central challenge of natural monopoly pricing: minimizing deadweight loss while ensuring the recovery of average costs (Brown and Sibley 1986, Kahn 1988, Braeutigam 1989, Borenstein 2016). Here the main concern has been the inclusion of fixed and sunk costs in volumetric prices, potentially driving prices above marginal cost.[1] Various solutions have been proposed and at least partially implemented, including price discrimination with linear tariffs (Ramsey 1927, Boiteux 1960, Boiteux 1971), two-part pricing (Feldstein 1972, Littlechild 1975), and more sophisticated non-linear pricing (Wilson 1997, Laffont et al. 1998). Yet, despite a plethora of complex rate structures in use, there is a general perception that utility rates do not closely approximate (private) marginal costs (Friedman 1991, Puller and West 2013). Davis and Muehlegger (2010) estimate marginal tariff rates for natural gas utilities and find that they do not adjust fully to fluctuations in wholesale gas supply costs, while Borenstein and Davis (2012) examine the equity effects of these departures from marginal cost pricing of natural gas and discuss the potential equity and efficiency effects of changing fixed charges. We are not aware of any comprehensive effort to measure the departure from marginal cost of retail electricity prices.

A second literature on electricity pricing is concerned with the variation of costs over time, particularly those driven by scarcity or capacity constraints. Early theory focused on forms of peak-load, or capacity, pricing that could at least

[1] Low inframarginal costs can drive average cost below marginal cost in some circumstances, most notably utilities that have access to limited quantities of cheap power, such as from federal hydroelectric generation.

partially capture scarcity effects in otherwise static tariff structures (Boiteux 1960, Steiner 1957, Joskow 1976, Oren et al. 1985, Crew and Kleindorfer 1976). The advent of advanced metering technology made feasible the prospect of dynamic electricity pricing (Borenstein 2005, Joskow and Wolfram 2012) that could capture scarcity costs through frequently varying linear prices. However, despite a growing literature on its practical effectiveness (Jessoe and Rapson 2014), dynamic pricing is still quite rare. As we describe below, only 4% of residential US customers are on a time-varying price, and the bulk of those customers are on static time-of-use prices. The lack of dynamic retail pricing has been widely cited as a source of inefficiency in the electricity industry (Borenstein and Holland 2005, Borenstein 2005, Joskow and Wolfram 2012, Puller and West 2013).

The most recently active strand of literature on the efficiency of electricity prices concerns their relationship with the external costs of electricity production and consumption (Cullen 2013, Graff Zivin et al. 2014, Novan 2015, Holland et al. 2016, Callaway et al. 2018). The environmental impacts of electricity supply, particularly with respect to climate change, are significant and have been the focus of policy activity for at least two decades. Environmental economists have generally advocated for the pricing of external costs, through either Pigouvian taxation or cap-and-trade systems, in this and other industries. However, alternative approaches, such as subsidies for clean energy through either tax credits or performance standards, and non-market interventions relating to energy efficiency have been more common in practice than the pricing of externalities.[2] These latter programs have been criticized by economists on several grounds.

Several papers have addressed the optimality of environmental policies with respect to consumer incentives. These studies have raised concerns about policies that limit the pass-through of externality costs. For example, the impact of intensity standards for limiting carbon emissions (Bushnell et al. 2017), the use of output-based allocation of allowances in cap-and-trade systems (Fowlie 2011, Fischer and Fox 2012), and energy efficiency interventions (Allcott and Greenstone 2017). A common theme is that many "green" policies tend to promote overconsumption as they fail to properly reflect marginal environmental damages in electricity costs (Borenstein 2012). However, with the exception of Allcott and Greenstone (2017), these papers address the design of optimal externality policies from an underlying assumption that retail prices accurately reflect private (but not social) marginal cost. To the extent that pre-existing distortions to retail prices, due to natural monopoly pricing for example, have already distorted retail prices, the optimal environmental policy can look very different from the one applied in a system with prices reflecting private marginal costs.

[2]For example, the Obama-era EPA regulatory initiative known as the Clean Power Plan offered States several options for compliance, including an intensity standard or direct subsidies of zero-carbon generation sources, as alternatives to carbon pricing (Fowlie et al. 2014).

II. Residential Electricity Pricing

The challenge in constructing data on residential electricity pricing is to accurately characterize the marginal price that a customer faces. While data on aggregate revenues and quantity sales to residential customers by utility are available, those data alone only allow inference about the average price paid by residential customers. Optimizing consumers, however, would respond to the marginal price of electricity, not the average price. Thus, we must adjust the analysis in order to get a more accurate measure of marginal price.

Our primary source of utility sales data is the Energy Information Administration's Form EIA-861 survey (Energy Information Administration 2017a). The EIA-861 is an annual survey of electric utilities that covers many aspects of their commercial activities.[3] The EIA-861 data include for every utility-state annual total revenues from residential customers, total number of customers, and total kilowatt-hours (kWh) sold. Dividing total revenues by total kWh yields an average price.[4]

However, many utilities have monthly fixed charges. In order to calculate the marginal price, we remove the fixed charges. The utility fixed charges for residential customers come from the National Renewable Energy Laboratory's Utility Rate Database (URDB) (National Renewable Energy Laboratory 2017b). The URDB is described in more detail in the appendix. It includes many residential rates for each utility. For each utility we chose what appeared to be the primary or basic rate (the process of determining this rate is described in the appendix) and took the fixed charge from that rate. We used this fixed charge to approximate fixed revenues – total customers multiplied by fixed charge – and subtracted that amount from the total residential revenues. We divided the remainder by kWh sold to get the average variable rate, which we take as our measure of marginal price.

In most of the country, the same company is responsible for procuring or producing electricity on behalf of residential customers and physically delivering it through local distribution lines. In some parts of the country, however, the electricity sector has been restructured such that customers can choose their "retail provider" from among many companies that buy power wholesale and sell to the customer at retail. About 32% of residential customers in the US are eligible

[3]To be precise, our sample contains 2,104 utility-state combinations. Utilities report their operations separately by state to the EIA. For each utility-state combination, we calculate each measure separately for each year and then for the maps we take the average across the years for which the utility-state is in the dataset (which is 2014, 2015, and 2016 for almost all utility-states). See the appendix for further details. For simplicity, we refer to the unit of observation as a utility-state. A smaller number of major utilities are surveyed monthly, covering about two-thirds of the household customers in the annual survey (Energy Information Administration 2017b). In the appendix, we discuss a robustness check that we carry out using the monthly survey. We find very small seasonal changes in retail rates.

[4]Throughout this paper, we use the standard measurement of electricity quantity, kilowatt-hours. The average US household uses about 900 kWh per month. For further context, 1 kWh is enough electricity to run a residential central air conditioning system continuously for 10-15 minutes, a hairdryer or electric kettle for about 40 minutes, or a microwave oven for an hour.

to choose their retail provider, and just under half of those customers choose a retailer that is different from their local distribution company. Full details of how we handle these cases can be found in the appendix.

Removing the fixed component of customers' bills still does not fully capture marginal rates if those rates vary with the level of consumption, such as from increasing-block or decreasing-block pricing – under which marginal price rises or falls in steps as a household's consumption increases. Thus, some customers of a given utility are likely to have a higher marginal rate, and others a lower marginal rate, than the one we use. Based on the 1743 utilities with rates in the URDB, about 58% of residential customers are served by a utility for which it appears that the marginal price in the primary residential tariff varies with consumption, of which about 37% face increasing-block pricing and about 21% face decreasing block pricing.[5]

Similarly, we do not capture differences in static rates across customers of a utility. Many utilities, for instance, have lower rates for low-income households. But it could also occur if a utility charges rates that vary by geographic region. It is worth noting, however, that the failure to reflect variations in marginal rates across customers that are not based on marginal cost is very likely to lead to understated estimates of the deadweight loss associated with residential rates. This is because deadweight loss increases more than proportionally with the difference between price and marginal cost. Thus, for linear pricing, if all customers have the same demand elasticity, deadweight loss is minimized by charging all customers the same linear price.

In all cases, we also have assumed that the primary residential rate had no time-varying component, including no time-of-use variation, no critical peak pricing, no demand charges, and no real-time pricing. The prevalence of these kinds of tariffs is very low among residential customers. During 2014-2016 about 4% of customers were on some form of time-varying pricing, and just under 6% of customers were part of some form of demand response rebate program.[6]

Our final dataset on residential electricity pricing covers an average of 128.2 million residential customers during 2014-2016, with average annual total sales of 1.384 trillion kWhs and revenues of $174 billion. After incorporating our estimates of fixed charges we were able to calculate the average variable per-kWh price faced by just over 93% of residential customers and kWh sales.

[5] The share of *quantity* sold on non-linear pricing is somewhat smaller, as the retail providers utilizing increasing-block pricing serve smaller average residential demand per customer. Overall, providers serving larger numbers of customers are more likely to use increasing-block pricing. Of the 1743 retail electricity providers in our URDB sample, about 39% utilize non-linear marginal pricing, with about 15% using increasing-block pricing and about 24% using decreasing block pricing in their primary residential rates.

[6] The EIA-861 data that are the source of these figures do not allow one to calculate the overlap between these two sets of customers, but it is probably significant. Furthermore, a very large share of the customers on time-varying pricing are on simple peak/off-peak rates with fixed time periods and fairly small differentials between peak and off-peak.

A. Is marginal price the correct measure?

A number of papers, most prominently Ito (2014), have challenged the belief that electricity consumers respond strictly to marginal price.[7] Ito finds that in the context of steeply increasing-block electricity pricing at two large utilities in California, consumers are more accurately characterized as responding to the average price they face, rather than the marginal price.

These analyses, however, do not address the extent to which consumers are able to separate recurring fixed charges from volume-based charges.[8] Understanding and distinguishing a monthly fixed charge from volumetric pricing seems likely to be much less difficult than diagnosing which step of an increasing-block marginal price schedule the household is likely to end up on at the end of the month. Ito and Zhang (2020) is the only work of which we are aware that addresses the former question. They find strong evidence that consumers respond to changes in marginal prices apart from changes in fixed charges.[9]

Luckily, for our analysis, the three large utilities in California that have steep increasing-block electricity price schedules, where the steps differ by more than 4 cents per kWh, are outliers in the US as a whole. Out of the 1743 utilities we study that are in the URDB, there are 673 with non-constant marginal price. Among those 673, the median absolute difference between the lowest and highest tier across all US utilities was 1.9 cents per kWh, with 75% of the rates showing a difference of less than 3.7 cents per kWh.[10]

To the extent that customers respond to the average variable price they face in an increasing-block price schedule – a possible third "wrong" – they would perceive a lower incremental cost than the actual marginal price. This would somewhat reduce their response to the very high marginal prices in California and a few other areas, however perceived prices in most of those areas with significant increasing-block pricing would still substantially exceed social marginal cost. Decreasing-block price schedules would have the opposite effect, raising the perceived incremental cost of consumption above the true marginal price. Nonetheless, the existence of marginal pricing that changes with consumption quantity should be recognized in interpreting our results.

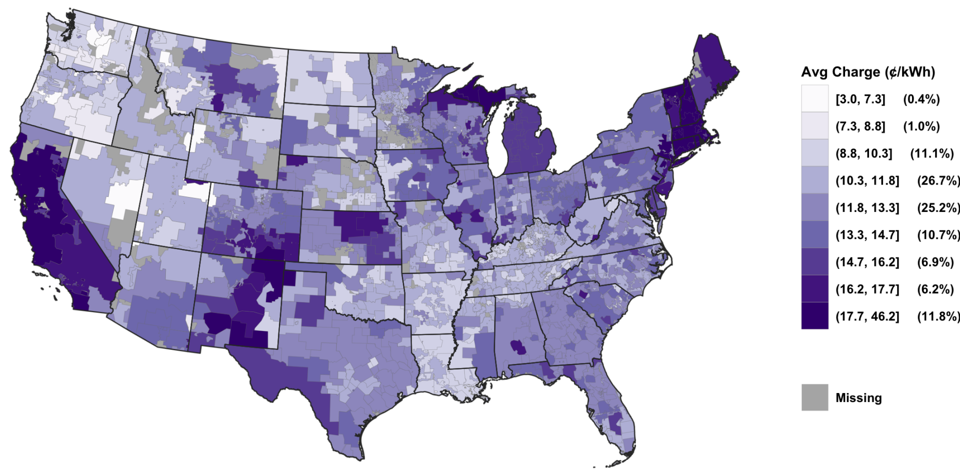

Figure 1: Average Price per kWh

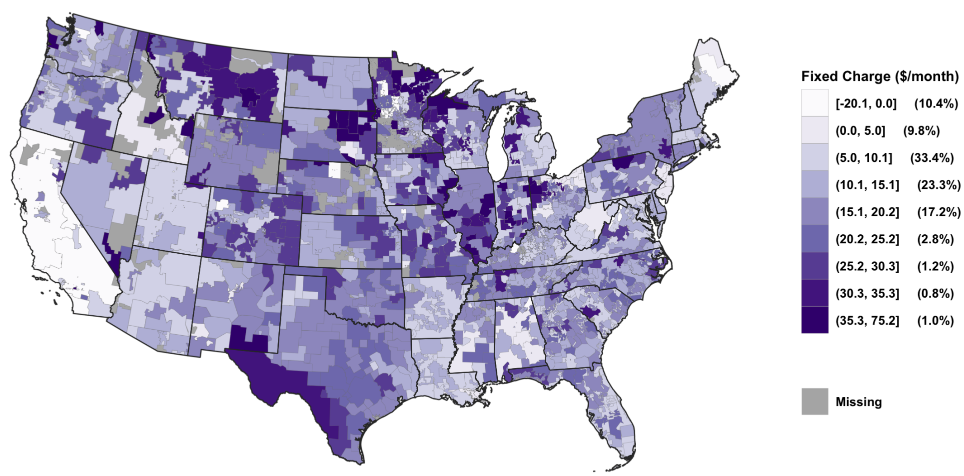

Figure 2: Fixed Monthly Charge

B. Residential Electricity Pricing Results

We present many results in this analysis graphically at the utility-state level for the 48 contiguous US states. To do this, we use utility service territory shape files from Homeland Infrastructure Foundation-Level Data (HIFLD).[11] In all of the maps, the percentages in parentheses in the legend are the percent of residential customers in each category, and areas of the map with no data are represented by a dark gray shade, such as in part of the Northern border area of Montana. Values represented for each utility-state are the sales-weighted average across the three years of data.

Figure 1 presents the average price per kilowatt-hour (kWh), over 2014-2016, by utility-state. It shows, for instance, that California has among the highest average prices per kilowatt hour for residential customers, but that the very highest prices are in the Northeast. The lowest prices can be found in much of the Northwest and the South. It also shows that even in fairly high-priced states like California and New York, there are some areas with substantially lower prices.

Figure 2 presents monthly fixed charges as discussed above. Much of California has zero or slightly negative fixed charges – which occurs because of a semi-annual "climate rebate" that each residential customer gets as part of the state's cap and trade program – while some utilities in the center of the country have fixed charges of $30 per month or higher.

Figure 3 shows the results from adjusting the average price for the monthly fixed charges to get an average variable price. Nationally, the adjustment for monthly fixed charges lowers the marginal price by about 1 cent per kWh. We would expect this to be an accurate indicator of the marginal price that consumers face if the utility uses a simple two-part tariff. For those utilities that utilize increasing-block or decreasing-block pricing, as discussed earlier, this captures the average variable price across customers.[12] The average variable prices illustrated in this figure are used in our calculation of the gap between marginal price and social marginal cost.

Table 1 presents the quantity-weighted summary statistics on average price,

[7] See also Shin (1985) and Borenstein (2009).

[8] The customers in Ito's sample faced increasing-block pricing, but no fixed charge.

[9] If customers are not able to distinguish fixed from marginal prices and make their decisions based on average price, this would raise perceived prices most for customers in the middle of the country, as shown in figure 2, where we find that price is below social marginal cost. However, it would also raise prices in the Northeast, where we generally find prices are already substantially above social marginal cost. Though fixed charges are common, they are not typically a large share of the total bill, averaging about 10% as suggested by the difference between average price and variable price in table 1.

[10] Furthermore, even in California the variation in marginal price across the steps has shrunk significantly in the last decade for the vast majority of households, from a ratio of more than 3 to 1 in 2014, to a ratio of less than 1.4 to 1 in 2017. There remains a higher "superuser" rate that applies for usage over 400% of the baseline quantity, but that is relevant for just a few percent of households.

[11] https://hifld-geoplatform.opendata.arcgis.com/datasets/electric-retail-service-territories. The map captures more than 99% of all residential customers, but omits a few small municipal and coop utilities.

[12] How closely this reflects the average of the marginal prices faced by customers depends on the distribution of customers across the tiers of the block pricing. See Borenstein (2009) and Ito (2014) for further discussion.

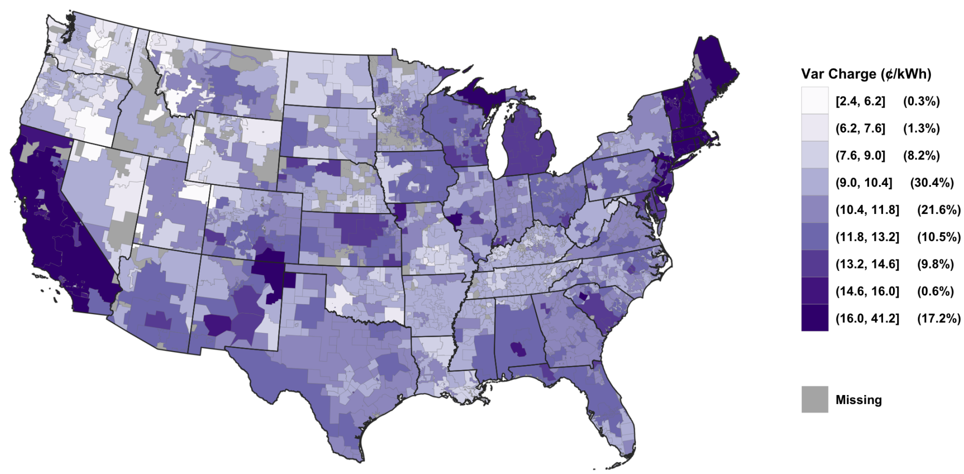

Figure 3: Marginal Price per kWh

	Mean	StDv	Min	P10	P90	Max
Retail Fixed Charge (\$/month)	10.78	7.65	-26.11	2.53	20.00	75.53
Retail Variable Price (¢/kWh)	11.49	3.07	2.36	8.79	16.29	48.22
Retail Average Price (¢/kWh)	12.61	3.01	2.96	9.83	16.65	53.31

N=6215 (utility-state-years). Statistics are sales-weighted

Table 1: Summary Statistics of Residential Rates

fixed charge and average variable charge across the 6,215 utility-state-years in the entire sample. For the maps, we calculate the statistics separately for each utility-state-year it is in the data set, and then take the quantity-weighted average of those years.

III. Private Marginal Costs

Provided that wholesale electricity markets are competitive,[13] the primary component of the private marginal cost of supplying electricity is captured in the wholesale price. We collected wholesale prices from regions that are part of Independent System Operator (ISO) control areas. ISOs calculate and report locational marginal prices (LMPs), which reflect the marginal cost of electricity generation plus high-voltage transmission congestion and line losses (due to heat dissipation of energy during transport). These prices literally represent the derivative of total system production cost with respect to a change in consumption at a given location (node), accounting for all relevant transmission, operating reserve, and unit-level operating constraints considered by the system operators. Electricity prices are extremely volatile over time and geography due to a combination of volatile demand, a lack of economic storage, and frequently binding limits on low-cost production and/or transmission capacity. Technical constraints also limit the flexibility of some production resources and the ability to freely dispose of power. In addition, wind generation receives tax subsidies that are based on quantity produced, effectively lowering their short-run private marginal cost to below zero. These factors combined to occasionally result in negative prices in regions with inflexible supply that is unable to export its surplus.

Some parts of the country, particularly the Southeast, have large areas that are not covered by ISOs. In those areas, we collected data that grid operators are required to file to the Federal Energy Regulatory Commission as part of the FERC Form-714 survey (Federal Energy Regulatory Commission 2017). This survey includes a requirement to report the "system lambda", which is the engineering calculation of the shadow cost of changing production by one unit. Thus, ideally, it would correspond with the marginal cost, as reflected by a competitive market price, in the ISOs. In practice, however, much of the Form-714 data are obviously unreliable, exhibiting many consecutive hours of identical values and zero values where they are not plausible. As described in the appendix, we incorporate data for those areas where the Form-714 data seem to be most reliable. Nonetheless, the Form-714 data may understate the true private marginal cost, both because system lambda likely does not fully incorporate marginal transmission losses and congestion costs and because system lambdas may not fully incorporate scarcity rents in constrained hours.[14]

[13] Each organized electricity market operator estimates and reports a price-cost margin based upon its estimate of the marginal cost of supply. In the last decade these reports indicate very low to negative margins in most electricity markets (Federal Energy Regulatory Commission 2016).

[14] The wholesale prices in areas with ISOs are also imperfect measures, because they likely incorporate

Region	Location	Mean	Min	P10	P90	Max
CA	California	33.86	-150.00	17.47	52.38	1658.94
FRCC	Florida	25.87	-32.69	15.91	37.28	1043.18
MRO	Upper Midwest	25.94	-150.00	13.42	38.91	1858.24
NPCC	Northeast	40.95	-150.00	13.27	76.17	1446.06
RFC	Great Lakes	34.90	-150.00	17.95	52.66	1938.75
SERC	Southeast	30.63	-150.00	17.09	41.75	2726.81
SPP	Oklahoma/TX	27.11	-150.00	14.97	38.43	4655.87
TRE	Texas	28.24	-110.47	15.20	40.15	4708.40
WECC	Non-CA West	30.85	-150.00	15.28	48.07	2770.26

Weighted by Retail Sales.

Table 2: Wholesale Power Prices by NERC Region ($/MWh)

We calculate private marginal cost based for each ZIP Code based on LMP prices and/or system lambda values that are closest to the ZIP Code, which allows those costs to include transmission losses and transmission congestion costs, and then aggregate up to utility levels. Full details of this calculation can be found in the appendix. Table 2 summarizes the wholesale power cost, weighted by hourly consumption, by regions designated by the North American Electricity Reliability Council (NERC), which correspond to fairly integrated grids. The one exception is that California is part of the Western Electricity Coordinating Council (WECC), but operates a wholesale electricity market that was not well integrated with the rest of the WECC during these years. The data allow us to break out California, so we do.[15] As we discuss later, average prices are below levels generally considered sufficient to cover long-run average cost of a modern combine-cycle natural gas power plant, even at the very low gas prices that existed at this time. These averages, however, mask significant heterogeneity in prices both regionally and over time. When wholesale markets have experienced either scarcity conditions or high natural gas prices, wholesale prices have risen to extremely high levels. Each of our ISO-based markets experienced prices in individual hours well above $1000/MWh. This supports the viewpoint that market prices are capable of reflecting marginal costs that include significant scarcity rents when applicable, and that the relatively low average prices are reflective of a lack of scarcity, rather

some market power in some hours, although analysis by oversight divisions suggests very modest if any market power averaged over all periods (Bushnell et al. 2017). Unfortunately, comparing system lambdas to wholesale prices where they exist does not help to reveal the magnitude of these biases, because utilities in these areas typically report the market price as the system lambda for their region.

[15]We Winsorize hourly prices at -$150, because that is the minimum bid allowed in most ISO markets. A few observations of much lower prices appear in the data, but it is unclear whether they are data errors. Including all prices has a very small effect on average price calculations and the deadweight loss from price deviating from average SMC. But for a few utilities, extremely negative prices cause larger deadweight loss calculations from hourly SMC variation. We also did all calculations with hourly prices Winsorized at $0, which has very little effect on any of the calculations compared to a -$150 cutoff.

than a systemic suppression of wholesale price below short-run marginal cost.

A. Distribution Losses

The private marginal costs calculated based on wholesale prices do not include the cost of local electricity distribution on low-voltage lines. The primary marginal cost of distribution is electricity losses. Losses from low-voltage distribution lines fall into two categories: a smaller share is attributed to "no-load" losses that occur in transformers, and a larger component is "resistive" losses that are a function of the flow on the line. No-load losses are fairly constant for a utility and vary across utilities as a function of the size of their systems. Resistive losses change constantly scaling with the square of the flow on a line.[16] On average, around 25% of distribution losses are no-load with the remainder attributed to resistive losses.

A range of factors affect the magnitude of losses, including the distance electricity must be carried (approximately the inverse of geographic demand density), the density of load on circuits, the use of equipment to optimize voltage, and the volatility of demand. Demand volatility increases losses for a given average demand level due to the quadratic relationship between flow and resistive losses. Many of these factors are likely to differ between residential customers and commercial or industrial customers. Importantly, many industrial and some commercial customers take power from the distribution system at higher voltages than residential customers, which can substantially reduce the level of line losses.

Unfortunately, the only systematic data available on distribution line losses are reported on an annual basis by utility in the EIA-861, with no breakdown by class of customers, or by hour. As we describe in the appendix, we approximate hourly losses for service to residential customers by first estimating an OLS regression equation for annual average losses, controlling for the factors mentioned in the previous paragraph. Using these parameter estimates, for each utility, we then predict the average losses for service to residential customers.[17] The average loss rate is then converted to a time-varying hourly marginal loss rate on the assumption that 25% of losses are independent of flow (not marginal) and the remaining 75% are marginal losses allocated across hours in proportion to the square of load. The details are presented in the appendix.

To do this, however, we need data on the pattern of hourly consumption by residential customers, which don't exist for most utilities. FERC Form-714 provides hourly data on total consumption of all customers from groups of utilities, known as planning areas. We use that load profile, scaled by the share of total demand that comes from residential customers, to approximate the residential

[16]Lazar and Baldwin (1997) have a very accessible discussion of distribution line losses.

[17]As discussed in the appendix, the vast majority of our predicted average distribution losses for serving residential customers fall between 4% and 8%, with a mean of around 6.5%. Thus, the effect of marginal distribution losses on social marginal cost is overall about one-third larger than it would be if we had simply assumed that every utility had the national average loss rate of about 5.0%, but this approach also allows us to capture variation across utilities.

demand in each hour. This is not ideal. The alternative, however, is to use data produced with an engineering model of residential energy use patterns, which also is highly imperfect. As discussed in the appendix, we conduct a sensitivity using engineering-model based data and it does not materially affect our results.

Distribution losses turn out to be significant in the overall analysis. Figure 4a presents the spread of average annual distribution losses from residential customers for the utilities in our analysis. Table 3 shows that on a sales-weighted basis the estimated average distribution loss rate is 6.2%. Furthermore, because the externalities associated with electricity consumption take place upstream from the distribution losses, the loss rate scales up both the private marginal cost and the external marginal cost. After assuming that 25% of losses are non-marginal and the other 75% vary with the square of load, figure 4b presents the distribution of marginal hourly distribution losses from residential service that we estimate. These average about 8.9%, but vary greatly hourly with load.

	Mean	StDv	Min	P10	P90	Max
Avg. Total Losses (%)	4.90	1.33	0.55	3.36	6.55	10.44
Avg. Res. Dist. Losses (%)	6.20	1.26	2.58	4.85	7.85	12.58
Marg. Res. Dist. Losses (%)	8.87	1.82	3.75	6.95	11.15	18.23

N=6215 (utility-state-years). Statistics are sales-weighted

Table 3: Summary Statistics of Distribution Losses

B. Other private cost considerations

The energy costs captured by the LMP and system lambda data used in this analysis constitute the great majority of the wholesale electricity costs that must be covered by customers over the year. The remainder is made up of capacity costs, ancillary services costs and other side payments known as "uplift". Across the seven ISOs, energy costs comprised between 74% and 98% of the total wholesale cost of electricity in 2015, as shown in table 4.[18] More detail on the source and interpretation of these costs is provided in the appendix.

Capacity payments are regulatory-mandated payments from buyers to suppliers based upon their installed effective production capacity. We do not include capacity costs in our calculation of short-run private marginal cost. In some markets, such as ERCOT or SPP, there are no explicit capacity payments. In other markets that do have capacity requirements, the standards have to be adjusted in the medium or long run in response to variation in demand. These costs can sometimes be substantial. In 2015 capacity costs comprised between 4% and 22%

[18]Our geographic results are grouped by NERC region.The geographic footprints of ISOs overlap those of the NERC regions, so it is difficult to make one-to-one mappings from and ISO to a NERC region.

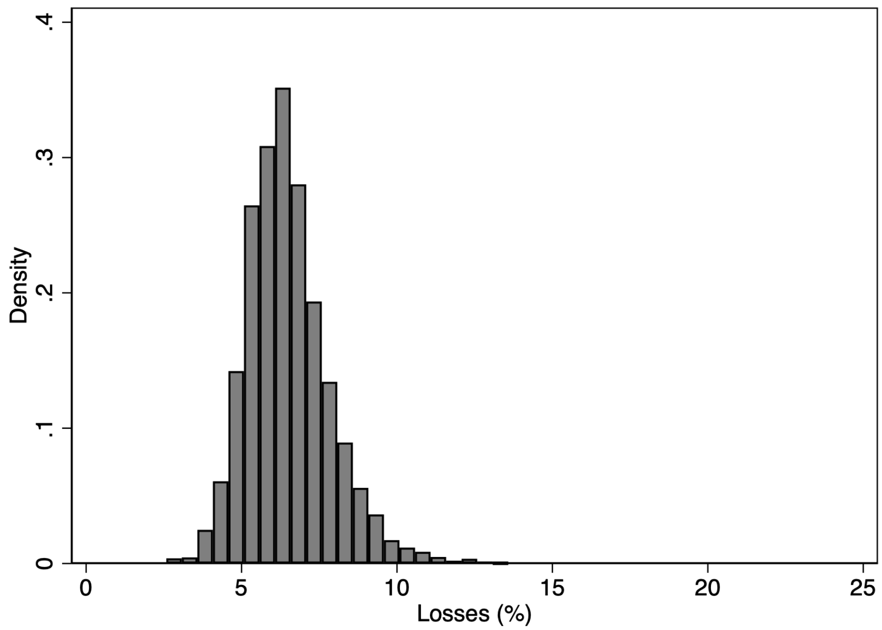

(a) Average annual residential distribution losses

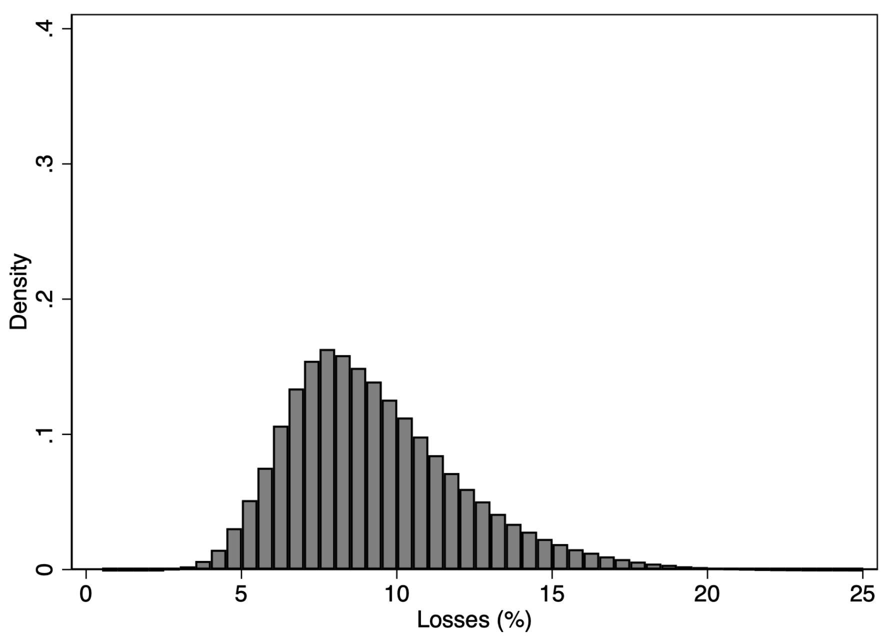

(b) Marginal hourly residential distribution losses

Figure 4: Estimates of residential distribution losses

	Location	Energy	Capacity	Ancillary	Uplift
CAISO	California	89%	9%	1%	1%
PJM	Mid-Atlantic	74%	23%	2%	1%
ISO-NE	New England	81%	15%	3%	1%
NYISO	New York	74%	22%	3%	1%
ERCOT	Texas	92%	-	4%	4%
SPP	Great Plains	98%	-	1%	1%
MISO	Midwest	95%	4%	0%	1%

Note: Percentages may not sum to 100 due to rounding

Table 4: Estimates of the composition of total wholesale costs by ISO

of the total wholesale cost of electricity at the five ISOs that make these payments. Importantly, these revenues move inversely to energy market revenues. When energy prices, reflecting short-run marginal costs, are high, capacity payments to generation implicitly or explicitly adjust to reflect the fact that resources are recovering more of their fixed costs through energy prices. In other words, capacity payments partially smooth the difference between long-run and short-run marginal cost.

The link between incremental consumption in a given hour and the capacity requirement is complex. However, conditioned upon the capacity at any point in time, the wholesale energy market price, adjusted for distribution losses, should reflect the true marginal resource cost of delivering one more kWh. Thus, from a strict economic efficiency vantage, longer-run investments triggered by current demand would not be a short-run marginal cost.[19]

We also do not incorporate short-run operating reserve, or "ancillary service", costs into our marginal cost calculation. LMPs are calculated in a process that simultaneously optimizes for meeting demand and reserve requirements. The LMP therefore already reflects the shadow costs imposed through reserve requirements. The primary marginal impact of reserves is reflected in the energy prices or system lambda values used to reflect cost. This is because most reserves operate as stand-by resources and do not incur marginal cost unless a contingency event occurs. The main cost impact of an expansion of reserves arises when lower cost units are held back to provide reserves, while more expensive units are deployed to supply energy in their place. However this effect is captured in the marginal energy price when the more expensive units set those prices. In any event, these costs are relatively small, even in aggregate. In 2015 ancillary service costs at the

[19]One complication to this interpretation of short-run marginal cost arises when there is scarcity of supply. When electricity systems experience short-term violations of operating constraints, such as unit ramping or transmission flow constraints, prices include penalty values to reflect the cost of the scarcity of appropriate supply. To the extent these values do not reflect the true underlying value of electricity to end-users, they are rough approximations of the short-run marginal costs in these periods. There were relatively few such periods during 2014-2016.

seven ISOs comprised between less than 1% and 4% of the total wholesale cost of electricity.

Finally, some non-convex incremental costs, such as generator "start-up" costs, that are incurred to supply energy are at times not captured in the energy price and are instead paid as out-of-market side-payments, or "uplift," payments to specific units. The funds for these side-payments are usually raised through a small fee applied to all demand.[20] We do not currently adjust our costs for these considerations. Again though, these costs are very small. In 2015 "uplift" payments range from less than 1% to 4% of the total wholesale cost of electricity.

Including all of the non-energy wholesale electricity costs would have a modest effect on the average wholesale price of electricity, and therefore on the gap between the marginal retail price and the average social marginal cost. It could, however, have a significant effect on the SMC during peak hours if reserve costs were considered marginal and were attributed entirely to the highest-demand hours. In that case, SMC would be more volatile than our analysis suggests and the deadweight loss of static pricing would be greater.

C. *Private Marginal Cost Results*

Figure 5 presents the private marginal cost calculations. Summary statistics on private marginal cost are presented in table 5 in the next section along with external marginal costs and total social marginal cost. As discussed above, these private marginal cost levels are below what many consider to be the long-run average cost of power supply. In part, that reflects the fact that much of the country had excess capacity during 2014 to 2016 due to a combination of mistakes or bad luck in planning and policies of carrying large quantities of excess capacity. Consistent with such policies, this also reflects the fact that in most deregulated markets, power plant owners receive revenues from capacity payments as well as energy payments. Regardless of whether such capacity payments are appropriate, they do not reflect marginal cost and therefore can distort consumption when reflected in marginal consumer prices.

There is also significant variation over time in these levels. Figure 6 summarizes the monthly average wholesale private marginal cost by NERC region.[21] During winter periods of high demand and gas prices, such as the 2014 polar vortex, prices rose to extremely high levels, raising monthly averages above $0.15/kWh in parts of the Mid-Atlantic (RFC) and northeast (NPCC) regions. This pattern reflects, on a longer time-scale, many of the issues raised in discussions of short-run dynamic electricity pricing. Marginal costs in power markets are quite volatile, even on a monthly or annual basis.[22] The electricity industry has experienced

[20] These costs are also not considered to be technically marginal, which is why they are not incorporated into market prices.

[21] We have combined California with the rest of the WECC and Florida (FPCC) with the neighboring SERC region in order to make the figure more readable. The regions that we combined have very similar price patterns.

[22] In contrast to other energy markets, such as oil, the level and volatility of wholesale electricity prices

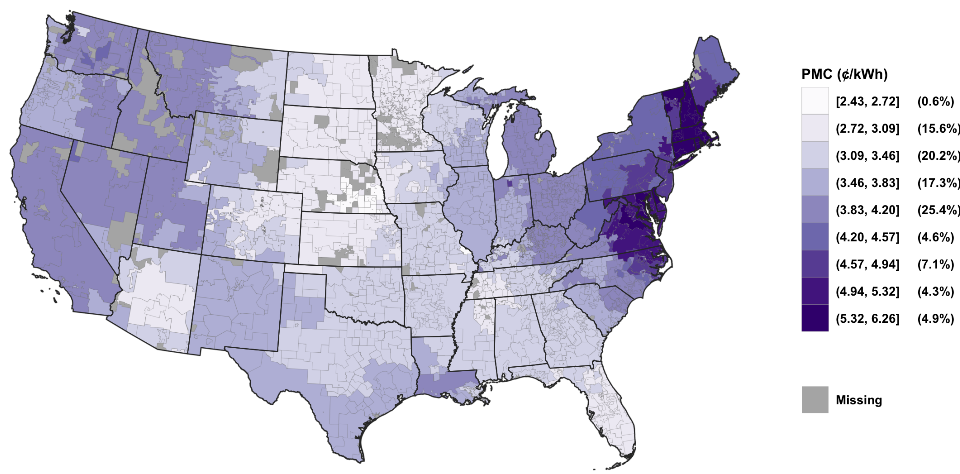

Figure 5: Average Private Marginal Cost per kWh

repeated cycles where marginal costs move dramatically relative to average cost (Borenstein and Bushnell 2015), while retail prices, which are strongly linked to historical average cost, are significantly more rigid.

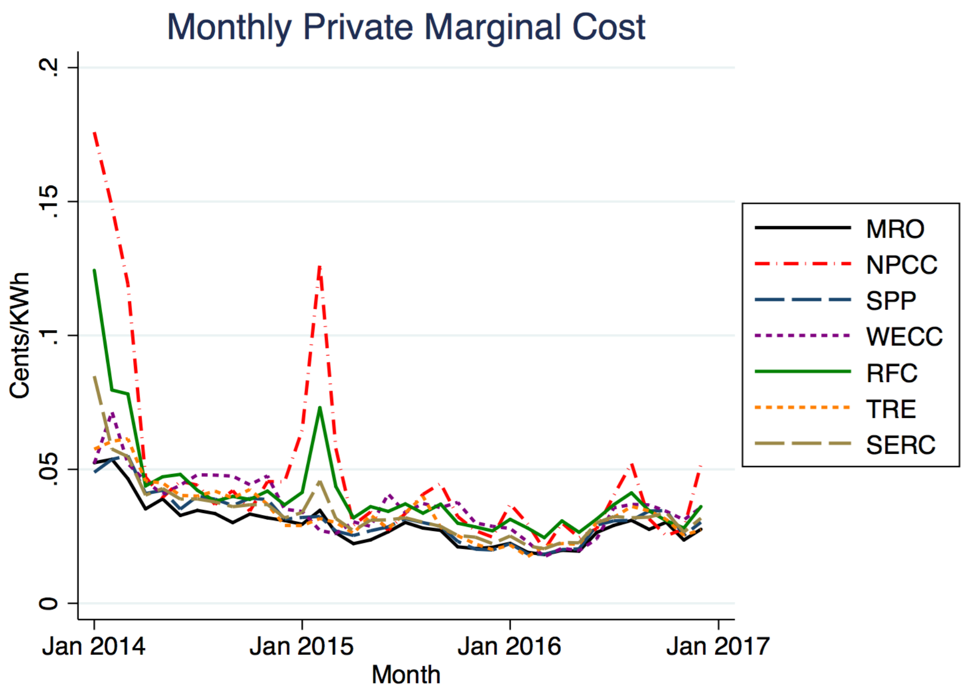

Figure 6: Monthly Private Marginal Cost by NERC Region

Wholesale prices (and implied private marginal costs) that remain for long periods below levels necessary to cover long-run average cost are certainly a concern for generators and policymakers. However, even if measured accurately, such a shortfall does not have direct bearing on our analysis of the efficiency of residential retail marginal prices and their deviation from SRSMC. Economic theory dictates that if short-run marginal costs are indeed quite low, then efficient pricing should reflect that, even if such prices are not sufficient to cover average cost.[23] Furthermore, even if policymakers believe that additional revenue must

can differ substantially between regions because transmission losses, limits on transmission capacity, and high storage costs limit the ability to arbitrage wholesale electricity prices and costs between regions and over time. In the case of Texas, a higher cap on spot prices also leads to a more extreme distribution of prices.

[23] And, conversely, if the marginal generation costs are quite high, yielding very high profits for producers (but without exercise of any market power), then efficient retail prices should reflect those high short-run marginal costs.

be raised in order to cover the past investments of suppliers, such revenues need not come from marginal energy prices. Fixed charges, subscription charges (*e.g.,* based on the customer's circuit breaker capacity), demand charges (*e.g.,* based on the customer's maximum hourly demand), or government subsidies are among the alternatives that can be used to increase revenue collection without raising marginal price, though these alternatives can also create distortions.

IV. External Marginal Costs and Total Social Marginal Costs

For external marginal cost, we build on Graf Zivin, Kotchen and Mansur (2014) and Holland, Mansur, Muller, and Yates (2016), as well as the newer Air Pollution Emission Experiments and Policy (AP3) pollution damage model (see Clay, Jha, Muller and Walsh (2018)) to estimate the marginal damages associated with a change in load. We do this analysis for the same nine U.S. regions introduced in section III. The details of the estimation are in the appendix. In brief, for each of the four major pollutants from electricity generation (CO_2, SO_2, NO_x, and $PM_{2.5}$), we create a variable that is total emissions damages by hour of the three-year sample for each of the nine regions, incorporating the operations of each fossil fuel power plant and the damages associated with emissions from each plant, based on the AP3 damage model for 2014.[24] In this analysis, we assume that the "social cost of carbon" (SCC) – the net present value of damages resulting from additional emitted CO_2 – is \$50 per metric ton. In the appendix, we present the results of an analysis with the SCC assumed to be \$100 per ton.

We then regress each pollutant damage variable on piecewise linear functions of the load within the same region and the load in the other regions that are part of the same grid interconnect (Western, Eastern, and Texas). The regressions are estimated in 24-hour differenced form, so identification is based on the change in emissions from day to day in response to a change in load. The coefficients of these regressions can be interpreted as estimates of the marginal damage from a change in load in one region as a function of the load level in that and interconnected regions.

We use these coefficients to construct the damage associated with marginal electricity consumption in each of the nine regions for each hour of the sample. We do make two small adjustment to these damage estimates. The first involves scaling up the calculations of pollution associated with a marginal end-use kWh to account for distribution losses as discussed above. The second involves adjusting down our estimates of external costs to account for any policies that incorporate

[24]The AP3 model links county-level emissions of criteria air pollutants and greenhouse gases to an air transport model and physical damages in order to establish costs of emissions by pollutant/county. It incorporates assumptions about dose-response of various pollutants and value of a statistical life. For more information, see https://public.tepper.cmu.edu/nmuller/APModel.aspx. Our external marginal cost analysis does not incorporate emissions upstream of the power plant, in either fossil fuel extraction or transportation. The extent of externalities from extraction and transportation, particularly for natural gas, is the subject of widely divergent views. See Alvarez et al (2018) and Jha and Muller (2018). Including these upstream externalities would increase external marginal costs and decrease price minus social marginal cost.

externality costs into electricity prices, such as carbon cap-and-trade programs. The resulting aggregate damages from incremental quantity consumed in each region are shown in table A1 in the appendix. The results show that CO_2 is consistently a major contributor to the external marginal cost, but in areas with high levels of coal-fired generation, SO_2 tends to impose as much or more cost.

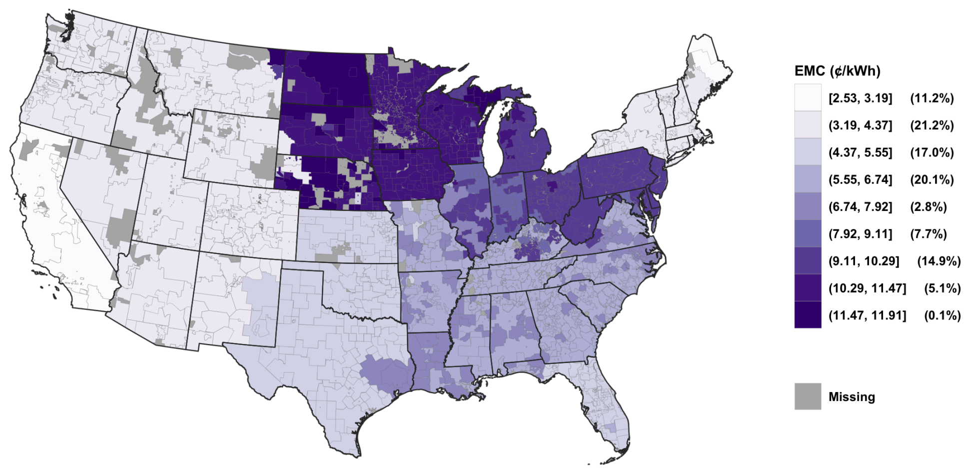

Figure 7: Average External Marginal Cost per kWh

A. External Marginal Cost Results

In figure 7, we show the average externality cost per kWh. The figure shows the average dollar-value externality cost associated with a marginal kWh of demand change in each location. The figure illustrates some coarseness in these data, because the analysis assumes that the same plants are marginal for any incremental demand within each of the nine regions for a given hour of the sample regardless of the location of the incremental demand in the region. Still, the figure demonstrates that externality costs vary widely and are particularly large in the areas where coal-fired power plants are most prevalent. Comparing the scales of figure 5 and figure 7 also indicates that the majority of the social marginal cost in our calculations in most locations is due to externalities, rather than the private marginal cost of generation.

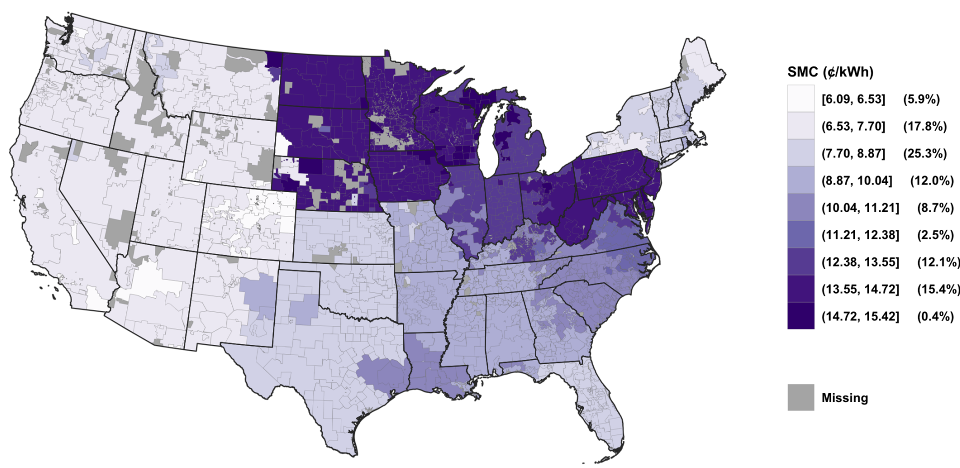

Figure 8: Average Social Marginal Cost per kWh

B. Total Social Marginal Cost Results

Figure 8 then aggregates the data in figures 5 and 7 to present the average social marginal cost. Though California has among the higher private marginal costs, the external marginal cost associated with that generation is much lower than in most of the U.S., causing it to have among the lowest SMCs. In contrast, the upper Midwest has low PMC, but such high EMC that it exhibits a very high SMC. Table 5 shows that the average quantity-weighted social marginal cost is 9.9 cents per kWh, nearly two-thirds of which is due to external marginal costs.

	Mean	StDv	Min	P10	P90	Max
Retail Variable Price (P, ¢/kWh)	11.49	3.07	2.36	8.79	16.29	48.22
Private Marginal Cost (¢/kWh)	3.72	1.15	2.16	2.59	5.10	8.22
External Marginal Cost (¢/kWh)	6.21	2.38	2.50	3.04	9.38	12.12
Social Marginal Cost (SMC, ¢/kWh)	9.93	2.67	5.14	6.51	13.72	17.71
P - SMC (¢/kWh)	1.56	4.21	-9.39	-2.82	6.74	35.89
(P - SMC) / P	0.08	0.31	-3.47	-0.28	0.51	0.81

N=6215 (utility-state-years). Statistics are sales-weighted

Table 5: Annual Averages of Prices and Marginal Costs

Note: Each observation is the hourly-sales-weighted average value of the variable for a utility-state-year. These summary statistics are weighted across observations by the utility-state annual sales.

V. Mispricing and Deadweight Loss

Figure 9 presents the marginal price minus average social marginal cost map. The bluer areas are pricing above average SMC, while the redder areas are pricing below average SMC. Much of the country has fairly light colors, indicating that the static marginal price that residential customers pay is fairly close to average SMC. California and parts of New England are notable for price being well above SMC, while parts of the Dakotas, Nebraska and Minnesota exhibit the largest price deviations below SMC. In the legend, the percentages in parentheses are the share of residential customers in each category.

Figure 9, however, captures only part of the story, because it does not include variation in SMC over time. The static price might reflect the average SMC well, but still create significant inefficiency because the SMC varies substantially hour-to-hour. Figure 10 shows histograms by state of the hourly markup, $\frac{P-SMC}{P}$, illustrating that SMC varies quite widely in some states, while it varies much less in others.

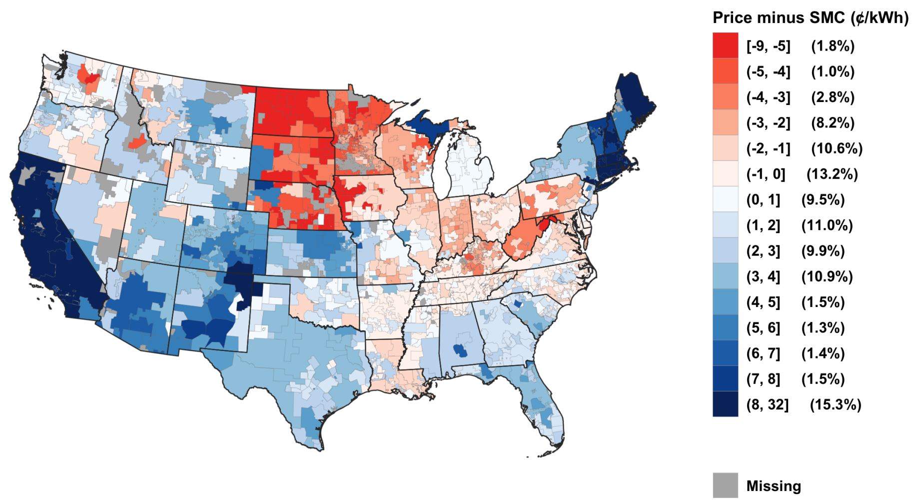

Figure 9: Marginal Price minus Average Social Marginal Cost per kWh

A. Analyzing Deadweight Loss

In order to study the social cost of electricity mispricing, we next analyze deadweight loss (DWL) directly. As discussed above, residential retail electricity pricing departs from efficiency both by charging a single time-invariant price – rather than allowing price to change hourly or more frequently as social marginal cost changes – and by setting that single price at a level substantially different from the average social marginal cost. We decompose mispricing into these two components and allow different demand elasticities, a "short run" elasticity that reflects consumer response to high-frequency price variation and a "long run" elasticity that reflects consumer response to changes in the long-run average price they face. A detailed explanation of these calculations is in the appendix.

Estimates of demand elasticity vary substantially in the literature, but there seems to be widespread agreement that demand elasticity is dependent on technology and the longevity of price variation. In the last decade, these elasticities very likely have increased with improved alternatives, from end-user batteries to advances in technology for automated price response (Gerke et al. (2020)). This trend is likely to continue with higher levels of electric vehicle adoption. Long-run electricity demand elasticity has generally been thought to be greater than short-run elasticity (Zhu et al (2018)), but that could change in the future as technology allows time-shiftable uses to be moved away from the highest priced hours, in-

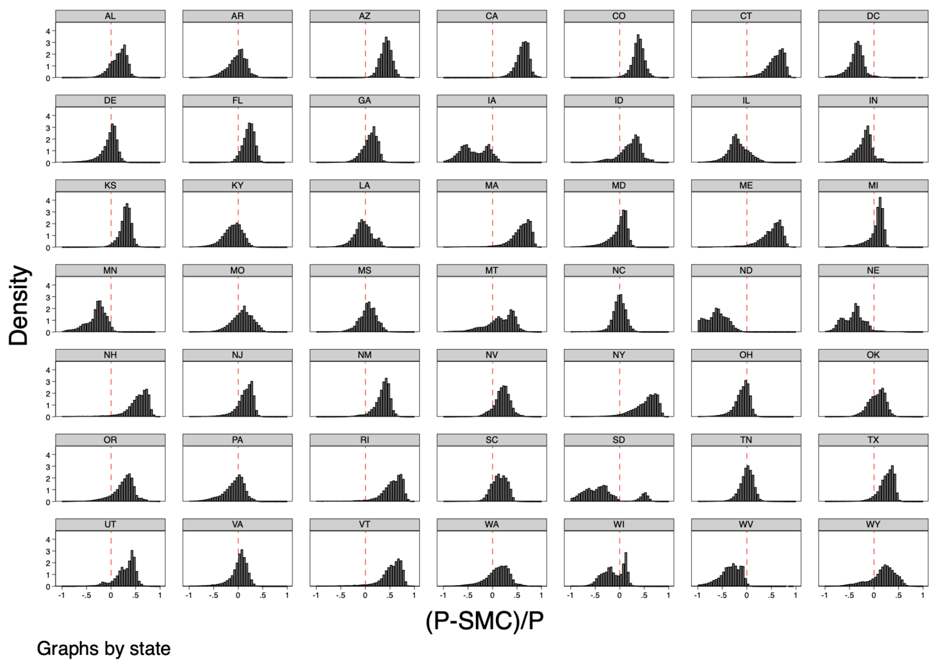

Figure 10: Hourly Price Markup ($\frac{P-SMC}{P}$) by State

cluding water heating, dish and clothes washing and drying, and to some extent even space heating and cooling (through pre-heating and pre-cooling).[25] Rather than base our conclusions completely on one set of demand elasticity estimates, we present results for a base illustration of a short-run elasticity of -0.2 and a long-run elasticity of -0.5 – which are approximately the median estimates found in a recent meta-analysis (Zhu et al (2018)) – but in the appendix we include results for a wide range of possible elasticities.

DWL calculations are also a function of how SMC changes with quantity. Consistent with discussion of externalities above, we have estimated an SMC elasticity with respect to quantity at the NERC region level, which we detail in the appendix. We find, however, that the range of quantity changes that result from the price variation and assumed demand elasticities are associated with miniscule changes in SMC and have almost no effect on our DWL calculations. Thus, for computational simplicity, we assume that SMC is constant over the range of the quantity changes considered.

DWL from mispricing electricity is also potentially influenced by mispricing of other goods in the economy, most importantly substitutes for electricity such as natural gas and gasoline (Davis and Muehlegger 2010, Borenstein and Davis 2012). In follow-on work (Borenstein and Bushnell 2021), we investigate distortions in pricing these substitute energy sources. Though we find that US gasoline is generally underpriced relative to SMC and natural gas is generally overpriced, on an energy-equivalent basis the deviations for both are far smaller in absolute value than for electricity.[26]

We calculate total DWL based on the departure from efficient pricing equal to the hourly SMC. The component of DWL that is due to charging a static price equal to the quantity-weighted average SMC at all times, rather than hourly SMC, is calculated using the short-run demand elasticity.[27] The component of DWL due to charging a static price that deviates from the quantity-weighted average SMC is then calculated using the long-run demand elasticity.

Figure 11 presents the histogram of total DWL per kilowatt-hour. To avoid biasing the calculation due to the endogeneity of quantity to $\overline{P}$, however, each utility's quantity of kilowatt-hours is normalized to the quantity it would have sold at $\overline{P} = \overline{SMC}$ given the assumed long-run elasticity. The histogram is also weighted by these normalized quantities. The results presented show that most utilities exhibit fairly low DWL per unit demand, compared to an average SMC of about \$0.10 per kWh, but a few show much higher DWL. Figure 12 shows the histogram of utility $\frac{DWL}{SMC}$. This ratio also is quite low for most utilities, but

[25] As this discussion suggests, a complete analysis of DWL from electricity mispricing would require estimation of the full matrix of own- and cross-price elasticities across hours, an extremely challenging task that we leave for future research.

[26] We exclude congestion and accident externalities from the gasoline calculations, assuming those would be approximately unchanged if vehicles were powered with electricity.

[27] Borenstein and Holland (2005) show that the DWL-minimizing static price is the quantity-weighted average SMC if demand elasticity is static or if time-varying demand elasticity is not correlated with time-varying SMC.

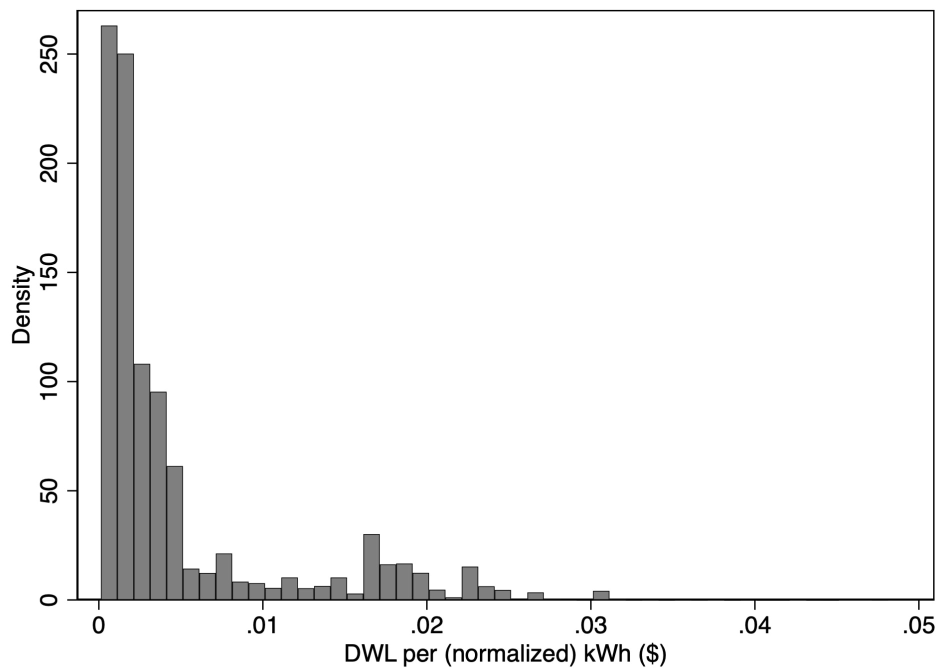

Figure 11: Distribution of Deadweight Loss per Normalized Kilowatt-Hour

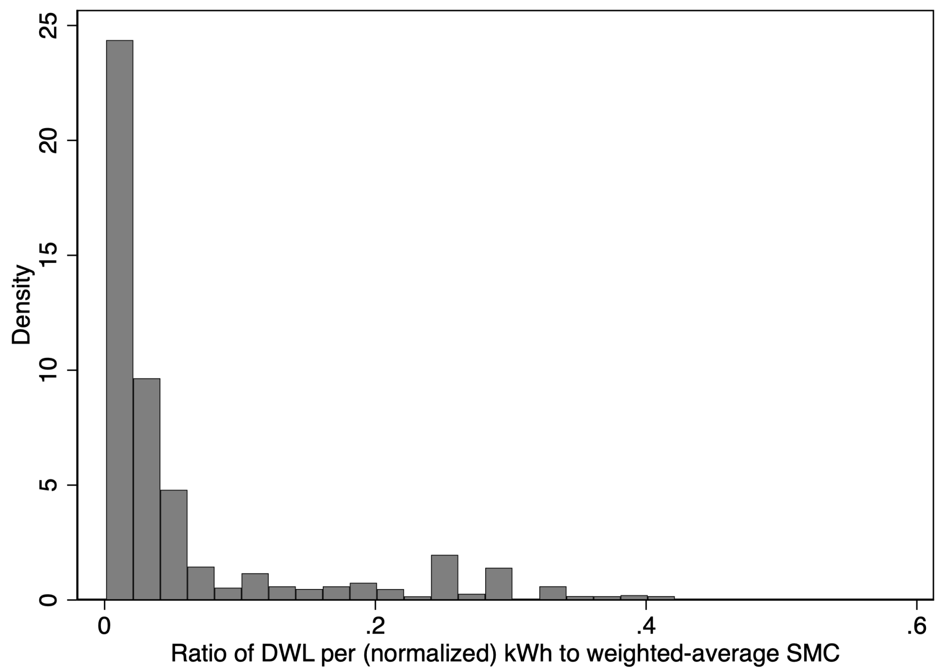

Figure 12: Distribution of DWL/SMC Ratio

for some it is 0.2-0.4, suggesting that, at these demand elasticities, mispricing dissipates as much surplus as would a 20%-40% increase in the social marginal cost of supplying electricity.[28]

For this base case, about 77% of the total DWL is a result of the long-run response, that is, the failure to set $\overline{P} = \overline{SMC}$, with the remainder attributable to the absence of real-time price variation. The larger assumed long-run elasticity is partially responsible for such a high share. If the short-run and long-run elasticities are set equal to one another, in the -0.1 to -0.9 range, then 57%-62% of the DWL would be attributed to the long-run response.

So, the answer to our title question is that two pricing wrongs do indeed make a (mostly) right for most utilities. One might ask whether this is because each pricing wrong is relatively small or whether it is a serendipitous result of offsetting "wrongs." To address this question, we carry out a similar DWL analysis as above, but with each of the pricing wrongs separately. First, we consider how large DWL would be if there were no environmental externalities so that SMC were equal to PMC in every hour.[29] Second, we consider how large DWL would be if there were environmental externalities, but utilities set a static volumetric price equal to $\overline{PMC}$, perhaps by imposing substantial monthly fixed charges.

Figure 13 presents the (normalized) quantity-weighted distribution of DWL for the base case and these two alternative cases with only one of the pricing wrongs. It makes clear that the low DWL in the base case is the result of a lucky coincidence of the two offsetting distortions. With both distortions, the mean DWL per normalized quantity is \$0.005 per kWh, but with either distortion alone, it is about five times higher, \$0.023 with no environmental externalities and \$0.027 with environmental externalities and price set equal to $\overline{PMC}$. With either distortion alone, the DWL from mispricing dissipates on average about as much surplus as would a 25% increase in SMC.

VI. Applications and Implications

Having calculated estimates of both the marginal prices and marginal social costs of electricity, we now consider some policy areas where such information ideally would be considered, and the implications of our calculation for such policies. One area where our calculation has potential relevance, but has received limited policy attention in the U.S., is the application of carbon pricing to the electricity sector. As discussed above, policy debates over the design of carbon pricing policies periodically invoke the Pigouvian ideal of capturing the marginal externality costs of greenhouse gases in consumer energy prices. Mechanisms such

[28] California's three large investor-owned utilities are among these outliers. Together they make up 8% of the total US residential normalized quantity (at a long-run elasticity of -0.5), but are responsible for 31% of the DWL.

[29] We don't adjust retail price, because nearly all of the deviation of retail price from $\overline{PMC}$ is unrelated to expenses associated with environmental control or remediation. Even in California, Borenstein, Fowlie, and Sallee (2021) find that a small share of the gap between residential retail price and SMC is attributable to environmental mandates, though it is difficult to know how much is due to climate change adaptation.

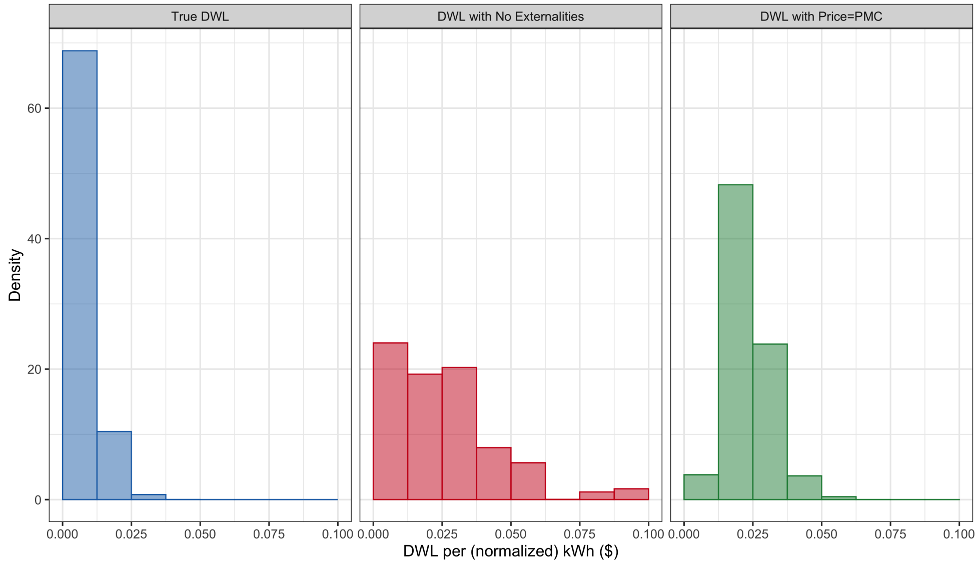

Figure 13: Distribution of DWL per Normalized kWh with One or Both Pricing Wrongs

as output-based updating of allowance allocation, and the application of intensity standards, have been criticized on the grounds that they dilute the externality cost faced by consumers (Holland et al. 2009, Fowlie 2011).

However, if marginal prices are already above social marginal cost, the additional externality signal only pushes prices further away from first best. It is worth noting that in the United States, carbon pricing - in the form of cap-and-trade - is currently applied to electricity only in California and the northeastern states comprising the Regional Greenhouse Gas Initiative. However, these are the collection of states where we have found average retail prices to be well above social marginal cost.

Still, it is important to recognize that our analysis focuses only on the distorted consumption incentives when residential retail price deviates from social marginal cost. We have not studied commercial and industrial rates, which are more complex, with greater use of time varying pricing and demand charges that determine (and distort) customer incentives. More importantly, our analysis does not consider the effect of market mechanisms for greenhouse gases and other pollution externalities on the mix of generation, between coal-fired generation, gas-fired generation, nuclear power, renewable generation and other sources. The efficiency value of pricing emissions at the wholesale level seems likely to be quite significant. Our findings, however, suggest that the argument for passing through those costs to residential rates is much weaker in some parts of the country.

Our findings also have direct implications for two other areas that have received considerable attention in the energy and economics literature: energy efficiency and distributed energy resource policy. We explore each of these in turn. We do not attempt here to perform a detailed calculation of the welfare implications of these policies, but rather present suggestive evidence that efforts in both areas may be significantly geographically misaligned with the benefits they can provide.

A. Energy Efficiency

The subject of energy efficiency in general, and its role in the electricity industry in particular, has been a topic of debate among economists and technologists for decades. Much of the debate has focused on whether these programs deliver the reduced energy consumption claimed by the utilities that implement them (Joskow and Marron 1992, Auffhammer et al. 2008). Economists have also examined the specific behavioral, regulatory, and market channels that could justify energy efficiency policies (Allcott and Greenstone 2012, Gillingham and Palmer 2014). However, much of the literature on the "efficiency gap" has focused on what Gerarden, Newell and Stavins (2017) call the "private energy-efficiency gap" - the question of whether customers are making individually rational economic choices. They note that the more policy-relevant question of the social energy-efficiency gap hinges on many factors, including the relationship of energy prices to social marginal cost, a question they identify as a "relatively high priority" for further research. Indeed, well-informed consumers who face retail prices that are significantly above social marginal cost are already being given too much incentive to adopt energy efficiency measures. If consumers are able to make privately optimal energy-efficiency decisions, government programs to promote improved energy efficiency would be best aimed at areas where price is below social marginal cost, and could be welfare-reducing where price is well above social marginal cost.

Several recent papers have attempted to address aspects of the relationship between energy efficiency programs and the social benefits they provide. Both Novan and Smith (2016) and Boomhower and Davis (2017) examine the impact of energy efficiency programs on the hourly profile of energy use, and compare those impacts to wholesale power costs and environmental impacts.

Using state-level data from the Consortium for Energy Efficiency,[30] we examine per-customer reported expenditures on residential energy efficiency programs.[31] This includes both energy efficiency programs run through utilities and those run through non-utility organizations, which play a significant role in New York, Oregon, Vermont, and parts of California, for instance. Other efficiency measures, such as appliance and building standards, impose costs on firms and consumers that are also not captured in these data. Still the data presented here are strongly reflective of the relative emphasis that different jurisdictions place upon energy efficiency measures. Figure 14 illustrates the state-level expenditures per customer

[30] https://www.cee1.org/annual-industry-reports

[31] Our thanks to Hunt Allcott for suggesting this comparison.

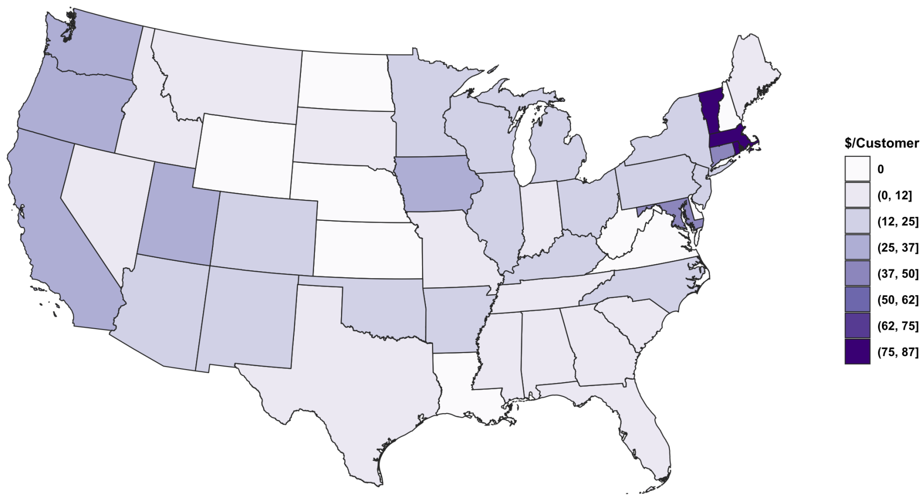

Figure 14: Electric Utility Expenditures on Energy Efficiency Programs

of electric utilities on energy efficiency programs. The largest expenditures are focused on the coasts, with particular intensity in California and the northeast. According to our calculations, these are the regions where marginal energy efficiency expenditures provide the least, possible even negative, social value. Clearly, the distribution of spending on energy efficiency within the US is suboptimal at best.[32]

B. Distributed Energy Resources

Another area of energy policy that is directly impacted by the relationship between retail prices and marginal cost is the deployment of small-scale distributed energy resources. Small scale generation resources, currently overwhelming comprised of rooftop solar photovoltaic (PV) installations, are deployed "behind the meter." Most of these systems are generally eligible for "net metering," meaning that when a customer's production exceeds consumption in an hour, the excess production is allowed for billing purposes to offset excess demand in other hours. This allows residential customers with distributed generation to offset the full

[32]This discussion assumes that the purpose of energy efficiency programs is to reduce consumption because $P < SMC$ gives incentives for excessive consumption. To the extent that these programs have other goals – such as correcting non-optimizing private consumption choices caused by imperfect information, myopia or irrationality – these other goals would also need to be considered.

retail price of electricity, rather than the marginal replacement cost of the energy that is produced. Where retail variable prices substantially exceed the marginal cost, residential solar is considerably more attractive for consumers. In California, Borenstein (2017) calculates that the gap between retail and wholesale marginal electricity prices means that net metering provides about as large an incentive for residential solar as the 30% federal investment tax credit.

Drawing again from the EIA Form-861, we aggregate the capacity of distributed resources that is subject to net metering by utility reporting area. Figure 15 illustrates the capacity of distributed generation (in Watts) per customer for the utility systems that report this statistic to the EIA. California, with over 40% of the residential solar capacity in the nation, again dominates this calculation.

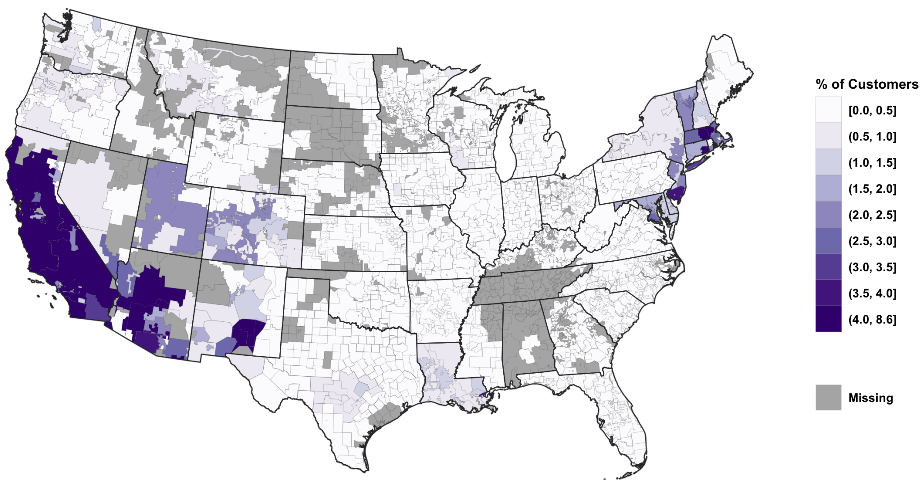

Figure 15: Percentage of Residential Customers with Rooftop Solar

Note: This figure shows the proportion of residential customers under a "net metering" tariff, which allows a customer to sell their own generation to the grid at retail prices. Virtually all customers with rooftop solar are covered by net metering and virtually all residential net metering is associated with rooftop solar.

The map reflects the union of at least three sets of attributes: significant solar incentives (*e.g.,* New Jersey), solar potential (desert southwest), and high retail prices. Comparing figure 15 to figure 10, the strong relationship between high retail prices and solar deployment again stands out. A full calculation of the welfare implications of retail tariffs on distributed generation would require a decomposition of rate effects from other incentives, as well as estimates of the relative efficiency of solar deployment in different locations. However, figure 15 does suggest that expenditures on distributed solar are strongly associated with retail price incentives that greatly exceed the social value of distributed generation.

The deployment of distributed energy resources, and the resulting reduction in metered consumption, known as "load defection," is a growing threat to the finances of distribution utilities who have been recovering capital cost though volumetric rates. Critics and proponents of small-scale distributed generation have battled over net metering policies, but net-metering policies lose their relevance if the marginal retail rate reflects social marginal cost. Recognizing this fact, utilities are increasingly seeking to adjust their rate structures to increase monthly fixed charges and reduce their volumetric prices. While not a panacea (Borenstein 2016) a shift toward larger fixed fees, particularly in states like California where they are modest to non-existent, would partially insulate utilities from the loss of customer load and reduce the marginal private reward of solar deployment for customers.

Consideration of distributed generation also raises questions of their potential impact on distribution losses and other costs associated with distribution networks, such as voltage support. As discussed above, marginal distribution losses can be significant, reaching over 20% at times, which distributed generation could mitigate or exacerbate depending on location and timing of production. More generally, the degree to which optimized location and control of distributed resources could change the cost of distribution remains an important area of research. Collection of distribution-level data with higher temporal and locational resolution could help address these questions.

VII. Conclusion

Most policy recommendations from economists for responding to the challenge of climate change focus on "getting the prices right." But in retail electricity markets, the prices are wrong for many reasons beyond greenhouse gas emissions. In this paper, we have analyzed the direction and degree of mispricing in residential electricity.

We find that with the current generation capacity and remuneration mechanisms for generation, the short-run private marginal cost was quite low during 2014-2016, a (quantity-weighted) average of less than 4 cents per kWh, which is below most estimates of the long-run average cost that generation must cover to support new investment. Estimates of the average externalities associated with generation are nearly twice the level of private marginal costs. We show that distribution-level marginal line losses significantly increase both of these costs, by about 9% on average. Accounting for private and external marginal costs, and adjusting for distribution line losses, we find large variation in full societal marginal cost from a (quantity-weighted) 10th percentile of 6.5 cents per kWh to a 90th percentile of 13.7 cents per kWh.

Somewhat surprisingly, we find that across the country about 39% of residential sales at a time-invariant marginal electricity price are below the utility's average social marginal cost of providing electricity. But we find wide variation, with prices well above average SMC in California and the Northeast, and below in much of the Midwest and the South.

That comparison, however, captures only part of the inefficiency, because social marginal cost varies hour to hour while price does not for nearly all residential customers. Using median short-run and long-run demand elasticity estimates from a recent meta-study, we find that about three-quarters of the total DWL from electricity mispricing is due to setting a time-invariant price at an inefficient *level* while the other one-quarter is attributable to the failure to change retail prices dynamically as SMC changes. These proportions, however, are highly dependent on relative short-run and long-run elasticities, which we note are likely to change as consumer technologies evolve.

Nonetheless, the largest DWL today likely results from a small number of utilities, mostly in California, setting prices well above average SMC. In most of the country, we find that retail price is fairly close to SMC, creating relatively little DWL. We show, however, that this is a serendipitous outcome of the two pricing wrongs referenced in our title, with underpricing due to unpriced externalities on average about balancing out overpricing due to the recovery of significant fixed costs through volumetric prices. Our analysis suggests that if either there were no unpriced externalities or fixed costs were recovered through other charges or revenues so that price equaled private marginal cost, (quantity-weighted) average DWL would be about five times greater.

Our findings have implications not just for standard deadweight loss analysis of consumption, but also for common related policies on residential energy efficiency and distributed generation. Many states have aggressive programs to encourage such investments, but if prices already exceed social marginal cost, the value of additional investments beyond those that well-informed individuals would already choose to make is questionable. It is perhaps not politically surprising, but nonetheless economically concerning, that we find these programs are most prevalent in areas where retail prices are already substantially above social marginal cost.

Appendix for Online Publication

The data used in this analysis come from a diverse range of sources. The construction of the data necessary for this analysis can be divided into the following categories:

- The annual sales of electricity to residential customers
- The marginal retail price paid by residential customers
- The location of residential customers as determined by utility service territories
- The private marginal costs of serving electricity demand
- The external marginal costs of serving electricity demand
- The hourly load shapes to distribute annual residential demand throughout the year
- The losses associated with distributing electricity from the transmission grid to residential customers

Each of these categories is covered by a section below. All results were converted to constant 2016 dollars using Consumer Price Index data (US Census 2018). The last section covers the details of the deadweight loss calculations.

1. Residential Electricity Sales

The starting point for this analysis was the Form EIA-861 survey published by the US Energy Information Administration (EIA) (Energy Information Administration 2017a). This survey collects a range of valuable annual data on every electric utility in the US. Of primary interest for this work was the dataset on "Sales to Ultimate Customers" which contains annual data on kilowatt-hour sales of electricity, numbers of customers and retail revenues. These data are broken down by state, so there can be multiple entries for a single utility if it has customers in multiple states. These data are also broken down by customer class, such that the sales, revenues and customer numbers are reported separately for residential, commercial and industrial customer types.[33] There is also some other key information available through the EIA-861 including data on the ownership structure of a utility (*e.g.,* Investor Owned, Municipal, Cooperative, etc.); the various regulatory regimes each utility belongs to (*e.g.,* reliability regions or balancing authorities); the counties that are part of a given utility's service territory; and operational data such as the peak load in each utility's service territory, numbers of distribution circuits and line losses.

[33]Strictly speaking a Transportation customer class is also included, although during our analysis period this represents a negligible volume and so is largely ignored.

The analysis here is focused on residential customers, so all information on industrial and commercial customers was dropped. Only utility-state pairs serving at least some residential customers were retained. The analysis here also focuses on the continental 48 states and the District of Columbia because the necessary private and external marginal cost data are not available for Hawaii, Alaska or the US territories. We also opted to drop the very small number of utilities that were classed as "Behind the Meter" as we are interested in comparing residential customers receiving a standard electricity service throughout the US.

Finally, the data were reformatted to appropriately deal with the different ways that residential customers receive their electricity. Roughly 85% of customers still receive their electricity through a vertically integrated utility that provides "bundled" service. This means the utility that is procuring the electricity that customers consume is also the company that owns and operates the distribution network that delivers the electricity to customers homes. However, in some states the electricity sector has been restructured such that customers can choose their electricity provider. In this case the service has been "unbundled" such that one company provides the electricity procurement service (*i.e.,* the "energy" service) and another company distributes the electricity to the customer (*i.e.,* the "delivery" service). The company providing the energy service is subject to competition from other providers, and will be referred to here as the "retail choice provider". The utility providing the delivery service continues to be a public or regulated monopoly and will be referred to hear as the "local distribution company". Various states take different approaches to handling which of these two entities is in charge of the other aspects of electricity service, such as billing and customer service. Roughly 32% of customers have the option to receive their electricity this way, although only about half of these actually do have a retail provider that is not integrated with their local distribution company. A large number of these customers are concentrated in a few states such as Texas, Ohio, Pennsylvania and New Jersey.

To ensure these customers can be correctly incorporated into the analysis, the data were reformatted such that each entry had a "delivery" utility and an "energy" utility. For vertically integrated utilities providing "bundled" service these two entries were the same. For customers receiving "unbundled" electricity service these two entries would necessarily differ. Unfortunately, the EIA-861 data do not include information on how a given retail choice provider's customers and sales are divided among the various local distribution companies that are providing delivery-only service in a given state. As such, new entries were created for all possible state-by-state combinations of retail choice providers and local distribution companies. The sales and customer numbers were then allocated proportionally. In the limited cases where we had prior knowledge about the operations of a retail provider this was included before any proportional allocation.[34] Where there were discrepancies between the state totals for energy-only

[34]For example, Marin Clean Energy is effectively a retail choice provider in California and there

and delivery-only customer numbers and sales the convention was adopted that the energy service totals were correct and the delivery service totals were re-scaled accordingly. In general any discrepancies were relatively small and likely due to errors in reporting.

One final wrinkle in completing this reformatting was the approach taken to reporting in the EIA-861 by utilities in Texas. Unfortunately, the Texas utilities do not break out their reporting between "energy" and "delivery" service. Instead, the retail choice provider reports the sales, customer numbers and revenues as if they were providing a complete "bundled" service. This also means that the six local distribution companies that offer delivery services to the retail choice providers in Texas do not report any information in this part of the survey.[35] To remedy this and make the data for Texas consistent with the other retail choice states, additional data were collected from the Texas Public Utilities Commission on the residential customer numbers, sales and revenues for these six missing local distribution utilities (Public Utility Commission of Texas 2017b). These data were then matched with the retail choice providers using the same proportional allocation process used for the other retail choice states.

2. Residential Marginal Retail Prices

Once the EIA-861 data were collected and reformatted, it was then straightforward to calculate the annual average retail price paid by every residential customer. To do this, total revenues were divided by total kWh sales to get the average cents per kWh price. However, this is almost certainly not a good reflection of the marginal retail price faced by each customer for three reasons. First, electricity tariffs are usually designed as two part tariffs, with a fixed monthly charge and a variable per-kWh charge. Because fixed charges are so prevalent and can comprise a substantial portion of customers' bills, simply using the average price would overstate the marginal rate customers actually face. Second, for many utilities, there is variation in the variable per-kWh price individual customers pay even after accounting for fixed charges. The most common reason is that the per-kWh price a customer pays depends on the amount that a customer consumes (i.e. tiered rates where prices increase or decrease in discrete blocks of cumulative consumption). Less common reasons are that the price may vary by time of day (*i.e.,* "time-of-use" or "dynamic" pricing), or time of year (*i.e.,* seasonal pricing where winter and summer rates differ). Third, the structure of retail tariffs themselves are also not static over time and are updated as utilities' new regulatory cases are approved, as changes in certain costs are automatically passed through to customers or as retail choice providers alter their tariffs in an

are many local distribution companies that provide delivery service in the state. However, Marin Clean Energy's operations are limited to Marin County and nearby counties, so delivery service is only provided to its customers by Pacific Gas & Electric.

[35]These six utilities are Oncor Electric Delivery Company LLC, CenterPoint Energy, AEP Texas Central Company, AEP Texas North Company, Texas-New Mexico Power Company and Sharyland Utilities LP.

effort to win new customers.

To deal with fixed charges, we have collected information on the retail tariffs actually offered by utilities and extracted the monthly fixed charges. Our main source for this information is the National Renewable Energy Laboratory's Utility Rate Database (URDB) (National Renewable Energy Laboratory 2017b). This is an open-access repository for rate structure information for utilities operating in the US. The fixed charges for residential tariffs active during our analysis period were extracted, and the utility names were cleaned so that their corresponding identifiers and states matched those in the EIA-861 data. At the time of writing, the URDB only contained tariffs for utilities providing "bundled" service. This presented us with a similar challenge to the EIA-861 data in dealing with the roughly 15% of customers with a retail choice provider that differs from their local distribution company. To resolve this, we manually collected additional fixed charge information for the largest retail choice providers in the states with substantial numbers of retail choice customers (Public Utility Commission of Texas 2017a).[36]

Once we had finished collecting all the necessary data on fixed charges, we found that it was almost always the case that a given utility operating in a given state had many different residential tariffs. The average fixed charge paid by a given utility's residential customers must therefore be some weighted average of the fixed charges in each of these tariffs, with the weights determined by the number of customers on each tariff. Unfortunately we know of no comprehensive data source that could give us this breakdown of customers by tariff. As such we summarized the fixed charges in these tariffs by identifying the standard tariffs that were most likely to have many customers on them, as compared to the more niche non-standard tariffs that few customers were likely to be on. We did this by searching for keywords in the names of the tariffs. Tariffs containing words like "vehicle", "solar", "medical" or "three-phase" were identified as non-standard. This tended to leave us with a set of more standard tariffs with names containing words like "default", "residential" and "general". Full details of the keywords used can be found in the accompanying code. Once these standard tariffs had been identified, we took the median, giving us a single estimate of the residential fixed charge for each utility-state pair. We considered other approaches to combining these (e.g. mean or mode), but this did not significantly affect our results. It was also often the case that utilities had similar or identical fixed charges on many or all of their residential tariffs. We checked our selection of standard rates for 166 large utilities that report in FERC Form 1 the number of customers on each paragraph. For those utilities, 92% of our selected standard rates matched the rate with the most customers on FERC Form 1. Over 95% differed by less than $2 per month in absolute value.

[36]In collecting these data we sought to capture whether the fixed charges offered by a given retail choice provider varied depending on the local distribution company whose service territory their customer was located in. In general though we found very little evidence of utilities having much variation in their fixed charges for this reason.

Once this exercise was complete, these rates were matched with the utility-state pairs in our reformatted version of the EIA-861 data. At this point it was now possible to estimate the second part of the two part tariff - namely the average variable per kWh price. To do this, the fixed charge was multiplied by the number of residential customers to get fixed revenues, these were subtracted from total revenues to get variable revenues, and these were then divided by total kWh sales to get the average variable cents per kWh price.

The second issue in identifying the marginal retail price was dealing with the fact that utility tariffs often do not contain just a single flat per-kWh variable price. This could mean that the average variable per kWh price calculated using the fixed charge information described above does not reflect the actual marginal price paid by customers. The URDB does in fact contain some information on the structure of the per kWh prices in each tariff (e.g. tier sizes and prices for increasing- or decreasing-block rates, or peak vs off-peak rates and timings for time-of-use pricing). However, these data are necessarily complex, and they are less complete than the fixed charge information we had already extracted. As already noted, these data also don't cover retail choice providers, so significant additional manual collection would be required to make these data complete. Furthermore, to properly use this information we would need to know both how many customers are on each tariff and the consumption patterns of the customers on each tariff. To the extent that these data are held by individual utilities they are confidential.

Thus, we have opted here to conduct the analysis assuming that all utilities charge a single flat variable per kWh price. While this is obviously not strictly true, we believe it is not an unreasonable assumption for the purposes of our analysis. To look at the issue of variation in prices due to seasonal factors changing flat or tiered rate structures we calculated monthly estimates of the variable per kWh rate. To do this we used the EIA-861M survey which is a monthly version of the annual EIA-861 survey that covers a sample of the complete population of utilities (Energy Information Administration 2017b).[37] For this subset of utilities, we found that the variation is fairly small compared to average variable prices, with the vast majority of monthly implied average variable prices within 10% of the annual average variable price. Given the cost drivers and regulatory arrangements in the electricity sector, it is unclear whether accounting for more frequent retail rate changes would align retail rates with contemporaneous marginal cost more closely. To look at the possibility of hourly variation in retail prices during the day we examined evidence from the "Demand Response" and "Dynamic Pricing" sections of the EIA-861 survey. These sections provide data on the numbers of customers participating in demand response programs or subject to some form of dynamic pricing tariff. We find that around 4% of residential customers in the US are on tariffs with time-varying prices. This includes time-of-use, real

[37]In 2015 the EIA-861M contained information on utilities accounting for 67% of residential customers and sales.

time, variable peak and critical peak tariffs. Demand response programs are also limited in scope with less than 6% of customers enrolled in a demand response rebate program during 2014-2016. There is also likely substantial overlap in the customers exposed to these two forms of price variability. Roughly three quarters of the customers on tariffs with time-varying prices or in demand response programs are served by the same set of 96 utilities.

A closely related issue for many utilities is that a share of customers are on low-income rates, which in many cases are lower marginal rates than the standard tariff. Our analysis captures the average variable payment (assuming that we have correctly characterized the fixed charges), but it is possible that some customers pay a higher marginal rate and others pay a lower marginal rate. We are not able to capture such variation in marginal rates across customers. It is worth noting, however, that because DWL increases with the square of the price deviation, such variation would almost certainly mean that our analysis understates the deadweight loss associated with marginal rates deviating from average SRSMC.

3. Utility Service Territories

To match up our data on retail rates with information on social marginal costs, we had to represent the spatial distribution of residential customers. For this we used information on the service territories of the local distribution companies that distribute electricity to end consumers.

Our main source for this was a lookup file provided as part of the URDB (National Renewable Energy Laboratory 2017a). This provides a list of ZIP Codes served by each local distribution company. These lookups were created using a proprietary set of shape files detailing the actual service territories of major electric utilities, which were converted to a list of ZIP Codes falling within those service territories. Unfortunately the ZIP Code lookups did not cover all the utilities in our dataset. To fill in any gaps we relied on the "Service Territory" section in the EIA-861 survey. This provides a list of counties served by each local distribution company. For consistency these were converted to ZIP Code lookups by assuming any local distribution company serving a given county also served all the ZIP Codes in that county. Our spatial data on US ZIP Codes were downloaded from Environmental Systems Research Institute and included polygons for 30,105 ZIP Code areas, and central coordinates for the full universe of 40,552 ZIP Codes (Environmental Systems Research Institute 2017).[38] These data were used as they were more comprehensive than the Zip Code Tabulation Area data available from the US Census Bureau.

To increase the accuracy of our geographic allocation of residential customers within a given service territory we also collected data on population counts by ZIP Code. The vast majority of these data were from the ESRI spatial data

[38]The latter is larger because it includes ZIP Codes that have no associated area such as post office box ZIP Codes and single site ZIP Codes (e.g. government, building, or large volume customer).

we downloaded, as this also included estimates of population for each ZIP Code based on ESRI's analysis of US Census Bureau data. However, there were a few ZIP Codes where the population data were missing but where we were confident that people lived. To remedy this, county-level population data were downloaded from the US Census Bureau, along with spatial data on US counties and a set of lookups from counties to ZIP Codes (US Census 2017a, US Census 2017b, US Census 2017c). The ZIP Codes with missing data were then assumed to have a population density equivalent to the county they belonged to. Missing ZIP Code population counts were then calculated as the county-level population density multiplied by the ZIP Code area.

The matching of utility service territories to ZIP Codes, or counties, was used to assign LMPs and load to zip codes before aggregating the zip codes to the utility level (described below). For the final mapping of the data, we use utility service territory boundary shapefiles from HIFLD, as described in the paper.

affects only the visual representation of results in maps. It does not affect any of the underlying results by utility, or the calculations of deadweight loss and its decomposition.

4. Private Marginal Costs

The primary source of the data for calculating private marginal costs was the price information provided by the seven major US Independent System Operators (ISOs).[39] These are Electric Reliability Corporation Texas (ERCOT), the New England ISO (ISO-NE), the New York ISO (NYISO), the California ISO (CAISO), the Southwestern Power Pool (SPP), the Midcontinent ISO (MISO) and the PJM Interconnection (PJM). Each of these manages the operation of the electricity transmission grid over a large geographic area, most encompassing multiple states. These organizations calculate wholesale locational marginal prices (LMPs) for major locations in their covered territories, reflecting the value of electricity supplied at different points in the power grid. Each ISO has LMPs for thousands of pricing nodes within their system, such that across all seven ISOs there are in excess of 30,000 nodes with hourly price data available.[40] We did not consider it necessary to utilize data from all these nodes in our analysis. This was in part because prices at nodes located very close to one another are usually very highly correlated, so selecting a smaller number should still allow us to create a sufficiently robust picture of the main spatial and temporal variation. In light of this we selected a total of 157 key LMPs. All of these were aggregated "zonal" LMPs that represent averages of many individual nodal prices. In selecting these we were also mindful that different nodes can refer to a range of important locations in the power grid, such as power stations, load substations

[39]Strictly speaking some of these, such as PJM, are classed as Regional Transmission Organizations (RTOs) but for the purpose of this paper the distinction is largely immaterial, so we refer to all as ISOs.

[40]Often pricing data are available at even finer temporal resolutions (*e.g.,* 15 minute) but for this analysis we have used hourly data as they are consistently available across all seven ISOs.

or major interconnection points with neighboring systems. Wherever possible our selection focused on zones that were aggregates of load nodes or were used by regulators in their determinations of utilities' wholesale costs for supplying their customers. This clearly fits with our interest in finding the marginal cost of serving residential customer demand. These data were downloaded from SNL Financial (SNL Financial 2017b). This is a proprietary source of financial data and market intelligence and includes a convenient centralised database of LMP data from all seven ISOs.[41] All data were converted to Eastern Standard Time (EST) for consistency.

These seven ISOs cover large parts of the US. However, their coverage is not complete and they are most notably absent from the most of the Southeastern U.S. To remedy this and provide a secondary source of corroborating data we also used data from the Federal Energy Regulatory Commission's Form-714 survey (Federal Energy Regulatory Commission 2017). This survey collects data from electric utility balancing authorities (or control areas) in the United States. The seven ISOs are also classed as balancing authorities, so their aggregate system-wide data appear in the FERC-714 data. Importantly though, balancing authorities also include approximately 200 additional utilities and regulatory entities that undertake a similar electricity system operation role. This includes major utilities in the Southeastern U.S. The FERC-714 data used are the hourly system lambda data. Here respondents are supposed to report hourly values of the incremental cost of energy in their system. In principal this seems ideal. In practice, a check of the data reported by the ISOs shows that ISOs simply report LMPs as the system lambdas at various locations. Unfortunately, visual inspection of the system lambda data provided by the other balancing authorities reveals a range of suspect data, including respondents providing no data, respondents providing all zeros, respondents providing data that remain unchanged over long periods, and respondents providing data that differ substantially from LMPs at nodes in nearby ISOs. To deal with these weaknesses in the system lambda data, each series was individually inspected to determine if it was sufficiently robust to be included. This left just 19 balancing authorities (besides the seven ISOs) with reliable system lambda data. Fortunately this still included a number of balancing authorities in Southeastern states such as Florida and Alabama. As with the ISO data, all series were converted to EST for consistency. Unfortunately, the quality of the reporting of time zones and daylight saving time for these data is often unreliable such that it is not always clear what time format these data are in. In some cases respondents even left the time zone section blank. Where there were clear errors or gaps we sought to identify the reporting time zone and the presence of daylight saving time by visual inspection and the location of the reporting entity. We then manually corrected for this and adjusted to EST as appropriate. Lastly, the system lambda data do not account for transmission

[41] It should be noted that these data are freely available directly from each ISO. We have opted to utilize SNL Financial's database purely due to ease of accessing and compiling the data.

losses, while LMP data implicitly do. To remedy this, all system lambda prices were increased by an assumed transmission loss rate of 2%.

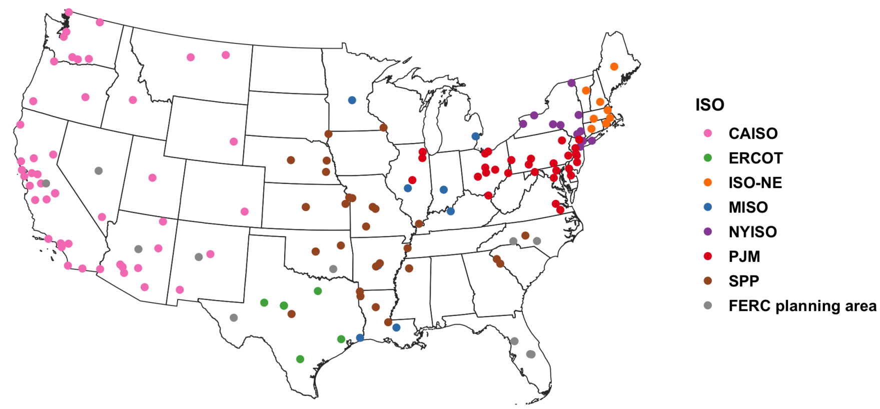

Figure A1: Locations of ISO zonal price points and Balancing Authority area system lambdas in 2015

Once the ISO and balancing authority data had been collected, we next sought to use these data to calculate hourly ZIP Code level estimates of the marginal private costs of supplying electricity. We chose to do this at the zip code level in order to accurately assign wholesale prices within a utility service territory. We then aggregated the zip code level prices by taking population-weighted averages of the wholesale prices across the zip codes within a utility service territory. To begin this process of creating ZIP Code-level prices we first had to determine where each ISO zone or balancing authority area was located. Unfortunately, we were unable to get access to the necessary spatial polygon data files detailing the areas covered by the ISO zones. Instead SNL Financial were able to provide us with a list of coordinates they use to represent the location of each ISO node, including the zonal nodes we had chosen for this analysis (SNL Financial 2017a). Strictly speaking, the ISO zonal nodes are themselves representing many individual nodes, but for our purposes the central coordinates of these zones are likely sufficient. For consistency we also represented the locations of the FERC-714 balancing authorities using the central coordinates of their respective network areas. These coordinates were calculated using the polygon centroid from spatial data on electricity balancing authorities downloaded from the Homeland Infrastructure Foundation-Level Data website, which is part of the US Department of Homeland Security (Department of Homeland Security 2017a). These spatial coordinates can be seen in Figure A1. Once these had been collected we calculated the distance to each ZIP Code centroid.[42] The price for each ZIP Code was then calculated as the inverse distance-weighted average of the prices at the three closest price nodes.[43] While the system lambda data can be considered a less accurate measurement of private marginal costs, less than 10% of utility-states (weighted by load) rely exclusively on system lambda pricing points. As another check, we also dropped all of the system lambda values and set prices for every utility purely on ISO pricing hubs. The result are almost indistinguishable from figure 9, though there are slight changes in some states, such as Florida, Georgia, South Carolina, Washington, Oregon, Idaho and Arizona. Comparing the results with and without system lambda values on a utility-by-utility basis, the (quantity-weighted) mean absolute difference in PMC is 0.07 cents, the 95th percentile is 0.49 cents and the maximum is 1.32 cents.

Average wholesale electricity costs are made up of energy costs, capacity costs, ancillary services costs and other uplift payments. Our use of LMP and system lambda data captures the energy cost component. Table 4 shows the relative contributions of each of these four categories across the seven ISOs (Electric Reliability Council of Texas 2015, California Independent System Operator 2016, Independent System Operator New England 2016, Midwest Independent System Operator 2015, New York Independent System Operator 2016, PJM Interconnection 2016, Southwest Power Pool 2016).[44]

The end product of the private marginal cost data collection process was a dataset of hourly estimates for each US ZIP Code. These data were then merged with the reformatted retail rates data using the information on the ZIP Codes served by each local distribution company. The hourly price assigned to a utility-state was an average of each of the ZIP Code prices, weighted by the total population of each ZIP Code.

5. *External Marginal Costs*

The AP3 model (see (Clay et al. 2018)) provides estimated marginal damage by county/pollutant/smoke stack height for 2014. This is an updated version of the AP2 model used in Holland, Mansur, Muller, and Yates (2016). The model

[42]This was done using the geodesic on a WGS84 ellipsoid to properly account for the curvature of the earth.

[43]Prior to calculating these averages we winsorized any extermely negative prices at a cutoff of -\$150/MWh. This only affected prices at a few nodes in a small number of hours and was done to avoid the calculations of deadweight loss being distorted by unusual outliers.

[44]These values are taken from the annual reports of each ISO. The one exception to this is capacity costs in the CAISO. Capacity payments in California are primarily agreed through bilateral contracts overseen by the CPUC's Resource Adequacy program, so do not show up as capacity costs levied by the ISO. To account for this we have calculated capacity costs using data from the CPUC's Resource Adequacy Report (California Public Utilities Commission 2015). This yields an additional capacity cost of approximately \$4/MWh, or approximately 9% of total wholesale costs.

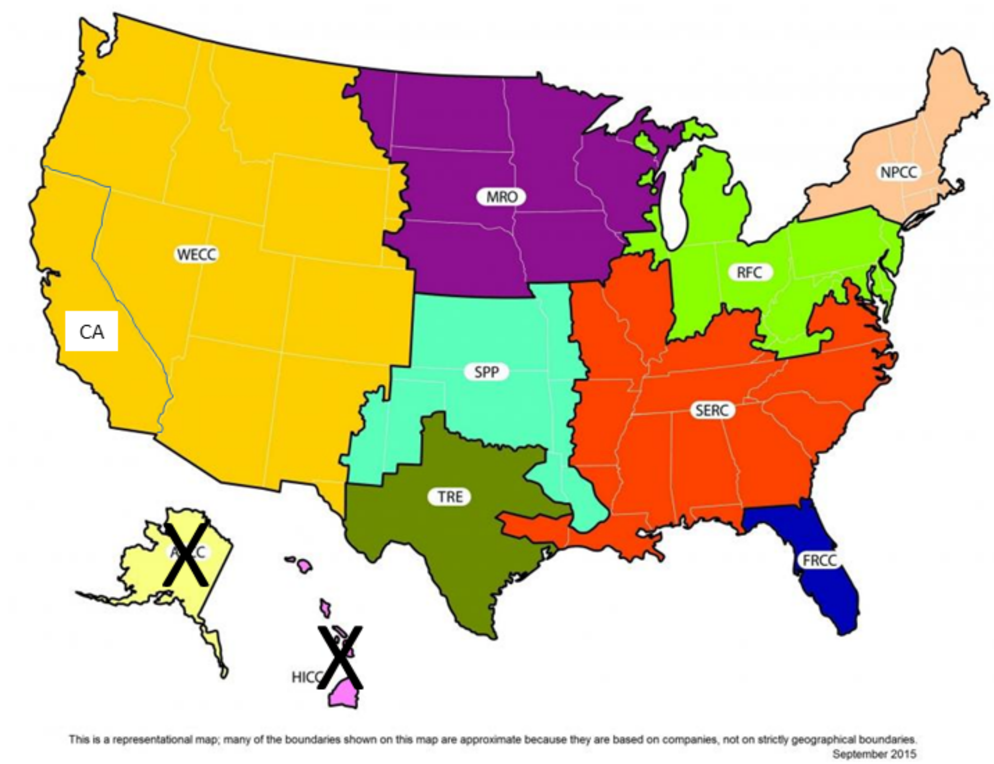

Figure A2: NERC Subregions used in analysis does not differentiate marginal damage by season or time of day, or by location within county. The data contain estimates of the environmental externality costs in dollars per marginal ton for four pollutants associated with the generation and supply of electricity: particulate matter (PM), nitrogen oxides (NO_x), sulphur dioxide (SO_2) and carbon dioxide (CO_2). Baseline damages assume pollutants are emitted at a height of 200-500m. This is classed as a "medium" height in the model and is in line with the smoke stack height for most fossil fuel power plants. The dataset also then has individual plant-specific marginal damage values for a small number of large power plants that have "tall" smoke stacks.

The data on power plant emissions are from the Environmental Protection Agency (EPA) Continuous Emissions Monitoring System (CEMS) (Environmental Protection Agency 2018a). The data are comprised of hourly emissions of NO_x, SO_2 and CO_2 from large stationary sources. For our purposes this includes more than 90% of the (output-weighted) fossil fuel power plants in the US. As well as emissions, the CEMS data also include hourly information on fuel energy inputs and electricity generated. These data do not include hourly emissions of PM. To resolve this we follow an approach suggested by Holland, Mansur, Muller, and Yates (2016). We use annual total emissions data by power plant from the EPA's National Emissions Inventory (NEI) (Environmental Protection Agency 2018c). We divide annual PM emissions by annual fuel energy inputs to get a PM emissions rate for each power plant. We then use the hourly fuel energy inputs information in the CEMS data to calculate hourly PM emissions, thereby assuming the annual rate is constant throughout the year. To match plants to counties and NERC regions (show in figure A2) we use plant characteristics data from EPA's Emissions & Generation Resource Integrated Database (eGRID) (Environmental Protection Agency 2018b).

The data on hourly load are from the FERC-714 survey described earlier (Federal Energy Regulatory Commission 2017). It contains hourly load data for planning areas in the US. These planning areas have a regulatory responsibility to ensure resources are available to meet customer load. There is considerable overlap with the balancing authorities discussed above for the system lambda data. The coverage and quality of the planning area load data are much better than for the balancing authority system lambda data, resulting in 122 planning areas with usable load data. Again we converted all data to EST using the same approach as the one set out above for the price and system lambda data. We then divided the contiguous U.S. into nine regions, in line with the approach taken by Holland, Mansur, Muller, and Yates (2016). These correspond to the eight reliability regions of the North American Electric Reliability Coorporation (NERC), with the exception of the Western Interconnection region which is split into a California region and a non-California region. Each planning area was then assigned to one of the nine regions - the regions cover the Eastern Interconnection (NPCC, RFC, MRO, SERC, SPP, FRCC), the Western Interconnection (CA, non-CA-WECC) and Texas (TRE). Each planning area was assigned to one of the nine regions.

The one exception here was MISO which actually spans several regions in the Eastern interconnect. To deal with this we collected data on kWh sales from the EIA-861 survey described earlier. We then identified both whether a given utility was in MISO, and also which of our nine regions it was in. We then used this to proportionally allocate the hourly MISO load across our nine regions. This primarily resulted in MISO being split fairly evenly between MRO, RFC and SERC.

To run the regressions to estimate marginal dollar per kWh damages we first combine the hourly emissions data for each plant with the relevant dollar per ton of marginal damages. For most plants this merge is done based on the county the plant is located in. For the small number of large plants with taller smoke stacks this is done using a plant-specific identifier. We then multiply emissions in each hour by marginal pollutant damages to get hourly dollar damages by pollutant for each plant. Next we sum together damages by pollutant for all plants in a given region, yielding a total dollar damages value for each region in each hour. We aggregate damages by pollutant to the region level, because we do not differentiate the location of load within a region for the marginal generation and emissions.

The basic externality regressions are, for each region and pollutant, a regression of the dollar-value damage of the pollutant in a given hour on the level of load in that region and the aggregate level of load in all other regions that are in the same interconnect (except for Texas, where there is no other region in the interconnect). In all, we estimated four regressions – one for each pollutant – for each of the nine NERC regions. Because the generation technology that provides marginal output varies systematically with the level of output in the region, we allowed the marginal pollution damages to be a nonlinear function of the "own region" and "other region" loads. We split the "own" region load data into terciles and created three variables that allow us to estimate a piecewise linear response to own-region load with separate slopes for the lowest, middle, and highest terciles of load. To be precise, if we define Q^{33}_{own} and Q^{67}_{own} as the 33.3rd and 66.7th percentiles in the distribution of "own region" load, and Q_{own} as the own region load, then the three variables used to estimate a piecewise linear function are:

$Q_1 = min\{Q_{own}, Q^{33}_{own}\}$
$Q_2 = 0\ if\ Q_{own} < Q^{33}_{own}\ else\ Q_2 = min\{Q_{own} - Q^{33}_{own}, Q^{67}_{own} - Q^{33}_{own}\}$
$Q_3 = 0\ if\ Q_{own} < Q^{67}_{own}\ else\ Q_3 = Q_{own} - Q^{67}_{own}$

We handle the terciles of "other-region" load differently, because we assume that own-region load is the primary determinant of the impact of incremental generation in a region on pollution. Thus, other region load is assigned to one of three variables depending on the tercile into which own-region load falls in that hour:

$Q_4 = Q_{other}\ if\ Q_{own} < Q^{33}_{own}\ else\ Q_4 = 0$
$Q_5 = Q_{other}\ if\ Q^{33}_{own} \leq Q_{own} < Q^{67}_{own}\ else\ Q_5 = 0$

$Q_6 = Q_{other}\ if\ Q_{own} \geq Q_{own}^{67}\ else\ Q_6 = 0$

While we could estimate each of the 36 regressions separately, we instead estimate the coefficients in a single "stacked" regression. That is, define y_{prt} to be the damage from pollutant p released in region r in sample hour t. And define I_r to be an indicator variable that is equal to 1 if the dependent variable is from region r. Then the regression can be written as:

$$y_{prt} = \sum_{r=1}^{9} I_r \cdot (\sum_{j=1}^{6} \beta_{jpr} Q_{jrt})$$

where j indexes the quantities defined above and their associated coefficients. We then cluster the standard errors on the day of sample, thereby accounting for correlated errors within the day for a given region/pollutant, and correlated errors across regions/pollutants on a given day. We do this because the errors are almost certain to be correlated across the regions/pollutants. Furthermore, as explained below, we need to construct estimates and standard errors of parameters that are linear functions of coefficients from different regressions. Unbiased estimates of the standard errors require accounting for the error correlation across regressions and the covariances of the coefficient estimates, which is straightforward to do in a stacked regression.

Once the dependent and independent variables are constructed in this manner, we 24-hour difference the data.[45] We estimate the linear regressions with three years of hourly observations (26,304 hours) for each of the 36 region/pollutants.[46] Due to the 24-hour differencing, the regression does not include a constant term, or hour-of-day or month-of-sample fixed effects.

From this regression, we then construct the marginal pollution damage due to marginal load in each region for each hour. The marginal pollution damage from marginal load in region A is the sum of the marginal pollution caused by that load from generation in region A and the marginal pollution caused by that load from generation in all other regions that are part of the same interconnect as region A. In each case, the appropriate coefficient is determined by the tercile of the load in which the generation resides.

For instance, assume that in hour h the CA region load is in tercile 1 (lowest) and the $WECC$ region load is in tercile 2. Then the marginal damage from marginal load in CA from that pollutant would be $\beta_1^{CA} + \beta_5^{WECC}$, where the superscripts indicate the regressions from which each coefficient is taken.[47]

[45] So for example, the dependent and all of the independent variables for hour 3 today are differenced with their values from hour 3 yesterday.

[46] One minor modification, as noted earlier: the Texas region is not interconnected with any other regions, so the variables Q_4, Q_5, Q_6 are zero when the dependent variable is a pollutant in Texas.

[47] To possibly belabor the point, but hopefully avoid confusion in the much more complicated Eastern interconnect, assume that in hour h, the $NPCC$ region is in tercile 1 (lowest), $FRCC$ is in tercile 2, MRO is in tercile 3, RFC is in tercile 1, $SERC$ is in tercile 2, and SPP is in tercile 3. Then the marginal

The estimated marginal damage (in dollars) from marginal load with demand in each tercile is presented in table A1.

These calculations produced values for the dollar-value marginal external damage per kWh for each region for each hour. We made a small set of adjustments to our estimates of external marginal costs to avoid double counting. This can arise where the private marginal costs data already incorporate some portion of external marginal costs due to environmental policies that put a price on externalities. The two main instances of this that are relevant here are California's Cap and Trade Program and the Regional Greenhouse Gas Initiative (RGGI) that covers nine states in the Northeastern US. Our external marginal cost estimates were created using a social cost of carbon (SCC) of \$50/ton of CO_2. The California and RGGI carbon prices in 2014-2016 averaged \$12.70/ton and \$6.00/ton respectively. We therefore multiply the \$/kWh external damages from CO_2 by approximately $(\$50 - \$12.70)/\$50 = 75\%$ for the state of California and by approximately $(\$50 - \$6.00)/\$50 = 88\%$ for the states that participate in the RGGI.[48]

Another potential complication is the impact of zero-carbon renewable energy resources that produce intermittently and with seasonally varying patterns. Note that our primary specification differences our variables of interest with the value from 24 hours prior. This accounts for any systematic seasonal variation in renewable energy output. However, to confirm that our analysis was not being affected by fluctuations in renewable generation we also gathered data on hourly renewables (wind and solar) for each of our nine regions. First we downloaded monthly generation data by plant from the EIA-923 survey (Energy Information Administration 2018). This includes generation from all plants including wind and solar (unlike the CEMS data which is focused on fossil fuel plants). We then matched information on the state and NERC region each plant is located in to aggregate the plant-level values and get monthly total wind and solar generation for our nine regions. Next, we used hourly data on renewable generation from the ISOs to allocate this monthly generation across the hours of each month and get our desired estimates of hourly renewable generation by region (Electric Reliability Council of Texas 2018, California Independent System Operator 2018, Midwest Independent System Operator 2018, Southwest Power Pool 2018, New York Independent System Operator 2018, PJM Interconnection 2018, Independent System Operator New England 2018). For each region we identified the most relevant ISO (or combination of ISOs) with which to do this within-month allocation.[49] Once

damage from marginal load in $NPCC$ from the pollutant would be $\beta_1^{NPCC} + \beta_5^{FRCC} + \beta_6^{MRO} + \beta_4^{RFC} + \beta_5^{SERC} + \beta_6^{SPP}$. And the marginal damage from marginal load in SPP from the pollutant would be $\beta_4^{NPCC} + \beta_5^{FRCC} + \beta_6^{MRO} + \beta_4^{RFC} + \beta_5^{SERC} + \beta_3^{SPP}$.

[48] These are Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New York, Rhode Island and Vermont.

[49] The CA region used CAISO for wind and solar. The TRE region used ERCOT for wind and solar. The SPP region used SPP for wind and solar. The MRO region used MISO for wind but solar was assumed negligible. For the SERC region both wind and solar were assumed negligible. The RFC region used PJM for wind and solar was assumed negligible. For the FRCC region both wind and solar were assumed

Region	Tercile	Marginal Damages				
		CO_2	NO_X	SO_2	$PM_{2.5}$	Total
CA	1	17.25	1.29	0.82	4.36	23.72
		(0.92)	(0.19)	(0.21)	(1.62)	(2.08)
CA	2	16.87	1.40	0.83	5.29	24.40
		(0.82)	(0.19)	(0.21)	(1.16)	(1.60)
CA	3	20.66	1.84	0.91	9.57	32.98
		(0.77)	(0.23)	(0.21)	(1.15)	(1.60)
FRCC	1	25.76	2.51	14.48	2.78	45.53
		(0.54)	(0.17)	(1.37)	(0.09)	(1.82)
FRCC	2	25.82	3.26	15.57	2.72	47.37
		(0.53)	(0.19)	(1.34)	(0.09)	(1.80)
FRCC	3	26.39	4.88	16.49	2.91	50.67
		(0.53)	(0.20)	(1.43)	(0.09)	(1.86)
MRO	1	58.85	9.78	43.66	2.08	114.38
		(3.71)	(0.71)	(2.76)	(0.14)	(6.89)
MRO	2	52.45	9.77	43.55	1.88	107.65
		(3.67)	(0.78)	(2.92)	(0.14)	(7.08)
MRO	3	42.19	8.71	31.75	1.69	84.33
		(2.84)	(0.64)	(2.65)	(0.13)	(5.80)
NPCC	1	15.25	1.58	0.72	2.80	20.34
		(0.63)	(0.22)	(1.43)	(0.13)	(2.03)
NPCC	2	17.52	2.30	4.40	2.97	27.19
		(0.62)	(0.24)	(1.39)	(0.13)	(1.96)
NPCC	3	20.59	5.32	11.48	3.67	41.06
		(0.62)	(0.34)	(1.44)	(0.15)	(2.14)
RFC	1	29.88	5.99	44.04	3.99	83.90
		(0.77)	(0.25)	(2.52)	(0.15)	(3.23)
RFC	2	29.39	6.21	44.55	4.24	84.40
		(0.92)	(0.26)	(2.69)	(0.19)	(3.58)
RFC	3	27.01	6.69	39.44	4.97	78.11
		(0.65)	(0.26)	(2.28)	(0.16)	(2.81)
SERC	1	27.12	3.92	22.53	2.00	55.56
		(0.65)	(0.22)	(1.70)	(0.12)	(2.28)
SERC	2	28.41	4.42	26.34	2.20	61.38
		(0.64)	(0.23)	(1.78)	(0.12)	(2.40)
SERC	3	28.67	5.65	24.23	2.24	60.78
		(0.67)	(0.38)	(2.05)	(0.13)	(2.89)
SPP	1	27.72	4.30	18.05	1.44	51.51
		(1.34)	(0.25)	(1.51)	(0.10)	(2.71)
SPP	2	25.80	4.64	15.57	1.40	47.42
		(1.41)	(0.27)	(1.47)	(0.10)	(2.76)
SPP	3	22.46	5.01	12.34	1.33	41.14
		(1.06)	(0.25)	(1.37)	(0.10)	(2.29)
TRE	1	26.86	2.14	17.23	2.03	48.25
		(1.26)	(0.12)	(1.15)	(0.11)	(2.45)
TRE	2	24.88	1.84	13.38	1.84	41.95
		(0.96)	(0.11)	(0.95)	(0.08)	(1.92)
TRE	3	24.69	3.14	6.38	1.60	35.82
		(0.71)	(0.12)	(0.69)	(0.06)	(1.29)
WECC	1	26.48	4.94	4.80	1.23	37.45
		(0.91)	(0.20)	(0.26)	(0.72)	(1.46)
WECC	2	22.82	4.13	3.97	1.14	32.06
		(1.04)	(0.23)	(0.32)	(0.73)	(1.65)
WECC	3	18.82	2.91	2.61	0.98	25.32
		(0.96)	(0.22)	(0.28)	(0.73)	(1.53)

Table A1: Marginal External Costs by Region and Load Tercile we had assembled these data on renewables we conducted a sensitivity analysis by subtracting from hourly total load to get load net of renewables (i.e. "net load"). We then repeated our regression analysis using net load instead of load. Reassuringly this did not meaningfully alter our estimates of marginal dollar per kWh damages, so the analysis presented here just uses load as the independent variable in all regressions.

As of 2021, it appears that the most common estimates of the SCC, around \$50/ton, may be revised upward significantly as our understanding of climate change continues to evolve. For comparison purposes, we have recalculated externalities and the gap between price and SMC based on a \$100/ton SCC. Figure A3 presents the gap under this assumption. Though it obviously shifts the colors to be redder than in figure 9, the change in the bluer areas is more modest, because these are areas with relatively low marginal CO_2 emissions to begin with.

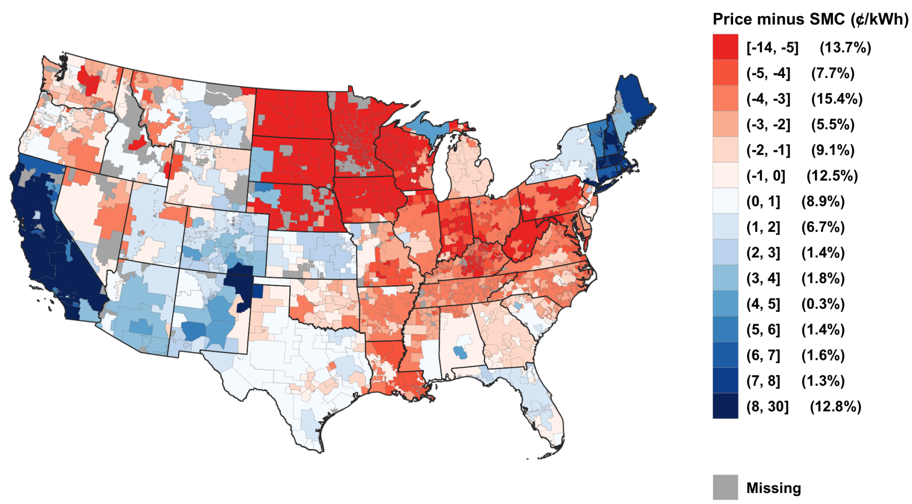

Figure A3: Marginal Price minus Average SMC per kWh with SCC=\$100/ton

6. Hourly Load Shapes

Residential customer demand for electricity is not constant, nor is the deviation between residential retail price and the social marginal costs of supplying

negligible. The NPCC region used ISONE for wind (2014-2015) and combined NYISO/ISONE for wind (2016) but solar assumed was negligible. The non-CA-WECC region used combined CAISO/MISO for wind and combined CAISO/SPP for solar.

electricity. In fact, it is likely the case that these will sometimes be strongly correlated (*e.g.,* periods of peak wholesale electricity prices tend to coincide with peak residential electricity demand). It is therefore important to be able to determine how annual residential sales are distributed across the hours in our analysis period. The ideal dataset for this would likely be some form of hourly metered consumption data for the universe of residential households in the US. Clearly such a dataset does not exist - customers' meter data are confidential and held by their individual utility, and many residential households still do not even have meters that can record this information at an hourly frequency. To tackle this challenge our preferred approach involved using hourly load data from a selection of ISO zonal nodes and planning areas. These data were used to represent the shape of hourly residential load profiles at the ZIP Code level up to a scale factor, and then once again we used our dataset of ZIP Code service territory lookups to match these up to utilities.

To do this, we again used the ISO zonal data from SNL Financial (SNL Financial 2017b). Unlike pricing nodes, load is only available for a limited number of zonal nodes, and is not available for the many thousands of individual load nodes where LMPs are calculated. Fortunately many of these are the same nodes that we already chose to use in our selection of LMPs. In total this gave us load data for 66 ISO zonal nodes. The FERC-714 survey was then used to supplement this with additional hourly load data for planning areas. All series were then normalized to hourly shares of annual load by dividing each hour by the annual total for that ISO zone or planning area.[50] On average this would mean the load share in a single hour should be 1/8760, or 0.0114%. Above average hours (*e.g.,* 6pm on weekdays) should be above this and below average hours (*e.g.,* 3am on weekends) should be below this. Normalizing the data in this way helped account for the fact that ISOs and planning areas differ massively in size (as measured by total load) and is also consistent with our intended use of these data to apportion annual kWh sales across each hour of the year. As with the private marginal cost data, these shares of annual load needed to be assigned to the utility-state entries in our reformatted retail rates dataset. We employ the same approach as for the private marginal costs analysis. This involves assigning each ISO zone or planning area series to a central coordinate (SNL Financial 2017a, Department of Homeland Security 2017b). These spatial coordinates can be seen in Figure A4.[51] We then calculated load shares for each ZIP Code using the inverse distance-weighted averages of the three nearest load points.

The end product of the residential load profile data collection process was a

[50] There were some series with data missing for some hours of the year. If an ISO zone or FERC balancing authority had more than 10% of the hours in a year missing, shares were not calculated and that series was dropped. The concern here was that shares calculated using a subset of the hours in the year may not produce accurate shares if the hours for which there were missing data were not representative of all hours. This only led to data for 3 planning areas being dropped.

[51] The figure depicts selected load points for ISO-NE (orange), NYISO (purple), PJM (red), MISO (blue), SPP (brown), ERCOT (green), CAISO (pink) and FERC planning areas (grey)

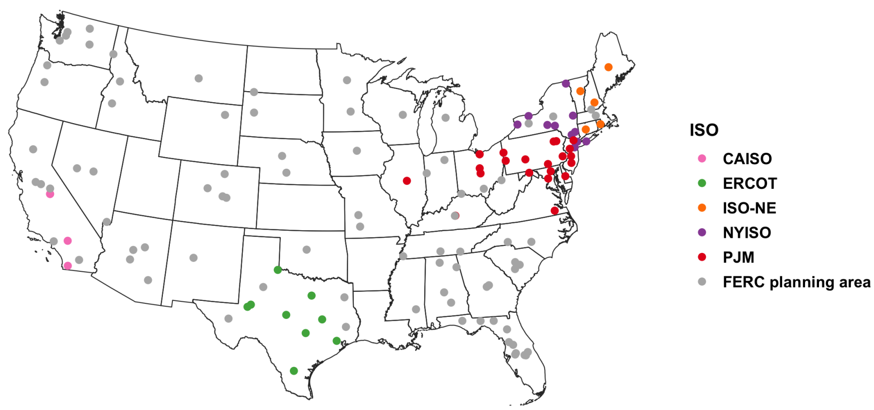

Figure A4: Locations of ISO load zones and load Planning Areas in 2015

dataset of estimates of hourly shares of annual residential electricity demand for each US ZIP Code. These data were then merged with the reformatted retail rates data using the information on the ZIP Codes served by each local distribution company. Where a utility served multiple ZIP Codes in a given state, we again weighted the ZIP Code values for the load shares by the total population of each ZIP Code. A final adjustment was made to ensure that each of the newly created series correctly summed to one over the year.

It is important to note that our preferred approach of using system load profiles as a proxy for residential load profiles has a clear drawback in that it likely underestimates the peakiness of residential load. This is because system load is made up of all demand for electricity from residential, commercial and industrial customers. Differences in the load profiles of residential versus commercial and industrial customers mean that the combination of these three customer classes tends to lead to a smoother total system load profile. It is true that residential customers make up the largest customer class, accounting for over 37% of all kWh sales in 2015, so are an important driver of total system load. Even so, where commercial and industrial customers have significantly different load profiles to residential customers and where they make up a significant portion of total load, our hourly allocation of residential load will almost certainly be biased towards less volatility.

To test the robustness of using these system load profiles as a proxy for residential load profiles, we conducted a sensitivity analysis using an alternative source of residential load profile data. For this, we collected modeled residential load profiles produced by NREL (National Renewable Energy Laboratory 2013). This dataset uses an engineering model to estimate hourly residential electricity demand profiles for a set of representative residential households at different locations throughout the US. To construct the dataset NREL classified the US into five climate zones and made assumptions about building characteristics that varied by climate zone (*e.g.,* source of space heating, presence of air conditioning, square footage, construction materials etc.). NREL also made additional assumptions about operational conditions, such as occupancy rates and weather. An hourly weather profile was used based on NREL's "typical meteorological year" (TMY3) dataset. This provides hourly averages for a range of weather variables (*e.g.,* temperature, humidity, precipitation etc.) based on up to 30 years of historical data from 1976 to 2005. The engineering model then takes these assumptions and weather data and estimates a residential electricity demand profile at over 1,400 TMY3 locations throughout the US (National Renewable Energy Laboratory 2008). The clear advantage of the NREL dataset is that it is a more explicit measure of fluctuations in *residential* load, rather than system load. The main disadvantages are twofold. First, the dataset is comprised of estimates of residential load based on a 2008 engineering model that necessarily makes strong assumptions about building performance, customer behavior and the nature of the housing stock. As such this may be a poor proxy for the performance of the actual housing stock in our analysis period. Second, the dataset is produced using averaged weather data from well before our chosen period of analysis. As such the weather profile used may differ substantially from the actual weather that prevailed during our analysis period.

To conduct our sensitivity analysis we carried out the same processing steps described earlier to get a second set of estimates of residential load profiles for each US ZIP Code, in this case based on the NREL simulation data. To assess the actual performance of the load profiles based on the NREL dataset relative to our load profiles based on observed system load we compared both approaches against the very few datasets of actual metered residential load we were able to find. In general we found that the load profiles based on system load understated the peakiness of residential load and the load profiles based on the NREL modeling data overstated the peakiness of residential load. We also found some limited evidence that the profiles based on system load were more strongly correlated with the actual residential load data. Finally, we conducted the entire analysis using both approaches to estimating the residential load profile to see how this would move the results. We found that the choice of residential load profile had a very small impact on the final results (*e.g.,* on the extent of estimated deadweight loss) so we have opted throughout to use the approach based on system load.

7. Distribution Losses

Our estimation of private and external marginal costs gives the marginal cost of electricity delivered in the high-voltage transmission system. However, our analysis is concerned with the marginal costs of serving residential customers. It is therefore important that we account for losses incurred as power is carried through the low-voltage distribution system to residential households. We estimate average annual residential distribution losses for each local distribution company using data in the EIA-861 survey. Unfortunately, the only data on losses that are available report total losses for a given utility across all types of customers (*i.e.*, residential, commercial and industrial). This is problematic because losses to residential customers are likely higher than for any other customer type. This is because residential customers are located at the furthest ends of the distribution network at the lowest voltage levels. Industrial customers, on the other hand, likely have the lowest losses because they are connected to more centralized portions of the distribution network at higher voltage levels. Sometimes industrial customers are even connected directly to the transmission network, so incur zero distribution losses. A second issue with these data on total losses is that they are not exclusively distribution system losses; some utilities own and operate both transmission and distribution system infrastructure, so their reported losses cover both these parts of the power grid.

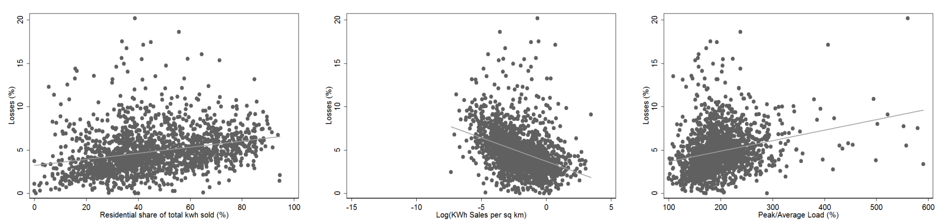

(a) Residential share (b) Density of load (c) Peak/average load

Figure A5: Losses plotted against three key covariates

To address these shortcomings, we estimate average annual residential distribution losses. We compile data on the following variables for each local distribution company, i: total losses in kWh, L_i; total sales in kWh, Q_i, sales for residential customers in kWh, $Qres_i$, commercial customers, $Qcom_i$, and industrial customers, $Qind_i$; the density of customer load, D_i, as measured by the log of total kWh sales divided by the service territory area in square kilometers; the share of distribution circuits with voltage optimization, $VoltOpt_i$, and the ratio of peak load to average load, P_i.[52] We also created dummies for each state, $State_{si}$, utility

[52]The log of the density of kWh sales was used as it provided a much better fit, likely due to the very

type, $UtilityType_{ui}$, and a dummy variable representing whether the utility is involved in electricity transmission, $Transmission_i$.[53] Table A2 presents summary statistics on these variables.

	Mean	StDv	Min	Max	N
Avg. Proportion Total Losses	0.05	0.03	0.00	0.27	5088
Share of Sales (Residential)	0.46	0.21	0.00	1.00	5796
Share of Sales (Commercial)	0.30	0.17	0.00	1.00	5796
Share of Sales (Industrial)	0.24	0.23	0.00	1.00	5796
Log(Sales per sq. km)	-2.29	2.02	-12.73	3.44	5791
Share of Circuits w. Volt. Optim.	0.23	0.39	0.00	1.00	5761
Ratio of Peak to Average Load	1.97	0.49	1.00	5.90	5184
Transmission	0.17	0.38	0.00	1.00	5274

5001 out of 5796 observations have complete information (observations are utility-state-years)

Table A2: Summary Statistics of Variables in the Distribution Losses Regression

The equation for annual losses of a utility could be written as

$$
\begin{aligned}
(1) \qquad L_i &= \alpha_0 Qtot_i + \alpha_1 Qres_i + \alpha_2 Qcom_i + \alpha_3 Qtot_i Density_i \\
&+ \alpha_4 Qtot_i VoltOpt_i + \alpha_5 Qtot_i (Qpeak/Qavg_i) \\
&+ \alpha_6 Qtot_i Transmission_i \\
&+ \sum_{u=1}^{U} \gamma_u UtilityType_{ui} Qtot_i + \sum_{s=1}^{S} \beta_s State_{si} Qtot_i + \epsilon_i
\end{aligned}
$$

where the Qs are total, residential, and commercial electricity delivered, $Density$ is $\log(Qtot/area)$, $VoltOpt$ is the share of circuits with voltage optimization equipment, $Qpeak/Qavg_i$ is the ratio of the utility's peak to average load, and $Transmission_i$ is an indicator that the utility also owns transmission lines (and reported losses include losses from transmission). The equation includes fixed effects for type of utility (investor-owned, municipal, cooperative, etc.) and state. The coefficient α_0 alone would represent the losses associated with an additional unit of electricity delivered to an industrial customer. The derivative of equation (1) with respect to $Qres$ (recognizing that $dQtot/dQres = 1$) would then give the

large range of density values in the data.

[53]All utilities in our sample were involved in distribution. We also chose to aggregate the State, Federal and Political Subdivision utility types into a single "Other Public" category as some of these classifications only contained a very small number of observations. The Retail Power Marketer utility type was also not relevant for this analysis because we are focused on local distribution companies. This left us with four utility type categories for our distribution losses analysis: Investor Owned, Cooperative, Municipal, Other Public.

change in annual losses from delivering one additional unit of electricity.

$$
\begin{aligned}
(2) \qquad dL_i/dQres_i &= \alpha_0 + \alpha_1 + \alpha_3 Density_i \\
&+ \alpha_4 VoltOpt_i + \alpha_5(Qpeak/Qavg_i) \\
&+ \alpha_6 Transmission_i \\
&+ \sum_{u=1}^{U} \gamma_u UtilityType_{ui} + \sum_{s=1}^{S} \beta_s State_{si} + \epsilon_i
\end{aligned}
$$

Equation (1), however, would be highly heteroskedastic in the form shown, so we normalize (1) by total quantity and estimate

$$
\begin{aligned}
(3) \qquad Lavg_i &= \alpha_0 + \alpha_1 Qres_i/Qtot_i + \alpha_2 Qcom_i/Qtot_i + \alpha_3 Density_i \\
&+ \alpha_4 VoltOpt_i + \alpha_5(Qpeak/Qavg_i) \\
&+ \alpha_6 Qtot_i Transmission_i \\
&+ \sum_{u=1}^{U} \gamma_u UtilityType_{ui} + \sum_{s=1}^{S} \beta_s State_{si} + \epsilon_i
\end{aligned}
$$

where the interpretation of the coefficients is the same as in (1) and (2).

We estimate (3) on annual observations for the 1669 distribution utilities for which these data are available for the years 2014 through 2016. A few of the utilities are not in the data for all three years, so the total number of observations is 5001. The results, presented in table A3, suggest that distribution to residential customers exhibits about 3 percentage point higher losses than to industrial customers, and that higher geographic density of customers significantly lowers distribution losses. Voltage optimization also lowers distribution losses, while more volatile load raises distribution losses for a given average level of load. Utilities that also own transmission may exhibit somewhat higher losses, though that effect is not estimated precisely.[54]

From this regression, we then impute average distribution losses for residential customers of all utilities in the dataset by calculating the predicted value of $Lavg_i$ with $Qres_i/Qtot_i = 1$, $Qcom_i/Qtot_i = 0$ and $Transmission_i = 0$.[55] The vast majority of our estimates fall between 4% and 8%, as can be seen in the histogram below.

Clearly, this is an imperfect approximation to average distribution losses for residential customers. It assumes implicitly that the relative losses of residential versus commercial and industrial customers are the same for all utilities. Fur-

[54] Even though this is a (short) panel, it is worth noting that identification of the parameters in this regression comes almost entirely from the cross-sectional variation. If one includes utility fixed effects, only the density effect remains statistically significant.

[55] We predict losses for all utilities in the data set. For those for which some of the right-hand side variables are not available, we use the average value of the variable from the 1669 utilities in the regression.

	$L_{\text{i}}/Qtot_{\text{i}}$
Share of Sales (Residential)	0.0284^{***}
	(0.0064)
Share of Sales (Commercial)	0.0059^{*}
	(0.0034)
Log(Sales per sq. km)	-0.0065^{***}
	(0.0006)
Share of Circuits w. Volt. Optim.	-0.0019^{*}
	(0.0010)
Ratio of Peak to Average Load	0.0076^{***}
	(0.0020)
Transmission	0.0022
	(0.0015)
R^2	0.2916

Standard errors in parentheses

N=5001 (observations are utility-state-years)

Dependent Variable: Avg. Proportion Total Losses

Fixed Effects: State, Utility Type and Year

Cluster Variable: State

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3: Estimates of Average Distribution Losses

thermore, we have no information on the extent to which voltage optimization or variation in hourly sales relates to residential circuits. Without making very strong assumptions about the correlates of residential losses, it is unclear how to improve on this estimate.

Once we had estimates for average annual distribution losses for residential customers, the final step was to convert these to marginal losses and account for how losses vary throughout the year. As explained in the paper, we use the common characterization that 25% of losses are independent of flow on the line – and therefore not associated with any marginal losses from increased consumption – and the engineering result that the other 75% resistive losses increase with the square of flow on the line.[56]

We adapt the approach taken in Borenstein (2008) and assume that utility i's losses in each hour are:

$$L_{it} = \alpha_{i1} + \alpha_{i2} Q_{it}^2 \tag{4}$$

[56]See Lazar and Baldwin (1997) and Southern California Edison's methodology for calculating Distribution Loss Factors, as set out in filings to the California Public Utilities Commission (California Public Utilities Commission 1997).

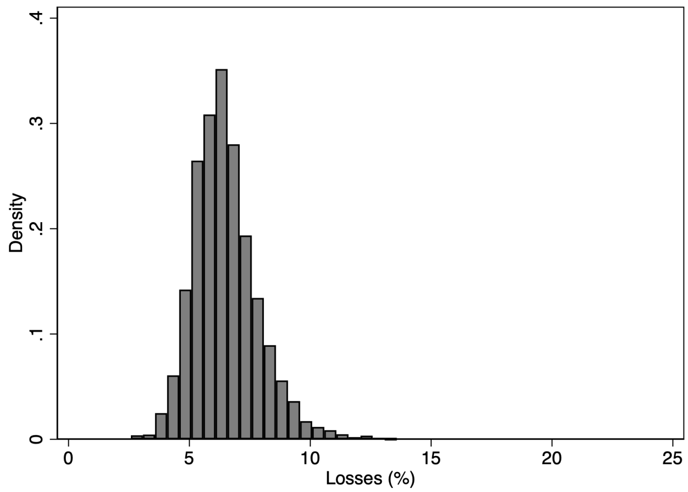

Figure A6: Histogram of Predicted Average Residential Distribution Losses

We have already estimated average annual losses for each local distribution company, which we call γ_i. Because the α terms are constant across all hours we can convert the equation to annual sums and substitute for L_{it}. If we also assume that the static no-load losses, as represented by the α_{i1} term, constitute a quarter of a utility's total losses, we can then solve for α_2 for each local distribution company.

$$\sum_{t=1}^{T} L_{it} = \gamma_i \sum_{t=1}^{T} Q_{it} = \alpha_{i1} + \alpha_{i2} \sum_{t=1}^{T} Q_{it}^2 \iff \alpha_{i2} = (1 - 0.25)\gamma_i \frac{\sum_{t=1}^{T} Q_{it}}{\sum_{t=1}^{T} Q_{it}^2} \tag{5}$$

Finally, our interest is in marginal losses so we take the derivative of our original losses expression such that:

$$\frac{dL_{it}}{dQ_{it}} = 2\alpha_{i2} Q_{it} \tag{6}$$

Thus, equation (6) produces our estimate of marginal line losses as a fraction of energy that enters the distribution system of utility i in hour t. For each hour, private and external marginal costs were then scaled up by $\frac{1}{1-dL_{it}/dQ_{it}}$ to give our complete estimate of the social marginal cost of residential electricity consumption.

8. Calculation of Deadweight Loss in the Short Run and the Long Run

To evaluate DWL while recognizing that short-run and long-run demand elasticities for electricity may differ substantially, we consider a two-stage consumer decision process. In the first stage, the consumer chooses the devices to buy, energy efficiency investments to make, and household habits for using the devices, all based on the average price they expect to face.[57] We refer to all of these choices collectively as the consumer's investment in devices. In the second stage, the customer uses the devices, responding to hourly prices, which will generally deviate from the average price.

Figure A7 illustrates the short-run demand functions that a household might have during a specific hour for the electricity to use individual devices 1, 2, 3, and 4. The household makes the long-run investment decisions – adding, removing or shifting a short-run demand function for a type of device – by comparing the price, energy efficiency and other device attributes with the consumer surplus that the household expects to receive by owning it. The consumer's gross consumer surplus from a device is calculated as the area under the short-run demand out to the quantity consumed, aggregated over the life of the device. We assume that the long-run demand elasticity reflects the household's optimized response to different average prices through the device investments they make and the extent to which they use them on average. The short-run demand elasticity reflects the household's change in hourly usage in response to changing hourly prices.

Efficient purchase and hourly usage of devices results when the hourly price is set equal to hourly social marginal cost, which implies that the average price is the average social marginal cost. In figure A7, assume that when that occurs the household purchases devices 1, 3, and 4, and has a household short-run demand function of D_{134} for this specific hour. The household's demand function varies hour to hour, but in the short run it will always reflect owning devices 1, 3, and 4.

The change in total surplus from a change in pricing regime is equal to:

$$\begin{aligned}\Delta\ Total\ Surplus = {} & \Delta\ Gross\ Consumer\ Surplus \\ & - \Delta\ Variable\ Costs - \Delta\ Investment\ Costs\end{aligned}$$

where Δ *Gross Consumer Surplus* occurs as a result of changing usage of current devices in the short run and changing the household's device investments in the long run. Measurement of the first two terms is fairly straightforward given assumptions about electricity price and costs, along with long-run and short-run demand elasticities. Recognizing that consumers will make investments only

[57]We implement this using the quantity-weighted average price for the utility. A sophisticated buyer facing time-varying prices could do a more granular calculation, taking into account the timing of their expected device usage and ability to shift consumption in response to price differences, but we abstract from this for simplicity, and because none of the long-run demand elasticity estimates in the literature reflect such optimization.

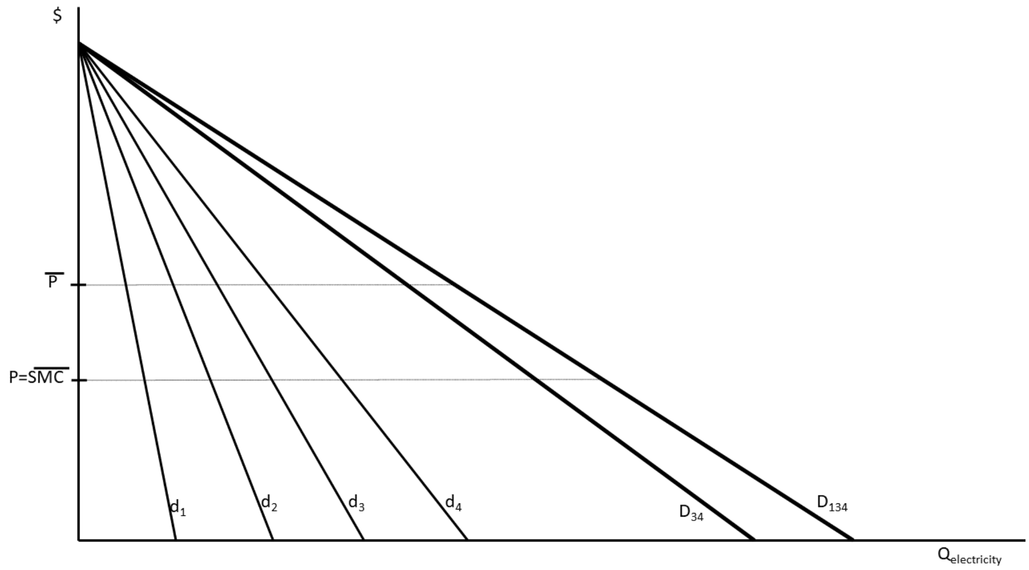

Figure A7: Illustration of Consumer Choices of Electrical Devices and Their Usage

if the additional consumer surplus from use is at least as large as the cost of the investment implies that long-run demand can be used to analyze the cost of adopting a device. For example, if this consumer chooses not to adopt device 1 when $P = \overline{P}$, but does choose to adopt it when $P = \overline{SMC}$, then for device 1

$$\sum_{h \in H} CS_{P=\overline{P}} < Cost\ of\ Adoption < \sum_{h \in H} CS_{P=\overline{SMC}}$$

where the appropriately discounted consumer surpluses are summed over all hours in which the device will be used. Adopting device 1 would then change this consumer's short run demand function in the illustrated hour from D_{34} to D_{134}.[58] Applying this approach to a continuum of devices, the marginal consumer surplus associated with the long-run response to changes in the average electricity price traces out the cost of incremental investments.

Though figure A7 illustrates the analysis with linear demand curves, we implement the calculations using modified constant-elasticity demand. Linear demand has the unfortunate property that an elasticity at $\overline{P}$ that is in the range of common empirical estimates yields implausibly low choke off prices and potentially understates price responsiveness at extremely low prices. On the other hand,

[58]Similarly, the fact that this consumer does not adopt device 2 at a price of $P = \overline{SMC}$ implies that the sum of discounted future consumer surpluses for device 2 at $P = \overline{SMC}$ does not exceed its cost of adoption.

constant-elasticity demand yields implausibly high, or infinite, quantities as price goes to zero or below, and unbounded willingness to pay for small quantities. In an attempt to model a more realistic demand setting, we take demand to be constant elasticity up to a price of \$2/kWh, about 20 times the average retail price, and then horizontal at \$2/kWh down to zero quantity. We also modify the function at low prices, imposing a quantity cut off at twice the quantity demanded when $P = \$0.05/kWh$ for each hour, which is about half the sample average retail price. In every hour, the demand function is determined by the assumed elasticity and the observed point $(\overline{P}, \overline{Q})$, through which both the short-run and the long-run demand curves are assumed to run.

DWL calculations are also potentially a function of how SMC changes with quantity. Following the discussion of externalities above, we have estimated the slope of SMC at the NERC region level. For each region, we regressed hourly average region SMC on hourly region quantity and month-of-sample dummies to capture variation in availability of wind, solar, and hydroelectric supply. The potential endogeneity of quantity is not a significant concern here, because virtually all customers face prices that are invariant to market conditions, as discussed earlier.[59] For eight of the nine regions, the estimated slope of the SMC(Q) function is positive and statistically significant, but economically extremely small, particularly compared to the slope of the demand function. For one region, the estimated slope is still positive, but even smaller and not statistically significant. For the utilities in our sample, a one standard-deviation increase in the quantity supplied is estimated to increase SMC by a quantity-weighted average of \$0.000006. For 99% of utilities, the change is less than \$0.0015, and for no utility is the change larger than \$0.005. Thus, for computational simplicity, we assume that SMC is constant over the range of the quantity changes considered.[60]

Table A4 presents the average DWL per normalized quantity for the U.S., the annual total DWL per customer, and the share of DWL attributable to the long-run mispricing ($\overline{P} \neq \overline{SMC}$) for combinations of short-run and long-run elasticities. The table suggests that the change in each type of deadweight loss is approximately linear in the elasticity, and that the basic finding that average DWL is small – compared to SMC or retail price – is robust to even fairly large elasticities.

[59]There is a second concern that high demand hours might be associated with larger or smaller production from wind and solar – likely larger for solar and smaller for wind – therefore shifting the SMC function. As discussed in appendix section 5 on external marginal costs, however, we did not find that accounting for supply from intermittent resources meaningfully changed the analysis of external marginal costs, so we did not make further adjustments to this regression.

[60]Incorporating non-constant SMC implies that there is no closed-formed solution for the intersection of SMC and the constant-elasticity demand functions, so requires an approximation to the intersection. When we solved for the intersection using a linear approximation of the constant-elasticity short-run demand functions around the value of SMC at $\overline{P}$, nearly all of the quantity changes in the DWL calculations were within 1% of the changes that result from assuming a constant SMC.

			LR Elasticity					
			-0.1	-0.2	-0.3	-0.5	-0.7	-0.9
SR Elasticity	-0.1	DWL_{total} per kWh (¢)	0.131	0.207	0.283	0.435	0.587	0.738
		DWL_{total} per customer ($)	14.39	23.04	32.08	51.41	72.65	96.10
		DWL_{LR}/DWL_{total}	0.565	0.721	0.794	0.863	0.896	0.915
	-0.2	DWL_{total} per kWh (¢)	0.180	0.256	0.332	0.486	0.639	0.790
		DWL_{total} per customer ($)	19.78	28.56	37.74	57.41	79.06	103.00
		DWL_{LR}/DWL_{total}	0.412	0.582	0.675	0.773	0.823	0.853
	-0.3	DWL_{total} per kWh (¢)	0.226	0.303	0.379	0.534	0.688	0.840
		DWL_{total} per customer ($)	24.81	33.72	43.05	63.04	85.08	109.49
		DWL_{LR}/DWL_{total}	0.328	0.493	0.591	0.704	0.765	0.803
	-0.5	DWL_{total} per kWh (¢)	0.311	0.388	0.465	0.621	0.778	0.932
		DWL_{total} per customer ($)	34.06	343.22	52.81	73.42	96.19	121.48
		DWL_{LR}/DWL_{total}	0.239	0.384	0.482	0.604	0.676	0.724
	-0.7	DWL_{total} per kWh (¢)	0.387	0.465	0.544	0.702	0.860	1.017
		DWL_{total} per customer ($)	42.46	51.85	61.69	82.89	106.36	132.47
		DWL_{LR}/DWL_{total}	0.192	0.320	0.413	0.535	0.612	0.664
	-0.9	DWL_{total} per kWh (¢)	0.459	0.537	0.617	0.776	0.937	1.096
		DWL_{total} per customer ($)	50.28	59.89	69.98	91.73	115.87	142.79
		DWL_{LR}/DWL_{total}	0.162	0.277	0.364	0.484	0.562	0.616

Note: For each elasticity pair, the DWL_{total}, DWL_{LR}, number of customers, and normalized quantities are summed across the 6,215 utility-state-years, then the relevant sums are divided to obtain each statistic.

Table A4: Average Deadweight Loss Measures Under Alternative Elasticity Assumptions

*

REFERENCES

Allcott, Hunt and Michael Greenstone, “Is there an energy efficiency gap?,” *Journal of Economic Perspectives*, 2012, *26* (1), 3–28.

_ **and** _ , “Measuring the Welfare Effects of Residential Energy Efficiency Programs,” Working Paper #23386, National Bureau of Economic Research May 2017.

Alvarez, Ramón A, Daniel Zavala-Araiza, David R Lyon, David T Allen, Zachary R Barkley, Adam R Brandt, Kenneth J Davis, Scott C Herndon, Daniel J Jacob, Anna Karion et al., “Assessment of methane emissions from the US oil and gas supply chain,” *Science*, 2018, *361* (6398), 186–188.

Auffhammer, Maximilian, Carl Blumstein, and Meredith Fowlie, “Demand-side management and energy efficiency revisited,” *The Energy Journal*, 2008, pp. 91–104.

Boiteux, M., "Peak-Load Pricing," *The Journal of Business*, 1960, *33* (2), 157–179.

Boiteux, M, "On the management of public monopolies subject to budgetary constraints," *Journal of Economic Theory*, 1971, *3* (3), 219 – 240.

Boomhower, Judson P and Lucas W Davis, "Do Energy Efficiency Investments Deliver at the Right Time?," Technical Report, National Bureau of Economic Research 2017.

Borenstein, Severin, "The Long-Run Efficiency of Real-Time Electricity Pricing," *The Energy Journal*, 2005, *26* (3), 93–116.

_ , "The Market Value and Cost of Solar Photovoltaic Electricity Production," Center for the Study of Energy Markets Working Paper #176, University of California Energy Institute, 2008.

_ , "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing," Center for the Study of Energy Markets Working Paper #195, University of California Energy Institute, 2009.

_ , "The Private and Public Economics of Renewable Electricity Generation," *Journal of Economic Perspectives*, February 2012, *26* (1), 67–92.

_ , "The economics of fixed cost recovery by utilities," *The Electricity Journal*, 2016, *29* (7), 5–12.

_ , "Private net benefits of residential solar PV: the role of electricity tariffs, tax incentives, and rebates," *Journal of the Association of Environmental and Resource Economists*, 2017, *4* (S1), S85–S122.

_ **and James Bushnell**, "The US electricity industry after 20 years of restructuring," *Annu. Rev. Econ.*, 2015, *7* (1), 437–463.

_ **and** _ , "Headwinds and Tailwinds: Implications of Inefficient Retail Energy Pricing for Energy Substitution," Energy Institute at Haas Working Paper #318, Energy Institute at Haas, UC Berkeley 2021.

_ **and Lucas W. Davis**, "The equity and efficiency of two-part tariffs in US natural gas markets," *The Journal of Law and Economics*, 2012, *55* (1), 75–128.

_ **and Stephen Holland**, "On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices," *RAND Journal of Economics*, 2005, *36* (Autumn), 469–493.

_ **, Meredith Fowlie, and James Sallee**, "Designing Electricity Rates for An Equitable Energy Transition," Energy Institute at Haas Working Paper #314, Energy Institute at Haas, UC Berkeley 2021.

Braeutigam, Ronald R., "Optimal policies for natural monopolies," in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization*, 1 ed., Vol. 2, Elsevier, 1989, chapter 23, pp. 1289–1346.

Brown, Stephen J. and David Sumner Sibley, *The Theory of Public Utility Pricing*, Cambridge University Press, 1986.

Buchanan, James, "External Diseconomies, Corrective Taxes, and Market Structure," *American Economic Review*, 1969, *59* (1), 174–77.

Bushnell, James B., Stephen P. Holland, Jonathan E. Hughes, and Christopher R. Knittel, "Strategic Policy Choice in State-Level Regulation: The EPA's Clean Power Plan," *American Economic Journal: Economic Policy*, May 2017, *9* (2), 57–90.

California Independent System Operator, "Annual Report on Market Issues and Performance," 2016.

_ , "Renewables Watch Reports," 2018.

California Public Utilities Commission, "Distribution Loss Factors, Supplement to the July 25th 1997 Workshop Report on Retail Settlement and Information Flows," 1997.

_ , "Resource Adequacy Report," 2015.

Callaway, Duncan S, Meredith Fowlie, and Gavin McCormick, "Location, location, location: The variable value of renewable energy and demand-side efficiency resources," *Journal of the Association of Environmental and Resource Economists*, 2018, *5* (1), 39–75.

Clay, Karen, Akshaya Jha, Nicholas Z Muller, and Randy Walsh, "The External Costs of Shipping Petroleum Products by Pipeline and Rail: Evidence of Shipments of Crude Oil from North Dakota," *Energy Journal*, 2018, p. forthcoming.

Costello, Kenneth W. and Ross C. Hemphill, "Electric Utilities Death Spiral: Hyperbole or Reality?," *The Electricity Journal*, 2014, *27* (10), 7–26.

Crew, Michael A. and Paul R. Kleindorfer, "Peak Load Pricing with a Diverse Technology," *Bell Journal of Economics*, 1976, *7* (1), 207–231.

Cullen, Joseph, "Measuring the environmental benefits of wind-generated electricity," *American Economic Journal: Economic Policy*, 2013, *5* (4), 107–33.

Davis, Lucas W and Erich Muehlegger, "Do Americans consume too little natural gas? An empirical test of marginal cost pricing," *The RAND Journal of Economics*, 2010, *41* (4), 791–810.

Department of Homeland Security, "Electric Control Area boundary files," 2017.

_ , "Electric Planning Area boundary files," 2017.

Electric Reliability Council of Texas, "State of the Market Report," 2015.

_ , "Generation by Fuel," 2018.

Energy Information Administration, "Form EIA-861 "Annual Electric Power Industry Report"," 2017.

_ , "Form EIA-861M "Monthly Electric Power Industry Report"," 2017.

_ , "Form EIA-923 "Power Plant Operations Report"," 2018.

Environmental Protection Agency, "Continuous Emissions Monitoring Data," 2018.

_ , "Emissions & Generation Resource Integrated Database," 2018.

_ , "National Emissions Inventory," 2018.

Environmental Systems Research Institute, "USA ZIP Code Areas," 2017.

Federal Energy Regulatory Commission, "Performance Metrics for Regional Transmission Organizations, Independent System Operators and Individual Utilities for the 2010-2014 Reporting Period," August 2016.

_ , "Form-714 - Annual Electric Balancing Authority Area and Planning Area Report," 2017.

Feldstein, Martin S, "Distributional equity and the optimal structure of public prices," *The American Economic Review*, 1972, *62* (1/2), 32–36.

Fischer, Carolyn and Alan Fox, "Comparing policies to combat emissions leakage: Border carbon adjustments versus rebates," *Journal of Environmental Economics and Management*, 2012, *64* (2), 199–216.

Fowlie, Meredith, "Updating the Allocation of Greenhouse Gas Emissions Permits in a Federal Cap-and-Trade Program," in D. Fullerton and C. Wolfram, eds., *The Design and Implementation of US Climate Policy*, University of Chicago Press, 2011, pp. 157–171.

_ **, Lawrence Goulder, Matthew Kotchen et al.**, "An economic perspective on the EPA's Clean Power Plan," *Science*, 2014, *346* (6211), 815–816.

Friedman, Lee S., "Energy utility pricing and customer response: the recent record in California," in R.J. Gilbert, ed., *Regulatory Choices: A Perspective on Developments in Energy Policy*, University of California Press, 1991, chapter 2, pp. 10–62.

Gerarden, Todd D, Richard G Newell, and Robert N Stavins, "Assessing the energy-efficiency gap," *Journal of Economic Literature*, 2017, *55* (4), 1486–1525.

Gerke, Brian F, Giulia Gallo, Sarah Josephine Smith, Jingjing Liu, Shuba V Raghavan, Peter Schwartz, Mary Ann Piette, Rongxin Yin, and Sofia Stensson, "The California Demand Response Potential Study, Phase 3: Final Report on the Shift Resource through 2030," Working Paper DOI: 10.20357/B7MS40, Lawrence Berkeley National Laboratory July 2020.

Gillingham, Kenneth and Karen Palmer, "Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence," *Review of Environmental Economics and Policy*, 2014, *8* (1), 18–38.

Holland, Stephen P, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates, "Are there environmental benefits from driving electric vehicles? The importance of local factors," *American Economic Review*, 2016, *106* (12), 3700–3729.

_ **, Jonathan E Hughes, and Christopher R Knittel**, "Greenhouse gas reductions under low carbon fuel standards?," *American Economic Journal: Economic Policy*, 2009, *1* (1), 106–46.

Independent System Operator New England, "State of the Market Report," 2016.

_ , "Hourly Historical Renewables," 2018.

Ito, Koichiro, "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing," *American Economic Review*, February 2014, *104* (2), 537–63.

_ **and Shuang Zhang**, "Reforming Inefficient Energy Pricing: Evidence from China," Working Paper #26853, National Bureau of Economic Research March 2020.

Jessoe, Katrina and David Rapson, "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use," *American Economic Review*, April 2014, *104* (4), 1417–38.

Jha, Akshaya and Nicholas Z Muller, "The local air pollution cost of coal storage and handling: Evidence from US power plants," *Journal of Environmental Economics and Management*, 2018, *92*, 360–396.

Joskow, Paul L., "Contributions to the Theory of Marginal Cost Pricing," *The Bell Journal of Economics*, 1976, *7* (1), 197–206.

_ **and Catherine D. Wolfram**, "Dynamic Pricing of Electricity," *American Economic Review*, May 2012, *102* (3), 381–85.

Joskow, Paul L and Donald B Marron, "What does a negawatt really cost? Evidence from utility conservation programs," *The Energy Journal*, 1992, pp. 41–74.

Kahn, Alfred Edward, *The economics of regulation: principles and institutions*, Vol. 1, MIT Press, 1988.

Laffont, Jean-Jacques, Patrick Rey, and Jean Tirole, "Network Competition: II. Price Discrimination," *RAND Journal of Economics*, 1998, *29* (1), 38–56.

Lazar, Jim and Xavier Baldwin, "Valuing the Contribution of Energy Efficiency to Avoided Marginal Line Losses and Reserve Requirements," Regulatory Assistance Project discussion paper August 1997.

Lipsey, Richard G and Kelvin Lancaster, "The general theory of second best," *The review of economic studies*, 1956, *24* (1), 11–32.

Littlechild, Stephen C, "Two-part tariffs and consumption externalities," *The Bell Journal of Economics*, 1975, *6* (2), 661–670.

Midwest Independent System Operator, "State of the Market Report," 2015.

_ , "Historical Generation Fuel Mix," 2018.

National Renewable Energy Laboratory, "National Solar Radiation Data Base, 1991- 2005 Update: Typical Meteorological Year 3 Stations Metadata," 2008.

—, "Commercial and Residential Hourly Load Profiles for all Typical Meteorological Year 3 Locations in the United States," 2013.

—, "U.S. Electric Utility Companies and Rates: Look-up by ZIP Code," 2017.

—, "Utility Rate Database," 2017.

New York Independent System Operator, "State of the Market Report," 2016.

—, "Real-Time Fuel Mix," 2018.

Novan, Kevin, "Valuing the wind: renewable energy policies and air pollution avoided," *American Economic Journal: Economic Policy*, 2015, *7* (3), 291–326.

— **and Aaron Smith**, "The Incentive to Overinvest in Energy Efficiency: Evidence From Hourly Smart-Meter Data," *Journal of the Association of Energy and Resource Economists*, 2016, *Forthcoming.*

Oren, Shmuel, Stephen Smith, and Robert Wilson, "Capacity Pricing," *Econometrica*, 1985, *53* (3), 545–566.

Pigou, Arthur C., *The Economics of Welfare*, Palgrave Macmillan, 1920.

PJM Interconnection, "State of the Market Report," 2016.

—, "Data Miner Wind Generation," 2018.

Public Utility Commission of Texas, "Average Annual Rate Comparison for Residential Electric Service," 2017.

—, "Report Cards on Retail Competition and Summary of Market Share Data," 2017.

Puller, Steven L. and Jeremy West, "Efficient Retail Pricing in Electricity and Natural Gas Markets," *The American Economic Review*, 2013, *103* (3), 350–355.

Ramsey, Frank P., "A Contribution to the Theory of Taxation," *The Economic Journal*, 1927, *37* (145), 47–61.

Shin, Jeong-Shik, "Perception of Price When Price Information Is Costly: Evidence from Residential Electricity Demand," *The Review of Economics and Statistics*, 1985, *67* (4), 591–598.

SNL Financial, "ISO LMP Node Coordinates," 2017.

—, "ISO Price and Load Data Repository," 2017.

Southwest Power Pool, "State of the Market Report," 2016.

—, "Generation Mix Historical," 2018.

Steiner, Peter O, "Peak loads and efficient pricing," *The Quarterly Journal of Economics*, 1957, *71* (4), 585–610.

US Census, "Cartographic County Boundary Shapefiles," 2017.

—, "County Population Estimates," 2017.

—, "Zip Code Tabulation Area to County Relationship File," 2017.

— , “Annual Average Bureau of Labor Statistics Consumer Price Index Research Series,” 2018.

Wilson, Robert, *Nonlinear Pricing*, Oxford University Press, 1997.

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

Zivin, Joshua S Graff, Matthew J Kotchen, and Erin T Mansur, “Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies,” *Journal of Economic Behavior & Organization*, 2014, *107*, 248–268.