The Use and Misuse of Average and Marginal Energy Prices: Implications for Climate Policy*

Jonathan T. Hawkins-Pierot
Consumer Financial Protection Bureau

Katherine R. H. Wagner
University of British Columbia

July 8, 2024

Click here for latest draft

Abstract

Many governments and institutions use marginal price signals, such as carbon taxes, to meet emission reduction targets. However, the responses of firms and individuals to these price signals are often calculated using average prices commonly reported in economic data, rather than marginal prices. This paper first documents that the marginal price of electricity paid by manufacturing plants is 50% lower than the average price. To do so, we construct a novel dataset of both average and marginal electricity prices, and the wedge between these, using plant-level microdata from the U.S. Census and utility-level electricity rate schedules for over two hundred utilities. Second, we provide guidance on when average prices are an appropriate proxy for marginal prices by identifying economic and geographic characteristics that predict variation in this wedge. Overall, the magnitude of this wedge suggests that the standard use of average energy prices to calculate responses to carbon taxes may underestimate the energy price increases needed to meet emissions targets by 50%.

*We thank Brian Copeland, Lucas Davis, Jenya Kahn-Lang, Severin Borenstein, Stephanie Weber, and participants at the UC Berkeley Energy Camp for helpful discussions, Andrew Tang and Benji Reade Malagueno for excellent research assistance, and the Cowles Foundation, the Yale Economic Growth Center, and the Stanford Institute for Economic Policy Research for funding. Any views expressed are those of the authors and not those of the U.S. Census Bureau, the Consumer Financial Protection Bureau, or the United States. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2309. (CBDRB-FY21-P2309-R8908)

1 Introduction

Many governments have put policies that change relative energy prices, through carbon pricing or clean energy subsidies, at the center of their plans to meet emission reduction targets. To set a carbon price at the level necessary to meet these targets, policymakers require reliable estimates of how emissions respond to marginal changes in input prices. However, the responses of firms and individuals to these prices are often estimated using average, rather than marginal, prices. Most surveys conducted by governments and research agencies, such as the U.S. Energy Information Administration and by the U.S. Census Bureau, collect data on average energy prices only. Both academic research and government reports flag that this distinction may reduce the usefulness of standard data for economic analysis (Boyd and Lee, 2018; Ganapati et al., 2020). However, the extent to which average and marginal energy prices diverge, and whether this distinction is economically meaningful, is not well-understood.

This paper provides guidance on when it is appropriate for policymakers to make decisions on climate policy based on the data on average energy prices that they collect. The divergence between marginal and average costs is relevant for policy: most academic and policy models use average price elasticities to predict responses to marginal energy price changes arising from counterfactual carbon taxes (Nordhaus and Boyer, 2000). In instances where marginal energy prices are below average prices (due to e.g., transmission costs, bulk discounts, or fixed costs of production), a given carbon tax represents a larger percentage change in marginal cost. But any behavioral response is attributed to a smaller change in average costs, leading to overestimates of energy price elasticities and resulting reductions in energy use. By extension, using these average price elasticities to predict energy use reductions will underestimate the carbon price needed to meet emissions targets.

Our empirical analysis focuses on quantifying the wedge between average and marginal electricity prices in the U.S. manufacturing sector and studying the implications for climate policy. This analysis exploits a panel of plant-level average and marginal electricity prices that we construct using restricted-access microdata from the U.S. Census and utility-level electricity rate schedules for over two hundred utilities. Our primary sources of data from the U.S. Census Bureau are administrative records on annual plant-level inputs and outputs from the Census of Manufactures (CMF) for the universe of U.S. manufacturing plants in 2017, the most recently available Census year. The marginal electricity prices come from utility-level industrial rate data contained in the OpenEI's Utility Rate Database, which we compiled for all utilities surveyed by the U.S. Department of Energy.

Using the combined dataset, we summarize differences in average and marginal electricity prices and discuss if average prices are an appropriate proxy for marginal prices by identifying economic and geographic characteristics that predict variation in this wedge. We examine heterogeneity in the extent to which average prices accurately reflect marginal prices across different geographies, industries, and plant sizes. If fixed costs and demand charges are important, we expect systematic differences between large and small plants, with small plants paying higher per unit energy prices, all else equal. If transmission constraints are important, then we expect heterogeneous prices by geography, with manufacturing plants that are less subject to transmission bottlenecks obtaining lower prices. Overall, we identify variables that are indicative of the useful of average energy costs for policy decisions.

There are two primary findings. First, we document that the marginal price of electricity paid by manufacturing plants is approximately 50% lower than the average price. This pattern is evident in aggregate across all manufacturing establishments and also consistent across eleven distinct industry groups. This result is concerning because it suggests that using average price elasticities to predict responses to marginal price changes may overestimate responsiveness to a carbon price by 50%. Given that average price elasticities are typically used in models to predict the behavioral response to carbon pricing, this result suggests that policymakers should calculate the corresponding change in average prices resulting from a carbon tax, rather than marginal prices, to obtain more accurate predictions of its effects. In addition, this result suggests that inefficient demand charges used by utilities to recover fixed costs are a substantially higher portion of industrial average electricity prices than has been previously recognized (Borenstein, 2016).

The second main result is that electricity use is the largest predictor of the wedge between average and marginal electricity prices. Principled variable selection methods using least absolute shrinkage and selection operator (Lasso) regressions reject the informativeness of many other plant characteristics, including revenue, labor and capital inputs, carbon dioxide intensity of the local electricity grid, among others (Tibshirani, 1996). This result suggests that while caution is warranted in using average price elasticities to interpret responses to marginal price changes, the wedge between average and marginal costs shrinks as electricity use increases. However, we find that average prices are rarely a good proxy for marginal prices.

These results contribute to three main literatures. To the best of our knowledge, this paper provides the first estimate of the wedge between average and marginal electricity prices in the industrial sector. This result has important policy implications because it suggests that the standard use of average energy price elasticities to calculate responses to carbon taxes may lead to policies that are insufficiently stringent to meet emissions targets. Understanding the incentives that industry faces to reduce their energy use is important in its own right because this sector accounts for almost all of the predicted increase in U.S. energy use in the next decade (EPA, 2021). This result contributes to a literature on the responses of households and firms to changes in electricity prices. Existing work provides estimates of average electricity price elasticities for residential customers (Alberini and Filippini, 2011; Deryugina et al., 2020; Fell et al., 2014) and industry (Blonz, 2021; Hawkins-Pierot and Wagner, 2023; Linn, 2008; Paul et al., 2009); these average price elasticities are used to predict responses to carbon pricing. A smaller literature compares elasticities of marginal and average electricity prices for residential customers, suggesting that households may respond to changes in average electricity prices even when the price changes themselves are marginal (Ito, 2014; Shaffer, 2020). Industrial customers are plausibly more attentive to price schedules than residential customers due to their higher sensitivity to prices and the close attention paid to fuel costs (Csereklyei, 2020; Hayes and Hafstead, 2020); however, the magnitude of the difference between average and marginal prices that we find suggests that the distinction may be important.

These results also contribute to the literature on Integrated Assessment Models (IAMs) and their use in policy analysis. IAMs are climate-economy models that incorporate interactions between economic activity and the global carbon cycle, and the constraints imposed by policy. Such models are widely used by national and international government agencies to predict economy-wide responses to proposed climate policies (e.g., carbon taxes), including the U.S. Environmental Protection Agency (EPA) and the Intergovernmental Panel on Climate Change (IPCC) (Nordhaus and Boyer, 2000; Pindyck, 2020). These models are also used in academic and policy research to calculate the social cost of carbon (Carleton and Greenstone, 2021). Our contribution is to highlight the importance of carefully considering the model inputs used to simulate the outcomes of counterfactual climate policy. Specifically, if average price elasticities are used to predict emissions reductions resulting from carbon pricing, then the carbon prices themselves should be measured relative to average costs, rather than marginal ones.

Finally, we contribute to a literature on optimal electricity contract design. This literature studies how electric utilities, which are typically considered natural monopolies, can recover fixed costs while minimizing deadweight loss. Research dating back to Ramsey (1927) proposes price discrimination and nonlinear tariffs as one potential solution, but the actual rates implemented vary widely in their structure and efficiency (Borenstein, 2016; Borenstein and Bushnell, 2022; Puller and West, 2013). This research highlights that "economics does not support the use of demand charges" that bill customers for their hour of peak consumption during a payment cycle and "it is unclear why demand charges still exist" (Borenstein, 2016, p.5, p.10), and yet we find that these charges comprise almost half of average prices. We are not aware of any comprehensive effort to measure the contribution of what are considered inefficient demand charges to overall industrial electricity prices.

The rest of this paper proceeds as follows. Section 2 provides a brief overview of electricity use in U.S. manufacturing to contextualize the analysis and Section 3 describes the data and how we analyze it. Section 4 outlines our main results. Section 5 discusses the implications of our findings for climate policy.

2 Electricity Contracts and Use in U.S. Manufacturing

2.1 Energy Use in Manufacturing

Manufacturing accounts for about one-quarter of both total U.S. energy consumption and total U.S. greenhouse gas emissions. Most manufacturing energy is consumed as electricity: in our plant-level data, electricity expenditures accounted for 75% of total energy expenditures and 95% of thermal energy consumed.[1] Plants' activities determine the amount and type of energy used. Common uses include powering production machinery, fueling boilers, on-site transportation, heating, cooling, and lighting. Improving the energy intensity of these processes often requires replacing energy-inefficient equipment or machinery, which can involve significant disruption to production. Despite these costs and frictions, there is strong evidence that manufacturing plants do respond to changes in electricity prices by changing their operations to become more or less energy intensive (Hawkins-Pierot and Wagner, 2023; Linn, 2008).

2.2 Electricity Contract Structure

The market for industrial and commercial electricity in the U.S. is complex and there are significant differences in state regulation. Borenstein and Bushnell (2015) provides a more detailed overview of

[1] In some industries, fuels such as natural gas are used directly as feedstock materials and not included as energy sources here. In other industries, such as cement, the production processes themselves emit greenhouse gasses.

electricity markets in the U.S. Generally, electric utilities have distinct regional markets that cover one or a few states. The areas served by utilities often but do not always overlap. The median manufacturing plant in our data has a single utility serving their ZIP code, but the mean number of available utilities is approximately three.

Economies of scale in electricity distribution require electric utilities to include several components in electricity prices above the marginal cost of supplying electricity to recover their total costs. Electricity contracts often include fixed charges that are independent of electricity quantity purchased, coincident or usage charges that depend on the total electricity used, and demand charges that depend on the maximum electricity used within each billing period. These charges may also be nonlinear: utilities often charge a lower price for electricity after usage has exceeded some threshold. This contract structure drives a wedge between marginal and average electricity prices.[2]

Industrial and commercial electricity contracts differ from residential contracts in several ways. Importantly, while fixed monthly charges are common for residential customers, industrial and commercial contracts rarely include them because these customers are much more heterogeneous in size. Tiered pricing is also less common for industrial and commercial customers for the same reason. Instead, these contracts typically include demand charges based on the highest consumption hour of the billing cycle, typically one month. These one-time charges are based on the maximum amount of power required and are in addition to the volumetric charges determined by total consumption over the billing cycle.[3] More aggregate data suggests that, on average, demand charges typically account for over 30% of an industrial customer's bill, with the remainder based almost entirely on volumetric charges; pure fixed charges common for residential customers are typically negligeable for industrial ones (Borenstein, 2019). To offset the demand charge portion of the bill, the volumetric rate is typically lower for industrial consumers than residential ones. These demand charges are intended to approximate a customer's contribution to the system load, but they are considered inefficient because a plant's single highest consumption hour need not be coincident with overall peak demand periods (Borenstein, 2016).

[2]In addition, some industrial customers in select states are on direct access or retail choice contracts, so that the marginal prices that we collect may in fact underestimate the wedge for some industrial customers (NREL, 2017).

[3]In comparison with a purely volumetric bill, the cost of electricity is now partly driven by peak consumption, in addition to overall consumption; the demand charge acts as a fixed charge that potentially varies by billing cycle. As a result, two industrial customers with identical total consumption will pay different amounts for electricity when the consumption of one is more highly concentrated than the other, say if they run several machines simultaneously rather than sequentially.

3 Data and Methodology

We draw on restricted-use microdata from the U.S. Census Bureau on plant-level manufacturing inputs and outputs, as well as energy data from publicly available government sources. Additional data details are in Appendix A.

U.S. Manufacturing Census—Our primary data sources are administrative records on annual establishment-level inputs and outputs from the Census of Manufactures (CMF) from the U.S. Census Bureau for the 2017 Census year.[4] The CMF is conducted in years ending with 2 or 7 and surveys all manufacturing plants in the United States. The survey reports quantity of electricity purchased, in kilowatt-hours, and total expenditures on electricity and other energy sources (e.g. diesel fuel) separately. We calculate each plant's annual average electricity price as reported total electricity expenditure divided by electricity purchased. We validate our calculated average electricity prices against average price data from the Energy Information Administration's State Energy Data System (SEDS).[5]

We supplement the CMF data with information from the Manufacturing Energy Consumption Survey (MECS). The MECS is a probabilistic survey of approximately 15,000 manufacturing plants per year, conducted by the U.S. Census Bureau every four years beginning in 1998; we draw on information from 2018, the closest available year to the 2017 Census. This survey includes more detailed energy information that is not available in the CMF, including breakdowns of expenditure on and quantity consumed of fuels other than electricity. We calculate plant-level total carbon dioxide (CO2) produced and total British thermal units (BTU) of energy consumed by calculating industry-level CO2 and BTU per dollar of fuel expenditure from the MECS and multiplying by plant-level fuel expenditure in the CMF.

Electricity Rate Schedules—OpenEI provides a platform for creating, validating, and disseminating energy data. OpenEI's Utility Rate Database (URDB) contains rate structure information from over 200 utilities servicing industrial customers, including over 85% of electricity consumption in the U.S. It is compiled based on the authoritative list of U.S. utility companies maintained by the Energy Information Administration (EIA) of the U.S. Department of Energy (DOE). Rate schedules are verified and updated annually by the National Renewable Energy Laboratory (NREL) and

[4]The U.S. Census Bureau defines an establishment as a single physical location at which business is conducted or services or industrial operations are performed. For clarity, we use "establishment" interchangeably with "plant."

[5]Hawkins-Pierot and Wagner (2023) find energy price elasticities that are similar in sign, magnitude, and precision using the CMF and using SEDS.

Illinois State University additionally conducts more frequent quality control checks on behalf of the DOE; industry experts consider these data to be highly reliable (NREL, 2012).

Using the URDB, we compile the electricity rate schedules available in each ZIP code and match them to individual plants. Industrial electricity rates may include fixed monthly charges, usage rates per kilowatt-hour which may vary by season, time of day, or total usage, and demand charges which generally depend on the maximum usage within a billing cycle. We calculate a manufacturing plant's marginal cost of electricity in each year by first matching them to the industrial electricity rate schedules available in their ZIP code. We then assign them to a marginal usage rate based on their reported electricity use. For example, in August 2023 Kansas City Power and Light Co. offered a rate schedule that charges \$0.05073 per kilowatt-hour for the first 180 kilowatt-hours, \$0.03579 per kilowatt-hour for the next 180 kilowatt-hours, and \$0.02572 thereafter.

In our main analysis, we assume that manufacturing plants choose the lowest cost schedule based on their usage when multiple schedules are available in their ZIP code and that electricity usage is smooth across the year and the business day. For robustness, we show that our results are similar for markets with only one available rate or without seasonal or time-of-day charges. We then calculate the difference between the plant-level average electricity price reported in the Census data and the utility-level marginal electricity price available to each plant.

We summarize the distribution of average and marginal electricity prices and the wedge between them for the universe of manufacturing plants in the U.S., separately for eleven distinct industry groups, and separately for each state. To capture these distributions and preserve the anonymity of plants in the data, we present kernel density estimates of plant-level average price, marginal price, and the wedge trimmed at the 5th and 95th percentiles in accordance with Census disclosure requirements. Motivated by our regression results showing that total electricity usage is the strongest predictor of the electricity price wedge, we estimate these three kernel densities for the plants above and below median electricity usage.

We also analyze which plant characteristics are informative about the difference between average and marginal prices. We use the following regression specification:

$$(P_{avg} - P_{marg})_{isk} = X_{isk}\beta + \alpha_s + \gamma_k + \epsilon_{isk}$$

where X_{isk} includes plant characteristics such as revenue, inputs, energy intensity, and the number of available utilities, α_s and γ_k are state and industry fixed effects, respectively, and ϵ_{isk} is the error term, for plant i in industry k and state s. We estimate this specification with both OLS and using the least absolute shrinkage and selection operator (Lasso) to measure the relationship between plant characteristics and the average-marginal electricity price wedge. Due to the relatively large number of potentially informative plant characteristics, we use the commonly used Lasso estimator to improve efficiency and select the most informative correlates of the electricity price wedge in a principled way (Zou, 2021). In practice, this machine learning algorithm assesses whether each variable has additional explanatory power after the other covariates are included.

4 Results

The first key finding is that the marginal price of electricity paid by manufacturing plants is substantially lower than the average price. This pattern is intuitive given that demand and fixed charges will drive average prices above marginal prices, but the magnitude of the wedge is surprising. Table 1 summarizes the means of the average and marginal prices, as well as the difference between them and marginal price as a share of the average. Across all plant-years, the mean difference between the average and marginal price of electricity is 5.16 cents per kWh. This is about half the mean average price of 9.98 cents per kWh and slightly more than the mean marginal price of 4.82 cents per kWh. For 80% of plant-years, the marginal electricity price is between 22% and 79% of the average price reported by that plant.

Figure 1 shows the kernel density estimate of the full distribution of average and marginal electricity prices across all manufacturing establishments, as well as the difference between the prices. The figure again highlights that the mean of the distribution of marginal prices is roughly half of the mean of the distribution of average prices, but also that the distribution of average prices is almost entirely to the right of the distribution of marginal costs. As a result, the wedge between average and marginal prices is equal on average to the marginal price itself.

This pattern of substantially higher average electricity prices holds across states and industries. Figure 3 shows the mean wedge as a share of average electricity prices across plants in each state. Although marginal prices are closer to average prices in some states, such as Nevada, than others, the difference varies between 30% and 70% of the average. Appendix Figure A.2 shows that the magnitude and sign of the wedge between average and marginal prices is similar across eleven comprehensive industry groups.

There are two main policy implications of the finding that average prices exceed marginal prices by a substantial margin. First, the magnitude of the wedge means that using average price elasticities to predict responses to marginal price elasticities may overestimate energy use reductions by 50%. A given carbon price or clean energy subsidy is a small change in average electricity prices, but the change in marginal prices is twice as large. This distinction is important for understanding the magnitude of energy price changes needed to achieve the reductions in energy usage that are necessary to meet emissions targets. In a standard Dynamic Integrated Climate Economy (DICE)-like IAM, reducing the initial rate of decarbonization by 50% results in around an additional $0.25°C$ of warming in 2100 (Chikhani and Renne, 2023).[6]

The importance of the distinction between marginal and average prices can be observed using a simple example. Since the average price is the average of the marginal price and other, inframarginal components, we have: $p_{avg} = \frac{p_{other}+p_{marg}\times q}{q}$, where q is electricity quantity and we abstract from tiered pricing for illustration. The observed elasticity with respect to the average price is:

$$\begin{aligned}
\epsilon_{q,p_{avg}} &= \frac{\partial q}{\partial p_{avg}} \times \frac{p_{avg}}{q} = \left(\frac{\partial q}{\partial \left(\frac{p_{other}}{q}\right)} + \frac{\partial q}{\partial p_{marg}} \right) \times \left(\frac{p_{other} + p_{marg} \times q}{q} \right) \times \frac{1}{q} \\
&= \left(\frac{\partial q}{\partial \left(\frac{p_{other}}{q}\right)} \right) \times \left(\frac{p_{other} + p_{marg} \times q}{q} \right) \times \frac{1}{q} + \frac{\partial q}{\partial p_{marg}} \times \frac{p_{other}}{q^2} + \frac{\partial q}{\partial p_{marg}} \times \frac{p_{marg}}{q}
\end{aligned}$$

If $p_{avg} = 2p_{marg}$ and if firms respond only to changes in the marginal price (i.e., $\frac{\partial q}{\partial\left(\frac{p_{other}}{q}\right)} = 0$), then solving for p_{other} and replacing in the equation above yields $\epsilon_{q,p_{avg}} = 2\frac{\partial q}{\partial p_{marg}} \times \frac{p_{marg}}{q} = 2\epsilon_{q,p_{marg}}$. Hence, with the wedge we observe in our data, we would overestimated responsiveness to a marginal price change by a factor of 2.[7]

The magnitude of this wedge suggests that demand charges are a significant component of the total electricity price paid by industry. Demand charges bill industrial customers for their peak usage during a billing period, as an approximation of their imposition on the grid during peak times. The usage charge and fixed charge components of prices are small for industrial and commercial customers, unlike for residential contracts. The goal of these types of charges is to recover utilities' fixed costs, which include the costs of building and maintaining the distribution system, and the common historical solution for industrial customers was to allocate these costs based on an industrial customer's maximum demand since the utility had to build the distribution

[6] This reduction is readily observable by halving the initial decline in carbon intensity in the model of Chikhani and Renne (2023) that builds on DICE (Nordhaus and Boyer, 2000).

[7] If $\frac{\partial q}{\partial\left(\frac{p_{other}}{q}\right)}$ approaches $\frac{\partial q}{\partial p_{marg}}$, the wedge between the elasticities falls.

network to meet this demand. The importance of demand charges as a share of average electricity prices is surprising because these charges are considered an inefficient way to recover fixed costs due to their poor approximation of a customer's contribution to the need for generation. With today's technology, smart meters permit time-varying rate schedules that can efficiently charge customers for the load that they impose on the grid (Borenstein, 2016).

The second result is that electricity use is the largest predictor of the wedge between average and marginal electricity prices. Total electricity expenditure is spread over more kWh for higher utilization manufacturing plants, whether higher utilization is defined in terms of electricity usage (kWh) or intensity (kWh per dollar revenue). However, principled variable selection methods using Lasso regressions reject the informativeness of many other plant characteristics, including revenue, labor and capital inputs, carbon dioxide intensity of the local electricity grid, among others (Tibshirani, 1996). Table 2 shows that other variables have limited information content in the Lasso regressions and are largely statistically insignificant in the OLS regressions as well. This result continues to hold when we consider plants that have only one utility contract available to them or when we consider inverse-hyperbolic sine transformations of the outcome variables, rather than levels (Appendix Tables A.1 and A.2).

Motivated by these results, we estimate the kernel density of the distribution of average and marginal electricity prices. Figure 2 shows these results. Average and marginal prices are closer together for the high utilization plants because demand charges are spread over more units, but the difference, while statistically significant, is not economically meaningful as the distributions overlap for most of their support. The wedge between average and marginal prices is large for almost all plants, even conditional on high electricity usage. As result, marginal prices seem to be rarely a good proxy for average prices in U.S. manufacturing.

5 Discussion and Conclusion

This paper summarizes how and when the differences between average and marginal electricity prices are relevant for climate policy. We show that marginal electricity prices are half of average electricity prices, on average in the U.S. manufacturing sector and separately for eleven major industry groups. We find that while the wedge between average and marginal electricity prices is smaller for plants that use more electricity, the reduction in the wedge is limited. As a result, caution is always warranted in using average price elasticities to interpret responses to marginal price changes—even for the largest plants.

The policy implications of these results are consequential: since marginal energy prices are below average prices, energy price elasticities based on average price changes will be higher than the elasticity with respect to changes in the marginal price of energy. Therefore, the energy price changes needed to achieve desired emissions reductions will look too low. Using average price elasticities to estimate responses to marginal energy price changes may therefore lead to energy policy that falls short of constraining climate change to "safe" levels. However, the use of either average or marginal costs to calculate behavioral responses to energy price changes does not change the optimal Pigouvian tax needed to internalize unpriced emissions externalities, which is derived independently of these elasticities as the social cost of carbon; what changes is the extent to which a carbon tax equal to the social cost of carbon will actually reduce industrial energy consumption and, by extension, emissions.

What can policymakers do to ensure that the average price data collected in administrative surveys are used appropriately? A first step is ensuring model calibrations are internally consistent: if average price elasticities are used to predict behavioral responses, then the change in average price resulting from a carbon price should be inputted, rather than the change in marginal price. Continuing to use average price elasticities is likely more straightforward than collecting marginal price data, which are difficult to observe. Existing models can therefore continue to be applied to assess the effectiveness of policies to avoid dangerous climate change.

References

Alberini, A. and Filippini, M. (2011). Misunderstanding nonlinear prices: Evidence from a natural experiment on residential electricity demand. *Energy Economics*, 33:889–895.

Blonz, J. A. (2021). Making the best of the second-best: Welfare consequences of time-varying electricity prices. Energy Institute at Haas Working Paper 275.

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

Borenstein, S. (2019). Are demand charges fair?

Borenstein, S. and Bushnell, J. B. (2015). The u.s. electricity industry after 20 years of restructuring. *Annual Review of Economics*, 7:437–463.

Borenstein, S. and Bushnell, J. B. (2022). Do two electricity wrongs make a right? cost recovery, externalities, and efficiency. *American Economic Journal: Economic Policy*, 14(4):80–110.

Boyd, G. and Lee, J. (2018). Relative effectiveness of energy efficiency programs versus market based climate policies in the chemical industry. Working Paper.

Carleton, T. and Greenstone, M. (2021). Updating the united states government's social cost of carbon. Working Paper.

Chikhani, P. and Renne, J.-P. (2023). An analytical framework to price long-dated climate-exposed assets. Working Paper.

Csereklyei, Z. (2020). Price and income elasticities of residential and industrial electricity demand in the european union. *Energy Policy*, 137:111079.

Deryugina, T., MacKay, A., and Reif, J. (2020). The long-run dynamics of electricity demand: Evidence from municipal aggregation. *American Economic Journal: Applied Economics*, 12(1):86–114.

EPA (2021). Sources of greenhouse gas emissions. Technical report, U.S. Environmental Protection Agency.

Fell, H., Li, S., and Paul, A. (2014). A new look at residential electricity demand using household expenditure data. *International Journal of Industrial Organization*, 33:37–47.

Ganapati, S., Shapiro, J., and Walker, R. (2020). Energy cost pass-through in u.s. manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12(2):303–342.

Hawkins-Pierot, J. and Wagner, K. (2023). Technology lock-in and costs of delayed climate policy. Working Paper.

Hayes, K. and Hafstead, M. (2020). Carbon pricing 103: Effects across sectors. Technical report, Resources for the Future.

Ito, K. (2014). Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–563.

Linn, J. (2008). Energy prices and the adoption of energy-saving technology. *Economic Journal*, 118:1986–2011.

Nordhaus, W. D. and Boyer, J. (2000). *Warming the World: Economic Models of Global Warming*. MIT Press.

NREL (2012). National utility rate database. Technical report, National Renewable Energy Laboratory.

NREL (2017). An introduction to retail electricity choice in the united states. Technical report, National Renewable Energy Laboratory.

Paul, A., Myers, E., and Palmer, K. (2009). A partial model of u.s. electricity demand by region, season, and sector. Technical report, Resources for the Future.

Pindyck, R. S. (2020). What we know and don't know about climate change, and implications for policy. NBER Working Paper 27304.

Puller, S. L. and West, J. (2013). Efficient retail pricing in electricity and natural gas markets. *American Economic Review*, 103(3):350–355.

Ramsey, F. P. (1927). A contribution to the theory of taxation. *Economic Journal*, 37(145):47–61.

Shaffer, B. (2020). Misunderstanding nonlinear prices: Evidence from a natural experiment on residential electricity demand. *American Economic Journal: Economic Policy*, 12(3):433–461.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288.

Zou, E. Y. (2021). Unwatched pollution: The effect of intermittent monitoring on air quality. *American Economic Review*, 111(7):2101–2126.

Figures and Tables

Figure 1: Distribution of Average and Marginal Electricity Prices: All Industries

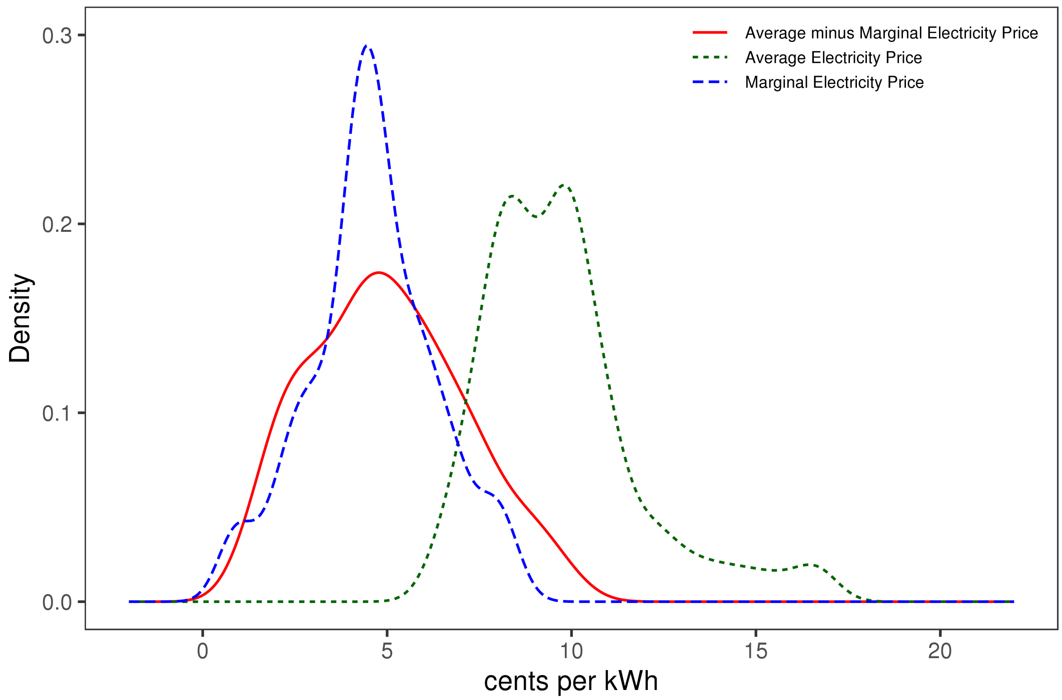

Notes: This figure shows the kernel density estimates of the distribution of average electricity prices, marginal electricity prices, and average minus marginal electricity prices for U.S. manufacturing industries. Prices are in cents per kWh.

Figure 2: Distribution of Average and Marginal Prices: High and Low Electricity Usage Plants

(a) Average and Marginal Prices

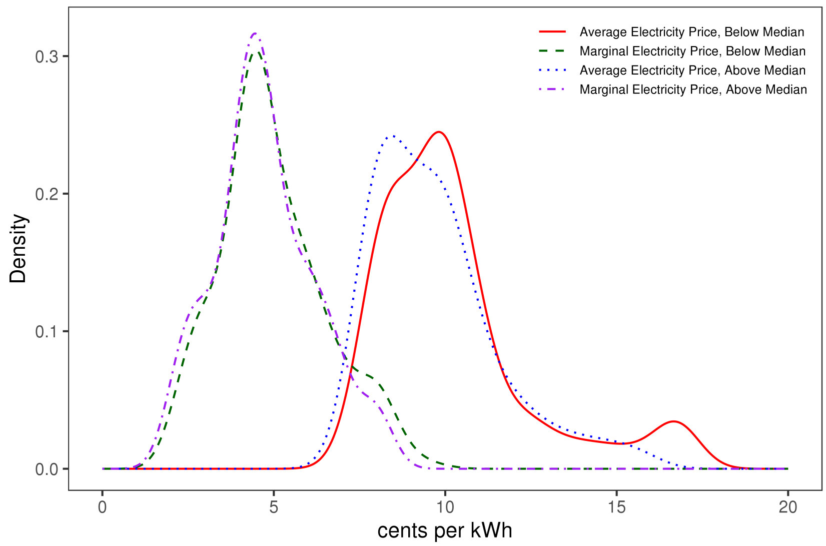

(b) Difference Between Average and Marginal Prices

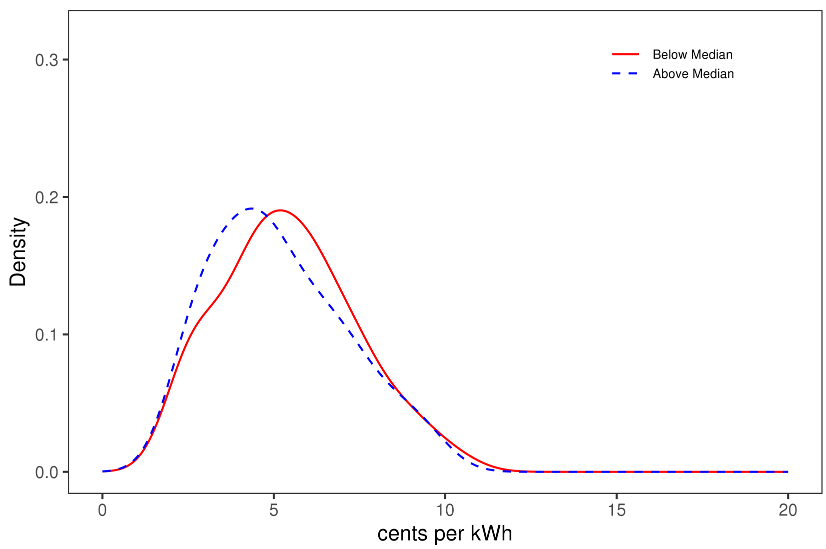

Notes: This figure shows the kernel density estimates of the distribution of average electricity prices and marginal electricity prices (Panel A) and average minus marginal electricity prices (Panel B) for U.S. manufacturing industries. Above and below median plants refer to plants with above and below median electricity usage. Prices are in cents per kWh.

Figure 3: Average - Marginal Electricity Price Wedge by State

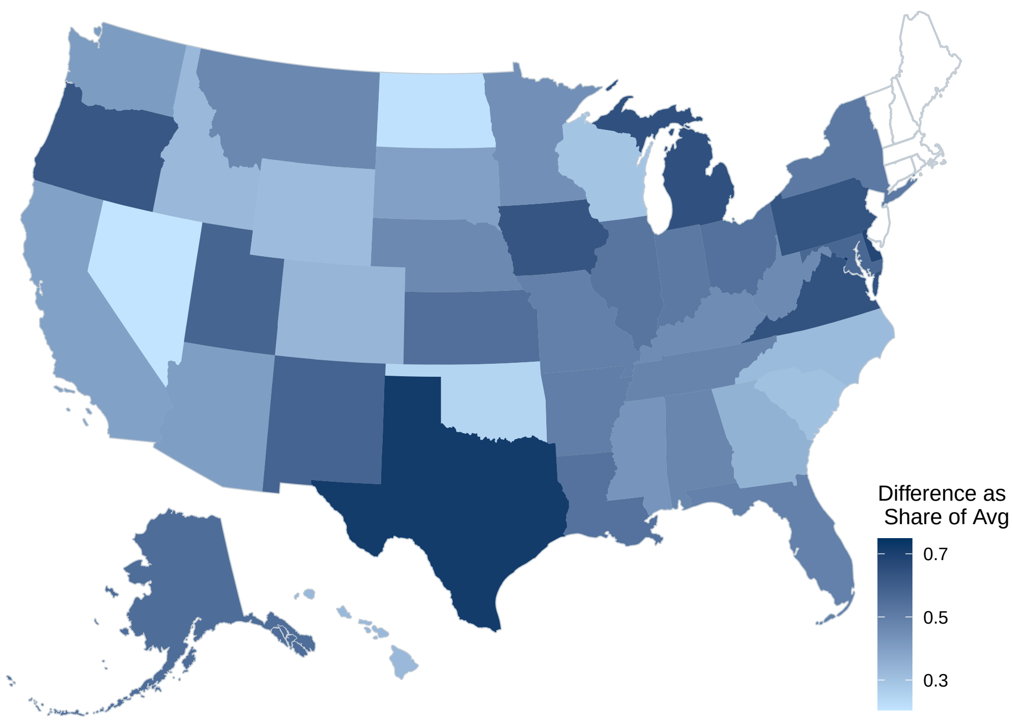

Notes: This figure shows the difference between average and marginal electricity prices as a share of the average price by state for the U.S. manufacturing industry. States with no data have insufficient manufacturing plants to meet Census Bureau disclosure requirements.

Table 1: Summary Statistics

	Mean	Std	Percentiles 10th	Median	90th
Average Price (cents / kWh)	9.98	3.12	7.11	9.58	14.24
Marginal Price (cents / kWh)	4.82	2.40	2.23	4.59	7.67
Average - Marginal Price (cents / kWh)	5.16	2.78	1.90	4.89	8.69
Marginal Price / Average Price	0.51	0.21	0.22	0.58	0.79
Plant Revenue (million USD)	30.84	255.20	0.44	3.74	51.43
Electricity Intensity (kWh / $ revenue)	0.13	0.34	0.02	0.07	0.27
Electricity Expenditure (million USD)	0.28	2.12	0.00	0.02	0.48
Electricity Usage (million kWh)	4.06	38.11	0.02	0.24	5.57
Utilities Available (N)	3.12	3.45	1.00	1.00	7.00
N = 108000					

Notes: This table shows summary statistics for the universe of manufacturing plants in the 2017 U.S. Census. All dollar values are in 2017 USD.

Table 2: Correlations of Electricity Prices and Plant Characteristics

	Average - Marginal Price		Average Price		Marginal Price	
	OLS	LASSO	OLS	LASSO	OLS	LASSO
Revenue (10 mil $)	-0.001		-0.001		0.000	
	(0.002)		(0.002)		(0.000)	
Electricity Intensity (kWh/$)	-1.366***	-0.573***	-1.518***	-0.639***	-0.152***	-0.063**
	(0.193)	(0.097)	(0.201)	(0.105)	(0.035)	(0.023)
CO2 Intensity (kg/$)	-1.313**		-1.269*		0.044	
	(0.584)		(0.652)		(0.302)	
BTU Intensity (BTU/$)	0.553***		0.591***		0.038	
	(0.116)		(0.128)		(0.040)	
Electricity Costs (10 mil $)	0.580***		0.784***		0.204	
	(0.211)		(0.259)		(0.123)	
Electricity Used (10 mil kWh)	-0.044***	-0.022***	-0.056***	-0.024***	-0.012*	-0.004*
	(0.012)	(0.008)	(0.015)	(0.007)	(0.006)	(0.002)
Labor (10 mil hours)	-3.508***	-3.533***	-3.735***	-3.598***	-0.227	
	(0.875)	(0.786)	(0.894)	(0.736)	(0.232)	
Capital Stock (10 mil $)	-0.002		-0.003*		-0.001	
	(0.002)		(0.001)		(0.001)	
Materials Costs (10 mil $)	0.003*		0.003*		0.000	
	(0.002)		(0.002)		(0.001)	
Fuels Costs (10 mil $)	-0.137		-0.178**		-0.041*	
	(0.086)		(0.082)		(0.022)	
Utility Count (N)	0.014		-0.001		-0.015	-0.015
	(0.025)		(0.010)		(0.027)	(0.028)
N	108000	108000	108000	108000	108000	108000
AdjR2	0.247	0.242	0.546	0.541	0.646	0.645

Notes: Outcome variables are measured in cents per kWh. Columns 1, 3, and 5 are estimated using OLS and columns 2, 4, and 6 are estimated using Lasso. All regressions include state and industry fixed effects, which are not penalized in Lasso regressions. The sample includes the universe of manufacturing plants in the 2017 U.S. Census. All dollar values are in 2017 USD.

Appendix

A Data

This section includes additional details on the construction of the dataset.

Manufacturing data—We draw mainly on data from the U.S. Census Bureau that includes inputs and outputs for the universe of manufacturing plants in 2017, the most recently available Census year. Our steps to clean the data broadly follow other papers using the same data (e.g., Ganapati et al., 2020; Hawkins-Pierot and Wagner, 2023). First, we exclude observations that have missing or negative values for average electricity prices, electricity intensity, revenue, or inputs including capital stock, labor, materials, electricity and fuel expenditures. Second, we exclude observations that have zero values for average electricity prices, revenue, labor costs, material costs, or electricity expenditures. Third, we drop imputed administrative records. Finally, we exclude remaining outliers for capital stocks, revenue, labor costs, materials costs, electricity expenditures, or raw fuels expenditures that exceed 100 times the 99th percentile of the respective distributions of these variables.

To calculate CO2 and BTU per dollar revenue, we use the closest available year of data from the Manufacturing Energy Consumption Survey, which is 2018. The MECS includes about 15,000 plants across industries and oversamples large plants. This survey includes information on quantities and expenditures of individual fuels that are unavailable in the CMF, which includes only total expenditures on electricity and aggregate raw fuels, as well as quantity of electricity. We calculate six-digit North American Industry Classification System (NAICS) energy intensity averages (million BTU per USD of fuel expenditures and kg CO2 per USD fuel expenditures), which we then use with the CMF to calculate CO2 intensity and BTU intensity at the plant level. We multiply the industry averages from the MECS by the raw fuels expenditures available at the plant level from the CMF to obtain plant-specific estimates of CO2 and BTU from raw fuels. To obtain total CO2 and BTU, we calculate embodied emissions and energy content from electricity using conversion factors from the EPA's eGrid product, which includes emissions factors by state that capture the energy mix of each state's upstream electricity grid. In this process, we exclude feedstock fuels and process emissions from these calculations, as well as "by-product" fuels that aren't associated with well-defined CO2 and BTU conversion factors. Throughout, we deflate all monetary values to 2017 dollars using the input- and industry-specific price indices available from the National Bureau of

Economic Research-Census of Economic Studies (NBER-CES) Productivity Database.

Electricity rate schedules—We merge marginal price data with the average price data available in the CMF. We match marginal prices from OpenEI to the plants in the CMF at the ZIP code level. We obtain a match rate of 78% of plants to utilities. We use the electricity quantities reported in the CMF to determine each plant's effective marginal price in instances where utilities offer a tiered rate structure. When more than one rate schedule is available, we calculate the minimum marginal price and the total number of utilities available to each plant by ZIP code. In the rare instances where utilities offer industrial time-of-day pricing, we calculate the mean marginal price across business hours. We find that our results are robust to including and excluding plants with multiple rate schedules available to them. We exclude rate schedules that have negative prices and outliers outside of the top and bottom 1% of the difference between average and marginal costs. Throughout, we deflate marginal prices to 2017 dollars using the average of the energy deflators across industries from the NBER-CES Productivity Database.

B Figures and Tables

Figure A.1: Distribution of Average and Marginal Prices: Plants Serviced by One Utility Only

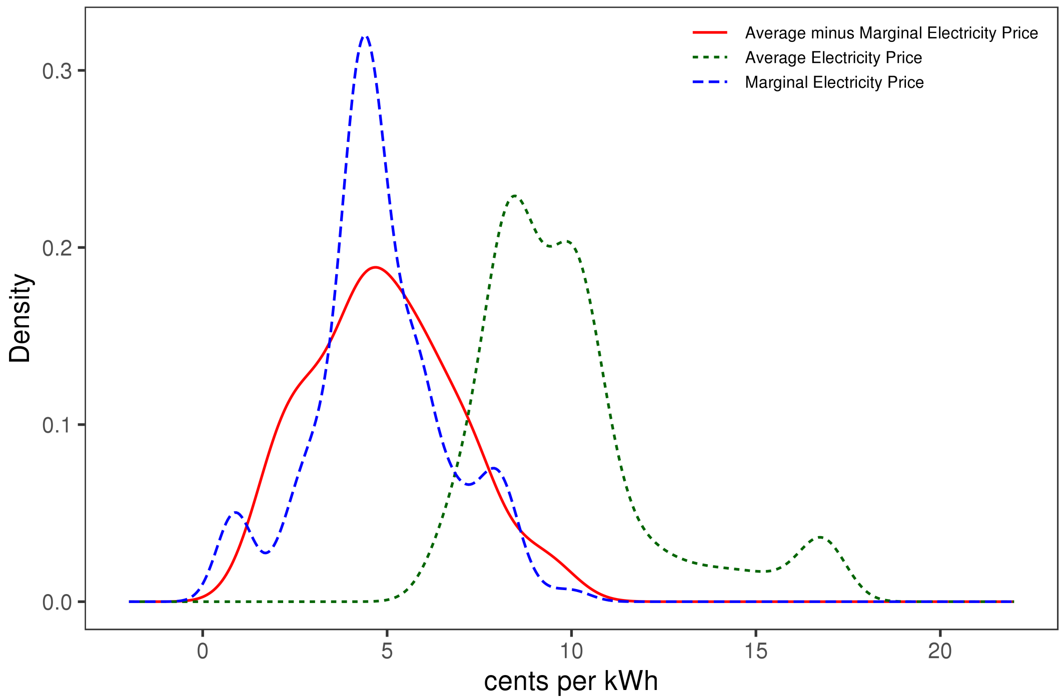

Notes: This figure shows the kernel density estimates of the distribution of average electricity prices, marginal electricity prices, and average minus marginal electricity prices for U.S. manufacturing industries. The sample includes only plants serviced by one electric utility. Prices are in cents per kWh.

Figure A.2: Distribution of Average and Marginal Prices: by Industry

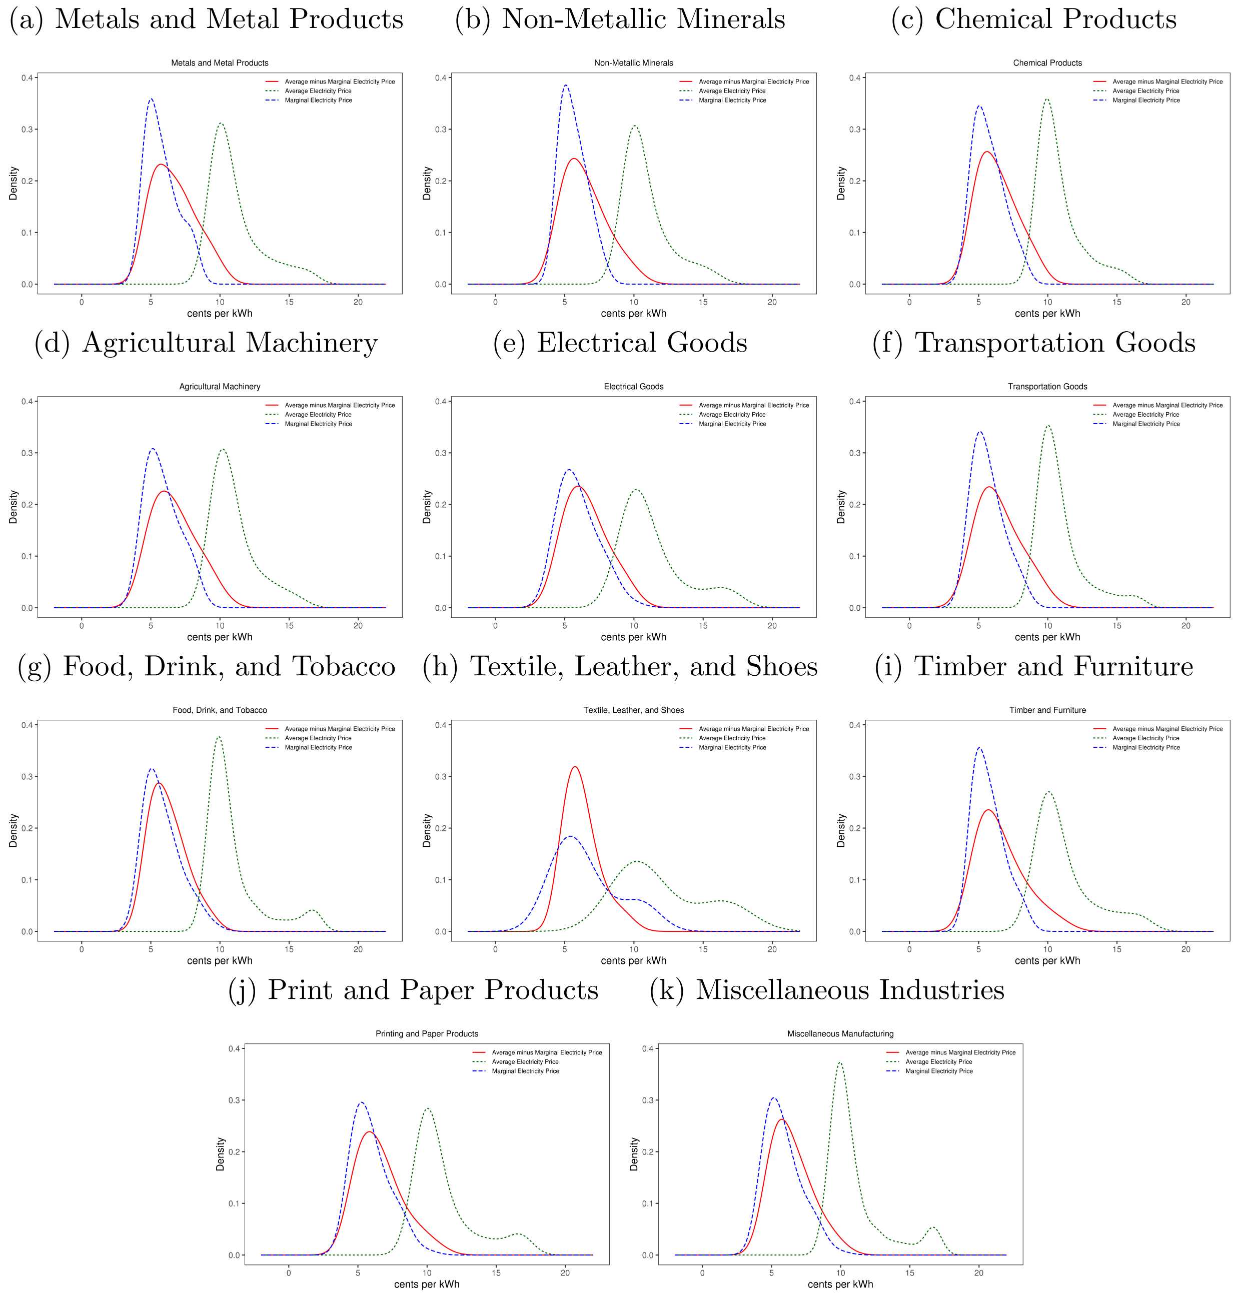

Notes: This figure shows the kernel density estimates of the distribution of average electricity prices, marginal electricity prices, and average minus marginal electricity prices separately for 11 distinct manufacturing industry groups. Prices are in cents per kWh.

Table A.1: Correlations of Electricity Prices and Plant Characteristics, One Utility

	Average - Marginal Price		Average Price		Marginal Price	
	OLS	LASSO	OLS	LASSO	OLS	LASSO
Revenue (10 mil $)	-0.001		-0.001		0.000	
	(0.004)		(0.004)		(0.001)	
Electricity Intensity (kWh/$)	-1.282***	-0.664***	-1.382***	-0.723***	-0.100***	-0.049**
	(0.191)	(0.103)	(0.211)	(0.112)	(0.0306)	(0.022)
CO2 Intensity (kg/$)	-1.642**		-1.465*		0.177	
	(0.751)		(0.861)		(0.295)	
BTU Intensity (BTU/$)	0.523***		0.526***		0.003	
	(0.121)		(0.134)		(0.035)	
Electricity Costs (10 mil $)	0.605***		0.702***		0.097**	
	(0.200)		(0.202)		(0.040)	
Electricity Used (10 mil kWh)	-0.033**		-0.040***		-0.007**	-0.004***
	(0.013)		(0.014)		(0.003)	(0.001)
Labor (10 mil hours)	-3.138***	-3.433***	-3.272***	-3.597***	-0.135*	
	(0.997)	(0.982)	(0.999)	(0.955)	(0.077)	
Capital Stock (10 mil $)	-0.002		-0.002		0.000	
	(0.002)		(0.002)		(0.000)	
Materials Costs (10 mil $)	0.002		0.003		0.001	
	(0.004)		(0.003)		(0.001)	
Fuels Costs (10 mil $)	-0.313***		-0.359***		-0.046**	
	(0.109)		(0.109)		(0.018)	
N	61500	61500	61500	61500	61500	61500
AdjR2	0.264	0.260	0.656	0.653	0.789	0.789

Notes: Outcome variables are measured in cents per kWh. Columns 1, 3, and 5 are estimated using OLS and columns 2, 4, and 6 are estimated using Lasso. All regressions include state and industry fixed effects, which are not penalized in Lasso regressions. The sample includes manufacturing plants serviced by one electric utility only. All dollar values are in 2017 USD.

Table A.2: Correlations of Electricity Prices and Plant Characteristics

	Average - Marginal Price		Average Price		Marginal Price	
	OLS	LASSO	OLS	LASSO	OLS	LASSO
Revenue (10 mil $)	0.004	0.010	0.009		0.001	
	(0.020)	(0.013)	(0.007)		(0.014)	
Electricity Intensity (kWh/$)	-0.806***	-0.727***	-0.471***	-0.485***	-0.123***	-0.040**
	(0.077)	(0.072)	(0.031)	(0.030)	(0.030)	(0.017)
CO2 Intensity (kg/$)	0.502**	0.415***	0.241***	0.259***	-0.020	
	(0.093)	(0.096)	(0.043)	(0.039)	(0.063)	
BTU Intensity (BTU/$)	0.275***	0.252***	0.177***	0.170***	0.059***	
	(0.036)	(0.032)	(0.012)	(0.011)	(0.017)	
Electricity Costs (10 mil $)	1.434***	1.203***	0.786***	0.746***	0.120***	
	(0.173)	(0.142)	(0.062)	(0.064)	(0.040)	
Electricity Used (10 mil kWh)	-0.557***	-0.542***	-0.274***	-0.247***	-0.014	
	(0.048)	(0.048)	(0.011)	(0.014)	(0.022)	
Labor (10 mil hours)	0.067		-0.012		-0.051	
	(0.074)		(0.026)		(0.046)	
Capital Stock (10 mil $)	0.006		0.002		-0.002	
	(0.006)		(0.002)		(0.006)	
Materials Costs (10 mil $)	0.006		0.001		0.002	
	(0.011)		(0.004)		(0.006)	
Fuels Costs (10 mil $)	-0.393***		-0.271**	-0.264***	-0.073*	
	(0.090)		(0.037)	(0.036)	(0.036)	
Utility Count (N)	0.006		-0.002		-0.065	-0.065
	(0.027)		(0.004)		(0.046)	(0.046)
N	108000	108000	108000	108000	108000	108000
AdjR2	0.360	0.358	0.554	0.554	0.520	0.520

Notes: Outcome variables are measured in cents per kWh. Dependent and independent variables are all inverse hyperbolic sine transformations. Columns 1, 3, and 5 are estimated using OLS and columns 2, 4, and 6 are estimated using Lasso. All regressions include state and industry fixed effects, which are not penalized in Lasso regressions. The sample includes the universe of manufacturing plants in the 2017 U.S. Census. All dollar values are in 2017 USD.