Adjustment on the Margin: Evaluating emissions reduction policies in the face of short-run adjustment costs*

James Archsmith†
University of Maryland, College Park

Draft Date: May 13, 2019

Preliminary and incomplete. Please do not circulate or cite without permission of the Author.

*I would like to thank James Bushnell, Kenneth Gillingham, Joshua Linn, Erin Mansur, and David Rapson for insightful comments on this work. I am grateful to the UC Davis Office of Graduate Studies, College of Letters and Science: Division of Social Sciences, and the UC Davis Department of Economics for financial support during the initial stages of this project. José Eguiguren-Cosmelli provided excellent research assistance. Any errors are my own.

†Department of Agricultural and Resource Economics, University of Maryland, College Park, MD 20742 and Department of Economics. Email: archsmit@umd.edu, URL: https://econjim.com

Abstract

Subsidies for renewable electricity generation frequently target the quantity of energy produced but ignore emissions resulting from *changes* in demand. I estimate the impact of policies that alter the generation or demand for electricity on polluting emissions, paying special attention to the impact on the rate and frequency with with fossil fuel plants are required to adjust output (ramp) to meet demand. Failing to account for the impacts of ramping overstate the emissions benefits of new solar PV and vastly understate the benefits of electricity storage technology. These results suggest policies targeting pollution reductions from electricity generation should account for the level and pattern of changes in supply and demand.

JEL: Q51, Q52, Q53

Keywords: Electricity, Environment, Marginal Emissions Factors, Electricity Storage, Renewable Energy

1 Introduction

There is a substantial literature dedicated to the analysis of the emissions of both local criteria pollutants and greenhouse gases from the electricity generating sector. This research is critical in developing policy to reduce exposure to the harmful effects of poor air quality and mitigate climate change at the lowest possible cost. Importantly – e.g., Graff Zivin, Kotchen, and Mansur (2014) or Holland, Mansur, Muller, and Yates (2016) – an analysis of policy that will induce incremental changes to electricity supply or demand should assess impacts on the behavior of the marginal generator.

These works, and many that followed, use information on demand and realized emissions to compute an *ex post* estimate of the emissions factor of the marginal producer. These *ex post* analyses are informative to policy for a number of reasons. Importantly, forward-looking estimates of marginal emissions factors derived from engineering models and dispatch simulations will not accurately compute marginal emissions factors (MEFs) if actual dispatch and operation of generating units differs from engineering models.[1] *Ex post* analyses, however, are based on the observed behavior of electricity generators and do not hinge on engineering assumptions.

Given the importance of reliable estimates of emissions from electricity generation for policy analysis, this paper presents several innovations in the *ex post* estimation of MEFs, described in detail in Section 4. First, I do not treat the MEF as a static function of current electricity demand. In reality, changes in the almost perfectly inelastic demand for dispatchable electricity generation cause fossil fuel generators to engage in costly ramping of output. Startup and safe shutdown of fossil fuel plants consumes fuel, and generates emissions, that do not result in useful output. I allow the MEF to be a function of not only current demand but also the past and expected future levels of demand. As noted by Holland and Mansur (2008), ramping has non-trivial impacts on the emissions from fossil fuel generators and ignoring the impact of ramping has the potential to vastly understate the emissions benefits of reducing variability in demand falling on fossil fuel generators, such as through electricity storage.

Empirical models of the MEF include myriad controls and choices of specification. Rather than make seemingly arbitrary decisions on the precise relationship between electricity demand and power plant emissions, I define a set of candidate models and use model selection algorithms from the machine learning literature to choose a preferred specification based on out-of-sample forecast performance. The objective of this process is to develop parsimonious models that accurately capture the complicated relationship between power plant emissions and past, present, and expected future demand. I demonstrate empirical models that achieve substantially lower out-of-sample forecast error with fewer free parameters than have been previously presented in the literature. Further, it is clear that the level of and rate of change in fossil fuel demand (ramp), as opposed to factors such as the hour of the day, are the relevant determinants of power plant emissions on the

[1] Actual dispatch of electricity generating units could deviate from engineering models due to unmodeled engineering details such as unit failures or transmission constraints, or economic market failures such as the exercise of market power.

margin.

Electricity is not only produced by fossil fuel-fired plants. Numerous other generators participate in electricity markets, supplying energy without emitting carbon dioxide (CO_2) or other gaseous pollutants. Much of the previous literature has treated the supply of non-fossil fuel generation as exogenous to demand at any given moment. In many regions, however, this is not an accurate representation of electricity supply. Large hydroelectric dams and, to a lesser degree, solar thermal and electricity storage facilities can adjust output on demand, and often do in response to periods of high demand.[2] Ignoring the supply response of dispatchable renewables will overstate the impact of marginal changes in demand on emissions.

Finally, I rely on a new and highly detailed dataset of electricity demand from the Energy Information Administration (EIA). This dataset, EIA-930, provides hourly electricity demand, forecast demand, imports, and exports for every electricity balancing authority in the continental United States. Importantly, reporting in these data accurately reflect the structure of the United States electric grid and potential constraints to transmission between regions. Further, unlike previously available data, EIA-930 provides consistent, reliable forecasts of future demand, which are an important determinant of the MEF.

While these ramping effects account for a small portion of overall (or average) emissions, they are a large component of changes to emissions on the margin. By way of example, in the ERCOT region during 2018, median hourly load was 38,770 MW, and the median absolute hour-to-hour change in load was 1,079 MW. A transient 10 MW increase in load represents a 0.025% increase in total demand, but nearly a 1% increase in the total amount of ramp required to meet that change in demand. Importantly, generation technologies such as solar PV, wind turbines, and electricity storage will necessarily induce changes not only to the level of demand falling on fossil fuel generation, but the rate and frequency with which plants on the margin must adjust output.

I find the ramping effects can be a large component of the MEF in any given hour. Further, there is substantial heterogeneity in both space and time in the magnitude of these ramping effects. Analyses described in Section 5, show the effects are largest in regions with heavy penetration of solar PV systems. For example, in California, a 1 MW increase in demand for electricity sustained for four hours starting at 7 AM, when the bulk of solar PV generation is coming online, has half the impact on emissions of a similar increase starting at 8 PM, when solar production is falling off and fossil fuel plants must quickly increase output to compensate, substantially larger than the estimated difference when not accounting for the impacts of ramping.

Applying these results, in Section 6, I show ramping effects have large implications for policies promoting installation of solar, wind, and electricity storage systems. Accounting for ramping emissions, the CO_2 emissions benefits per MW of installed solar capacity in California is 16% lower than the estimate when ignoring ramping. This is fully a consequence of the timing of solar generation, which drive much of the need for ramping by fossil fuel generators. Solar PV begins

[2] Archsmith (2017) shows hydroelectric dams in California increase output in periods where residual demand of fossil fuel plants is rapidly increasing and marginal costs are high. Further, when the ability of these dams to flexibly adjust output is restricted, fuel consumption and emissions from fossil fuel generators increase.

supplying electricity at a time when fossil fuel plants are being forced to ramp down and the marginal benefit to reducing fossil fuel load is very small, but solar generation falls off at time when fossil fuel plants are quickly ramping up and there are large emissions impacts from marginal increases in the required ramp rate.

Failing to account for ramping effects also substantially understates the emissions benefits of electricity storage. In the case of a battery storage facility with 80% efficiency operated by a CO_2-minimizing social planner in Texas, estimated CO_2 emissions reductions are nearly double when accounting for the ramping effect. Energy storage systems can consume additional power as demand on fossil fuel generation is falling off, reducing the rate at which fossil fuel plants need to ramp down. They can then discharge power later in the day, softening the rate at which those same plants must ramp back up.

Absent a price on the emissions externality, however, a profit maximizing operator of energy storage will deviate substantially from emissions minimization. Failing to account for ramping effects, one would conclude profit-maximizing electricity storage might increase emissions.[3] Accounting for avoided ramp, however, even profit maximizing energy storage has CO_2 emissions benefits.

2 Motivation

Electricity supply in the United States is provided by a heterogeneous set of generation technologies. These technologies differ in attributes including energy source (fossil fuels, nuclear, gravity, solar, etc.) and corresponding externalities (GHG or local criteria pollutant (LCP) emissions, radiation hazard, etc.), capital-operating cost tradeoff, the ability to be activated and deactivated on demand (dispatchability), and output adjustment capability and cost.

Demand for electricity is derived from all end-users, including residential consumers, industry, and commercial operations. The bulk of electricity consumers in the United States face retail electricity tariffs wholly unreflective of supply scarcity at any given moment.[4] Tariffs are generally set through a rate-making process months or years in advance and, if they have a time-varying component at all, only coarsely approximate the actual cost of electricity supply. Since even large shifts in marginal cost are not communicated into retail tariffs, consumers do not adjust consumption in response and demand for electricity is nearly perfectly inelastic.

Electricity supply, as measured by the energy fed into the electric grid, must always match the quantity demanded with very little discrepancy. Mismatch between supply and demand can damage generators,transmission equipment, or devices that consume electricity and can cause blackouts or brownouts. The United States is divided into regions where a single entity, called a Balancing

[3] e.g., Carson and Novan (2013) find electricity storage participating in the ERCOT energy market could increase CO_2 emissions but decrease nitrogen oxides (NO_X) emissions. Babacan et al. (2018) demonstrate residential energy storage systems have the potential to increase greenhouse gas (GHG) emissions.

[4] In 2017 the EIA reported in Form EIA-861 (EIA-861) about 4.6% of United States residential electricity customers were using some form of time-varying pricing.

Authority (BA), is responsible for ensuring supply meets demand at all times while minimizing the total cost of suppling electricity.

Amongst dispatchable generators of electricity, there is heterogeneity in the marginal cost of operation. A BA minimizing the total cost of supplying electricity would rank-order generators in terms of marginal cost, with low marginal cost generators being called on first followed by higher cost generators. Ranking generators by marginal cost and intersecting with demand identifies a marginal generator. This is the generator that would be called on by the BA to increase output in the face of an increase in demand, or decrease output if demand were to fall.

This paper will focus on estimating the impact of changes incremental to the pattern of supply of or demand for electricity on the level of emissions from electricity generating plants. There are many candidates that could cause such changes, e.g., changes in the type or characteristics of the generators supplying electricity, demand-side changes such as time-varying pricing or demand management, or the deployment of electricity storage facilities. These incremental changes to supply and demand will have little impact on electricity supply and demand as a whole, but will alter the behavior of suppliers on the margin. Thus, any analysis of the impact on emissions of incremental changes to electricity markets should consider the behavior of those suppliers on the margin.

Extending previous work in this space, I note many of these changes alter not only the intersection of supply and demand for electricity at any given moment, but how that intersection changes over the course of hours or days. Since that demand is nearly perfectly inelastic and supply and demand must balance at all times, changes in this intersection often requires that some dispatchable supplier adjust output in response to changes in demand or willingness to supply by other firms. For fossil fuel generators, who are often the supplier on the margin, these adjustments are costly. Increasing or decreasing output reduces plant efficiency, consuming more fuel and generating more emissions, than would be produced under a constant level of output.

In particular, I will consider how incremental changes to the electricity infrastructure, such as adding a solar array or electricity storage facility, will impact the profile of power plant emissions. This analysis is similar in spirit to many others. Siler-Evans, Azevedo, and Granger Morgan (2012), who use regression models to compute *ex post* estimates of MEFs. Graff Zivin, Kotchen, and Mansur (2014) apply similar methods to show how spatial and temporal heterogeniety in MEFs have impacts for policies promoting electric vehicles, distributed solar or energy storage. Archsmith, Kendall, and Rapson (2015) deploy regression models to identify the marginal fuel source for electricity generation in an analysis of the life-cycle benefits of electric vehicles. Holland, Mansur, Muller, and Yates (2016) combine plant-level estimates of emissions and an atmospheric transport model to determine comprehensive environmental benefits of electric vehicles compared to those with internal combustion engines. Novan (2015) considers the impact of variability in wind generation in emissions from fossil fuel plants in Texas and Oklahoma.

More recent work has combined both *ex post* analyses of emissions with data on prices. Fang, Asche, and Novan (2018) combine marginal prices and MEFs to determine the hours of lowest social cost for charging electric vehicles. Borenstein and Bushnell (2018) use retail and wholesale prices combined with marginal damage estimates to compare the retail and social cost of electricity throughout the United States.

Here I will extend techniques for estimating MEFs to analyze the emissions benefits of policies which alter the pattern of supply or demand for electricity. These include adding either wind or solar PV capacity, which generate power in predictable cycles that vary throughout the day, and electricity storage. A critical value of electricity storage technology is that it allows electricity generated from non-dispatchable sources, such as solar PV or wind turbines, to be collected and dispatched on demand. In the thought experiment considered here, some electricity storage technology will consume electricity in some periods, marginally increasing electricity demand in those periods, and then discharge it later, marginally decreasing the supply required from other sources.

Importantly, this behavior by hypothetical electricity storage provider will cause adjustments in output by the generator on the margin of electricity supply. Since demand for electricity can vary substantially throughout the course of a day, the identity of the marginal generator will change as well. As noted by Graff Zivin, Kotchen, and Mansur (2014) and others, the impact from a unit of electricity storage on power plant emissions depends on the emissions profiles of the marginal producer both when the storage operator is storing and discharging electricity.

The importance of identifying the marginal supplier is clear from examining the average CO_2 emissions rates of various types of fossil-fuel fired power plants, shown in Table 1. Emissions per MWh of electricity generated vary substantially by source. Further, emissions can vary within a given plant as it adjusts output to changing levels of demand.

The bulk of previous empirical estimates of marginal emissions factors treat emissions from electricity generation as a static attribute, *i.e.*, the marginal emissions factor depends only on currently electricity demand and on past changes or expected future changes in demand. In reality, adjusting power plant output is costly. In particular, a power plant may consume significant quantities of fuel during its startup phase without generating any appreciable output.[5] Electricity storage could provide significant emissions benefits by reducing variability in the residual demand falling on fossil fuel generators and the number of costly startup-shutdown cycles.

Grid operators will anticipate future changes in electricity demand and may begin to ramp up generation resources hours before higher levels of demand are realized. Likewise, plants may take time to reach peak efficiency as they increase output to meet rising demand. A static analysis, as has been employed in previous empirical estimates of marginal emissions factors, will misattribute emissions that are a direct response to past or expected future changes in demand as a response to current demand, biasing the estimated marginal emissions factor.

This effect can be seen comparing the marginal emissions factors from Graff Zivin, Kotchen,

[5] For example, Cullen (2015) shows non-convexities due to adjustment costs impact the decisions of fossil fuel generators in Electric Reliability Council of Texas (ERCOT). Bushnell and Novan (2018) show increasing solar penetration led to lower midday but higher "shoulder" hour prices.

and Mansur (2014) the average emissions rates in Table 1.[6] Looking at the Western Electricity Coordinating Council (WECC) Interconnection, the marginal emissions factor falls below the average emissions rate for an efficient combined cycle gas turbine (CCGT) in several hours. Conversely, the marginal emissions factor in coal-heavy MRO is frequently much larger than the average emissions rate of an inefficient coal plant. These discrepancies are likely the result of plants adjusting output to coincide changes in the level of demand.

This is not the first analysis to consider the emissions benefits of electricity storage using *ex post* estimates of electricity emissions. Beyond Graff Zivin, Kotchen, and Mansur (2014) described above, Babacan et al. (2018) rely on previous estimates of MEFs to determine the benefits of residential electricity storage. Carson and Novan (2013) compute MEFs for the ERCOT electricity market and combine with a two-period model of the generation decision by electricity producers to find energy storage would increase CO_2 and sulfur dioxide (SO_2) emissions, as fossil fuel generation would shift from cleaner peaking natural gas plants to coal baseload plants.

A further literature considers long-run impacts from the entry of electricity storage. Holladay and LaRiviere (2018) show changes in natural gas prices impact the market and non-market benefits of electricity storage. Linn and Shih (2016) find that the emissions benefits of electricity storage depend on the supply responsiveness of alternative technologies. Abrell, Rausch, and Streitberger (2019) show as the capacity of storage increases the marginal emissions benefits of storage decrease rapidly over day-to-day horizons and are very small over monthly horizons.

3 Data

I have assembled a suite of data on electricity plant and grid operations from 01 January 2016 to 31 December 2018. Data sources include:

Electricity Generator Operations: I obtain hourly operations at electricity generators, including electricity generated, fuel consumed, and emissions of CO_2 and local criteria pollutants from the Environmental Protection Agency (EPA)'s Continuous Emissions Monitoring System (CEMS). These data cover all electricity generators covered by the EPA's Air Markets Program Data (AMPD) with a nameplate capacity of 25 MW or greater.

Electricity Demand: Data on hourly demand, forecast day-ahead demand, and net trades of electricity by balancing authority are available in Form EIA-930 (EIA-930). These data provide uniformly-reported and comprehensive, hourly observations and are available from 2016 to the present in near-real time.

Electricity Prices: I have compiled hourly prices by zone (ERCOT and CAISO) and by pricing node (CAISO) for 2016 to the present in both day ahead and real-time markets from the respective

[6] Marginal emissions factors in Graff Zivin, Kotchen, and Mansur (2014) are presented in lbs CO_2/kWh, which are double the value for tons CO_2/MWh shown on Table 1.

system operators.

Additional Data Sources: Some types of electricity generators inconsistently report generation into CEMS.[7] I adjust output at these plants to match monthly reported fuel consumption and net electricity generation reported in EIA-923 when available. I additional obtain details on electricity generators, such as their geographic location and balancing authority, from Form EIA-860 (EIA-860).

4 Methods

4.1 Model Selection and Evaluation

A central goal of this paper is to develop models which accurately capture the relationship between marginal changes in electricity demand and marginal changes in the emissions of power plants supplying electricity. These MEFs are the most relevant parameters for analyzing policies affecting electricity markets. However, as described in Section 2, I am not the first to argue MEFs are the economically important primitive.

Researchers interested in this relationship have developed a range of models to extract MEFs from observational data. I will take two approaches to modeling MEFs. First, similar to past literature, I will specify linear regressions and use fixed effects to control for unobserved or unmodeled factors influencing both demand and the emissions from electricity generation.

As an alternative to the fixed-effects approach, I will also estimate MEFs using local linear forests (LLFs) proposed in Athey, Tibshirani, and Wagner (2019). This approach allows for rich heterogeneity in marginal emissions factors across the space of covariates and can reveal important patterns in the relationship between demand and emissions.

Under both of these modeling approaches, properly accounting for changes in demand requires a number of decisions on the appropriate empirical specification. Rather than attempt to justify potentially arbitrary modeling choices, I will rely on data-driven procedures, described below, to evaluate models and to compare them to approaches previously deployed in the literature with an eye on selecting models with high predictive power, a minimum of free parameters, and avoiding overfitting.

4.1.1 Cross Validation Procedure

I select my preferred empirical specification and evaluate it against other models utilized in the literature using k-fold cross validation, which is a common model selection algorithm in the machine learning literature.[8] I first randomly divide my estimation data into $k = 10$ similarly-sized and

[7] For example, many CCGTs report generation only from their combustion turbine and omit the steam turbine from CEMS. Generation from both cycles are recorded in Form EIA-923 (EIA-923).

[8] k-fold cross validation has been suggested as a model selection algorithm robust to overfitting since at least Stone (1974).

mutually-exclusive groups, or folds. Demand, emissions, and the state of the electricity generating infrastructure are affected by shocks that are correlated over time and any validation procedure should be robust to that serial persistence. I assign pairs of full weeks, starting at 6 AM UTC on Monday and continuing through 5 AM UTC of the following Monday.[9]

For each candidate model, I estimate the parameters using $k-1$ of the folds. I then generate predictions and compute mean square prediction error using those parameters for observations in the excluded fold. I repeat this procedure k times, excluding a different fold each time. I will prefer models with lower mean square prediction error on the out-of-sample observations averaged across all k sets of estimates. Finally, I reestimate the model across all k folds to estimate the final set of model parameters.

This model selection algorithm has several advantages. It provides a uniform metric for comparing models with very different empirical specifications, will tend to select specifications with the largest predictive power and, since models are evaluated using only out-of-sample predictions, will avoid over-fitting on the estimation sample.[10] Further, by dividing the data into k folds and repeating the estimation and out-of-sample forecasting procedure, I make more efficient use of the available data for the analysis of each model's predictive power.

4.2 MEF Model Specification

I model the relationship between marginal changes in electricity demand and MEFs in the following stages. First I estimate models in the spirit of MEFs using empirical strategies from previous research, including Graff Zivin, Kotchen, and Mansur (2014) and Archsmith, Kendall, and Rapson (2015) using a more recent timeframe (January 2016 through December 2018) and more comprehensive total demand data from Form EIA-930 instead of FERC 713, where applicable. For each of these methods, I compute the mean square forecast error using the cross validation procedure described in Section 4.1.1. Where previous research has made choices with obvious alternatives in their empirical specification, I select the "best" set of possible choices, again using cross-validation.

I next develop an empirical approach to estimating MEFs which attempts to model the process underlying the generation of emissions by electricity generating facilities. This specification differs from the previous literature, accounting the inefficiency of adjusting output at fossil fuel electricity generators by including controls for changes in demand leading up to the current time and expected changes in future demand.

4.3 MEF Local Linear Forest Estimates

As an alternative to fixed-effect models, I also estimate MEFs using LLF described in Athey, Tibshirani, and Wagner (2019). This method is an extension of the random forest estimator,

[9]I convert timestamps in hourly datasets to UTC to provide a consistent measure of time without the complications of time zones or daylight saving time. Generally, when presenting results over time, I will convert to standard time for the appropriate region. Time zones in the continental United States range from UTC-04 (EDT) to UTC-08 (PST).

[10]Stone (1977) demonstrates for models fit by maximum likelihood, model selection using cross-validation is asymptotically equivalent to model selection using the Akaike information criterion (AIC).

replacing the mean moment condition with a linear regression. As a broad outline, estimation of a LLF proceeds as follows:[11]

1. For some dependent variable Y and covariates X
2. Grow tree i
   (a) Randomly select a subsample of observations as training data (in-bag data)
   (b) Randomly select a subset of variables $X^i \subset X$
   (c) Find a partition of some variable in X^i that minimizes the mean squared error of the regression $Y = X^i\beta + \varepsilon$ on the out-of-bag data in each partition
   (d) Repeat this process of partitioning X until some stopping criterion is reached
3. Grow many trees
4. Predicted values are the average of predictions from all trees

LLFs have several desirable properties. First, the average value of a parameter can vary across the full covariate space. This can reveal heterogeneity in parameters without exploration or data mining by the econometrician. Second, the partitioning process places observations with similar parameter values in the same leaves. This is analogous to a locally-weighted least squares procedure where the regression weights are determined using a random-forest tree splitting algorithm.

4.3.1 Random Forest Forecasts

There are cases, described below, where I seek to forecast a value, e.g., expected future demand, at some time $t + s$ as a flexible function of information available at time t. Rather than specify a parametric relationship between information and the forecast, I predict a value using random forests, as described in Breiman (2001).

Random forests use an ensemble of regression trees to approximate even highly nonlinear functions with many interactions between variables and tend to outperform other methods of generating out-of-sample forecasts with minimal tuning of free parameters. Further, Athey, Tibshirani, and Wagner (2019) demonstrate that, under relatively straightforward conditions, random forest forecasts are consistent with asymptotically normal errors.

While I desire accurate forecasts, I am careful to avoid overfitting when using random forests. First, random forest estimation algorithms include regularization components, such as pruning or honest splitting, to reduce overfitting on idiosyncrasies in the data. Second, I always generate forecasts using out-of-sample data. For each of my cross-validation folds, I train a random forest model on the other 9 folds then make predictions on that fold using the model trained out-of-sample.

After estimating LLF models I compute marginal effects of each parameter across the covariate space. For each observation, I perturb a covariate around its true value and compute new predictions. I then approximate the marginal effect using centered finite differences.

[11] I compute LLF estimates using *grf: Generalized Random Forests* from Tibshirani et al. (2018).

4.4 Modeling Non-fossil Fuel Generation

Electricity is generated by both fossil fuel-fired power plants (which generate gaseous GHG and LCP emissions) and other non-fossil fuel generators such as nuclear, hydroelectric dams, solar photovoltaic (PV), or wind turbines. While all of these technologies differ in marginal cost and dispatchability, their end products, electricity fed into the grid, are perfect substitutes.

Much previous work, e.g., Graff Zivin, Kotchen, and Mansur (2014), considers how a marginal change in fossil fuel generation will impact emissions. When evaluating the impact of policies or technology that will change the demand for electricity, however, one should consider how a marginal change in electricity demand (not the portion of demand supplied by fossil fuel plants) will impact emissions. This distinction is important because

1. There may be a supply response from non-fossil fuel sources (such as hydroelectric dams) in response to changes in demand. Considering only fossil fuel generation treats these dispatchable sources as perfectly inelastic supply and will overstate the true impact of changes in electricity demand on emissions.

2. As noted by Archsmith, Kendall, and Rapson (2015) unexpected changes in supply by non-dispatchable generators – e.g., a cloud moving over a PV array – may be correlated with unobserved factors which influence the thermal efficiency of fossil fuel plants. It is important to isolate this type of variation from expected changes in demand for fossil fuel generation.

In each hour and region, I predict current period and expected future electricity generation from non-fossil fuel sources as a flexible function of lagged non-fossil fuel generation (each hour over the previous seven days), the previous hour's expectation of demand in the current hour, the hour of day, the day of year, and the absolute days until/since the closest summer solstice. To fit the model, I train a random forest, of 200 trees, predicting out-of-sample as described in Section 4.3.1.

4.5 Modeling Forward Demand Forecasts

An important component of the MEF model proposed here is accounting for how grid operators respond to past and anticipated future changes in the demand for electricity. Current and past levels of demand are recorded in EIA-930, as are day-ahead forecasts of future demand. Grid operations are not, however, decided strictly on a day-ahead schedule. Significant information on the path of future demand is revealed after the release of the day-ahead forecast as that hour approaches.

Importantly, deviations between forecast and actual demand may be correlated with other factors influencing the emissions rate of the marginal generator. To account for additional information that is revealed leading up to each hour, I generate revised forecasts that incorporate information revealed after the day-ahead forecast is released.

This is a pure forecasting exercise, refining the day-ahead forecast using additional information as it becomes available. Specifically, I estimate models of future demand using random forests. My revised forecast in hour $t + s$ is the day-ahead forecast demand for hour $t + s$ plus a "shock" to demand in that hour, computed as a flexible function of the level of demand in hour t, deviations of actual demand from the day-ahead forecast for that hour for the previous five hours, the hour of day, and the day of the year.

$$E_t[D_{t+s}] = DA_{t+s} + f_s(D_t, D_{t-1} - DA_{t-1}, D_{t-2} - DA_{t-2}, \ldots, D_{t-5} - DA_{t-5}, hod(t), doy(t))$$

As described in Section 4.3.1, I train the random forests on data from the other nine cross validation folds to make predictions on the tenth fold. I treat these revised forecasts as the grid operator's best forecast of future demand. I estimate a separate model for each value of $s \in \{1, 2, 3, 4\}$. This allows the updated forecast to improve as $s \to 0$. *i.e.*, at 9 AM, the revised forecast for 12 PM will incorporate information available as of 9 AM. At 10 AM, a different model using additional information revealed in the past hour generates a new forecast for 12 PM.

In Figure 1 I have plotted forecast values of demand on the horizontal axis against realized values of demand on the vertical axis. Day-ahead forecasts from EIA-930 are shown in black and the revised forecast in orange. Each panel represents a different forecast horizon. The revised forecasts are substantially lower variance than the day-ahead forecasts, reflecting that more information about the future state of demand is available as that hour approaches. Further, the variance of forecast errors from each model decreases as the forecast horizon decreases to zero. Finally, the updated forecasts reduce systematic error in the day-ahead forecasts. Day-ahead forecasts generally over-estimate demand in periods of low demand and under-estimate realized demand in periods of high demand. The updated forecasts reduce this bias and the forecast horizon becomes smaller.

4.6 Model Components

The empirical models of MEFs will include components to account for the level and rate of change in demand and non-fossil fuel generation over time, described in the sections below. When estimating linear fixed effects models, I will also include hour-by-month-of-year and year fixed effects to control for unobserved components correlated with both the demand and the emissions from electricity generation.

4.6.1 Demand Level

As an initial step, I model MEFs on an hourly basis. The basic model considers emissions to be a function of current electricity demand and non-fossil fuel generation. Let E_t be the mass of total emissions, Q^D_{rt} the demand for electricity in region r, and Q^N_{rt} the supply of electricity from non-fossil fuel sources at time t. For each region $r \in R$, each hour-of-day $h \in H$, I the level component of CO_2 emissions at time t (E_t) can be expressed as:

$$\sum_{h \in H} \sum_{r \in R} \left(\beta^0_{rh} Q^D_{tr} - \gamma_{rh} \mathbf{E}_{t-1} \left[Q^N_{tr} \right] \right) \tag{S.1}$$

Here β_{rh} can be interpreted as the marginal effect of a change in demand in region r and hour-of-day h on the total emissions.

4.6.2 Forward Ramp

I consider extending this basic model above by accounting the effects anticipated future changes in demand and non-fossil fuel generation. Here, for some number of leading hours $s \in \{1, 2, \ldots, S\}$ I estimate:

$$\sum_{h \in H} \sum_{r \in R} \sum_{s=1}^{S} \left(\beta_{rh}^{s} \mathbf{E}_t \left[\Delta Q_{t+s,r}^{D} \right] - \gamma_{rh}^{s} \mathbf{E}_t \left[\Delta Q_{t+s,r}^{N} \right] \right) \tag{R.1}$$

Here the coefficient β_{rh}^{s} is the marginal effect on a unit of expected ramp required by hour $t+s$ on total emissions in hour t.

It is important to consider how transient changes in the pattern of demand might impact emissions. Suppose the model considers only 1-hour expected changes in emissions. In the case of a 1-unit anticipated increase to demand in period t, emissions in period t would increase by β_{rh}^{0} units due to the direct effect. This increase in demand in period t would also decrease the expected ramp needed moving to period $t+1$, thus emissions would decrease by β_{rh}^{1}. Finally, since this shock to demand is fully anticipated, additional ramp would be required in period $t-1$ and emissions would increase $\beta_{r,h-1}^{1}$ as a result of that additional ramp. Thus, the full marginal effect of a 1-unit increase in demand in period t would be:

$$\frac{\partial E}{\partial Q_{tr}^{D}} = \beta_{rh}^{0} - \beta_{rh}^{1} + \beta_{r,h-1}^{0}$$

4.6.3 Past Ramp

I also consider models where past changes in demand leading up to the present, either anticipated or unexpected, may impact the current MEF. It is reasonable that the impact on emissions may differ between anticipated demand changes, which can be served by slow or quick adjusting supply, and unanticipated shocks, which can only be served by fast-adjusting sources. Here, I divide the change in demand over the previous hour into an anticipated shift to demand (A_t^D) – defined as the difference between realized demand at $t-1$ and the $t-1$ forecast of demand at time t – and an unanticipated shock to demand (S_t^D) or

$$Q_t^D - Q_{t-s}^D = \underbrace{Q_t^D - \mathbf{E}_{t-s}[Q_t^D]}_{\text{Unanticipated Shock } (S_t^D)} + \underbrace{\mathbf{E}_{t-s}[Q_t^D] - Q_{t-s}^D}_{\text{Anticipated Change } (A_t^D)}$$

In these models, I likewise account for adjustments in the non-fossil fuel generating sector by similarly decomposing changes in non-fossil fuel generation into expected shifts and an unanticipated shock. I then estimate:

$$\sum_{h\in H}\sum_{r\in R}\left(\phi^*_{rh}S^D_t+\psi^*_{rh}S^N_t\right) \tag{R.1}$$
$$+\sum_{h\in H}\sum_{r\in R}\sum_{s=1}^{S}\left(\beta^s_{rh}\mathbf{E}_t\left[\Delta Q^D_{t+s,r}\right]-\gamma^s_{rh}\mathbf{E}_t\left[\Delta Q^N_{t+s,r}\right]+\phi^{-s}_{rh}A^D_t+\psi^{-s}_{rh}A^N_t\right)$$

In this class of models the impact of a change in demand on emissions depends on whether the change was anticipated in the previous hour. Both anticipated and unanticipated demand changes may be relevant for analyzing the impact of policies impacting electricity supply and demand. For example, a household adding a rooftop solar PV array will steadily and predictably reduce its net demand for electricity in the morning hours as the sun rises, and then increase (as compared to the no-solar counterfactual) as the sun begins to set in the afternoon. This new pattern or residual demand can be fully anticipated and would not alter the shock components S^D_t or S^N_t. In this case, the marginal impact of a 1-unit increase in demand at time t for one hour is the sum of anticipated demand effects less the additional expected ramp required in the future

$$\frac{\partial E}{\partial Q^D_{tr}}=\beta^0_{rh}+\sum_{s=1}^{S}\left(\phi^{-s}_{rh}+\beta^s_{rh}\right)-\sum_{s=1}^{S}\left(\beta^s_{r,h+s}+\phi^{-s}_{r,h+s}\right)$$

This same solar array, however, may also generate unpredictable changes in demand, e.g., if a cloud passes over and reduces output. In this case, the household's net demand for electricity may unexpectedly increase, raising S^D_t. Unanticipated shocks to demand at time t, however, do not alter expected demand prior to realization of the shock, but do reduce the need for additional ramp in the future. Thus, in the case of a 1-hour shock to demand in period t, the marginal impact on emissions is

$$\frac{\partial E}{\partial Q^D_{tr}}=\beta^0_{rh}+\phi^*_{rh}-\sum_{s=1}^{S}\left(\beta^s_{r,h+s}+\phi^{-s}_{r,h+s}\right)$$

4.7 Local Linear Forest Models

I will allow for similar dependent variables in the local linear forest models. I will also allow the algorithm to split leaves based on hour-of-day, month-of-year, day-of-week, and year.[12] Since LLF estimates a separate regression in each leaf, it is not necessary to include hour-of-day interactions with the independent variables. If hour-of-day is an important component in modeling demand, the tree-splitting algorithm will partition leaves on values of the hour-of-day variable.

[12] The tree-splitting algorithm will only make partitions of the form $x > x^*$. The hour-of-day, month-of-year, and day-of-week are most naturally thought of as rings, not lines. I therefore include a second set of variables where the value of the original variable is offset by the median value modulo the range of the variable. This allows, for example, a split on month-of-year to include both December and January in the same leaf.

5 Results

In the following section, I will describe results of both the fixed effects and LLF modeling procedures. In each case, I start by by evaluating the out-of-sample forecast accuracy of a range of models, selecting a preferred model that is both parsimonious and which minimizes forecast errors. I then graphically demonstrate the range of estimated MEFs from these preferred models.

5.1 Fixed Effects Estimates

First, I present results of fixed effects estimates of the MEFs.

5.1.1 Model Fit

The set of criteria for MEF models described above lead to a large space of candidate specifications. It is *a priori* unclear whether or how many leads of expected forward ramp, lags of past ramp, or demand shocks should be included in a proper empirical model of MEFs. Given the large space of parameters, accepting or rejecting models due to the presence of statistically significant parameters is likely to lead to including spurious factors in the model.

As an alternative, I estimate a range of models implied by the above specifications and compute out-of-sample fit for each using ten-fold cross validation. Overfit models or those with extraneous parameters will have poorer out-of-sample forecast performance than a properly-specified model. Further, out-of-sample forecast error provides a consistent measure of model fit, regardless of the underlying specification or estimation routine.

Model fit statistics for the set of linear models using hourly coefficients are presented in Table 2 with models of the ERCOT in Panel 2a and the WECC in Panel 2b. For both interconnections, including past ramp substantially reduces forecast error and forward ramp reduces it further. Estimating a unique set of parameters for each quarter again substantially reduces forecast error in WECC interconnection, but has minimal impact, or increases, out-of-sample forecast error in the ERCOT interconnection. This is consistent with the large seasonality in renewables generation (both solar PV and hydroelectric) in WECC. Finally, likelihood-ratio tests reject the null that models including ramp are nested within the no-ramp models.

The improvement in out-of-sample forecast accuracy may seem small compared to the mean level of CO_2 emissions. These effects account for approximately 6% of the forecast error of the static model. This is a large effect considering the bulk of the ramping effect occurs on the margin. The bulk of emissions from electricity generators are produced by inframarginal units. Hour-to-hour changes in demand are small compared to the mass of inframarginal generation, but any policy that causes a transitory 1 MW increase in the level of demand will also generate an additional 1 MW of ramp.

5.1.2 Marginal Emissions Factors

From the out-of-sample fit statistics, it is clear accounting for changes in demand – and the corresponding ramp in fossil fuel plants – explains total emissions better than models that ignore these effects. It is useful to compare these estimated MEFs and ramping effects to models that ignore ramp.

Figure 2 and Figure 3 show estimated MEFs in the ERCOT and CAL regions for each hour of the day from models ignoring ramping effects (top) and including ramping effects (bottom). The blue bars represent the marginal effect of a steady 1 MW increase in electricity demand on CO_2 emissions. In the bottom panel, the orange bar shows the ramping effect, specifically the marginal effect of a 1 MW increase in demand from the previous hour to the current on emissions. A hypothetical increase in demand now means that less ramp up (or more ramp down) will be required in the future when the hypothetical increase ends. The gray bars show the effect of avoided future ramp assuming the change in demand lasts 1 hour. The black line shows the sum of marginal effects of this hypothetical 1-hour increase in demand with the 95% confidence interval shaded gray.

In Figure 2 you can see a static model finds only small variation in the MEF across hours. Looking at Panel 2b, however, ramping effects become evident. The effect of transient increases in load are small in the evening and early morning when onshore wind generation is the highest and fossil fuel generators are ramping down in response.

Similarly, Panels 3a and 3b of Table 3 shows estimated MEFs in California from models ignoring and including ramp, respectively. Here, ramping effects are primarily driven by solar generation. They are negative as the sun rises in the morning and then become positive in the late afternoon as the sun sets.

Figure 2 and Figure 3 show the effect of transient changes in electricity demand lasting one hour. Policies intended to alter the sources or electricity supply or change patterns of demand, however, can have impacts that last many hours. For example, a solar PV panel will start generating electricity as the sun rises, steadily increase output until it reaches peak output, then decrease output as the sun sets. This creates a pattern of electricity output impacting both the level and rate of change in demand falling on fossil fuel plants. Likewise, an electricity storage facility could charge for several consecutive hours when fossil fuel plants are ramping down and renewable generation is plentiful then discharge as fossil fuel plants are ramping back up again.

The impacts of ramping on emissions are particularly evident when considering changes in supply or demand that persist over many hours. Figure 4 shows the marginal impact by hour of day for a 1 MW increase in electricity demand that persists for 4 or 8 hours in the CAL region. Here the impact of solar PV-induced ramp is quite pronounced. A 4-hour increase in demand initiated at 7 AM leads to half the CO_2 emissions of the same change in demand initiated at 5 PM.

These graphs also highlight the importance of examining the margin when considering ramping effects. Ramping impacts emissions when the hypothetical change in demand starts and ends. For loads sustained for long periods, the sustained effect substantially outweighs the ramping effects.

However, for transient changes in supply or demand, as may be provided by electricity storage systems or non-dispatchable renewables, ramping effects are a significant component of emissions on the margin.

5.2 Local Linear Forests Estimates

As an alternative to ordinary least squares (OLS)-based fixed-effects models, I also estimate marginal emissions factors using LLF with results described in the following section.

5.2.1 Model Fit

Estimating models of fossil fuel emissions using local linear forests generally leads to superior out-of-sample forecast performance over traditional linear models. The out-of-sample RMSE of several local linear forest models are shown in Table 3 below.[13]

The model including 2-hour leads and lags in changes of demand (and non-fossil fuel supply) leads to the best out-of-sample fit. Further, out-of-sample fit from all of these models is substantially lower than the OLS-based fixed effects models. Analyses and simulations that follow will be based on estimates from the 2-hour ramp model.

5.2.2 Marginal Emissions Factors

LLF estimate a local linear regression at each observation, using a tree-splitting algorithm to determine weights of "nearby" observations. This allows for the estimated marginal effects to vary substantially across the covariate space and reveal substantial heterogeniety in marginal emissions factors across both levels and changes in demand over time.

Panel (a) of Figure 5 shows the marginal effect of a change in demand on CO_2 emissions as a function of the level of demand. The marginal effect of increasing demand is large at very low levels of and falls to a minimum for moderate levels of demand before increasing again at very high levels of demand. Both the pattern and levels of marginal emissions factors are consistent with dispatch of generators in the ERCOT region in order of marginal costs. At low levels of demand, coal-fired generators – with high emissions – are the marginal generators. For much of the observed space of demand natural gas generators – with lower emissions rates – are marginal. Finally, at high levels of demand, less efficient peaker plants are the plants adjusting on the margins.

Panels (b) and (c) of Figure 5 show the marginal impact on emissions of anticipated changes in demand over the next hour (hour t to $t+1$) or two hours (hour $t+1$ to $t+2$) in the future. Importantly, the marginal effects are always positive and can be large in magnitude. Typically, an additional 1 MWh of demand will lead to an additional 0.6 tons of CO_2 emissions in a steady state. A 1 MW change in demand, however, will cause at least an additional 0.2 tons of CO_2 emissions

[13] While estimating LLFs is computationally simple, generating predictions and computing marginal effects from these models is very computationally demanding. I currently present results for only the ERCOT interconnection, but I am in the process of extending these results to the WECC and Eastern interconnection.

as plants ramp up to meet that additional demand. On the margin, the emissions from ramping are comparable in magnitude to the steady-state emissions rate.

6 Analysis

Using the model selection procedure in Section 5, I can identify the parametric specification that most accurately captures the relevant determinants of emissions from electricity generation. Using parameters from this preferred model, I can then simulate the marginal impact of policies that change the electricity generation mix.

6.1 Fixed Effect Model Simulation Details

For these simulations, I will use the fixed effects model specification including hourly measures of current demand, the previous hour shock to demand, and the expected change in demand over the previous hour ("Shock+Past Ramp"). For the ERCOT region, I use coefficients estimated for each hour of day. In the WECC region the coefficients vary by hour of day and quarter of year. I will revise these preferred models if new specifications yield better out-of-sample fit.

I simulate the emissions benefits of non-dispatchable generation by computing a typical output profile for either solar PV or onshore wind. I derive the output profile using actual solar and wind generation data provided by California Independent System Operator (CAISO). I also simulate the emissions benefits of an electricity storage facility. In these simulations, I assume the storage facility has 80% roundtrip efficiency, must charge one hour for each hour it discharges, does not store electricity for more than 24 hours, and operates to minimize CO_2 emissions.

6.2 Simulation Results

The simulated CO_2 emissions reductions for the marginal unit of onshore wind, solar, and battery storage are shown in Table 4 to Table 7. In the ERCOT region (Table 4) similar units are on the margin throughout the day and ramping effects are relatively small. Here, there is little difference in predictions on the emissions reductions from onshore wind or solar PV generation from MEF models ignoring the ramping effect and those accounting for it. Note, however, that models ignoring ramping predict a CO_2-minimizing storage facility would never operate in ERCOT, there are positive emissions benefits from storage when ramping is taken into account.

Results are markedly different in regions within the WECC interconnection. While the emissions benefits from each approach are similar for onshore wind generation, ignoring ramping substantially over estimates the emissions benefits of solar. These regions currently have high penetration of solar PV generation, which displaces fossil fuel generation during day, requiring substantial ramp-up of fossil fuel generation as the sun sets. Additional units of solar exacerbate the need for late-afternoon ramp, diminishing their emissions benefits. While the marginal unit of solar generation still reduces CO_2 emissions, estimates ignoring ramp effects overstate the benefits by over 16% in both California and the Pacific Northwest.

By the same token, ignoring ramp substantially understates the emissions benefits of electricity storage. Ignoring ramping effects, energy storage appears to have small emissions benefits throughout the WECC. Accounting for ramping effects substantially increases these benefits. Energy storage reduces the need for fossil fuel ramp quickly in the evening – on the margin the emissions benefits are large. Further, storage units can charge in the morning as solar generation comes online, reducing the rate at which fossil fuel plants must ramp down.[14] On net, failing to account for ramping underestimates the CO_2 emissions benefits of storage by around 90%.

6.3 Electricity Storage Simulations

The simulation in the previous section demonstrates accounting for ramping effects are critical for evaluating policies designed to alter the electricity generation mix. The emissions benefits of electricity storage, however, assume it is operated with a CO_2-minimizing objective. The behavior of a profit-maximizing storage operator, however, could be entirely different.

The LLF models described in Section 4.7 provide a rich platform for simulating the behavior of a hypothetical electricity storage provider. In each hour of 2017, I compute the marginal impact of level and ramp changes in electricity demand in the ERCOT interconnection. Then using day-ahead prices in the ERCOT market, I solve for the perfect-foresight behavior of a hypothetical small electricity storage operator. The technical properties of the electricity storage are the same as in Section 6.2. Energy capacity is ten times power capacity with 80% round-trip efficiency.

Under each scenario, I compute the impact on CO_2 emissions at the margin from a storage operated under under two objective functions. First, similar to Section 6.2, storage is operated to minimize CO_2 emissions from fossil fuel generation. I then compare this to the behavior of a private storage provider where the storage operator has perfect foresight of day-head prices and chooses the path of charge and discharge behavior that maximizes profits while respecting constraints that the quantity of stored electricity is never less than zero or greater than the maximum capacity.

Results of these simulations are shown in Table 8. The first three columns show the net change in CO_2 emissions per MWh of electricity stored for a profit-maximizing operator. Ramping effects are clearly important; failing to account for ramp predicts electricity storage would *increase* total CO_2 emissions. However, accounting for ramping effects, even profit-maximizing storage leads to modest decreases in emissions.

Emissions reductions from CO_2-minimizing storage are somewhat larger than under the fixed-effects simulations. There are two source for this difference. First, there is substantially larger variability in hour-to-hour MEFs when applying LLF estimates to each hour of data compared to average hourly MEFs estimated using fixed-effects models. Second, this simulation computes the dynamically optimal path over a full year, allowing for much more flexibility in charge/discharge behavior. Nevertheless, failing to account for ramping effects still underestimates the benefits of CO_2-minimizing electricity storage by 54%.

[14]Emissions benefits in CAL are consistent with avoided ramping of a gas peaker plant, in SW with a ramping coal plant, and a CCGT in the NW region. These are all consistent with the electricity generating stock in these regions.

In Figure 6, I plot the probability in each hour of the day a hypothetical electricity storage operate would charge discharge under profit-maximization and emissions minimization. Panel (a) plots storage in the Houston region, where transmission is frequently constrained due to high local demand, and the West region in Panel (b) which has large wind generation resources but low demand.

From these diagrams it is clear that prices send market signals counter to an emissions minimizing objective. Particularly, an emissions minimizer would charge during the afternoon hours when demand is relatively stable and efficient natural gas units are on the margin, then discharge in the early evening and early morning when plants are generally ramping to meet increased demand. A profit-maximizer, however, would generally charge at night when prices and demand are low but coal plants are marginal, then discharge in the afternoon, offsetting relatively cleaner natural gas generation.

7 Conclusion

Electricity generation in the United States is impacted by myriad policies designed to alter the sources of supply and patterns of demand for electricity with the goal of improving ambient air quality and reducing GHG or other LCP emissions. Subsidies for generating renewable energy, renewable portfolio standards, and storage mandates are often used as alternatives to the first-best, but politically-infeasible Pigouvian tax on emissions. Understanding precisely how these second-best policies impact emissions is critical to achieving environmental goals at the lowest possible cost.

This paper extends the literature computing *ex post* estimates of MEFs for electricity generation in the United States. Importantly, I account for the fact that emissions from fossil fuel power plants depend not only on the level of demand, but the rate at which demand is changing over time. Many renewable energy and storage technologies provide energy in predictable patterns, altering the level of demand and systematically altering the rate of change in demand over time.

Accounting for these changes in demand (ramp) are critical to understanding the benefits of new renewable generation or electricity storage. Ignoring ramp overstates the CO_2 emissions reduction of a solar PV array in the Western United States by 16%. Further, emissions reductions from electricity storage systems optimized to minimize CO_2 are at least double is predicted from models ignoring ramping effects.

Introducing and properly accounting for ramping effects requires estimating models with hundreds of free parameters and, potentially, dozens of decisions by the econometrician on the measures of ramp to include. Borrowing from the machine learning literature, I choose a preferred model from a large space of candidate specifications by selecting the model that minimizes out-of-sample forecast error using cross-validation. This procedure will select features that most accurately model power plant emissions and guard against overfitting. Further, out-of-sample fit provides a consistent metric for comparing diverse, non-nested models.

Work presented in this paper is ongoing. LLF estimates reveal substantial heterogeneity in MEFs across levels of demand and ramping rates, and other factors, such as hour-of-day, have little impact on predicted emissions. These models show promise for providing further insight into the factors driving electricity emissions. While estimating LLF models is straightforward, computing predictions and marginal effects are computationally demanding. Results presented here for the ERCOT interconnection are a proof of concept and I plan to extend these methods to the remaining US interconnections.

Results presented in the manuscript currently exclude the Eastern Interconnection of the United States. This interconnection is comprised of many more balancing authorities than ERCOT or WECC, and requires additional attention to properly match supply and demand data. I currently only consider GHG emissions from fossil fuel generation. These plants also emit other LCPs, namely NO_X and SO_2. I will extend the analysis to estimate impacts on the emissions of LCPs in space and time.

Finally, the results presented here and my ongoing work suggest it is important to consider ramping effects when setting policy in electricity markets. Costly ramping will tend to increase prices when plants are ramping up and decrease prices when plants are ramping down. Non-dispatchable renewables, such as solar PV, by the intermittent nature of their production profile, will induce ramping at fossil fuel plants. Since renewable production is correlated with the induced ramp, markets may provide an incomplete signal of ramping costs to renewable generators, distorting investment incentives. Further, since fossil fuel ramping generates additional emissions, typical market instruments for addressing externalities may not be able to achieve first best. I plan to further explore this relationship between ramping costs, emissions, prices, and investment incentives.

References

Abrell, Jan, Sebastian Rausch, and Clemens Streitberger (2019). "Buffering Volatility : Storage Investments and Technology-Specific Renewable Energy Support." In: pp. 1–21.

Archsmith, James (2017). "Dam Spillovers: The Direct and Indirect Costs from Environmental Constraints on Hydroelectric Generation." In: *Ssrn.* DOI: 10.2139/ssrn.3046246.

Archsmith, James, Alissa Kendall, and David Rapson (2015). "From Cradle to Junkyard: Assessing the Life Cycle Greenhouse Gas Benefits of Electric Vehicles." In: *Research in Transportation Economics* 52, pp. 72–90. ISSN: 07398859. DOI: 10.1016/j.retrec.2015.10.007. URL: http://dx.doi.org/10.1016/j.retrec.2015.10.007.

Athey, Susan, Julie Tibshirani, and Stefan Wagner (2019). "Generalized random forests." In: *The Annals of Statistics* 47(2), pp. 1148–1178.

Babacan, Oytun, Ahmed Abdulla, Ryan Hanna, Jan Kleissl, and David G. Victor (2018). "Unintended Effects of Residential Energy Storage on Emissions from the Electric Power System." In: *Environmental Science and Technology* 52(22), pp. 13600–13608. ISSN: 15205851. DOI: 10.1021/acs.est.8b03834.

Borenstein, Severin and James Bushnell (2018). "Are Residential Electricity Prices Too High or Too Low? Or Both?" In: *NBER Working Paper*(24756), pp. 1–51. DOI: 10.3386/w24756. URL: http://www.nber.org/papers/w24756.

Breiman, Leo (2001). "Random Forests." In: *Machine Learning* 45, pp. 5–32. DOI: 10.1023/A:1010933404324.

Bushnell, James and Kevin Novan (2018). "Setting with the Sun : The Impacts of Renewable Energy on Wholesale Power Markets." In: *NBER Working Paper* 24980. DOI: 10.3386/w24980.

Carson, Richard T. and Kevin Novan (2013). "The private and social economics of bulk electricity storage." In: *Journal of Environmental Economics and Management* 66(3), pp. 404–423. ISSN: 00950696. DOI: 10.1016/j.jeem.2013.06.002. URL: http://dx.doi.org/10.1016/j.jeem.2013.06.002.

Cullen, Joseph A. (2015). "Dynamic Response to Environmental Regulation in the Electricity Industry." In: (Eia 2008), pp. 1–40. URL: http://www.josephcullen.com/resources/dynamicresponse.pdf.

Fang, Yingkai, Frank Asche, and Kevin Novan (2018). "The costs of charging Plug-in Electric Vehicles (PEVs): Within day variation in emissions and electricity prices." In: *Energy Economics* 69, pp. 196–203. ISSN: 01409883. DOI: 10.1016/j.eneco.2017.11.011. URL: https://doi.org/10.1016/j.eneco.2017.11.011.

Graff Zivin, Joshua S., Matthew J. Kotchen, and Erin T. Mansur (2014). "Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies." In: *Journal of Economic Behavior and Organization* 107, pp. 248–268. ISSN: 01672681. DOI: 10.1016/j.jebo.2014.03.010. URL: http://dx.doi.org/10.1016/j.jebo.2014.03.010.

Holladay, Scott and Jacob LaRiviere (2018). "How does welfare from load shifting electricity policy vary with market prices? Evidence from bulk storage and electricity generation." In: *Energy Journal* 39(6), pp. 235–272. ISSN: 01956574. DOI: 10.5547/01956574.39.6.jhol.

Holland, Stephen P. and Erin T. Mansur (2008). "Is Real-Time Pricing Green? the Environmental Impacts of Electricity Demand Variance." In: *The Review of Economics and Statistics* 90(August), pp. 550–561.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates (2016). "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors." In: *American Economic Review* 106(12), pp. 3700–3729. ISSN: 0002-8282. DOI: 10.1257/aer.20150897. URL: http://econpapers.repec.org/RePEc:aea:aecrev:v:106:y:2016:i:12:p:3700-3729.

Linn, Joshua and Jhih-Shyang Shih (2016). "Does Electricity Storage Innovation Reduce Greenhouse Gas Emissions?" In: DOI: 10.2139/ssrn.2851952.

Novan, Kevin (2015). "Valuing the wind: Renewable energy policies and air pollution avoided." In: *American Economic Journal: Economic Policy* 7(3), pp. 291–326. ISSN: 1945774X. DOI: 10.1257/pol.20130268.

Siler-Evans, Kyle, Lima Azevedo, and M Granger Morgan (2012). "Marginal Emissions Factors for the U.S. Electricity System." In: *Environ. Sci. Technol* 46, p. 4748. DOI: 10.1021/es300145v.

Stone, M. (1974). "Cross-Validatory Choice and Assessment of Statistical Predictions." In: *Journal of the Royal Statistical Society. Series B (Methodological)* 36(2), pp. 111–147.

Stone, M. (1977). "Assessment and Propagation of Model Uncertainty." In: *Journal of the Royal Statistical Society. Series B (Methodological)* 39(1), pp. 44–47. URL: https://www.jstor.org/stable/2984877.

User-contributed software

This section lists user-contributed, non-commercial software used in analyses supporting this research.

Baum, Christopher F and Mark E Schaffer (2015). *AVAR: Stata module to perform asymptotic covariance estimation for iid and non-iid data robust to heteroskedasticity, autocorrelation, 1- and 2-way clustering, and common cross-panel autocorrelated disturbances.* URL: https://ideas.repec.org/c/boc/bocode/s457689.html.

Baum, Christopher F, Mark E Schaffer, and Steven Stillman (2002). *IVREG2: Stata module for extended instrumental variables/2SLS and GMM estimation.* URL: https://ideas.repec.org/c/boc/bocode/s425401.html.

Correia, Sergio (2014). *REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects.* URL: https://ideas.repec.org/c/boc/bocode/s457874.html.

Tibshirani, Julie et al. (2018). *grf: Generalized Random Forests.* URL: https://cran.r-project.org/web/packages/grf/index.html.

Public Data Sources

This section lists publicly-available, non-commercial datasets prepared by individuals or researchers used in this research.

NREL (2019). *PVWatts*. Golden, CO. URL: https://pvwatts.nrel.gov/pvwatts.php.

Figure 1: Comparison of day-ahead and estimated hourly demand forecasts - CAL region

(a) 4-hour ahead forecast

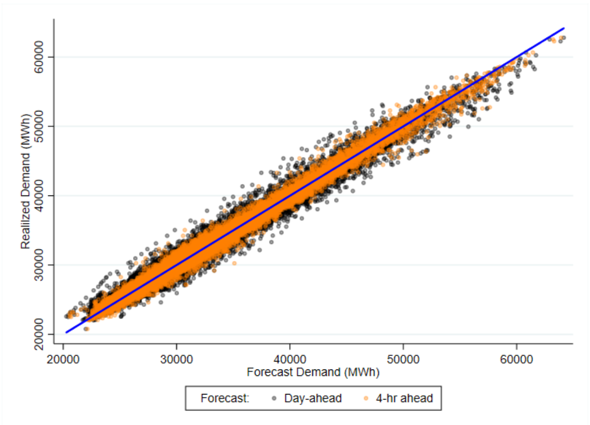

(b) 3-hour ahead forecast

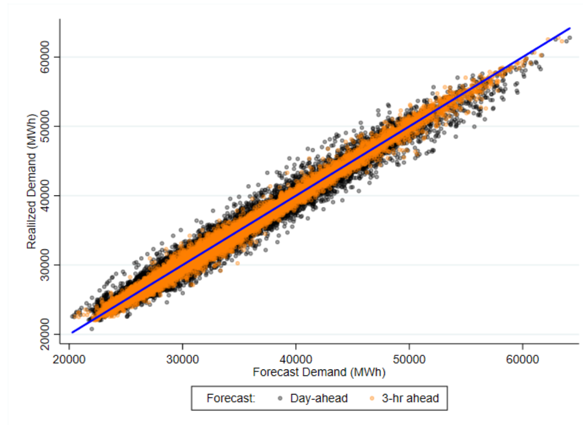

(c) 2-hour ahead forecast

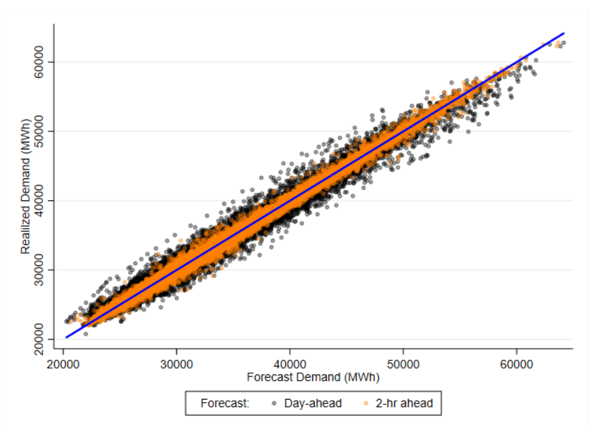

(d) 1-hour ahead forecast

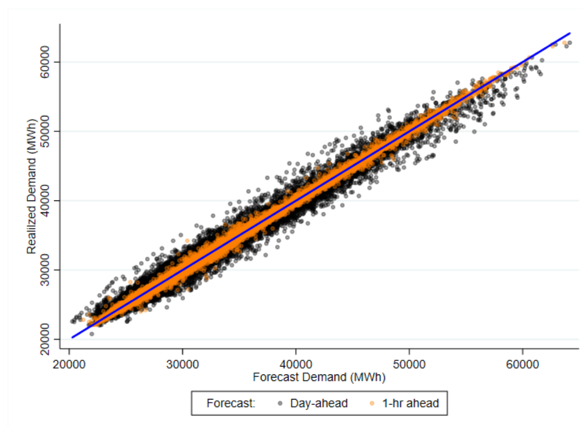

Scatter plot of forecast demand against realized demand. Black dots use the day-ahead forecast. Orange dots show the updated forecast from 4,3,2 and 1 hour ahead. Updated forecasts are computed using out-of-sample predictions from a random forest model using deviations from the forecast demand for the previous five hours, the hour of day, and the day of the year. The blue line represents a perfect forecast. Day-ahead forecasts have an R-square of 0.978. Updated forecasts have R-squares of 0.989, 0.992, 0.995, and 0.998 for 4,3,2, and 1-hour ahead forecasts, respectively.

Figure 2: Estimated MEF for ERCOT Region

(a) Ignoring Ramping Effects

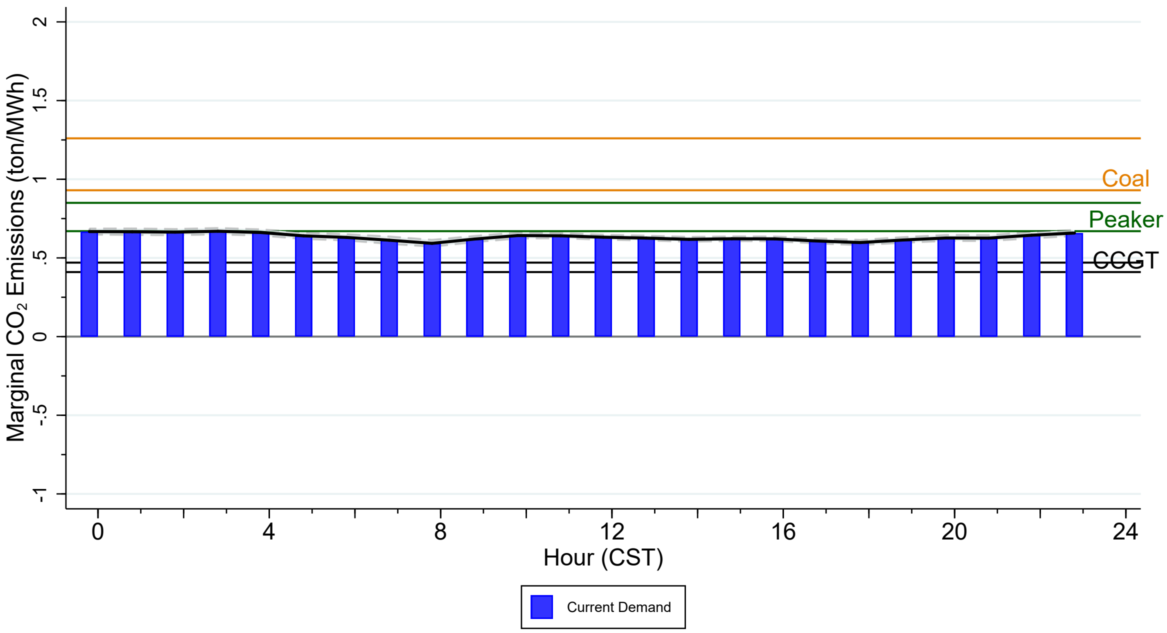

(b) Including Ramping Effects

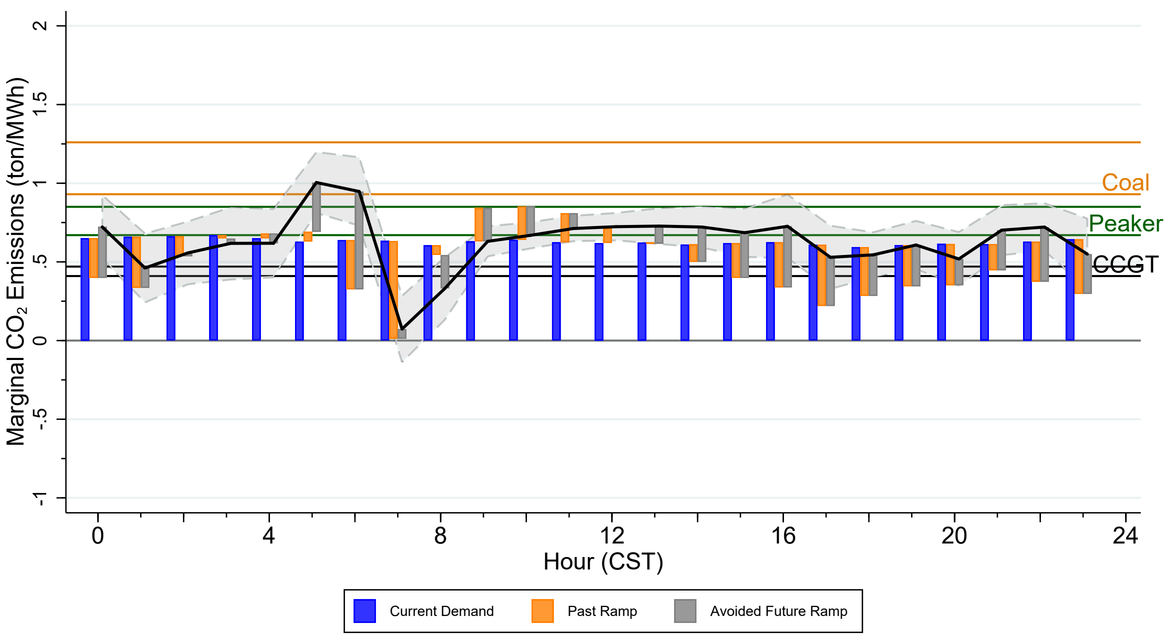

Marginal effect of level and change in electricity demand on CO_2 emissions by hour in the ERCOT region. Panel 2a excludes controls for ramping effects. Panel 2b adds controls for expected and unanticipated changes in demand over the previous hour. All models include hour-by-month of year and year fixed effects and control for non-fossil fuel generation. The gray shaded area shows the 95% confidence interval of the total effect of a 1-hour increase in demand of 1 MW, robust to heteroskedasticity and autocorrelation of up to 24 hours. Bars in the background represent the typical range of MEF from CCGTs, natural gas or petroleum peaker plants, and coal generation.

Figure 3: Estimated MEF for CAL Region

(a) Ignoring Ramping Effects

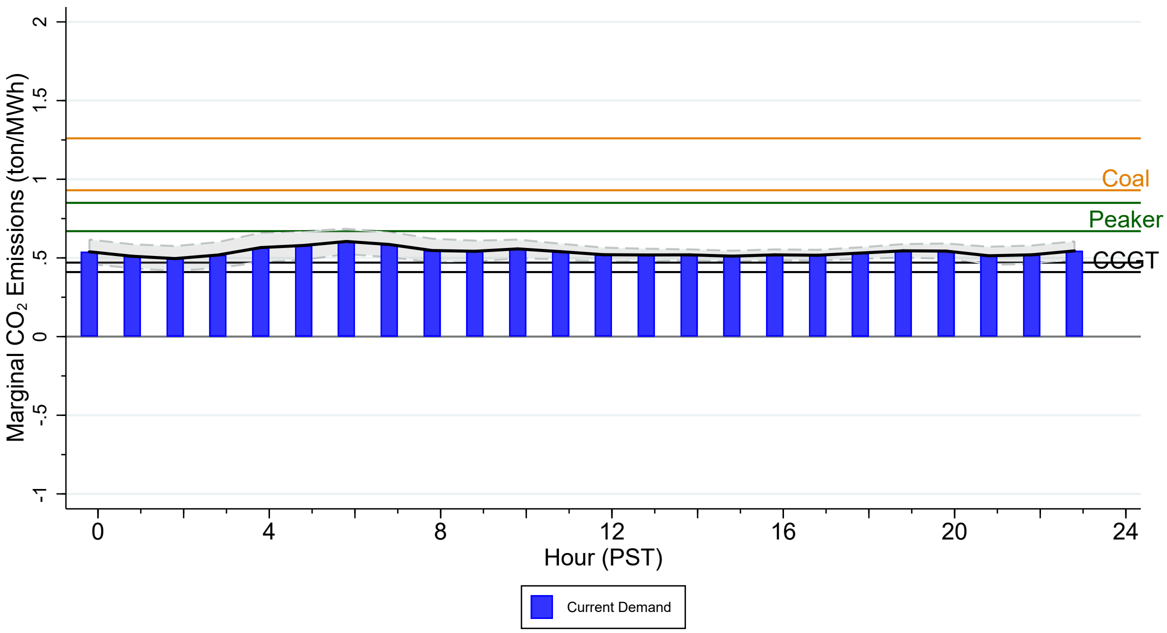

(b) Including Ramping Effects

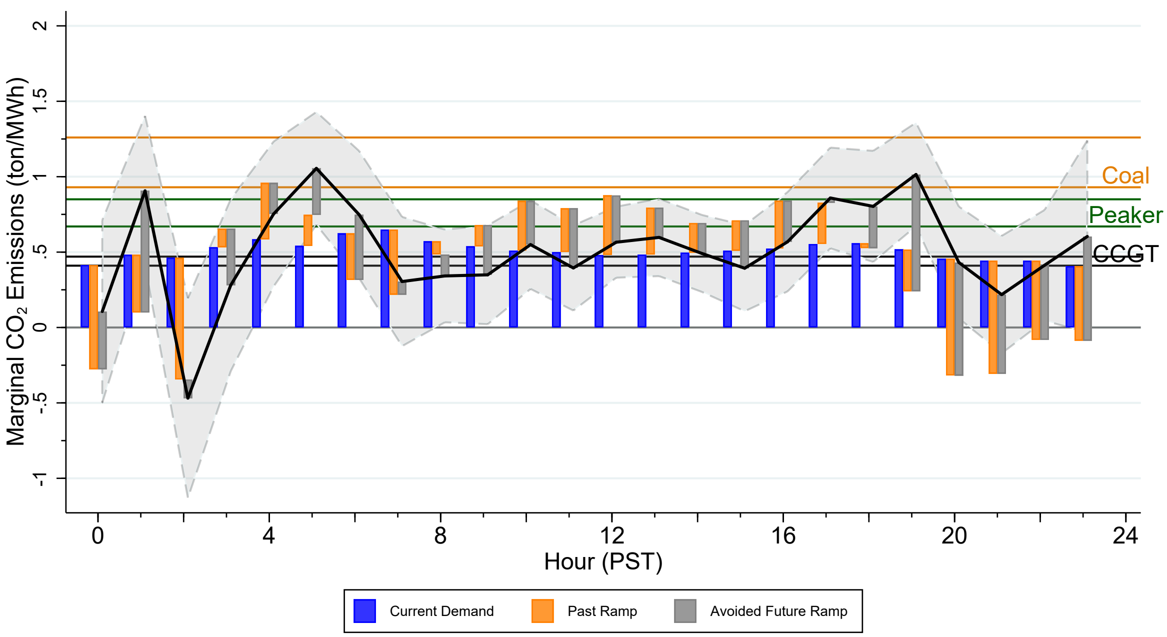

Marginal effect of level and change in electricity demand on CO_2 emissions by hour in the CAL region. Panel 3a excludes controls for ramping effects. Panel 3b adds controls for expected and unanticipated changes in demand over the previous hour. All models include hour-by-month of year and year fixed effects and control for non-fossil fuel generation. The gray shaded area shows the 95% confidence interval of the total effect of a 1-hour increase in demand of 1 MW, robust to heteroskedasticity and autocorrelation of up to 24 hours. Bars in the background represent the typical range of MEF from CCGTs, natural gas or petroleum peaker plants, and coal generation.

Figure 4: Estimated MEF for CAL Region

(a) 4 hours of sustained demand

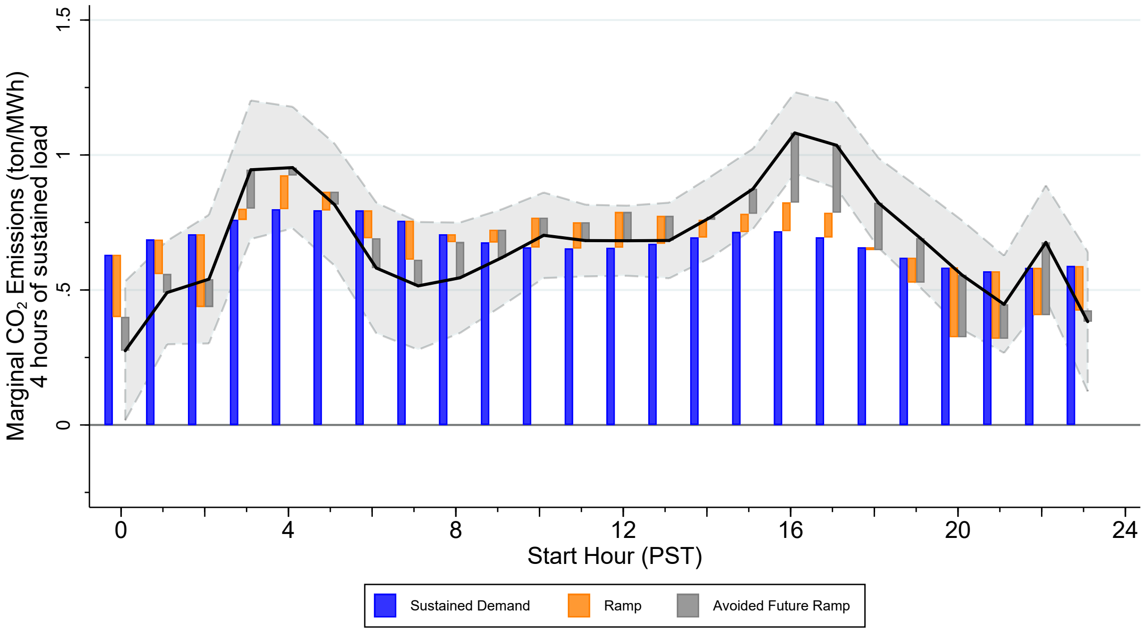

(b) 8 hours of sustained demand

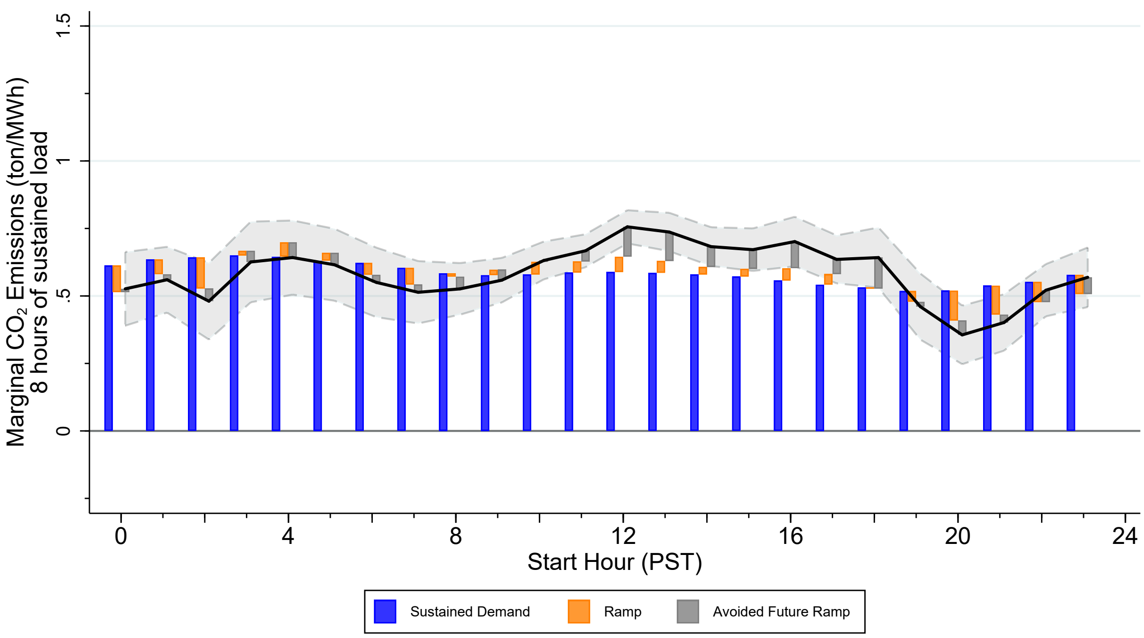

Marginal effect of level and change in electricity demand on CO_2 emissions by hour in the CAL region for an increase in demand persisting for 4 and 8 hours. All models include hour-by-month of year and year fixed effects and control for non-fossil fuel generation. The gray shaded area shows the 95% confidence interval of the total effect of a 1-hour increase in demand of 1 MW, robust to heteroskedasticity and autocorrelation of up to 24 hours. Bars in the background represent the typical range of MEF from CCGTs, natural gas or petroleum peaker plants, and coal generation.

Figure 5: Estimated CO_2 MEF for ERCOT region

(a) Demand Level

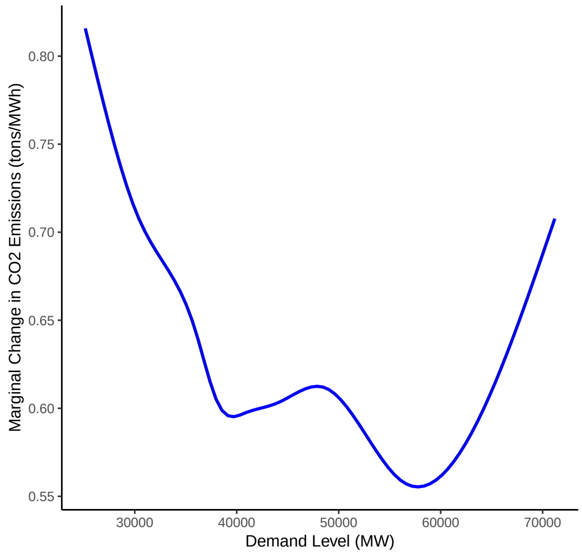

(b) Anticipated 1-hour Change in Demand Next Hour

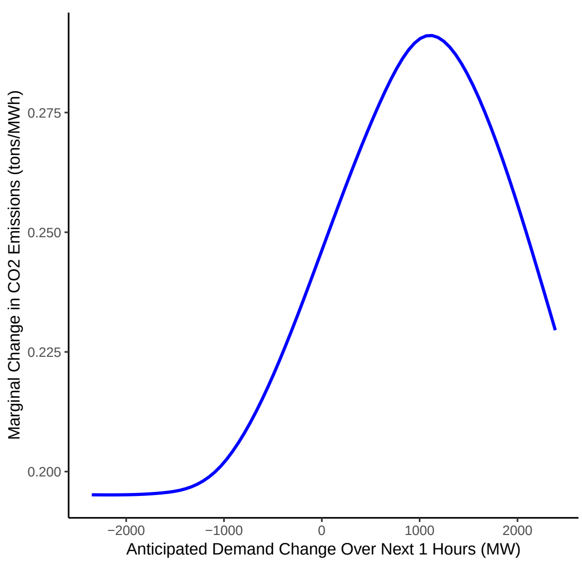

(c) Anticipated 1-hour Change in Demand Two Hours Hence

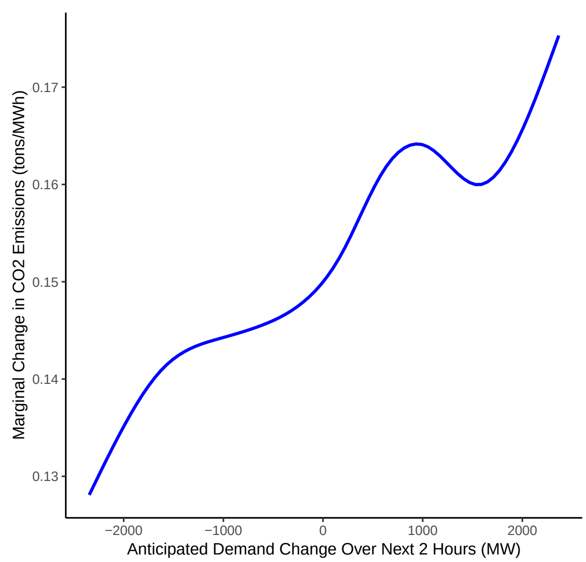

Marginal change in CO_2 emissions from a marginal change in the level of demand (Panel (a)) or anticipated 1-hour changes in demand over the next hour (Panel (b)) or hours $t+1$ to $t+2$ (Panel (c)). Marginal effects from a local linear forest as described in Section **??**. In each graph the horizontal axis is limited to the 5th to 95th percentile range of actual values observed from 2017 to 2018.

Figure 6: Simulated Electricity Storage Behavior - ERCOT Interconnection

(a) Houston Zone

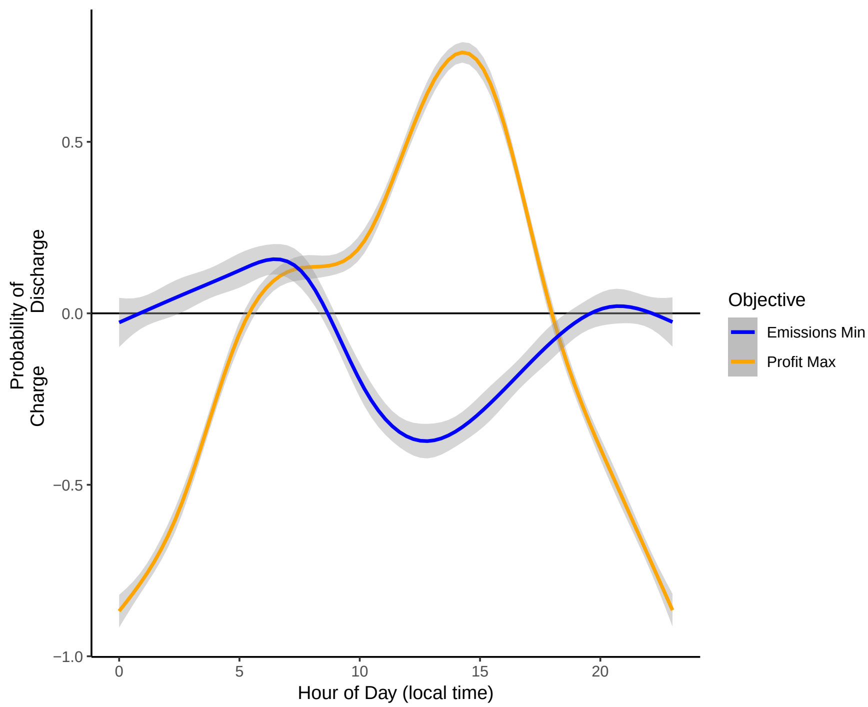

(b) West Zone

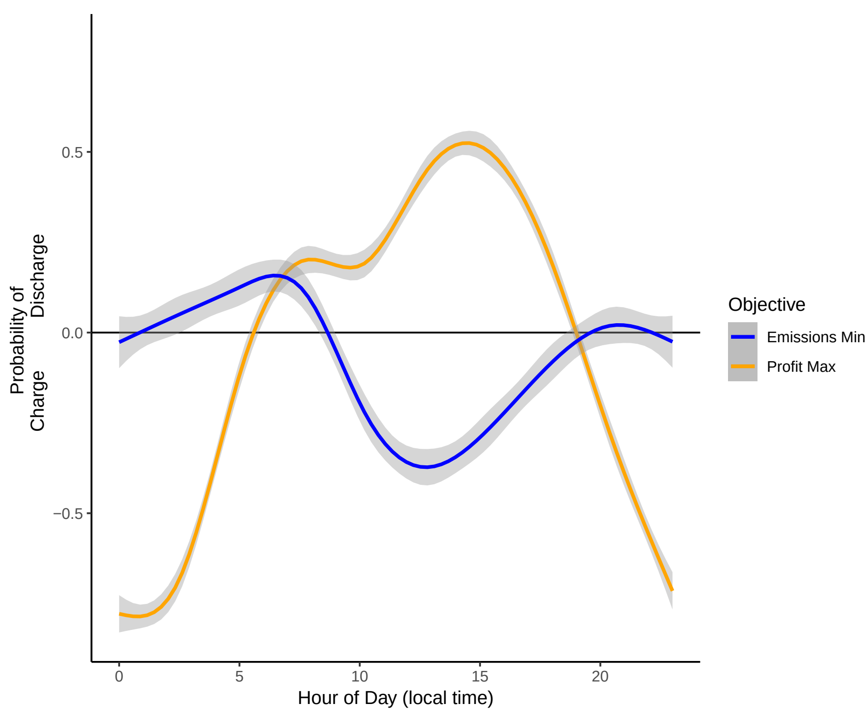

Predicted charge/discharge likelihood of a hypothetical electricity storage provider under profit-maximizing and emissions-minimizing assumptions. Simulations assume storage operators participate in the day-head market and have perfect foresight over prices and emissions. Operator chooses the dynamically optimal path over actual 2017 prices and emissions respecting constraints on the quantity of electricity stored.

Table 1: Typical average CO_2 emissions rate for electricity generation

Fuel	Description	Heat Rate (mmBTU/MWh)	Avg. Emissions Rate (tons CO_2/MWh)	Avg. Emissions Rate (lbs CO_2/MWh)
Coal (anthracite)	Inefficient Coal Plant	11.0	1.26	2.51
Coal (bituminous)	Efficient Coal Plant	9.0	0.93	1.85
Natural Gas	Inefficient CCGT	8.0	0.47	0.94
Natural Gas	Efficient CCGT	7.0	0.41	0.82
Natural Gas	Typical CT	11.5	0.67	1.35
Petroleum	Typical ICE	10.5	0.85	1.69

Typical heat rates and average CO_2 emissions rates for fossil-fueled electricity generation in the United States. Heat rates and emissions factors from EIA Electric Power Monthly.

Table 2: Summaries of hourly CO_2 MEF models

(a) ERCOT Interconnection

Model Name	Coef. Freq.	Demand Shock	Past Ramp	Future Ramp	Fcst Error RMSE	Fcst Error SD	Number of Free Params	LR Test p-value
Fixed Effects Only	N/A	N	-	-	3,116.69	486.92	290	
Static	Single	N	-	-	1,118.80	123.31	292	
Static	hr	N	-	-	1,116.65	122.47	338	
Future Ramp	hr	N	-	1 hr	1,112.18	129.94	386	0.0000
Past Ramp	hr	N	1 hr	-	1,093.50	129.29	386	0.0000
Past+Future Ramp	hr	N	1 hr	1 hr	1,091.06	132.34	434	0.0000
Demand Shock	hr	Y	-	-	1,096.32	125.08	386	0.0000
Shock+Future Ramp	hr	Y	1 hr	-	1,095.62	130.76	434	0.0000
Shock+Past Ramp	hr	Y	-	1 hr	1,091.36	132.85	434	0.0000
Shock+Past+Future Ramp	hr	Y	1 hr	1 hr	1,091.05	133.79	482	0.0000
Demand Shock	hr× qtr	Y	-	-	1,115.49	118.55	674	0.0000
Shock+Future Ramp	hr× qtr	Y	1 hr	-	1,123.02	119.41	866	0.0000
Shock+Past Ramp	hr× qtr	Y	-	1 hr	1,118.36	121.12	866	0.0000
Shock+Past+Future Ramp	hr× qtr	Y	1 hr	1 hr	1,125.80	123.17	1,058	0.0000

Cross-validated model fit parameters from hourly CO_2 MEF models. All models include month-of-year by hour-of-day fixed effects. Future ramp models control at time t for the expected increase in demand at time $t+s$. Past ramp models additionally control for the anticipated changes in demand between time t and time $t-s$. Demand shock models include controls for unanticipated changes in demand between time $t-1$ and time t. RMSE and SD the mean and standard deviation of out-of-sample forecast RMSE across the 10 cross validation iterations. Number of free parameters denotes the number of estimated slope and fixed effect parameters in a given model. The final column is the p-value from a likelihood ratio test of the specified model against the static model with hourly coefficients. Mean CO_2 emissions were 20,021 and 26,454 tons per hour ERCOT and WECC, respectively.

Table 2: Summaries of hourly CO_2 MEF models (Continued)

(b) WECC Interconnection

Model Name	Coef. Freq.	Demand Shock	Past Ramp	Future Ramp	Fcst Error RMSE	Fcst Error SD	Number of Free Params	LR Test p-value
Fixed Effects Only	N/A	N	-	-	2,998.93	748.46	290	
Static	Single	N	-	-	1,328.62	256.98	298	
Static	hr	N	-	-	1,306.60	195.44	434	
Future Ramp	hr	N	-	1 hr	1,275.65	197.64	578	0.0000
Past Ramp	hr	N	1 hr	-	1,291.34	195.48	578	0.0000
Past+Future Ramp	hr	N	1 hr	1 hr	1,254.44	198.84	722	0.0000
Demand Shock	hr	Y	-	-	1,282.56	200.82	578	0.0000
Shock+Future Ramp	hr	Y	1 hr	-	1,264.19	212.91	722	0.0000
Shock+Past Ramp	hr	Y	-	1 hr	1,257.80	198.95	722	0.0000
Shock+Past+Future Ramp	hr	Y	1 hr	1 hr	1,246.58	206.90	866	0.0000
Demand Shock	hr× qtr	Y	-	-	1,203.31	198.47	1,442	0.0000
Shock+Future Ramp	hr× qtr	Y	1 hr	-	1,197.76	202.49	2,018	0.0000
Shock+Past Ramp	hr× qtr	Y	-	1 hr	1,192.26	191.81	2,018	0.0000
Shock+Past+Future Ramp	hr× qtr	Y	1 hr	1 hr	1,204.42	206.58	2,594	0.0000

Cross-validated model fit parameters from hourly CO_2 MEF models. All models include month-of-year by hour-of-day fixed effects. Future ramp models control at time t for the expected increase in demand at time $t+s$. Past ramp models additionally control for the anticipated changes in demand between time t and time $t-s$. Demand shock models include controls for unanticipated changes in demand between time $t-1$ and time t. RMSE and SD the mean and standard deviation of out-of-sample forecast RMSE across the 10 cross validation iterations. Number of free parameters denotes the number of estimated slope and fixed effect parameters in a given model. The final column is the p-value from a likelihood ratio test of the specified model against the static model with hourly coefficients. Mean CO_2 emissions were 20,021 and 26,454 tons per hour ERCOT and WECC, respectively.

Table 3: Summaries of hourly CO_2 MEF LLF models

Model Name	Past Ramp	Future Ramp	Fcst Error RMSE	Fcst Error SD
No Ramp	-	-	1029.59	70.42
1 Hour Ramp	1 hr	1 hr	1011.98	82.36
2 Hour Ramp	2 hr	2 hr	990.44	78.78
3 Hour Ramp	3 hr	3 hr	1003.52	79.96
4 Hour Ramp	4 hr	4 hr	1015.87	79.45

Cross-validated model fit form hourly CO_2 MEF models fit using LLF. All specifications allow for independent year, month-of-year, and hour-of-day effects. Future ramp is the additional level of demand (or non-fossil fuel supply) expected at hour $t + s$ expected in hour t. Past ramp in the change in demand (or non-fossil fuel supply) between hour t and $t - s$. RMSE and SD denote the mean and standard deviation of out-of-sample forecast RMSE across all 10 cross-validation iterations.

Table 4: Marginal CO_2 Emissions Change from Generation Changes - ERCOT

Generation Technology	Marginal Change tons CO_2 per MWh		Percent Difference
	No Ramp	With Ramp	
Onshore Wind	-0.634	-0.632	0.2%
Solar PV	-0.625	-0.626	-0.2%
Storage	0.000	-0.252	-

Marginal change in CO_2 emissions per MWh of electricity generated from adding a marginal unit of the specified type. Percent difference is the amount emissions benefits are overstaged by ignoring ramping. Effects computed accounting for unexpected demand shocks and the anticipated level of change in demand over the previous hour. Marginal emissions factors vary by and hour. Solar output and wind output profile derived from hourly grid-scale solar output in the CAISO region during 2018. Energy storage is assumed to have 80% roundtrip efficiency, can fully charge in 12 hours, and is operated to maximize CO_2 emissions.

Table 5: Marginal CO_2 Emissions Change from Generation Changes - CAL

Generation Technology	Marginal Change tons CO_2 per MWh No Ramp	Marginal Change tons CO_2 per MWh With Ramp	Percent Difference
Onshore Wind	-0.588	-0.564	4.2%
Solar PV	-0.537	-0.461	16.3%
Storage	-0.144	-1.735	-91.7%

Marginal change in CO_2 emissions per MWh of electricity generated from adding a marginal unit of the specified type. Percent difference is the amount emissions benefits are overstaged by ignoring ramping. Effects computed accounting for unexpected demand shocks and the anticipated level of change in demand over the previous hour. Marginal emissions factors vary by region, quarter of year, and hour. Solar output and wind output profile derived from hourly grid-scale solar output in the CAISO region during 2018. Energy storage is assumed to have 80% roundtrip efficiency, can fully charge in 12 hours, and is operated to maximize CO_2 emissions.

Table 6: Marginal CO_2 Emissions Change from Generation Changes - NW

Generation	Marginal Change tons CO_2 per MWh		Percent
Technology	No Ramp	With Ramp	Difference
Onshore Wind	-0.515	-0.565	-9.0%
Solar PV	-0.514	-0.443	16.2%
Storage	-0.085	-1.022	-91.7%

Marginal change in CO_2 emissions per MWh of electricity generated from adding a marginal unit of the specified type. Percent difference is the amount emissions benefits are overstaged by ignoring ramping. Effects computed accounting for unexpected demand shocks and the anticipated level of change in demand over the previous hour. Marginal emissions factors vary by region, quarter of year, and hour. Solar output and wind output profile derived from hourly grid-scale solar output in the CAISO region during 2018. Energy storage is assumed to have 80% roundtrip efficiency, can fully charge in 12 hours, and is operated to maximize CO_2 emissions.

Table 7: Marginal CO_2 Emissions Change from Generation Changes - SW

Generation	Marginal Change tons CO_2 per MWh		Percent
Technology	No Ramp	With Ramp	Difference
Onshore Wind	-0.623	-0.665	-6.3%
Solar PV	-0.711	-0.687	3.6%
Storage	-0.328	-2.958	-88.9%

Marginal change in CO_2 emissions per MWh of electricity generated from adding a marginal unit of the specified type. Percent difference is the amount emissions benefits are overstaged by ignoring ramping. Effects computed accounting for unexpected demand shocks and the anticipated level of change in demand over the previous hour. Marginal emissions factors vary by region, quarter of year, and hour. Solar output and wind output profile derived from hourly grid-scale solar output in the CAISO region during 2018. Energy storage is assumed to have 80% roundtrip efficiency, can fully charge in 12 hours, and is operated to maximize CO_2 emissions.

Table 8: Marginal Emissions Change from the Marginal Change in Storage

(a) ERCOT Interconnection

	Marginal Change in CO_2 Emissions per MWh					
	Profit Max			Emissions Min		
ERCOT Region	No Ramp	With Ramp	Difference	No Ramp	With Ramp	Difference
HOUSTON	0.1730	-0.1605	207.8%	-0.2911	-0.6277	53.6%
SOUTH	0.1793	-0.1674	207.1%	-0.2911	-0.6277	53.6%
WEST	0.1721	-0.1634	205.3%	-0.2911	-0.6277	53.6%

Marginal impact of electricity storage on CO_2 emissions per MWh of electricity stored. MEFs computed using the LLF models from Section 5.2. "No Ramp" estimates do not account for ramping effects. "Ramp" models include two hours of forward and lagged ramp. Profit-maximizing storage chooses the allocation of charge/discharge behavior that maximizes profits in the day-ahead market. Emissions-minimizing chooses the allocation that minimizes total CO_2 emissions.