

Ramping Costs and Coal Generator Exit*

Miguel Borrero[†] Gautam Gowrisankaran[‡] Ashley Langer[§]

August 30, 2024

Abstract

Declines in the cost of generating electricity with natural gas and the rise of low marginal cost renewables have contributed to coal generators moving from operating nearly continuously to only generating during periods of relatively high demand. This has led to substantially more “ramping” where generators change generation levels rather than operating continuously at maximum efficiency. This paper provides new approaches to estimating ramping and operations and maintenance costs for fossil fuel generators. Our approaches incorporate the dynamic linkages between generation decisions over time with simple and tractable estimation strategies. Our estimator recovers ramping costs from generator choices of ramping and deramping, controlling for the future option value by matching hours where generators are currently at different operation levels but where the future expected values for a given level of generation are similar. We find that ramping costs for coal generators are substantial, with costs of \$132,000 for 100-300MW generators for starting up and ramping to minimum generation and further costs of \$75,000 for ramping to maximum generation. We plan to investigate the extent to which costs contributed to the decline in coal profits starting around 2009. Our results suggest that incorporating ramping costs is important to understanding the exit of coal generation over this period.

JEL Codes: Q48, Q52

Keywords: coal generators, operations and maintenance costs, startup costs

*We thank James Archsmith, Louis Preonas for comments and help with the Continuous Emissions Monitoring System data. We also thank Allan Collard-Wexler, Liran Einav, and Harrison Fell for helpful comments.

[†]Columbia University

[‡]Columbia University, CEPR, and NBER

[§]University of Arizona and NBER, alanger@gmail.com

1 Introduction

The twenty-first century has seen a massive transition in how the United States generates electricity. The combination of combined-cycle natural gas technology and decreases in the cost of natural gas stemming from the development of hydraulic fracturing (“fracking”) means that producing electricity with natural gas has become less expensive than with coal, and the U.S. now generates more electricity with natural gas than with coal. Combined with investments in zero marginal cost renewable energy capacity, these changes have moved coal from “base load” generation that runs nearly continuously as the lowest cost generation technology to a higher cost generation technology that only runs when demand is sufficiently high. As an example, in 2008, merchant coal generators in the Eastern Interconnection averaged 31.4 hours at their maximum generation level each time they ramped to this level, but this number dropped to 21.0 hours by 2017 (?).

This increase in ramping has the potential to dramatically change coal generators’ profits. Engineers have long understood that ramping and starting up is costly to generators both because it increases their fuel costs and also because it causes large thermal and pressure stresses to their principal components, damaging them and eventually causing them to fail (?). Economists have used electricity market bidding behavior to demonstrate that coal generators are substantially more expensive to ramp than natural gas generators ?, and that the excess fuel consumed in ramping has the potential to drastically increase pollution and the social cost of coal electricity generation (?).

Because of the multiple sources of ramping and start-up costs and their realization over time, it is difficult to directly measure all of the costs associated with electricity generator ramping. It is therefore also difficult to understand the extent to which ramping for coal generators has contributed to their decreased profits and ultimately exit. This paper develops new methods to estimate ramping costs—for both coal and natural gas generators—that rely on generation choices made by generators, rather than engineering estimates. By using a revealed preference approach, we are able to capture not only the fuel costs of ramping, but instead a combined estimate of all sources of ramping costs, including notably depreciation of equipment and increased maintenance costs (e.g. labor). We develop a similar revealed

preference approach to estimate operations and maintenance (O&M) costs, which we define to be the marginal costs that are proportional to generation and not captured by direct fuel costs.

Ramping costs create dynamic connections between hourly generation decisions: generators must account for future economic conditions in determining when to ramp or deramp. For instance, a generator may choose to start up in a given hour not because it recovers its ramping costs in that particular hour, but because it anticipates higher profits in future hours from having already ramped to a higher production level. We estimate ramping and O&M costs with new, tractable methods that incorporate these dynamic linkages between hours.

Specifically, we develop a conceptual experiment that compares sets of hours with similar future electricity prices and hence similar relative values from generating at any level, but where the generation level varied in the previous hour. We recover ramping costs with a discrete choice regression of the chosen generation level on revenues, lagged generation, and controls for future value. We interpret the coefficient on lagged generation as the ramping costs in utility units and divide it by the coefficient on revenues to obtain a measure of ramping costs in dollars.

Our conceptual experiment identifies ramping costs under two main assumptions. The first is that we can sufficiently control for the difference in the relative value of generation at each generation level with a set of regressors that does not include the current wholesale electricity price or lagged generation. In other words, the controls must be sufficient to eliminate any correlation between the residuals and current electricity price or lagged generation. The second assumption is that generators do not affect wholesale electricity market prices with their individual ramping decisions. Thus, our approach is best suited to markets where there are relatively many generators.

We use a related conceptual experiment to estimate O&M costs. Here, we choose sets of hours around when a ramp or deramp occurs. Under the assumption that generation will change sometime in the window around when it actually did change, the future value at the end of this window and the ramping costs incurred are not relevant to the decision of when to ramp. Thus, we can identify O&M costs by uncovering the hour at which ramping or

deramping actually occurred, conditional on it occurring during this window. We implement this analysis with a multinomial logit model that considers the hour of ramping or deramping within a 6 hour window on either side of the actual ramp or deramp.

Combining our ramping and O&M cost estimates with observed panels of generation choices allows us to recover generator profits and estimate the relationship between profits and underlying dynamic states. This in turn allows us to understand the extent to which increased ramping costs were a large contributor to coal generator exit at the beginning of the twenty-first century. While our ramping and O&M cost estimation leverages the assumption that generators do not affect prices, these profit calculations use observed generator behavior and hence are consistent with more general models of firm behavior.

Our current preliminary estimates use the analysis sample from ?, which consists of merchant generators in the Eastern Interconnection. Using these data, we find that coal startup and ramping costs are substantial, with costs of \$132,000 for generators sized 100-300MW for starting up and ramping to minimum generation and further costs of \$75,000 for ramping to maximum generation. We also find that O&M costs are approximately \$15 per MWh, which is similar to estimates from the EIA’s National Energy Modeling System.

We are in the process of estimating our ramping and O&M cost model on a broader sample of coal and gas generators throughout the U.S., and conducting an extensive set of robustness checks. These estimates will allow us to calculate the change in coal generators’ costs from increased ramping. This will in turn allow us to better understand the relative contribution of ramping costs and decreased generation quantities to the decline in coal generators’ profits.

Our paper extends the literature on estimating ramping and O&M costs. In particular, our dynamic model of generator ramping builds on the work of ? and ? by introducing a different, and more tractable, estimation approach. Our work also relates to ? and ? who develop different, and complementary, approaches to estimating ramping, startup, and O&M costs. Finally, we contribute to a broader literature that estimates structural models of the electricity market, e.g. ?, ?, ?, ?, ?, and ?. Two papers in this literature are by an overlapping set of coauthors. This paper provides inputs that are used in ? in their dynamic oligopoly model of coal plant retirements. ?, which considers the role of regulation in utilities’ investment, retirement, and operations decisions, uses a similar model to control for ramping

costs but different estimation methods.

The remainder of the paper is organized as follows. Section ?? presents the model of generators’ hourly decisions. Section ?? discusses how we estimate our model and explains the intuition behind our identification strategies. Section ?? presents the data we use for estimation and relevant descriptive statistics. Section ?? presents our results, and Section ?? concludes.

2 Model

We model ramping and O&M costs by developing and estimating a dynamic model of hourly operations. Consider a generator j that is faced with an infinite horizon hourly problem with hourly discount factor β . Each generator is endowed with a heat rate (the fuel burned per unit of electricity production) $heat_j$ and a capacity K_j . It faces a fuel price f_j . In each hour h , the generator must choose a generation quantity q_{jh} , conditional on a state. It chooses each q_{jh} to maximize expected future profits. Profits in any hour are the sum of revenues minus three costs: fuel, ramping, and O&M. Generator j ’s fuel costs per MW of production are the product of the fuel price, f_j , and $heat_j$.

In each hour, we model the generator as choosing between generating at its capacity, K_j , generating at a minimum generation level, $L_j K_j$ for $L_j \in (0, 1)$, and not generating.¹ In order to specify the generator’s Bellman equation, we assume that firms that own multiple generators make independent state-contingent profit maximization decisions for each generator.

At any hour h , let $q_{jh} \in \{0, L_j K_j, K_j\}$ denote the current generation level, \tilde{q}_{jh} denote the level in the previous hour, and ω_{jh} denote the generator’s state, which includes information that affects current and future wholesale electricity prices. For instance, ω_{jh} might include load, the time of day, weather, and rival generators’ availability. Further, let $r_{\tilde{q},q}$ be the cost of ramping from \tilde{q} to q for $\tilde{q} < q$,² om be the per MWh O&M costs, and $p(\omega, q)$ be the

¹if generators take hourly electricity prices as given, bear a startup cost, have other ramping costs that are proportional to the increase in generation, and have a minimum generation level of $L_j K_j$, then they will only produce at one of these three generation levels. This formulation of generators’ choice sets also follows ?.

²For simplicity, our base specification does not allow for “deramping” costs, although we provide robustness

expected wholesale electricity price given ω and q . Combining these terms, generator j 's Bellman equation is:

$$V_j(\tilde{q}_j, \omega_j) = \max_q \left\{ \pi_j(\tilde{q}_j, \omega_j, q) + \beta E[V_j(q, \omega'_j | \omega_j)] \right\}, \quad (1)$$

where per-hour profit is:

$$\pi_j(\tilde{q}_j, \omega_j, q) \equiv q \times [p(\omega_j, q) - \text{heat}_j \times f^C - \text{om}] - \mathbb{1}\{\tilde{q}_j < q\} r_{\tilde{q}_j, q} + \sigma \varepsilon_{jq}. \quad (2)$$

Our model specifies that om consists of the costs that are proportional to the production level of the generator other than fuel costs. We assume that the unobservable ε_{jq} is type 1 extreme value and *i.i.d.* across generation levels, hours, and generators. We further assume that a generator observes its own ω_j and ε_{jq} at the start of each hour, but forms expectations about future values of ω_j based on the current state. Because operating profits are measured in dollars, we include a parameter, σ , that scales ε_{jq} .

3 Estimation and Identification

While it is theoretically possible to estimate the ramping and operations and maintenance costs in equations (??) and (??) with a full-solution dynamic structural model (?), the high frequency of the decisions and the complicated state space embodied in ω make this extremely difficult in practice. Instead, we develop “conceptual experiments” in order to identify ramping and O&M costs in a more tractable way.

Before turning to our estimation approach, we provide an overview of the data we have available, which we discuss in more detail in Section ???. Our model depends fundamentally on the revenues that generators earn in hourly electricity markets and the costs that they incur. We calculate revenues directly from our data by multiplying wholesale electricity price, p_h , with quantity supplied, q_h . Our data also include generator heat inputs, production, and fuel prices, which allow us to calculate fuel costs.

results that allow for this possibility.

3.1 Ramping Costs

In principle, ramping costs are identified by the extra revenue that generators expect to earn when they increase their generation level. The issue is that ramping is fundamentally dynamic: generators may increase generation in an hour in order to capture the option value of remaining at a high generation level in future hours. Our estimation approach involves finding instances where the generator finds itself in different sets of situations with identical information about future prices, ω , but where generation in the previous hour varies across the sets. The difference in the probability of generation levels across these sets identifies the ramping costs.

We implement our approach with a trinomial choice model of hourly generation levels that follows from equations (??) and (??):

$$u(q_h|\tilde{q}_h, \omega_h) = \underbrace{\frac{1}{\sigma}[q_h p_h]}_{\text{Revenues}} - \underbrace{\frac{1}{\sigma} r_{\tilde{q}_h, q_h} \mathbb{1}\{\tilde{q}_h < q_h\}}_{\text{Ramping costs}} - \underbrace{\psi x(q_h, \omega_h)}_{\text{Other costs and relative continuation value}} + \varepsilon_{qh}. \quad (3)$$

There are four key parameters for each generator: the standard deviation of the unobservable, σ , which scales revenues, the costs of ramping from off to minimum generation ($r_{0,LK}$), off to maximum generation ($r_{0,K}$), and minimum to maximum generation ($r_{LK,K}$). In this regression, the x_h variables serve as controls that capture anything that might be included in non-ramping costs and the continuation value relative to not generating, which ties this regression to our “conceptual experiments.” Regression controls include a flexible functional form of generation quantity, year relative to first year and its square, fuel price ratio and its square, relative generator capacity, fuel cost per MW, weather, hour-of-day, month, and current load. We explicitly do not include \tilde{q}_h in these controls. We classify each hour into off, minimum, and maximum generation bins as we describe in Section ?? below.

Comparing equation (??) to equations (??) and (??), we have made three changes. First, we collapsed the future value term in (??) and the fuel cost and operations and maintenance cost terms from (??) into $x(q_h, \omega_h)$. Second, we divided each coefficient by the standard deviation of the unobservable. This change allows us to estimate generation decisions in a multinomial choice framework, where we are now estimating a ramping cost parameter of

$r_{\tilde{q}_h, q_h} / \sigma$. We therefore recover an estimate of ramping costs in dollars, $r_{\tilde{q}_h, q_h}$, by dividing our estimate of the ramping cost parameter by the estimated coefficient on revenues. Third, we assume that generators do not affect wholesale electricity prices with their individual generation choices. This implies that $p(\omega_j, q)$ is a function only of ω_j and not q . We then write p_h as a shorthand for $p(\omega_h)$.

Identification of ramping costs hinges on the ability of x_h to appropriately capture the generator’s relative state-contingent continuation values. This relies on three exclusion restrictions: that current revenues, $q_h p_h$, and the lagged generation level, \tilde{q}_h , do not enter into x_h , and that (as noted above) units take wholesale electricity prices on the hourly market as given. Thus, x_h must capture the generator’s expectations of relative future values sufficiently well that current revenues and lagged generation do not provide additional information on these expectations. To ensure that x_h is sufficiently rich, we estimate it with a flexible functional form that captures expected future generator costs and revenues.³

Fundamentally, our approach is based on comparing generation across hours with identical prices and continuation values but different lagged generation levels. An alternative approach would be a conditional choice probability (CCP) estimator (e.g. ??), which uses a log transformation of the differences in probabilities across generation level by state. Given the infrequency with which generators change generation levels, many of these probabilities would be very close to 0, making it difficult to estimate the future values accurately. In contrast, our approach only requires that we have sufficient information to appropriately control for generators’ continuation values and sufficient observations to have “matches” where continuation values are very similar, but lagged generation levels differ.

3.2 Operations and Maintenance Costs

O&M costs, $om \times q$, are proportional to the quantity generated. In this way, they act much like an option-specific constant term in equation (??), entering via $\psi x(q_h, \omega_h)$. However, this

³We have estimated alternative specifications where we interact x_h with functions of prices over the following 20 hours. These specifications should very accurately control for continuation values, but since they include future prices, they incorporate more information than generators would actually have available. We obtain broadly similar ramping cost estimates with either specification, so use the ones based on available information.

term also includes fuel costs and the relative continuation value, so we cannot recover O&M costs directly from this regression. We can control for fuel costs in a straightforward way since we observe heat rates and fuel prices. The key challenge is in separating O&M costs from continuation values.

We recover O&M costs by focusing on the production decision within a set of hours where the continuation values are very similar. In particular, we estimate O&M costs with regressions on a subsample of the data consisting of windows of hours around an actual ramp or deramp. We assume that the generator knows that it will be ramping or deramping exactly once during this window. This implies that, within a subsample, the generator’s continuation value will be the same regardless of which hour it chooses to change generation within the window. The fact that the continuation values are the same regardless of the ramp hour then allows us to separate O&M costs from the continuation values.

We operationalize this idea by using a window around ramping events in our hourly data that includes the 6 hours before, the observed hour, and the 5 hours after a generation change. We focus on the subsample of windows around when generators ramp from minimum to maximum generation or vice versa.⁴ We further assume that the generator has perfect information about the wholesale electricity prices for every hour in this window.

Given these assumptions, the generator chooses the hour in which to change generation that maximizes the sum of profits over this window. We explain our approach mathematically for windows where generator j is ramping up to maximum generation; windows where the generator is ramping down to minimum generation are similar, but in reverse. We estimate a multinomial logit model with 12 choices, one for each hour in the window. For a window that begins in hour \underline{h} , we can write the relative profits in the window, w , from ramping at hour $h \in \{\underline{h}, \dots, \underline{h} + 11\}$ as:

$$\pi_{jh}^w = \sum_{\tilde{h}=\underline{h}}^{\underline{h}+11} (p_{\tilde{h}} - heat_j \times f^C)(K_j - K_j L_j) - om \times (\underline{h} + 12 - h) \times (K_j - K_j L_j) + \sigma^w \varepsilon_{jh}^w, \quad (4)$$

⁴These changes are better for identifying O&M costs than hours when generators turn off or on because those decisions are more likely to be affected by unobservable factors such as required maintenance, and so may not fully be responding to immediate market price incentives.

where the generator earns revenues on the additional sales from generating at maximum rather than minimum generation starting at hour h , but also needs to pay fuel and O&M costs. The unobservable ε_{jh}^w will capture the conditional distribution of the underlying structural residuals ε that stem from the decision to ramp at hour h , and have a standard deviation, σ^w .⁵ As an approximation, we use a type 1 extreme value distribution for ε_{jh}^w . Because the generator bears the same ramping costs and continuation values regardless of its choice of hour in which to ramp, these terms do not enter its maximization in (??).

Our model identifies om from the increase in hourly profits in the hour the generator changes generation. In our ramping window subsample, we always observe the generator changing generation in the seventh hour of the window. Thus the O&M costs will essentially let the generator break even in this seventh hour of the ramping window.

Figure ?? uses our O&M cost estimation subsample to provide intuition behind the identification of our estimates. The figure shows the hourly profit per MW before O&M costs (so revenues minus fuel costs) in the hours around each ramp from minimum to maximum generation or back.⁶ We indicate the fact that the generator always chooses to ramp in between hour 6 and hour 7 with a grey vertical dashed line. The O&M costs that rationalize this choice are higher than the profit bar in hour 6, but lower than the profit bar in hour 7. The mean of these estimates is approximately \$17, as indicated with the horizontal grey dotted line.

4 Data and Descriptive Statistics

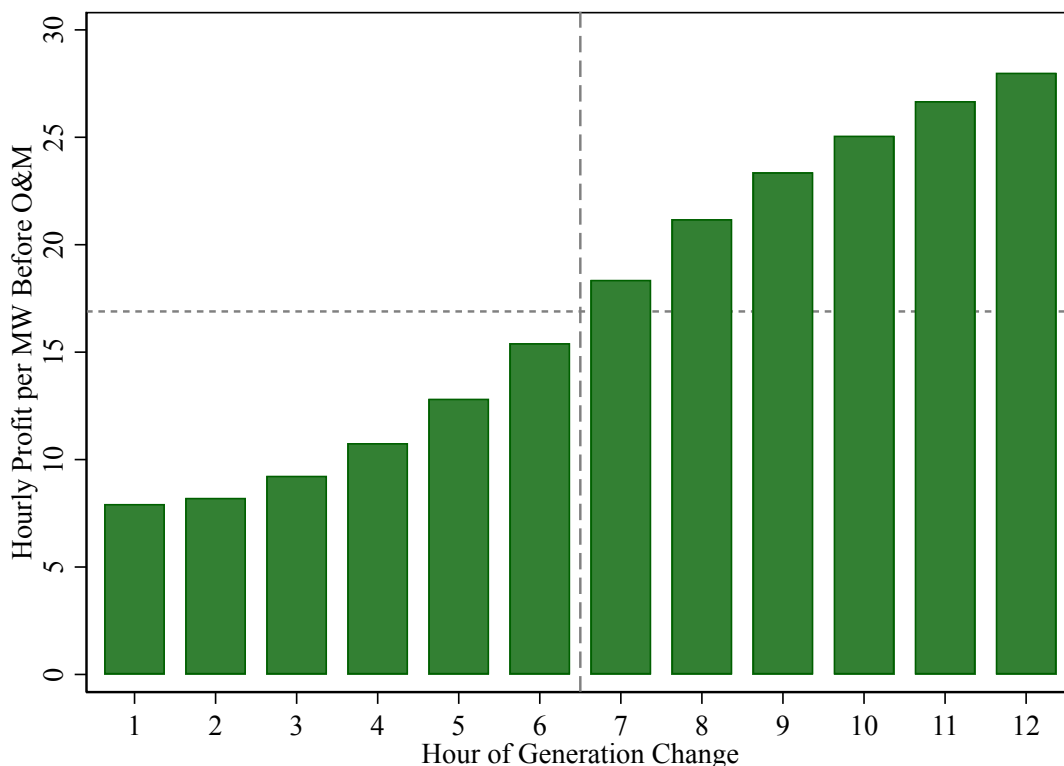
4.1 Data Sources

Our analysis data are at the generator-hour level. We use generators (what the EPA calls units and the EIA calls boilers) rather than power plants—which can be a collection of generators—because these are the level at which ramping and startup costs are generally

⁵The deramping case sums backwards over the relative profits from hour $\underline{h} + 11$ to hour h .

⁶For hours when the generator ramps from minimum to maximum generation, the hours are listed in order, while for hours when the generator ramps from maximum to minimum generation, the hours are reversed so that hour 12 occurs first and hour 1 occurs last.

Figure 1: Increase in Generating Profit Net of Operations and Maintenance Costs



Note: Created from the sample of 12-hour windows around generator ramping events into or out of maximum generation. Hourly profits per MW before operations and maintenance are equal to electricity price minus heat rate times generator price. The horizontal axis represents the number of hours at the *lower* generation level. The vertical line indicates that generators always choose to ramp up in the 7th hour and down in the 6th hour. The horizontal line indicates the mean potential O&M cost that would rationalize this behavior.

borne. Our data sets include information on generators’ production, fuel efficiency, emissions, fuel costs, demand, and electricity prices.

The analysis sample in this paper are currently closely related to the analysis sample in ?. Specifically, we use merchant coal generators in the Eastern Interconnection. This means that we use independent power producers, rather than generators that are owned by regulated utilities. We plan to extend our sample to include other coal generators and natural gas combined cycle generators.

The EPA’s Continuous Emissions Monitoring System (CEMS) database provides the backbone of our analysis data. Observations in the CEMS data include information on the

heat input of the fuel used (in MMBtu) and electricity production (in MWh). We currently use coal generators for our estimation though we plan to estimate raping costs for combined cycle natural gas generators. We determine the primary fuel for each generator using EPA data on generator characteristics. We define a generator as using coal if the reported fuel type includes the word “COAL.” We define a generator as a combined cycle natural gas generator if the reported fuel type is “NG” and the reported prime mover includes the term “Combined Cycle.”

Calculating revenues from hourly wholesale electricity markets requires wholesale electricity prices. We use hourly electricity prices by U.S. state from nodes in each Regional Transmission Organization (RTO) or Independent System Operator (ISO) in the Eastern Interconnection.⁷ We average across nodes for states with multiple nodes and use prices from the nearest node for those states without a reported electricity price (e.g. Georgia). We deflate these prices to January, 2006 dollars using the Bureau of Labor Statistics’ chain-weighted consumer price index for urban consumers.

To calculate the fuel costs of generating electricity, we combine heat inputs from CEMS with fuel prices. We recover annual, U.S. state-level natural gas and coal prices from EIA Form 423. While Form 423 reports fuel prices at the generator-year level, we instead use the mean state-level fuel price weighted by each generator’s annual production in order to better capture the opportunity cost of burning fuel.

We use load at the state level to control for relative continuation values. To understand local load we use U.S. state-level electricity load from the Public Utility Data Liberation (PUDL) database, which relies upon the Federal Energy Regulatory Commission (FERC) Form 714. We use PUDL’s version of load that scales hourly load to match the total annual load at the state level from EIA Form 861.

Finally, we use county-level weather data from PRISM as further controls for relative continuation values. Following ?, we recover daily heating degrees, which measure the amount by which population-weighted state average daily temperature exceeds 65 degrees (if at all), and analogously cooling degrees.

⁷Specifically, we retrieve electricity prices from the New England ISO, New York ISO, PJM, Midcontinent ISO, and the Southern Power Pool.

We use these data to construct minimum and maximum generating capacity and heat rate, following the approach in ?. Specifically, we define a generator’s maximum generation—and its capacity—as the 95th percentile of observed generation since generators can produce above their maximum generation for short periods of time at potentially high cost. We define the minimum generation as the generator’s modal generation between the 5th and 60th percentile of capacity. We define a generator’s efficiency, or “heat rate” as its mean MMBtus of fuel input to generate one MW of electricity when operating at maximum generation. This means that any excess fuel consumed at lower generation levels will be considered to be part of ramping costs.

4.2 Evidence on Importance of Ramping Costs Over Time

Table 1: Change in Ramping and Natural Gas Prices Over Time

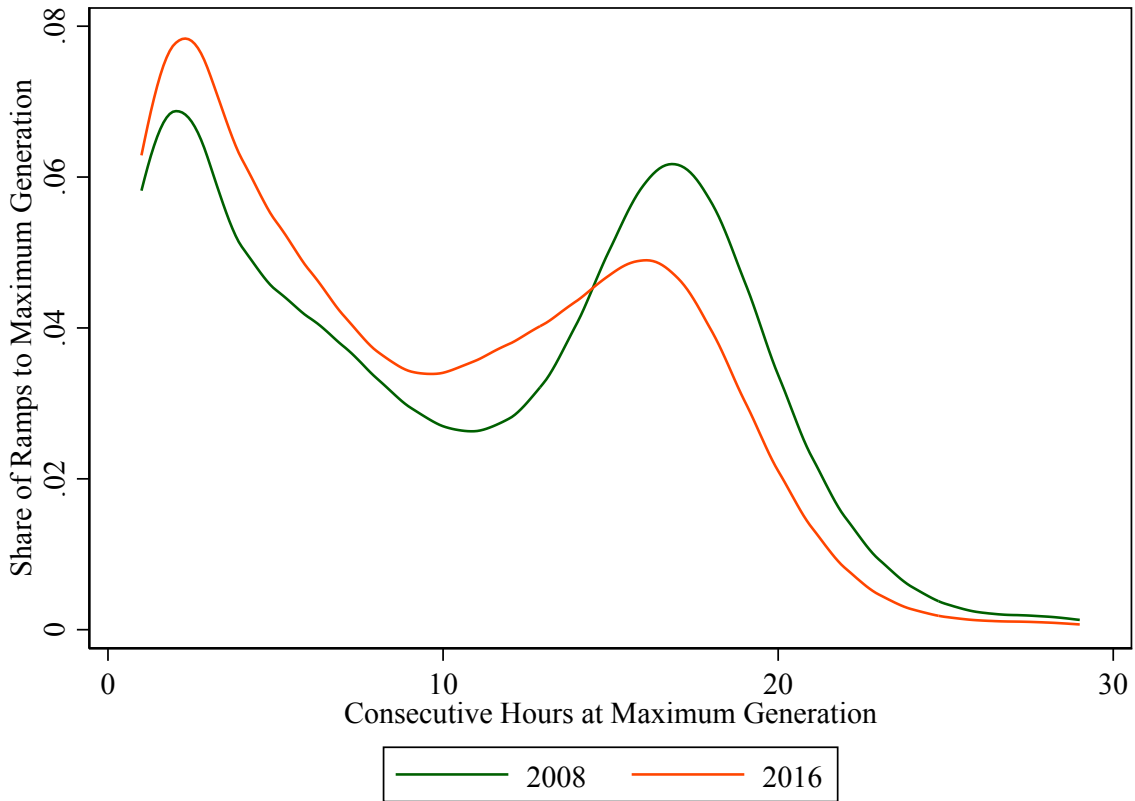
Year	Hours at Max Generation Per Ramp	Natural Gas Price Over Coal Price
2006	27.99	4.62
2007	32.96	4.39
2008	30.42	4.82
2009	23.08	2.35
2010	26.70	2.27
2011	24.68	1.94
2012	20.73	1.32
2013	22.76	1.82
2014	24.41	2.28
2015	17.57	1.41
2016	17.17	1.35
2017	19.94	1.63

Note: Authors’ calculations based on analysis samples of IPP coal generators. Each observation in the second column pertains to one observed ramp to maximum generation. Each observation in the third column pertains to one observed generator in one year.

Table ?? shows changes over time in the hours that coal generators spend at maximum generation each time they ramp to maximum generation alongside the ratio of the natural gas price to the coal price. Starting in 2009, both natural gas prices and hours at maximum

generation declined sharply. This change was largely caused by the rise of fracking, which led to a substantial drop in the price of natural gas. In many cases, fracking caused coal generators to be replaced by combined-cycle natural gas sources as the lowest-cost fossil-fuel electricity source. Coal therefore went from generating regardless of electricity prices to only generating when electricity prices were particularly high.

Figure 2: Distribution of Consecutive Hours at Maximum Generation



Note: Authors' calculations based on analysis sample of IPP coal generators. Each observation is one observed ramp to maximum generation. The green line displays a kernel density of the number of hours at maximum capacity for each ramp to maximum capacity in 2008. The orange line displays the same information for 2016. Both densities are truncated at 30 hours per ramp.

Figure ?? provides further detail on the change in hours at maximum generation. Specifically, it shows the density of hours at maximum generation per ramp for 2008 and 2016. The 2016 distribution has substantially more weight at lower levels of hours per ramp. These patterns demonstrate that accurately measuring ramping costs is necessary to understand

how the operating profits of coal generators have changed over time.

5 Results

Our results are currently based on the set of generators and years used in ?, which focuses on coal generators in the Eastern Interconnection. We are in the process of constructing a more complete data set of coal and combined-cycle natural gas generators across the U.S. Once we have estimated ramping and O&M costs on those data, we will calculate how ramping cost changed between 2000 and 2023 and how this contributed to the overall decline in coal generator profits during this period.

We estimate ramping and O&M costs using our hourly generation model combined with equations (??) and (??). We estimate ramping costs completely independently for generators of different sizes. Table ?? provides the ramping cost estimates for generators between 100 and 300MW of capacity, which includes the modal capacity generator in our sample. The first column presents results without any controls for continuation value, while the second column includes our full set of controls. Controlling for continuation value appears to be extremely important for estimating ramping accurately. With controls, operating revenue is less important by an order of magnitude, and this increases the estimated ramping costs similarly. With controls, we estimate that ramping from off to minimum generation costs these generators \$132,000 while ramping from minimum to maximum generation costs \$75,000. Further, we estimate that ramping from 0 straight to maximum generation costs these generators \$233,000. Thus, ramping directly from off to maximum generation costs \$26,000 more than dividing the ramp across two hours. This is consistent with it being costly for generators to ramp quickly to full generation.

Table ??, Column 3 allows for the possibility that deramping, or reducing generation, is also costly. The results in Column 3 show ramping costs that are substantially lower than Column 2 and deramping costs that are similar in magnitude to ramping costs. However, a comparison of Columns 2 and 3 shows that the total cost of ramping from any level to another level and then deramping back to the original level has approximately the same costs across models. Since the number of ramps will be approximately equal to the number of deramps

Table 2: Estimates of Ramping Costs: 100–300MW Capacity

	No Controls	With Controls	With Controls
Operating Rev. (Mill. \$)	763.87*** (0.46)	76.47*** (0.84)	76.58*** (0.84)
Ramp 0 to Min	−9.03*** (0.01)	−10.13*** (0.01)	−5.21*** (0.18)
Ramp Min to Max	−6.06*** (0.00)	−5.77*** (0.00)	−2.83*** (0.19)
Ramp 0 to Max	−37.82*** (0.16)	−17.82*** (0.06)	−9.97*** (0.21)
Deramp Min to 0	−	−	−4.87*** (0.18)
Deramp Max to Min	−	−	−2.94*** (0.19)
Deramp Max to 0	−	−	−8.74*** (0.21)
N	28,686,219	28,686,219	28,686,219
Pseudo R^2	0.7612	0.8570	0.8571

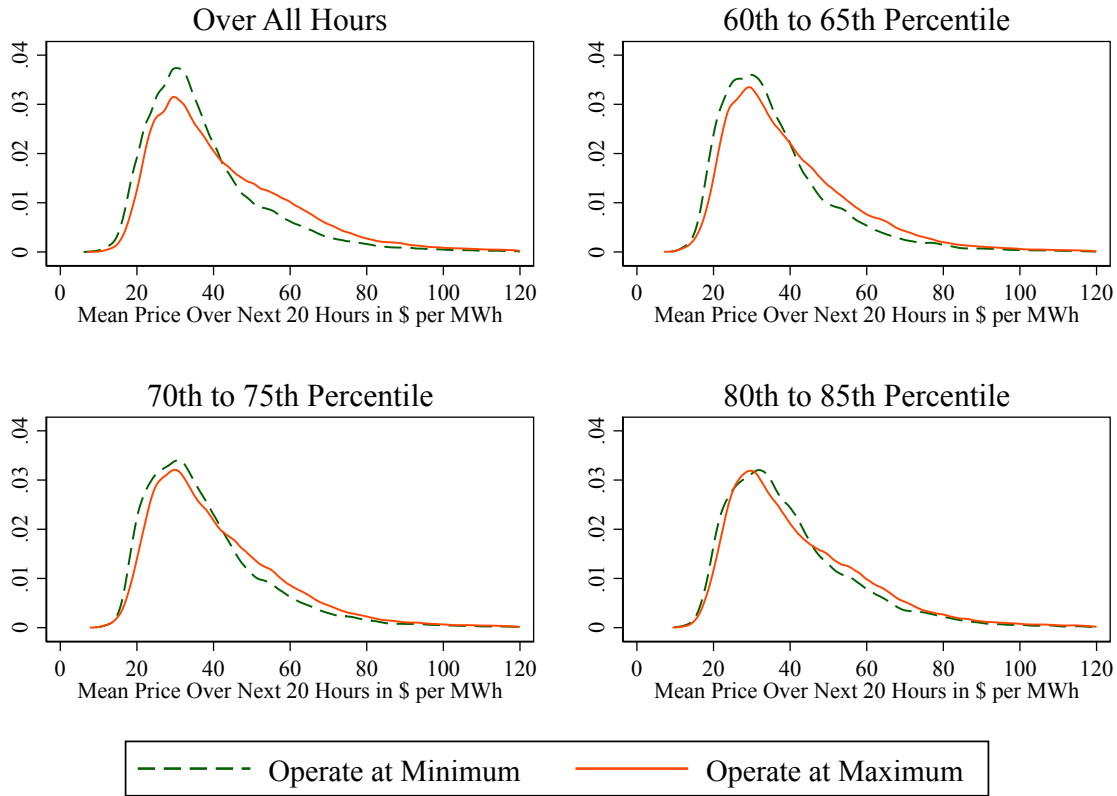
Note: Regression controls include a flexible functional form of generation quantity, year relative to first year and its square, fuel price ratio and its square, relative coal capacity, abatement technology, fuel cost per MW, weather, hour-of-day, month, and current load. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

over a year, this means that estimates of annual equilibrium generator ramping costs will be very similar with either of these specifications of ramping costs.

Overall, we find that starting up from off is particularly costly for generators, while ramping from minimum to maximum generation is substantially less costly. Ramping directly from off to maximum generation in a single hour is more costly than ramping from off to minimum in one hour and then minimum to maximum in a later hour. We find that this pattern is very consistent across generators of different sizes, with costs increasing with generator capacity. The largest generators in our sample (those with capacity over 700MW) need to spend approximately \$1.05 million to ramp from off to minimum generation. This is somewhat higher than engineering estimates that place start-up costs for large coal plants as high as \$500,000 (?). This difference likely comes from the fact that engineering estimates derive from models of excess fuel and maintenance costs rather than the actual optimizing decisions of generators, and therefore may omit some dimensions of wear and tear or labor costs. In the economics literature, estimates using revealed preferences similarly find very large start-up costs (?). Using generators' bidding data for electricity auctions in Spain, ? finds start-up costs of €15-20k for a similarly-sized (150MW) coal generator and €30k for larger, 350MW coal generator.

Our identification argument for ramping costs relies on our ability to accurately capture expected relative continuation values with our controls. For our ramping cost estimates to be consistent, future prices should not predict current actions, conditional on our controls. In order to provide evidence on the suitability of our controls, we compare future mean electricity prices across generators' actual generation decisions, conditioning on the difference in our predicted net continuation values between maximum and minimum generation.

Figure 3: Comparison of Future Electricity Prices By Generation Level



Note: Each panel plots price densities, separately by hours when the generator chooses minimum or maximum generation for generators with 100–300 MW of capacity. The top left panel shows this for all hours. The other panels show hours where the difference in other costs and relative continuation value— ψ/x —from operating at maximum minus minimum generation is within particular quantiles.

Specifically, Figure ?? conditions on three ventiles of the relative continuation value distribution and examines how the densities of the future price distribution over the subsequent

20 hours differ across chosen minimum or maximum generation.⁸ The green dashed line in each picture shows the future mean distribution of electricity prices for generators who chose to produce at their minimum, and the orange solid line shows the distribution for generators who chose to produce at their maximum. The top-left panel shows the two distributions unconditional on relative continuation value. As we would expect, when future prices are low, generators are choosing minimum generation more often. However, once we condition on relative continuation value, the lines are quite similar.

Table 3: Estimates From O&M Regression

Benefit per Million Dollars of Variable Profit	10.55*** (0.45)
Cost per TW of Additional Production	-160.18*** (9.07)
Observations	384,672
Pseudo R^2	0.0043

Note: Multinomial logit regression of choice of number of hours produced within a 12 hour window surrounding each ramping event into or out of maximum generation. Standard errors are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

We estimate O&M costs from equation (??). Table ?? presents our estimates. By dividing the cost per additional TW of production by the benefit per million dollars of variable profit, we find that generators on average pay O&M costs of \$15.18 per MWh of generation, or just over \$1,500 for a 100MW generator at maximum generation.

Our structural estimates here coincide with Figure ??, which provides the reduced form identification behind our results. Specifically, Figure ?? showed that generators did not ramp before the sixth hour, where average revenues net of fuel costs were just under \$15/MWh, but did ramp before the 7th hour when average revenues net of fuel costs were over \$18/MWh. Further, our results are similar to the EIA’s National Energy Modeling System estimates, which find that O&M costs are approximately \$14/MW. ? uses the EIA’s estimates to center their own structural estimates.

⁸We chose these ventiles of the future price distribution because we would expect generators to be making a choice between minimum and maximum generation when the relative continuation value is fairly high.

6 Conclusions and Further Discussion

In this paper we estimate ramping and operations and maintenance (O&M) costs using new approaches and identification arguments. Properly accounting for ramping costs is particularly important during our sample period since large declines in natural gas prices meant that coal generators switched to producing for shorter periods at maximum generation, thereby significantly increasing ramping costs. We find that ramping costs for coal generators in the U.S. Eastern Interconnection are substantial, with start-up costs reaching as high as \$1.05 million for the largest generators.

We are in the process of quantifying the impact of these ramping costs on coal generators' profits. We anticipate that this analysis will allow us to better understand the substantial rate of coal generator retirement over the twenty-first century.