DISCUSSION PAPERS IN ECONOMICS

Working Paper No. 19-04

Rate of Return Regulation, Competition, and Fossil Fuel Retirements

Cameron Duff
University of Colorado Boulder

October 2019
Revised October 22, 2019

Department of Economics

University of Colorado Boulder
Boulder, Colorado 80309

Rate of Return Regulation, Competition, and Fossil Fuel Retirements

Cameron Duff

Updated October 22, 2019

Current Version

Abstract

Electricity deregulation aimed to lower costs by introducing competition to the generation sector, but more than half of US generation capacity is still regulated. In this paper, I investigate how rate-of-return regulation affects the incentives to shut down uncompetitive fossil fuel generators, a potentially important driver of grid efficiency. I find that deregulated generators are more likely to shut down when prices are disadvantageous, while regulated generators are mostly unresponsive. Regulated generators tend to remain operable at much lower capacity factors than their deregulated counterparts, which may be due to utilities' unwillingness to shut down assets that are not fully depreciated. A counterfactual analysis suggests that an additional 700-2,500 MW of natural gas capacity would have shut down each year if all regulated generators were to be deregulated.

1 Introduction

Expensive and inefficient excess generating capacity motivated a wave of state-level electricity deregulation laws in the late 1990s (Borenstein & Bushnell 2015). Proponents of deregulation argued that it would reduce long-run costs by promoting competition in the generation sector (Joskow 2008). However, momentum towards deregulation faltered in the early 2000s, resulting in a mix of regulated and deregulated generation in almost every state. Today, more than half of the power in the US is still produced by utility-owned generators, which may be insulated from competition. In this paper, I examine how traditional utility regulation distorts the utility's incentives to shut down uncompetitive fossil fuel plants, a potentially important driver of long-run efficiency.

The short-run impacts of deregulation have been well-studied. Deregulation has been associated with: increased fuel efficiency for divested plants due to changed incentives (Bushnell & Wolfram 2005); a reduction in plant-level labor and non-fuel costs (Fabrizio et al. 2007); higher efficiency and utilization of nuclear plants (Zhang 2007); shorter downtime for nuclear plant outages (Davis & Wolfram 2012); lower cost of fuel procurement and less capital-intensive emissions abatement strategies (Cicala 2015); lower generation costs due to re-allocation of power output (Cicala 2017).[1] Less is known about long-run effects of deregulation, especially when it comes to decisions affecting the stock of generating assets.[2] In this

[1] On the other hand, Borenstein & Bushnell (2015) make the case that deregulation has been a disappointment due to lackluster retail price reductions for end-use customers, arguing that changes to retail prices are driven more by input prices and improving technology, and less by deregulation.

[2] Csereklyei & Stern (2018) show that technology choices in initial investment are more sensitive to expected input prices for states with access to wholesale markets.

paper, I consider the role of utility regulation in shaping generator-level retirement decisions, relative to their deregulated counterparts.

The historical pace of technology growth in electricity generation has been slow[3], and regulated utilities could confidently invest in plants with the expectation that they would remain useful for the entirety of their "book life."[4] Early retirements (relative to a generator's book life) for purely economic reasons are uncommon. However, in 2019 the electricity industry is in a state of transition, as concerns about carbon emissions continue to grow, and emerging technologies threaten to displace existing generation capacity. Decarbonization of electricity, by replacing fossil fuels with zero-emission sources, is widely considered among the first and most important steps towards decarbonization (Williams et al. 2012). The stock of operable generators is a practical constraint on the carbon intensity of the electric grid. Thus, it is important to consider how the traditional utility structure could be a roadblock to energy transition if regulated generators are insensitive to competition.

In this paper, I present a conceptual model of retirement for fossil fuel generators in the US, based on the concept of economic dispatch. Generators are dispatched in rising order of marginal cost to meet electricity demand, which fluctuates according to daily and seasonal cycles. As a given generator's marginal cost increases because of age or input prices, it is less likely to be dispatched. I describe generators as being "competitive" if they are dispatched

[3]The biggest recent advances in fossil fuel generation have been super-critical coal and combined-cycle gas generators, which are marginally more efficient than their predecessors, but also more capital-intensive.

[4]As emphasized in Rode et al. (2017), book life is defined as the time period over which fixed costs are assumed to be recovered for accounting purposes. This is in contrast to physical life, which is influenced by deterioration, or economic life, which is influenced by market forces. For the purposes of rate-of-return regulation, the book life and depreciation rate of a regulated generator are determined in rate cases, which must be approved by regulators.

frequently enough to justify continued expenditure on fixed operating costs. That is to say, the decision to shut down a generator is largely a function of its capacity factor.[5] Generators which are no longer competitive represent a cost burden on firms (deregulated) or customers (regulated) until they are shut down. As shown in Figure 4, regulated natural gas generators rarely shut down before ~ 40 years, after which there is a large jump in the frequency of shutdowns. The pattern is not present for deregulated generators. Taken together, this implies a potentially significant difference in the way regulated and deregulated firms consider consider the shutdown decision.

To identify this difference, I use plausibly exogenous variation in the price spread (P^{Δ}) between coal and natural gas to estimate a natural gas-fired generator's probability of shutdown in a given year. While age, size, and technology are useful predictors of a generator's marginal cost, its frequency of dispatch is also dependent on how it compares to the marginal cost of other nearby generators, which may use a different fuel, or none at all.[6] I jointly estimate how regulated and deregulated[7] generators respond separately to changes in price conditions. My analysis is focused on natural gas generators, though the results may generalize to coal generators as well.

I find that deregulated gas generators are more likely to shut down when coal is cheap and gas is expensive, and less likely when the opposite is true. For deregulated generators in a given year, a \$1 decrease in P^{Δ} is associated with an increase in the probability of shutdown by

[5]Capacity factor is the ratio of production to potential production over a given period of time.

[6]The coal-gas price spread is highly predictive of the quantity produced by generators in a given year; see Table 6.

[7]I use "deregulated" in this paper to mean "not regulated." It may be the case that a deregulated generator was once regulated, but it need not be.

0.375 percentage points, from a baseline of around 1.2 percentage points. This is consistent with the original motivation for competition as a driver of efficiency in the generation sector, as firms are incentivized to stop using uncompetitive generators. On the other hand, I find that regulated generators are generally unresponsive to price conditions, which is consistent with the theoretical incentives faced by regulated utilities under rate-of-return regulation. Utilities would not want to retire a generator before it is fully depreciated because it may be disallowed from the utility's rate base, which is the foundation of the utility's profits. My results indicate that this difference is likely driven by regulated utilities' willingness to continue operating uncompetitive generators, rather than by regulated generators supplying power out of the merit order. A back of the envelope counterfactual implies that an additional 700-2,500 MW (0.3-1.44 percent) of regulated gas capacity would have shut down each year from 2006-2017, if all regulated generators became deregulated. This result suggests that regulators are generally unsuccessful in exerting pressure for early shutdowns, to the extent they are appropriate. Regulated gas generators tend to shut down only once they have crossed a certain age threshold, typically 35-40 years (see Figure 4).

This paper contributes to the literature measuring the effects of electricity deregulation in the US. I am the first to empirically analyze how the shutdown decision varies according to market structure, which has implications for both consumer welfare and the environment. My analysis highlights another mechanism by which deregulation increases efficiency, as owners of deregulated generators in competitive markets are more motivated to adapt to changing conditions via shutdowns. I also contribute to the literature on efficient pathways to decarbonization. My results emphasize the importance of industry-specific policies and the limitations of catch-all policies such as carbon pricing. If utilities are insensitive to fuel costs, then a carbon price would not have any direct effect on their incentive to transition away from fossil fuels. Instead, policies such as renewable portfolio standards are more likely to succeed because they introduce binding constraints on utility generation portfolios.

2 Background

Section 2.1 describes the unique features of regulated utilities. Section 2.2 explains how the incentives to shut down a generator vary for regulated and deregulated firms. Section 2.3 connects generator shutdowns, especially natural gas, to pathways for decarbonization.

2.1 Regulation and deregulation

Traditional regulated utilities are granted monopoly franchise over service territories. State-level public utility commissions (PUCs) are responsible for regulating the price they are allowed to charge their captive customer base.[8] Prices are set according to the "revenue requirement" formula, such that:

$$E_t[R_{t+1}] = E_t[C_{t+1}] + r_t B_t \tag{1}$$

where R is total revenue, C is operating costs, r is the allowed rate of return, and B is the rate base. Both r and B are determined in periodic "rate cases" where utilities and

[8] This is the sense in which plants and generators are either regulated or deregulated in this paper

regulators negotiate over new capital outlays. Regulators are responsible for keeping prices reasonably low while ensuring grid reliability, so they must balance the wishes of the utility company against those of customers. The key feature of the revenue requirement formula for this paper is that the utility's input costs are passed on to customers, and do not directly affect the utility's profit.[9] As a result, utilities face little direct incentive to minimize costs.

In the late 1990s, many US states began to pass laws aimed at restructuring their electricity industries, often referred to as deregulation. One important component of deregulation was the push for competition in electricity generation.[10] Deregulated states required owners of transmission capacity to provide open access to this crucial infrastructure required to sell power. In many cases, control over transmission lines was given to Independent System Operators (ISOs), which organize the dispatch of power across large regions of the US and Canada.

Open access to transmission facilitated the entry of non-utility firms called Independent Power Producers, who would compete with other suppliers on the basis of cost. These deregulated entrants would have a greater incentive to innovate than their regulated counterparts because they could capture the benefits from lower production costs. In 2017, almost half of the operating generator capacity in the US was deregulated.

[9] Lags between rate cases provide a small incentive for cost minimization, but cost expectations can always be revised in the next rate case, so the overall impact is limited (Fabrizio et al. 2007).

[10] The other segments of the electricity market, transmission and distribution, remained regulated, as they still have characteristics of natural monopoly.

2.2 Shutdown incentives

Regulated and deregulated generators produce the same good, but face very different incentives. Deregulated generators can be categorized in three ways. First, they can sell into centrally-organized competitive wholesale markets. Second, they can sell to utilities and sometimes retail customers, either through long-term PPAs or through shorter term provisional contracts. Third, they can be owned by the end-user of the power they generate, for instance by universities or manufacturers. The shutdown decision for deregulated generators comes down to profitability.[11]

In the case of wholesale markets, owners of generation submit bids to the independent system operator (ISO) which then dispatches generators in ascending order of marginal cost, subject to transmission constraints. The marginal generator's bid determines the wholesale price received by all inframarginal suppliers. Deregulated generator bids are constrained from below by fuel costs, which comprise a large portion of operating costs. A deregulated generator's profitability depends both on 1) the spread between electricity price and marginal cost, and 2) the amount of power sold into the wholesale market. When a generator's marginal cost increases, both of these factors can reduce revenue. If revenue earned falls below a certain level, then the generator is not recouping its fixed operations and maintenance costs, such as salaries and regularly scheduled maintenance. In this case, it may be justified to shut down the generator.

[11] In the analysis, I focus on first-time shutdowns rather than retirements. Shutdowns can be temporary or permanent, and the factors driving this choice are outside the scope of this paper (see Fleten et al. (2017)). For the purposes of this paper, I consider temporary shutdowns to be a strong precursor to permanent shutdowns (see Table 1). At the very least, a generator going on standby predicts a significant decrease in its long-run production, an outcome that does not hinder the interpretation of this paper's results.

For deregulated generators that sell power to utilities or retail customers, profitability depends on a generator's ability to sign a bilateral contract. Thus it still matters how a generator's marginal cost compares to its competitors, even though competition happens much less frequently. Generators which are unable to find contracts are a financial burden to the firm until they are shut down.

Finally, deregulated generators owned by the end-users of power are a profitable investment to the extent that they lower their owner's electricity bill.[12] If the outside option for electricity approaches the marginal cost of the end-user's generator, then it may no longer be justified to operate it.

The shutdown decision for regulated utilities is more complicated, especially for early retirements. Utilities earn a stream of revenues from capital expenditures in their rate base, which generally have predetermined depreciation schedules. In this context, an "early" generator retirement is a non-fully depreciated asset, which may include retrofits or other capital investments after the initial startup date. There are several options for regulators to adjust the utility's rate base when a generator shuts down early (Lehr & O'Boyle 2018).

In the simplest case, regulators could designate the generator as a "regulatory asset," and allow it to remain in the rate base according to the original depreciation schedule. On the opposite end of the spectrum, the regulators may disallow the remainder of the asset's value from the rate base, although this is usually viewed as an especially punitive measure.

[12] Entities that regularly purchase large amounts of power may find it cost effective to own some generation of their own because they can avoid the costs associated with transmission infrastructure, which are usually passed on to retail customers. They also may be valued for backup power in the case of outages, which are usually due to sudden loss of transmission capacity.

Regulators usually would like to preserve the utility's credit-worthiness because it saves money for customers if the utility has a lower cost of capital.

The intermediate case is when a generator is designated as a regulatory asset, but not according to the original depreciation schedule. Regulators can accelerate the rate of depreciation to meet early retirement dates, revise downward the remaining book value, disallow a regulatory asset at a later date, lower the rate of return on the regulatory asset, or a combination of these options. Regulators are tasked with managing utilities according to public interest, and customers may be unhappy to allow full recovery of investments when they have proven to be less useful than originally thought. Customers may also fight against the accounting of a regulatory asset at the original negotiated rate of return r, on the argument that r already reflects equity risk.

Thus, the decision to shut down a regulated generator involves substantial negotiation between utilities and regulators, the latter of which is mandated to keep prices low for customers while at the same time preserving the credit-worthiness of the utilities they regulate. Utilities have little incentive to push for early retirement of their assets because of regulatory uncertainty, which has little upside and substantial downside. Because utilities pass on costs to customers, they would rather keep plants operating until they are fully depreciated, regardless of the input costs. As shown in figure 4, it is quite rare for regulated generators to shut down before 35 years of age.[13]

[13] Another potential barrier to regulated shutdowns, relative to deregulated, is concern for lost jobs. When an entire plant shuts down, it marks the end of multi-decade careers for potentially hundreds of local workers. While this obviously applies to both regulated and deregulated plants, it probably weighs more heavily on the regulator's decision, who may view it as an additional argument against shutting down a plant.

2.3 Decarbonization policy

While the relationship between utility regulation and shutdown affects prices for customers, it is potentially even more important for determining the carbon intensity of the electricity grid. As support for aggressive decarbonization grows, coal and even gas plants are likely to be displaced by zero-carbon power sources. Kefford et al. (2018) highlight the challenge of stranded generating assets in the IEA's Two Degrees Scenario. While the vast majority of projected early retirements are coal, they still find some role for early gas retirement in the US. There are a few reasons why this should be considered a best case scenario for owners of gas plants.

First, it assumes that gas is irreplaceable for supplementing variable renewable generation, an assumption that may not hold over time as new storage technologies reach maturity. Second, it does not consider the role of upstream methane leakage, which can quickly reduce the carbon advantage of gas compared to coal. Third, it does not incorporate the UNFCCC doctrine of common but differentiated responsibility across countries, wherein wealthier countries are expected to decarbonize more quickly than less wealthy countries. Fourth, it targets a level of emissions consistent with a 50% chance of limiting average temperature increases to 2 degree Celsius. Many organizations are pushing for a stricter carbon budget than this, especially for wealthy countries like the US.

For all of these reasons, it is worth considering the effect of price pressure on the shutdown decisions for both coal and natural gas plants, as most decarbonization pathways consider early retirements of both. Carbon pricing is a relatively straightforward policy intervention that could significantly re-order the merit order for electricity (in the short run) and the levelized cost of energy across different technologies (in the long run). However, fossil fuel plants can last for decades, so the existing stock of generators operating at the inception of a carbon policy must be targeted by policy as well. If regulated generators are unmoved by rising input prices, then it is unclear how a simple carbon price would affect the stock of generators under the traditional utility model.

3 Conceptual Model

In order to make an appropriate comparison between regulated and deregulated generators, whose owners likely have very different objective functions, my conceptual model is based on the concepts of capacity factor and economic dispatch, which are relevant features of both sectors. In this conceptual model, I focus on natural gas and coal-fired generators. The price spread P^{Δ} is defined as the difference between the price of coal P_c and price of natural gas P_g, measured at the state-year level in dollars per million British thermal units ($/mmBtu).

$$P^{\Delta} = P_c - P_g \tag{2}$$

This price spread affects the position of a generator on the dispatch curve, relative to generators of the opposite fuel type. For instance, when P^{Δ} increases (natural gas becomes relatively cheaper compared to coal), some gas generators may overtake some coal generators in the merit order so that gas is dispatched at lower levels of demand, as in Cullen & Mansur (2017) and Fell & Kaffine (2018).

The capacity factor of a generator j in a given year t is a function of the probability of dispatch[14], which is itself a function of the following terms:

$$CF_j = f(Pr(dispatch)) = g(Q_d, \sum_i Q_s^g(P^\Delta), \sum_k Q_s^c(P^\Delta)) \tag{3}$$

Where Q_d is inelastic electricity demand, $\sum_i Q_s^g$ is the quantity produced by other gas resources at a lower marginal cost, and $\sum_k Q_s^c$ is the quantity produced by other coal resources at a lower marginal cost.

For both coal and gas generators, the following relationships hold. First, $\partial CF_j / \partial Pr(dispatch) > 0$ for reasons discussed above. $\partial CF_j / \partial Q_d > 0$, because more total generation is needed when demand is high. $\partial CF_j / \partial \sum_i Q_s^g < 0$ and $\partial CF_j / \partial \sum_k Q_s^c < 0$ because competitors with lower marginal cost reduce the probability that generator j will be dispatched, holding all else equal.

For gas generators only, $\partial \sum_i Q_s^g / \partial P^\Delta = 0$, because a change in P^Δ does not affect the ordering of gas generators with respect to each other, on the assumption that fuel costs are linear in production. On the other hand, $\partial \sum_k Q_s^c / \partial P^\Delta < 0$, since an increase in P^Δ leads to fewer coal generators below generator j in the dispatch order. Similarly, for coal generators only, $\partial \sum_k Q_s^c / \partial P^\Delta = 0$ and $\partial \sum_i Q_s^g / \partial P^\Delta < 0$ since the opposite relationship is true for P^Δ from a coal generator's perspective.

A generator j is considered to be competitive in period t if its capacity factor CF_{it} is high enough to justify continued operation. When a generator shuts down, it avoids some fixed

[14] Generators usually operate at close to either 0% or 100% capacity at any given moment. Thus, a generator with an annual capacity factor of 0.4 can be interpreted as operating about 40% of the time during that year.

operating cost C_j. The manager's decision thus boils down to:

$$V_{var}(CF_{jt}(P^{\Delta}_{jt})) + V_{fixed}(C_j) \geq V_{shutdown}(S_{jt} = 1) \longrightarrow \text{operate} \tag{4}$$

$$V_{var}(CF_{jt}(P^{\Delta}_{jt})) + V_{fixed}(C_j) < V_{shutdown}(S_{jt} = 1) \longrightarrow \text{shut down} \tag{5}$$

For both regulated and deregulated generators, $V_{shutdown}$ includes legally-mandated decommissioning and site cleanup costs, net of potential scrap value. However, in the regulated case it also includes potential financial losses in the process of conversion to a regulatory asset, as discussed in Section 2. When a generator is fully depreciated, this additional component falls to zero. All else equal, a non-fully depreciated regulated generator will continue operating at a lower capacity factor than an equivalent deregulated generator.

There are a few advantages to using capacity factor as a basis for comparison across regulated and deregulated generators.

First, I am able to avoid the misuse of terms such as profit, which would be appropriate for deregulated firms, but not for regulated utilities. Utilities view shutdown decisions in the context of a larger resource picture, outlined in Integrated Resource Plans. The profit impact of an individual shutdown is not deterministic from the utility's perspective because of regulatory uncertainty. Profit impacts depend on how the utility replaces lost production, if at all.[15] If a utility replaces a retired generator with another generator that contributes to their rate base, they would be allowed to raise future prices according to the revenue requirement equation. However, the regulators may not allow this, instead pushing the utility to meet

[15] One important consideration is that the US electricity grid is technically continuous over large regions, and new transmission capacity is always possible. Therefore it is hard to say that any single generator replaces another.

demand from outside sources. For instance, this could be a power purchase agreement (PPA) with an IPP or another utility. It could also mean participation in a restructured wholesale market, if that option is available.

Second, electricity supply is organized by dispatch curves in essentially every market, and there is strong theoretical justification for this on efficiency grounds. In the short run, the least-cost method of meeting any level of electricity demand is by dispatching available generation capacity in ascending order of marginal cost. Whenever a generator is dispatched, the price it ultimately receives (from customers or other buyers) is greater than the marginal cost of production. This is necessarily true for regulated generators, which charge retail prices according to the "cost plus" formula discussed in Section 2. This is also a reasonable assumption for deregulated generators, which have no reason to operate at a loss in the short term. Thus a higher capacity factor implies that a generator is more frequently among the least-cost sources of power to supply the grid.

Third, capacity factor has helpful empirical properties. There is a substantial amount of variation in capacity factor across fossil fuel generators, as shown in Figure 1.

It should be noted that some technologies are designed to specialize as peaking generation ("peakers"), which means they only operate when demand is higher than usual. Capacity factor is less useful for evaluating the merit of peakers because they are primarily intended for reliability rather than low-cost bulk generation. Peakers in wholesale markets generate revenue either through capacity market payments or very high peak prices. Peakers are generally small natural gas or oil-fired generators. I do not focus on peakers in this paper for a few reasons. Compared to bulk fossil fuel generators, peakers are less environmentally important, less responsive to average prices, and less likely to be pushed into early retirement by carbon pricing or renewables. In fact, natural gas peakers will likely be the last type of fossil fuel generator to become obsolete in the long run because they complement intermittent renewables so well (Rode et al. 2017).

In the empirical analysis below, I aim to explain how prices and capacity factor explain the pattern of shutdowns observed in the data. Section 5.1 describes the effect of prices on shutdown decisions, while Section 5.2 is a decomposition of the shutdown decision, via the effect of prices on capacity factor and capacity factor on shutdowns. Section 5.3 provides further support for the conceptual model by examining how the state-level resource mix affects the sensitivity to the price effect in 5.1.

4 Data

4.1 Sources

I use the form EIA-860 to describe annual (2006-2017) generator and plant-level characteristics, including operating status. Owners of generators submit the form early in the subsequent year, so the reported operating status is in retrospect. There are three relevant operating statuses that could be reported for a generator; operating (OP), standby (SB)[16], or retired (RE). "Operating" indicates that a generator is in service (commercial operation) and is producing some electricity. This includes peaking units that produce power situationally.

[16] Standby category also includes the "out of service" statuses, which are similar

Standby indicates that a plant did not operate in a given year, but could open again at a future date. Retirement means that a plant is shut down permanently, and may be in the process of physical decommissioning.

Table 1: Transition Matrix

	OP_{t-1}	SB_{t-1}	RE_{t-1}
OP_t	38,669	190	8
SB_t	234	2,119	12
RE_t	258	185	2,721

The distinction between standby and retirement is not always clear, since the EIA-860 form is self-reported, and classification may be subjective for generators that are on a clear path to retirement. For my main analysis, I combined operating statuses SB and RE into a single status I refer to as "shutdown". As shown in Table 1, about 5% of generators return to OP from SB, in roughly equal proportion for both regulated and deregulated generators.[17]

For the purposes of this analysis, I focus on the changes in market conditions that lead a generator to shut down for the first time. I argue that this is the appropriate outcome variable for measuring a generator's competitiveness. Although some shutdowns are not permanent, utilities likely do not know ex-ante whether or not the generator will be re-started. Fleten et al. (2017) investigate the shutdown and startup decisions in a real options model, finding that regulatory uncertainty plays an important role in whether to re-start or abandon generators that are shut down.

There is heterogeneity in reporting practices in form 860. Some generators report being

[17] For generators that remain in SB for multiple periods, the probability of re-starting declines but levels off at around 4% in each period. The ratio of generators moving to OP/RE after some number of periods of SB levels off around 2:3

on standby or retired for several years in a row, while others go directly from operating to out of sample. For the former, I drop observations after the first year of shutdown. For the latter, I drop all observations, as I cannot conclusively determine if they have shutdown. This eliminates a source of bias from false negatives or false positives. Other important variables from EIA-860 include plant-level regulation and location, and generator-level age, capacity, and technology type.

I also use EIA data for annual, state-level natural gas prices. The EIA provides an incomplete range of natural gas prices used for electric power, due to confidentiality. Instead, I use the "citygate" prices of natural gas.[18] Citygate price should be a relatively unbiased estimate of the price of natural gas for electric power, as the latter is largely determined upstream of the power plant.

EIA's annual state-level data for observed coal prices are similarly limited due to confidentiality. Instead, I use data from EIA's Annual Energy Outlook (AEO) which reports average coal prices used for electricity across nine EIA-defined regions in the US.[19] I also use the AEO price forecasts as a proxy for expected prices of both natural gas and coal, which is necessarily regional as well.

Finally, I use power production data from EIA-923, which is reported at the plant-year-fuel level. I apportion this data across generators according to the method described in Section 5.2.

[18]Citygate: A point or measuring station at which a distributing gas utility receives gas from a natural gas pipeline company or transmission system.

[19]East North Central, East South Central, Middle Atlantic, Mountain, New England, Pacific, South Atlantic, West North Central, West South Central

My analysis focuses on natural gas generators. The biggest reason for this choice is that coal shutdowns are motivated by a broader array of concerns which fall outside the conceptual model. These include environmental regulations and public pressure campaigns from activists, in addition to all of the factors in this analysis. As an example, owners of coal generators facing binding pollution constraints commonly invest in capital-intensive pollution equipment as an alternative to shutdown. Fowlie (2010) finds that regulated plants are more likely to invest in this type of equipment, which would be a complication to my identification strategy. Coal generators also sometimes convert to run on gas, whereas the reverse is not true. As another empirical difficulty, there are far fewer coal generators to observe, and they have less variation in observed capacity factor.

4.2 Descriptive Statistics

Regulated and deregulated generators have fairly similar physical characteristics, as shown in Table 2. Regulated gas generators tend to be somewhat smaller and older. Figure 2 shows age distributions and Figure 3 shows size distributions for generators that are operating in 2017.

Perhaps most notable difference between regulated and deregulated generators is in the age at first shutdown, as shown in Figure 4. Deregulated generators display a wide range of retirement ages, peaking in frequency from 20 to 30 years, but with many lasting for decades longer. On the other hand, regulated generators show a sharp increase in retirements around 40 years. This pattern is potential evidence of the effect of rate-of-return regulation; utilities

Table 2: Summary Statistics

	Deregulated	Regulated
(in 2017)		
generators	1,519	1,821
combined cycle	57%	40%
steam	8%	11%
other technology	35%	49%
size	145 MW	136 MW
age	20 years	22 years
(all years)		
capacity factor	36%	23%
P^{Δ}	-$3.67	-$3.43
shutdowns	231	224

Summary statistics for operating generators over 25MW

have little incentive to shut down a generating asset before it is fully depreciated.

The other important difference between the two is in average capacity factor, with deregulated generators operating at around 13% higher annual capacity factor. The difference persists even without the observations for which capacity factor is zero. Part of this is explained by the higher prevalence of combined cycle generators owned by deregulated firms, which tend to be the youngest and most fuel efficient among gas generators. However, regulated generators still have a lower average capacity factor controlling for generator characteristics. This observation also suggests a rate-of-return regulation may distort incentives, where regulated generators are profitable for the utility to operate for reasons unrelated to their ultimate competitiveness. There is also a small difference in average P^{Δ}, indicating that regulated gas generators faced slightly better prices over the sample period. This suggests the difference in average capacity factor is likely even larger when controlling for the price spread. Figure 6 shows the distribution of P^{Δ} across all observations in the sample.

Taken together, and without controlling for generator characteristics, the summary statistics indicate a potentially important difference in the way that regulated generators are managed. The empirical analysis in this paper aims to causally identify this difference in the decision to shut down a generator for the first time.

5 Methodology

5.1 Effect of prices on shutdowns

The unit of observation in this model is the generator-year.[20] My main analysis uses a linear probability interaction model to estimate the probability of shutdown. For a generator j in state s and year t,

$$status_{jt} = P^{\Delta}_{st}(\beta_1 + \beta_2 reg_j) + \boldsymbol{X}_{jt}(\beta_3 + \beta_4 reg_j) + \beta_5 reg_j + \alpha_t + \phi_s \tag{6}$$

[20] It is also possible to run the analysis at the plant-year level. Many plants are composed of multiple generators, which may run in tandem or separately. Ultimately, the long-run operating status of each generator within a plant is determined separately. Running the analysis at the generator-year level makes use of additional variation that would otherwise be lost in the case of within-plant differences in operating status. Additionally, generators within plants may have different ages and capacities, so it is not straightforward to aggregate these to the plant level. I include a specification with plant-year clustered standard errors as a robustness check to account for potentially correlated outcomes.

The dependent variable $status_{jt}$ takes a value of 1 if the generator shuts down for the first time in year t.[21] The coefficients of interest are β_1, the effect of prices on the probability of shutdown for deregulated generators, and $\beta_1 + \beta_2$, the price effect for regulated generators. A negative value for either indicates that as the price spread increases (i.e. gas becomes cheaper relative to coal), a generator is less likely to shut down. β_2 is identified by comparing regulated generators to otherwise similar deregulated generators in the same region and year, which face the same prices and generally supply power in the same markets.

P^{Δ}_{st} represents the coal-gas price spread[22], which is plausibly exogenous with respect to shutdown decisions. Most of the variation in the coal-gas price spread is driven by the price of natural gas. Variation in gas prices in this sample period (2006-2017) was primarily driven by the shale gas boom, which is determined almost entirely by geographic location of shale formations (Fell & Kaffine 2018). Natural gas prices are also affected by other upstream uses, such as industrial power and heating for buildings. Weather affects demand for gas used for heating, so gas prices are also affected by weather-related exogenous shocks to availability. Decisions about when and where to build natural gas-fired plants are likely related to expectations about prices, but these are very difficult to forecast over the time frame that most gas plants exist. Thus it is unlikely that prices are meaningfully affected by the shutdown decision.

[21] A small number of generators return to operation in later years after shutting down, and an even smaller number then shut down again. I drop all years after the first year of shutdown. As a result, this specification and all subsequent specifications should be interpreted as the factors that cause a generator to shut down for the first time.

[22] P^{Δ}_{st} in principle could be constructed from contemporaneous, lagged, or forecast prices

The variable reg_j is a generally time-invariant, plant-level, binary indicator for regulation, taking the value of one if a generator sells power for prices that are set by regulators, according to the formula discussed in Section 2.1. I argue that reg_j is also plausibly exogenous to shutdown decisions, as shutdowns reflect a strategic response to conditions which are very difficult to predict more than a few years ahead of time. This variable would be endogenous if generators changed ownership in expectation of a future status change, which would be hard to measure. However, there are only a few changes in reg_j over my sample period from 2006-2017, so it is unlikely to cause interpretation problems.[23]

$\boldsymbol{X}_{jt}$ includes age, capacity, and technology indicators.[24] Age is obviously exogenous to shutdown decisions. Capacity and technology are almost entirely determined at the beginning of the generator's lifespan, so they are plausibly exogenous to shutdown decisions for the same reason as reg_j discussed above. $\boldsymbol{X}_{jt}$ also includes region-annual electricity demand, Q_d. This term helps to control for the potentially confounding effect of changing reserve margins on the grid. If a region-year has higher demand, it should reduce the probability of shutdown for a given generator, all else equal, as generators in higher-demand regions are more likely to be dispatched. I include year fixed effects α_t and regional[25] fixed effects ϕ_s which are not interacted with reg_j. I limit my analysis to generators larger than 25MW in an effort to avoid including peaking generators, which tend to be smaller.[26]

[23]Mandated divestiture of regulated assets was much more common from 1998-2001 (Bushnell & Wolfram 2005).

[24]Technology indicators are for 1) combined cycle and 2) single-cycle steam, with the omitted category mostly comprised of internal combustion turbines

[25]Nine regions as described in the EIA's Annual Energy Outlook

[26]25MW is also the lower-bound cutoff for generator-level production in the EPA CEMS data. In all

5.2 Mechanisms affecting probability of shutdown

While P^{Δ} represents a shock to input costs which may affect the propensity to shut down, capacity factor (CF) is the key mechanism described in the conceptual model. As P^{Δ} increases, all gas generators become more efficient relative to all coal generators, leading some gas generators to surpass some coal generators in the merit order. This increases the probability of dispatch (and thus CF) for gas generators, as they are the least cost option at a lower level of demand. The price effect on shutdowns can be decomposed according to the following equation:

$$\frac{\mathrm{d}Pr(S)}{\mathrm{d}P} = \frac{\partial Pr(S)}{\partial CF}\frac{\partial CF}{\partial P} \tag{7}$$

Knittel et al. (2019) find that $\frac{\partial CF}{\partial P}$ is greater for plants in vertically integrated markets than for plants in restructured wholesale markets. However, because my analysis looks at a plant-level definition of regulation (whether or not the asset is part of a utility's rate base), their results do not necessarily shed light on this decomposition. Instead, I show in Table 6 that $\frac{\partial CF}{\partial P}$ is conditionally positive for both regulated and deregulated generators, and the difference is not significant.

The conceptual model predicts that a lower generator-level capacity factor increases the probability of shutdown. However, generators that shut down have a capacity factor of zero by definition. Thus, a simple regression of status on capacity factor is biased because capacity factor is endogenous to operating status. To overcome this, I use an instrumental variables-

specifications, I also drop observations from two plants which shut down because of damage from Hurricane Sandy.

style approach, predicting $\widehat{CF}$ with P^{Δ}. This allows me to impute the capacity factor for all generators in the sample, even those that shut down.

Similar to the specification in the previous section, for a generator j in state s and year t,

$$status_{jt} = \beta_1 \widehat{CF}_{st}(P^{\Delta}_{st}, \boldsymbol{X}_{jt}) + \beta_2 \widehat{CF}_{st}(P^{\Delta}_{st}, \boldsymbol{X}_{jt}) * reg_j + \tag{8}$$

$$\boldsymbol{X}_{jt}(\beta_3 + \beta_4 reg_j) + \beta_5 reg_j + \alpha_t + \phi_s$$

P^{Δ} and $P^{\Delta} * reg$ are excluded instruments for $\widehat{CF}$ and $\widehat{CF * reg}$, and the vector of covariates $\boldsymbol{X}_{jt}$ are included instruments. In order for P^{Δ} to be a valid instrument in this context, they need to satisfy two requirements. First, they need to be correlated with capacity factor, conditional on other covariates. Second, they need to be exogenous in respect to operating status, conditional on other covariates. The coal-gas price spread satisfies the first requirement very well, as it is a very strong predictor of capacity factor (see Table 6). This is consistent with the conceptual model, where changes in generator-level marginal cost can rearrange the dispatch order. As discussed in the previous section, it also plausibly satisfies the second assumption.

I hypothesize that β_1 is negative, meaning that a higher capacity factor for deregulated generators is associated with a lower probability of shutdown, all else equal. On the other hand, I hypothesize that β_2 is positive, indicating that regulated generators are less likely to shut down when their capacity factors are low, relative to their deregulated counterparts.

In order to create a generator-level measure of annual production, I use electricity production data from EIA-923, which is converted to a measure of generator-level capacity factor. EIA-

923 reports annual generation data, broken down by plant, year, and fuel type.[27] I apportion the quantity produced across generators of the same plant-year-fuel in proportion to their capacity, for operating generators. I then divide the production by the generator's maximum possible production in a year, measured by the generator's nameplate capacity multiplied by the number of hours in a year. This method relies on two assumptions.

First, gas generators within the same plant must have the same capacity factor. This assumption is less likely to hold if generators within a plant have different characteristics, such as age and size. Younger and larger generators tend to be used more often (see Table 6), so to the extent that this assumption is violated, it should bias the coefficients on $\boldsymbol{X}_{jt}$ toward zero. Because generators within the same plant experience the same price conditions, this assumption is unlikely to bias the coefficients of interest, β_1 and β_2.

Second, generators that shut down in the middle of the year are assumed to have a capacity factor of zero in that year. This is consistent with the operating status definitions given in the EIA-860 instructions, which require a generator to report as OP if it produced any power at all, and SB or RE if it did not. It is unclear whether plant managers adhere strictly to these definitions. This assumption would be violated for a generator that reports being shut down during year t, but really was only shut down for part of the year. Violating this assumption would bias the coefficients on $\boldsymbol{X}_{jt}$ away from zero, in contrast to violations of the first assumption. However, the coefficients of interest would likely not be biased, for the

[27] EIA-923 also reports at the "prime mover" level, e.g. steam turbine, internal combustion engine, or either part of a combined cycle unit. I chose not to use this additional dimension in my capacity factor variable because the merge with the EIA-860 data was flawed, which would have introduced measurement error.

same reason as before.

6 Results

6.1 Effect of prices on shutdowns

The LPM results are shown in Table 3, with my preferred specification in column 5. I find a statistically significant negative coefficient on P^{Δ}, which indicates that deregulated generators are less likely to shut down when prices are advantageous, and more likely to shut down when they prices are disadvantageous. This is consistent with the conceptual model, where owners of deregulated generators are incentivized to shut down when their capacity factor falls below some threshold. This estimate implies that a \$1 per mmBtu decrease in P^{Δ} is associated with an increase in the probability of shutdown by 0.375 percentage points, from a baseline of around 1.2 percentage points.

In contrast, regulated generators essentially do not respond to changes in price conditions. Column (6) shows the preferred specification with P^{Δ} measuring the price effect for regulated generators, which is not significantly different from zero. This is consistent with the pattern shown in Figure 4, where it appears that the probability of shutdown for natural gas generators is very low until a certain age threshold, beyond which it becomes much more common. This result is consistent with criticisms of rate of return regulation, where input costs are passed through to customers, and utilities are not directly incentivized to minimize them.

Unsurprisingly, age is also an important factor determining the operating lifespan of a generator before its first shutdown. This holds roughly true for both regulated and deregulated generators. Larger regulated generators are less likely to shut down, all else equal, although this effect does not seem to hold for their deregulated counterparts. This result supports the idea that regulators and utilities are averse to concentrated job losses (see footnote in section 2.2). There are other potential explanations for this, including economies of scale in non-fuel variable costs such as labor. It may be more efficient from an organizational perspective to shut down two small generators than a single large one, although it is less clear why these explanations would not also hold for deregulated firms. I find a negative but insignificant effect for regional electricity demand Q_d, implying that deregulated generators are slightly less likely to shut down when demand is higher. Regulated generators on the other hand respond in the opposite direction. Regardless, the inclusion of Q_d and Q_d*reg in the model has a negligible effect on the coefficients of interest, P^{Δ} and P^{Δ}*reg, so it is unlikely that the need to preserve generators in times of higher demand is affecting the price sensitivity of shutdowns.

6.2 Effect of predicted capacity factor on shutdowns

Table 4 shows the results of the instrumental variables-style specification, again with column 5 as the main specification. I use P^{Δ} and $P^{\Delta} * reg$ as a vector of excluded instruments to predict $\widehat{CF}$ and $\widehat{CF} * reg$, which I use to estimate probability of shutdown. I find a statistically significant negative coefficient on capacity factor, which indicates that uncompetitive deregulated generators are more likely to shut down. A 10% decrease in capacity factor for deregulated generators is associated with about a 1.25% increase in probability of shutdown. I also find that the coefficient on the interaction term more than overtakes the main effect, indicating that regulated generators actually respond in the opposite direction, with advantageous prices increasing the probability of shutdown. However, interacting CF with an indicator for "deregulated" instead of reg shows that the coefficient for regulated generators is not significantly different from zero.

It is also worth noting that regulated generators tend to have a significantly lower capacity factor across my sample than deregulated generators, even controlling for prices and generator characteristics (see Table 6). The coefficient on reg in Table 4 indicates that regulated generators are significantly less likely to shut down at the same predicted capacity factor as deregulated generators. This is consistent with the conceptual model, where regulated generators incur an extra penalty from early retirement that is not present for deregulated generators, which is the potential for regulators to disallow some portion of future profits associated with a given capital asset. Once again, age is an important predictor of shutdowns, although this time the effect is significantly larger for regulated generators. The results for capacity are qualitatively similar to those in Table 3.

Because the endogenous variable, CF is interacted with the plausibly exogenous indicator reg, there are two endogenous variables, and there is not a straightforward two stage least squares interpretation of the IV-style specification. Instead, Table 6 shows the first stage of the (hypothetical) non-interacted version of this model. As shown in Table 6, P^{Δ} strongly predicts capacity factor. However, this relationship does not appear to vary consistently across regulated and deregulated generators.

Taken alongside the results from Table 4, this suggests that the difference in shutdown behavior is driven by regulated utilities' willingness to continue operating generators with low capacity factors. The results do not support an alternative explanation, that utilities support these uncompetitive generators by ignoring prices and artificially inflating their capacity factors.

6.3 Price effect in high-coal regions

In support of my conceptual model, I also analyze how the response to P^{Δ} varies according to its competitors. As discussed in the conceptual model, P^{Δ} influences a gas generator's position on the dispatch curve relative to its coal-fired competitors, but not relative to other gas generators. In regions where most of its competitors are powered by gas, P^{Δ} should theoretically have a smaller effect on a gas generator's competitiveness, since there are fewer potential coal generators to displace in the merit order. On the other hand, gas generators in high-coal regions should be even more sensitive to P^{Δ} for the opposite reason.

To test this hypothesis, I constructed a variable HC_s indicating whether a generator is located in a high-coal state. I used EIA data describing the state-year consumption of coal and natural gas for electricity generation, and averaged the coal/gas consumption ratio for each state across the years in my sample. For the main heterogeneous effects specification, I defined high-coal states as those in the top half of the rankings, according to this measure. I estimated equation (5) again, this time interacting the P^{Δ}_{st} and reg_j with the indicator variable HC_s:

$$status_{jt} = P^{\Delta}_{st}(\beta_1 + \beta_2 reg_j + \gamma_1 HC_s + \gamma_2 HC_s reg_j) + \boldsymbol{X}_{jt}(\beta_3 + \beta_4 reg_j) + \beta_5 reg_j + \alpha_t + \phi_s \tag{9}$$

Results shown in Table 5 support this hypothesis. The first specification includes observations all states, the second drops the middle ten, and the third drops the middle twenty. The negative coefficient on P^{Δ}*hicoal indicates that deregulated generators in high-coal states are more likely to shut down in response to unfavorable price conditions, relative to deregulated generators in low coal states. As I drop more of the middle states, this coefficient becomes larger and more significant, which further supports the hypothesis.

On the regulated side, P^{Δ}*hicoal*reg is the triple interaction term, showing the additional probability of shutdown for regulated generators in high-coal states in response to price conditions. The term is not significantly different from zero, which indicates that once again, regulated generators are unresponsive to price conditions, even in high-coal areas where the price conditions should theoretically matter more for their probability of dispatch.

6.4 Robustness checks

Tables 7 and 8 show the estimates from Tables 3 and 4, respectively, with robust standard errors in column 1, clustered at the plant level in column 2, generator level in column 3, state level in column 4, and AEO region-year level in column 5. In both specifications, the precision is lower when clustered at the plant level, as there are several plants with multiple generators that close down all at once. This suggests there is room in the conceptual model for additional consideration of within-plant dynamics.[28] Clustering at the generator level in column 3 accounts for serial correlation, but still treats within-plant observations as independent. Comparing columns 2 and 3 to column 1 indicates that within-generator serial correlation is likely not as important as within-plant correlation. Column 4 allows for correlation within a state but not over time, and column 5 allows for correlation within a region-year, which may capture related outcomes from grid-level constraints about the total amount of capacity needed at a given point in time.

Table 9 shows alternative LPM specifications. Column 1 is the same as Table 4. Column 2 includes age squared and its interaction with *reg*. This specification accounts for the

[28]For instance, generators within a combined-cycle gas unit operate in parallel, but plant managers are ultimately able to shut down the component generators separately, and they can be converted into or away from combined cycle units. Another example of within-plant dynamics is that some plants feature multiple identical generators that operate in tandem. Theoretically, this demonstrates a trade-off between the efficiency of fewer, larger generators and the logistical advantages of more, smaller modular generators. For instance, a single small generator can be taken offline for maintenance with lower disruption to the overall power supply, and it can enter or exit service independently of the rest.

possibility that the probability of shutdown does not increase linearly with age, but rather exponentially. This would be consistent with a model in which the oldest generators are disproportionately likely to shut down. Results are very similar to the main specification in column 1, suggesting that the coefficients on P^{Δ} and $P^{\Delta} * reg$ are not biased by the differences in age distribution across regulated and deregulated generators.

Column 3 of Table 9 interacts technology indicators with age. The technologies considered are combined cycle, single-cycle steam, and "other", which are mostly combustion turbines. It is reasonable to suppose that the expected lifespans of these three categories are not the same, so this specification tests the sensitivity of the results to the distribution of ages across technology types. Results are similar to the main specification, although there is a slight shift in explanatory power away from the price variables and towards the age variables.

Column 4 of Table 9 shows the same specification as in Table 3, but without year fixed effects. This allows identification through inter-temporal variation in the covariates. Omitting year FEs unlocks significant variation in P^{Δ} over time, which is helpful for identification, but it also introduces bias from common trends. My sample period includes several significant common trends, including the rise of renewable generation, decline in the use of coal, and the 2008 recession. Regardless, the estimated coefficients of interest are qualitatively similar, so this appears not to be an important dimension.

7 Counterfactual

In the results above, I find that regulated generators are much less responsive to prices in the shutdown decision. In this section, I present a counterfactual scenario aiming to describe how much the stock of regulated gas generators would be different had there been total deregulation. The results in 5.1 are from a linear probability model, which provide the describe the average treatment effect, so I avoid characterizing which individual plants would have shut down, focusing instead on the aggregate capacity of first time shutdowns induced by a change in regulatory status.[29]

7.1 Within sample period

In this section, I use the estimates from Table 4 to describe the predicted effect of deregulation on first-time shutdowns for generators within the sample period. I estimate the predicted change in shutdown probability by switching the interaction term on price spread to zero for all generators. Thus the counterfactual must be interpreted as the effect of prices on regulated generators, if they responded to prices as if they were deregulated.

Table 10 shows the results of the procedure. The second column, Δcap, shows the year by year predicted changes in capacity if regulated generators responded to prices in the same way as deregulated generators. Yearly estimates range from 706 MW to 2,508 MW of extra

[29]As an important caveat, I am aiming to describe the total capacity of generators that would have shut down for the first time, which is an important precursor to retirement, rather than retirement itself.

shutdowns, which represent between 0.3 and 1.44 percent of operable regulated gas capacity (shown in column 4). For context, the capital cost of a new natural gas combined cycle plant is around $1 million per MW (Lazard 2018). The average natural gas generator in my sample is around 140 MW in capacity, so the estimates imply that total deregulation would result in an additional 5-18 average sized generators shutting down for the first time each year.

In order to characterize the effect of deregulation on retirement, and not just shutdowns, the simplest method is to extrapolate based on observed status changes in the data. From 2006 to 2017, 258 generators moved from operable to retired, while 234 generators moved from operable to standby. Of the generators that went on standby for at least one year in the sample, 53 were operable as of 2017, implying that 11% of the shutdowns were not permanent. There were 69 generators still on standby as of 2017, or 14% of all shutdowns. Thus first-time shutdowns as measured in this analysis ended up being permanent between 75-89% of the time. Regardless of permanence, each shutdown represents a significant decline in the expected lifetime production of a given generator, especially given the relatively old age of generators that come back online from standby (around 33 years at time of first shutdown). While the estimates shown in Table 10 (for first time shutdowns) are likely greater than for permanent shutdowns, they are consistent in the short term, only gradually declining over time as a few generators may come back online from standby.

8 Conclusion

The results of this paper support an early argument for deregulation, that forcing owners of generation capacity to compete based on costs would ultimately increase the efficiency of the electricity grid. I find that deregulated generators respond to input prices in a way that is consistent with the conceptual model described above. When input prices are favorable, deregulated generators are less likely to shut down, and the inverse is true as well. A $1 increase in the price spread is associated with a decrease in probability of shutdown by 0.375 percentage points, from a baseline of around 1.2 percentage points. On the other hand, deregulated generators were found to be mostly unresponsive to prices.

I also find that deregulated generators are also responsive to predicted changes in capacity factor, which indicate they are more likely to shut down when they are less profitable. A 10% increase in predicted capacity factor lowers the probability of shutdown by about 1.25 percentage points. However, regulated generators, which are shown to have a much lower average capacity factor, may actually respond in the opposite direction to changes in predicted capacity factor, indicating a clear disconnect between competitiveness and probability of shutdown. Further decomposition of this difference indicates that while prices have an important effect on capacity factor, this effect is does not vary on average across regulated and deregulated generators. This supports the argument that regulated generators are continuing to operate at lower capacity factors than would be tolerable for deregulated generators.

Counterfactual estimates indicate that full deregulation of existing regulated generators would have caused an additional 700 to 2,500 MW of first-time shutdowns each year of the sample, as the least efficient generators are incentivized to stop operating. At a capital cost of about $1 million per MW for new gas capacity, this implies that full deregulation would have shut down an additional $700 million - $2.5 billion per year of capital assets.

These results are driven by generators changing relative position along the dispatch order. Thus, I argue that my findings may generalize to other situations in which fossil fuel generators are pushed out of the merit order. This includes the introduction of competition from low-marginal cost renewable sources, and policies such as carbon taxes that raise the marginal cost of fossil fuel sources relative to others. The electricity industry is in a complicated transition phase, where many markets are still figuring out how to incorporate new technologies into the generation mix while still promoting competition and innovation. To the extent that the optimal solution involves replacing these relatively high-marginal cost fossil fuel generators, the role of utility regulation in keeping these generators operating must be considered as a potentially important roadblock to progress.

References

Borenstein, S. & Bushnell, J. (2015), 'The us electricity industry after 20 years of restructuring', *Annu. Rev. Econ.* **7**(1), 437–463.

Bushnell, J. B. & Wolfram, C. (2005), 'Ownership change, incentives and plant efficiency: The divestiture of us electric generation plants'.

Cicala, S. (2015), 'When does regulation distort costs? lessons from fuel procurement in us electricity generation', *American Economic Review* **105**(1), 411–44.

Cicala, S. (2017), Imperfect markets versus imperfect regulation in us electricity generation, Technical report, National Bureau of Economic Research.

Csereklyei, Z. & Stern, D. I. (2018), 'Technology choices in the us electricity industry before and after market restructuring', *The Energy Journal* **39**(5).

Cullen, J. A. & Mansur, E. T. (2017), 'Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution', *American Economic Journal: Economic Policy* **9**(3), 106–33.

Davis, L. W. & Wolfram, C. (2012), 'Deregulation, consolidation, and efficiency: Evidence from us nuclear power', *American Economic Journal: Applied Economics* **4**(4), 194–225.

Fabrizio, K. R., Rose, N. L. & Wolfram, C. D. (2007), 'Do markets reduce costs? assessing the impact of regulatory restructuring on us electric generation efficiency', *American Economic Review* **97**(4), 1250–1277.

Fell, H. & Kaffine, D. T. (2018), 'The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions', *American Economic Journal: Economic Policy* **10**(2), 90–116.

Fleten, S.-E., Haugom, E. & Ullrich, C. J. (2017), 'The real options to shutdown, startup, and abandon: Us electricity industry evidence', *Energy Economics* **63**, 1–12.

Fowlie, M. (2010), 'Emissions trading, electricity restructuring, and investment in pollution abatement', *American Economic Review* **100**(3), 837–69.

Joskow, P. (2008), 'Lessons learned from electricity market liberalization. the energy journal, special issue', *The Future of Electricity: Papers in Honor of David Newbery* .

Kefford, B. M., Ballinger, B., Schmeda-Lopez, D. R., Greig, C. & Smart, S. (2018), 'The early retirement challenge for fossil fuel power plants in deep decarbonisation scenarios', *Energy policy* **119**, 294–306.

Knittel, C. R., Metaxoglou, K. & Trindade, A. (2019), 'Environmental implications of market structure: Shale gas and electricity markets', *International Journal of Industrial Organization* **63**, 511–550.

Lazard (2018), *Lazard's Levelized Cost of Energy Analysis-Version 12.0*.
URL: *https://www.lazard.com/media/450784/lazards-levelized-cost-of-energy-version-120-vfinal.pdf*

Lehr, R. & O'Boyle, M. (2018), 'Depreciation and early plant retirements', https://energyinnovation.org/wp-content/uploads/2018/12/Depreciation-and-Early-Plant-Retirements-Brief_12.3.2018.pdf. Accessed: 2019-7-26.

Rode, D. C., Fischbeck, P. S. & Páez, A. R. (2017), 'The retirement cliff: Power plant lives and their policy implications', *Energy Policy* **106**, 222–232.

Williams, J. H., DeBenedictis, A., Ghanadan, R., Mahone, A., Moore, J., Morrow, W. R., Price, S. & Torn, M. S. (2012), 'The technology path to deep greenhouse gas emissions cuts by 2050: the pivotal role of electricity', *science* **335**(6064), 53–59.

Zhang, F. (2007), 'Does electricity restructuring work? evidence from the us nuclear energy industry', *The Journal of Industrial Economics* **55**(3), 397–418.

9 Figures and tables

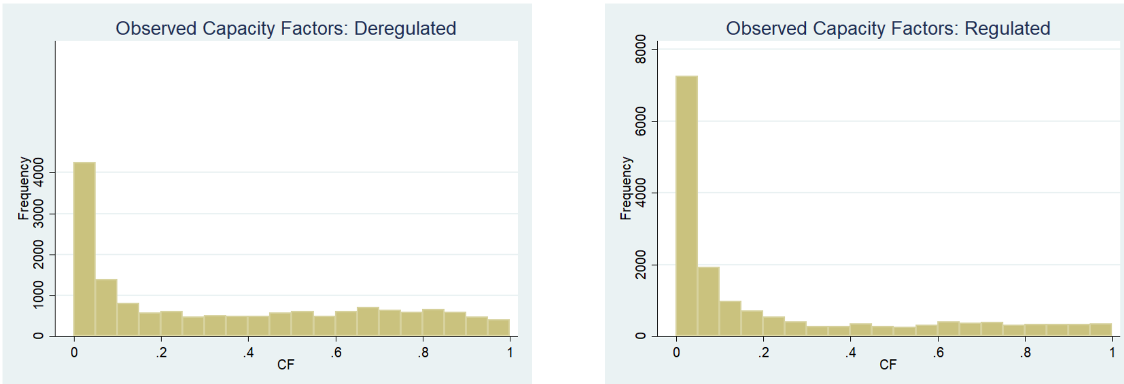

(a) (b)

Figure 1: Observed capacity factors for operable generators, 2006-2017

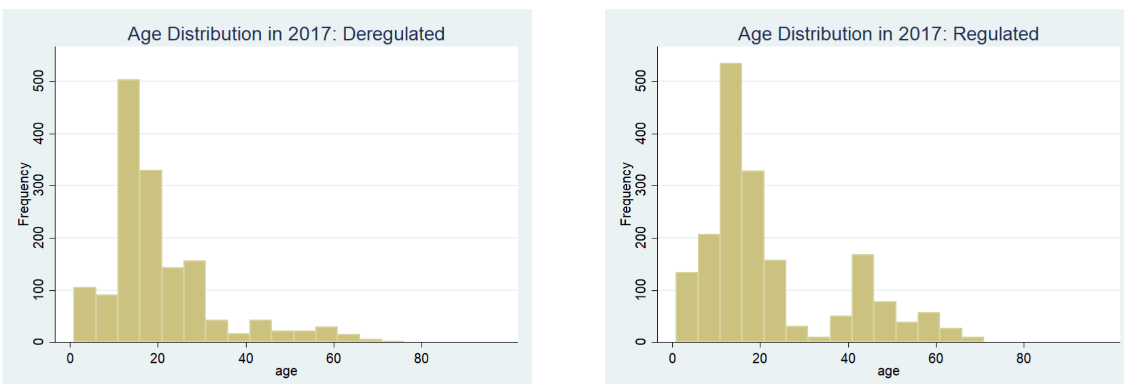

(a) (b)

Figure 2: Generator age distribution in 2017

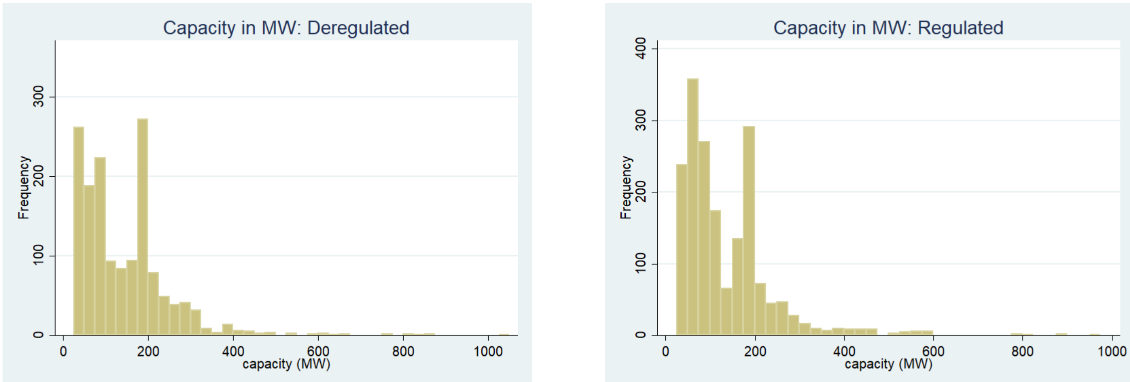

(a) (b)

Figure 3: Generator size distribution in 2017

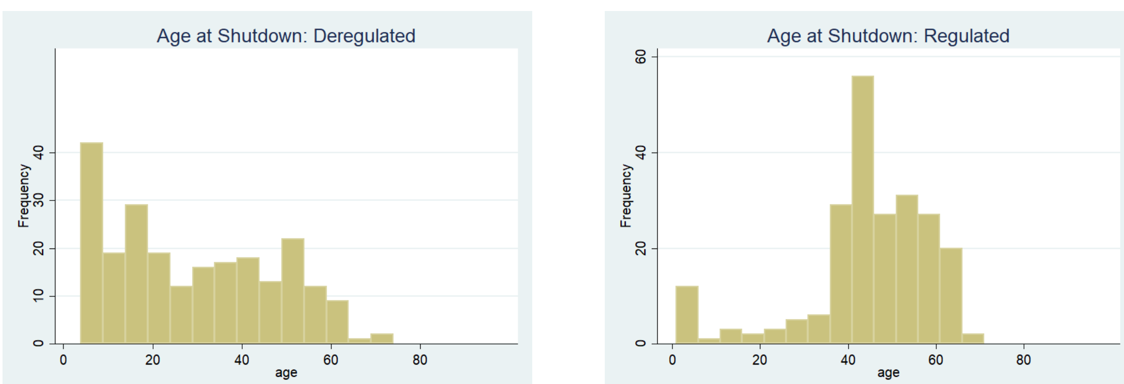

(a) (b)

Figure 4: Generator age at shutdown

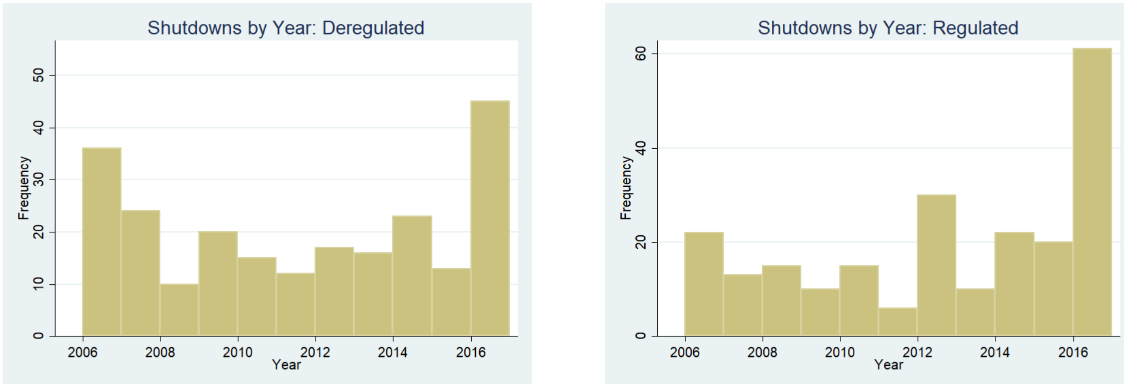

(a) (b)

Figure 5: Year of first shutdown

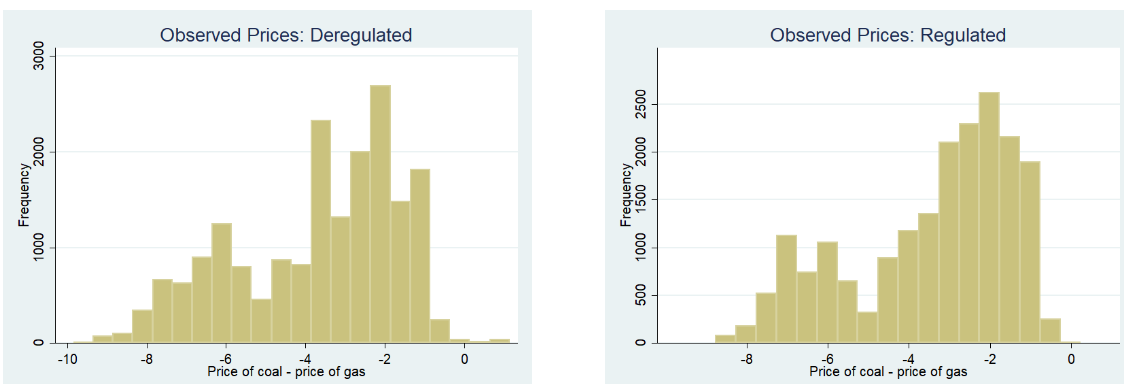

(a) (b)

Figure 6: Observed prices over sample period

Table 3: Effect of prices on shutdowns

	(1)	(2)	(3)	(4)	(5)	(6)
P^{Δ}	-0.000187	-0.00142**	-0.00283*	-0.00315**	-0.00375**	-0.00173
	(-0.26)	(-2.03)	(-1.94)	(-2.15)	(-2.49)	(-1.21)
P^{Δ}*reg	0.000622	0.00157	0.00158	0.00178	0.00203*	
	(0.52)	(1.34)	(1.37)	(1.51)	(1.88)	
reg	0.00149	0.0000991	-0.000915	-0.00163	-0.0194**	-0.0194**
	(0.30)	(0.02)	(-0.21)	(-0.37)	(-2.16)	(-2.16)
age		0.00104***	0.00105***	0.000719***	0.000859***	0.000859***
		(10.21)	(10.41)	(5.06)	(3.73)	(3.73)
capacity (MW)		-0.0000195*	-0.0000189*	-0.0000359***	-0.0000154	-0.0000154
		(-1.95)	(-1.89)	(-3.15)	(-0.97)	(-0.97)
Q_d		1.14e-11**	1.13e-11	8.10e-12	-1.36e-11	-1.36e-11
		(2.15)	(1.24)	(0.91)	(-1.27)	(-1.27)
age_reg					-0.000214	-0.000214
					(-0.73)	(-0.73)
cap_reg					-0.0000463**	-0.0000463**
					(-2.11)	(-2.11)
Q_d*reg					3.90e-11***	3.90e-11***
					(3.00)	(3.00)
P^{Δ}*dereg						-0.00203*
						(-1.88)
Observations	38297	38297	38297	38048	38048	38048
Tech indicators	-	-	-	X	X	X
Year, Region FEs	-	-	X	X	X	X

t statistics in parentheses

standard errors clustered at the plant level

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 4: Effect of predicted capacity factor on shutdowns

	(1)	(2)	(3)	(4)	(5)	(6)
CF	0.769	-0.783	-0.698	-0.179	-0.125*	0.0264
	(0.13)	(-0.43)	(-0.40)	(-1.53)	(-1.78)	(0.76)
CF*reg	-0.706	0.874	0.674	0.235*	0.151**	
	(-0.12)	(0.46)	(0.47)	(1.77)	(2.14)	
reg	0.265	-0.299	-0.220	-0.0695*	-0.0343*	-0.0343*
	(0.13)	(-0.45)	(-0.46)	(-1.80)	(-1.85)	(-1.85)
age		-0.000475	-0.00122	0.000474**	0.000141	0.000141
		(-0.22)	(-0.26)	(2.37)	(0.55)	(0.55)
capacity (MW)		0.0000289	-0.0000113	-0.0000282*	-0.00000697	-0.00000697
		(0.22)	(-0.12)	(-1.76)	(-0.52)	(-0.52)
Q_d		-5.67e-11	-1.93e-10	-4.81e-11	-2.35e-11*	-2.35e-11*
		(-0.39)	(-0.44)	(-1.62)	(-1.66)	(-1.66)
age_reg					0.000327	0.000327
					(0.91)	(0.91)
cap_reg					-0.0000179	-0.0000179
					(-0.89)	(-0.89)
Q_d*reg					1.33e-11	1.33e-11
					(1.08)	(1.08)
CF*dereg						-0.151**
						(-2.14)
Observations	32315	32315	32315	32126	32126	32126
Tech indicators	-	-	-	X	X	X
Year, Region FEs	-	-	X	X	X	X

t statistics in parentheses

standard errors clustered at the plant level

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 5: Prices → capacity factor

	(1) CF	(2) CF	(3) CF	(4) CF
P^{Δ}	0.0111***	0.00972***	0.0288***	0.0296***
	(16.23)	(9.30)	(14.44)	(14.08)
reg	-0.0324***	-0.102***	-0.00824***	-0.108***
	(-10.93)	(-11.52)	(-2.77)	(-12.31)
age	-0.00102***	-0.000637**	-0.00130***	-0.00200***
	(-6.97)	(-2.34)	(-8.89)	(-7.38)
capacity (MW)	0.000120***	-0.0000358*	0.000101***	-0.0000767***
	(8.74)	(-1.84)	(7.25)	(-3.83)
P^{Δ}*reg		0.00108		-0.00205
		(0.79)		(-1.53)
age_reg		-0.0000915		0.00153***
		(-0.28)		(4.74)
cap_reg		0.000374***		0.000403***
		(12.63)		(13.75)
Observations	32126	32126	32126	32126
Tech indicators	X	X	X	X
Year, Region FEs	-	-	X	X

robust t statistics in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 6: Linear probability model with high-coal region indicators

	(1)	(2)	(3)	(4)
P^{Δ}	-0.00375**	-0.00287**	-0.00245*	-0.00317*
	(-2.49)	(-1.97)	(-1.83)	(-1.76)
P^{Δ}*reg	0.00203*	0.00320**	0.00222*	0.000644
	(1.88)	(2.06)	(1.75)	(0.34)
reg	-0.0194**	-0.0158*	-0.00778	-0.00656
	(-2.16)	(-1.74)	(-0.77)	(-0.45)
P^{Δ}*hicoal		-0.00249	-0.00355	-0.00409
		(-1.33)	(-1.49)	(-1.60)
hicoal*reg		-0.00334	0.00825	0.0159
		(-0.35)	(0.87)	(1.56)
P^{Δ}*hicoal*reg		-0.00136	0.00122	0.00283
		(-0.52)	(0.42)	(0.86)
hicoal		-0.0191***	-0.0158*	-0.0160*
		(-2.69)	(-1.87)	(-1.69)
age	0.000859***	0.000838***	0.000814***	0.000966***
	(3.73)	(3.64)	(3.42)	(3.31)
capacity (MW)	-0.0000154	-0.0000156	-0.0000111	-0.00000634
	(-0.97)	(-0.98)	(-0.66)	(-0.31)
Q_d	-1.36e-11	-1.30e-11	-5.84e-12	4.96e-12
	(-1.27)	(-1.19)	(-0.50)	(0.20)
Observations	38048	38048	30362	24146
Tech indicators	X	X	X	X
Year, Region FEs	X	X	X	X
Sample	full	t25/b25	t20/b20	t15/b15

t statistics in parentheses

standard errors clustered at the plant level

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 7: Price regressions- Clustered SEs

	(1)	(2)	(3)	(4)	(5)
P^{Δ}	-0.00375***	-0.00375**	-0.00375***	-0.00375**	-0.00375**
	(-4.83)	(-2.49)	(-4.73)	(-2.56)	(-2.13)
P^{Δ}*reg	0.00203***	0.00203*	0.00203***	0.00203	0.00203**
	(3.53)	(1.88)	(3.51)	(1.62)	(2.14)
reg	-0.0194***	-0.0194**	-0.0194***	-0.0194	-0.0194*
	(-3.61)	(-2.16)	(-3.63)	(-1.62)	(-1.95)
age	0.000859***	0.000859***	0.000859***	0.000859***	0.000859***
	(5.91)	(3.73)	(6.16)	(3.55)	(4.52)
capacity (MW)	-0.0000154	-0.0000154	-0.0000154	-0.0000154	-0.0000154
	(-1.21)	(-0.97)	(-1.25)	(-0.94)	(-0.92)
Observations	38048	38048	38048	38048	38048
Clusters	robust	plant	generator	state	region-year

t statistics in parentheses

All specifications include tech indicators, year and region FEs, and controls interacted with *reg*

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 8: Capacity factor regressions- Clustered SEs

	(1)	(2)	(3)	(4)	(5)
CF	-0.125***	-0.125*	-0.125***	-0.125	-0.125*
	(-4.60)	(-1.78)	(-3.96)	(-1.64)	(-1.83)
CF*reg	0.151***	0.151**	0.151***	0.151*	0.151*
	(4.46)	(2.14)	(4.08)	(1.81)	(1.72)
age	0.000141	0.000141	0.000141	0.000141	0.000141
	(1.50)	(0.55)	(1.10)	(0.56)	(0.86)
capacity (MW)	-0.00000697	-0.00000697	-0.00000697	-0.00000697	-0.00000697
	(-0.92)	(-0.52)	(-0.72)	(-0.49)	(-0.70)
reg	-0.0343***	-0.0343*	-0.0343***	-0.0343*	-0.0343*
	(-5.10)	(-1.85)	(-4.58)	(-1.79)	(-1.73)
Observations	32126	32126	32126	32126	32126
Clusters	robust	plant	generator	state	region-year

t statistics in parentheses

All specifications include tech indicators, year and region FEs, and controls interacted with *reg*

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 9: Alternative specifications

	(1)	(2)	(3)	(4)
P^{Δ}	-0.00375**	-0.00160**	-0.00359**	-0.00365**
	(-2.49)	(-2.31)	(-2.44)	(-2.36)
P^{Δ}*reg	0.00203*	0.00212*	0.00192*	0.00173
	(1.88)	(1.94)	(1.81)	(1.57)
reg	-0.0194**	-0.0180**	-0.0114	-0.0194*
	(-2.16)	(-2.07)	(-1.14)	(-1.86)
age	0.000859***	0.000861***	0.000230	0.00121***
	(3.73)	(3.72)	(0.54)	(3.14)
capacity (MW)	-0.0000154	-0.0000153	-0.0000163	-0.00000439
	(-0.97)	(-0.96)	(-1.03)	(-0.27)
Q_d	-1.36e-11	-1.22e-11	-1.61e-11	-1.43e-11
	(-1.27)	(-1.15)	(-1.51)	(-1.34)
age*reg	-0.000214	-0.000205	-0.00152**	-0.000411
	(-0.73)	(-0.70)	(-2.52)	(-0.85)
cap*reg	-0.0000463**	-0.0000461**	-0.0000288	-0.0000390*
	(-2.11)	(-2.09)	(-1.32)	(-1.70)
Q_d*reg	3.90e-11***	3.71e-11***	4.00e-11***	4.04e-11***
	(3.00)	(2.99)	(3.06)	(3.13)
age^2			0.0000133	
			(1.33)	
age^2*reg			0.0000258*	
			(1.85)	
Observations	38048	38048	38048	38048
Tech indicators	X	X	X	*age
Region FEs	X	X	X	X
Year FEs	X	-	X	X

t statistics in parentheses

standard errors clustered at the plant level

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 10: Counterfactual estimates in MW of capacity

Year	Δcap	reg cap	%Δreg cap
2006	2106	156486	0.0135
2007	1997	162234	0.0123
2008	2508	173727	0.0144
2009	1511	179839	0.00840
2010	1398	187476	0.00750
2011	1262	192568	0.00660
2012	951	210794	0.00450
2013	919	217394	0.00420
2014	1287	220600	0.00580
2015	820	234031	0.00350
2016	706	237118	0.00300
2017	1007	248200	0.00410