

NBER WORKING PAPER SERIES

SPILOVERS FROM ANCILLARY SERVICES TO WHOLESALE POWER MARKETS:
IMPLICATIONS FOR CLIMATE POLICY

Jesse F. Buchsbaum
Catherine Hausman
Johanna L. Mathieu
Jing Peng

Working Paper 28027
<http://www.nber.org/papers/w28027>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2020, Revised July 2021

We are grateful to Jim Archsmith, Severin Borenstein, Karl Dunkle Werner, Ken Gillingham, Akshaya Jha, Jenya Kahn-Lang, Justin Kirkpatrick, Gordon Leslie, Matt Tarduno, and various conference and seminar attendees for helpful comments. This research was funded by the Alfred P. Sloan Foundation. The authors do not have any other financial relationships that relate to this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Jesse F. Buchsbaum, Catherine Hausman, Johanna L. Mathieu, and Jing Peng. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Spillovers from Ancillary Services to Wholesale Power Markets: Implications for Climate Policy

Jesse F. Buchsbaum, Catherine Hausman, Johanna L. Mathieu, and Jing Peng

NBER Working Paper No. 28027

October 2020, Revised July 2021

JEL No. L21,L94,L98,Q42,Q48,Q53,Q58

ABSTRACT

In electricity markets, generators are rewarded both for providing energy and for enabling grid reliability. The two functions are compensated in separate markets: energy markets and ancillary services markets. We provide evidence of changes in the generation mix in the energy market that are driven by exogenous changes in an ancillary services market. We provide quasi-experimental evidence and a theoretical frame-work for understanding the mechanism, showing that it results from the multi-product nature of conventional power plants combined with discontinuities in costs. While research in economics typically focuses solely on the energy market, our results suggest that spillovers between markets are important as well. Furthermore, policy changes relating to grid operations, grid reliability, or climate change could have unintended effects.

Jesse F. Buchsbaum
University of California at Berkeley
jesse_buchsbaum@berkeley.edu

Catherine Hausman
Gerald R. Ford School of Public Policy
University of Michigan
735 South State Street
Ann Arbor, MI 48109
and NBER
chausman@umich.edu

Johanna L. Mathieu
University of Michigan
4116 EECS
1301 Beal Avenue
jlmath@umich.edu

Jing Peng
University of Michigan
1301 Beal Avenue Ann
Arbor, MI 48109
jingpeng@umich.edu

1 Introduction

In electricity markets, generators are rewarded both for providing energy and for enabling grid reliability. The two functions are compensated in separate markets – energy provision is compensated in the energy market, while grid reliability is compensated in the ancillary service market. To date, the economics literature has largely focused on energy provision markets, even as other academic literatures, policymakers, and grid regulators have more carefully considered ancillary service markets. Ancillary service markets are interesting and important in their own right: they procure services that prevent brownouts and blackouts and ensure power quality. Moreover, changes in ancillary services markets can impact the behavior of generators in the much-larger energy market. While these market interactions have been extensively studied by engineers using optimization models and simulations, quasi-experimental evidence is largely lacking. In this paper, we show that exogenous policy changes implemented in the ancillary services portion of a large East Coast electricity market have changed the behavior of conventional (coal and natural gas) generators in the energy market, and as a result have had unintended environmental impacts. Changes to ancillary services markets are being considered as policymakers look for ways to lower the carbon emissions of electricity supply – as such, unintended CO₂ emissions changes in this context are particularly policy-relevant.

Specifically, we focus on the frequency regulation market in PJM. PJM is the largest wholesale electricity market in the US, serving major population centers on the East Coast and dispatching nearly one fifth of all generation capacity in the lower 48 states. Frequency regulation refers to the short-timescale balancing of supply and demand by grid operators; we describe it in depth below. We leverage policy-induced quasi-experimental variation in the amount of frequency regulation required by grid operators. As a result, we identify how changes in the provision of frequency regulation by conventional power plants impact the electricity market as a whole. We show that increases in the required frequency regulation capacity induce changes in the composition of power plants providing energy generation. For a 100 MW increase in the frequency regulation requirement, we estimate an additional 390 MWh from combined cycle plants over the course of one hour in the energy market and a corresponding decrease from boiler units.¹ These results are qualitatively similar across a broad suite of robustness checks, with estimated increases of combined cycle units of around

¹We use megawatts (MW) throughout to measure capacity, and megawatt-hours (MWhs) to measure generation. A 1 MW unit operating at its full capacity for one hour generates 1 MWh of electricity. The regulation requirement is a market-wide capacity or power requirement and is thus measured in MWs. Electricity generation is sold as a function of energy provided over the course of some time frame and is thus measured in MWhs.

300 to 490 MWh.

Because combined cycle plants primarily use natural gas whereas boiler units use coal, we also see a change in the fuel used to provide energy. Specifically, we find that increases in the regulation requirement lead to an increase in natural gas usage and a decrease in coal and oil. As a result, for every 100 MW of increased frequency regulation, CO₂ emissions fall by 250 metric tons per hour in our sample, with robustness checks showing a range of 180 to 350 tons. This represents a decrease of 0.5% of CO₂ emissions, or an annual total of 2.2 million tons of reduced CO₂. This is roughly equal to the total emissions of around one to four additional coal-fired units (at around one or two plants) in PJM. Valued at the IWG Social Cost of Carbon, the short-term generation changes from an increase in the regulation requirement would reduce climate damages by nearly \$100M per year.² Recent peer-reviewed estimates would place the value even higher (Pindyck, 2017; Moore et al., 2017; Ricke et al., 2018; Bastien-Olvera and Moore, 2020). For the time period we study, the PJM market implemented several changes to *decrease* the regulation requirement. While doing so can reduce the system-wide private cost of electricity provision, in this case it inadvertently led in the short term to higher CO₂ emissions and increased climate damages.

At first glance, the magnitude of the generation mix change is surprising – it is a greater than one-to-one effect. We next provide a simple model to analyze the mechanisms behind this. Key features of the model include (1) conventional power plants are multi-product suppliers; (2) conventional power plants tend to operate within a somewhat narrow operational range defined by non-zero minimum and maximum constraints; (3) policy and market changes can cause conventional power plants to switch from fully off to operating at minimum load. Importantly, the minimum load at many generators (a physical constraint) is non-negligible – while the related literature frequently ignores this constraint, it may be as high as 50 percent or more of capacity. To summarize, changes along the extensive margin – which units are turned on – can lead ancillary services markets to have outsized impacts on generation markets. We discuss qualitatively whether the changes in the generation mix that we observe would be expected to occur in other markets or other time periods, arguing that the mechanisms we study will generally be relevant, although the magnitudes (and the direction of CO₂ changes) are context-specific.

We then provide empirical evidence of extensive margin changes. We show that increasing the regulation requirement causes boilers to be dispatched at lower *levels* of generation in our sample (i.e., a change along the intensive margin). However, combined cycle plants are dispatched more *frequently* (i.e., a change along the extensive margin). These effects are

²We use the 2015 SCC from the Interagency Working Group on Social Cost of Greenhouse Gases (2016) and convert to 2019 dollars using the CPI, implying a value of \$44 per metric ton.

consistent both with the overall generation impacts and with the stylized model.

These results make several contributions to the energy and environmental economics and the industrial organization literatures. First, we contribute to a still-small empirical literature in energy economics that analyzes electricity markets other than the energy market. One strand of the economics and engineering literature explores ancillary services markets with optimization models or simulation approaches (Hirst and Kirby, 1998; Just and Weber, 2008; Yu and Foggo, 2017). However, empirical analysis of ancillary services markets, particularly in how they interact with energy markets, is limited,³ despite there being a large empirical literature on wholesale electricity markets.⁴ We show that ignoring the ways that energy markets interact with these other markets can lead to incorrect conclusions about the impacts of policy changes.

Second, we contribute to a strand of the electricity literature that emphasizes the importance of understanding how technical constraints impact power plant behavior, especially when considering generators as multi-product firms. Perhaps most closely related is Mansur (2008), which shows that ignoring intertemporal constraints gives an inflated estimate of the welfare impact that restructuring electricity markets had in the late 1990s. Previous empirical papers had focused on merit-order dispatch without considering intertemporal constraints such as minimum load, startup costs, and ramping. Also related is Reguant (2014), which estimates markups using firms' bidding behavior in the presence of these intertemporal constraints. Similarly to these papers, we show that minimum load constraints can have a large impact on plant behavior. We additionally show that the existence of multiple markets can have significant impacts on plant behavior. We argue that power plants should be considered multi-product firms, and that limiting attention to just one of the markets might lead to incorrect or incomplete conclusions about plant behaviors.

The behavior of multi-product firms has been the focus of a growing body of work in industrial organization and in international trade. Researchers have shown that examining how firms optimize *across* different products plays a key role in understanding productivity differences across firms; the behavior of firms with market power; and the impacts of trade policy, exchange rate movements, demand shocks, and more (Johnson and Myatt, 2006;

³The primary exceptions are: Doraszelski, Lewis and Pakes (2018), which estimates models of firm learning and convergence to equilibrium in a newly deregulated frequency response market in the UK; Knittel and Metaxoglou (2008), which examines ancillary services in the context of the California electricity crisis; and Schwenen (2015), which examines New York's capacity market. Jha and Wolak (2020) examines the impact of "explicit virtual bidding" on the cost of electricity provision, focusing on fuel costs but also incorporating the costs of ancillary services provision.

⁴See Borenstein, Bushnell and Wolak (2002); Fabrizio, Rose and Wolfram (2007); Reguant (2014); Borenstein and Bushnell (2015); Cicala (2015); Cullen and Mansur (2017); Davis and Hausman (2016); Holland et al. (2016a); Hortacsu et al. (2019), among many others.

Eckel and Neary, 2010; De Loecker, 2011; Chatterjee, Dix-Carneiro and Vichyanond, 2013). Nonetheless, studies of the electricity market have tended to treat firms as providers of a single good – electricity generation – rather than multiproduct firms, with the exception of an older literature on the optimal regulation of natural monopolies that provide multiple goods (for instance, Mayo, 1984). Our ability to model the production of electricity with an engineering-based model makes clear how the degree of complementarity versus substitutability in production matters for outcomes in multi-product firms. Given the large number of industries with multi-product features (refining, airlines, freight transportation, manufacturing, to name a few), these mechanisms are of widespread relevance.

Third, this paper contributes to policy discussions about several ongoing developments in electricity markets: changes in the way frequency regulation is procured and compensated, the introduction of utility-scale batteries, and the increasing deployment of renewable electricity (MIT Energy Initiative, 2011; Department of Energy, 2013, 2016; Hledik et al., 2017). Across the country, frequency regulation markets have seen multiple changes in recent years. In 2011, the Federal Energy Regulatory Commission (FERC) issued Order 755 (Federal Energy Regulatory Commission, 2011), which required grid operators to change their frequency regulation compensation mechanisms; we give details below. Various electricity markets across the US have responded with differing changes to their frequency regulation markets (Department of Energy, 2013; Tabari and Shaffer, 2020). These compensation mechanisms favor some resource types more than others, so supply in the regulation market (and therefore dispatch in the energy market) is likely to be affected. The extent to which these changes in compensation mechanisms have impacted electricity markets has not been thoroughly analyzed in the energy economics literature.

Moreover, there has been a growing interest in energy storage devices, such as utility-scale batteries, to provide frequency regulation and other grid support services. Within the PJM market we study, batteries have largely been deployed for frequency regulation, rather than for the intertemporal arbitrage⁵ potential explored in energy economics papers (Carson and Novan, 2013; Antweiler, 2021; Holladay and LaRiviere, 2018; Kirkpatrick, 2018; Ambec and Crampes, 2019; Linn and Shih, 2019; Castro, 2020).⁶ Worldwide, a primary use of battery storage is for frequency regulation (International Renewable Energy Agency, 2017; Deloitte, 2018). The behavior of batteries providing ancillary services may differ tremendously from that of batteries providing arbitrage – both because of the way the relevant markets are

⁵For batteries, arbitrage involves charging from the grid when the price is low and selling electricity to the grid when the price is high. As we discuss later, this is not the primary use of most grid-connected batteries in PJM.

⁶There is also a related literature on the use of hydroelectric facilities for storage, again focusing on storage for arbitrage purposes; see e.g. Liski and Vehvilainen (2020).

designed and because of the timescale of use (e.g., seconds versus hours).

Our empirical results on the impact of changes to the regulation requirement provide a useful analogy for the entry of batteries in the regulation market, since their entry will lower the residual demand for regulation from conventional generators. Many proponents argue that the entry of batteries for grid reliability will lead to lower emissions (Federal Energy Regulatory Commission, 2011). In contrast, engineering models and simulations show the possibility of increased emissions when batteries provide frequency regulation (Ryan et al., 2018), for which our analysis provides the first empirical corroboration.⁷ However we caution that the magnitude of the fuel mix changes (and therefore a CO₂ emissions change) is likely to change as secular coal retirements continue.

Finally, the rise of intermittent renewables such as wind and solar generation can increase the amount of frequency regulation required in electricity markets (Kirby, 2004; MIT Energy Initiative, 2011).⁸ As generating technologies continue to evolve and batteries and renewable resources play a larger role in both markets, the interactions between energy and ancillary services markets are likely to continue to be important.

2 Background

2.1 Ancillary Services and Frequency Regulation

Electricity markets are actually made up of numerous separate markets: energy, capacity, and several types of ancillary services markets.⁹ The energy market is the most studied and best understood by economists – this is where firms are compensated for generating electricity to be used by residential, commercial, and industrial customers. In ancillary service markets, generators are compensated for providing services that enable grid reliability – for example, frequency regulation and other types of reserves. Because the same generators participate in both the energy market and ancillary services market, the structure of ancillary service markets may have important spillovers in the energy provision market. Despite this, there

⁷Our paper thus also relates to the literature on spillover effects from other non-fossil technologies, such as from solar to fossil power plants (Bushnell and Novan, 2021) and from hydro to fossil power plants Archsmith (2018).

⁸Ovaere and Gillingham (2019) examines the empirical impact of renewables on the cost of ancillary services provision.

⁹By energy market, we refer to the sale of electricity, in MWh, in a wholesale market. In PJM, this is called simply the “energy market.” In this paper, we sometimes refer to “energy provision” to distinguish it from energy markets more generally such as natural gas and oil markets. Since the electricity economics literature frequently does not cover ancillary services markets, it frequently uses terms like “electricity market,” “power market,” and “wholesale market” without specifying whether the market is for providing energy or capacity or ancillary services or some combination thereof.

has been very little empirical research into ancillary service markets. In this paper, we focus on frequency regulation, motivated by ongoing policy changes in this market.

Electricity markets are unique in that demand must constantly equal supply, a responsibility that falls on grid operators. However, there are frequent fluctuations in demand and supply, creating small mismatches between the two. When supply exceeds demand, frequency (the number of cycles per second of the alternating current) rises above the nominal frequency (i.e., 60 Hz in North America); when demand exceeds supply, frequency falls below the nominal frequency. If the grid frequency departs enough from the nominal level, it can cause damaged equipment or brownouts and blackouts for customers (Federal Energy Regulatory Commission, 2011).

To prevent this from happening, system operators have created markets to regulate the frequency of the grid, a service called “frequency regulation,” “regulation reserves,” “load frequency control,” “secondary frequency control,” or simply “regulation.”¹⁰ In the typical market, the system operator sets a frequency regulation requirement – this is the total capacity (in MW) that must be set aside for provision of frequency regulation. It is sometimes time-invariant, sometimes a function of forecasted demand, or possibly also a function of renewables forecasts. Generators can then bid a portion of their capacity to be available to grid operators to either increase or decrease generation (relative to their set point) at any time (within the frequency regulation contract duration), depending on the needs of the grid. The independent system operator¹¹ sends out a signal (automatic generation control, or AGC) to participating units, to which they automatically make small adjustments in their generation to balance supply and demand.¹² Typically, these small adjustments are made within seconds (Zhou, Levin and Conzelmann, 2016). An example signal is shown in Figure 1.

Some system operators use separate signals for “regulation up” versus “regulation down.” The PJM market that we study does not use separate signals for up and down movements. For the time period we analyze, this signal is energy-neutral within a short time-frame (15 minutes), so that units always return to their initial set point (Monitoring Analytics, 2018).

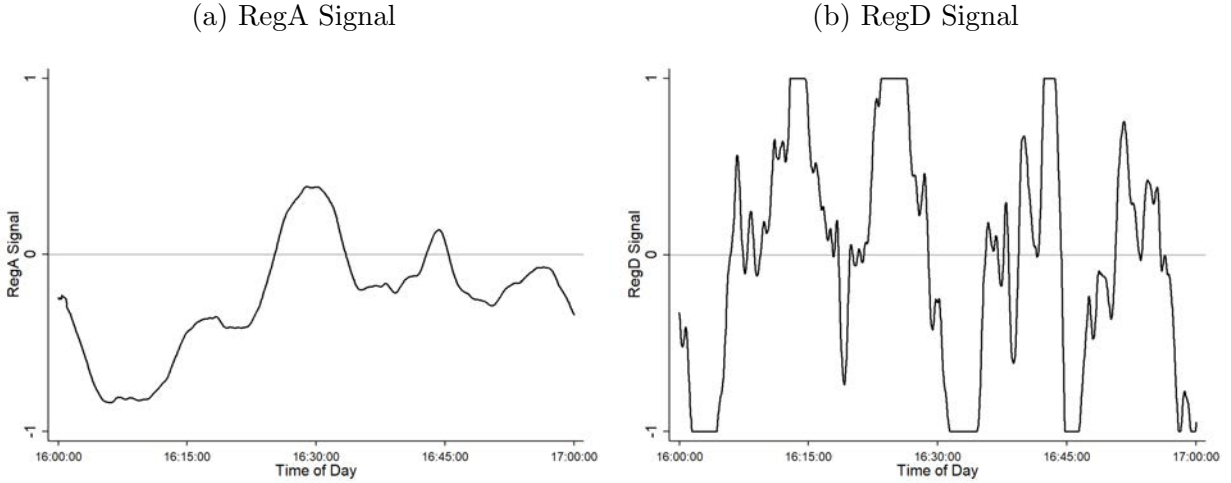
To participate in the market, a power plant must have the technical capability to follow the operator’s regulation signal. It must also be dispatched at a non-zero level of generation, with headroom and footroom to follow the signal. That is, it cannot be operating at its

¹⁰Background on frequency regulation and other ancillary services is provided in Hirst and Kirby (1997); Kirby (2004); Hummon et al. (2013); Tacka (2016); Zhou, Levin and Conzelmann (2016).

¹¹An independent system operator is non-profit entity that operates the wholesale electricity market.

¹²Specifically, FERC Order 755 defines frequency regulation as “the capability to inject or withdraw real power by resources capable of responding appropriately to a system operator’s automatic generation control signal in order to correct for actual or expected Area Control Error needs” ((Federal Energy Regulatory Commission, 2011), p 67266).

Figure 1: Example Regulation Signals



Note: This figure shows the RegA and RegD signals in PJM from 4 pm to 5 pm on July 19, 2019. Data are from PJM.

minimum or maximum constraints because it must be able to move up and down in response to the operator’s signal (Kirby, 2004). The resource mix contributing to frequency regulation varies across regions. In PJM, it is a mix of coal, natural gas, hydro, and battery storage (Monitoring Analytics, 2018).

Generators take several cost considerations into account when deciding whether and how to bid in a regulation market. Small fluctuations around the generator’s set point impact the plant’s heat rate, thus changing fuel costs. Providing regulation also imposes wear and tear on the plant (Hirst and Kirby, 1997; Hummon et al., 2013) and can change SO_2 and NO_x emissions, which in some markets are priced.¹³

Over time, system operators have moved towards rewarding regulation providers for both the *capacity* committed to regulation and the *quality* of regulation services provided.¹⁴ Specifically, some units, such as coal-fired boilers, have significant physical inertia that prevents them from responding quickly to regulation signals. In contrast, units such as batteries, hydro generators, and some natural gas generators are able to respond very quickly to the signal, which gives the system operator greater flexibility and speed in restoring the system-wide frequency to its desired level. To incorporate this difference across suppliers,

¹³The exact cost associated with additional wear and tear is not generally known for individual generators; simulations typically assume a cost that varies across fuel types (see, e.g., Hummon et al. (2013)).

¹⁴This has been spurred by FERC Order 755 (FERC 2011), which requires that independent system operators change their frequency regulation compensation mechanisms to “pay for performance” systems that recognize the differential speed and accuracy with which different resources respond to the regulation signal. Each independent system operator has designed its pay for performance compensation mechanisms differently (Department of Energy, 2013; Tabari and Shaffer, 2020).

PJM uses a more complicated set of payments.¹⁵ First, PJM sends out two separate regulation signals, one termed “RegA” (for slower-responding units) and one termed “RegD” (for faster-responding units). Moreover, units receive payments both for their capability (i.e., the quantity of MWs offered) and for their performance (the accuracy with which the unit responds to the operator’s signal). In Section 5, we discuss how the compensation of both capability and performance might impact our results. We explore the entry of utility-scale batteries for frequency regulation in Section 7.

2.2 Related Literature

One strand of the economics and engineering literature explores ancillary services markets with optimization models and simulations (Hirst and Kirby, 1998; Just and Weber, 2008; Yu and Foggo, 2017). In particular, Hirst and Kirby (1997) notes that the minimum and maximum constraints of individual generators, when combined with the need to provide headroom and footroom for regulation provision, can lead to a dispatch of units across the system that would not appear least-cost if one only considered the marginal cost of energy provision, and it can also lead to significant complexity in which units are dispatched. Simulations showing changes in dispatch to meet regulation requirements are also given in Hummon et al. (2013). This informs our stylized model in Section 5.

A small number of ex-post empirical papers examine how additional operational constraints can lead to out of merit dispatch (Mansur, 2008; Reguant, 2014); these papers focus on dynamic constraints related to startup costs and ramping. However, empirical analysis of ancillary services markets, particularly in terms of how they interact with energy markets, is limited, despite there being a large literature on wholesale electricity markets (Reguant (2014); Borensten and Bushnell (2015); Hortacsu et al. (2019) to name a few prominent recent papers).

Several papers use optimization models and/or small-scale simulations to show how frequency regulation markets and other ancillary services markets interact with energy markets. Notably, one report (Atanacio et al., 2012) simulates the impact of storage providing regulation in the PJM system.¹⁶ Specifically, it looks at the emissions changes expected to result across the PJM system when fast-acting storage devices provide 10 percent, 25 percent, or 50 percent of frequency regulation services. It is not clear if the study’s proprietary simulation model includes the unit commitment and/or economic dispatch algorithms used by PJM. So, it is not clear if storage providing frequency regulation can impact commitment in the model,

¹⁵This payment mechanism in PJM was established in October 2012, following FERC Order 755.

¹⁶Papers from the empirical economics literature exploring other aspects of the PJM market include Mansur (2008); Mansur and White (2012); Abito et al. (2018).

nor how exactly it impacts dispatch. The simulation results show small emissions reductions from storage entry, because conventional plants operate less efficiently when providing regulation services. However, the report noted that the interaction between the energy market and the frequency regulation market are complicated and could limit the emissions benefits of storage, and the study’s analysis of California’s system showed the potential for increased CO₂ emissions (Atanacio et al., 2012).

Two additional papers that examine emissions impacts of using storage for frequency regulation are Lin, Johnson and Mathieu (2019) and Ryan et al. (2018). Ryan et al. (2018) uses a small test system to show that changes in the frequency regulation market can lead to changes in the fuel mix and therefore emissions changes. Indeed, the authors find that “[c]hanges in generator commitment and dispatch caused by the addition of energy storage were the most significant contributors to the energy storage system’s environmental impact” (p 10172). Hummon et al. (2013) also explores the interaction between reserve markets and energy markets using a simulation approach, although the paper does not examine emissions outcomes. Cho and Kleit (2015) explores the optimal bidding strategy for a storage device that can provide ancillary services. And finally, in Yu and Foggo (2017), “[s]imulation results with a realistic battery storage system reveal that the majority of the market revenues comes from frequency regulation services” (p 177).

While several of these papers examine ancillary service markets, and even consider spillovers into the energy provision market, this paper is the first to our knowledge to show quasi-experimental evidence of these spillover effects, which lead to fuel use changes and therefore changes in emissions.

3 Data

We collect data from the Environmental Protection Agency’s Continuous Emissions Monitoring System (CEMS) on hourly generation (MWh) and CO₂ emissions (metric tons)¹⁷ at the generator level, for fossil-fuel-fired units.¹⁸ From CEMS, we also observe the primary

¹⁷The CEMS-reported CO₂ emissions are missing for approximately 9% of observations with non-zero heat input data, representing 2% of generation. In place of these missing values, we assume an emissions rate (per mmBtu of fuel used) equal to the median rate at the unit, typically around 0.093 metric tons per mmBtu for coal-fired units and 0.054 for natural gas fired units. In the Appendix, we show results using CEMS-reported CO₂ emissions.

¹⁸CEMS reports *gross* generation rather than *net*, i.e. not accounting for the generation used by the plant itself (or instance, to run pollution control equipment). Net generation is the variable of interest, since that is what is sold in the electricity market. Following Cicala (2017), we scale each unit’s generation down from gross to net using monthly generation data from the Energy Information Administration’s (EIA) form 923 dataset. This approach also solves a problem we see with some combined cycle units: they do not report the full value of their electrical output in the CEMS data. Finally, some units report only steam load in CEMS,

fuel source used by a generator each year (coal, natural gas, or oil) and the technology type of the generator (boiler, combined cycle, or combustion turbine). From EIA-860 data, we observe whether units are located in PJM. We also observe in EIA-860 whether units are operated by electric utilities, independent power producers, or as part of a commercial or industrial operation. We drop all commercial and industrial units, as they are unlikely to sell into the electricity market. We report our primary results for four technology categories of PJM units: boilers, combined cycle (CC) units, combustion turbine (CT) units, and all other PJM CEMS units aggregated. We also report results across four fuel types: coal, natural gas, oil, and others aggregated.¹⁹

The CEMS data do not provide information on non-fossil units (e.g., nuclear, hydro, wind, municipal solid waste), nor does CEMS cover fossil-fuel fired units with capacity less than 25 MW. (For context, units smaller than 25 MW are quite small; the average capacity in EIA-860 data is over 350 MW for steam units in PJM.) To observe the behavior of these units, we calculate a residual category of generation (in MWh), equal to the difference between total demand reported by PJM and total generation reported in CEMS. This residual variable thus covers PJM units not in CEMS as well as net imports into PJM from other regions.

We collect hourly data, from PJM, on regulation market activity. Our primary explanatory variable is the hourly regulation requirement, in MW. The regulation requirement refers to a pre-determined quantity of regulation capacity that the independent system operator announces it will purchase. As we describe below, this amount varies over time because of several policy changes. We also observe the quantity of frequency regulation provided by slow-responding resources (RegA) versus fast-responding resources (RegD), two variables to which we return later in our analysis.

We collect several additional variables that serve as controls. From EIA, we observe the daily price of natural gas, measured at Henry Hub and the daily price of oil (West Texas Intermediate).²⁰ We observe the hourly requirement (in MW) for other types of ancillary services: synchronized and non-synchronized reserves. We observe total hourly electricity demand across PJM, in MWh. We observe the forecasted peak and valley demand for each day, also from PJM. We collect hourly data on wind generation (MWh, from PJM). Also, we observe daily average temperature at the Philadelphia airport (degrees Fahrenheit, from NOAA), a relatively central location within PJM.

but report non-zero net generation in EIA-923. We similarly scale from steam load to net generation for these units; they account for 4% of our final net generation variable. Details are in the Appendix.

¹⁹The category “other technology types” is composed entirely of a small number of units that changed their technology type over this time period. The category “other fuel types” includes, for instance, wood-fired units, as well as a small number of units that changed their fuel type over this time period. Details are in the Appendix.

²⁰Coal prices paid by power plants were essentially flat over this time period.

For placebo tests, we collect hourly generation and CO₂ emissions data at the generator level for two other types of units: fossil-fuel-fired units not in PJM but in surrounding states; and units in CEMS data that are not categorized by CEMS as electrical generating units (e.g., refineries).

Finally, we note that the last decade has seen a secular decline in coal generation and rise in natural gas, a product primarily of the fracking revolution (Linn and McCormack, 2019). As such, we collect coal plant retirement dates and capacities from the EIA-860 dataset. We also collect retirement dates for other plant types (for instance, natural gas and oil), although they are smaller in magnitude.

Summary statistics are provided in Appendix Tables A1, A2, and A4. A time-series of generation by fuel types and technology types is provided in Appendix Figure A2. The sample is characterized by a reliance on coal generation and on natural gas generation. Other fuel types are present (oil-fired generations; etc.), but are a very small portion of the energy market.

As expected, Appendix Figure A2 shows that the period 2012-2018 is characterized by a decline in coal generation and a rise in natural gas. By the end of the sample, the two are responsible for roughly equal quantities of generation. However, within the main sub-period which we will focus on (2012-2014), coal-fired generation holds roughly constant, so we anticipate that secular trends in the mix of gas versus coal will not be a major problem for identification.

4 Empirical Evidence on Generation and Emissions

4.1 Identifying Variation

We are interested in how exogenous changes in the frequency regulation market spill over into the behavior of generators in the energy market. As such, we leverage policy changes in PJM's regulation market. We begin by focusing on a short time window with limited battery capacity: October 1, 2012 to December 31, 2014. Over this period, we leverage quasi-experimental variation in the *total* regulation requirement set by the independent system operator, which changes several times. The benefit of focusing on this period is twofold: battery capacity for frequency regulation is limited and conventional generator entry and exit is small. This allows us to focus on the total regulation requirement and its impact on conventional generators. In Section 7, we explore variation in battery capacity.

As of October 1, 2012, the regulation requirement was set at 0.78% of forecasted peak load. On November 22, 2012, the requirement was reduced to 0.74% of forecasted peak

load, then to 0.70% on December 18, 2012. More important than these minor changes is the change on December 1, 2013: the regulation requirement was changed from a percentage of the forecasted peak load to a requirement of 700 MW of effective regulation during peak hours and 525 effective MW during off-peak hours.²¹ This variation in the regulation requirement can be seen in Figure 2. A PJM report from 2011 describes the motivation behind these changes: “[d]ecreasing regulation requirements reduces regulation payments” and “[f]ewer resources providing regulation means more resources available for the energy market.”²²

In Section 2, we discuss how the regulation *signal* at any point in time is a function of mismatch between supply and demand, and how the signal varies at a very fine scale (e.g., two seconds). Thus the signal itself is endogenous to market activity. However, the regulation *requirement* is a pre-determined capacity procurement – set in advance by policy – and thus is not endogenous to day-to-day or hour-to-hour market activity, once we control for the variables used in the pre-determined requirement-setting process (e.g., forecasted peak load).

Unfortunately, we do not observe individual participation in the regulation market, as this is considered sensitive market information and is not published by PJM. However, we can use CEMS data to observe how generation behavior in the energy market changes over this period as a function of the total regulation requirement.

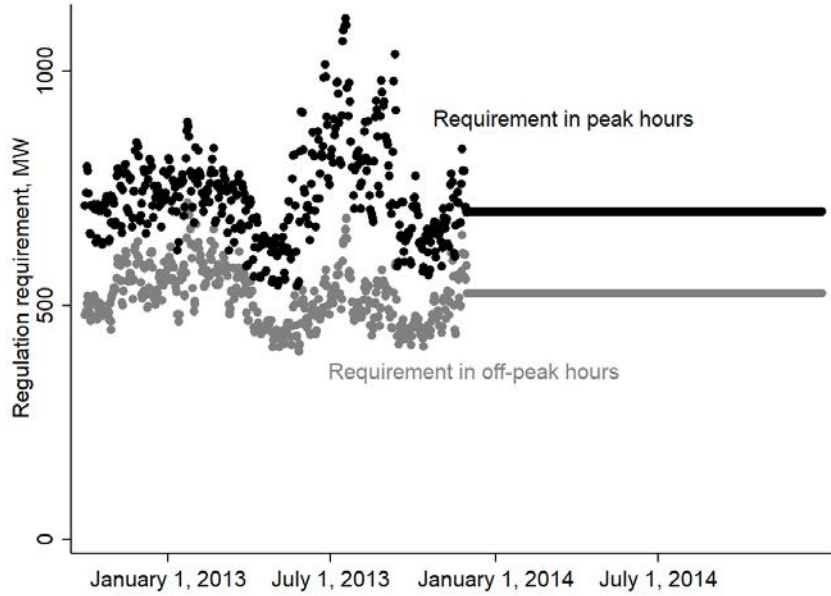
For expositional purposes, we focus on generation aggregated up to a prime mover type (e.g., boiler) or a fuel type (e.g., coal), but the regression analysis could also be done at the individual unit level. We treat the regulation requirement as random variation, but we control for observables that may be correlated with the regulation requirement. Most importantly, we control for forecasted peak load, as it directly determines the regulation requirement in the first part of the sample. The policy change thus allows us to separately identify the impact of forecasted peak load from the impact of the regulation requirement.

Essentially, the second half of our sample (when the regulation requirement is fixed) helps identify the effects of control variables that may be correlated with the regulation requirement in the first half of the sample. In the Appendix, we provide a simulation of this identification strategy, where we directly control the data-generating process. Finally, we note that also apparent in Figure 2 is the substantial variation in the regulation requirement over this time period, from less than 500 to around 1000 MW.

²¹Source: http://www.monitoringanalytics.com/reports/pjm_state_of_the_market/2013/2013-som-pjm-volume2.pdf. “Effective” regulation is a measurement that takes into account the performance of the units providing regulation and the substitutability across RegA and RegD units; see <https://pjm.com/~media/documents/manuals/m11.ashx>.

²²<https://www.pjm.com/~media/committees-groups/committees/mrc/20110915/20110915-item-13-rpstf-update-presentation.ashx>

Figure 2: Regulation Requirement in PJM



Note: The regulation requirement changes across hours within a day, hence the two different levels plotted on each day. Peak hours are defined 4 a.m. to midnight, and off-peak hours as midnight to 4 a.m. In the raw data, one hour (on 4/2/2013) is listed as having a regulation requirement of zero; this hour is dropped from the regressions. Data source is PJM.

The regression takes the form:

$$G_{i,t} = \alpha_i + \beta_i R_t + X_t \Theta_i + \varepsilon_{i,t}, \quad (1)$$

where $G_{i,t}$ is generation at unit type i in hour t , R is the regulation requirement, and X_t is a vector of controls. The parameters are allowed to vary by unit type. Standard errors are clustered by sample week.²³

Note this is estimated as a single time-series, but we index G with unit type i , since we can separately estimate the regression for multiple unit types. While we do not have panel variation in the variable of interest R , we do observe a policy change mid-way through our sample, clearly depicted in Figure 2. Moreover, the policy change does not appear correlated with a simple time trend; there are levels of the regulation requirement in the pre-period that lie both above and below the post-period levels. We include a dummy for the policy change in how the regulation requirement was set, which occurred on December 1, 2013 – but again, this is not expected to be critical for identification, since the policy change was

²³Clustering by week accounts for serial correlation in unobservables, such as an exogenous multi-day outage at a plant.

not simply a level shift in the regulation requirement, and since we are not aware of any other policy changes on that date.

We include a number of additional controls, both to minimize potential bias from correlated observables and to improve precision. Most importantly (see the simulations in the Appendix), we control for daily forecasted peak and valley levels.²⁴ In the first part of the sample, these directly influence the regulation requirement. It is possible that generators behave differently when the peak load is forecasted to be high rather than low, so we include these controls to minimize potential bias.

We also control for the amount of reserves required in the two other reserve markets in PJM: synchronized reserves and primary reserves. PJM sets requirements both for the territory as a whole and for the Mid-Atlantic Dominion area; we control for both sets of requirements. None of these variables were directly tied to policy changes on December 1, 2013, and they are not generally correlated with the frequency regulation requirement (the correlation coefficient between each of these variables and the regulation requirement is less than 0.1), so these controls are not expected to be important for identification.

We also control for the total generation by PJM units appearing in the CEMS data. This follows Davis and Hausman (2016) and helps to isolate compositional effects *within* conventional units. It is equivalent to controlling for total demand net of nuclear, solar, wind, biomass, and other renewable or non-CEMS generation. We also control for total demand in PJM. In combination with the total CEMS generation variable, this is akin to controlling for the aggregation of non-CEMS (nuclear, solar, wind, etc.) generation. Given that we are controlling for total CEMS generation, this is less critical for bias – it essentially treats nuclear and renewable generation as exogenous, and it may aid with precision.

Following the related empirical literature, we add controls for the Henry Hub natural gas price and West Texas Intermediate oil price. These are not expected to be directly correlated with the treatment variable of interest, but they may aid with precision and may also help with bias if they happen to be correlated with the regulation requirement. Similarly, we control for weather, measured as daily heating degree days and cooling degree days in Philadelphia, a central location in our geographic region.

We control for a quadratic time trend. Again, this is not expected to be crucial for identification since the identifying variation in Figure 2 is not following a simple trend. We include a vector of time effects: day of week, hour, and month effects to reduce noise from, e.g., seasonality.

²⁴Specifically, we use the day-ahead forecasted peak load in hours for which the peak regulation requirement applies (4 a.m. to midnight) and set the variable equal to zero in off-peak hours. Similarly, we use the day-ahead forecasted valley load in hours for which the off-peak regulation requirement applies (midnight to 4 a.m.) and set the variable equal to zero in peak hours.

Table 1: The Impact of the Regulation Requirement on the Energy Market

	Boiler (MWh)	CC (MWh)	CT (MWh)	Other tech. (MWh)
Regulation requirement, 100 MW	-363.8** (180.2)	390.5** (149.9)	-0.1 (184.7)	-26.7 (20.0)
Load controls	Yes	Yes	Yes	Yes
Other reserves controls	Yes	Yes	Yes	Yes
Retirement controls	Yes	Yes	Yes	Yes
Quadratic time trend	Yes	Yes	Yes	Yes
Fuel price controls	Yes	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes
Month, day of week, and hour effects	Yes	Yes	Yes	Yes
Observations	19,680	19,680	19,680	19,680
Within R ²	0.92	0.71	0.49	0.56
Mean of dep. var.	40,188	14,293	849	554

Note: This table shows estimates from four separate time-series regressions. In all columns, the dependent variable is total MWh of electricity generated per hour in the PJM market by units that appear in CEMS data (where each column aggregates across all units of a particular type). Coefficients on control variables are shown in the Appendix, Table A5. The unit of analysis is an hour. Standard errors are clustered by sample week.

Finally, we include controls for the total capacity retired, by fuel type, over this time period. As shown in Table A4, this period saw a substantial amount of generation capacity retired; primarily from coal-fired and oil-fired generators. Again, given that the identifying variation in the regulation requirement is policy-induced and not following any particular trend, these controls are not expected to matter for bias.

4.2 Regression Results

We first show regression results aggregated by technology type, in Table 1 (coefficients on control variables are shown in the Appendix, Table A5). We see that when more regulation is required, the unit types providing services in the energy market change, with decreased generation from boiler units and increased generation from combined cycle units. Specifically, for each 100 MW of additional regulation capacity required by the system operator, boiler units decrease their generation by 360 MWh and combined cycle units increase their generation by 390 MWh, both statistically significant at the five percent level. The difference between these two comes from a small number of units with other technology types (right-most column).

Because different technology types use different primary fuel sources, the results in Table 1 will have implications for fuel use and therefore environmental impacts. We next show

Table 2: The Impact of the Regulation Requirement on the Energy Market

	Coal (MWh)	Natural gas (MWh)	Oil (MWh)	Other fuel (MWh)	CO ₂ (tons)
Regulation requirement, 100 MW	-313.9 (215.8)	441.3** (181.9)	-101.4 (63.5)	-26.1 (19.9)	-246.1*** (81.8)
Load controls	Yes	Yes	Yes	Yes	Yes
Other reserves controls	Yes	Yes	Yes	Yes	Yes
Retirement controls	Yes	Yes	Yes	Yes	Yes
Quadratic time trend	Yes	Yes	Yes	Yes	Yes
Fuel price controls	Yes	Yes	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes	Yes
Month, day of week, and hour effects	Yes	Yes	Yes	Yes	Yes
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.90	0.82	0.46	0.56	0.98
Mean of dep. var.	39,821	14,910	571	581	45,220

Note: This table shows estimates from five separate time-series regressions. For the first four columns, the dependent variable is total MWh of electricity generated per hour in the PJM market by units that appear in CEMS data (where each column aggregates across all units of a particular fuel type). For the right-most column, the dependent variable is CO₂ emissions (tons) per hour for all PJM units in CEMS (i.e. combining across the four unit types from the first four columns). Coefficients on control variables are shown in the Appendix, Table A6. The unit of analysis is an hour. Standard errors are clustered by sample week.

regression results aggregated by fuel type, in Table 2. Most PJM boiler units burn coal, although some burn natural gas or oil. Similarly, most combined cycle units and combustion turbines use natural gas, but some use oil. Consistent with the boiler results in Table 1, we see a decrease in coal-fired generation and oil-fired generation in Table 2. Similarly, consistent with the combined cycle results, we see an increase in natural gas generation.

To understand the implications of this change in the generation mix, we can estimate a similar regression with CO₂ emissions as the dependent variable. We now aggregate CO₂ emissions across all units in CEMS. We see in the right-most column of Table 2 that CO₂ emissions fall when the regulation requirement is raised: for every 100 MW increase in the regulation requirement, we estimate 250 fewer tons of CO₂ are emitted by units participating in the energy market, statistically significant at the one percent level. This is in line with the fuel use changes described above. Specifically, the PJM-wide emissions rate for natural gas units (combining CC and CT units) in our sample is 0.42 tons CO₂ per MWh; the emissions rate for coal and other units combined is 0.95. This difference implies that a shift of 440 MWh from coal to gas would lead to a decrease in CO₂ emissions of around 230 tons, close to our estimate of 250 tons in Table 2.²⁵

²⁵The difference between this back-of-the-envelope calculation and our estimate may be due to differential heat rates of marginal versus average units, and/or to a heat rate effect of regulation provision that we describe below.

4.3 Robustness Checks

In the Appendix, we explore a suite of robustness checks and placebo regressions. First, we show (Figure A5) sensitivity of the generation mix changes and CO₂ emissions results to the removal of various control variables. One might worry about oversaturation in our time-series regression, but these results provide reassurance that the effects we estimate are not driven by the inclusion of too many controls. The estimates are generally centered around the main results we report in Tables 1 and 2. Moreover, we show that the sparsest possible specification (which controls only for peak and valley forecasts, which are a determinant of the regulation requirement in the first part of the sample) shows qualitatively similar results.²⁶

Next, we limit the sample to hours for which there is overlap in the peak and valley forecasts between the first and second half of the sample (Table A7). These peak and valley forecast controls are important as the regulation requirement is a direct function of them for some of the sample. One might worry about extrapolating the impact of these covariates beyond the range observed in the second half of the sample. Results are, however, qualitatively similar. Across all nine outcome variables (four technology, four fuel types, and CO₂ emissions), some of the individual coefficients do change, but we continue to see increases in combined cycle generation and decreases in emissions.²⁷

We also show that the CO₂ results are similar across a suite of additional control variables. Given the identification strategy, which relies primarily on the peak and valley forecast controls (as shown in the Appendix simulations), these are not expected to be important identification checks - and indeed, our results are similar across this broad range of specification checks. Specifically, the additional specifications are as follows. We include linear or cubic time trends, rather than quadratic. Next, we more flexibly control for system-wide load using a spline. Then we add weather controls for Chicago, whereas the main specification uses only Philadelphia weather. Next, we use a more flexible fuel price control: splines with three knots for natural gas prices and for oil prices. We then control more flexibly for PJM load and the natural gas price using non-parametric controls. Finally, we control for the standard deviation of the regulation requirement over the previous 72 hours in two different ways.²⁸

²⁶An exception is the “other technology” and “other fuel” results, for which most estimates with fewer controls are very similar, although the sparsest possible specifications yield somewhat different results. However, this does not change any of our main conclusions about the changes to boilers versus combined cycle units, nor about the changes to coal versus natural gas.

²⁷As the model in the next section shows, some heterogeneity across the regulation requirement, and therefore across the peak and valley forecast controls, is to be expected.

²⁸One might be concerned that our main results are driven by the change in the variance of the regulation requirement from the first to second half of our sample, visible in Figure 2. Some of this will be absorbed by

In addition to these specifications with additional controls, we examine several other specification checks. First, we collapse to a daily specification to allow for across-hour effects as in Bushnell and Novan (2021). Second, we use only CEMS units that do not retire during our sample period, to ensure that the main results are driven by these non-retiring units as opposed to units that retired (perhaps for secular reasons) over this time period. Next, we calculate alternative standard errors, specifically Newey-West standard errors with a maximum lag length of 168 hours (one week).²⁹

Finally we examine a number of checks to ensure that our data cleaning measures do not drive the results. These are expected to yield very similar point estimates, and indeed they do. First, we use a constructed regulatory requirement variable, from PJM reports on policy changes, rather than the reported regulatory requirement variable (these two variables are equivalent in most hours, but occasionally the reported regulatory requirement variable deviates from policy). Second, we use reported gross generation rather than our re-scaled net generation. Third, we ignore steam load when re-scaling to net generation. Fourth, we use the raw variables originally reported in Eastern Prevailing Time rather than correcting to Eastern Standard Time, to ensure that the associated data cleaning does not drive our results.

Across all of these specifications, we estimate fuel use shifts (coal to natural gas), with magnitudes and statistical significance comparable to what we display in Tables 1 and 2. We also estimate a negative and statistically significant impact of the regulation requirement on CO₂ emissions. Overall, across these sixteen robustness checks, we estimate qualitatively similar fuel use shifts (increased generation by natural gas units of 260 to over 500 MWh) and CO₂ reductions (180 to 350 tons).

We also show CO₂ results using reported rather than constructed emissions (Table A8). With this variable, we again estimate statistically significant emissions reductions.

Additional regressions in the Appendix show estimated effects at a number of placebo units (Table A9). We run our primary regression using generating units in nearby states that do not participate in the PJM wholesale energy market.³⁰ We also separately use CEMS-

our post-change dummy. Nonetheless, we also estimate specifications that control directly for the standard deviation of the regulation requirement. Including the standard deviation does not, however, qualitatively change our results - providing reassurance that this change in the variance of the requirement does not drive our estimated impacts. The change in variance may be of interest in its own right, which we leave for future researchers to analyze.

²⁹If there is cross-day correlation, then a rule of thumb like $N^{1/4}$ (in this case = 12 hours) would be insufficient, so we use a more conservative lag length.

³⁰Specifically, our full dataset contains data on all CEMS units in Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, North Carolina, New Jersey, Ohio, Pennsylvania, Tennessee, Virginia, Washington DC, and West Virginia. These states are covered in part by PJM. However, in states such as Indiana, only a minority of the plants participate in the PJM market, with the rest primarily participating in the MISO

reporting units that are classified as part of commercial or industrial operations; this includes generation from facilities such as hospitals and petroleum refineries. We also estimate our main specification with wind generation as the dependent variable. Finally, we consider a “residual” category of generation as the dependent variable. Following Davis and Hausman (2016), we estimate the effects of the regulation requirement on the amount of generation that would be needed to satisfy total demand, after accounting for the generation quantity reported in CEMS. This category accounts for nuclear generation, hydro generation, net imports, and small units not reported in CEMS. Across all of these categories, we estimate only very noisy effects of the regulation requirement.

We also estimate a series of two-way fixed effects regressions, in which the plants in nearby states serve as “control” units, allowing us to remove day-of-sample effects (Table A10). Results are qualitatively similar, with comparable point estimates but larger standard errors.³¹

Finally, we note that our results incorporate an additional effect on CO₂ emissions. When a power plant supplies frequency regulation, its heat rate is impacted – the amount of fuel it must use per unit of electricity sold. This is for two reasons. First, the heat rate at an individual generator depends on its generation level; it is non-linear (and frequently modeled as quadratic). Thus because generators are operating at new set points (the point around which they move in response to the regulation signal), their heat rate could change. Second, the generator must move up and down around its operating set point, rather than holding steady at a given level of output. This will worsen the heat rate, i.e. require greater heat input (and therefore more CO₂ emissions) per unit of electricity sold (Hirst and Kirby, 1997; Hummon et al., 2013). Our regressions implicitly incorporate these two effects. Because our CO₂ emissions rate is time-varying, our left-hand side variable in Table 2, Column 5 will vary as the heat rate changes. These two effects do not appear to be the main drivers of our results, given the magnitude of the generation mix changes we observe and how closely our back-of-the-envelope CO₂ calculations line up, above.

Overall, our empirical estimates show that when more frequency regulation is needed, substantial fuel use changes occur in the energy market. In particular, coal units sell less in the energy market while natural gas units sell more. This leads to an overall decrease in

wholesale market. Our placebo sample consists therefore of, e.g., MISO-participating units in states such as Indiana.

³¹This two-way fixed effects specification does not serve as our main specification for a couple of reasons, as described in the Appendix. Even in this specification, our primary source of identification is the switch from a varying regulation requirement in the first half of the sample to a flat requirement in the second half. Thus “control” units in other states are not expected to provide much additional help with identification. The time-series specifications are a more transparent way to implement this identification strategy. The two-way fixed effects can help if there additional regional or national secular trends that are not adequately controlled for, and as such it is reassuring that they yield qualitatively similar estimates.

CO₂ emissions, holding generation constant. This magnitude is quite large – why would a 100 MW change in the regulation market lead to a roughly 400 MWh change in the energy market? To understand the magnitude and the mechanism behind it, we turn to a simplified model.

5 Stylized Model of the Electricity Market

5.1 The System Operator’s Optimization Problem

In this section, we develop a stylized model of the regulation and energy markets, following Kirschen and Strbac (2004). The goal is not to fully represent all the complexities of the electrical grid, but to show some of the mechanisms by which an increase in the regulation requirement will change generation decisions in the energy market. Specifically, we elucidate how a system operator optimally procures energy and regulation services, especially in the presence of constraints over minimum generation. We use a stylized version of the system operator’s problem, which is essentially a single-period unit commitment problem.³² While we focus on regulation provision, the model has features that apply to reserves markets in general.

We assume that there are multiple heterogeneous thermoelectric generating units participating in an energy market and a regulation market. We evaluate how the plants operate before and after the regulation requirement increases. We use the following notation:

- x_i - generation for energy market for unit i ³³
- y_i - one-sided capacity committed to the regulation market for unit i ³⁴
- p_x - energy market price per MWh
- p_y - regulation market price per MW
- m_i - marginal cost of generation for unit i , a function of fuel use and wear and tear

³²Note we do not model the bidding behavior of individual plants. We are essentially assuming that there is no market power, and so plants bid their marginal costs. This is realistic if there are many firms and/or if regulators are able to observe marginal cost and thus punish anti-competitive bidding. The latter is especially likely to be true in our context. Annual market monitoring reports state that the exercise of market power has not generally been observed in PJM’s frequency regulation and energy markets (for instance, https://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2014/2014-som-pjm-volume2-sec1.pdf, and other years’ reports).

³³This is the set point around which regulation will be provided, if the generator offers frequency regulation.

³⁴By one-sided, we mean the capacity available in either direction. The generator must be available to deviate up or down from its set point x_i by the amount y_i .

- n_i - marginal cost of regulation for unit i , a function of wear and tear
- M_i - minimum generation for unit i , a physical constraint
- C_i - maximum generation for unit i , a physical constraint
- r_i - maximum regulation for unit i , a physical constraint³⁵

Additionally, we make the following assumptions:

- Firms are sufficiently small that their actions do not influence the market price in either the regulation or the energy market (i.e. no market power).
- Constant marginal costs of generation and of regulation for each unit, with no fixed costs. These are heterogeneous across plants, a function of the technology installed (fuel choice, prime mover type, etc).³⁶
- $r_i < C_i - M_i$ for each unit i .³⁷

Focusing on the energy market and leaving aside (for now) minimum constraints, this yields a short-run marginal cost curve that is a step function, as shown in the left-hand panel of Figure 3. The height of each step is the marginal cost m_i for each fuel and technology combination, and the width is the maximum generation C_i for each fuel and technology combination. Typically natural gas combined cycle units and coal-fired units are the cheapest to operate, and natural gas combustion turbines operate at much higher cost and with smaller capacities.


Many papers in the economics literature focus on the energy market and elide minimum constraints, yielding a supply curve much like the left-hand panel in Figure 3 (Borenstein, Bushnell and Wolak, 2002; Davis and Hausman, 2016). Dispatch order would then follow a least-cost framework, in which the lowest cost units are dispatched, up until demand (exogenously determined) has been fulfilled. In such a supply curve, when demand exogenously changes, one can examine whether a different unit is on the margin to examine price impacts as well as emissions impacts. This kind of framework is also used in the literature to examine

³⁵A typical conventional power plant can commit at most 10 to 20 percent of its capacity to regulation (Makarov et al., 2008; Atanacio et al., 2012).

³⁶A fully realistic model of the electric grid would allow for non-linear costs within each unit, as in Hirst and Kirby (1997). However, constant marginal costs are frequently assumed in the electricity economics literature (Borenstein, Bushnell and Wolak, 2002; Mansur, 2008; Davis and Hausman, 2016) and are sufficient to illustrate the mechanisms at play in our model.

³⁷This is reasonable to assume in our empirical setting. As noted above, a power plant can typically commit 10 to 20 percent of capacity to regulation. Typical minimum constraints are around 30 to 50 percent of maximum capacity. For further discussion of this constraint, see Kirschen and Strbac (2004).

Figure 3: Short-Run Marginal Cost Curves



Note: These figures show stylized marginal cost curves for a hypothetical energy market and regulation market, with four types of participating generating units.

what happens when fuel price changes lead to a re-ordering of the dispatch, i.e. a change in which units are least-cost.

Focusing on the regulation market, our framework implies a similar short-run marginal cost curve, again a step function (right-hand panel of Figure 3). Here the height of each step is the marginal cost n_i for each fuel and technology combination, and the width is the maximum regulation r_i for each fuel and technology combination.

Solving for the competitive equilibrium in the energy market when there is no regulation market and there are no minimum constraints is simple, as described above. However, when the system operator is minimizing the cost across these two markets and when minimum constraints are incorporated, we have a more complicated mixed integer linear programming problem:

$$\begin{aligned}
 \min_{x_i, y_i} & \left(\sum_{i \in (1, 2, \dots, I)} m_i x_i + n_i y_i \right) & s.t. & \sum x_i = \text{demand}; \\
 & & & \sum y_i = \text{regulation requirement}; \\
 & & & x_i + y_i \leq C_i \quad \forall i; \\
 & & & x_i - y_i \geq M_i \quad \text{or} \quad x_i = y_i = 0 \quad \forall i; \\
 & & & 0 \leq y_i \leq r_i \quad \forall i;
 \end{aligned}$$

The operator minimizes the total cost of generation and regulation provision, subject to a

number of constraints. The total demand and regulation requirements must be satisfied.³⁸ Units cannot commit more than their total capacity across the two different markets; they must have sufficient headroom if they offer regulation. The system operator can choose not to dispatch any particular unit, but if the unit operates, it must be at least at its minimum constraint. As a result, if it commits non-zero capacity to the regulation market, it must be at its minimum constraint *plus* its regulation provision in the energy market, i.e. $x_i - y_i \geq M_i$ (i.e. have footroom). These minimum constraints are frequently elided in empirical papers, but they are nontrivial. The typical unit in our data has a minimum constraint of 30 to 50 percent of its maximum capacity.³⁹ The constraint is related to technical restrictions – operating a plant below minimum load can damage plant equipment.⁴⁰

In this model, corner solutions are possible for many individual generating units. For instance, a unit might operate at maximum capacity in the energy market ($x_i = C_i$) and not participate in the regulation market ($y_i = 0$). It might instead operate just below maximum capacity to fully participate in the regulation market, i.e. with $x_i = C_i - r_i$ and with $y_i = r_i$. It might similarly operate just above minimum capacity to fully participate in the regulation market, with $x_i = M_i + r_i$ and with $y_i = r_i$. There may also be marginal units with generation levels between minimum and maximum capacity, and/or with regulation commitments between 0 and r_i .

Now suppose that the regulation requirement is exogenously increased, and that the change is large enough that the marginal unit cannot provide the additional regulation. The system operator will procure regulation from an additional unit. Suppose this does not require a change in which plants are dispatched in the energy market (an unrealistic assumption to which we return momentarily). Then the system operator might change the energy and regulation procurement from an individual generating unit in multiple ways, as shown in Table 3.

To be able to provide regulation services, it is possible that a unit that had previously been operating at maximum capacity C_i would need to back down from maximum. This could occur if the regulation price change is large enough to outweigh the lost revenues from participating less in the energy market. This outcome is demonstrated in the first row of Table 3. It is also possible, however, that a unit that had not been participating in

³⁸The magnitude of the regulation requirement refers to the one-directional capacity needed across all units.

³⁹The median PJM combustion turbine in our data has a minimum constraint at 50 percent of its maximum; the median non-CT in our PJM data has a minimum constraint at 30 percent of its maximum capacity.

⁴⁰Minimum load constraints are sometimes determined by environmental compliance, if emissions rates are very high at low levels of generation. Cost considerations can also be a factor, if fuel efficiency is very low at low levels of generation.

Table 3: Some Potential Effects of an Increase in Regulation Requirement

Pre-period x_i	Post-period x_i	Change in x_i	Pre-period y_i	Post-period y_i	Change in y_i
$x_i = C_i$	$x_i = C_i - r_i$	$-r_i$	$y_i = 0$	$y_i = r_i$	r_i
$x_i = 0$	$x_i = M_i + r_i$	$M_i + r_i$	$y_i = 0$	$y_i = r_i$	r_i

Note: This table shows two potential effects of an increase in the regulation requirement on a generating unit participating in the energy market. In the first row, the generating unit backs down from maximum capacity C to have enough headroom to provide regulation. In contrast, in the second row, the generating unit enters the energy market to have enough footroom to offer regulation. In both cases, the generating unit increases its regulation provision (from zero to r). However the cases show effects in opposite directions and of differing magnitudes in the energy market. Other outcomes are possible as well – for instance if the generating unit had previously been offering some quantity between zero and C in the energy market, or because of market-wide re-dispatch in the energy market.

either market could be induced to enter *both* markets. In this scenario, a unit would move from zero generation in the energy market to at least a bit above its minimum operational constraint, operating at $M_i + r_i$ (or more) to be able to sell r_i regulation services. This could occur if its marginal energy cost m_i is above the market clearing energy price p_x , but the additional revenues in the regulation market make up for losses in the energy market. This outcome is demonstrated in the second row of Table 3. In short, an increase in the regulation requirement could lead an individual unit to either increase or decrease the energy it sells in the energy market. Moreover, it is possible for the change in the energy market to be *larger* than the change in the regulation market, if a unit is induced to move from not generating at all to generating above its minimum constraint.

We must also consider the follow-on changes for *other* units in the energy market. Since energy demand is exogenous and inelastic, any changes induced by one unit, as described in Table 3, must be offset by an equal amount across all other units (neglecting changes in losses due to changes in power flows, which we have not modeled). That is, the change in the regulation requirement could induce not only changes in p_y but also changes in p_x and therefore a different set of plants committed, and different dispatch levels for those plants, in both the regulation and energy markets. The system as a whole could move to a different equilibrium with different inframarginal units, and with some units changing by more than the regulation requirement change. How the system changes will depend on the ways short-run operating profits in one market (e.g., $p_x - m_i$) compare to short-run operating profits in

the other market (e.g., $p_y - n_i$) across the entire set of generators.

5.2 Simulated Market

We construct a four-unit model, showing that a regulation requirement change can have a wide range of impacts in the energy market. We then solve for the equilibrium, exogenously changing the regulation requirement to show how changes occur along the extensive margin for various plants. The four units represent three baseload units, with differing marginal costs of energy and regulation, and one peaker unit with higher marginal cost for both services (details in Appendix). All four units have minimum operational constraints.⁴¹ All four units are capable of following the regulation signal, which is an energy-neutral signal (units return to their initial set point within a specified time frame). The four units combined must meet an exogenous perfectly inelastic demand requirement as well as a perfectly inelastic regulation requirement.

Figure 4 shows how each unit changes its generation (top row) and regulation provision (bottom row) as the regulation requirement is scaled up (results are also shown in table form in the Appendix, Table A11). In the top row, note the axes are scaled differently across units.

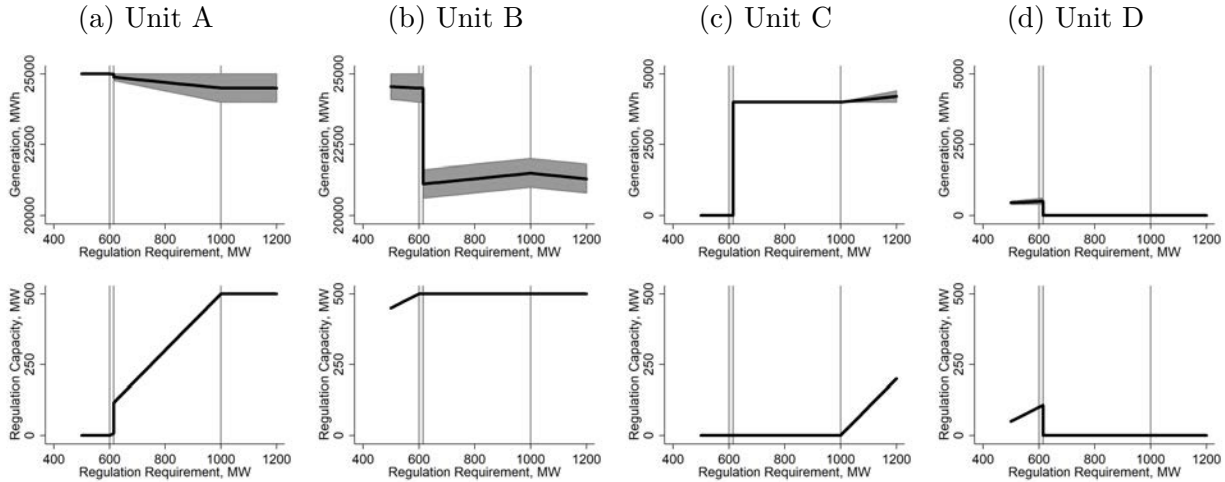
If there were no minimum constraints and no regulation requirement, Units A and B would provide energy at their maximum (25,000 MWh each) as they have the lowest marginal costs of energy provision. However, the minimum constraints and the regulation requirements change the equilibrium in qualitatively important ways.

Consider first the setting where the regulation requirement is set at 500 MW. Unit A produces the maximum possible energy (25,000 MWh), but Unit B provides only 24,550. It then uses its remaining capacity (450 MW) to provide regulation. Unit C is not dispatched to fill in the remaining 450 MWh to satisfy energy demand. Instead, the most expensive unit, Unit D, provides 450 MWh of energy and 50 MW of regulation services. This is because Unit D, a peaker, has a lower minimum generation requirement than does Unit C – in fact, it operates just at its minimum generation requirement. This illustrates how the minimum generation constraints can alter the dispatch order. (We caution that the results are particularly “lumpy” in that the minimum generation has an especially pronounced impact because there are only four generators; results would be less lumpy in a large market.)

Next, consider what happens as the regulation requirement increases from 500 to 550 or 600 MW. Unit B decreases its energy a bit, to be able to provide additional regulation

⁴¹The minimum is the same for the three low-cost plants. The high-cost unit has a smaller minimum operational constraint, representing the fact that the peaking portion of the electricity market is made up of many small peaker units that can each be dispatched at quite small levels of generation.

Figure 4: Simulated Four-Unit Model



Note: This figure displays the equilibrium results for a four-unit model with energy and regulation output. The top row shows generation outcomes with a black line, surrounded by the regulation band in grey. Generation is in MWh provided (over one hour). The bottom row shows the capacity committed to regulation. Units are ordered left to right from least to most expensive, with the ranking the same across the generation and regulation markets (Unit A has the lowest marginal cost for both services; Unit D the highest). All four units face minimum and maximum constraints. Energy demand is held constant, while the regulation requirement varies exogenously across the x-axis. Discontinuities and kinks are shown with vertical grey lines, at a regulation requirement of 600, 615, and 1000 MW. Quantities are given in table form, along with cost and constraint details, in Appendix Table A11.

– it must decrease energy provision to do so, since it had been operating at its maximum constraint when combining both energy and regulation. Unit D also provides a bit more regulation, but to do so, it must *increase* its generation, to maintain status above its minimum constraint. This illustrates how a unit wishing to provide additional regulation could conceivably increase *or* decrease its energy provision, depending on which (if any) of its constraints are binding.

Once the regulation requirement is increased to 601 MW, Unit A begins to decrease its energy, to be able to provide regulation. This is because Unit B is providing as much regulation as possible (500 MW) and cannot provide additional regulation. As the regulation requirement continues to increase, up to 615 MW, additional 1-unit changes in the regulation requirement lead to Unit A decreasing its energy provision and increasing its regulation provision.

The most interesting change occurs when the regulation requirement increases to 616 MW. At this point, the least-cost solution involves dispatching Unit C at its minimum generation. Thus Unit D exits both markets. Previously, this had not been least-cost because of Unit C’s minimum constraint. However, it is now least-cost for Unit C to operate at its minimum constraint, rather than to be turned off. So a change in the regulation

requirement from 615 MW to 616 MW leads to a very different equilibrium across all units. Unit A decreases its energy provision by around 100 MWh and increases its regulation provision by around 100 MW. Unit B decreases its energy provision by around 300 MWh and maintains the same amount of regulation provision. Unit C increases from 0 to 4000 MWh of energy, but does not provide any regulation. Unit D drops from over 500 MWh to 0 MWh of energy, and also exits the regulation market. This illustrates how a small change in the regulation requirement can lead to oversized changes in the energy market if it changes whether a unit switches between zero generation and operating at its minimum constraints.

Finally, the last change to occur in this figure is when the regulation requirement goes from 1000 to 1001 MW. Unit C provides the marginal regulation. To do so, it must increase its energy provision since it had been operating at its minimum constraint. Recalling that total generation is fixed, we see that Unit B decreases its energy provision as a result.

We make several caveats regarding the generalizability of our model. We have deliberately presented a stylized version of the energy and regulation markets, to show some of the ways that minimum constraints combine with the multi-product nature of power plants. In practice, there are two different regulation prices in the PJM market: units are rewarded separately for the quantity of MWs offered and for the accuracy with which they respond to the regulation signal. This will impact the mix of, for instance, natural gas combined cycle units versus coal units in the regulation market, since they have differing levels of accuracy. The mix of coal versus natural gas units in the regulation market will also be impacted by secular changes in fuel prices and thus the generation mix of the broader electrical grid.

Also, we have not modeled other features of the market that could interact with the minimum constraints. For instance, our model is static, and dynamic constraints such as minimum up and down times could matter. How exactly they would interact with frequency regulation provision would depend on demand profiles across the day, among other things. Moreover, in reality PJM runs multiple optimization algorithms because there are markets on different timescales: the system operator must make decisions at the day-ahead, hour-ahead, and real-time levels. How regulation interacts with unit commitment will be determined by what algorithms the system operator uses across these different timescales. Also, transmission congestion can matter; as Ryan et al. (2018) write, “we cannot make generalizations about the effects of congestion because, in practice, results would be strongly dependent on grid topology and parameters, generator sizing and location, and so on” (p 10172). Finally, our model consists of only four units; in practice PJM is of course much larger. More generating units will mean that supply is less “lumpy” and so one might expect less discontinuous changes than what is observed in Figure 4; however in a real-world market, transmission congestion could shrink the number of units that are able to respond.

Overall, there are four primary takeaways from this stylized model: (1) there are potential nonlinearities in the impacts of a regulation requirement change; (2) high marginal cost units can be dispatched over cheaper units because of minimum constraints; (3) an increase in the regulation requirement can either increase or decrease generation at a given unit; and (4) changes in generation at an individual unit can be bigger than the change in the regulation requirement. Thus we see that minimum constraints can be quite important here in that they create lumpiness in how the market responds to exogenous changes, although the economics literature tends to elide them for simplicity. Moreover, the two output markets can interact in surprising ways, with outsized impacts of the regulation market on the energy market. Finally, the specific changes that will be observed following a regulation market change will depend on a suite of parameters.⁴² As such the effects empirically observed in a real-world scenario will be sensitive to the particular time period and electricity market studied.

6 Evidence of Changes Along Extensive Versus Intensive Margins

To connect our modeling results to our empirical results, we next estimate various intensive versus extensive margin changes. We separate hourly generation into five bins for each unit: *Off*, *Below Minimum Constraint*, *At Minimum Constraint*, *Between Minimum Constraint and Maximum Capacity*, and *At Maximum Capacity*. Then we count the number of units of each fuel type in each bin in each hour, giving us a time series of bin-level counts for each fuel type. Details of the data and variable construction are in the Appendix.

We regress the count of generators falling into each bin on the regulation requirement and a vector of controls separately for each fuel/mover type. The regressions take the form:

$$N_{i,t} = \alpha + \beta R_t + X_t \Theta + \varepsilon_{i,t} \quad (2)$$

where $N_{i,t}$ is a count of units of fuel type i in hour t that have generation levels falling in a particular bin (e.g., the number of coal boiler units at 5 a.m. on November 1, 2012 with capacity factors below their minimum constraint). Again R_t is the regulation requirement and X_t is a vector of controls (the same controls as in the generation regressions above).

⁴²The fact that we see these non-linearities suggests that even *within* a given geographic region or broad time period, one might expect to see heterogeneity in the impact of a change to the regulation requirement across different levels of demand or different ex-ante levels of the regulation requirement itself. Unfortunately, we have insufficient power for exploring this heterogeneity ourselves. But, future work could look at this in depth with either additional sources of identifying variation or with additional simulations calibrated to various markets.

Standard errors are clustered by sample week.

Table 4 shows the results of these regressions for each of the three units types of interest: boiler, combined cycle, and combustion turbine (results for “other” units are shown in the Appendix, Table A12). These effects are in line with both the generation changes shown in Table 1 and with the stylized model. Panel A shows that boilers are less likely to be at their maximum (Column 5) and more likely to operate within their main operational range (Column 4) when the regulation requirement is higher. Specifically, a 100 MW increase in the regulation requirement causes one fewer units to be at maximum capacity. This is consistent with this unit providing additional regulation, and needing its set point to be below its maximum to have the flexibility for upward movements in response to a regulation signal. It is also consistent with the boiler decreasing its energy provision to accommodate additional energy provision by combined cycle plants.

Panel B shows that combined cycle units are more likely to be dispatched when the regulation requirement is higher. For every 100 MW additional regulation requirement, around two combined cycle units are more likely to be dispatched and in their main operational range (Column 4). This could be consistent either with dispatching with positive generation to be able to themselves provide frequency regulation, or with needing to fill the gap left by reduced boiler generation, shown in Panel A.

Panel C shows that combustion turbines are also more likely to be operating within their main operational range, some combination of units being less likely to be off (Column 1) and less likely to be at their maximum capacity (Column 5). This could be because they move to the middle of their range to provide frequency regulation, because they are newly dispatched to fill in for lost boiler generation, or some combination of both.

As shown in the Appendix, results are robust to an alternative construction of the minimum constraint variable. Results are also robust to using ten bins, identically spaced across the capacity of each unit (0 to 10 percent of capacity, 11 to 20 percent of capacity, etc.), rather than a minimum constraint definition.

Overall, the primary effect we see when the regulation requirement is higher is that more units operate within their main operational range, rather than being off or at the maximum constraint. This is consistent with the stylized model in the previous section. The magnitudes are also consistent with a back-of-the-envelope calculation of the number of units that would be needed to provide 100 MW of regulation. Recall that the typical unit can commit 10-20 percent of its capacity to regulation. In our sample, the average boiler or combined cycle unit has a capacity of around 230-280 MW, implying that two to four units would be needed to provide 100 MW of regulation. Combustion turbine units in our sample have a capacity of around 80 MW, so more of these plants would be needed to provide the

Table 4: The Regulation Requirement and Extensive Versus Intensive Margins

Panel A. Boilers	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	0.30 (0.97)	0.16 (0.15)	-0.09 (0.35)	0.66 (0.81)	-1.03** (0.39)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.77	0.17	0.41	0.50	0.79
Mean of dep. var.	199	5	11	99	29
Panel B. Combined Cycle Plants	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-1.71** (0.80)	-0.46* (0.27)	0.13 (0.21)	2.40*** (0.65)	-0.36* (0.21)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.65	0.28	0.13	0.62	0.35
Mean of dep. var.	56	13	9	55	2
Panel C. Combustion Turbines	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-1.63 (2.98)	-0.04 (0.42)	0.29 (0.29)	2.55 (2.13)	-1.17** (0.49)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.48	0.31	0.28	0.45	0.22
Mean of dep. var.	410	3	1	8	1

Note: This table shows estimates from 15 separate regressions. The dependent variable is a variable representing the count of units of each type generating at each level in PJM. The unit of analysis is an hour. Effects for other unit types are shown in the Appendix, Table A12. Standard errors are clustered by sample week.

same amount of regulation.

We see empirically that generators behave in intuitive ways along both the intensive and extensive margins when the regulation requirement is exogenously changed. Because they are multi-product firms, they adjust their outputs in multiple markets. Indeed, recall that one of the motivations behind the policy changes to the regulation requirement was to allow more resources to move their availability from the regulation market to the energy market (Section 4). Moreover, minimum and maximum constraints can lead to changes in the energy market that are outsized in comparison with the change in the regulation requirement.

7 Implications for Battery Deployment

Our results on frequency regulation markets have important implications for utility-scale battery storage. There is a growing interest in using energy storage, including batteries, flywheels, and loads coordinated to behave like storage, to help operate the electrical grid. While the energy economics literature has focused on the use of batteries for arbitrage, batteries can also be used to provide ancillary services such as frequency regulation (Department of Energy, 2013; International Renewable Energy Agency, 2017; Deloitte, 2018; Ryan et al., 2018). Utility-scale batteries, typically sized at 1 to 10 MW,⁴³ can charge and discharge according to an operator’s command. This can be used for intertemporal arbitrage, for instance to charge when demand is low (at night) and discharge when demand is high (the late afternoon). It can also be used at a faster timescale, to provide grid reliability services. Batteries are often better-suited to provide frequency regulation than are conventional generators, since they are able to very quickly and accurately respond to the system operator’s regulation signal.

PJM has been at the forefront of incorporating storage into ancillary service markets; storage providers found the RegD system particularly lucrative when it was first introduced.⁴⁴ Indeed, nearly all storage capacity in PJM is built for the provision of frequency regulation. While the academic literature in economics has focused on storage for energy market arbitrage, 22 of PJM’s 27 facilities list “frequency regulation” as a service; only five facilities list “load management” and just one lists “arbitrage” as a service, according to Energy Information Administration data.⁴⁵

Storage for frequency regulation or other ancillary services has been investigated in the academic literature through an engineering lens;⁴⁶ via a techno-economic model;⁴⁷ and through a theoretical economics lens.⁴⁸ These papers have generally focused on the private incentives of operators using storage for frequency regulation, and have not documented the emissions impacts of storage used for frequency regulation. Lin, Johnson and Mathieu (2019); Ryan et al. (2018) examine emissions impacts of using storage for frequency reg-

⁴³EIA-860 data for 2018 give the 25th and 75th percentiles nationwide as 1 and 10 MW. For PJM, the 25th and 75th percentiles are 2 and 20 MW. The largest battery listed is 40 MW.

⁴⁴See Maloney (2017), “Is the bloom off the RegD rose for battery storage in PJM?” in *Utility Dive*, <https://www.utilitydive.com/news/is-the-bloom-off-the-regd-rose-for-battery-storage-in-pjm/503793/>.

⁴⁵Source: EIA-860 data for 2018.

⁴⁶See Castillo and Gayme (2013); Mège, Mathieu and Andersson (2013); Xi, Sioshansi and Marano (2014); Cho and Kleit (2015); Mège, Mathieu and Andersson (2015*a,b*); Moreno, Moreira and Strbac (2015); Wu et al. (2015); De Sisternes, Jenkins and Botterud (2016); He et al. (2016); Basic, Hashmi and Meyn (2017); Namor et al. (2018); Shi et al. (2018); Watson et al. (2018); Kern, Johnson and Mathieu (2019).

⁴⁷See Stephan et al. (2016).

⁴⁸See Cho and Kleit (2015).

ulation using simple test systems, not calibrated to the PJM market, rather than ex-post empirical evidence. Thus our results may be informative for providing novel empirical evidence on the interaction between frequency regulation and energy markets, with implications for battery deployment.

For instance, with respect to the operations of conventional generators, changes to the regulation requirement over this period are analogous to the entry and exit of batteries used for frequency regulation, where a decrease in the regulation requirement can be thought of as the entry of a battery. Consider a battery entering the market in order to participate in the regulation market. Suppose this battery is inframarginal – batteries generally have high fixed costs but low marginal costs – and suppose that the battery does not participate in the energy market (recall from above that batteries in PJM generally do not provide arbitrage in the energy market). If this battery participates in each period, its entry represents a reduction in the residual regulation requirement faced by conventional generators, by a magnitude equal to the capacity of the battery. Hence, examining changes to the regulation requirement will provide insights into how conventional generators are likely to respond to changes to their residual regulation requirement, whether those changes are the result of policy or the entry and exit of batteries.

Thus we expect, based on our analysis of the PJM regulation market, that the entry of batteries in PJM would lead to generation mix changes and emissions changes in the energy market. For the time period we study (2012-2014), we could infer that battery entry (akin to a reduction in the regulation requirement) would lead to *increased* CO₂ emissions, with a gas to coal shift. However, caveats are in order. First, note that reducing the regulation requirement for other generators could lead to a lowering of *private* generation costs; whether or not this outweighs the external climate damages is an empirical question.

Moreover, recall that in the stylized model, we see that the changes in regulation can have many potential impacts on the energy market. Thus our regression results are not necessarily *externally* valid: the results cannot be simply extrapolated to alternative time periods in PJM nor to other system operators. Different fuel costs or the secular retirements of power plants could mean very different impacts of changes to the regulation market.⁴⁹ In particular, the retirement of coal plants may imply less fuel switching in recent years than what we estimate for the 2012-2014 period in Table 2.

⁴⁹Note we also expect two additional impacts of battery entry. Battery entry can impact the heat rates of conventional generators, similar to how frequency regulation impacts the heat rate, a point we discuss in Section 2. Furthermore, we note that batteries are net users of electricity (they do not have 100% round-trip efficiency), and so their entry impacts the amount of conventional power plant generation required (Department of Energy, 2016). This mechanism will not be captured with the regression approach we have taken, which conditions on total quantity demanded across the system.

Ideally, we would next empirically estimate the impact of battery entry into the frequency regulation market in PJM. Empirical analysis is, unfortunately, limited in several ways. First and most importantly, battery capacity over this time period essentially follows an upward trend. Moreover, coal plants are retiring and there is a secular shift from coal to gas nationwide in the electricity market. The simple correlation of the battery capacity variable with a linear time trend is 0.9, as is the correlation between battery capacity and coal retirements in PJM. Without additional variation, this will make it difficult to identify the causal effect of changes in battery capacity.

One could instead leverage discontinuities in battery entry, similar to the approach taken in Davis and Hausman (2016). However, the entry dates for new battery capacity differs somewhat across the two government sources we observe, making us less confident in the exact entry dates. Moreover, some of the entry dates coincide with coal plant retirement dates. To address this uncertainty, we explore using a variable available from PJM: the amount of self-scheduled RegD participating in the frequency regulation market. If one is willing to assume that batteries self-schedule (and hydro may as well), but that combustion turbines bid a non-zero marginal cost; and if one is willing to assume that batteries (and hydro) participate in RegD whereas boilers and combined cycle units participate in RegA, then the amount of self-scheduled RegD participating may serve as a useful proxy for battery and hydro participation that crowds out fossil fuel generator participation.

A final identification challenge is that we have reason to suspect that the impact of battery entry over the 2012-2018 time period is time-varying, based on documentation from the PJM market monitor. The 2017 State of the Market describes that RegA was being used to *offset* RegD movements.⁵⁰ That is, for the latter part of our sample, it is possible that increased battery capacity could lead to *more* regulation provided by conventional units, with batteries and conventional plants moving in opposite directions to follow their respective regulation signals.

All of these together make identification extremely challenging. With that caveat in mind, in the Appendix we provide a series of empirical results. We regress generation by fuel type (as we did in the previous sections) on battery capacity. We show results using both indicators of battery capacity entry, which have slightly different dates, and the self-scheduled RegD measure. Not surprisingly, the results are unstable in sign, magnitude, and statistical significance across specifications. Moreover, in placebo regressions using generators in nearby

⁵⁰Specifically, the report states: “PJM’s current regulation market design is severely flawed and does not follow the appropriate basic design logic... RegA is now explicitly used to support the conditional energy neutrality of RegD. The RegD signal is now the difference between ACE and RegA. RegA is required to offset RegD when RegD moves in the opposite direction of that required by ACE control in order to permit RegD to recharge” (pp 472-473).

states, we see specifications with similar magnitudes. This is not surprising, given the identification challenges described above.

In sum, our results have qualitatively important implications for battery deployment. The generation mix changes that could result will have impacts on CO₂ emissions. Future empirical work in this area will be important as more batteries are deployed.

8 Discussion and Conclusion

Overall, we see that changes in the structure and makeup of the frequency regulation market impact conventional generators that participate in the energy market. Specifically, we find that for every additional 100 MW *reduction* in frequency regulation required of plants in PJM, there is an approximately 400 MWh increased use of natural gas power plants in the energy market, and a corresponding decrease in the use of coal and other units. This leads to an increase in carbon emissions of 250 tons per hour, implying climate damages on the order of \$100 million or more per year. In this context, relaxing a constraint on the system (the regulation requirement) can allow for a lower cost solution to the optimization problem – but notably, the optimization problem in practice ignores external costs such as CO₂ emissions. The results are directionally robust to considering alternative controls and various alternative specifications. Results also pass a series of placebo tests using generators outside of the PJM market.

We use a simple model to demonstrate the mechanisms behind the observed generation mix change. The model considers generators as multi-product suppliers, and makes the realistic assumption that units are constrained by non-zero minimum and maximum constraints. With these assumptions, we show that policy and market changes can cause conventional power plants to move from fully off to operating at non-negligible minimum load, and vice-versa. These changes along the generators’ extensive margins can explain the magnitude of our results. Moreover, our model and results suggest that generating units should be thought of as multi-product firms participating in multiple markets, and that the concept of a “marginal” firm is much more complex after considering both minimum constraints and multiple markets.

Furthermore, we consider how changes in the frequency regulation market are related to the entry and exit of batteries. While the academic economic literature has focused on energy storage as an energy market arbitrage opportunity, the vast majority of batteries within PJM (and the country) are used for frequency regulation and not for arbitrage. Changes to the regulation requirement within PJM can be thought of as analogous to changes in battery capacity available for frequency regulation. While identification challenges prevent us from

directly measuring the impacts of battery entry and exit on conventional generators, the generation mix change that occurs due to changes in the regulation requirement suggests that battery entry might lead to *increased* CO₂ emissions.

The direction of these results matches evidence in the engineering literature about the entry of batteries in frequency regulation (Ryan et al., 2018). That research uses a unit commitment and dispatch model of a small power system to show the life-cycle impacts of batteries, and finds similar fuel switching. The magnitude of our results is larger than that in Ryan et al. (2018). The magnitudes in this paper makes sense given what we find for extensive margin changes in the number of plants participating in each market. Future engineering work could build test systems that more closely match PJM parameters to explore the difference between our magnitudes and those in Ryan et al. (2018). Effects of the sort described above are expected to be sensitive to the size of the system and the number of generators modeled in a unit commitment model. Such a test system could also allow for simulations under alternative grid conditions, such as increased coal retirements, increased deployment of renewables, etc.

Note that the magnitude and potentially the direction of these effects is context-specific, as described in depth in Ryan et al. (2018). In different regions of the country, different fuel types might be inframarginal and changes to the regulation requirement might have very different impacts.⁵¹ In markets with carbon pricing, one would also expect very different impacts. In addition, the long-term impact of batteries could be different from the short-term impacts we measure. As batteries crowd conventional generators out of the regulation market, the profitability of fossil-fuel fired plants could decline and some of these conventional plants could be more likely to retire (and/or new conventional plants could be less likely to enter). Whether this increases or decreases CO₂ emissions in the long run of course depends on whether profitability is impacted at coal or natural gas units. Future work could investigate plant retirements in this context using modeling of the sort used in Linn and McCormack (2019). Future work could also investigate how battery storage enables greater deployment of intermittent renewables such as wind and solar generation.

However, this paper does demonstrate that ancillary markets and energy markets are far more intertwined than economics researchers might have previously thought. The structure and policy details of ancillary service markets have important impacts for generators and for the energy market, and more careful research across the country is necessary to better understand these complexities in different settings. Future economics research could use structural methods to estimate power plant costs in both markets and evaluate market

⁵¹For related mechanisms in the context of wind energy, see Callaway, Fowle and McCormick (2018); in the electric vehicle context, see Holland et al. (2016*a,b*).

power after accounting for the existence of both minimum constraints and multiple markets.

Finally, while our results point to the possibility that batteries lead to increased CO₂ emissions, the results are not to say that batteries are in general harmful for climate change. First, as described above, the results are specific to a particular set of costs across fossil fuel fired plants. Moreover, they are specific to a second-best world in which CO₂ emissions are not priced. Finally, batteries may be desirable for supporting renewables integration and thus decreasing emissions. However the results do point towards the complexities inherent in designing second-best greenhouse gas abatement policy. As policy-makers continue to grapple with whether and how to support batteries and other new technologies, the realities of the electrical grid and its markets cannot be ignored.

References

- Abito, Jose-Miguel, Christopher R. Knittel, Konstantinos Metaxoglou, and Andre Trindade.** 2018. “Coordinating Separate Markets for Externalities.” *NBER Working Paper 24481*.
- Ambec, Stefan, and Claude Crampes.** 2019. “Decarbonizing Electricity Generation with Intermittent Sources of Energy.” *Journal of the Association of Environmental and Resource Economists*, 6(6): 1105–1134.
- Antweiler, Werner.** 2021. “Microeconomic models of electricity storage: Price forecasting, arbitrage limits, curtailment insurance, and transmission line utilization.” *Energy Economics*, 101.
- Archsmith, James.** 2018. “Dam Spillovers: The direct and indirect costs from environmental constraints on hydroelectric generation.” *Working paper*.
- Atanacio, Manuel, Richard Fioravanti, Warren Katzenstein, and Khoi Vu.** 2012. “Emission II Study of Advanced Storage used for Frequency Regulation.” *Department of Energy – Sandia National Labs, Project Number: Sandia Negotiation 96362*.
- Bastien-Olvera, Bernardo A., and Frances C. Moore.** 2020. “Use and non-use value of nature and the social cost of carbon.” *Nature Sustainability*.
- Borenstein, Severin, James B. Bushnell, and Frank A. Wolak.** 2002. “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market.” *American Economic Review*, 92(5): 1376–1405.
- Borensten, Severin, and James Bushnell.** 2015. “The US Electricity Industry After 20 Years of Restructuring.” *Annual Review of Economics*, 7: 437–464.
- Bushnell, James, and Kevin Novan.** 2021. “Setting with the Sun: The Impacts of Renewable Energy on Conventional Generation.” *Journal of the Association of Environmental and Resource Economists*, 8(4): 759–796.
- Busic, Ana, Md Umar Hashmi, and Sean Meyn.** 2017. “Distributed Control of a Fleet of Batteries.” *2017 American Control Conference (ACC)*, 3406–3411. doi: 10.23919/ACC.2017.7963473.
- Callaway, Duncan S., Meredith Fowlie, and Gavin McCormick.** 2018. “Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources.” *Journal of the Association of Environmental and Resource Economists*, 5(1): 39–75.
- Carson, Richard, and Kevin Novan.** 2013. “The Private and Social Economics of Bulk Electricity Storage.” *Journal of Environmental Economics and Management*, 66: 404–423.
- Castillo, Anya, and Dennice F. Gayme.** 2013. “Profit Maximizing Storage Allocation in Power Grids.” *52nd IEEE Conference on Decision and Control*, 429–435. doi: 10.1109/CDC.2013.6759919.

- Castro, Miguel.** 2020. “Intermittent Renewable Energy, Hydropower Dynamics and the Profitability of Storage Arbitrage.” *IDB Working Paper 1107*.
- Chatterjee, Arpita, Rafael Dix-Carneiro, and Jade Vichyanond.** 2013. “Multi-Product Firms and Exchange Rate Fluctuations.” *American Economic Journal: Economic Policy*, 5(2): 77–110.
- Cho, Joohyun, and Andrew N. Kleit.** 2015. “Energy Storage Systems in Energy and Ancillary Markets: A Backwards Induction Approach.” *Applied Energy*, 147: 176–183. doi: 10.1016/j.apenergy.2015.01.114.
- Cicala, Steve.** 2015. “When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation.” *American Economic Review*, 105(1): 411–444.
- Cicala, Steve.** 2017. “Imperfect Markets versus Imperfect Regulation in U.S. Electricity Generation.” *NBER Working Paper 23053*.
- Cullen, Joseph A., and Erin T. Mansur.** 2017. “Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution.” *American Economic Journal: Economic Policy*, 9(3): 106–133.
- Davis, Lucas, and Catherine Hausman.** 2016. “Market Impacts of a Nuclear Power Plant Closure.” *American Economic Journal: Applied Economics*, 8(2): 92–122.
- De Loecker, Jan.** 2011. “Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity.” *Econometrica*, 79(5): 1407–1451.
- Deloitte.** 2018. “Supercharged: Challenges and opportunities in global battery storage markets.”
- Department of Energy.** 2013. “Grid Energy Storage.” <https://www.energy.gov/sites/prod/files/2014/09/f18/Grid%20Energy%20Storage%20December%202013.pdf>.
- Department of Energy.** 2016. “Demand Response and Energy Storage Integration Study.”
- De Sisternes, Fernando J., Jesse D. Jenkins, and Audun Botterud.** 2016. “The Value of Energy Storage in Decarbonizing the Electricity Sector.” *Applied Energy*, 175: 368–379. doi: 10.1016/j.apenergy.2016.05.014.
- Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes.** 2018. “Just Starting Out: Learning and Equilibrium in a New Market.” *American Economic Review*, 108(3): 565–615. <https://doi.org/10.1257/aer.20160177>.
- Eckel, Carsten, and J. Peter Neary.** 2010. “Multi-Product Firms and Flexible Manufacturing in the Global Economy.” *The Review of Economic Studies*, 77: 188–217.
- Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram.** 2007. “Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency.” *American Economic Review*, 97(4): 1250–1277.

- Federal Energy Regulatory Commission.** 2011. “Order No. 755, Frequency Regulation Competition in Organized Wholesale Power Markets.” *Federal Register*, Vol. 76, No. 210.
- He, Guannan, Qixin Chen, Chongqing Kang, Pierre Pinson, and Qing Xia.** 2016. “Optimal Bidding Strategy of Battery Storage in Power Markets Considering Performance-Based Regulation and Battery Cycle Life.” *IEEE Transactions on Smart Grid*, 7(5): 2359–2367. doi: 10.1109/TSG.2015.2424314.
- Hirst, Eric, and Brendan Kirby.** 1997. “Creating Competitive Markets for Ancillary Services.” *Oak Ridge National Laboratory ORNL/CON-448*.
- Hirst, Eric, and Brendan Kirby.** 1998. “Simulating the Operation of Markets for Bulk-Power Ancillary Services.” *The Energy Journal*, 19(3): 49–68.
- Hledik, Ryan, Roger Lueken, Colin McIntyre, and Heidi Bishop.** 2017. “Stacked Benefits: Comprehensively Valuing Battery Storage in California.” *Brattle Group*.
- Holladay, J. Scott, and Jacob LaRiviere.** 2018. “How Does Welfare from Load Shifting Electricity Policy Vary with Market Prices? Evidence from Bulk Storage and Electricity Generation.” *The Energy Journal*, 39(6).
- Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates.** 2016a. “Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors.” *American Economic Review*, 106(12): 3700–3729.
- Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates.** 2016b. “Distributional Effects of Air Pollution from Electric Vehicle Adoption.” *Journal of the Association of Environmental and Resource Economists*, 6(S1): S65–S94.
- Hortacsu, Ali, Fernando Luco, Steven L. Puller, and Dongni Zhu.** 2019. “Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets.” *American Economic Review*, 109(12): 4302–4342.
- Hummon, Marissa, Paul Denholm, Jennie Jorenson, David Palchak, Brendan Kirby, and Ookie Ma.** 2013. “Fundamental Drivers of the Cost and Price of Operating Reserves.” *National Renewable Energy Laboratory Technical Report NREL/TP-6A20-58491*.
- Interagency Working Group on Social Cost of Greenhouse Gases, United States Government.** 2016. “Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.”
- International Renewable Energy Agency.** 2017. “Electricity Storage and Renewables.”
- Jha, Akshaya, and Frank Wolak.** 2020. “Can Forward Commodity Markets Improve Spot Market Performance? Evidence from Wholesale Electricity.” *Working Paper*.
- Johnson, Justin P, and David P Myatt.** 2006. “Multiproduct Cournot Oligopoly.” *RAND Journal of Economics*, 37(3).

- Just, Sebastian, and Christoph Weber.** 2008. "Pricing of reserves: Valuing system reserve capacity against spot prices in electricity markets." *Energy Economics*, 30: 3198–3221.
- Kern, Abigail, Jeremiah X. Johnson, and Johanna L. Mathieu.** 2019. "Environmental Impacts of Using Energy Storage Aggregations to Provide Multiple Services." *Proceedings of the Hawaii International Conference on Systems Science*.
- Kirby, Brendan J.** 2004. "Frequency Regulation Basics and Trends." *Oak Ridge National Laboratory ORNL/TM-2004/291*.
- Kirkpatrick, Justin.** 2018. "Estimating Congestion Benefits of Batteries for Unobserved Networks: A Machine Learning Approach." *Working paper*.
- Kirschen, Daniel, and Goran Strbac.** 2004. *Fundamentals of Power System Economics*. John Wiley & Sons, Ltd.
- Knittel, Christopher R., and Konstantinos Metaxoglou.** 2008. "Diagnosing Unilateral Market Power in Electricity Reserves Market." *The Journal of Energy Markets*.
- Linn, Joshua, and Jhih-Shyang Shih.** 2019. "Do lower electricity storage costs reduce greenhouse gas emissions?" *Journal of Environmental Economics and Management*, 96: 130–158.
- Linn, Joshua, and Kristen McCormack.** 2019. "The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions." *RAND Journal of Economics*, 50: 733–767.
- Lin, Yashen, Jeremiah X. Johnson, and Johanna L. Mathieu.** 2019. "Emissions Impacts of Using Energy Storage for Power System Reserves." *Applied Energy*, 168: 444–456.
- Liski, Matti, and Iivo Vehvilainen.** 2020. "Gone with the Wind? An Empirical Analysis of the Equilibrium Impact of Renewable Energy." *Journal of the Association of Environmental and Resource Economists*, 7(5): 873–900.
- Makarov, Y. V., S. Lu, J. Ma, and T. B. Nguyen.** 2008. "Assessing the Value of Regulation Resources Based on Their Time Response Characteristics." *Pacific Northwest National Laboratory, Prepared for California Energy Commission*.
- Mansur, Erin T.** 2008. "Measuring Welfare in Restructured Electricity Markets." *The Review of Economics and Statistics*, 90(2): 369–386.
- Mansur, Erin T, and Matthew W White.** 2012. "Market Organization and Efficiency in Electricity Markets." *Working paper*.
- Mayo, John W.** 1984. "Multiproduct Monopoly, Regulation, and Firm Costs." *Southern Economic Journal*, 51(1): 208–218.

- Mégel, Olivier, Johanna L Mathieu, and Goran Andersson.** 2013. “Maximizing the Potential of Energy Storage to Provide Fast Frequency Control.” *IEEE PES ISGT Europe 2013*, 1–5. doi: 10.1109/ISGTEurope.2013.6695380.
- Mégel, Olivier, Johanna L Mathieu, and Goran Andersson.** 2015a. “Scheduling Distributed Energy Storage Units to Provide Multiple Services Under Forecast Error.” *International Journal of Electrical Power & Energy Systems*, 72: 48–57. doi: 10.1016/j.ijepes.2015.02.010.
- Mégel, Olivier, Johanna L Mathieu, and Goran Andersson.** 2015b. “Stochastic Dual Dynamic Programming to Schedule Energy Storage Units Providing Multiple Services.” *2015 IEEE Eindhoven PowerTech*, 1–6. doi: 10.1109/PTC.2015.7232775.
- MIT Energy Initiative.** 2011. “Managing Large-Scale Penetration of Intermittent Renewables.”
- Monitoring Analytics.** 2018. “State of the Market Report for PJM.”
- Moore, Frances C., Uris Baldos, Thomas Hertel, and Delavane Diaz.** 2017. “New Science of Climate Change Impacts on Agriculture Implies Higher Social Cost of Carbon.” *Nature Communications*, 8(1): 1607.
- Moreno, Rodrigo, Roberto Moreira, and Goran Strbac.** 2015. “A MILP Model for Optimising Multi-service Portfolios of Distributed Energy Storage.” *Applied Energy*, 137: 554–566. doi: 10.1016/j.apenergy.2014.08.080.
- Namor, Emil, Fabrizio Sossan, Rachid Cherkaoui, and Mario Paolone.** 2018. “Control of Battery Storage Systems for the Simultaneous Provision of Multiple Services.” *IEEE Transactions on Smart Grid*, 10(3): 2799–2808. doi: 10.1109/TSG.2018.2810781.
- Ovaere, Marten, and Kenneth Gillingham.** 2019. “The Heterogeneous Value of Solar and Wind Energy: Empirical Evidence from the United States and Europe.” *Working paper*.
- Pindyck, Robert S.** 2017. “Coase Lecture—Taxes, Targets and the Social Cost of Carbon.” *Economics*, 84(335): 345–364.
- Reguant, Mar.** 2014. “Complementary Bidding Mechanisms and Startup Costs in Electricity Markets.” *Review of Economic Studies*, 81: 1708–1742.
- Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni.** 2018. “Country-Level Social Cost of Carbon.” *Nature Climate Change*.
- Ryan, Nicole A., Yashen Lin, Noah Mitchell-Ward, Johanna L. Mathieu, and Jeremiah X. Johnson.** 2018. “Use-Phase Drives Lithium-Ion Battery Life Cycle Environmental Impacts When Used for Frequency Regulation.” *Environ. Sci. Technol.*, 52: 10163–10174.

- Schwenen, Sebastian.** 2015. “Strategic bidding in multi-unit auctions with capacity constrained bidders: the New York capacity market.” *RAND Journal of Economics*, 46(4): 730–750.
- Shi, Yuanyuan, Bolun Xu, Di Wang, and Baosen Zhang.** 2018. “Using Battery Storage for Peak Shaving and Frequency Regulation: Joint Optimization for Superlinear Gains.” *IEEE Transactions on Power Systems*, 33(3): 2882–2894. doi: 10.1109/TPWRS.2017.2749512.
- Stephan, A., B. Battke, M. D. Beuse, J. H. Clausdeinken, and T. S. Schmidt.** 2016. “Limiting the Public Cost of Stationary Battery Deployment by Combining Applications.” *Nature Energy*, 1. doi: 10.1038/nenergy.2016.79.
- Tabari, Mokhtar, and Blake Shaffer.** 2020. “Paying for performance: The role of policy in energy storage deployment.” *Energy Economics*, 92.
- Tacka, Natalie.** 2016. “RTO/ISO Regulation Market Comparison.” *Regulation Market Issues Senior Task Force, April 13*.
- Watson, David, Cameron Hastie, Brady Gaudette, and Marianne Rodgers.** 2018. “Demonstrating Stacked Services of a Battery in a Wind R&D Park.” *IEEE Transactions on Power Systems*, 33(2): 1411–1419. doi: 10.1109/TPWRS.2017.2718512.
- Wu, Di, Chunlian Jin, Patrick Balducci, and Michael Kintner-Meyer.** 2015. “An Energy Storage Assessment: Using Optimal Control Strategies to Capture Multiple Service.” *2015 IEEE Power & Energy Society General Meeting*, 1–5. doi: 10.1109/PESGM.2015.7285820.
- Xi, Xiaomin, Ramteen Sioshansi, and Vincenzo Marano.** 2014. “A Stochastic Dynamic Programming Model for Co-optimization of Distributed Energy Storage.” *Energy Systems*, 5: 475–505. doi: 10.1007/s12667-013-0100-6.
- Yu, Nanpeng, and Brandon Foggo.** 2017. “Stochastic valuation of energy storage in wholesale power markets.” *Energy Economics*, 64: 177–185.
- Zhou, Zhi, Todd Levin, and Guenter Conzelmann.** 2016. “Survey of U.S. Ancillary Services Markets.” *Argonne National Laboratory ANL/ESD-16/1*.

A1 Appendix

A1.1 Data Appendix

A1.1.1 Gross to Net Conversion

As described in the main text, we must re-scale the CEMS-reported hourly generation to account for both in-house load and incomplete reporting of combined cycle units. Specifically, we do as follows.

In the EIA-923 dataset, we observe annual generation by plant. While EIA-923 reports monthly generation, it is imputed for some units. Thus we focus on the annual generation variable, which is not imputed. EIA-923 reports generation at a somewhat finer scale: prime mover by fuel type within a plant (e.g., aggregating across all coal boilers within a plant). However, we are most confident in the matching at the plant level as opposed to the prime mover by fuel type level, since there may be some differences in the reporting of technology between EIA and CEMS.

We merge annual CEMS data with annual EIA data at the plant level. For each plant-year, we calculate the ratio of net to gross generation. At plant-year combinations with small generation quantities, this may lead to outliers, so we take the median across years for each plant. We also winsorize the upper and lower 2% to deal with outliers – the 2nd percentile is 0.4 and the 98th percentile is 2.3. Across all electrical generating units in PJM, the median is 0.95, fairly consistent with (Cicala, 2017). The median for boilers is 0.92. The median for combustion turbines is 0.98. The distribution for combined cycles is bimodal, with one mass at around 0.97 (consistent with reporting both cycles) and one mass at around 1.5 (consistent with reporting only one cycle).

A1.1.2 Minimum Constraints

First, we estimate the minimum constraint for each generator, using EIA-860 data on minimum operational constraints. We observe reported minimum operational constraints for the years 2013-2014; they are not reported in the 2012 EIA-860. Unfortunately, a comprehensive merge between EIA-860 and CEMS at the unit level does not exist. However, merging at the plant level, or even at the plant by prime mover by fuel type level, is straightforward. Accordingly, we bring in minimum operational data as follows. For around half of generator-year combinations at electrical generating units in PJM, the minimum operational constraint is the same across all units within a plant (when expressed as a percentage of maximum capacity), so merging at the plant level is appropriate. For the remaining generator-year combinations, we use the median operational constraint within the plant at the prime mover

by fuel type level. Some units (representing 3% of generation) do not appear in the minimum constraints data in EIA-860, and for these units we use the median constraint by prime mover and fuel type across all PJM plants.

Example plots of hourly capacity factors show that these minimum constraints are visible in hourly data (Figure A3). Here we show nine histograms – one unit at three large plants for each of our three main technology types. A vertical black line depicts the minimum operational load in EIA-860 data. For most of these units, the vertical line is close to a discontinuity in the hourly histogram.

However, in a robustness check, we construct an alternative minimum operational load using the unit-level observed behavior, as follows. We calculate the portion of hours a plant is generating at a capacity factor of 0, a capacity factor between 0 and 10 percent, between 10 and 20 percent, etc. We then use as the minimum operational load whatever is the smallest bin in which at least 5 percent of non-zero generating hours fall. This is a proxy for the discontinuities observed visually in the histograms. We generally calculate minimum operational loads of around 40 to 60 percent for the boilers and CC plants, although we also observe units with a very small minimum constraint (0-10% of capacity), especially for the CT units. (Regression results using this alternative minimum constraint measure are shown in Table A13.)

Once we have a measure of minimum constraints for each unit, we proceed as follows. We calculate the capacity factor of each unit in each hour, defined as net generation divided by maximum observed generation. We then place each unit-hour observation into one of five bins: *Off* (capacity factor of zero), *Below Minimum Constraint* (capacity factor between 0% and less than 5% of the minimum constraint to maximum capacity ratio), *At Minimum Constraint* (capacity factor within 5% of the minimum to maximum ratio), *Between Minimum Constraint and Maximum Capacity*, and *At Maximum Capacity* (capacity factor between 95% and 100%).

A1.1.3 RegD Data for Battery Regressions

While our data do not allow us to directly identify battery participation in the regulation market, we can draw some inferences based on the technical capabilities of batteries. First, PJM documentation implies that batteries participate exclusively in RegD rather than RegA.⁵²

⁵²See, e.g., the 2015 presentation “Performance, Mileage and the Mileage Ratio” at <https://www.pjm.com/-/media/committees-groups/task-forces/rmistf/20151111/20151111-item-05-performance-based-regulation-concepts.ashx>, or Figure 10-19 in the 2015 State of the Market report, at http://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2015/2015-som-pjm-volume2-sec10.pdf.

Second, because batteries typically have very low marginal costs and because there is no evidence of market power exertion in regulation markets,⁵³ we infer that batteries typically use self-scheduling, rather than pool procurement. It is possible that units will switch in and out of self-scheduling and pool procurement, as a function of activity in the energy market. To address this, we use the daily maximum of self-scheduled RegD provision as our independent variable. As Figure A4 shows, this measure experiences a number of discrete large jumps over this time period, consistent with the entry of new batteries. However, we note that the results in Table A15 should be taken as only suggestive because of the data limitations in this setting.

A1.1.4 Fuel Types and Unit Types

From CEMS, we observe fuel types and unit types. The raw CEMS data lists 37 unique primary fuel types. The most common are coal, pipeline natural gas, and diesel oil. Less common categories include, e.g., “residual oil” “process gas,” “wood,” etc., as well as combinations of these fuels, e.g., “coal, natural gas.” We generate four categories: “coal” (which aggregates across coal as well as a small number of units using “coal refuse” or “petroleum coke”), “pipeline gas” + “natural gas,” “oil” (diesel, residual, or other oil), and “other,” where “other” aggregates across, e.g., wood, units listing combinations of fuels, and units for which we do not have a fuel type.

The raw CEMS data similarly lists 22 different technology types, with the most common being “combustion turbine”, “dry bottom wall-fired boiler,” and “combined cycle.” We generate four categories: “boiler” (an aggregation of all boilers, stokers, and tangentially-fired units), “combined cycle,” “combustion turbine,” and “other.” The latter includes a small number of other technology types, a small number with unreported technology type, and some units that changed technology over this 2012-2014 sample.

For our 2012-2014 sample, total gross generation by category is shown in Table A1.

A1.1.5 CEMS Versus EIA Generation Data

In addition to the net-versus-gross distinction described above, the CEMS and EIA data differ in their coverage across plants. EIA data include hydro, nuclear, solar, wind units, etc. EIA data also include small coal, gas, and oil units not in CEMS. Total generation by fuel type can be compared in Tables A1 and A2. The difference between CEMS and EIA data is accounted for by the “residual” generation variable we construct, equal to the

⁵³See, e.g., the 2019 State of the Market report, at http://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2019/2019-som-pjm-sec10.pdf.

Table A1: Total Annual Generation by Unit Type, CEMS Data, 2012-2014

Unit Type	Generation, TWh
Coal, Boiler	347
NG, CC	127
NG, CT	8
Switch	4
Oil, Boiler	3
Oil, CC	2
NG, Boiler	<1
Oil, CT	<1
Other, Boiler	<1

Note: This table shows annual generation over 2012-2014 for the aggregations of fuel by technology type that we have used. Data coverage is all CEMS-reporting electrical generating units in PJM. Data source is CEMS for generation, fuel type, technology type; and EIA for electrical generating unit designation and PJM designation.

difference between total demand reported by PJM and total generation reported in CEMS. This residual variable thus captures the behavior of nuclear, etc. units; as well as in-house load and imports and exports between PJM and other ISO/RTOs.

A1.1.6 Hour Naming Conventions

PJM data are reported in both Coordinated Universal Time (UTC) and Eastern Prevailing Time (EPT). CEMS data, in contrast, are reported in local, standard time (Central or Eastern, depending on the plant's location). We convert all PJM data to Eastern Standard Time (EST). For CEMS units in Illinois and parts of Indiana, Kentucky, Michigan, and Tennessee, we convert from Central Standard Time (CST) to Eastern Standard Time. Thus all regressions use variables in Eastern Standard Time. Regression results (Table A7, Columns 6 and 11) are similar if one uses the raw data, mixing EPT, CST, and EST across variables and plants.

A1.1.7 Other

We drop one hour (5 a.m. on April 2, 2013) when the regulation requirement is listed as zero. This represents less than 0.01 percent of our sample (19,689 hours).

Table A2: Total Annual Generation by Fuel Type, EIA Data, 2012-2014

Unit Type	PJM Generation, TWh
Coal	336
Nuclear	276
Natural Gas	142
Wind	16
Waste Coal	9
Hydroelectric Conventional	7
Biogenic Municipal Solid Waste and Landfill Gas	5
Distillate Petroleum	2
Other (including nonbiogenic MSW)	1.8
Petroleum Coke	1
Wood and Wood Waste	0.9
Other Gases	0.7
Solar PV and thermal	0.6
Residual Petroleum	0.3
Waste Oil	0.2
Other Renewables	<0.01
Hydroelectric Pumped Storage	-2

Note: This table shows annual generation over 2012-2014. Data coverage is all PJM units in EIA-923 data operating as independent power producers or electric utilities. Data source is EIA for generation, fuel type, sector, and PJM designation.

A1.2 Simulated Dataset to Illustrate the Identification Strategy

As is shown in Figure 2 in the main text, and discussed in Section 4.1, the bulk of our identifying variation comes from a policy change mid-way through our sample, in which the regulation requirement changes from being a function of forecasted peak and valley load (which change daily) to a flat requirement (albeit with separate levels in peak versus off-peak hours). In addition, as a secondary source of variation, we leverage two policy changes that modified the *multiplier* used to convert from forecasted peak and valley load to the regulation requirement

Thus, the second half of our sample, during which the regulation requirement does not vary across days, allows us to identify the effects of control variables (including the forecasted peak and valley load, as well as other things that may be correlated with these forecasts) separately from the effects of the regulation requirement, our variable of interest.

To illustrate how this works, we conduct a simulation in which we directly control the data-generating process. We set our sample size to 20,000, roughly equal to the sample size in Tables 1 and 2. We construct a peak forecast variable, normally distributed with mean zero and standard deviation equal to one.⁵⁴ We then construct a treatment variable, equal to the peak forecast variable in the first half of the sample and equal to zero in the second half, as shown in Figure A1. The outcome variable is a function of a constant, the peak

⁵⁴For simplicity, we use only a peak forecast variable and not separate peak and valley forecasts across hours.

forecast, the treatment variable, and random noise (also normally distributed with mean zero and standard deviation equal to one):

$$y_t \equiv 1 + 1 \cdot \text{peak}_t + 1 \cdot \text{treatment}_t + \varepsilon_t.$$

A successful identification strategy will thus recover a coefficient on the peak variable equal to one, and a coefficient on the treatment variable equal to one. Table A3 illustrates a series of regressions. In Column 1, the regression is correctly specified, including the entire sample and controlling for the peak forecast variable. As expected, all three coefficients are estimated to be 1.0, with a high degree of precision.

Columns 2 through 5 illustrate how identification is achieved by displaying specifications that are *not* identified. In Column 2, only the first half of the sample is included. Thus the effects of the peak forecast and treatment variable cannot be estimated. The software has dropped the coefficient on the treatment variable because of perfect collinearity, and the effects of both variables have been rolled into the coefficient on “Peak,” which is now biased upwards.


In Column 3, only the second half of the sample is included. The software has again dropped the coefficient on the treatment variable because of perfect collinearity. The effect of the peak forecast variable can be correctly estimated, but the treatment effect of interest cannot be recovered.

In Column 4, the entire sample is included, but the crucial “Peak” control has been left out by the researcher. Again, the effects of both variables have been rolled into the coefficient on “Peak,” which is now biased upwards.

Columns 2 through 4 thus show how having the policy change as well as the peak control variable are the crucial components for identification. The second half of the sample allows the researcher to estimate the “Peak” effect, which can then be controlled for to allow the researcher to estimate the “Treatment” effect. This is comparable to our main specification in Tables 1 and 2, for which we observe a policy change mid-way through the sample, and where we know (both based on policy documentation and what we observe in the data itself) that the treatment variable is a direct multiplier of the peak variable.

Finally, Column 5 illustrates that identification is not achieved via just a simple pre/post comparison. In this example, the mean level of the treatment variable has not changed from the pre-period to the post-period, and indeed a simple regression on a post-period dummy would not uncover the coefficient of interest.

Figure A1: Treatment Variable for Simulation of Identification Strategy



Note: This figure shows the treatment variable constructed for the simulation exercise. It is normally distributed with mean zero and standard deviation equal to one for the first part of the sample, and it is equal to zero for the second part of the sample.

Table A3: Simulated Dataset to Illustrate the Identification Strategy

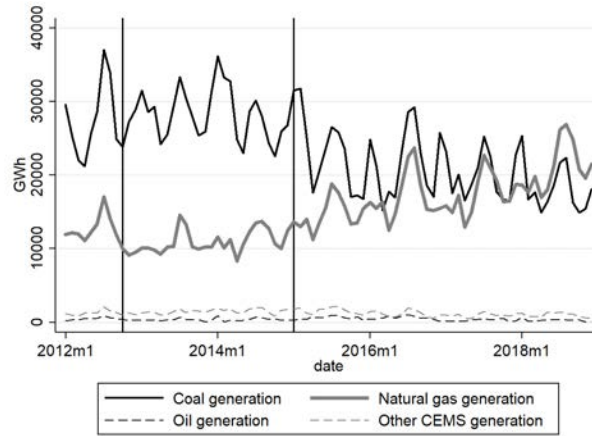
	Correctly specified	First half	Second half	Dropping peak control	Only dummy
Treatment	0.99*** (0.01)			1.99*** (0.01)	
Peak	1.01*** (0.01)	1.99*** (0.01)	1.01*** (0.01)		
Post dummy					-0.01 (0.03)
Constant	1.01*** (0.01)	1.02*** (0.01)	1.01*** (0.01)	1.01*** (0.01)	1.01*** (0.02)
Observations	20,000	10,000	10,000	20,000	20,000
R ²	0.72	0.80	0.50	0.57	0.00

Note: This table shows five regressions using a simulated dataset constructed by the researchers. The true data-generating process is $y_t \equiv 1 + 1 \cdot \text{peak}_t + 1 \cdot \text{treatment}_t + \varepsilon_t$, where “peak” and ε are each normally distributed with mean zero and standard deviation equal to one. The “treatment” variable is equal to “peak” in the first half of the sample and equal to zero in the second half. The first Column includes the entire sample and both the Treatment variable of interest and the Peak control variable; it correctly uncovers all three coefficients. This column mimics the identification strategy used in the main text for the impact of the regulation requirement on the generation mix. Dropping either half of the data or not including the “Peak” control leads to mis-specification, as shown in the second through fourth columns. Estimation using a only a post-period dummy is not possible, as the mean value of the Treatment variable is constant throughout the sample, as shown in the last column. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

A1.3 Additional Tables and Figures

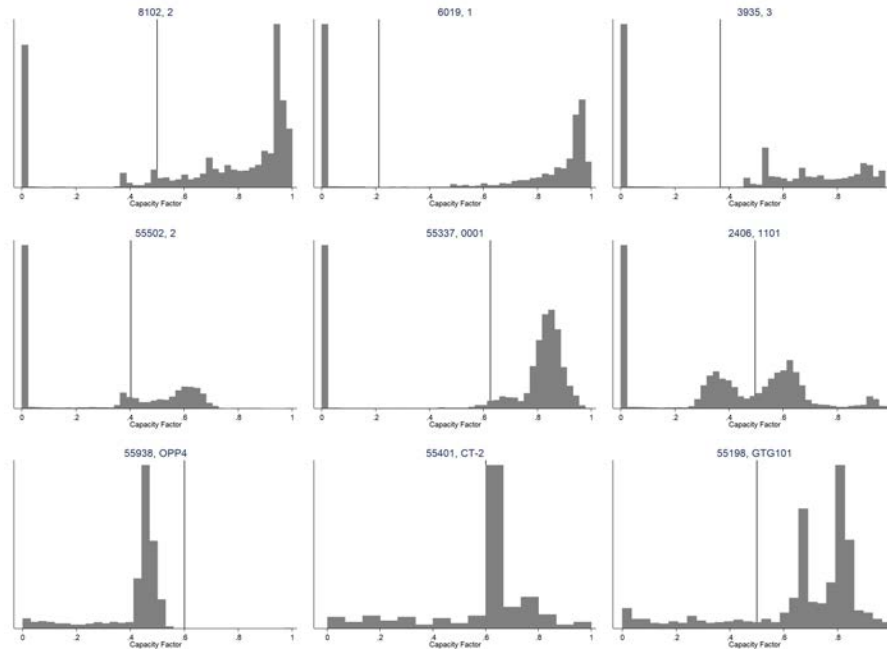
This section contains additional tables and figures referenced in the text, including summary statistics, robustness checks, etc.

Figure A2: Monthly Generation



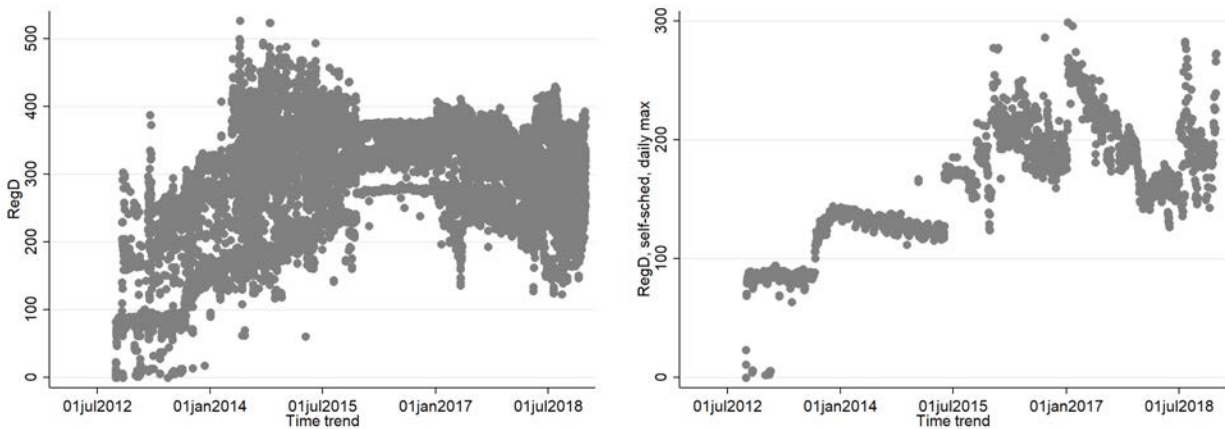
Note: This figure shows monthly generation by unit type for PJM units that appear in CEMS data. Vertical lines display our primary sample window (October 2012 through December 2014).

Figure A3: Minimum Constraints



Note: These are nine units at large plants, three for each technology type. The top row shows three coal-fired boilers, with the plant id and unit id given at the top of each histogram. The second row shows natural gas combined cycle plants, and the bottom row shows natural gas combustion turbines. In the bottom row, zeros are not displayed – because CTs operate infrequently, displaying zeros makes it difficult to visualize the non-zero portion of the histogram. Vertical lines are placed at the minimum operating constraint constructed from EIA data (which in some cases is a plant-level proxy, rather than measured at the individual unit level - that may be why some panels appear to show measurement error). See Appendix text for details.

Figure A4: RegD



Note: The left-hand panel show the total amount of self-scheduled and pool-procured RegD at the hourly level. While discrete jumps (perhaps from policy changes or from battery entry) and a general trend upwards (consistent with increasing battery deployment) are observed, it is clear that total RegD is not driven solely by the contribution of new batteries. The right-hand panel plots the daily maximum MW of self-scheduled RegD, for which the discrete jumps are clearer. Data source is PJM.

Table A4: Summary Statistics

	Mean	Std. Dev.	N
Regulation requirement, 100 MW	6.78	1.06	19693
Generation, by technology:			
Boilers, MWh	40184.7	8233.9	19728
Combined cycle, MWh	14289.3	3635.1	19728
Combustin turbine, MWh	847.9	1923.4	19728
Other unit types in CEMS, MWh	553.4	294.7	19728
Generation, by fuel type:			
Coal generation, MWh	39818.1	7934.4	19728
Natural gas generation, MWh	14906.0	4836.5	19728
Oil generation, MWh	570.1	699.0	19728
Other fuel types in CEMS, MWh	581.0	295.1	19728
CEMS CO2 emissions, tons	45215.2	9378.2	19728
Wind generation, MWh	1721.4	1118.9	19723
Generation not in CEMS, MWh	34044.3	4568.1	19728
PJM load, MWh	89919.5	15730.8	19728
Peak forecast, in peak hours, MWh	104021.8	15040.0	15618
Valley forecast, in off-peak hours, MWh	74386.2	10460.2	4106
Coal steam retirements, MW	2536.5	2027.1	19728
Natural gas CT retirements, MW	38.6	65.2	19728
Other retirements, MW	412.7	305.0	19728
Henry hub natural gas price	3.97	0.64	19728
Oil price	94.7	10.3	19728
Cooling degree days in Philadelphia	3.15	5.29	19728
Heating degree days in Philadelphia	13.1	13.5	19728

Note: Data cover the period October 1, 2012 through December 31, 2014. Unit of observation is one hour. Data sources: PJM, EPA, and EIA. Peak and valley forecasts apply only in the peak (4 a.m. to midnight) and valley (midnight to 4 am) hours, respectively. A small number of observations (<1%) are missing for the regulation requirement, wind generation, and peak/valley forecast variables.

Table A5: Displaying Control Coefficients: The Impact of the Regulation Requirement on the Energy Market

	Boiler	CC	CT	Other
Regulation requirement, 100 MW	-363.77** (180.17)	390.50** (149.86)	-0.08 (184.69)	-26.65 (19.95)
PJM load, MWh	-0.17*** (0.02)	0.02 (0.02)	0.16*** (0.02)	-0.00 (0.00)
CEMS units generation, MWh	0.77*** (0.02)	0.23*** (0.02)	-0.01 (0.01)	0.01** (0.00)
Peak forecast, in peak hours, MWh	0.03*** (0.01)	-0.01 (0.01)	-0.02*** (0.01)	0.00** (0.00)
Valley forecast, in off-peak hours, MWh	0.04*** (0.01)	-0.00 (0.01)	-0.04*** (0.01)	0.00* (0.00)
Coal steam retirements, MW	-0.21 (0.45)	0.12 (0.38)	-0.03 (0.19)	0.12** (0.06)
Natural gas CT retirements, MW	45.04* (24.70)	-61.24*** (21.54)	12.36 (10.66)	3.84 (2.79)
Other retirements, MW	-16.90** (7.15)	21.02*** (6.09)	-3.23 (3.29)	-0.89 (0.84)
Time trend	13.08*** (3.92)	-16.34*** (3.65)	2.50 (1.83)	0.76 (0.49)
Time trend, quadratic (centered)	-0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	-0.00 (0.00)
Primary reserve req., RTO-wide, MW	1.27 (0.94)	-0.14 (0.93)	-1.34*** (0.44)	0.22* (0.11)
Synchronized reserve req., RTO-wide, MW	0.04 (0.62)	-0.25 (0.59)	0.27 (0.22)	-0.06 (0.06)
Primary reserve req., MAD sub-zone, MW	-0.05 (1.83)	-2.94 (2.04)	3.29*** (0.87)	-0.31** (0.15)
Synchronized reserve req., MAD sub-zone, MW	-0.76 (1.25)	2.31 (1.47)	-1.67*** (0.59)	0.11 (0.11)
Dummy, equal to one beginning December 1, 2013	-1541.75 (1536.36)	2485.22* (1340.05)	-618.14 (775.01)	-325.33 (211.25)
Henry hub price	714.91*** (265.76)	-788.21*** (246.78)	134.91 (135.15)	-61.61** (24.44)
WTI price	18.89 (23.33)	-31.10 (19.03)	8.64 (11.69)	3.57 (3.38)
Cooling degree days in Philadelphia	6.41 (24.07)	-26.92 (19.73)	21.63 (16.57)	-1.11 (3.46)
Heating degree days in Philadelphia	20.63** (10.32)	6.57 (9.96)	-29.14*** (9.25)	1.94 (2.00)
Observations	19,680	19,680	19,680	19,680
Within R ²	0.92	0.71	0.49	0.56

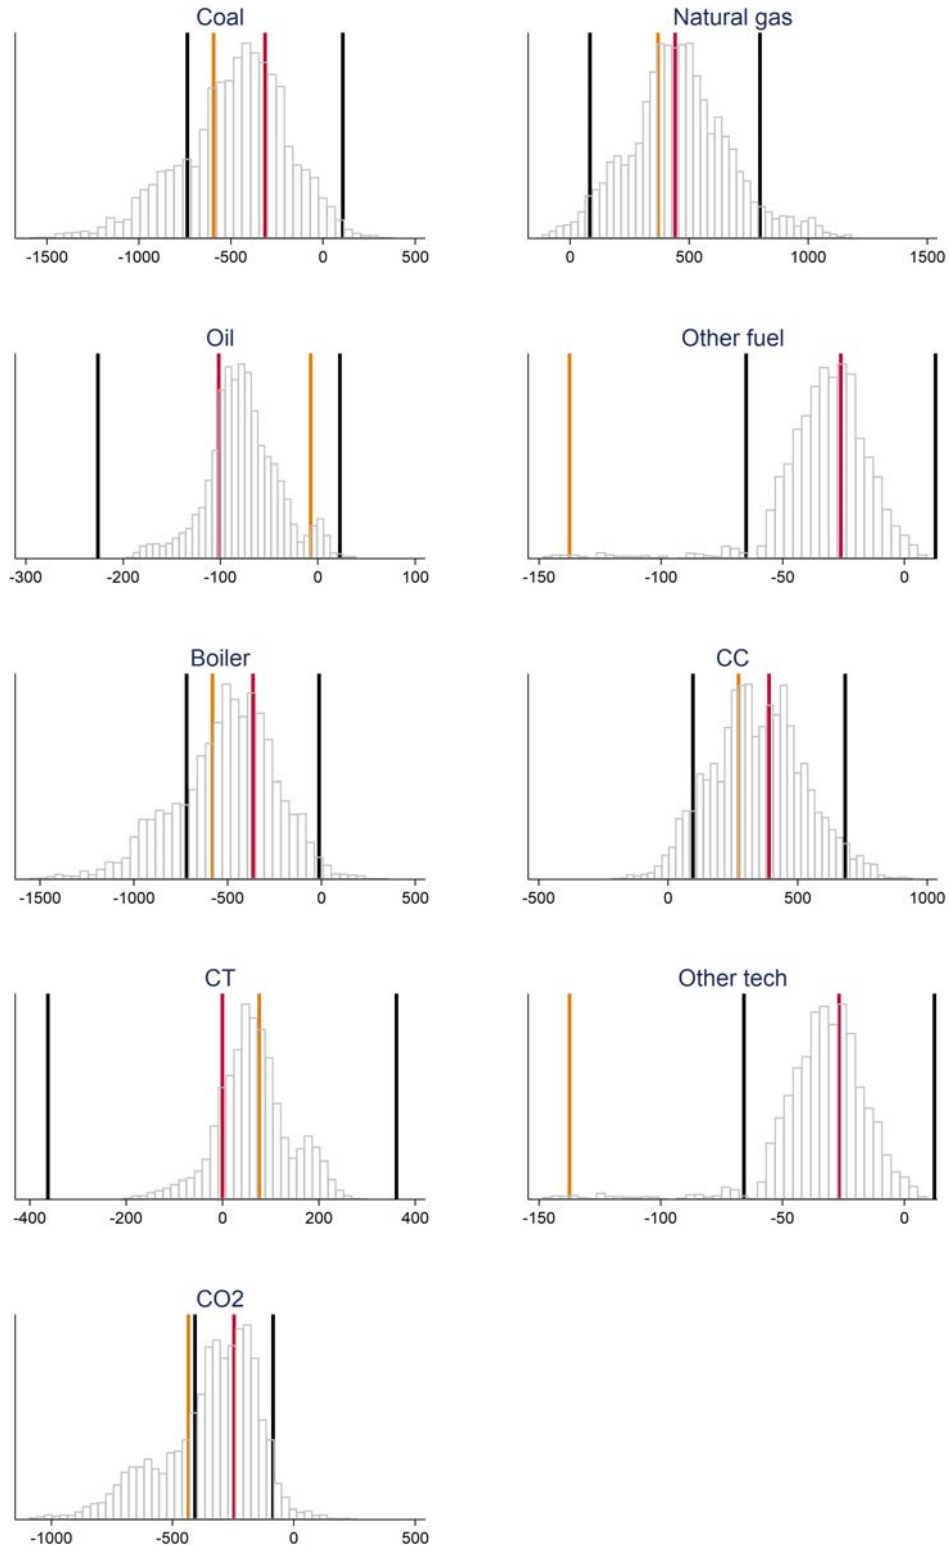
Note: This table shows coefficients on the control variables for the regression results shown in the main text in Table 1. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A6: Displaying Control Coefficients: The Impact of the Regulation Requirement on the Energy Market

	Coal	NG	Oil	Other	CO2
Regulation requirement, 100 MW	-313.87 (215.79)	441.35** (181.94)	-101.43 (63.50)	-26.05 (19.86)	-246.10*** (81.83)
PJM load, MWh	-0.21*** (0.02)	0.17*** (0.02)	0.04*** (0.01)	-0.00 (0.00)	-0.06*** (0.01)
CEMS units generation, MWh	0.78*** (0.03)	0.22*** (0.02)	-0.01* (0.01)	0.01** (0.00)	0.80*** (0.01)
Peak forecast, in peak hours, MWh	0.02 (0.01)	-0.03*** (0.01)	0.01*** (0.00)	0.00** (0.00)	0.03*** (0.00)
Valley forecast, in off-peak hours, MWh	0.02** (0.01)	-0.04*** (0.01)	0.01*** (0.00)	0.00* (0.00)	0.03*** (0.01)
Coal steam retirements, MW	-0.31 (0.47)	0.16 (0.46)	0.04 (0.11)	0.12** (0.06)	-0.07 (0.25)
Natural gas CT retirements, MW	35.04 (25.49)	-46.36* (24.86)	7.51 (5.15)	3.81 (2.80)	35.74*** (13.13)
Other retirements, MW	-14.50* (7.50)	17.16** (7.28)	-1.78 (1.52)	-0.87 (0.84)	-12.52*** (3.94)
Time trend	12.85*** (3.98)	-13.52*** (3.99)	-0.09 (0.84)	0.76 (0.49)	8.32*** (2.21)
Time trend, quadratic (centered)	-0.01 (0.00)	0.01 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
Primary reserve req., RTO-wide, MW	1.73* (0.94)	-1.81* (1.01)	-0.13 (0.18)	0.21* (0.11)	0.90* (0.53)
Synchronized reserve req., RTO-wide, MW	-0.02 (0.62)	0.14 (0.67)	-0.06 (0.10)	-0.05 (0.06)	-0.35 (0.34)
Primary reserve req., MAD sub-zone, MW	-1.18 (1.61)	0.62 (1.80)	0.88* (0.47)	-0.31** (0.15)	0.49 (0.95)
Synchronized reserve req., MAD sub-zone, MW	-0.15 (1.08)	0.49 (1.27)	-0.46 (0.35)	0.11 (0.11)	-0.47 (0.67)
Dummy, equal to one beginning December 1, 2013	-1605.04 (1609.00)	1790.00 (1550.62)	142.99 (351.41)	-327.94 (211.12)	-1484.04* (865.12)
Henry hub price	755.19*** (260.94)	-633.47** (250.87)	-62.11 (72.25)	-59.61** (24.75)	596.01*** (152.98)
WTI price	26.45 (25.05)	-27.28 (23.64)	-2.80 (5.04)	3.63 (3.38)	18.61 (12.38)
Cooling degree days in Philadelphia	-22.81 (26.20)	5.04 (24.65)	18.77*** (6.86)	-1.00 (3.45)	28.72** (13.07)
Heating degree days in Philadelphia	28.78** (11.71)	-25.59** (10.70)	-5.10 (3.94)	1.92 (2.00)	6.43 (6.47)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.90	0.82	0.46	0.56	0.98

Note: This table shows coefficients on the control variables for the regression results shown in the main text in Table 2. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Figure A5: Robustness to Fewer Controls: The Impact of the Regulation Requirement on the Energy Market



Note: These figures show histograms of the estimated treatment effect of interest when various covariates are dropped, for comparison with the regressions displayed in the main text in Tables 1 and 2. All specifications control for peak and valley forecasts, but we loop over every other combination of control variables. To simplify somewhat, we group together some of the covariates (for instance, if we include cooling degree days, we also include heating degree days). We estimate nearly 25,000 specifications for each outcome variable. Red lines display the estimates from the main text, Tables 1 and 2. Black lines show the 95 percent confidence intervals from those specifications. Orange lines show the estimate from the sparsest possible specification, which controls only for peak and valley forecasts.

Table A7: Robustness to Alternative Specifications: The Impact of the Regulation Requirement on the Energy Market

Panel A. Boiler																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-156	-471***	-359**	-374**	-344*	-284	-425***	-382**	-352**	-336*	-471***	-364*	-394**	-353*	-388**	-457**
	(151)	(176)	(179)	(159)	(175)	(174)	(148)	(177)	(173)	(202)	(177)	(193)	(180)	(184)	(183)	(182)
Panel B. Combined Cycle																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	463***	487***	379**	391***	404***	296**	410***	462***	424***	352**	391**	391***	409***	326**	366**	407***
	(139)	(152)	(147)	(140)	(150)	(136)	(126)	(161)	(154)	(169)	(150)	(148)	(153)	(144)	(149)	(155)
Panel C. Combustion Turbines																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-277**	17	7	8	-34	10	29	-57	-77	14	4	-0	13	6	2	74
	(107)	(171)	(187)	(99)	(171)	(190)	(118)	(188)	(181)	(173)	(181)	(197)	(186)	(191)	(185)	(187)
Panel D. Other Units																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-30	-33	-26	-25	-26	-22	-14	-23	6	-29	-23	-27	-27	21***	20***	-25
	(21)	(21)	(20)	(21)	(20)	(19)	(21)	(22)	(22)	(25)	(20)	(20)	(21)	(7)	(6)	(20)
Panel E. Coal																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-61	-396*	-318	-346*	-282	-244	-376**	-288	-257	-306	-415*	-314	-349	-301	-339	-422*
	(173)	(208)	(216)	(178)	(207)	(214)	(173)	(211)	(204)	(241)	(211)	(229)	(216)	(225)	(218)	(218)
Panel F. Natural Gas																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	255	538***	441**	459***	417**	361**	488***	446**	387**	417**	440**	441**	478***	377**	419**	527***
	(157)	(177)	(181)	(164)	(177)	(176)	(152)	(179)	(172)	(206)	(180)	(192)	(182)	(181)	(180)	(185)
Panel G. Oil																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-164***	-110*	-98	-88*	-109*	-96	-98*	-136**	-136**	-82	-101	-101	-102	-97	-100	-81
	(55)	(62)	(64)	(51)	(62)	(64)	(55)	(63)	(60)	(73)	(63)	(65)	(64)	(69)	(63)	(64)
Panel H. Other Units																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-30	-32	-26	-25	-25	-21	-14	-22	6	-28	-23	-26	-27	21***	20***	-24
	(21)	(21)	(20)	(21)	(20)	(19)	(21)	(22)	(22)	(25)	(20)	(20)	(20)	(7)	(6)	(20)
Panel I. CO2 Emissions																
Reg. req.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	-208**	-266***	-243***	-237***	-252***	-184**	-264***	-287***	-253***	-247**	-345***	-246***	-271***	-278***	-339***	-268***
	(82)	(84)	(81)	(84)	(82)	(73)	(76)	(82)	(77)	(98)	(79)	(84)	(82)	(87)	(96)	(83)

Note: This table shows alternative specifications for the regressions displayed in the main text in Tables 1 and 2, using various alternative controls, variable definitions, and sub-samples. Column 1 limits the sample to time periods with overlap in the peak/valley forecasts between the pre and post-policy change periods. Column 2 uses a linear, rather than quadratic, time trend. Column 3 uses a cubic time trend. Column 4 uses a spline with five knots, rather than a linear control, for PJM-wide load. Column 5 adds weather controls for Chicago. Column 6 uses a spline with five knots, rather than a linear control, for the Henry Hub price. Column 7 adds flexible (binned) controls for PJM-wide load, CEMS generation, and the Henry Hub price. Column 8 controls for the standard deviation of the regulation requirement over the previous 72 hours. Column 9 controls for the standard deviation of each of the regulation requirement, PJM-wide load, CEMS generation, and the peak and valley forecasted load over the previous 72 hours. Column 10 collapses to the daily level. Column 11 restricts the sample to units that do not retire. Column 12 calculates Newey-West standard errors with a maximum lag length of 168 hours (one week). Column 13 uses a constructed regulation requirement variable as an instrument for the reported regulation requirement. Column 14 uses CEMS-reported gross generation rather than re-scaled net generation. Column 15 uses only CEMS-reported electrical generation, rather than also incorporating steam load in the net generation scaling. Column 16 uses PJM variables in their raw form, i.e., in Eastern Prevailing Time rather than Eastern Standard Time. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A8: Alternative CO₂ Measurement: The Impact of the Regulation Requirement on the Energy Market

Panel A. CO₂ Variable in Main Text, Metric Tons					
	Coal	NG	Oil	Other	Total
Regulation requirement, 100 MW	-331* (195)	156 (112)	-76 (51)	5 (7)	-246*** (82)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.90	0.79	0.44	0.43	0.98
Panel B. CO₂ as Reported, Metric Tons					
Regulation requirement, 100 MW	-336* (194)	128 (112)	-71 (48)	1 (7)	-278*** (83)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.90	0.78	0.45	0.47	0.98
Panel C. Using Unit-Level Emissions Rates, Metric Tons					
Regulation requirement, 100 MW	-333* (193)	185* (108)	-71 (50)	-0 (7)	-219*** (80)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.90	0.79	0.45	0.43	0.98

Note: Panel A shows the CO₂ emissions results by fuel type (Columns 1 through 4) and aggregated (Column 5), matching the specifications used in the main text, Table 2. Panel B shows analogous specifications, but using CEMS-reported CO₂ emissions (which are occasionally missing) rather than emissions constructed from the heat input variable. Panel C uses the unit-level emissions rate for all hours, not just hours with missing CO₂ data. All panels are reported in metric tons (i.e., in Panel B we convert CEMS-reported short tons into metric tons). ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A9: Placebo and Residual Units: The Impact of the Regulation Requirement on the Energy Market

	Non-PJM Coal	Non-PJM NG	Non-PJM Oil	Non-PJM Other	PJM Comm+Ind	PJM Wind	PJM Residual
Regulation requirement, 100 MW	310.1 (211.5)	25.9 (141.2)	-15.5 (17.3)	-2.7 (5.0)	-7.1 (8.0)	-125.9* (75.1)	243.9 (185.2)
Load controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other reserves controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Retirement controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quadratic time trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fuel price controls	Yes	Yes	Yes	Yes	Yes	Yes	No
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes	No
Month, day of week, and hour effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	19,680	19,680	19,680	19,680	19,680	19,679	19,680
Within R ²	0.71	0.68	0.14	0.32	0.32	0.18	0.53

Note: This table shows estimates from seven separate regressions, analogous to those presented in the main text, Tables 1 and 2. The dependent variable in the first four columns is MWh of electricity generated per hour for the electrical generating units that are located in PJM states but are *not* part of PJM; see footnote 30. The dependent variable in the fifth column is MWh of electricity generated by commercial and industrial units in PJM. The dependent variable in the sixth column is MWh of wind generation in PJM. The dependent variable in the seventh column is the difference between PJM-wide demand and the generation reported by electrical generating units in CEMS; this accounts for fuel types not in CEMS (nuclear, wind, solar, etc.), small units not in CEMS, and net imports. The unit of analysis is an hour. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A10: Specifications with Two-Way Fixed Effects

Panel A. Time Effects, Location Effects, and Peak/Valley Forecast Controls									
	Boiler	CC	CT	Other tech	Coal	NG	Oil	Other fuel	CO2
Reg. req., 100 MW, in PJM	-478 (374)	331* (190)	80 (151)	-125*** (27)	-471 (381)	371 (233)	20 (81)	-111*** (28)	-256 (370)
Observations	39,388	39,388	39,388	39,388	39,388	39,388	39,388	39,388	39,388
Within R ²	0.22	0.38	0.10	0.24	0.16	0.41	0.28	0.19	0.35
Panel B. Plus Additional Controls									
	Boiler	CC	CT	Other tech	Coal	NG	Oil	Other fuel	CO2
Reg. req., 100 MW, in PJM	-230 (145)	285* (171)	-45 (119)	-9 (18)	-111 (155)	245 (151)	-127* (65)	-7 (20)	-124 (114)
Observations	39,360	39,360	39,360	39,360	39,360	39,360	39,360	39,360	39,360
Within R ²	0.83	0.67	0.27	0.61	0.82	0.74	0.49	0.54	0.94

Note: Panel A shows specifications in which plants in nearby states (see footnote 30) serve as controls. The unit of observation is an hour in a region (PJM, or nearby states grouped together). The variable of interest takes on the value of the regulation requirement in PJM, and a value of zero in nearby states, as the regulation requirement does not directly affect them. Controls are: hour-of-sample effects, region effects, and two interaction variables. These latter controls are (1) peak forecasted load interacted with a PJM dummy, and valley forecasted load interacted with a PJM dummy; these may be important for removing bias as the regulation requirement is a direct function of these forecasts in the first half of the sample. Panel B includes the same controls, but also adds all the controls from the main specification in Tables 1 and 2 interacted with a PJM dummy for additional precision. As shown when moving from Panel A to Panel B, these additional controls aid with precision even in the two-way fixed effects specification. This is intuitive if the response to control variables such as the natural gas price varies between regions, in which case hour of sample effects will not fully account for the natural gas price effect across the two regions. This limitation of the two-way fixed effects specification – it is still aided by additional control variables – combined with the fact that identification is less transparent than it is in the time-series regression, is why our primary specification in the main text is a time-series regression. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A11: Four-Unit Model

Regulation Requirement	Generation in Equilibrium				Regulation in Equilibrium			
	Unit A, \$35/MWh	Unit B, \$37/MWh	Unit C, \$40/MWh	Unit D, \$60/MWh	Unit A, \$5/MW	Unit B, \$6/MW	Unit C, \$7/MW	Unit D, \$10/MW
500	25000	24550	0	450	0	450	0	50
550	25000	24525	0	475	0	475	0	75
599	25000	24500.5	0	499.5	0	499.5	0	99.5
600	25000	24500	0	500	0	500	0	100
601	24999.5	24500	0	500.5	0.5	500	0	100.5
602	24999	24500	0	501	1	500	0	101
610	24995	24500	0	505	5	500	0	105
615	24992.5	24500	0	507.5	7.5	500	0	107.5
616	24884	21116	4000	0	116	500	0	0
620	24880	21120	4000	0	120	500	0	0
624	24876	21124	4000	0	124	500	0	0
625	24875	21125	4000	0	125	500	0	0
650	24850	21150	4000	0	150	500	0	0
675	24825	21175	4000	0	175	500	0	0
700	24800	21200	4000	0	200	500	0	0
750	24750	21250	4000	0	250	500	0	0
800	24700	21300	4000	0	300	500	0	0
850	24650	21350	4000	0	350	500	0	0
900	24600	21400	4000	0	400	500	0	0
950	24550	21450	4000	0	450	500	0	0
999	24501	21499	4000	0	499	500	0	0
1000	24500	21500	4000	0	500	500	0	0
1001	24500	21499	4001	0	500	500	1	0
1050	24500	21450	4050	0	500	500	50	0
1100	24500	21400	4100	0	500	500	100	0
1150	24500	21350	4150	0	500	500	150	0
1200	24500	21300	4200	0	500	500	200	0

Note: Table A11 lists the equilibrium results for a four-unit model with energy and regulation output. Units A, B, and C face a maximum capacity of 25,000 MW each. They also face a minimum constraint, when generating, of 4,000 MW. Unit D faces a maximum capacity of 25,000 MW and a minimum when generating of 400 MW. (This lower minimum operational constraint is meant to represent the fact that the peaking portion of the electricity market is made up of many small peaker units that can each be dispatched at quite small levels of generation.) Marginal costs of energy and regulation provision are listed in the table. For both services, Unit A is cheapest, Unit B next cheapest, etc. Energy demand is held constant at 50,000 MWh, while the regulation requirement varies exogenously across rows. The equilibrium is found using the online tool <https://online-optimizer.appspot.com/>. We check whether results are global, not just local, solutions by forcing individual units on or off, finding that alternative solutions do not achieve a lower system-wide cost. We further explore whether the solutions are unique (as opposed to, e.g., having a flat objective function) by imposing additional constraints forcing an individual unit's generation or regulation to be $\epsilon = 0.001$ higher than the optimal solution in an effort to find other equal-cost solutions – however, for all cases we explored, doing so yields a higher total system cost (or no feasible solution) indicating that the reported solutions are likely unique.

Table A12: Showing Other CEMS Units: The Regulation Requirement and Intensive/Extensive Margins

Panel A. Boilers					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	0.30 (0.97)	0.16 (0.15)	-0.09 (0.35)	0.66 (0.81)	-1.03** (0.39)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.77	0.17	0.41	0.50	0.79
Panel B. Combined Cycle Plants					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-1.71** (0.80)	-0.46* (0.27)	0.13 (0.21)	2.40*** (0.65)	-0.36* (0.21)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.65	0.28	0.13	0.62	0.35
Panel C. Combustion Turbines					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-1.63 (2.98)	-0.04 (0.42)	0.29 (0.29)	2.55 (2.13)	-1.17** (0.49)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.48	0.31	0.28	0.45	0.22
Panel D. Other Units					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	0.34** (0.16)	-0.15*** (0.06)	-0.09*** (0.02)	-0.17 (0.13)	0.07 (0.04)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.58	0.06	0.07	0.59	0.08

Note: This table expands on Table 4 by showing results at other units. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A13: Alternative Minimum Constraints Data: The Regulation Requirement and Intensive/Extensive Margins

Panel A. Boiler					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	0.30 (0.97)	-0.84* (0.45)	-0.19 (0.62)	1.76** (0.88)	-1.03** (0.39)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.77	0.36	0.55	0.69	0.79
Panel B. Combined Cycle					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-1.71** (0.80)	-0.05 (0.32)	-0.04 (0.38)	2.16*** (0.72)	-0.36* (0.21)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.65	0.22	0.19	0.61	0.35
Panel C. Combustion Turbine					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-1.63 (2.98)	-0.01 (0.19)	0.09 (0.28)	2.73 (2.38)	-1.17** (0.49)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.48	0.21	0.42	0.44	0.22
Panel D. Other Tech					
	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	0.34** (0.16)	-0.16** (0.06)	0.18*** (0.06)	-0.43*** (0.14)	0.07 (0.04)
Observations	19,680	19,680	19,680	19,680	19,680
Within R ²	0.58	0.07	0.19	0.55	0.08

Note: This table is analogous to Table 4, but uses an alternative variable to construct the minimum constraint. Rather than EIA-reported minimum constraints, it uses the smallest bin with at least 5 percent of non-zero generating hours. Note this alternative definition does not impact the “off” or “at max” counts. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A14: Bins: The Regulation Requirement and Intensive/Extensive Margins

Panel A. Boiler											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	0.18 (1.00)	0.06 (0.10)	0.01 (0.14)	-0.04 (0.13)	-0.55* (0.29)	-0.06 (0.32)	-0.05 (0.33)	0.41 (0.27)	0.28 (0.25)	0.21 (0.37)	-0.57 (0.72)
Observations	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680
Within R ²	0.77	0.06	0.09	0.26	0.47	0.38	0.23	0.12	0.10	0.40	0.84

Panel B. Combined Cycle											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-1.69** (0.80)	-0.16*** (0.03)	-0.08*** (0.02)	0.03 (0.03)	-0.27*** (0.08)	-0.12 (0.13)	-0.10 (0.27)	1.18*** (0.35)	-0.41 (0.40)	1.44*** (0.42)	0.21 (0.41)
Observations	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680
Within R ²	0.64	0.06	0.05	0.07	0.16	0.10	0.21	0.14	0.09	0.56	0.49

Panel C. Combustion Turbine											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-1.75 (2.97)	-0.08 (0.12)	0.01 (0.09)	0.06 (0.08)	0.05 (0.10)	0.23 (0.20)	0.45* (0.26)	0.57 (0.39)	1.67** (0.84)	0.24 (0.73)	-1.57** (0.66)
Observations	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680
Within R ²	0.48	0.15	0.15	0.14	0.15	0.23	0.27	0.29	0.40	0.42	0.27

Panel D. Other Tech											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	0.34** (0.16)	-0.04 (0.03)	-0.03 (0.02)	-0.12*** (0.03)	0.01 (0.05)	0.24*** (0.07)	-0.07 (0.05)	-0.16** (0.06)	-0.38*** (0.09)	-0.08 (0.07)	0.29*** (0.09)
Observations	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680	19,680
Within R ²	0.58	0.02	0.03	0.11	0.06	0.11	0.08	0.27	0.40	0.21	0.23

Note: This table is analogous to Table 4, but rather than using data on minimum constraints, it simply counts the number of units generating at 0 percent of capacity, 0 to 10 percent of capacity, etc. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A15: Regressing Generation on Alternative Battery Capacity Variables, by Fuel Type

	RegD Self-Scheduled		EIA Battery Capacity		DOE Battery Capacity	
	PJM	Placebo	PJM	Placebo	PJM	Placebo
Coal	162.46 (673.65)	965.36 (948.95)	-1457.75* (767.03)	-2965.16*** (856.35)	-281.40 (697.80)	-2570.32*** (843.37)
NG	-169.78 (673.30)	-353.23 (387.05)	1457.64* (763.78)	1977.91*** (324.70)	269.79 (692.57)	1133.37*** (351.59)
Oil	55.93 (129.53)	-12.38 (35.03)	80.38 (223.96)	70.88 (50.70)	54.88 (191.00)	36.13 (45.05)
Other	-46.89 (130.98)	28.71 (62.22)	-27.27 (201.61)	-46.30 (92.23)	-113.16 (175.38)	-1.11 (65.37)

Note: This table shows estimates from 24 separate regressions. The dependent variable is CEMS generation. In Columns 1 and 2, the independent variable of interest is Self-Scheduled RegD, in 100 MWs. In Columns 3 and 4, the independent variable of interest is EIA-reported battery capacity, in 100 MWs. In Columns 5 and 6, the independent variable of interest is battery capacity, in 100 MWs, from a DOE storage database. Columns 1, 3, and 5 use our PJM sample of interest. Columns 2, 4, and 6 use a placebo sample: non-PJM electrical generating units in PJM states (e.g., MISO units in Indiana; see footnote 30). The sample is limited to the years 2014-2016. The unit of analysis is an hour. Control variables are the same as for Table 2. Standard errors are clustered at the monthly level, which is the level of variation for the EIA and DOE battery variables. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.