A1 Online Appendix

For Buchsbaum, Hausman, Mathieu, and Peng: "Spillovers from Ancillary Services to Wholesale Energy Markets"

This Appendix provides additional details on the stylized four-unit model, on the data used, and on the regression coefficients. It also provides additional robustness checks for the regression results.

Table A1: Four-Unit Model

	Generation in Equilibrium				Regulation in Equilibrium			
Regulation Requirement	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 A1 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, and 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 $\varepsilon = 0.001$ 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.

A1.1 Data Appendix

A1.1.1 Data Sources for Control Variables

From PJM, we observe total hourly electricity demand, in MWh. We also collect PJM data on the forecasted peak and valley demand for each day.

From the Energy Information Administration , we observe the daily price of natural gas (measured at Henry Hub), the daily price of oil (West Texas Intermediate), and the monthly price of coal paid by power plants. In regressions with month-of-sample effects, the coal price drops out. Also, we observe daily average temperature at the Philadelphia airport (degrees Fahrenheit, from NOAA), a relatively central location within PJM, which we use to calculate cooling and heating degree days. In some regressions we add the daily average temperature in Chicago, also from NOAA.

For some robustness checks, we include the hourly requirement (in MW) for other types of ancillary services: synchronized and non-synchronized 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 necessary for identification.

In one placebo specification, we use hourly data on wind generation (MWh, from PJM). For other placebo tests, we collect hourly generation and CO_2 emissions data at the generator level for two other types of units: CEMS fossil-fuel-fired units not in PJM but in nearby states; and units in CEMS data that that are not categorized by CEMS as electrical generating units (e.g., refineries).

A1.1.2 CO2 Data

The CEMS-reported CO_2 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. Below, we show alternative results using CEMS-reported CO_2 emissions.

A1.1.3 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. Although 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, because 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, 2022). 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).

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, 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.

A1.1.4 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 A15.)

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%).

A1.1.5 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 tangentiallyfired 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 A2.

Table A2: 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.

A1.1.6 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 A2 and A3. The difference between CEMS and EIA data is accounted for by the "residual" generation variable we construct, equal to the 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.

A1.1.7 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

Table A3: Total Annual Generation by Fuel Type, EIA Data, 2012-2014

Unit Type	PJM Generation, TWh
Coal	336
Nuclear	276
Natural Gas	143
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	0.8
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.

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 are similar if one uses the raw data, mixing EPT, CST, and EST across variables and plants.

A1.1.8 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 (19,693 hours).

A1.2 Simulated Dataset to Illustrate the Identification Strategy

As is shown in Figure 4 in the main text, and discussed in Section 5.1, 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 2 and 3. We construct a peak forecast variable, normally distributed with mean zero and standard deviation equal to one.[50] 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 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 A4 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. Column 2 shows that adding a post dummy does not affect our ability to estimate the treatment and peak effects; this is relevant as some of our specifications include time effects.

Columns 3 through 6 illustrate how identification is achieved by displaying specifications that are *not* identified. In Column 3, 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 4, 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 5, 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 3 through 5 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

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

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 2 and 3, 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 6 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

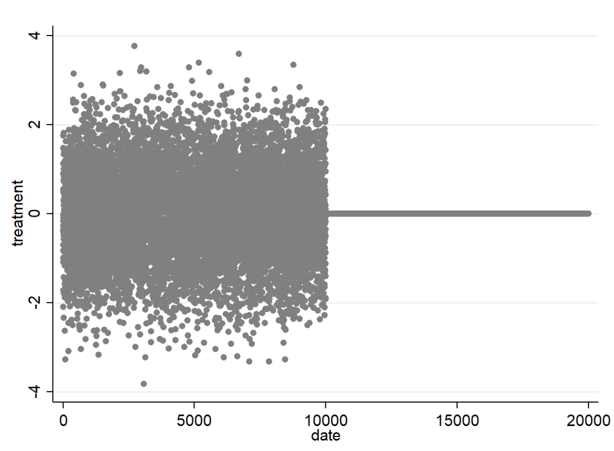

Note: This figure shows the treatment variable constructed for the simulation exercise. It is normally distributed with mean zero and standard deviation equal to one for the first part of the sample, and it is equal to zero for the second part of the sample.

Table A4: Simulated Dataset to Illustrate the Identification Strategy

	Correctly specified	Alternative	First half	Second half	Dropping peak control	Only dummy
Treatment	0.99***	0.99***			1.99***	
	(0.01)	(0.01)			(0.01)	
Peak	1.01***	1.01***	1.99***	1.01***		
	(0.01)	(0.01)	(0.01)	(0.01)		
Post dummy		-0.01				-0.01
		(0.01)				(0.03)
Constant	1.01***	1.02***	1.02***	1.01***	1.01***	1.01***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Observations	20,000	20,000	10,000	10,000	20,000	20,000
R^2	0.72	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. The second column shows that including a post-period dummy does not affect our ability to recover the treatment and peak effects. However, dropping either half of the data or not including the "Peak" control leads to mis-specification, as shown in the third through fifth 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. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

A1.3 Robustness Checks

Tables A8, A9, and A10 show additional robustness checks, mentioned briefly in the main text. Here we describe the rationale for examining these alternative specifications.

First, we examine more parsimonious specifications, dropping various control variables. In the main text, we argue that peak and valley forecast controls are necessary, as they directly impact the regulation requirement in some hours and as they may also directly impact generator behavior. We also argue that additional controls, common in the literature, may similarly be correlated with both the regulation requirement and generator behavior (for instance, weather). However, might worry about "oversaturation" in our regressions (for instance, that controlling for so many things leaves only measurement error, as in Fisher et al. 2012). In these parsimonious regression, we continue to control for peak and valley forecasts (and a peak versus valley hour dummy), following the logic of Section A1.2, but we drop all other control variables. These results provide reassurance that the effects we estimate are not driven by the inclusion of too many controls. The estimates from the parsimonious regression are generally similar to the main reported results, albeit with (not surprisingly) much less precision. An exception is the "other technology" and "other fuel" results, for which we estimate 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.

Table A10 next shows additional controls; more flexible non-parametric conrols; etc.[51] See table notes for details. Overall, across our robustness checks, we estimate qualitatively similar fuel use shifts (increased generation by natural gas units) and CO_2 reductions.

We also show CO_2 results using reported rather than constructed emissions (Table A11). With this variable, we again estimate statistically significant emissions reductions.

A1.4 Heat Rate Effect of Regulation Provision

We briefly note that our results incorporate an additional effect on CO_2 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

[51] In previous versions of this paper, we additionally controlled for some things like power plant retirements; these are subsumed in this version by month of sample effects.

steady at a given level of output. This will worsen the heat rate, i.e. require greater heat input (and therefore more CO_2 emissions) per unit of electricity sold (Hirst and Kirby, 1997; Hummon et al., 2013).

Our regressions implicitly incorporate these two effects. Because our CO_2 emissions rate is time-varying, our left-hand side variable in Table 3, 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 CO_2 calculations line up, in the main text.

A1.5 References

Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker. 2012. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment." *American Economic Review*, 102(7): 3749–3760.

A1.6 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

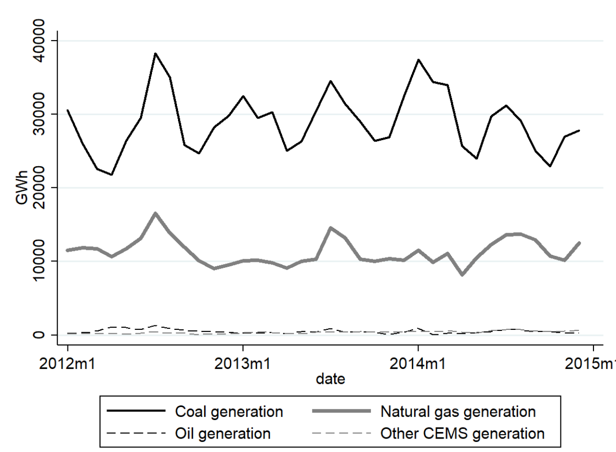

Note: This figure shows monthly generation by unit type for PJM units that appear in CEMS data.

Figure A3: Minimum Constraints

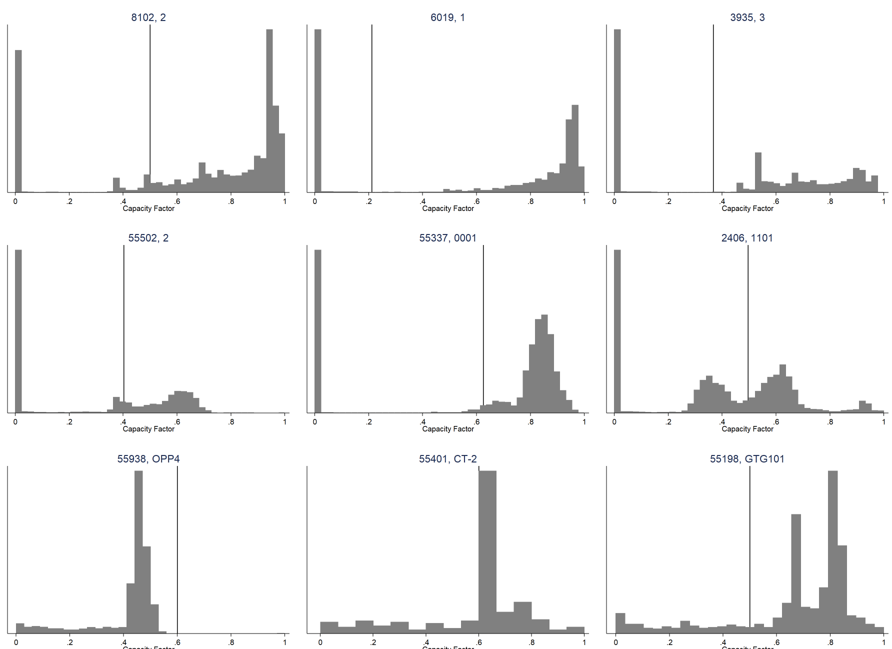

Note: These are nine units at large plants, three for each technology type. The top row shows three coal-fired boilers, with the plant id and unit id given at the top of each histogram. The second row shows natural gas combined cycle plants, and the bottom row shows natural gas combustion turbines. In 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 Online Appendix text for details.

Table A5: 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
Combustion 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
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
Henry hub natural gas price, dollars per mmbtu	3.97	0.64	19728
WTI oil price, dollars per barrel	94.7	10.3	19728
Coal price, dollars per mmbtu	2.34	0.040	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 and peak/valley forecast variables.

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

	Boiler	CC	CT	Other
Regulation requirement, 100 MW	-388.86**	357.11**	16.24	15.50
	(189.65)	(162.37)	(208.63)	(22.13)
PJM load, MWh	-0.15***	-0.01	0.16***	-0.00
	(0.02)	(0.02)	(0.02)	(0.00)
CEMS units generation, MWh	0.75***	0.27***	-0.02	0.01
	(0.02)	(0.02)	(0.01)	(0.00)
Peak forecast, in peak hours, MWh	0.02**	-0.01*	-0.01	0.00
	(0.01)	(0.01)	(0.01)	(0.00)
Valley forecast, in off-peak hours, MWh	0.08***	-0.01	-0.07***	0.00
	(0.01)	(0.01)	(0.02)	(0.00)
Henry hub price	556.01	-722.09**	165.99	0.09
	(351.02)	(302.18)	(148.48)	(22.15)
WTI price	5.31	-16.78	9.15	2.31
	(28.85)	(25.06)	(13.81)	(4.93)
Cooling degree days in Philadelphia	21.94	-29.40	7.78	-0.31
	(24.13)	(20.82)	(16.28)	(2.76)
Heating degree days in Philadelphia	20.11**	6.56	-28.23***	1.57
	(9.82)	(10.25)	(8.30)	(1.94)
Observations	19,693	19,693	19,693	19,693
Within R^2	0.91	0.71	0.49	0.13

Note: This table shows coefficients on the control variables for the regression results shown in the main text in Table 2. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A7: Displaying Control Coefficients: The Impact of the Regulation Requirement on the Energy Market

	Coal	NG	Oil	Other	CO2
Regulation requirement, 100 MW	-351.04	432.57**	-97.67	16.14	-242.37***
	(234.32)	(192.86)	(68.25)	(22.00)	(89.00)
PJM load, MWh	-0.19***	0.15***	0.04***	-0.00	-0.05***
	(0.02)	(0.02)	(0.01)	(0.00)	(0.01)
CEMS units generation, MWh	0.76***	0.24***	-0.01	0.01	0.78***
	(0.03)	(0.02)	(0.01)	(0.00)	(0.01)
Peak forecast, in peak hours, MWh	0.01	-0.02**	0.01***	0.00	0.03***
	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)
Valley forecast, in off-peak hours, MWh	0.07***	-0.07***	0.00	0.00	0.04***
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)
Henry hub price	561.19	-609.41*	45.82	2.40	580.30***
	(345.19)	(344.25)	(65.92)	(22.02)	(200.33)
WTI price	15.45	-16.05	-1.81	2.41	26.91*
	(29.96)	(28.68)	(6.00)	(4.97)	(15.56)
Cooling degree days in Philadelphia	-7.62	-11.45	19.27***	-0.21	31.40**
	(26.23)	(24.27)	(6.73)	(2.75)	(13.16)
Heating degree days in Philadelphia	28.43**	-24.14**	-5.84*	1.55	5.63
	(10.87)	(10.41)	(3.42)	(1.94)	(6.39)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.88	0.82	0.46	0.13	0.98

Note: This table shows coefficients on the control variables for the regression results shown in the main text in Table 3. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A8: Parsimonious Robustness Check: The Impact of the Regulation Requirement on the Energy Market

	Boiler (MWh)	CC (MWh)	CT (MWh)	Other tech. (MWh)
Regulation requirement, 100 MW	-637.8	280.3	101.6	-136.6***
	(410.2)	(204.0)	(239.0)	(25.3)
Observations	19,694	19,694	19,694	19,694
Within R^2	0.66	0.43	0.25	0.27

Note: This table shows a specification similar to that shown in the main text in Table 2, but controlling only for the peak and valley load forecasts. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A9: Parsimonious Robustness Check: The Impact of the Regulation Requirement on the Energy Market

	Coal (MWh)	Natural gas (MWh)	Oil (MWh)	Other fuel (MWh)	CO2 (tons)
Regulation requirement, 100 MW	-659.5	403.2	0.3	-136.5***	-460.8
	(442.3)	(275.0)	(84.1)	(25.4)	(333.9)
Observations	19,694	19,694	19,694	19,694	19,694
Within R^2	0.64	0.47	0.30	0.27	0.72

Note: This table shows a specification similar to that shown in the main text in Table 3, but controlling only for the peak and valley load forecasts. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A10: Robustness to Alternative Specifications: The Impact of the Regulation Requirement on the Energy Market

Panel A. Boiler																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	-408**	-381**	-389**	-348*	-425**	-407**	-340	-512***	-389*	-385**	-166	-358*	-403**	-406**	-337*	-376*
	(174)	(163)	(189)	(176)	(185)	(162)	(234)	(182)	(211)	(190)	(162)	(186)	(192)	(193)	(196)	(192)
Panel B. Combined Cycle																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	435**	414**	360**	375***	381**	363***	345*	357**	357**	356**	461***	369**	350**	368**	303*	348**
	(175)	(168)	(162)	(139)	(157)	(136)	(181)	(162)	(160)	(162)	(146)	(162)	(166)	(162)	(157)	(161)
Panel C. Combustion Turbines																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	-37	-61	17	-42	33	27	-20	20	16	13	-308**	-30	35	22	22	17
	(211)	(200)	(208)	(121)	(209)	(130)	(211)	(204)	(221)	(208)	(121)	(196)	(209)	(213)	(216)	(209)
Panel D. Other Units																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	10	29	13	16	11	16	15	18	16	17	13	18	17	16	12*	11*
	(23)	(23)	(21)	(23)	(21)	(21)	(29)	(23)	(21)	(22)	(23)	(22)	(22)	(22)	(7)	(6)
Panel E. Coal																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	-317	-280	-353	-318	-399*	-377*	-318	-467**	-351	-346	-66	-304	-374	-368	-296	-338
	(217)	(199)	(234)	(205)	(228)	(191)	(288)	(224)	(256)	(234)	(192)	(229)	(236)	(239)	(247)	(236)
Panel F. Natural Gas																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	445**	402**	435**	404**	474**	450***	387	429**	433**	427**	226	393**	449**	448**	376*	424**
	(177)	(167)	(192)	(183)	(187)	(167)	(238)	(190)	(211)	(193)	(168)	(189)	(195)	(196)	(193)	(189)
Panel G. Oil																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	-139**	-152**	-96	-102*	-87	-90	-85	-97	-98	-98	-173***	-107	-93	-97	-92	-98
	(69)	(64)	(69)	(54)	(66)	(56)	(81)	(67)	(70)	(68)	(59)	(67)	(68)	(69)	(75)	(68)
Panel H. Other Units																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	11	29	13	16	12	17	16	18	16	17	13	19	18	17	12*	11*
	(23)	(23)	(21)	(23)	(20)	(21)	(29)	(22)	(21)	(22)	(22)	(22)	(22)	(22)	(7)	(6)
Panel I. CO2 Emissions																
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Reg. req.	-274***	-239***	-243***	-232**	-247***	-242***	-276**	-365***	-242***	-240***	-207**	-247***	-252***	-247***	-201**	-266**
	(87)	(83)	(89)	(91)	(86)	(86)	(110)	(80)	(93)	(89)	(90)	(89)	(90)	(89)	(95)	(104)

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

Table A11: Alternative CO_2 Measurement: The Impact of the Regulation Requirement on the Energy Market

	Coal	NG	Oil	Other	Total
Panel A. CO_2 Variable in Main Text, Metric Tons					
Regulation requirement, 100 MW	-355*	168	-67	11	-242***
	(214)	(122)	(55)	(8)	(89)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.89	0.79	0.43	0.25	0.98
Panel B. CO_2 as Reported, Metric Tons					
Regulation requirement, 100 MW	-361*	142	-63	8	-274***
	(214)	(122)	(51)	(7)	(92)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.89	0.78	0.44	0.20	0.98
Panel C. Using Unit-Level Emissions Rates, Metric Tons					
Regulation requirement, 100 MW	-355*	197*	-61	6	-213**
	(212)	(118)	(53)	(8)	(87)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.89	0.79	0.44	0.21	0.98

Note: Panel A shows the CO_2 emissions results by fuel type (Columns 1 through 4) and aggregated (Column 5), matching the specifications used in the main text, Table 3. Panel B shows analogous specifications, but using CEMS-reported CO_2 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 CO_2 data. All panels are reported in metric tons (i.e., in Panel B we convert CEMS-reported short tons into metric tons). *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A12: 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	348.6*	40.2	-22.2	6.9	-3.1	-32.9	-12.1
	(196.8)	(141.8)	(21.3)	(4.4)	(6.7)	(67.1)	(153.0)
Observations	19,693	19,693	19,693	19,693	19,693	19,688	19,693
Within R^2	0.71	0.67	0.12	0.11	0.09	0.17	0.53

Note: This table shows estimates from seven separate regressions, analogous to those presented in the main text, Tables 2 and 3. 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 39. 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. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A13: Specifications with Two-Way Fixed Effects

	Boiler	CC	CT	Other tech	Coal	NG	Oil	Other fuel	CO2
Panel A. Time Effects, Location Effects, and Peak/Valley Forecast Controls									
Reg. req., 100 MW, in PJM	-498	335*	97	-123***	-496	389*	27	-110***	-263
	(373)	(191)	(142)	(27)	(377)	(226)	(77)	(27)	(371)
Observations	39,388	39,388	39,388	39,388	39,388	39,388	39,388	39,388	39,388
Within R^2	0.22	0.26	0.13	0.24	0.16	0.32	0.26	0.19	0.33
Panel B. Plus Additional Controls									
	Boiler	CC	CT	Other tech	Coal	NG	Oil	Other fuel	CO2
Reg. req., 100 MW, in PJM	-494***	413***	68	13	-419**	483***	-73	9	-303***
	(149)	(146)	(133)	(22)	(180)	(155)	(61)	(23)	(90)
Observations	39,386	39,386	39,386	39,386	39,386	39,386	39,386	39,386	39,386
Within R^2	0.75	0.40	0.25	0.11	0.71	0.55	0.40	0.07	0.92

Note: Panel A shows specifications in which plants in nearby states (see footnote 39) 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 fixed effects, and two interaction variables. These latter controls are (1) forecasted peak load interacted with a PJM dummy, and forecasted valley load interacted with a PJM dummy; these may be important for avoiding omitted variables 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 2 and 3 interacted with a PJM dummy for additional precision. We also include a control for the total CEMS generation within each region. 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. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A14: 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	-0.29	0.20	0.39	1.18	-1.48***
	(1.03)	(0.16)	(0.40)	(0.90)	(0.47)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.71	0.11	0.38	0.43	0.75

Panel B. Combined Cycle Plants	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.33***	0.03	0.40*	2.31***	-0.40
	(0.86)	(0.25)	(0.24)	(0.70)	(0.25)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.22	0.07	0.61	0.32

Panel C. Combustion Turbines	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.13	0.01	0.31	3.06	-1.26**
	(3.38)	(0.46)	(0.33)	(2.44)	(0.59)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.31	0.27	0.45	0.19

Panel D. Other Units	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.03	-0.09	-0.10***	0.14	0.07
	(0.17)	(0.06)	(0.03)	(0.15)	(0.04)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.19	0.03	0.03	0.14	0.04

Note: This table expands on Table 4 by showing results at other units. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A15: 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	-0.29	0.04	0.48	1.26	-1.48***
	(1.03)	(0.53)	(0.66)	(0.96)	(0.47)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.71	0.36	0.51	0.65	0.75
Panel B. Combined Cycle	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.33***	0.55	0.73**	1.46*	-0.40
	(0.86)	(0.34)	(0.36)	(0.80)	(0.25)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.22	0.15	0.59	0.32
Panel C. Combustion Turbine	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-2.13	0.01	0.20	3.18	-1.26**
	(3.38)	(0.20)	(0.32)	(2.72)	(0.59)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.20	0.43	0.44	0.19
Panel D. Other Tech	Off	Below min	At min	Above min	At max
Regulation requirement, 100 MW	-0.03	-0.18***	0.15***	-0.02	0.07
	(0.17)	(0.06)	(0.06)	(0.15)	(0.04)
Observations	19,693	19,693	19,693	19,693	19,693
Within R^2	0.19	0.03	0.05	0.18	0.04

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. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Table A16: 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	-0.25	0.12	0.18	0.04	-0.51	0.11	0.60	0.79***	0.34	-0.13	-1.26*
	(1.04)	(0.11)	(0.15)	(0.14)	(0.34)	(0.36)	(0.37)	(0.26)	(0.26)	(0.42)	(0.73)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.70	0.05	0.06	0.20	0.34	0.31	0.22	0.11	0.07	0.36	0.82

Panel B. Combined Cycle											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-2.33***	-0.14***	-0.07***	0.03	-0.02	-0.06	0.33	1.65***	-0.14	1.23***	-0.46
	(0.86)	(0.04)	(0.02)	(0.03)	(0.07)	(0.13)	(0.27)	(0.37)	(0.40)	(0.46)	(0.46)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.64	0.05	0.04	0.04	0.02	0.07	0.17	0.14	0.08	0.54	0.46

Panel C. Combustion Turbine											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-2.15	-0.08	0.02	0.06	0.07	0.27	0.57*	0.73	1.99**	0.14	-1.63**
	(3.39)	(0.14)	(0.09)	(0.08)	(0.12)	(0.23)	(0.30)	(0.44)	(0.96)	(0.84)	(0.78)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.48	0.14	0.14	0.14	0.14	0.22	0.26	0.29	0.39	0.42	0.24

Panel D. Other Tech											
	0	10	20	30	40	50	60	70	80	90	100
Reg. req., 100 MW	-0.03	-0.02	-0.02	-0.08***	-0.01	0.22***	-0.03	-0.10	-0.23**	0.01	0.28***
	(0.17)	(0.03)	(0.02)	(0.03)	(0.06)	(0.08)	(0.06)	(0.06)	(0.09)	(0.06)	(0.09)
Observations	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693	19,693
Within R^2	0.19	0.01	0.02	0.03	0.02	0.06	0.01	0.03	0.05	0.06	0.13

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. *** Statistically significant at the 1% level; ** 5% level; * 10% level.