Federal Reserve Bank of Dallas

What Fuels the Volatility of Electricity Prices?

Alessio Saretto, Anastasia Shcherbakova and Jeremy Lin

Working Paper 2408 **October 2024**

Research Department

https://doi.org/10.24149/wp2408

What Fuels the Volatility of Electricity Prices?*

Alessio Saretto†, Anastasia Shcherbakova‡ and Jeremy Lin§

October 8, 2024

Abstract

We use emergency outages of coal generators as an exogenous source of variation in the power generation stack to study how changes in marginal fuel affect real-time prices. Contrary to anecdotal evidence, we find that wholesale prices are less volatile when natural gas is on the margin more often.

Keywords: policy spillovers, electricity price volatility, fuel switching, environmental policy

JEL Classification: Q41

---

*The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Dallas, the Federal Reserve System, or their staffs.

†Corresponding author: Alessio Saretto, The Federal Reserve Bank of Dallas, Alessio.Saretto@dal.frb.org. Alessio Saretto declares that he has no relevant or material financial interests that relate to the research described in this paper.

‡Anastasia Shcherbakova, Coherent Economics, Anastasia.Shcherbakova@CoherentEcon.com. Anastasia Shcherbakova declares that she has no relevant or material financial interests that relate to the research described in this paper.

§Jeremy Lin, Transmission Analytics, jeremylin@transmissionanalytics.net. Jeremy Lin worked at PJM in the role of Senior Engineer from 2005 until 2017, during which time he obtained the data used in this paper. He currently works at Transmission Analytics in the role of Director, in which he provides consulting services to generator developers and financial institutions in the areas of transmission risk analysis, market analysis, curtailment analysis of renewable generators, and review of third-party studies. He has no material financial interests that relate to the research described in this paper.

1 Introduction

The mix of fuels used to generate electricity in the U.S. has changed dramatically over the last two decades. One of the most pronounced changes has been a switch from coal to gas-based generation, and this has led to improved emissions outcomes (U.S. Energy Information Administration, 2018). Researchers and market participants have expressed a concern that a shift from coal to gas generation has led to higher real-time electricity price volatility (Brown and Kodaka, 2014; Linn and Muehlenbachs, 2018; Deyette et al., 2015; Larson, 2017).[1] The implication is that coal-to-gas switching might have led to better environmental outcomes at the expense of greater financial strain.

In theory, it is *ex-ante* uncertain how a switch from coal to gas generation affects price volatility in the wholesale electricity market. Volatility is driven by generators' bidding behavior, which, in turn, is affected by marginal and fixed generation costs. Marginal generation costs are mainly represented by the cost of sourcing fuels (see for example Cicala, 2015). Because natural gas prices are more volatile than coal prices, replacing coal generators with gas generators can result in higher average bids to supply generation and therefore higher wholesale price volatility. The existence of a volatility pass-through is essentially the argument put forth in the literature (Brown and Kodaka, 2014; Linn and Muehlenbachs, 2018; Deyette et al., 2015; Larson, 2017). On the other hand, fixed costs are primarily startup costs. In PJM, if a generator is not selected for dispatch in the day-ahead market, it can submit a revised (lower) bid to participate in the real-time market. Generators with high startup costs are limited in their ability to revise bids. Because coal generators have higher startup costs than natural gas generators, replacing coal generators with gas generators can result in lower real-time market bids and consequently lower wholesale price volatility.

Disentangling these two channels is not straightforward, as neither competing explanation is directly testable. On one hand, measuring the volatility pass-through from the natural

[1] Technological breakthroughs that led to the U.S. shale revolution was an important factor in the shift toward gas generation. Because our analysis is based on the post-fracking time period, we take the technological impacts of fracking as given.

gas market to the electricity market would require precise knowledge of the fuel procurement strategy of each generator. Even if one could estimate the pass-through cost (see Chu et al., 2017), one would face an empirical challenge: the estimated pass-through would introduce an error-in-variable problem when used as an explanatory variable. A similar problem would arise if one attempted to estimate generators' startup costs (see Reguant, 2014). A further complication is that the factors that affect supply (i.e., producers' bidding behavior) likely also affect demand (i.e., load). To circumvent the estimation challenges, we use plant-level emergency outages of coal generators as a source of exogenous variation in the marginal fuel source to isolate the effect that fuel switching has on wholesale power price volatility in the PJM service territory for the years 2014-2016.

Emergency outages are unscheduled interruptions in generating capacity caused by unanticipated operational problems. As an instrument for the amount of natural gas generation, emergency coal outages appear to satisfy the set of conditions that grant a causal interpretation to the instrumental variable estimate (as in Imbens and Angrist, 1994; Angrist et al., 1996): they are largely independent of unobservable factors that might affect demand for natural gas or market conditions that determine price volatility; they produce a relatively monotonic effect on the marginal fuel source (i.e., eliminating a coal generator from the production stack is unlikely to put another coal generator on the margin); they are unlikely to affect price volatility directly, but do impact the ability of coal generators to participate in the real time market.

We focus on the volatility of electricity prices that are set in the real-time balancing market (rather than the day-ahead market), where load servicing entities have to bear most of the financial risk due to sudden shifts in supply or demand. After isolating the effects of exogenous variation in the marginal fuel source, we find that the percentage of time that gas generators spend on the margin is related to lower, not higher wholesale real-time price volatility. Specifically, an extra 30 seconds of natural gas on the margin reduces the hourly range (i.e., the difference between the hourly maximum and minimum price) of real-time electricity prices by 72 cents (or approximately two percent of the average price).

This result holds up to several robustness checks: alternative definitions of electricity price volatility and emergency coal outages, considering electricity in the day-ahead market, the amount of wind generation, inclusion in the sample of hours when multiple fuel sources (i.e., not just gas and coal) are on the margin, and controlling for price volatility in fuel sources. Our statistical results support the Reguant (2014) argument that fixed costs are essential in determining the dynamic bidding behavior of electricity generators. Because of lower startup costs, natural gas generators are able to respond to variation in demand and other market conditions more efficiently, and therefore, relative to coal generators, they marginally reduce electricity price volatility. Note however, that while we do not find evidence that support a *volatility pass-through* from the source fuel market into electricity, does not imply that there is no *cost pass-through*. As Chu et al. (2017) find, it takes some time for changes in the spot price of natural gas to translate into changes in generators procurement costs. This is likely due to the fact that procurement contracts are a type of swap contract, and are renegotiated only at coarse intervals. When such contracts are renegotiated, there would be a substantial amount of cost pass-through, but it would not translate in volatility pass-through as long as the frequency of renegotiation is low enough.

Our findings have implications for the policy debate surrounding the continuing gas integration in electricity markets. Holding the PJM generation stack fixed, we provide evidence that coal-to-gas switching at the margin does not create negative spillovers for immediate financial outcomes, at least in the form of increased volatility. Our results should therefore generalize to regions that currently split power generation between coal and natural gas: replacing some coal generation with natural gas generation at the margin should produce similar results as those reported in this paper. Figure 1 shows how the share of coal and gas generation has changed over time in U.S.'s regional power markets. California, New York, and New England have had negligible coal shares for some time, and gas has been the dominant fuel for several years in PJM, ERCOT, and the Southeast. In MISO, Southwest, Northwest, and SPP markets, however, gas is just starting to overtake coal. These markets cover 25 U.S. states and our results can inform the policy of these regions.

Figure 1: Share of coal and natural gas in total generation

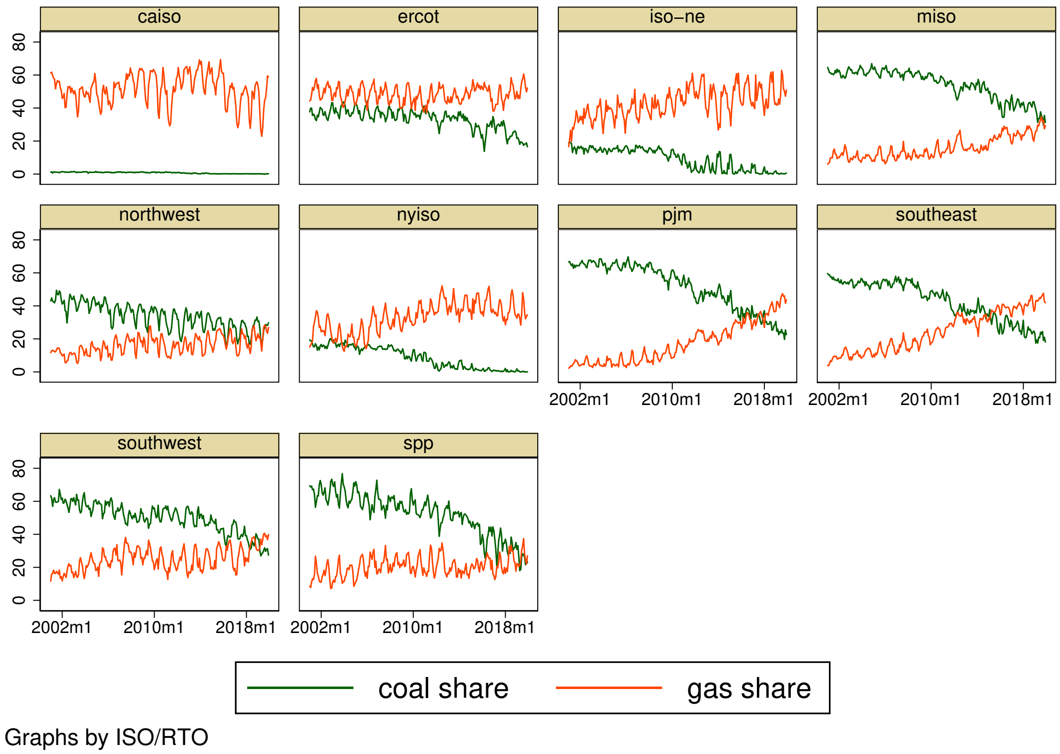

Note: The figure displays the average share of generation coming from coal and natural gas plants over time in each ISO/RTO. Data is from the EIA and covers the period from January 2001 through December 2019.

From a financial stand point, volatility represents a substantial concern for market participants, as it ultimately shapes their hedging demand. The more volatile the markets, the larger the pressure to engage in costly hedging activities (see for example, Aid et al., 2011; Boroumand and Zachmann, 2012; Dupuis et al., 2016; Boroumand et al., 2019) and creates an indirect costs in the form of capital reallocation to meet collateral requirements (see for example Rampini and Viswanathan, 2010; Rampini et al., 2014). Being regulated, electric utilities aim to avoid exposure to large risks in their operations (Wolak and Kolstad, 1991; Neuhoff and De Vries, 2004; Fehrenbacher, 2010). With the exception of infrastructure investments, regulated utilities are not allowed to make a profit, but are only allowed to recover their operating costs and, thus, cannot price unexpected shocks (large risks) into their retail rates. They have therefore an interest in minimizing large risks, such as short-term price volatility, which, according to Borenstein (2002), is a necessary consequence of the way electricity markets are structured. To hedge, utilities use long-term contracts to smooth out short-term price fluctuations. However, in absence of further hedging, they remain exposed to market volatility in proportion to how much capacity they purchase in the real-time market.

2 Related literature

We contribute to existing literature that studies the economic implications of natural gas integration into the power generation infrastructure. Linn and Muehlenbachs (2018) find that low natural gas prices in the U.S. have encouraged the switch from coal to natural gas generation and led to lower wholesale electricity prices, but that the price-reducing benefits of fuel switching were lower in regions with more switching. We examine one financial outcome of fuel-switching in one such region and show that coal-to-gas switching has an economically meaningful effect on the volatility of real-time electricity prices.

In another recent study, Chu et al. (2017) examine power plants' procurement costs and find that increases in natural gas spot prices pass through to power plants relatively quickly, with 85 percent of changes in spot prices of natural gas being reflected in natural gas plants' procurement costs within one month. This result gave some support to the idea that power markets with a large natural gas generation capacity could experience more volatile wholesale power prices (Brown and Kodaka, 2014; Deyette et al., 2015; ISO New England, 2018). The evidence offered by these studies, however, ignores potential endogeneity concerns that complicate the estimation of the relationship between plants' bidding strategies and market outcomes. We account for possible endogeneity and provide a more nuanced test of the relationship between the composition of the generation stack and electricity price volatility.

Our paper also complements the literature that examines the impact of deregulation in electricity generation. Fabrizio et al. (2007) and Cicala (2015) find that deregulation had a significant effect on the strategic operating choices of U.S. electricity generators, leading to reductions in operating costs and gains in efficiency. These studies provide economic benchmarks of potential efficiency gains that stem from a change in a regulatory structure (i.e., deregulation), while we contribute insights about potential financial benefits that could be gained by changing the mix of the physical generating infrastructure.

Finally, our results are directly related to the literature that studies measurement and forecasts of electricity price volatility (see for example Haugom et al., 2011; Hickey et al., 2012; Ullrich, 2012; Liu and Shi, 2013; Frommel et al., 2014; Ciarreta et al., 2017; Qu et al., 2018, among many others); the economic research on the impacts of renewable intermittency on electricity prices (Ketterer, 2014; Martinez-Anido et al., 2016; Kyritsis et al., 2017; Rintamaki et al., 2017; Masoumzadeh et al., 2018); and the finance literature that studies equilibrium electricity forward and spot prices, and hedging demands of market participants (see for example Bessembinder and Lemmon, 2002; Longstaff and Wang, 2004; Banerjee and Noe, 2006; Pirrong and Jermakyan, 2008; Bunn and Chen, 2013). Most of these studies model supply and demand effects, taking the marginal unit of generated power as exogenous. We take a step forward in understanding the relationship between prices and supply sources by endogenizing the relationship between price volatility and the generation stack and explicitly analyzing the impact of marginal fuel switching on the dynamics of real-time prices. Moreover, we extend the literature that studies volatility estimators based on high frequency data (for example, Ullrich, 2012; Haugom and Ullrich, 2012) by considering the case of negative prices.

3 Empirical framework

Uncovering the relationship between natural gas generation and price volatility is difficult, as unobservable factors that affect supply (i.e., producers' bidding behavior) likely also affect demand (i.e., load). We use emergency outages of coal generators to instrument for the amount of time that natural gas generators spend on the margin: taking one coal unit out of service will push a gas generator into the supply stack. We detail our econometric model and the assumptions needed for a causal interpretation in this section. Because bidding rules are essential to explain the economic mechanism that links outages to marginal generation, we review how the PJM market functions.

3.1 Wholesale bidding and market clearing in PJM

PJM's wholesale energy market is cleared using a two-tier mechanism: a day-ahead forward market where generators bid production for each hour of the day, and a real-time balancing market.

The day-ahead market is essentially organized as a forward market. According to the current PJM (Manual 11) rules, the morning ahead of each operating day (until 10:30AM), generators enter bids into the system for each hour of the day. Each bid must contain a quantity that is offered at a particular price along with a startup time. In addition, most generators are required to submit unit availability and non-binding generation bids for the following seven days. By 1:30PM, PJM determines a price that clears the load forecast for each hour with the quantities offered, and posts a dispatch schedule.

After the day-ahead market is cleared, units that are not scheduled to supply can modify their existing bids, until 2:15PM, for the hours in which they are not committed so that they might participate in the resource adequacy run. Generators can also enter new bids to participate in the real-time market, starting at 6:15PM and until 65 minutes before the operating hour. All bids that are accepted for the day-ahead market and not subsequently updated are considered binding for the real-time market. The real-time market consists of balancing settlements based on deviations of actual load from the quantities cleared in the day-ahead market. The real-time price is determined at five-minute intervals based on the residual bids from the day-ahead market and their subsequent revisions.

3.2 Econometric model

We model the fraction of time that natural gas is on the margin in hour h of day d as a function of emergency coal outages and other control variables, including proxies for demand, demand uncertainty, the amount of wind dispatched into the system, and various fixed effects,

$$log(1 + NaturalGasMargin_{d,h}) = \beta_0 + \beta_1 \ log(1 + CoalOutages_{d,h}) + \sum_{j=2}^{J} \beta_j X_{j,d,h} + \varepsilon_{d,h}, \quad (1)$$

where subscript j denotes the generic control ($j = 2 \dots J$). To facilitate the interpretation of the coefficient of interest, β_1, as an elasticity, we take the natural logarithm of both dependent and independent variables.

We observe the fraction of a given hour that a specific generation technology is on the margin. To identify the effect of fuel switching, we restrict our analysis to hours when only coal or gas is on the margin (relaxing this restriction does not alter our results.) In the data, $NaturalGasMargin_{d,h}$ reflects the share of an hour, between zero and one, during which the marginal generator is a natural gas plant. A value of 0.33 would mean that natural gas plants were the marginal producers for 33 percent of an hour, or 20 minutes. $CoalOutages_{d,h}$ captures the amount of coal capacity that is offline during an hour. Outages are measured in percent of total available coal generation capacity and reflect the capacity loss across all coal generators that experience an emergency outage during that hour. $\mathbf{X}_{d,h}$ is a vector of control variables. It includes demand (i.e., total load, in megawatt-hours during the hour) and a measure of demand uncertainty, expressed as the difference between actual realized load and PJM's forecasted load (both in log-MWh) for a given hour that is contracted for in the day-ahead market. $\mathbf{X}_{d,h}$ also captures intraday and seasonal variation in demand with hour and month fixed effects, and with indicator variables that isolate hours of extreme temperatures (i.e., above 25 and below –5 degrees Celsius, for heat and cold indicators, respectively). We include a weekend indicator in $\mathbf{X}_{d,h}$ to further differentiate between peak and off-peak demand hours. Lastly, we include a measure of total wind generation supplied into PJM during an hour to capture the fact that natural gas, as the main backup source for intermittent renewables, is often dispatched concurrently with wind.

In the second stage, we model real-time electricity price volatility in an hour as a function of the fitted estimate of natural gas margin in the same hour from equation (1), again controlling for variables that proxy for the total demand and supply, demand uncertainty, and the amount of wind dispatched into the system,

$$log(\sigma_{d,h}) = \gamma_0 + \gamma_1 log(1 + \widehat{NaturalGas}Margin_{d,h}) + \sum_{j=2}^{J} \gamma_j X_{j,d,h} + \eta_{d,h}. \quad (2)$$

$\sigma_{d,h}$ is the hourly volatility of the wholesale price in the real-time (five-minute-ahead) electricity market, defined as the intra-hour range (i.e., the difference between the maximum and minimum observed five-minute prices within the hour) scaled by the median price of previous 24 hours.

We focus on real-time volatility for two main reasons: first it allows us to tightly link the outcome variable to the exogenous variation in the fuel source (because we observe the timestamp of emergency outages). Second, load serving entities' exposure to market risk is greater in the real-time market, which presents greater potential for large financial losses and need for expensive hedging strategies. We provide evidence from the day-ahead market in the robustness section.

3.3 Causal interpretation

Angrist et al. (1996) list statistical assumptions that an instrumental variable must satisfy to serve as a valid instrument and provide a causal interpretation of the results. The first assumption is that the occurrence of an emergency outage is an "ignorable" event, which allows one to interpret an IV estimand as equivalent to a randomized study. The definition of an emergency outage (i.e., decision to stop a generating unit to prevent an outcome that would endanger lives or structures) indicates that it is an event that cannot be influenced by affected economic agents (i.e., other generators). For example, fires may be more likely during days with extremely high temperatures, but weather conditions are also essentially random, which guarantees that the mechanism that produces the outage is ignorable (see Rubin, 1978; Imbens and Rubin, 1997).

The second assumption is that an outage only affects the dynamics of electricity prices by affecting the marginal fuel source (i.e., how much natural gas generation contributes to marginal production). This is often referred to as the "exclusion restriction." We argue that mechanisms that would lead to a violation of the assumption are either implausible or economically irrelevant. For example, it would have to be the case that the event that is causing an emergency coal outage is also causing a large shift in demand that could cause large variations in electricity prices. Since most outages are related to fire and they are relatively more likely on hot days, weather could be one such event. Our data contains more outages in December and January than in July and August. In terms of displaced capacity, the largest outages tend to happen in March and April. Thus, despite an economically small relationship between outages and weather, weather is not driving variation in the incidence of emergency outages.

The third assumption requires that outages have a monotonic effect on the percentage of time that the marginal fuel is natural gas (i.e., the coefficient β_1 in equation 1 is different from zero). We describe the economic mechanism that leads an outage to affect generators' behavior. Recall that bids can be submitted no later than 65 minutes before the start of the operating hour. Within the first 30 minutes after an outage, PJM communicates the system's status to all market participants. For very short outages, or for the first hour of a long outage, generators will not be able to adjust their bids. Therefore, the regional transmission operator will just dispatch the first unit that was left out (the next lowest bid).[2] For outages longer than one hour, generators that are participating in the real-time market can update their bids. However, how the generator stack changes as a consequence of a coal outage depends on both submitted bids and a generator's ability to respond. 99 percent of the coal generators in our sample have a ramp-up time that exceeds one hour (with 40 percent of generators requiring more than 12 to ramp up). In contrast, about 90 percent of natural gas generators have a ramp-up time less than one hour. Hence, gas generators are the most likely to be chosen to respond to a shift in the power supply even if their bids are not the lowest.

The fourth assumption is that a generator's potential outcome (to shut down or not to shut down) is based only on its own treatment status, and not on the treatment status of other generators (i.e., the "Stable Unit Treatment Value Assumption.") In context of forced outages, a "treated" generator experiences adverse operating conditions that force it to shut down, while an "untreated" generator does not experience the conditions that would force it to stop producing electricity and shut down. Due to the unpredictable nature of emergency outages, there should be no spillover effects across generators.

Finally, we discuss how the instrumental variable estimation is affected by our focus on the marginal generator. We study the variation in the amount of natural gas at the margin of the generation stack, rather than considering the entire production function. Since emergency outages displace a relatively small amount of coal generating capacity (i.e., less than 1% of total capacity), it would be hard to argue that they have an equal effect on all generators, from the very inframarginal, such as wind and nuclear, to the very inefficient peaker units. Instead, because we focus on the marginal unit, it is much more likely that units affected by small variation in the supply of power from coal generators are either at or just outside the

[2] PJM's immediate response to an outage might be to call upon synchronized reserves. It is unlikely that this significantly affects our results: during our sample, synchronized reserves were called upon 76 times for an average duration of 21 minutes. Moreover, on rare occasions when balance is not restored, PJM could also import some interface resources. During our sample, interface imports were never on the margin.

margin.

4 Data

The data for our analysis span 2014 through 2016 and come from three main sources: PJM, the Energy Information Administration (EIA), and the National Centers for Environmental Information (NCEI). We observe several quantities of interest from PJM, which are sampled at different frequencies. We sample real time locational marginal prices (LMPs) for the wholesale market at five-minute intervals.[3] We use only the energy portion of LMP, which is common across all pricing nodes within PJM, and disregard the cost of transmission losses and congestion.

We observe hourly data on total generation supplied to the system, the hourly demand forecast generated by PJM one day in advance, total hourly wind generation, and the fraction of each hour during which each fuel type is on the margin.

For generator outages, we observe the start and end time of each outage, and the code entered into the system to describe the reason for the interruption. As defined previously, emergency outages are instances that are not scheduled and that are due to operational problems that occur while the unit is running. We observe 1,778 unique generators in PJM's outage data. We obtain these generators' characteristics from EIA's form 860. Form 860 is an annual survey of all electric generators in the U.S. with nameplate capacity of 1 MW or more. The survey collects information on ownership, location, capacity, and environmental characteristics of current, proposed, and retired generating assets.

For 2014-2016, the EIA reports information on 1,400 unique power plants and 4,048 unique generators in the PJM balancing authority. Unfortunately there is no unique generator identifier that is consistent between data sources, so in order to merge generator characteristics into PJM data, we must rely on matching by hand. We use plant name, generator name, and generator technology as matching parameters. Since plant and generator

[3] In PJM, the real-time LMP calculation includes optimization of ancillary services.

naming conventions also differ between EIA and PJM, we are not able to identify matches for all of our generators. We end up with matches for 1,556 generators at 450 plants, which represent over 87 percent of our PJM generator sample.

Finally, we obtain hourly air temperature information for all land-based weather stations in the PJM service area from the National Oceanic and Atmospheric Administration's (NOAA) National Centers for Environmental Information (NCEI). We identify the closest weather station for each PJM generator for which we are able to obtain an EIA match, matching latitude and longitude pairs in EIA data with those in NCEI data. The accuracy of information reported by weather stations tends to vary both within and across stations, and NCEI discloses the quality of each observation that appears in its datasets. In order to ensure accurate temperature representation for our data sample, we keep only those air temperature observations that have passed basic quality control tests (i.e., codes 0, 1, 4, 5, 6, and 9). For hours during which several intra-hour temperature observations are available, we average across these observations to obtain an hourly data series. Temperature is expressed in degrees Celsius.

We report summary statistics for PJM system outcomes in Table 1. During our sample, the average hourly load was around 90 gigawatts (GW), with peaks exceeding 150 GW. Coal was on the margin for an average of 45 percent of an hour, natural gas was on the margin on average for 39 percent of an hour, with renewables and petroleum-based generation together accounting for another twelve percent of an hour. A fuel source is rarely on the margin for 100 percent of the hour: coal is the only marginal fuel during 361 hours, natural gas during 291 hours.

The average real-time price was \$35.60 per MWh. In our sample, there are 158 hours during which at least one five-minute real-time price was negative. As one would expect, prices are extremely volatile: the average hourly price ranges between -\$242.70 and \$1,068.80 per MWh over the three-year period. Within the hour, the volatility is even higher with an intra-hour range varying between \$0 and \$1,858.35. The average range observed within an

Table 1: Generation summary statistics

	Mean	StDev	Min	Median	Max
Price – mean ($)	35.6	33.4	-242.7	27.7	1,068.8
Price – median ($)	33.9	32.1	-320.5	27.3	1,191.8
Price range ($)	21.6	56.2	0.0	5.3	1,858.4
Scaled price range	0.63	1.40	0.00	0.19	21.41
Price volatility ($)	7.5	20.3	0.0	1.7	762.8
Scaled price volatility	0.22	0.51	0.00	0.06	9.40
Coal on margin (% of one hour)	0.45	0.21	0.00	0.45	1.00
Gas on margin (% of one hour)	0.39	0.19	0.00	0.39	1.00
Nuclear on margin (% of one hour)	0.01	0.05	0.00	0.00	0.50
Petroleum on margin (% of one hour)	0.07	0.14	0.00	0.00	0.60
Renewables on margin (% of one hour)	0.05	0.09	0.00	0.00	1.00
Load (MW)	90,437	17,056	57,108	87,862	152,177
Day-ahead load forecast (MW)	91,561	17,021	57,436	88,982	152,117
Wind generation (MW)	1,871	1,297	0	1,627	6,249

Note: We report mean, standard deviation, minimum, median, and maximum values for variables used in the rest of the paper. All items are measured hourly. Volatility measures are constructed at the hourly interval from real-time prices sampled at five-minute frequency. The remaining 3% of marginal fuels is categorized as demand side, interface, miscellaneous, or missing. All data is from PJM and covers the period from January 1, 2014 through December 31, 2016.

hour is $21.60 per MWh, or 63 percent of the average hourly price level.

Figure 2 plots time series for two different measures of price volatility. The top panel depicts the hourly price range, defined as the difference between maximum and minimum prices observed within a given hour, scaled by the 24-hour average price. We observe several instances in which the within-hour range is larger than $500 per MWh, and in the winter of 2014, the intra-hour price variation peaked above $1,800 per MWh. The bottom panel of Figure 2 displays our second measure of price volatility: the standard deviation of five-minute prices within the hour, also scaled. This measure largely reflects the dynamics of the previous one. Indeed, the correlation between the two time-series is over 98 percent.

In Table 2, we summarize the relevant characteristics of generator outages observed in our data sample. In general, outages are quite frequent—there are more than eighty thousand in our overall sample, approximately 40 percent of which are recorded in coal generators, followed by 35 percent in gas and combined-cycle plants. For brevity, we limit the summary

Figure 2: Real-time electricity price volatility

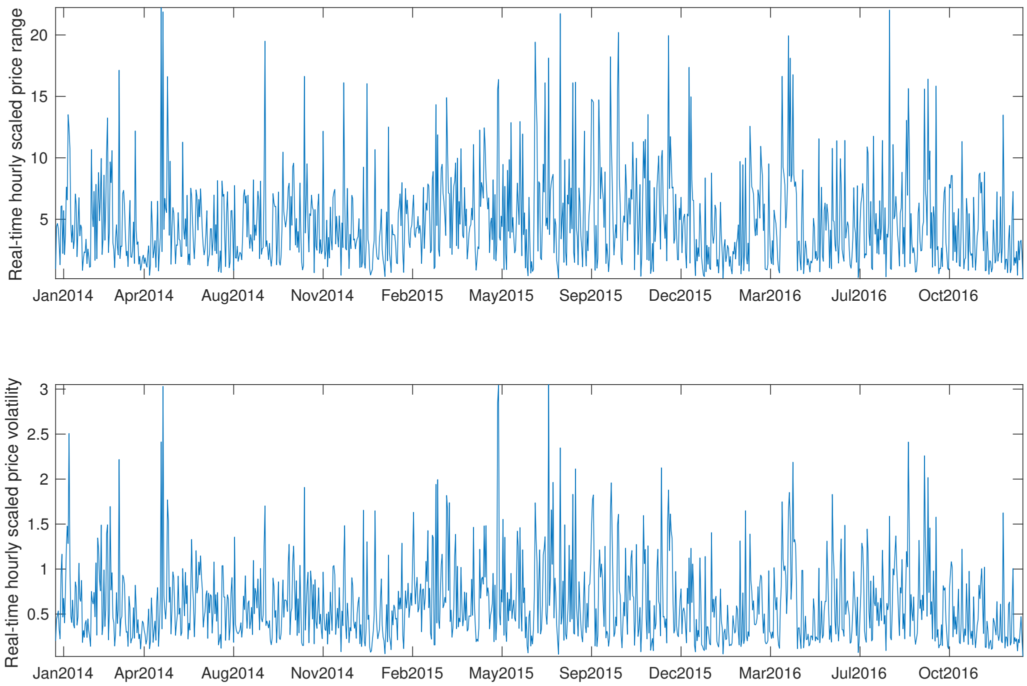

Note: We plot time series of two hourly measures of volatility, obtained from real-time wholesale 5-minute prices. The top panel displays the scaled price range for each hour, defined as the difference between the maximum and the minimum price observed within the hour. The bottom plot shows the scaled standard deviation of price levels within the hour. All data comes from PJM and covers the period from January 1, 2014 through December 31, 2016.

statistics in Table 2 and our subsequent discussion only to outages experienced by coal and natural gas generators. On average, coal and gas plant outages last 2.25 days and affect between 30 and 40 percent of a generator's capacity.

Emergency outages make up about 1.8 percent of all outages. They are disproportionally concentrated in coal power plants (about 90 percent of the cases). On average, emergency outages are about half a day longer than planned outages, and more frequently affect older natural gas units, but not older coal units.

Emergency outages can be divided into two groups: those related to a fire (according to PJM manuals, "operational outage caused by, or taken to alleviate concerns with, fire or

Table 2: Outages summary statistics

	Instances	Plants	MW Reduction	% Generator Reduction	Duration Days	Age Years
			All outages			
Coal	32169	177	125.4	0.28	2.25	44.5
Natural Gas	28915	423	59.7	0.38	2.24	15.1
			Emergency outages			
Coal	1030	91	86.3	0.22	2.66	41.2
Natural Gas	109	44	102.7	0.42	3.01	25.0
			Emergency outages – Non-fire related			
Coal	80	41	132.1	0.36	2.53	40.6
Natural Gas	64	32	82.2	0.47	2.41	20.4
			Emergency outages – Fire related			
Coal	950	88	82.4	0.21	2.67	41.2
Natural Gas	45	17	131.8	0.36	3.86	31.6

Note: The table reports summary statistics for all outages and for emergency outages. For each fuel source, we tabulate the number of instances, the number of plants affected, the average reduction in generation capacity in MW, the average reduction as percent of nameplate generation capacity, the average duration of an outage, in days, and average generator age, in years. All data is from PJM and covers the period from January 1, 2014 through December 31, 2016.

smoke"), and those that are not (described in PJM manuals as "operational outages that are taken for the purpose of avoiding risk to human life, damage to equipment, damage to property, or similar threatening consequences"). Fire-related outages represent more than 90 percent of all emergency outages that affect coal generators.

We plot the time series of emergency outages that affect coal power plants and the relative affected generation capacity in Figure 3. The total hourly MW capacity reduction due to outages is calculated as the sum of MW reductions across all units that are off line during a given hour. Consider two outages: one that begins at 3am on a Monday and lasts for three days, and another that begins at 3pm that same Monday and lasts for one day. The first outage affects a generator capacity for 30 MW, while the second produces a capacity reduction of 20 MW. Starting at 3am on Monday, the total reduction is 30 MW. At 3pm, when the second outage occurs, the total reduction increases to 50 MW. On Tuesday at 3pm, the second outage is restored, that generator returns to full capacity, and the total capacity reduction goes back to 30 MW until the first generator is returned to full capacity at 3am on Thursday.

Figure 3: Emergency outages of coal generators

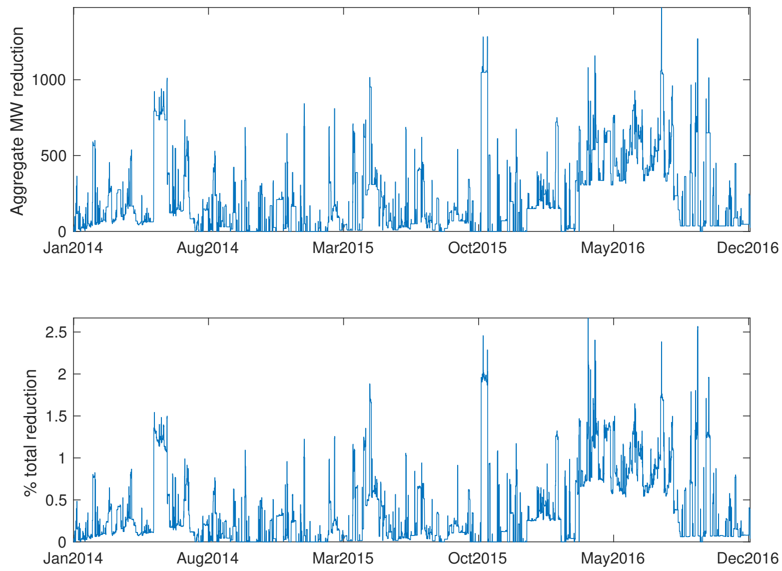

Note: The figure displays time series of the aggregate reduction in coal generation capacity due to emergency outages (top panel) and the aggregate reduction in coal generation capacity as percent of total *available* coal capacity, with available capacity defined as total installed capacity net of all *planned* outages (i.e., generators taken offline for scheduled maintenance and network upgrades) (bottom panel). Only emergency outages of coal generators are considered. All data is from PJM and covers the period from January 1, 2014 through December 31, 2016.

In the top panel, for each hour of the day, we display the total MW-reduction in coal generation capacity due to emergency outages. There is substantial variability in the data, with peak capacity reductions over 1,000 MW. The bottom panel plots outage-based capacity reductions as a percent of total available PJM coal capacity (i.e., total installed capacity minus total capacity that is not available due to planned outages), which is what we use in our empirical tests.

5 Volatility Estimation

One interesting consequence of increasing utilization of renewable energy sources is negative electricity prices. In our sample, about 5% of hours contain at least one negative price, but that fraction is likely to be much higher in more recent data. Negative prices complicate the estimation of volatility considerably: in particular, the classic definition of volatility as the standard deviation of percentage price changes becomes economically problematic. Consider the following example: over two consecutive five-minute intervals, electricity price goes from 3 to -2, and then to 1. The percentage return for the first five minutes is $(-2-3)/3 = -167\%$. This makes some intuitive sense, although it is not obvious what a loss greater than 100% means. The cash position involves buying electricity for 3 and selling it at a loss for -2 (the agent has to pay to inject electricity), so the total cashflow is -5 (a loss). Relative to the initial investment of 3, the asset has lost more than one and a half times its value. For the second five-minute interval, the initial price is -2, the final price is 1, and the return is $(1-(-2))/-2 = -150\%$. The calculated percentage return is negative, but the economic return is positive, since the cash flow is positive. Zero prices compound the challenge of calculating returns by giving rise to infinite returns. Moving from percentage returns to a natural logarithm of gross returns does not help. Thus, we rely on the price range as a measure of volatility.

The use of range is quite popular in economics, finance, and statistics. See for example, Feller (1951), Garman and Klass (1980), Rogers et al. (1991), Gallant et al. (1999), and Alizadeh et al. (2002). In our empirical analysis we estimate volatility as the ratio of the hourly range between maximum and minimum prices, divided by the average price over the previous twenty four hours. The numerator is a range in levels similar to traditional log range (see Garman and Klass, 1980). We scale the range by an average to provide economic grounding: a range of \$100/MWh has a different interpretation when the price hovers around \$30/MWh than when it hovers around \$800/MWh. We conduct a simulation exercise to show how our volatility estimator compares to other estimators, following the framework of Molnar (2012).

We begin with an asset price that follows a drift-less diffusion, $dp_t = \sigma dB_t$, and modify it to allow for price jumps. We start by simulating the log price, then exponentiate it to obtain a time-series of prices in levels. Next, we add jumps: with 0.2% probability the price might jump up by an average of \$150/MWh, and with 0.2% probability the price might jump down to an average price level of –\$50/MWh (the average of the negative prices in our data sample). Note that this is not exactly a jump diffusion process. See Fanone et al. (2013) for a more sophisticated price model.

We start with an initial price level of \$27/MWh (i.e., the median price in Table 1), and a diffusive volatility parameter $\sigma = 0.02/\sqrt{12}$, so the hourly diffusive volatility is 2%. We simulate the price process in discrete five-minute increments. We simulate 345,600 consecutive prices, obtaining approximately 28,800 hours or 1,200 days of data (roughly the length of our sample). We then construct "hourly" measures of volatility from twelve consecutive five-minute intervals. For each simulated hour, j, we compute the volatility estimators from Molnar (2012) as well as our proposed estimator:

- $v2_j = (H_j - L_j)/A_j$ is our scaled range estimator;
- $v1_j = \sqrt{\frac{1}{12}\sum_{t=1}^{12}(P_{j,t}/P_{j,t-1} - 1)^2}$ is the classic standard deviation of percentage return;
- $v3_j = \sqrt{\frac{1}{12}\sum_{t=1}^{12}(p_{j,t} - p_{j,t-1})^2}$ is the standard deviation of log returns;
- $v4_j = h_j - l_j$ is a log-range estimator;
- $v5_j = (h_j - l_j)/(2\sqrt{ln(2)})$ is Parkinson (1980) range-based estimator;
- $v6_j = \sqrt{0.5(h_j - l_j)^2 - (2ln(2) - 1)(c_j - o_j)^2}$ is Garman and Klass (1980) estimator that combines $v3_j$ and $v5_j$;
- $v7_j = \sqrt{(h_j - o_j)(h_j - c_j) - (l_j - o_j)(l_j - c_j)}$ is Rogers and Satchell (1991) estimator that allows for arbitrary drift.

Capital letters represent variables in levels, while lower-case letters represent logs. $P_{j,t}$ is the price of an asset (electricity) at the end of the five-minute interval t in hour j, H_j and L_j are the maximum and minimum electricity prices during hour j, A_j is the average electricity price of the 24 hours preceding hour j, and c_j and o_j are natural logarithms of the closing and opening electricity price for hour j.

We simulate the time-series of prices 1,000 times for each of 1,200 days. For each simulation we compute the six volatility measures for all 28,800 hours, and then compute the respective averages. We tabulate below the average and standard deviation across the 1,000 simulations (i.e., the average of an average, and the standard deviation of an average). We present three scenarios: the base case with no jumps in the top panel (i.e., the drift-less diffusion, in levels), the base case with only positive jumps in the middle panel, and the base case with positive and negative jumps in the bottom panel. Within each panel we also report the percentage of times when a simulation produces an hourly infinite or non-real volatility estimate (i.e., because of prices smaller than or equal to zero).

Consider the top panel. The standard deviation of price changes, in percentage ($v1$) or logs ($v3$), is pretty close to 2%. The two simple range estimators ($v2$ and $v4$) are close and slightly over-estimating volatility at 2.59%, while the rest of the estimators are under-estimating. The fact that the last three estimators are under-estimating volatility is in line with the results reported in Table 1 of Garman and Klass (1980), and due to the fact that in our simulations, as well as in the actual data, each volatility is estimated based on a small number of observations (i.e., 12 five-minute intervals). In their Table 1, Garman and Klass (1980) show that when prices are observed at "coarse" discrete intervals, their variance estimators are biased.[4]

Now consider the middle panel, where we introduce positive price jumps. Because there are no negative prices, all estimators are still feasible in every period. However, note that

[4] $v4$, $v5$, and $v6$ look better in Molnar (2012), but he uses 100,000 time intervals, not 12, to compute the volatility of one time period. Note that even at a daily level, with 12 x 24 time periods, volatility would be significantly under-estimated.

Table 3: Simulation of volatility estimators

	v1	v2	v3	v4	v5	v6	v7
	No jumps						
Average	1.96	2.59	1.96	2.59	1.55	1.43	1.32
Standard deviation	0.01	0.02	0.01	0.02	0.01	0.01	0.01
Percentage of missing	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	Positive jumps only						
Average	39.53	40.46	10.72	9.45	5.68	5.88	7.12
Standard deviation	31.17	28.80	2.21	1.72	1.03	1.10	1.44
Percentage of missing	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	Positive and negative jumps						
Average	52.15	49.69	10.72	9.45	5.68	5.88	7.12
Standard deviation	34.47	35.77	2.22	1.73	1.04	1.11	1.45
Percentage of missing	0.0	0.0	2.6	2.6	2.6	2.6	2.6

Note: The table shows comparisons of different volatility estimators in thee simulation scenarios. Each scenario reflects a different dynamic for the price of electricity. In the first scenario the price follows a driftless diffusion with an hourly diffusive parameter of 2%. In the second scenario with probability 0.2% the price can jump to an average of \$150. In the third scenario, with respective probability of 0.2%, the price to jump up to an average of \$150 or down to an average of –\$50. Each scenario represents 1,000 simulations of 28,800 hours (approximately the length of our actual sample). In each simulation hour the price is observed 12 times (i.e., the equivalent of 5-minute intervals). Using the 12 observations, for each generic hour j, we compute 7 measures of volatility: $v2_j = (H_j - L_j)/A_j$ is our scaled range estimator; $v1_j = \sqrt{\frac{1}{12}\sum_{t=1}^{12}(P_{j,t}/P_{j,t-1} - 1)^2}$ is the classic standard deviation of percentage return; $v3_j = \sqrt{\frac{1}{12}\sum_{t=1}^{12}(p_{j,t} - p_{j,t-1})^2}$ is the standard deviation of log returns; $v4_j = h_j - l_j$ is a log-range estimator; $v5_j = (h_j - l_j)/(2\sqrt{ln(2)})$ is Parkinson (1980) range-based estimator; $v6_j = \sqrt{0.5(h_j - l_j)^2 - (2ln(2) - 1)(c_j - o_j)^2}$ is Garman and Klass (1980) estimator that combines $v3_j$ and $v5_j$; $v7_j = \sqrt{(h_j - o_j)(h_j - c_j) - (l_j - o_j)(l_j - c_j)}$ is Rogers and Satchell (1991) estimator that allows for arbitrary drift. Capital letters represent variables in levels, while lower-case letters represent logs. $P_{j,t}$ is the price of an asset (electricity) at the end of the five-minute interval t in hour j, H_j and L_j are the maximum and minimum electricity prices during hour j, A_j is the average electricity price of the 24 hours preceding hour j, and c_j and o_j are natural logarithms of the closing and opening electricity price for hour j. For each of the 1,000 simulations, we then average the seven hourly volatilities across the 28,800 hours. The table reports the respective average and standard deviations across the 1,000 simulations, as well the percentage of simulated hours for which an estimator cannot be computed because of negative prices.

while $v2$ and $v4$ still track volatility measures based on squared returns, $v1$ and $v3$, respectively, the last three estimators produce a sizable bias.

In the bottom panel, adding negative prices complicates the interpretation and comparison of the numbers because the only volatility measure that does not have any mechanical problems is the one we suggest, $v2$. Even the simple standard deviation of returns has problems because it relies on returns that are not economically meaningful. Note that a 0.2%

probability of a negative jump does not mean that only 0.2% of the hours will have problems. In fact, because jumps are random, rarely is there more than one in an hour. However, just one negative price is enough to create a problem in the construction of a volatility estimator that uses log prices. The frequency of hours that have a problem in such cases will be much higher than 0.2%, around 2.6%. Intuitively, introducing negative jumps should increase price volatility, but we do not observe this for all log-price estimators because hours with negative prices need to be discarded. The volatility measured by our proposed estimator, however, does increase, and so does volatility based on the square of percentage returns.

Overall, this simulation shows that when prices are negative, most volatility estimators run into some problem, whether mechanical or logical. Hence our choice to focus on $v2$ in the rest of the empirical analysis.

6 Empirical analysis

6.1 Emergency outages

We begin our analysis by examining how emergency coal outages are related to the state of the market. We tabulate regression results of an hourly measure of outage severity, controlling for several proxies of electricity demand, demand uncertainty, and a measure of wind generation. The dependent variable is defined as the natural logarithm of the total reduction in coal generation capacity as percent of total available coal capacity (i.e., total installed capacity minus the aggregate reduction due to scheduled outages), multiplied by one hundred. All data are hourly. We also impose one important restriction on our sample: because we are interested in measuring the trade-off between coal and natural gas, we only look at hours during which the marginal generator is a gas and/or coal unit. We exclude hours when other fuel sources are also on the margin. Results of these regressions are shown in Table 4. All coefficient estimates are multiplied by 100 for ease of interpretation.

Emergency outages that affect coal generators are positively related to unforeseen demand

Table 4: Emergency coal outages and power network conditions

	All		Non-fire		Fire	
	(1)	(2)	(3)	(4)	(5)	(6)
Load	0.095	0.085	-0.035	-0.038	0.125	0.117
	(0.06)	(0.06)	(0.02)	(0.02)	(0.06)	(0.06)
Unforecasted load	0.527	0.569	0.004	0.015	0.522	0.553
	(0.22)	(0.22)	(0.06)	(0.06)	(0.20)	(0.20)
Wind generation	0.003	0.003	-0.003	-0.003	0.006	0.006
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)
Cold	-0.084	-0.074	-0.015	-0.012	-0.068	-0.061
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Hot	0.078	0.091	0.019	0.022	0.059	0.068
	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)
Weekend	0.002	-0.001	0.007	0.006	-0.005	-0.008
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)
Generator Age	0.145	0.145	0.026	0.025	0.093	0.090
	(0.11)	(0.11)	(0.01)	(0.01)	(0.05)	(0.05)
Past electricity volatility		-0.010		-0.003		-0.007
		(0.00)		(0.00)		(0.00)
R-squared	0.156	0.157	0.146	0.146	0.186	0.186
Observations	11826	11826	11826	11826	11826	11826

Note: The table reports estimated coefficients, standard errors and R-squared coefficients for linear regressions of outage-related reductions in coal generation capacity on measures of electricity demand, supply, weather conditions, and time. All variables are measured at hourly intervals. The dependent variable is the natural logarithm of the percentage reduction in PJM's total available coal generation capacity. Independent variables include the natural logarithm of hourly load, the difference between realized hourly load and hourly load forecast generated by PJM one day in advance, and total wind power generation. Among the independent variables, we also include an indicator variable equal to one when the average temperature in the counties serviced by PJM is less than negative five degrees Celsius (Cold), an indicator variable equal to one when the average temperature in the counties serviced by PJM is more than twenty five degrees Celsius (Hot), and an indicator variable equal to one for observations that fall on a weekend. Generator Age measures the years (as a fraction of a century) from the date of the outage to the ultimate construction of the plan. We also include the one-hour lag of electricity price volatility, defined as the intra-hour range between minimum and maximum prices (scaled by the average price in the previous twenty four hours). All estimated coefficients are multiplied by 100. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. For this table we only consider hours in which either gas and/or coal are on the margin. Data covers the period from January 1, 2014 through December 31, 2016.

for electricity and are more likely on hot days. This is not unexpected, as half of all emergency outages (and over 90 percent of emergency coal outages) are related to fires, which are aided by high temperatures and high operating capacities.

Outages are less likely to occur on cold days (but only those caused by a fire) and when electricity prices are more volatile. In all cases, the size of the effects is quite small. Coefficients on unforecasted load for all outages (columns 1 and 2) imply that a one percent increase in unforecasted load is associated with an increase in the coal capacity displaced by an outage of one half of a percent. Coefficient estimates on the Hot indicator suggest that on hot days the capacity reduction due to emergency coal outages is about one tenth of a percent larger than on an average temperate day. Thus, although some of our coefficients are statistically significant, in terms of economic significance, coal outages appear to be reasonably exogenous events in the context of our analysis. The negative and very small coefficient on lagged electricity price volatility also supports the idea that emergency coal outages cannot be predicted in a meaningful way by price data.

Looking at fire and non-fire outages separately offers some interesting insights. Emergency outages caused by operational factors other than a fire appear to be negatively related to load and are not related to the unforecasted component of demand. Fire-related outages, by contrast, are linked to both total demand and unanticipated demand spikes, both of which may make unit overloads more likely. Capacity losses from non-fire outages are also lower when wind generation is plentiful and on weekdays—the opposite, although not statistically significant, relationship from that observed for fire-related outages. Also interesting, the age of the generator is positively related to the amount of capacity being displaced, but the effect is stronger for non-fire outages.

These results suggest that there may be systematical differences between fire and non-fire emergency outages, but the impacts of these differences are small in an economic sense, and thus, not likely strong enough to violate the ignorability condition. Thus, for most subsequent analyses, we group fire and non-fire outages together and discuss differences between these two outage categories in the robustness section.

6.2 Outages, fuel switching, and real-time price volatility

Next, we investigate the relationship between price volatility and marginal fuel source, employing an instrumental variable approach described in Section 3. We report results of the first stage regressions in Table 5 and results of ordinary least squares and second stage IV regressions (2SLS) in Table 6. We calculate real-time electricity price volatility, henceforth volatility, as the natural logarithm of the ratio of the range between the minimum and the maximum prices (observed in the hour) and the average electricity price from the previous 24 hours. We use 24 to dampen the effect of negative prices, but results are robust to alternative choices (i.e., 1, 6, or 12 hour averages).

First stage 2SLS results displayed in Table 5 show that a one percent reduction in total available coal generation capacity increases the fraction of time during which natural gas is on the margin by 1.9 percentage points. These coefficient estimates are statistically significant and easily pass the weak instrument test with F-statistics above 36. The strength of the instrument supports the assumption that an outage that takes out an infra-marginal coal generator is unlikely to push another coal generator on the margin. Because the only way in which we could make an independent validation of the monotonicity assumption is by using generators' actual bids, which we do not have, we rely on the statistical strength of the instrument.

We can provide context to our main first-stage coefficient estimate by calculating the marginal effect of a one standard deviation change in our main explanatory variable. For emergency coal outages, one standard deviation in average capacity loss is 0.46 percent. Such increase in lost coal capacity would put gas generators on the margin for an extra $0.46 \times 1.9 = 0.9$ percent of an hour. Thus, emergency coal outages lead to fuel switching, which confirms the validity of our chosen instrumental variable.

Moving to Table 6, OLS regressions in columns (1) and (2) show a positive relationship between gas on the margin and volatility of real-time electricity prices. Columns (3) and (4) show that once we account for endogeneity of the marginal fuel source, the relationship

Table 5: Marginal shares of natural gas generators

	(1)	(2)
Coal outage MW reduction over total capacity	1.901	1.911
	(0.397)	(0.398)
	[36.69]	[37.04]
Load	-0.198	-0.196
	(0.021)	(0.021)
Unforecasted load	-0.097	-0.102
	(0.073)	(0.073)
Wind generation	0.006	0.006
	(0.002)	(0.002)
Cold	0.023	0.022
	(0.014)	(0.014)
Hot	0.021	0.020
	(0.006)	(0.007)
Weekend	-0.040	-0.040
	(0.004)	(0.004)
Past electricity volatility		0.001
		(0.001)
R-squared	0.069	0.069
Observations	11826	11826

Note: The table shows estimated coefficients, standard errors and R-squared statistics for linear regressions of the natural logarithm of one plus the time (percentage of an hour) during which natural gas generators are on the margin. The main independent variable is the natural logarithm of the relative reduction in available generating capacity due to a coal outage. Weak instrument F-tests are reported in brackets. Other independent variables are as described in Table 4. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. For this table we only consider hours in which either gas and/or coal are on the margin. Data covers the period from January 1, 2014 through December 31, 2016.

between marginal fuel and price volatility becomes negative and statistically significant. That is, electricity price volatility tends to be *lower* when more gas is on the margin.

In economic terms, switching the marginal generator from coal to natural gas for an additional one percent of an hour (36 seconds) decreases the scaled price range by 3.6 percent (from column 4). Evaluated at the mean, this fuel switching corresponds to a 5.7 percent reduction in the average power price volatility (from 63 percent to 59.4 percent), or a 77-cent reduction in the intra-hour range of electricity prices.

A negative relationship between electricity price volatility and gas on the margin is con-

Table 6: Real time electricity price volatility

	OLS		2SLS	
	(1)	(2)	(3)	(4)
Gas on margin	0.396	0.365	-5.612	-3.596
	(0.088)	(0.067)	(1.790)	(1.264)
Load	-0.723	-0.265	-1.900	-1.037
	(0.164)	(0.116)	(0.388)	(0.270)
Unforcasted load	4.480	2.574	3.959	2.215
	(0.591)	(0.409)	(0.609)	(0.425)
Wind generation	-0.017	-0.010	0.020	0.015
	(0.016)	(0.011)	(0.019)	(0.014)
Cold	0.964	0.515	1.092	0.595
	(0.108)	(0.073)	(0.113)	(0.077)
Hot	1.213	0.667	1.350	0.753
	(0.061)	(0.046)	(0.072)	(0.054)
Weekend	-0.283	-0.134	-0.524	-0.291
	(0.033)	(0.023)	(0.078)	(0.055)
Past electricity volatility		0.446		0.450
		(0.010)		(0.010)
R-squared	0.192	0.355	0.191	0.354
Observations	11826	11826	11826	11826

Note: The table shows estimated coefficients, standard errors and R-squared statistics for linear regressions of hourly real-time electricity price volatility on the percentage of time natural gas generators are on the margin in that hour. The dependent variable is the hourly range in real-time wholesale electricity price, defined as the difference between the maximum and minimum prices observed within a given hour, divided by the average price observed in the previous twenty four hours. Each specification is estimated by ordinary least squares and by 2SLS, where the proportion of time natural gas is on the margin is instrumented by the percentage reduction in total available generation capacity of coal due to an emergency outage. We only include observations where coal and/or natural gas are the marginal generation sources. Other independent variables are as described in Table 4. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. Data covers the period from January 1, 2014 through December 31, 2016.

sistent with the model of Reguant (2014) in which fixed costs are key determinants of the dynamic bidding behavior of electricity generators. Because of lower startup costs, natural gas generators are able to respond to variation in demand and other market conditions more efficiently, and therefore, relative to coal generators, they marginally reduce electricity price volatility. Our results thus cast significant doubt on the argument that fuel price volatility passed through in the generation process has a pervasive effect on electricity price volatility. At the minimum, the adverse impact of fuel price volatility appears to be secondary to the benefits of dispatching more flexible generating assets.

Our results however do not contradict those of Chu et al. (2017), for example, who find a substantial cost pass-through from natural gas prices, due to procurement contracts *periodically* readjusting to natural gas prices. It is likely that fuel procurement contracts are relatively long dated forwards. When such contracts are renegotiated, there would be a substantial amount of cost pass-through, but it would not translate in volatility pass-through as long as the frequency of renegotiation is low enough.

Although our analysis is based on PJM data, our results should hold for other electricity markets in which generators are able to revise bids dynamically to participate in the real time market.

In the next section we investigate how our results hold up to alternative definitions of our main dependent and independent variables, as well as alternative mechanisms behind fuel switching.

6.3 Robustness checks

6.3.1 Variable definitions and samples

Table 7 reports the results of several robustness checks. Changing the definition of volatility from scaled range to standard deviation of price, does not change the basic result: a one percent increase in natural gas on the margin leads to a 3.3 percent decrease in price volatility (compared to a 3.5 percent decrease in Table 6). Using unscaled emergency outages generation reduction in the first stage leads to similar second stage outcomes as in our main results. By splitting the sample according to the intensity of wind generation, we learn that gas on margin impacts volatility primarily during times of low wind generation: when little wind energy is being fed into the grid, a reduction in coal generation due to an outage produces a larger switch to natural gas. By contrast, when wind generation is plentiful, not only is gas more likely to be on the margin already, but because inframarginal wind generation pushes some coal units out of dispatch, more coal capacity is left available, reducing the need to dispatch additional flexible natural gas units. Our baseline result is also unaffected by considering all hourly observations, as opposed to only those in which marginal generation comes exclusively from natural gas or coal.

Table 7: Robustness checks

	Scaled Price Vol		MW Reduction		Low Wind		High Wind		Full Sample			
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Gas on margin	0.353	-3.338	0.365	-3.508	0.463	-7.962	0.309	-0.833	0.203	-5.176	0.426	-4.343
	(0.068)	(1.281)	(0.067)	(0.973)	(0.097)	(2.566)	(0.092)	(1.576)	(0.051)	(1.394)	(0.055)	(1.116)
Load	-0.262	-0.981	-0.265	-0.669	-0.504	-1.631	0.101	-0.240	0.008	-1.009	-0.010	-0.714
	(0.119)	(0.274)	(0.116)	(0.146)	(0.163)	(0.374)	(0.174)	(0.502)	(0.080)	(0.276)	(0.080)	(0.184)
Unforcasted load	2.647	2.312	2.574	1.978	3.494	3.776	1.639	1.422	3.158	2.384	3.052	2.815
	(0.422)	(0.438)	(0.409)	(0.442)	(0.632)	(0.648)	(0.546)	(0.621)	(0.305)	(0.364)	(0.303)	(0.310)
Wind generation	-0.011	0.012	-0.010	-0.005					-0.010	-0.027	-0.023	0.005
	(0.011)	(0.014)	(0.011)	(0.011)					(0.008)	(0.009)	(0.008)	(0.011)
Cold	0.522	0.597	0.515	0.524	0.743	1.051	0.385	0.421	0.553	0.583	0.529	0.632
	(0.075)	(0.079)	(0.073)	(0.073)	(0.131)	(0.159)	(0.084)	(0.097)	(0.043)	(0.044)	(0.043)	(0.049)
Hot	0.680	0.760	0.667	0.681	0.671	0.854	0.550	0.569	0.672	0.769	0.664	0.764
	(0.048)	(0.055)	(0.046)	(0.047)	(0.062)	(0.083)	(0.077)	(0.082)	(0.034)	(0.043)	(0.034)	(0.042)
Weekend	-0.129	-0.275	-0.134	-0.149	-0.136	-0.367	-0.120	-0.185	-0.129	-0.289	-0.114	-0.282
	(0.024)	(0.056)	(0.023)	(0.023)	(0.033)	(0.077)	(0.033)	(0.095)	(0.017)	(0.045)	(0.017)	(0.043)
Past electricity volatility	0.451	0.455	0.446	0.445	0.444	0.465	0.440	0.441	0.454	0.450	0.451	0.459
	(0.010)	(0.010)	(0.010)	(0.010)	(0.013)	(0.015)	(0.013)	(0.013)	(0.007)	(0.007)	(0.007)	(0.007)
Others on margin											0.543	-1.230
											(0.053)	(0.417)
R-squared	0.350	0.349	0.355	0.354	0.372	0.371	0.348	0.346	0.370	0.369	0.372	0.371
Observations	11826	11826	11826	11826	5913	5913	5913	5913	23285	23285	23285	23285

Note: The table presents robustness checks to the results presented in Table 6. In the first two columns we consider the natural volatility of price volatility (as opposed to price range) as the dependent variable. In columns three and four, we consider the natural logarithm of power generation reduction due to a coal emergency outage (as opposed to the relative power reduction) as the instrumental variable for the percentage of an hour that natural gas is on the margin. In columns five and six (seven and eight), we limit the sample to hours when wind generation is below (above) the median. In columns nine to twelve, we consider all hourly observation irrespective of which fuel source is on the margin (as opposed to only consider hours for which coal and/or natural gas are on the margin). Each specification is estimated by ordinary least squares and by 2SLS. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. Data covers the period from January 1, 2014 through December 31, 2016.

6.3.2 Non-fire emergency outages

Because emergency outages of coal generators are related to demand and temperature conditions, even if marginally, it is reasonable to wonder whether the units that are eventually forced to shut down because of a fire are driven to that critical level by market conditions, thus violating the orthogonality condition.

Hence, we repeat the analysis for fire- and non-fire-related outages separately. One caveat is that fire-related outages are overwhelmingly more frequent and thus displace a much larger portion of the existing coal generation capacity, even though a single fire-related emergency outage leads to a lower average capacity reduction than a single non-fire-related outage (see Table 2). Results are reported in Table 8.

Because of their relatively small number, non-fire-related emergency outages prove to be a weaker instrument than fire-related outages. Splitting emergency outages by cause has only a small effect on the first and second-stage results: a one standard deviation increase in lost coal capacity is associated with a 0.48 percent (non-fire) and 0.75 percent (fire) increase in natural gas on the margin; price volatility decreases by 2.1 and 2.6 percent as a result of non-fire and fire-related emergency outages, respectively.

6.3.3 Fuel price volatility

As Chu et al. (2017) point out, there is a certain amount of cost pass-through in the electricity generation industry. It is therefore reasonable to assume that volatility in fuel markets might transfer to the wholesale electricity market. We next check for evidence that price volatilities in the fuel and electricity markets are connected, and if so, whether that connection affects our main results.

Because we do not have access to high-frequency data on natural gas and coal prices, we construct measures of fuel price volatility by filtering daily percentage price changes (i.e., returns) through a GARCH(1,1) model.[5] We then match these daily volatilities to the rest of

[5]This is the model of best fit, though other GARCH specifications produce similar results.

Table 8: Fire versus non-fire related outages

	Non-fire outages		Fire outages	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
	(1)	(2)	(3)	(4)
Coal outage MW reduction over total capacity	4.347		1.880	
	(1.49)		(0.43)	
	[13.29]		[30.67]	
Gas on margin		-4.266		-3.477
		(2.16)		(1.39)
Load	-0.193	-1.167	-0.197	-1.013
	(0.02)	(0.44)	(0.02)	(0.29)
Unforcasted load	-0.091	2.154	-0.101	2.225
	(0.07)	(0.46)	(0.07)	(0.43)
Wind generation	0.006	0.019	0.006	0.014
	(0.00)	(0.02)	(0.00)	(0.01)
Cold	0.021	0.609	0.021	0.593
	(0.01)	(0.08)	(0.01)	(0.08)
Hot	0.021	0.768	0.020	0.750
	(0.01)	(0.07)	(0.01)	(0.05)
Weekend	-0.040	-0.318	-0.040	-0.286
	(0.00)	(0.09)	(0.00)	(0.06)
Past electricity volatility	0.001	0.450	0.001	0.450
	(0.00)	(0.01)	(0.00)	(0.01)
R-squared	0.067	0.354	0.068	0.354
Observations	11826	11826	11826	11826

Note: The table presents robustness checks to the results presented in Tables 5 and 6. Columns 1 and 2, display the results obtained using only non-fire related emergency outages as the instrumental variable. Column 1 contains the estimated coefficients from the first stage of the 2SLS procedure, where the dependent variable is the proportion of an hour that natural gas is on margin. Column 2, tabulates the results of the second stage, where real-time electricity price volatility is the dependent variable, and natural gas on margin is the fitted estimate from the first stage. In columns 3 and 4, we tabulate corresponding results obtained by only considering fire-related emergency outages. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. For this table we only consider hours in which either gas and/or coal are on the margin. Data covers the period from January 1, 2014 through December 31, 2016.

our hourly variables. Some caution in interpretation is merited. Because fuel price volatility measures are not hourly, they absorb some of the day-to-day variation in electricity prices, which in turn affects their coefficients and their impact on other coefficients. Table 9 displays the results of this analysis.

We find some evidence that price volatilities in the fuel and electricity markets are connected, but this relationship is negative. Given our empirical context and framework, despite the limitations presented by our measures of fuel price volatility, we do not find support for the hypothesis that high volatility in fuel markets is associated with high volatility in the real-time electricity market. (See the next subsection for evidence from daily day-ahead volatility). Coefficients on both fuel prices are negative, and those estimated for natural gas price volatility are statistically significant. Comparing these results to the baseline results reported in column (4) of Table 6, we find that the relationship between fuel and electricity price volatility does not affect our main result.

6.4 Daily electricity price volatility

One relevant question is whether our results only apply to hourly volatility or whether they are detectable at longer horizons. We adopt the hourly time-scale in the main analysis because it gives us the best identification. While hourly seems the more appropriate time dimension to think about the real-time market, it is true that the bulk of generation is transacted in the day-ahead market. Thus, keeping in mind that we lose the tight identification, we repeat our analysis at the daily frequency, computing volatility from day-ahead market data (i.e., 24 hour prices are used to compute one daily volatility). Because there are no negative prices in the day-ahead market, we can measure volatility as the standard deviation of percentage price changes. This offers a validation exercise for the range-based volatility measure that we adopt in our main analysis.

Results are reported in Table 10, and confirm the conclusions of our main analysis: more gas on margin is associated with lower volatility of electricity prices; there is no evidence of volatility pass-through from the fuel spot markets; without the complications of negative prices, scaled range and standard deviation of price changes produce very similar results.

We note that our evidence against volatility pass-through, does not invalidate the results of Chu et al. (2017), for example, who find a substantial cost pass-through from natural gas prices, due to procurement contracts *periodically* readjusting to natural gas prices. The

Table 9: Impact of fuel price volatility

	2SLS volatility regression – 2^{nd} stage			
	(1)	(2)	(3)	(4)
Gas on margin	-3.596	-4.368	-3.071	-3.909
	(1.264)	(1.457)	(1.621)	(1.884)
Load	-1.037	-1.120	-0.907	-1.012
	(0.270)	(0.288)	(0.356)	(0.386)
Unforcasted load	2.215	1.989	2.202	1.990
	(0.425)	(0.445)	(0.424)	(0.446)
Wind generation	0.015	0.019	0.013	0.018
	(0.014)	(0.014)	(0.014)	(0.015)
Cold	0.595	0.661	0.587	0.652
	(0.077)	(0.086)	(0.078)	(0.091)
Hot	0.753	0.757	0.742	0.748
	(0.054)	(0.054)	(0.058)	(0.059)
Weekend	-0.291	-0.317	-0.266	-0.296
	(0.055)	(0.061)	(0.071)	(0.080)
Past electricity volatility	0.450	0.450	0.449	0.449
	(0.010)	(0.010)	(0.010)	(0.010)
Gas volatility		-0.064		-0.061
		(0.025)		(0.027)
Coal volatility			-0.042	-0.034
			(0.041)	(0.043)
R-squared	0.354	0.354	0.355	0.355
Observations	11826	11826	11826	11826

Note: The table presents robustness checks to the results presented in Tables 5 and 6 by including a measure of volatility in natural gas and coal spot prices. We construct fuel price volatility measures by filtering daily percentage changes in spot prices (i.e., returns) through a GARCH model. We then expand those so that they pair with each hour of the corresponding day. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. For this table we only consider hours in which either gas and/or coal are on the margin. Data covers the period from January 1, 2014 through December 31, 2016.

coarser the timing of the readjustment, the less likely it is that a cost pass-through would produce a volatility pass-through.

Table 10: Day-ahead market volatility

	Scaled Range Volatility					
	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Gas on margin	-0.467	-0.424	-0.447	-3.147	-1.792	-1.311
	(0.082)	(0.083)	(0.084)	(1.255)	(0.517)	(0.595)
Load	0.515	0.524	0.456	0.782	0.647	0.599
	(0.091)	(0.117)	(0.113)	(0.153)	(0.093)	(0.097)
Unforcasted load	-0.436	-0.630	-0.373	-2.122	-1.270	-0.967
	(0.679)	(0.699)	(0.705)	(1.045)	(0.688)	(0.699)
Wind generation	0.001	0.001	0.001	-0.046	-0.022	-0.014
	(0.009)	(0.009)	(0.009)	(0.023)	(0.012)	(0.012)
Cold	0.118	0.148	0.118	0.002	0.061	0.081
	(0.041)	(0.043)	(0.042)	(0.071)	(0.054)	(0.056)
Hot	0.194	0.187	0.193	0.121	0.158	0.171
	(0.020)	(0.019)	(0.019)	(0.039)	(0.022)	(0.024)
Weekend	-0.073	-0.068	-0.072	0.035	-0.020	-0.039
	(0.014)	(0.014)	(0.014)	(0.053)	(0.024)	(0.027)
Past electricity volatility		0.041	0.053		-0.035	-0.052
		(0.042)	(0.042)		(0.039)	(0.041)
Gas volatility		-0.033			-0.104	
		(0.011)			(0.039)	
Coal volatility			-0.010			-0.091
			(0.017)			(0.041)
R-squared	0.408	0.415	0.409	0.391	0.398	0.392
Observations	1016	1016	1016	1016	1016	1016

7 Conclusion and policy implications

After accounting for endogeneity, more natural gas generation *per se* does not seem to be driving higher price volatility in PJM's real-time (or day-ahead) wholesale market for the years 2014 through 2016. All else equal, more natural gas generation on the margin is likely to lead to lower, not higher wholesale price risk for electricity market participants. This is important for electricity market planning in general and natural gas integration in particular because retail power prices in most regions of the U.S. are fixed and utilities are unable to pass through changes in risk exposure in wholesale markets to retail consumers. As the system

Table 10 continued	Standard Deviation of Percentage Price Changes					
	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Gas on margin	-0.364	-0.266	-0.280	-2.898	-3.524	-3.728
	(0.061)	(0.056)	(0.058)	(1.002)	(0.318)	(0.426)
Load	0.122	0.031	-0.007	0.375	0.437	0.458
	(0.074)	(0.079)	(0.073)	(0.128)	(0.070)	(0.071)
Unforcasted load	-0.395	-0.303	-0.149	-1.989	-2.384	-2.512
	(0.498)	(0.469)	(0.465)	(0.836)	(0.540)	(0.584)
Wind generation	-0.000	-0.005	-0.005	-0.045	-0.056	-0.059
	(0.007)	(0.006)	(0.006)	(0.019)	(0.008)	(0.009)
Cold	0.115	0.118	0.100	0.005	-0.022	-0.031
	(0.023)	(0.023)	(0.022)	(0.050)	(0.027)	(0.028)
Hot	0.049	0.045	0.049	-0.020	-0.037	-0.043
	(0.020)	(0.018)	(0.017)	(0.034)	(0.019)	(0.020)
Weekend	-0.024	-0.017	-0.019	0.078	0.103	0.111
	(0.012)	(0.011)	(0.012)	(0.042)	(0.018)	(0.021)
Past electricity volatility		0.302	0.315		0.075	0.082
		(0.042)	(0.040)		(0.023)	(0.026)
Gas volatility		-0.019			-0.127	
		(0.008)			(0.021)	
Coal volatility			-0.005			-0.133
			(0.012)			(0.021)
R-squared	0.168	0.252	0.246	0.142	0.232	0.225
Observations	1016	1016	1016	1016	1016	1016

Note: The table presents robustness checks to the results presented in Tables 5 and 6. The dependent variable is the daily volatility constructed from the 24 day ahead prices. Volatility is constructed using the same definition as in the paper (Maximum Price - Minimum Price)/Average Price (Scaled Range) or as the standard deviation of hourly returns (Standard Deviation). Gas on Margin is aggregated to daily from hourly data. It is computed as the average of the 24 measurements provided by PJM for each of the hours in the day, weighted by the load forecast for each hour. Load and Wind are summed across each hour of the day, while unforcasted load is the average of the hourly forecast errors. Daily capacity forced out by a coal outage is the daily average of the hourly variable used in the main analysis (i.e., the ratio of the capacity that is forced out and the difference between the total capacity and the capacity that was scheduled to be out). We construct fuel price volatility measures by filtering daily percentage changes in spot prices (i.e., returns) through a GARCH model. Standard errors are corrected for heteroskedasticity and serial correlation in residuals. A constant is estimated but not reported. Each specification includes month and hour fixed effects. Data covers the period from January 1, 2014 through December 31.

operator dispatches more flexible power generating units, the system's ability to respond to unanticipated changes in market conditions improves, which in turn leads to lower price volatility. By reducing response times and ramping constraints, more flexible natural gas generation appears to reduce bid markups created by generators' startup constraints.

Our results should hold for other markets with supply bidding rules and generation stacks similar to PJM. Our analysis does not, however, allow us to comment on the welfare implications of gas integration beyond fuel switching at the margin—that is, we cannot say how additional gas integration affects price volatility once gas is on the margin 100 percent of the time. The potential benefits of fuel-switching, however, should be taken into account by policy makers in power-sector market planning.

Acknowledgements

The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Dallas, or its staffs. We thank seminar participants at Texas A&M Energy Law and Policy workshop, 2017 USAEE conference, 2018, CU Environmental and Resource Economics workshop, 2018 EMEE workshop, 2018 APPAM fall research conference for useful comments and suggestions. All errors are our own.

References

Aid, R., G. Chemla, A. Porchet, and N. Touzi (2011). "Hedging and vertical integration in electricity markets" *Management Science* 57: 1438–1452.

Alizadeh, S., M. W. Brandt, and F. X. Diebold (2002). "Range-based estimation of stochastic volatility models" *The Journal of Finance* 57(3): 1047–1091.

Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996). "Identification of causal effects using instrumental variables" *Journal of the American Statistical Association* 91(434): 444–455.

Banerjee, S. and T. H. Noe (2006). "Exotics and electrons: electric power crises and financial risk management" *The Journal of Business* 79(5): 2659–2696.

Bessembinder, H. and M. L. Lemmon (2002). "Equilibrium pricing and optimal hedging in electricity forward markets" *The Journal of Finance* 57(3): 1347–1382.

Borenstein, S. (2002). "The Trouble With Electricity Markets: Understanding California's Restructuring Disaster" *Journal of Economic Perspectives* 16(1): 191–211 doi: 10.1257/0895330027175.

Boroumand, R.-H., S. Goutte, K. Guesmi, and T. Porcher (2019). "Potential benefits of optimal intra-day electricity hedging for the environment: The perspective of electricity retailers" *Energy Policy* 132: 1120–1129.

Boroumand, R.-H. and G. Zachmann (2012). "Retailers' risk management and vertical arrangements in electricity markets" *Energy Policy* 40: 465–472.

Brown, J. P. and A. Kodaka (2014). "U.S. electricity prices in the wake of growing natural gas production" *Main Street Economist* : 1–8.

Bunn, D. W. and D. Chen (2013). "The forward premium in electricity futures" *Journal of Empirical Finance* 23: 173–186.

Chu, Y., J. S. Holladay, and J. LaRiviere (2017). "Opportunity cost pass-through from fossil fuel market prices to procurement costs of the U.S. power producers" *The Journal of Industrial Economics* 65(4): 842–871.

Ciarreta, A., P. Muniain, and A. Zarraga (2017). "Modeling and forecasting realized volatility in German-Austrian continuous intraday electricity prices" *Journal of Forecasting* 36(6): 680–690.

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

Deyette, J., S. Clemmer, R. Cleetus, S. Sattler, A. Bailie, and M. Rising (2015). "The natural gas gamble: a risky bet on America's clean energy future" Technical report Union of Concerned Scientists.

Dupuis, D. J., G. Gauthier, and F. Godin (2016). "Short-term hedging for an electricity retailer" *The Energy Journal* 37: 31–59.

Fabrizio, K. R., N. L. Rose, and C. D. Wolfram (2007). "Do markets reduce costs? Assessing the impact of regulatory restructuring on U.S. electric generation efficiency" *American Economic Review* 97(4): 1250–1277.

Fanone, E., A. Gamba, and M. Prokopczuk (2013). "The case of negative day-ahead electricity prices" *Energy Economics* 35: 22–34.

Fehrenbacher, K. (2010). "Smart Grid 101: Utilities are risk averse" .

Feller, W. (1951). "The asymptotic distribution of the range of sums of independent random variables" *Annals of Mathematical Statistics* 22: 427–432.

Frommel, M., X. Han, and S. Kratochvil (2014). "Modeling the daily electricity price volatility with realized measures" *Energy Economics* 44: 492–502 ISSN 0140-9883 doi: https://doi.org/10.1016/j.eneco.2014.03.001.

Gallant, A. R., C. T. Hsu, and G. E. Tauchen (1999). "Using daily range data to calibrate volatility diffusions and extract the forward integrated variance" *Review of Economics and Statistics* 81: 617–631.

Garman, M. B. and M. J. Klass (1980). "On the Estimation of Security Price Volatilities from Historical Data" *The Journal of Business* 53: 67–78.

Haugom, E. and C. J. Ullrich (2012). "Forecasting spot price volatility using the short-term forward curve" *Energy Economics* 34: 1826–1833 doi: https://doi.org/10.1016/j.eneco.2012.07.017.

Haugom, E., S. Westgaard, P. B. Solibakke, and G. Lien (2011). "Realized volatility and the influence of market measures on predictability: Analysis of Nord Pool forward electricity data" *Energy Economics* 33(6): 1206–1215 ISSN 0140-9883 doi: https://doi.org/10.1016/j.eneco.2011.01.013.

Hickey, E., D. G. Loomis, and H. Mohammadi (2012). "Forecasting hourly electricity prices using ARMAX-GARCH models: An application to MISO hubs" *Energy Economics* 34(1): 307–315 ISSN 0140-9883 doi:https://doi.org/10.1016/j.eneco.2011.11.011.

Imbens, G. W. and J. D. Angrist (1994). "Identification and estimation of local average treatment effects" *Econometrica* 62(2): 467–475.

Imbens, G. W. and D. R. Rubin (1997). "Bayesian inference for casual effects in randomized experiments with non compliance" *Annals of Statistics* 25(1): 305–327.

ISO New England (2018). "ISO New England key grid and market stats" Online report.

Ketterer, J. C. (2014). "The impact of wind power generation on the electricity price in Germany" *Energy Economics* 44: 270–280 ISSN 0140-9883 doi: https://doi.org/10.1016/j.eneco.2014.04.003.

Kyritsis, E., J. Andersson, and A. Serletis (2017). “Electricity prices, large-scale renewable integration, and policy implications” *Energy Policy* 101: 550–560 ISSN 0301-4215 doi: https://doi.org/10.1016/j.enpol.2016.11.014.

Larson, A. (2017). “The future looks bright for natural gas-fired power generation but price volatility is a wild card” *Power* July: 1–4.

Linn, J. and L. Muehlenbachs (2018). “The heterogeneous impacts of low natural gas prices on consumers and the environment” *Journal of Environmental Economics and Management* 89: 1 – 28 ISSN 0095-0696 doi:https://doi.org/10.1016/j.jeem.2018.02.002.

Liu, H. and J. Shi (2013). “Applying ARMA-GARCH approaches to forecasting short-term electricity prices” *Energy Economics* 37: 152–166 ISSN 0140-9883.

Longstaff, F. A. and A. W. Wang (2004). “Electricity forward prices: a high-frequency empirical analysis” *The Journal of Finance* 59(4): 1877–1900.

Martinez-Anido, C. B., G. Brinkman, and B.-M. Hodge (2016). “The impact of wind power on electricity prices” *Renewable Energy* 94: 474–487 ISSN 0960-1481 doi: https://doi.org/10.1016/j.renene.2016.03.053.

Masoumzadeh, A., E. Nekouei, T. Alpcan, and D. Chattopadhyay (2018). “Impact of optimal storage allocation on price volatility in energy-only electricity markets” *IEEE Transactions on Power Systems* 33(2): 1903–1914 ISSN 0885-8950 doi:10.1109/TPWRS.2017.2727075.

Molnar, P. (2012). “Properties of Range-Based Volatility Estimators” *International Review of Financial Analysis* 23: 20–29.

Neuhoff, K. and L. De Vries (2004). “Insufficient incentives for investment in electricity generations” *Utilities Policy* 12(4): 253 – 267 ISSN 0957-1787 doi: https://doi.org/10.1016/j.jup.2004.06.002 infrastructure Regulation and Investment for the Long-Term.

Parkinson, M. (1980)). “The Extreme Value Method for Estimating the Variance of the Rate of Return” *The Journal of Business* 53: 61–65.

Pirrong, C. and M. Jermakyan (2008). “The price of power: The valuation of power and weather derivatives” *Journal of Banking and Finance* 32(12): 2520–2529.

Qu, H., Q. Duan, and M. Niu (2018). “Modeling the volatility of realized volatility to improve volatility forecasts in electricity markets” *Energy Economics* 74: 767–776 ISSN 0140-9883 doi:https://doi.org/10.1016/j.eneco.2018.07.033.

Rampini, A., A. Sufi, and S. Viswanathan (2014). “Dynamic risk management” *Journal of Financial Economics* 111: 271–296.

Rampini, A. and S. Viswanathan (2010). “Collateral, risk management, and the distribution of debt capacity” *Journal of Finance* 65: 2293–2322.

Reguant, M. (2014). "Complementary bidding mechanisms and startup costs in electricity markets" *The Review of Economic Studies* 81(4): 1708–1742.

Rintamaki, T., A. S. Siddiqui, and A. Salo (2017). "Does Renewable Energy Generation Decrease the Volatility of Electricity Prices? An Analysis of Denmark and Germany" *Energy Economics* 62: 270–82.

Rogers, L. and S. Satchell (1991). "Estimating Variance From High, Low and Closing Prices" *Annals of Applied Probability* 1: 504–512.

Rogers, L. C., E. Stephen, and E. Satchell (1991). "Estimating variance from high, low, and closing prices" *Annals of Applied probability* 1: 504–512.

Rubin, D. R. (1978). "Bayesian inference for casual effects: the role of randomization" *Annals of Statistics* 6(1): 305–327.

Ullrich, C. J. (2012). "Realized volatility and price spikes in electricity markets: The importance of observation frequency" *Energy Economics* 34(6): 1809–1818.

U.S. Energy Information Administration (2018). "Carbon dioxide emissions from the U.S. power sector have declined 28% since 2005" Today in Energy, October 29, 2018.

Wolak, F. A. and C. D. Kolstad (1991). "A Model of Homogeneous Input Demand Under Price Uncertainty" *The American Economic Review* 81(3): 514 ISSN 00028282.