NBER WORKING PAPER SERIES

INFERRING CARBON ABATEMENT COSTS IN ELECTRICITY MARKETS: A REVEALED PREFERENCE APPROACH USING THE SHALE REVOLUTION

Joseph A. Cullen
Erin T. Mansur

Working Paper 20795
http://www.nber.org/papers/w20795

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2014

We thank Luis Gautier, Matthew Kotchen, and seminar participants at Washington University, University of Chicago, Portland Energy Economics Conference, Dartmouth College, University of Colorado, International Industrial Organization Conference, IFN Electricity Markets Conference, the Energy Institute at Haas Conference, Middlebury College, Yale University, and the World Bank. We thank Wolfram Schlenker for providing weather data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach using the Shale Revolution
Joseph A. Cullen and Erin T. Mansur
NBER Working Paper No. 20795
December 2014, Revised May 2015
JEL No. Q4,Q5

ABSTRACT

This paper examines how much carbon emissions from the electricity industry would decrease in response to a carbon price. We show how both carbon prices and cheap natural gas reduce, in a nearly identical manner, the historic cost advantage of coal-fired power plants. The shale revolution has resulted in unprecedented variation in natural gas prices that we use to estimate the short-run price elasticity of abatement. Our estimates imply that a price of \$10 (\$60) per ton of carbon dioxide would reduce emissions by 4% (10%). Furthermore, carbon prices are much more effective at reducing emissions when natural gas prices are low. In contrast, modest carbon prices have negligible effects when gas prices are at levels seen prior to the shale revolution.

Joseph A. Cullen
Washington University in St. Louis
jacullen@gmail.com

Erin T. Mansur
Dartmouth College
100 Tuck Hall
Hanover, NH 03755
and NBER
erin.mansur@dartmouth.edu

1 Introduction

Over the past decade, regulators have implemented a myriad of policies to mitigate climate change. The US federal government's recent policies either indirectly address climate change (*e.g.*, weatherization and renewables programs) or mandate standards, like tightening the Corporate Average Fuel Economy standards or capping carbon dioxide (CO_2) emissions rates for new power plants. Meanwhile other regulators have set a carbon price, like British Columbia's carbon tax, EU's Emissions Trading System, northeastern states' Regional Greenhouse Gas Initiative (RGGI), and California's Cap-and-Trade Program.

This paper examines how a carbon price is likely to affect emissions from the US electricity sector, which accounts for about one third of US greenhouse gases (EPA 2013), and has been the focus of many recently proposed regulations.[1] Firms can respond to carbon prices immediately by altering the mix of power plants used to meet demand: this is known as fuel switching. Just by switching generation from operating coal plants to the available, underutilized capacity at natural gas plants could reduce CO_2 emissions from the electricity industry by 23 to 42 percent (Lafrancois 2012). Whether or not this is feasible is an empirical question. This paper measures the expected environmental benefits from fuel switching in response to a range of carbon prices.

In order to do this, we examine how a carbon price would affect the marginal cost of producing electricity. This paper shows how a carbon price provides similar incentives for fuel switching as does a change in the cost ratio: namely the price of coal (per unit of heat content) over the price of natural gas. Briefly, higher carbon prices make coal-fired power plants less competitive than natural gas-fired power plants. Other power plants (nuclear, hydroelectric, and other renewables) have low marginal costs and remain inframarginal.[2] Similarly, when the cost ratio rises, natural gas plants gain an advantage: some cheap baseload coal plants may be displaced by even cheaper combined-cycle natural gas plants. While broadly this is true of other pollutants, we discuss why the mapping from cost ratios

[1] For example, the proposed Clean Power Plan focuses exclusively on the power sector.

[2] Oil-fired power plants produced less than 1% of the electricity during our sample (EIA 2014).

Figure 1: US and European Natural Gas Prices

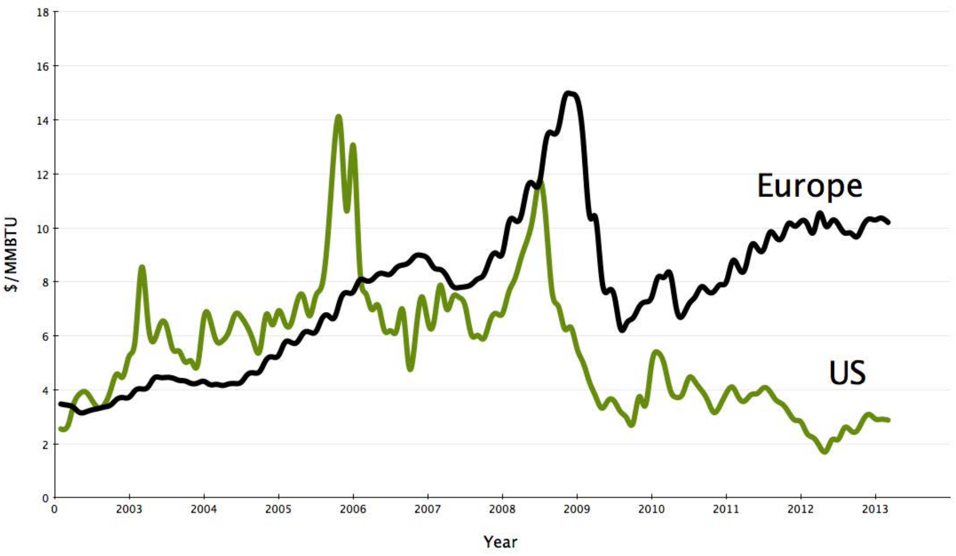

to carbon prices is substantially more precise. This mapping is important since we have no national carbon price that we could use to identify the short run marginal cost of abating carbon. Even where there are regional policies, there is limited variation in carbon prices.

On the other hand, we have recently observed abundant variation in natural gas prices. Technological advances in drilling (*i.e.*, hydrofracturing) have allowed firms to extract natural gas from shale formations. This "shale revolution" has resulted in a short run glut of gas: natural gas production has increased 26 percent from 2005 to 2012.[3] Furthermore, there are limited options to export substantial quantities of natural gas outside of North America. As a result, gas prices have dropped from over $12 per million British thermal units (mmBTU) to less than $2.[4] In 2012, gas in the US was less than a third of the cost of gas in Europe (see Figure 1).[5]

Using recent variation in fuel prices, we estimate the relationship between CO_2 emissions

[3]The Energy Information Administration (EIA) provides data on monthly natural gas production at www.eia.gov/dnav/ng/hist/n9010us2m.htm (accessed August 7, 2014).

[4]Henry Hub prices were $12.69 per mmBTU in June 2008 and $1.95 in April 2012. The EIA provides these monthly prices at www.eia.gov/dnav/ng/hist/rngwhhdm.htm (accessed August 7, 2014).

[5]The sharp drop in prices in 2008 reflects the recession. Since then, European prices have returned to levels seen before the recession while US prices remain low due to shale gas production (EIA 2012). These data are nominal prices from the World Bank Commodity Price Data (Pink Sheet).

and the coal-to-gas cost ratio using a flexible functional form. We control for several factors, including electricity load (the quantity consumed), temperature, generation from non-fossil sources, and net imports from Canada. In addition, we use time period fixed effects to proxy for macroeconomic shocks, other policies affecting the electricity sector, and power plant entry and exits. We find that when gas prices fall from \$6 to \$2, holding coal prices fixed, we predict a ten percent drop in aggregate CO_2 emissions.

We map this response curve into carbon prices. When baseline prices of natural gas are low, carbon prices are effective at reducing emissions. In particular, at the Energy Information Administration's expected fuel prices over the next decade (EIA 2012), we find that even a carbon price of \$10 (\$20) would reduce emissions about four (six) percent. A mandate of a ten percent reduction would be costly: the carbon price would need to be approximately \$60 per ton of CO_2.

In contrast, when coal holds a sizable cost advantage over natural gas, a marginal change in the cost ratio has no notable effect on emissions. Moving from a natural gas price of \$12 to \$6 has no effect on carbon emissions, holding coal prices constant. Thus, for high natural gas prices, even a moderate carbon price would have a limited impact on emissions. For example, if gas prices return to historic levels (like under a ban on hydrofracking or a significant build out of liquefied natural gas (LNG) exporting terminals), then even a price of \$20 per ton of CO_2 would reduce emissions by less than one percent. Even a \$60 per ton price would reduce emissions by only six percent. Finally, we show how a carbon price can result in co-benefits by reducing local emissions, in aggregate, in an approximately proportional manner.

Carbon prices provide incentives for a multitude of responses beyond fuel switching. Consumers facing higher electricity prices will conserve energy, for example by using energy-efficient technologies. Firms will build power plants that pollute few, if any, carbon emissions. In addition, companies may invest in order to make existing power plants operate more efficiently. These options are important to model in considering the overall effect of a carbon policy in the long run. However some carbon policies, like California's regulation and RGGI, are designed to protect consumers from rate increases. Furthermore, these other options take time (power plants are long lived) while policy tends to seek short-term performance.[6] Carbon pricing will lead to new investments and demand response in time, however the short-run response from fuel switching can be an important component from a political perspective.

Several recent papers directly examine the short-run effects of a carbon tax on emissions.[7] Newcomer, Blumsack, Apt, Lave & Morgan (2008) construct supply functions based on static, least-cost optimization: in other words, they must assume price-taking behavior and ignore technological constraints on operating power plants in order to model a static supply curve. For electricity markets in the mid-Atlantic (PJM), the upper-Midwest (MISO) and Texas (ERCOT), they find that a \$35/ton tax would result in a 2-2.5% carbon reduction due to fuel-switching.[8] Cullen (2013*a*) estimates a dynamic model of power plant production decisions and finds that a \$20/ton tax would have only a negligible effect on emissions in the Texas electricity market. Our study complements these papers by using a third method based on reduced-form estimates.

Another related literature examines how the low natural gas prices reduced emissions from the power sector.[9] Holladay & LaRiviere (2014) use hourly data for each of the eight North American Electric Reliability Corporation (NERC) regions to estimate the marginal emissions from regional fossil-fired gross generation. They show how their estimates have changed from a high gas price regime (2005-2008) to a low price regime (2009-2011). Linn, Muehlenbachs & Wang (2013) and Knittel, Metaxoglou & Trindade (2014) show how plant-level monthly production decisions change with cheap gas. Pratson, Haerer & Patio-Echeverri

[6]For example, the EU ETS was criticized for over-allocating permits making the policy ineffective in its first few years (Ellerman & Buchner 2007). Similarly, RGGI has had extremely low prices which has led regulators to tighten the cap in 2013.

[7]For example, Metcalf (2009) uses MIT's Emissions Prediction and Policy Analysis model and finds that a \$15 carbon price (in 2005 dollars per metric ton) would reduce US CO_2 emissions for all sectors by 8.4% in 2015.

[8]Overall they find reductions of about ten percent in PJM and MISO and about a third as much in ERCOT, but most of this is due to an assumed price response from consumers (with an assumed elasticity of -0.1). In practice, it remains unclear how much of the cost increases from carbon policies will be passed on to end users. For example, California grandfathers permits to utilities explicitly to protect customers from cost increases.

[9]Related government studies include Logan, Heath & Macknick (2012) and EPA (2013).

(2013) calculate the average cost of electricity generation for individual fossil plants and find that, under current regulations, coal plants maintain a cost advantage over gas plants if the coal-to-gas cost ratio lies below 0.56. Lu, Salovaara & McElroy (2012) regress coal production (as a share of monthly generation in a given census region) on the cost difference between natural gas and coal. They find that coal shares are responsive only to coal-to-gas cost ratios above 0.33 in most regions, and conclude that the drop in natural gas prices from 2008 to 2009 reduced CO_2 emissions from the US power sector by 4.3 percent, or half of the observed 8.8 percent reduction.[10] Finally, they use their analysis to analyze carbon taxes and find that a \$20/ton of CO_2 tax reduces annual electricity-sector emissions by seven percent. In contrast to this research, our paper accounts for the integrated grid across regions, estimates a flexible functional form of the daily cost ratio, and examines a longer time horizon with greater heterogeneity in natural gas prices.

We proceed with a brief discussion of the electricity industry in Section 2. Section 3 shows how cost ratios map into carbon prices. Sections 4, 5, and 6 describe the data, empirical model, and results, respectively, on the link between fuel costs and carbon emissions. Section 7 uses these estimates to examine the implications for carbon pricing and discusses caveats. We also examine the examine the co-benefits of carbon prices, namely reductions in local pollutants, in this section. Finally, we offer our conclusion in Section 8.

2 Background

Coal-fired power plants produce most of the electricity in the US (EIA 2014). On average, these baseload plants have low operating costs, are slow to adjust, and are costly to start up. However, there is substantial heterogeneity in the marginal cost of operating these plants. Some older, less efficient plants operate only during relatively high demand months. Most gas-fired generators fall into two categories: gas turbine peaker plants and combined cycle gas turbines (CCGT). Peaker plants have relatively low capital costs and high marginal costs.

[10]Linn, Mastrangelo & Burtraw (2014) find a similar result. When natural gas prices are high relative to coal prices, the effect of coal prices on electricity production from coal-fired power plants is smaller than when the prices are close together.

They operate during high demand hours, as power is prohibitively expensive to store and demand varies substantially over hours of the day and across seasons. In contrast, baseload CCGT plants are the most efficient fossil plants at turning the fuel's energy into power: *i.e.*, they have low heat rates (mmBTU/kWh). As such, some gas-fired power plants may have lower marginal costs than the most efficient coal plants *even* if coal costs less, per BTU, than natural gas.

Lower gas prices have been a boon for gas-fired generators in the US. Efficient gas power plants found themselves in the position to undercut coal-fired power plants. Figure 2 shows the monthly average electricity generation for power plants burning coal or natural gas from 2001 to the present.[11] While coal-based generation have generally been declining since the start of this century, a notable drop occurred in 2012 when natural gas briefly overtook it as the dominant fuel source. Note that this fuel switching primarily occurs across plants, not within a given plant.[12]

The degree to which production switches from coal to gas generation will depend on several factors. From a static dispatch framework, fuel switching depends on the relative fuel prices, the relative heat rates, the available capacity of gas plants, and the demand for electricity. In addition, intra-day fluctuations in electricity demand may be important as some generators are not well suited for starting and stopping production frequently. Start-up costs, ramping rates, minimum down times, and other intertemporal constraints limit firms' operation decisions (Mansur 2008, Cullen 2013*a*).

Furthermore, the transmission grid limits how much power plants can produce. As electricity is not stored, power supply and demand must equate at all times. This is subject to the network of transmission lines' capacity constraints, as well as the plants' intertemporal constraints (Mansur & White 2012). Therefore, optimal dispatch from a least-cost static model differs from the dynamic optimization.

Finally, observed production may differ still because power plants face forced outages

[11]This figure is based on EIA form 923 data that we describe in Section 4. See the web appendix for a figure on the generation shares by region, fuel type, and month.

[12]See Knittel et al. (2014) for a discussion of fuel switching within these dual-fuel plants.

Figure 2: Monthly Generation by Fuel Type

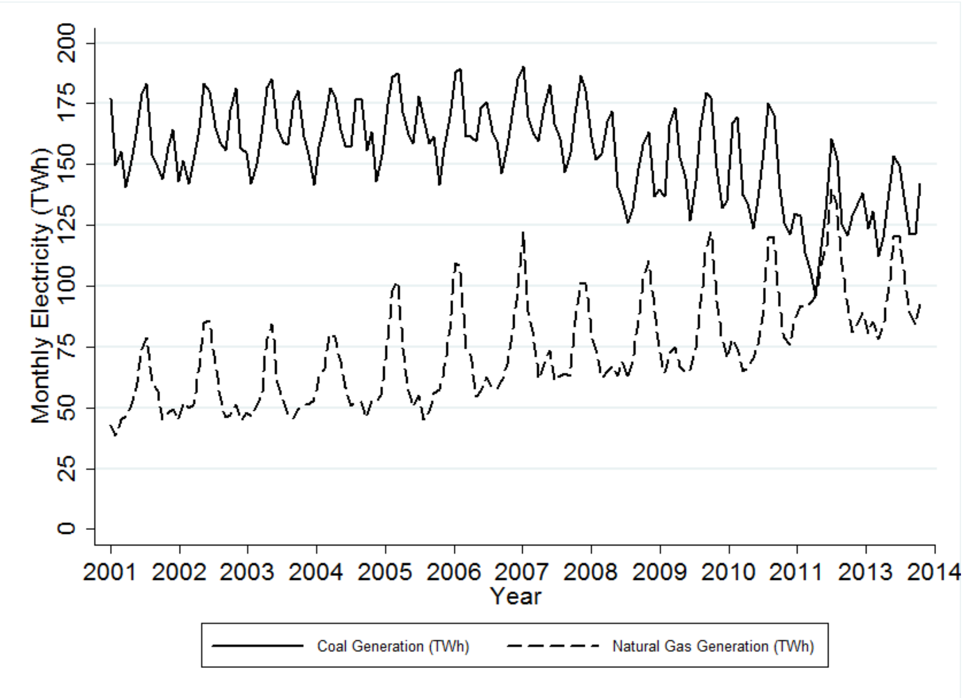

whereby they cannot operate when planned, firms may have imperfect information about trading opportunities (Mansur & White 2012), and firms may exercise market power (Borenstein, Bushnell & Wolak 2002, Mansur 2007, Puller 2007, Bushnell, Mansur & Saravia 2008). For these reasons, our analysis will use regressions to identify how firms actually respond to relative fuel prices.

The US electricity grid consists of three interconnections: East, West, and ERCOT (see Figure 3). Electricity produced in each interconnection is synchronized, allowing electricity to flow freely throughout the interconnection. Very little energy is transferred between interconnections due to the costs involved in transferring power between asynchronous grids. Analysis on a finer geographic scale is possible, but presents problems for measuring net emissions reductions in each area due to energy transfers between sub-regions within an interconnection. We revisit this issue when discussing co-benefits of carbon pricing from reduced local pollutants in Section 7.4.

Figure 3: NERC Interconnections

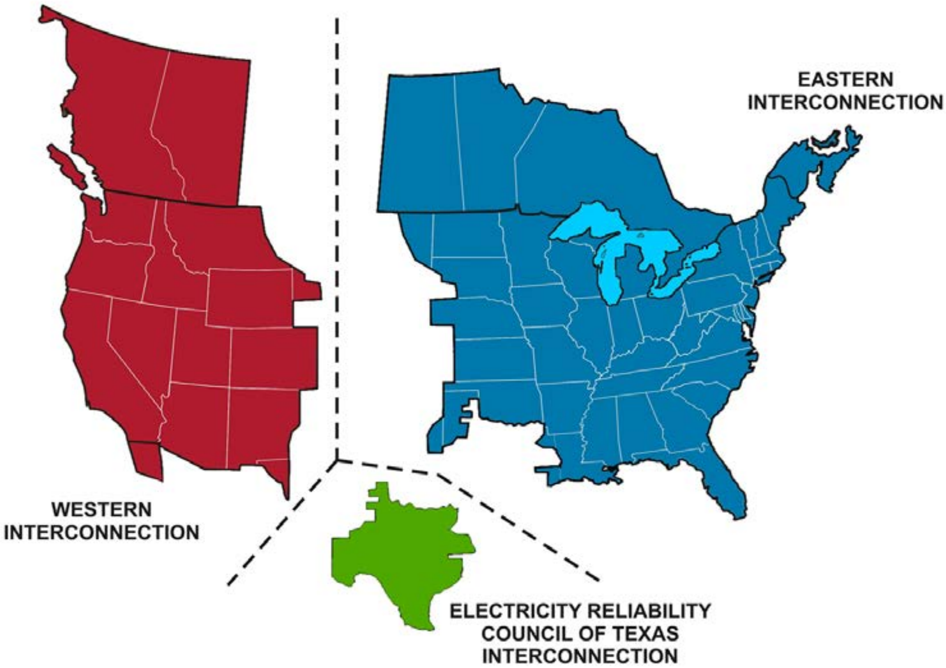

3 Mapping Carbon Pricing

Pricing carbon makes natural gas-fired generators more competitive with those burning coal. For a fossil-fired power plant, the marginal cost of producing electricity (MC) is a function of its variable operating and maintenance costs ($VO\&M$), heat rate (HR), price of fuel (P_{fuel}), carbon content($\frac{CO_2}{btu}$), and carbon price (P_{co2}):[13]

$$MC = VO\&M + HR \cdot P_{fuel} + HR \cdot \frac{CO_2}{btu} P_{co2}. \tag{1}$$

Although pricing carbon dioxide emissions increases marginal costs for both gas and coal plants, coal contains approximately twice as much CO_2 per unit of energy as natural gas. Thus, pricing carbon will affect the marginal costs of coal plants more than those of an equivalent gas plant. As previously mentioned, CCGT plants are more efficient than coal

[13]The $VO\&M$ costs include expenses for major overhauls, treating water, pumping water for cooling towers, replacing filters, *etc.* In addition to costs shown in equation (1), firms may have had to purchase, or forgo selling, pollution permits (*e.g.*, sulfur dioxide and nitrogen oxides permit prices). While these costs were extremely small relative to fuel costs during most of our sample, our analysis controls for them.

Figure 4: Carbon Prices and Generator Marginal Costs

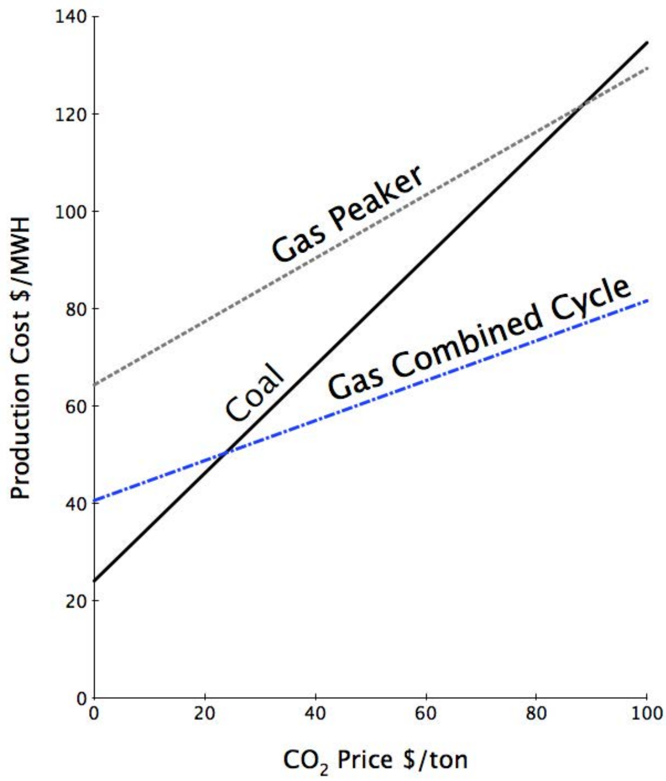

plants. These both lead to marginal costs rising more steeply with carbon prices for coal plants than gas plants. Figure 4 illustrates the change in marginal costs for an average coal plant relative to gas-fired technologies (CCGT and peaker) as the price of carbon increases.[14]

We examine the ratio of fuel costs, rather than the cost difference, for several reasons. First, the ratio captures how fuel prices translate into marginal costs. Suppose a gas plant is 25% more efficient than a coal plant. Then, assuming similar VO&M costs, the gas plant will have the same marginal cost when the cost ratio, P_C/P_G, equals 0.8. Production would switch between these generators as prices crossed this point. In fact, for a given cost ratio, the ordering of generators by marginal costs will be identical regardless of the level of fuel

[14]Typical heat rates for each technology were used to create this illustration. The heat rates used for technology are as follows: coal (10.73 mmBTU/MWH), gas combined cycle (7.07 mmBTU/MWH), gas peaker (11.2 mmBTU/MW). To calculate emissions costs, emissions factors for the carbon content of gas (117.0 lbs/mmBTU) and coal (210.86 lbs/mmBTU) were based on the rates reported by the Energy Information administration.

costs. To illustrate, consider a high cost and low cost scenario. In the low cost scenario, let $P_C = \$2$ and $P_G = \$4$. For the high cost scenario, let $P_C = \$3$ and $P_G = \$6$. In both cases, the cost ratio is the same. Now order the all the generators on the grid from lowest marginal cost to highest marginal cost to create an industry cost curve. From Equation (1), we see that the marginal cost of a generator depends on the product of its fuel cost and heat rate, the ordering of generators will be identical in each scenario. If generator A has 10% lower marginal costs than generator B in the low cost scenario, then it will also have 10% lower marginal costs in the high cost scenario. A second motivation for using the cost ratio is that it serves as a parsimonious function that translates the two dimensions of fuel costs (*i.e.*, coal and natural gas) into a single dimensional object that is simple to interpret.

Since coal costs are relatively constant over our time period, estimation by cost ratio is very similar to using other functional forms, such as differences and interactions of fuel costs. However, the cost ratio will be useful for connecting the estimated emission reductions to a counterfactual carbon price. We examine other functional forms in the appendix.

As mentioned previously, charging for carbon dioxide emissions increases the cost of burning coal more than burning gas. That is, it will increase the coal-to-gas cost ratio. For example, if coal were priced at \$2.25/mmBTU and gas were priced at \$5.75/mmBTU, this would imply a cost ratio of 0.39. Using this as a baseline, we can examine how a price on carbon would change the cost ratio. For instance, putting a \$20/ton price on CO_2 would change the cost ratio from 0.39 to 0.63.[15] Table 1 shows the mapping between carbon prices and cost ratios under the baseline for carbon prices up to \$100/ton CO_2 . High carbon prices can push the cost ratio above one; this means that a unit of energy form coal is now more expensive than a unit of energy from gas. Carbon prices are only one reason the cost ratio might change. Table 2 takes an alternative perspective. Rather than using a carbon price to change the baseline cost ratio, it looks at how lowering the price of gas could achieve a similar change in the cost ratio. We can replicate the cost ratios under carbon pricing by

[15] We use the following emissions factors for the carbon content of gas (117.0 lbs/mmBTU) and coal (210.86 lbs/mmBTU) based on the rates reported by the EIA (http://www.eia.gov/tools/faqs/faq.cfm?id=73&t=11). We weight bituminous, lignite, and sub-bituminous rates based on the aggregate annual fuel consumption of coal by power plants in 2011 (EIA form-923).

Table 1: Cost Ratios with Carbon Price

Carbon Price	Gas Cost					Coal Cost					Coal/Gas Cost Ratio
	Fuel	+	Carbon	=	**Total**	Fuel	+	Carbon	=	**Total**	
$0	5.75	+	0.00	=	**$5.75**	2.25	+	0.00	=	**$2.25**	0.39
$10	.	+	0.59	=	**$6.34**	.	+	1.05	=	**$3.30**	0.52
$20	.	+	1.17	=	**$6.92**	.	+	2.11	=	**$4.36**	0.63
$30	.	+	1.76	=	**$7.51**	.	+	3.16	=	**$5.41**	0.72
$40	.	+	2.34	=	**$8.09**	.	+	4.22	=	**$6.47**	0.80
$50	.	+	2.93	=	**$8.68**	.	+	5.27	=	**$7.52**	0.87
$60	.	+	3.51	=	**$9.26**	.	+	6.32	=	**$8.57**	0.93
$70	.	+	4.10	=	**$9.85**	.	+	7.38	=	**$9.63**	0.98
$80	.	+	4.68	=	**$10.43**	.	+	8.43	=	**$10.68**	1.02
$90	.	+	5.27	=	**$11.02**	.	+	9.49	=	**$11.74**	1.07
$100	5.75	+	5.85	=	**$11.60**	2.25	+	10.54	=	**$12.79**	1.10

Notes: Fuel costs are in $/mmBTU and carbon price is in $/ton of CO_2.

varying the price of gas between $2-$5/mmBTU.

Unlike other pollutants, the carbon content of a given fuel type is relatively homogeneous. Furthermore, there are no economically feasible end-of-pipe abatement technologies. In contrast, the ratio of nitrogen oxides emissions across plants, for example, varies widely because of differences in technologies and operational decisions. Thus while mapping cost ratios to carbon emissions is reasonable, we do not recommend using this approach for the pricing of other pollutants. We revisit this mapping in Section 7.

Table 2: Cost Ratios with Low Gas Price

Carbon Price	Gas Cost Fuel	Coal Cost Fuel	Coal/Gas Cost Ratio
$0	$5.75	$2.25	0.39
$0	$4.33	.	0.52
$0	$3.57	.	0.63
$0	$3.13	.	0.72
$0	$2.81	.	0.80
$0	$2.59	.	0.87
$0	$2.42	.	0.93
$0	$2.30	.	0.98
$0	$2.21	.	1.02
$0	$2.10	.	1.07
$0	$2.05	$2.25	1.10

4 Data

4.1 Data Sources

Our data are compiled from several public sources and cover January 2006 to December 2012. The Continuous Emissions Monitoring System (CEMS) of the Environmental Protection Agency (EPA) measures hourly output of CO_2, sulfur dioxide (SO_2), and nitrogen oxides (NO_x) from generators larger than 25 megawatts. We aggregate the hourly generator-level emissions information to construct daily CO_2 emissions. Generators are then aggregated by interconnection to create a measure of daily, regional CO_2 emissions. We aggregate the other pollutants by NERC region.

Second, we use data on electricity consumption (or load) provided by the Federal Energy Regulatory Commission (FERC). FERC Form 714 provides hourly information on electricity load by balancing area. We aggregate load to the daily level and sum across areas to arrive at daily electrical load by interconnection.

Third, we use EIA Form 923 data on production of electricity from non-fossil sources and prices paid for coal deliveries by power plants. EIA provides monthly electricity production by NERC region for nuclear and hydro power plants as well as for renewable sources (such as wind, solar and geothermal). We aggregate these data to the interconnection level for each type of non-fossil monthly electricity production: nuclear, water, and renewables. We also collect data from the National Energy Board of Canada on monthly net imports of power into each interconnection in the US.[16] We use data on permit prices for SO_2 from CantorCO2e and the EPA Clean Air Markets progress reports.

Fuel prices are aggregated by interconnection. In practice, there is some spatial heterogeneity in coal prices and, to a lesser degree, in natural gas prices. How much a power plant generates will depend on its own marginal costs as well as that of other plants: all fuel prices affect the order of dispatch. We simplify the vector of all power plants' fuel costs by looking at the average price of each fossil fuel.[17]

The EIA reports coal prices by transaction (plant, month, contract type, coal type, coal source, *etc.*).[18] We use this information to create a weighted-average price for each month and interconnection. In particular, we use data from 2001 to 2012 for spot prices only (dropping long term contracts over 12 months). For each interconnection, we regress coal costs on sulfur, ash, and BTU content, an indicator of surface mining, plant fixed effects, and indicator variables for each month of the sample. We estimate the model using weighted least squares, where we weight using a transaction's volume (in tons). The appendix reports the estimates (see Table A.1) and how they are used to construct a monthly coal price index for each region, holding coal composition fixed (see Figure A.1).

Finally, we use data from the Intercontinental Exchange (ICE) on the spot prices for natural gas at trading hubs around the country. ICE is an independent open-access electronic exchange for trading wholesale energy and metals commodities. For each gas hub, they report the average trading price for transactions on that day. For each interconnection, we weight the hub prices by the nameplate capacity of surrounding gas generators to arrive at a daily average spot price of natural gas. Although gas generators may have long term financial

[16] See http://www.neb-one.gc.ca/CommodityStatistics/Statistics.aspx?language=english, accessed August 27, 2014.

[17] Note that if these spatial differences are a constant percentage of the average price, then this heterogeneity will be captured in our model, much like the heterogeneity in heat rates.

[18] In reverse chronological order, the data sources are EIA-923, EIA-906, EIA-920, FERC 423, and EIA-423.

contracts for gas, the spot price for natural gas represents the opportunity cost to generators for using the gas to generate electricity versus selling it on the spot market. The general trends in the data are illustrated in Figure 1 using monthly averages.

Table 3 reports the mean and standard deviation for each interconnection. The East is the largest market by far with over four times the load in the West, which in turn is more than double ERCOT. The East is also the most carbon intensive with emissions over six times that in other markets.[19] The table also reports the summary statistics on the fuel prices for each region. All markets show substantial temporal variation in the cost ratio. While some of the variation in fuel prices is across regions, most of it is over time. The coal-to-gas cost ratio is 0.43 on average in the East and slightly smaller in the other markets.

Table 3: Summary Statistics

Variable	Units	Eastern	ERCOT	Western
CO_2 Emissions	1000s tons/day	5,005	527	802
		(768)	(89)	(119)
Load	GWh/day	7,456	866	1,835
		(879)	(159)	(168)
Emissions Rate	Tons/MWh	0.67	0.61	0.44
		(0.04)	(0.05)	(0.05)
Coal Price	\$/mmBTU	2.50	2.20	1.84
		(0.42)	(0.34)	(0.24)
Gas Price	\$/mmBTU	5.49	5.10	5.04
		(2.28)	(2.13)	(1.95)
Cost Ratio		0.43	0.41	0.35
		(0.88)	(0.81)	(0.76)
Observations		2,557	2,557	2,557

In the next section, we use these data to trace out the emission response of the electricity system to changes in input costs while controlling for important features of the market. However, we first calculate how much of a reduction in carbon emissions is feasible given the current stock of power plants.

[19] We report load in gigawatt-hours (GWh) per day and emissions rates in tons of CO_2 per megawatt-hour (MWh), where a GWh is one thousand MWh and one million kWh.

Table 4: Potential Shares of Carbon Emissions Reduced from Fuel Switching

Eastern	ERCOT	Western
0.42	0.37	0.40

As a simple back-of-the-envelope calculation, we examine whether there is sufficient capacity at natural gas facilities to have a substantial effect on carbon emissions. See the appendix for details on the methodology. Table 4 reports the share of carbon emissions from that could be reduced in 2012, assuming a derating factor of 90%. Similar to Lafrancois (2012), we see that we have enough unused gas capacity to reduce emissions by about 40%. These results vary over time.[20] In 2001, unused gas capacity was the limiting factor so that only 25% of emissions could be reduced. Starting in 2003, investment in gas capacity had grown such that unused gas capacity exceeded coal production in nearly every hour in the West and in ERCOT, and about 12-29% of the time in the East.[21]

5 Empirical Model

We aim to create a simple, yet flexible model that can trace out the response of emissions to changes in relative fuel that can accommodate the varied technologies on the grid and their complex interactions in electricity markets. The method used is similar to the literature that econometrically estimates the relationship between emissions and either electricity consumption(Graff Zivin, Kotchen & Mansur 2014), electricity generation (Siler-Evans, Azevedo & Morgan 2012, Holladay & LaRiviere 2014), and wind production (Callaway & Fowlie 2009, Cullen 2013*b*, Kaffine, McBee & Lieskovsky 2013, Novan forthcoming).

The model is a reduced-form regression with daily carbon dioxide emissions (CO_{2t}) in an interconnection as the dependent variable as shown below. We are relatively confident that the emissions reduction estimated for an interconnection come from that interconnection as

[20]The web appendix shows shares for each year from 2001 to 2012 for derating factors of both 90% and 80%.

[21]We use a second measure of hourly generation ($\widetilde{gen}_{ift}$) based on heat input data from CEMS to allocate $eiagen_{ifm}$ across hours in a month: $\widetilde{gen}_{ift} = eiagen_{ifm} \cdot (cemsheat_{ift} / \sum_{t \in m} cemsheat_{ift})$, where $cemsheat_{ift}$ is the hourly heat input reported by CEMS. The results are quite similar.

little power flows between them.[22] For day t, the estimating equation is:

$$CO_{2t} = f(CR_t) + g(X_t) + Z_t\psi + D_t\gamma + \epsilon_t, \tag{2}$$

where CR_t is the ratio of the coal price over the gas price, both measured in dollars per millions of BTU.[23] The matrices X_t, Z_t, and D_t are sets of controls described below.

For identification, we rely on exogenous shocks to natural gas prices. When selecting controls included in X_t, we need to include variables that would directly affect the interconnection emissions that might also be correlated with the variation in input fuel costs. The quantity of electricity demanded, or load, obviously meets this criteria. The quantity demanded on a given day, although driven by weather and day-specific demand shocks, may be correlated with the spot price for gas. This may be because electricity generators demand more gas when electricity demand is high or simply a correlation in the demand for electricity and the demand of gas outside the electricity sector, such as home heating. For example, lower electricity demand and emissions due to a negative macroeconomic shock would be correlated with low prices for natural gas due to the same shock. Failing to account for electricity demand would tend to overestimate the response of emissions to the price gap. Thus we include daily electricity demand in the interconnection as a control variable.

In X_t, we also include data on production from non-fossil fuel electricity production. Non-fossils electricity production includes wind, solar, hydro, and nuclear power generators. While these generators are not likely to change their production in response to gas or coal prices, they may be correlated with them. For example, wind power installations have been growing at the same time as technological innovation has led to more shale gas extraction. Likewise, seasonal variation in the availability of hydroelectric generating capacity may influence the spot prices of natural gas.

When implementing the estimating equation, we use flexible functional forms for $f(.)$ and $g(.)$ to trace out the emissions response of the system. Specifically, we use a cubic spline

[22] Novan (forthcoming) provides an example of some power plants in SPP capable of selling into ERCOT. In addition, DC lines do connect interconnections like from Quebec to New England.

[23] Note that gas prices change on a daily basis, while the information on coal prices changes monthly.

with six knot point for cost ratio and load. We test the robustness of our results to different numbers of knots and found the results to be stable with four or more.

We also control for other factors, Z_t, in more traditional parametric ways. We capture the within-day distribution of hourly load using the minimum, maximum, and standard deviation. We control for monthly net imports of electricity from Canada, non-fossil electricity production, and the permit prices of local pollutants (SO_2 and NO_x). Finally, we include a dummy variable (D_t) for each quarter in the time series to control for trends in generating capacity, macroeconomic shocks, as well as seasonality in generator availability. The resulting regression is:

$$CO_{2t} = s(CR_t|\beta) + s(load_t|\theta) + s(unfossil_t|\alpha) + s(temp_t|\omega) + Z\psi + D\gamma + \epsilon_t, \qquad (3)$$

where the function $s(\cdot)$ is a cubic spline. In order to account for serial correlation and heteroscedasticity, we use Newey-West standard errors allowing for a seven-day lag structure. With the estimated coefficients, we can trace out the emissions response of the electricity generating system to changes in the relative costs of coal and gas.

6 Results

6.1 Main Results

The results from the estimation for each interconnection are shown in Figure 5. We plot the percent change in carbon dioxide emissions against the price of natural gas, using the EIA (2012) expected future coal prices to construct the cost ratio.[24] Dashed lines show the 95% confidence interval for the estimates using Newey-West corrected standard errors. Control variables, such as demand and non-fossil electricity production, are held at their average levels in the sample.

The results show statistically insignificant changes in emission for high gas prices. That is, when gas prices are above \$6 (per mmBTU), changes in gas prices don't result in switching

[24]Changes are relative to the emissions given the EIA expected future natural gas and coal prices.

Figure 5: Estimated CO_2 Response to Fuel Prices

(a) Eastern Interconnection

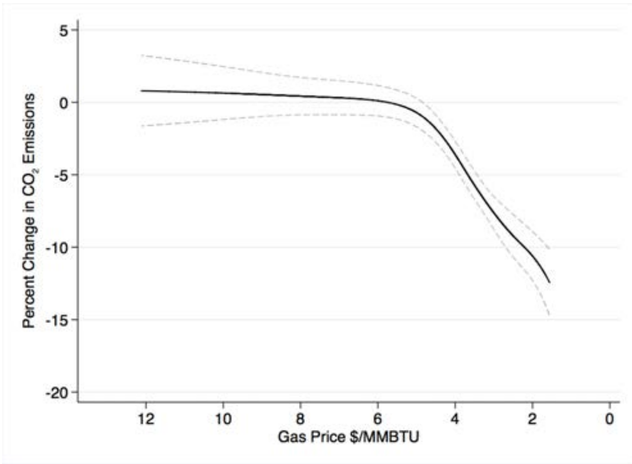

(b) ERCOT Interconnection

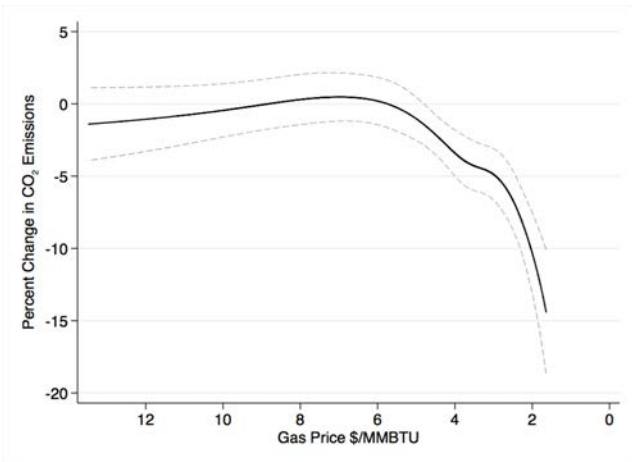

(c) Western Interconnection

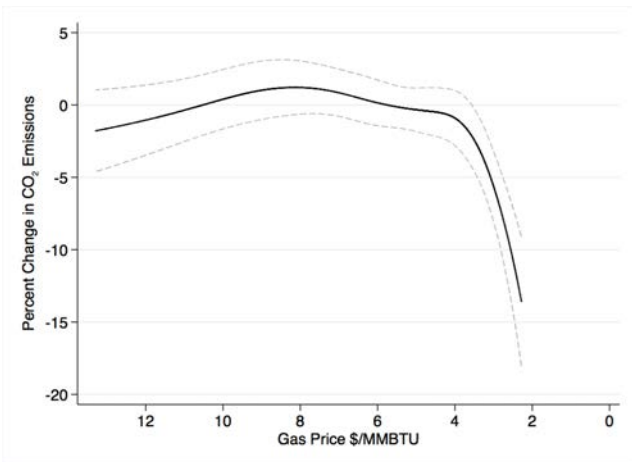

between high polluting plants and cleaner facilities. Not until the gas prices approach \$4-\$5, do emissions begin to fall. For the Eastern interconnection, emissions fall by about 10% when the gas price falls to \$2. For ERCOT and the Western interconnection, carbon emissions fall by about seven and twelve percent, respectively, at this gas price. Keep in mind that \$2 reflects *historically* low gas prices. This brings much of the gas-fired fleet on par with coal-fired generators. Though the reduction in emissions is significant, it does not begin to approach the 40% estimates in Section 4. This suggests that dynamics, transmission constraints, or other factors excluded from simpler models greatly reduce the emissions reductions possible from fuel switching.

6.2 Robustness

We examine the robustness of these results in the appendix. Figure A.7 examines the functional form that we use to compare coal and gas prices. Figure A.8 compares the main results using different number of knots for the cubic splines. In the web appendix, we examine additional robustness tests. We test how sensitive our results are to dropping some of our control variables. We also study how the response function changes with the length of time we use for fixed effects. Notably one specification includes month-of-sample fixed effects that control for all variation in our estimated coal prices. Even within a month, we observe modest fuel switching due to variation in daily natural gas prices.

We also consider concerns of potential endogeneity. We rely on exogenous variation in natural gas prices for identification. While we control for the effect that electricity demand and temperature might have on fuel prices, a random shock to daily emissions, *conditional* on electricity demand, could shift the market demand for fuels. For example, suppose a large coal fired-power plant is forced to shut down for a few days. All else equal, this increases demand for natural gas and while at the same time emissions fall. In theory, this could increase the price of natural gas and introduce bias into our the coefficient estimates. This would imply that our estimates are an underestimate of the true effect. However, we argue that these biases are likely to be small. First, we have included most of the factors that affect fuel choice like production from non-fossil power plants, net imports, *etc.* Second, the storage of electricity and fuels are dramatically different: While power is prohibitively expensive to store, fuels are storable commodities. Today's natural gas price reflects current both weather conditions and electricity load, but also expectations about future demand. Thus, one day's demand shock may have limited effect on prices.

7 Discussion of Carbon Price

7.1 Effect on Carbon Emissions

We use the electricity industry's experience with low gas prices to explore how the industry may respond to a carbon tax. As a first step, we need to choose a baseline level for fuel prices from which to compare the effect of various carbon prices. In this section, we assume that these prices are exogenously determined, with prices returning to the long run average costs of extracting and processing the fuels. The EIA (2012) forecasts that average delivered coal prices will be \$2.25/mmBTU and gas prices will be \$5.75/mmBTU in 2025. This implies a baseline cost ratio of 0.39, which will serve as our benchmark[25].

Next we map the emissions response curves that we estimated in Section 6 into carbon prices using the logic discussed in Section 3. As shown in Table 2, for any cost ratio observed in the data, there is a matching counterfactual carbon price with the same cost ratio given the baseline fuel prices. As previously discussed, the ordering of the generators in the industry marginal cost curve will be identical, whenever the fuel cost ratios are the same. The industry cost curve under a carbon tax will be proportional to the cost curve in the data with the same cost ratio.

With fixed baseline fuel costs, we can project our estimates of emissions reductions due to shocks to gas prices onto their equivalent carbon price. For each region, Figure 6 shows

[25] When calculating carbon tax from fuel prices, we don't account for the impact that an increasing carbon tax may have on the equilibrium price of fuels. That is, a higher carbon price will lead to increased demand for natural gas which could increase the price of gas. Rather, we assume that baseline fuel prices are fixed and exogenous. Given that incorporating any price response of fuels to a carbon tax would tend to decrease the emissions reductions for higher carbon taxes, our results represent generous estimates of emissions reductions from a carbon tax.

Figure 6: Imputed CO_2 Response to Carbon Prices

(a) Eastern Interconnection

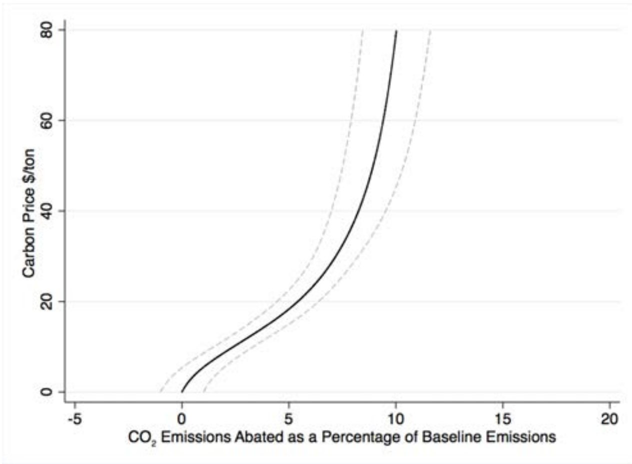

(b) ERCOT Interconnection

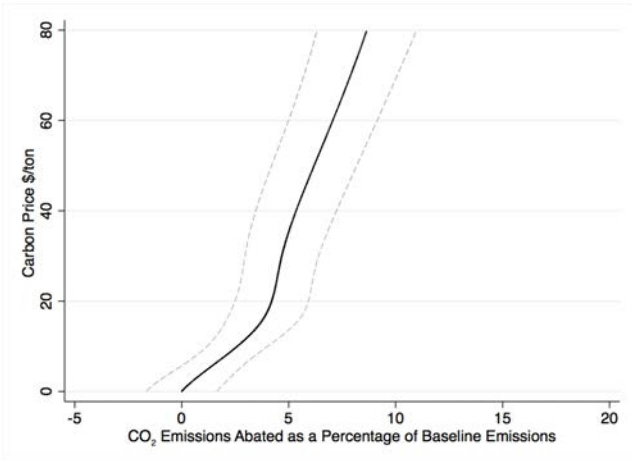

(c) Western Interconnection

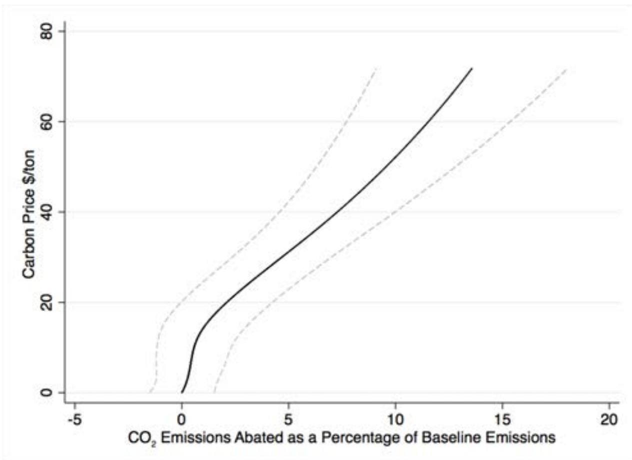

5

the estimated emissions reductions that would come from a carbon price (in \$/ton of CO_2) under these assumptions. The figures focus on the cost ratios that correspond to positive carbon prices under the baseline fuel costs. They show that emissions fall steeply at lower levels of carbon tax, but then the rate of change decreases for higher levels of carbon tax. These results indicate that much of the emissions reduction from technology switching can be captured with a relatively modest price on carbon. High price carbon on carbon do result in some further reduction in carbon dioxide emissions, but the large impact from high carbon price is likely to come from retooling the generating infrastructure. Specifically, we find that even a carbon price of \$10/ton would reduce emissions about four percent. However, to achieve a ten percent reduction, the carbon price would need to be closer to \$60/ton.

7.2 Effect on Carbon Emissions by Fuel Type

Table 5: CO_2 Emissions Abated at \$20/ton by Fuel Type*

	East	ERCOT	West
Coal	-5.97	-6.16	-3.29
	(0.45)	(1.09)	(1.02)
Gas	0.40	2.33	1.21
	(0.25)	(0.61)	(0.58)
Other	0.18	-0.37	–
	(0.06)	(0.08)	–
Total	-5.39	-4.21	-2.08

**Percentage of baseline emissions*

Table 6: CO_2 Emissions Abated at \$60/ton by Fuel Type*

	East	ERCOT	West
Coal	-11.10	-9.57	-13.31
	(0.72)	(1.38)	(1.94)
NG	1.64	2.74	1.84
	(0.33)	(0.70)	(1.13)
Other	0.06	-0.25	–
	(0.08)	(0.10)	–
Total	-9.40	-7.07	-11.47

**Percentage of baseline emissions*

7.3 Carbon Pricing Under High Gas Prices

We can also use our estimates to examine how effective carbon prices would be at reducing emissions under different initial prices. The results of Section 7.1 are based on EIA predicted fuel prices. Suppose that the price of natural gas ends up being much higher than expected. For example, new environmental regulations could increase the cost of shale gas production or even impose a ban on fracking. Alternatively, gas prices could rise if there is a significant increase in LNG exports.[26] Before shale gas became a major share of US gas production, the US and Europe had similar natural gas prices (see Figure 1). While the recession lowered European gas prices, they returned to pre-recession levels by 2012.

Figure 7 shows the effect of carbon pricing assuming that US fuel prices also returned to the levels seen in the spring of 2008.[27] The figure shows that, when natural gas prices are high, carbon prices are much less effective at reducing carbon emissions. For example, a $20/ton carbon price reduces carbon emissions in the East by about six percent in the base case (see Figure 6a), but by less than one percent when gas prices are high. In other words, in order to achieve a carbon cap-and-trade target, a much higher carbon permit price would be required if gas prices were high.

[26] Note that for prices to rise to the historic rates because of LNG exports would require much larger foreign prices as liquefying gas is costly.

[27] In April of 2008, fuel prices were $2.46/mmBtu for coal and $10.28/mmBTU for natural gas, implying a coal-to-gas cost ratio of 0.24.

Figure 7: Emission Response Limited in a High Gas Price Scenario

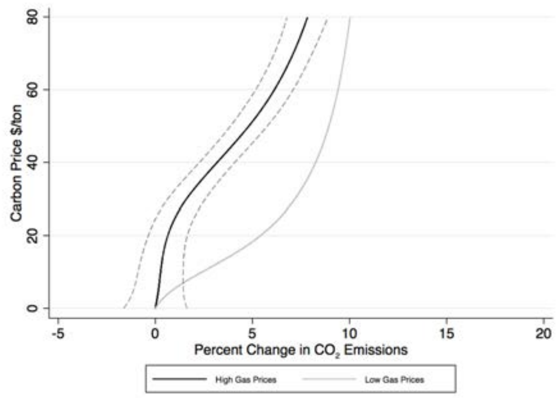

7.4 Co-Benefits of Carbon Price

By changing the dispatch of power plants, carbon prices are likely to reduce other pollutants like SO_2 and NO_x that have local and regional health effects.[28] For each interconnection, we replicate the estimation and simulation methods of Sections 5 and 7.1 where we replace the dependent variable in equation (3) with the daily emissions of either SO_2 or NO_x within a given interconnection. We estimate how these emissions depend on the coal-to-gas cost ratio, which we then convert into a carbon price as above.

Figure 8 shows the aggregate response of SO_2 emissions to carbon prices. We see that a $20/ton carbon price results in about a 6% drop in emissions in each region. However, the functions differ at other prices. For example, a 10% drop in SO_2 emissions would require a carbon price of $40 in the East or the West, but almost double that in ERCOT. Figure 9 shows the response curves for NO_x emissions by region. Here a large reduction in NO_x emissions would occur from a much larger carbon price in the East than in the other regions.

[28]Burtraw, Krupnick, Mansur, Austin & Farrell (1998) and Chestnut & Mills (2005) examine the environmental benefits of federal regulations of local pollutants. See Schmalensee & Stavins (2013) for a recent discussion of these policies.

Figure 8: SO_2 Response to Carbon Price by Interconnection

(a) Eastern

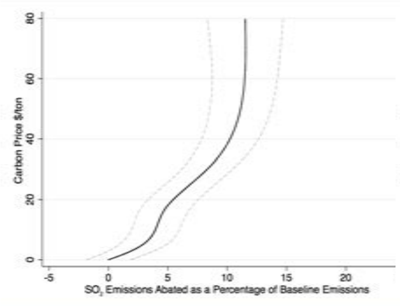

(b) ERCOT

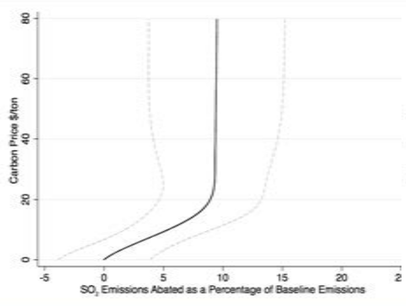

(c) Western

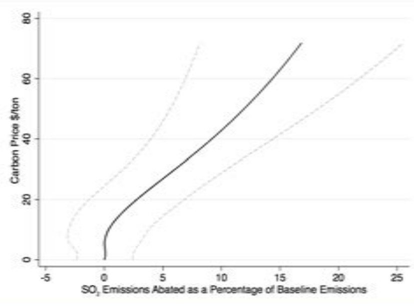

Figure 9: NO_x Response to Carbon Price by Interconnection

(a) Eastern

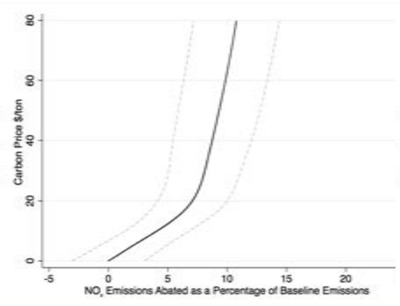

(b) ERCOT

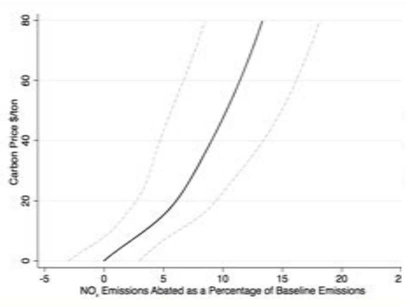

(c) Western

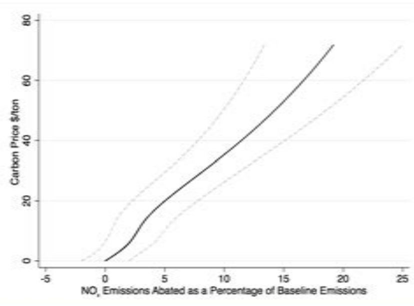

However, these figures mask important spatial variation. Unlike CO_2 emissions, the location of these local emissions matters for estimating the marginal damages. While a precise estimate of these marginal damages is beyond the scope of this paper, we do examine the spatial distribution of these emissions. In particular, we modify equation (3) by defining the dependent variable as the daily emissions (SO_2 or NO_x) within a given subregion of an interconnection.[29] Figure A.12 shows the regional variation in responses.

7.5 Caveats

There are several assumptions required in order for the behavior of generators and their associated emissions to be identical for observed and counterfactual scenarios with the same cost ratios. First, although the relative costs of coal and gas mayd be the same in the observed and counterfactual worlds, the level of costs are not. Coal and gas costs will be much higher under a carbon price than in the data where we observe average coal costs and very low gas costs. As such, equilibrium electricity prices will be higher under a carbon tax. The degree of demand response to these higher prices will also affect the change in carbon emissions. Only with a perfectly inelastic demand curve will a straight mapping based on cost ratios be accurate. However, the results of the estimation do highlight the degree to which technology switching can contribute to emissions reductions under a carbon tax.[30]

Second, even if the marginal cost curve under two scenarios imply the same dispatch order, the profits of an individual generator need not be identical. In particular, (inframarginal) generators reap higher markups due to the steeper slope of the marginal cost curve in the high-cost scenario. While this is not an issue when holding the generating infrastructure on the grid fixed, carbon prices do provide greater incentive for investment in clean generating technology.

Third, another issue arises from the dynamic considerations in operating a generator.

[29]The online appendix has a map of these NERC subregions as defined in the EPA's eGrid database. Note that the independent variables are still at the interconnection level.

[30]The analysis also assumes that the behavior of low marginal cost renewable and nuclear generators are unaffected by higher equilibrium electricity prices. Given that these generators already have incentives to operate at full capacity whenever possible, this assumptions seems reasonable.

Generators make operating decisions amid fluctuating intra-day demand. This necessitates that some generators shut down and restart. Generators will start production only if they expect to cover their start-up costs while operating. Since the profits of generators will be different under the two scenarios with the same cost ratios, firms may undertake different startup decisions. This is lessened if their start-up costs scale with the fuel prices. Although fuel costs are a central part of start-up costs, they are not the only component. Thus will not scale perfectly with fuel costs. The degree to which changes in dynamics affect the outcomes is difficult to judge given our approach. However, structural dynamic estimation of electricity markets indicates that start-up costs are not a driving factor for aggregate carbon emissions changes under a carbon tax (Cullen 2013*a*).

Finally, regulatory uncertainty is another way that these scenarios may differ. Firms facing natural gas prices can look at how the market is doing and predict where it likely to go given their model of the long run equilibrium. In contrast, predicting what a permit price will be in a cap-and-trade regulation, or how regulation may change for any carbon price is inherently challenging. There are no market forces pulling towards a clear regulatory outcome. As a result, firms making irreversible capital investments in generation technology, for example, are likely to behave differently given these differences in uncertainty.

A final comment concerns the interaction of multiple regulations. This paper predicts changes in emissions for CO_2 and co-benefits in response to carbon prices. If a cap-and-trade market already exists, for SO_2 for example, then an additional carbon tax cannot effect aggregate emissions (assuming the cap continues to bind). This does not mean that the market will be unaffected: a carbon tax reduces demand for SO_2 permits, the permit prices will fall, and the spatial distribution of emissions will change. Even for the direct effects on CO_2 emissions, the California and RGGI markets are now capping emissions in their respective areas. During our sample period, RGGI permits were quite low and California was just starting to trade, making it very unlikely that these policies affect our estimates. However, going forward, it is important to keep in mind how state, regional, and national carbon policies interact (Goulder & Stavins 2011).

8 Conclusion

This paper provides estimates, based on observed behavior rather than simulations, of the impact of carbon pricing on electricity-sector emissions. We show how lower gas prices and a carbon price can affect the relative costs of generators in similar ways. This paper exploits significant variation in natural gas prices that resulted from a rare combination of factors: a large recession, the start of the shale revolution, and limited capacity to export gas. In the near future, the federal government projects an end to these low prices as exports rise and the economy recovers (EIA 2012). In this paper, we use the recent price variation to estimate how the electricity sector's carbon emissions respond to fuel cost shocks, and examine conditions under which this response to relative fuel prices can inform us about how a price on carbon dioxide will change emissions in the short run. Understanding how a carbon price will affect polluting firms *in the short run* is an important step in demonstrating the effectiveness of such an instrument for use *in the long run*. On a longer time horizon, even greater emissions reductions could be expected as new generation could be built and consumers could adjust to new equilibrium electricity prices.

Our results indicate that carbon prices will result in a modest effect on emissions: even a price of $60 per ton of carbon dioxide will reduce emissions only 10%. However, much of the reduction in carbon dioxide emissions can be captured with a relatively modest carbon tax: a price of $20 reduces emissions by 6%. Furthermore, carbon prices are much more effective at reducing emissions when natural gas prices are low. In contrast, modest carbon prices have negligible effects when gas prices are at levels seen prior to the shale revolution. Finally, we show how a carbon price can result in co-benefits by reducing local emissions, in aggregate, in an approximately proportional manner.

Many emissions trading policies including the EU ETS, RGGI, and RECLAIM have been criticized for their effectiveness in reducing emissions in the short run.[31] While the overarching objective of climate policy is to reduce aggregate cumulative emissions of green-

[31] Tvinnereim (2014) discusses several reasons why many cap-and-trade policies have had lower permit prices than expected. Concerns over RECLAIM were due to high prices, non-compliance, and environmental justice (see Fowlie, Holland & Mansur (2012) for an analysis of these concerns).

house gases, much of the focus to date has been on the short-run impacts. We argue that understanding how firms fuel switch is important in knowing how effective a market-based instrument can be in the short run, where the policy debate over carbon pricing seems to focus.

References

Borenstein, S., Bushnell, J. B. & Wolak, F. A. (2002), 'Measuring market inefficiencies in California's restructured wholesale electricity market', *American Economic Review* **92**(5), 1376–1405.

Burtraw, D., Krupnick, A., Mansur, E. T., Austin, D. & Farrell, D. (1998), 'Cost and benefits of reducing air pollutants related to acid rain', *Contemporary Economic Policy* **16**(4), 379–400.

Bushnell, J. B., Mansur, E. T. & Saravia, C. (2008), 'Vertical arrangements, market structure, and competition: An analysis of restructured US electricity markets', *American Economic Review* **98**(1), 237–66.

Callaway, D. & Fowlie, M. (2009), Greenhouse gas emissions reductions from wind energy: Location, location, location?, Working paper.

Chestnut, L. G. & Mills, D. M. (2005), 'A fresh look at the benefits and costs of the U.S. acid rain program', *Journal of Environmental Management* **77**(3), 252–266.

Cullen, J. (2013*a*), Dynamic response to environmental regulation in the electricity industry, Working paper.

Cullen, J. (2013*b*), 'Measuring the environmental benefits of wind-generated electricity', *American Economic Journal: Economic Policy* **5**(4), 107–133.
URL: *http://www.aeaweb.org/articles.php?doi=10.1257/pol.5.4.107*

EIA (2012), Annual Energy Outlook 2012, Technical report, U.S. Energy Information Administration.

EIA (2014), Electric power monthly, Technical report.

Ellerman, D. & Buchner, B. (2007), 'The European Union Emissions Trading Scheme: Origins, allocation, and early results', , *Review of Environmental Economics and Policy* **1**(1), 66–87.

EPA (2013), Inventory of U.S. greenhouse gas emissions and sinks: 1990-2011, Epa 430-r-13-001, U.S. Environmental Protection Agency.

Fowlie, M., Holland, S. P. & Mansur, E. T. (2012), 'What do emissions markets deliver and to whom? evidence from Southern California's NOx Trading Program', *American Economic Review* **102**(2), 965–993.

Goulder, L. H. & Stavins, R. N. (2011), 'Challenges from state-federal interactions in US climate change policy', *American Economic Review: Papers & Proceedings* **101**(3), 253–257.

Graff Zivin, J. S., Kotchen, M. & Mansur, E. T. (2014), 'Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies', *Journal of Economic Behavior and Organization* **107**(Part A), 248268.

Holladay, J. S. & LaRiviere, J. (2014), The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy, Working paper.

Kaffine, D. T., McBee, B. J. & Lieskovsky, J. (2013), 'Emissions savings from wind power generation in Texas', *The Energy Journal* **34**(1), 155–175.

Knittel, C. R., Metaxoglou, K. & Trindade, A. (2014), Dash for gas: The sequel, Working paper.

Lafrancois, B. A. (2012), 'A lot left over: Reducing CO_2 emissions in the United States' electric power sector through the use of natural gas', *Energy Policy* **50**, 428–435.

Linn, J., Mastrangelo, E. & Burtraw, D. (2014), 'Regulating greenhouse gases from coal power plants under the clean air act', *Journal of the Association for Environmental and Resource Economics* **1**(1/2), 97–134.

Linn, J., Muehlenbachs, L. & Wang, Y. (2013), Short-run demand for natural gas in the U.S. electric power sector, Working paper.

Logan, J., Heath, G. & Macknick, J. (2012), Natural gas and the transformation of the U.S. energy sector: Electricity, Technical report nrel/tp-6a50-55538, Joint Institute for Strategic Energy Analysis.

Lu, X., Salovaara, J. & McElroy, M. (2012), 'Implications of the recent reductions in natural gas prices for emissions of CO2 from the US power sector', *Environmental Science & Technology* **46**(5), 3014–3021.

Mansur, E. T. (2007), 'Upstream competition and vertical integration in electricity markets', *Journal of Law and Economics* **50**, 125–781.

Mansur, E. T. (2008), 'Measuring welfare in restructured electricity markets', *The Review of Economics and Statistics* **90**(2), 369–386.

Mansur, E. T. & White, M. (2012), 'Market organization and efficiency in electricity markets', (Working Paper).

Metcalf, G. E. (2009), 'Designing a carbon tax to reduce U.S. greenhouse gas emissions', *Review of Environmental Economics and Policy* **3**(1), 63–83.

Newcomer, A., Blumsack, S., Apt, J., Lave, L. & Morgan, M. G. (2008), 'Short run effects of a price on carbon dioxide emissions from US electric generators', *Environmental Science & Technology* **22**(9), 3139–3144.

Novan, K. (forthcoming), 'Valuing the wind: Renewable energy policies and air pollution avoided', *American Economic Journal: Economic Policy* .

Pratson, L. F., Haerer, D. & Patio-Echeverri, D. (2013), 'Fuel prices, emission standards, and generation costs for coal vs natural gas power plants', *Environmental Science & Technology* **47**(9), 4926–4933.

Puller, S. L. (2007), 'Pricing and firm conduct in California's deregulated electricity market', *Review of Economics and Statistics* **89**(1), 75–87.

Schmalensee, R. & Stavins, R. N. (2013), 'The SO_2 allowance trading system: The ironic history of a grand policy experiment', *Journal of Economic Perspectives* **27**(1), 103–121.

Siler-Evans, K., Azevedo, I. L. & Morgan, M. G. (2012), 'Marginal emissions factors for the US electricity system', *Environmental Science & Technology* **46**(9), 4742–4748.

Tvinnereim, E. (2014), 'The bears are right: Why cap-and-trade yields greater emission reductions than expected, and what that means for climate policy', *Climatic Change* **127**(3-4), 447–461.

Appendix

Coal Price Regressions

Section 4 discusses the method we use to calculate a coal price index. In this appendix, we report the regression results for each region (see Table A.1). We also show the national average for comparison. While the type of coal a power plant purchases may change in response to natural gas prices, we are concerned that the heterogeneity in coal transactions within a plant may also reflect noise. So, we construct a coal price index based on the average coal characteristics in a region as of January 2001. We then add the monthly fixed effects to the base price, thereby keeping the coal composition constant for a region. This coal price index in shown in Appendix Figure A.1.

Table A.1: Coal Price Index Regression Results

Variable	National	East	West	ERCOT
Sulfur	-13.59***	-13.59***	-10.51*	1.40
	(1.38)	(1.36)	(5.77)	(12.86)
Ash	2.09***	2.69***	-1.53***	0.40
	(0.26)	(0.33)	(0.35)	(2.00)
Mine	-7.99***	-10.88***	19.31***	32.33***
	(1.81)	(1.88)	(5.99)	(7.13)
Btu	4.40***	4.37***	5.33***	5.24*
	(1.40)	(1.42)	(0.95)	(3.09)
Month-Year F.E.	Yes	Yes	Yes	Yes
Plant F.E.	Yes	Yes	Yes	Yes

Figure A.1: Monthly Coal Price Indices

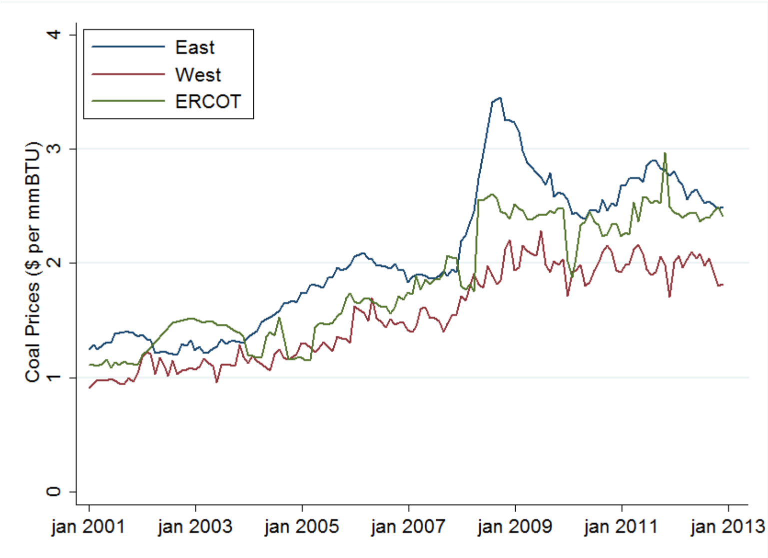

Gas Prices and Coal/Gas Cost Ratio

For comparison with the coal prices, Figure A.2 shows variation in gas prices over time for each of the three regions. Figure A.3 shows the variation in the ratio of the coal to gas price which is the variable of interest in our estimation procedure. Both gas prices and the ratio show substantial variation over time and across regions.

Figure A.2: Daily Gas Price Index by Region

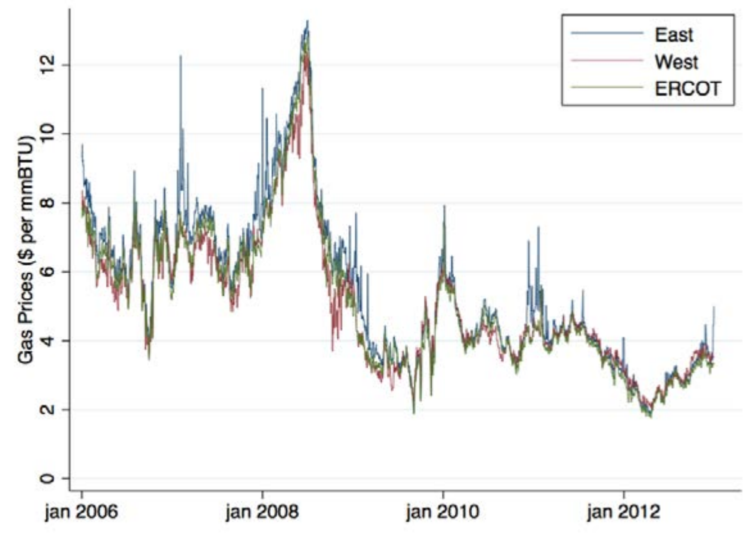

Figure A.3: Coal/Gas Price Ratio by Region

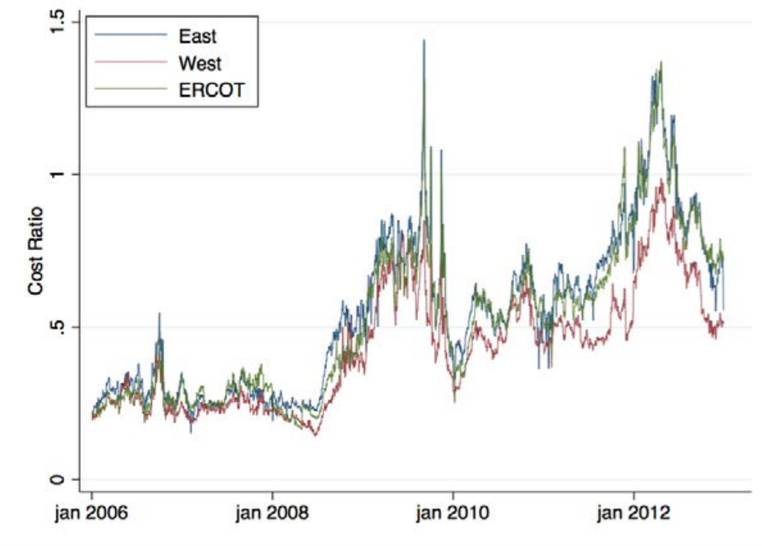

Simple Model of Potential Fuel Switching

This section of the appendix describes the methodology that we use to calculate the potential for fuel switching. First we calculate the total electricity generated (gen_{ift}) in interconnection i, fuel type f, and hour of sample t:

$$gen_{ift} = eiagen_{ifm} \cdot \frac{cemsgen_{ift}}{\sum\limits_{t \in m} cemsgen_{ift}},$$

where $eiagen_{ifm}$ is the aggregate monthly net generation reported in EIA form-923 and $cemsgen_{ift}$ is the hourly gross generation reported by CEMS. In other words, we use the variation within a month reported by CEMS to distribute the EIA monthly generation.

Next we calculate the nameplate capacity by fuel type, month, and interconnection. We define unused capacity as the difference between hourly generation and available capacity, where available capacity is nameplate capacity that is derated to account for the fact that power plants shut down for routine maintenance of because of forced outages. We test the robustness of our calculations to several different derating factors (75% to 100% for each 5% increment).

Finally, we calculate the carbon implications by comparing the unused capacity of natural gas plants within an interconnection and hour with the contemporaneous generation from coal-fired power plants. For each hour and interconnection, we calculate the generation from coal generators that could be reallocated to idle natural gas capacity. This to produces a measure of potential fuel switching. Surprisingly, in most hours, we find that there is substantial unused gas capacity to completely offset all coal generation, even for low derating rates. In order to convert generation into potential carbon reductions, we use the emissions factors mentioned in footnote 15 and the average heat rate by year, interconnection and fuel type: $(\sum_{m \in yr} eiaheat_{ifm} / \sum_{m \in yr} eiagen_{ifm})$, where $eiaheat_{ifm}$ is the aggregate monthly heat input reported in EIA form-923 for interconnection i, fuel type f, and month m in year yr.

Note that this calculation makes many assumptions about transmission capacity, power plant operation capabilities, information, and incentives that we argued in Section 2 were unreasonable and motivation for a more careful analysis that we revisit in the following section. Nonetheless, we report the simple model in order to get a sense of how much unused gas capacity is available.

Results and Data Distribution

In this section, we take the results shown in figures 5 and 6, and overlay histogram of variable of interest to show the density of data that identify the curve. We also show the location of the knot points used in estimation. These are shown for the East electricity generation region, but the pattern is similar in the West and ERCOT. Figure A.4, shows the results as gas prices decrease with the coal price fixed at the long run base case as in Figure 5a. The six knot points spaced according to the distribution of the cost ratio at the 5, 23, 41, 59, 77, 95 percentiles. Importantly, the histogram shows that the data is dense in the areas where the gas price is relatively low. Identifying the response of generators to low gas prices is what allows the model to make predictions about the response of generators to a price on carbon.

Figure A.5 transforms the results be a function of carbon prices as detailed in the body of the paper. The curve is identical to the one in Figure 6a, but again we have imposed the knot points and the histogram of the data onto the estimated response curve. There are many data points up through about \$60/ton after which the density of the data begins to be stretched out. Note that only 5 knot points show up in this graph. This because we only report the results for carbon price less than \$80/ton. There are implied carbon prices in excess of \$200/ton, but the data are sparse for higher carbon prices. Thus, we have to rely more on the function form to identify the behavior of generators for very high carbon prices. Also, these prices less interesting from a policy perspective. For comparison, the results over the full range of implied carbon prices is shown in Figure A.6.

Figure A.4: Estimated CO_2 Response to Fuel Prices

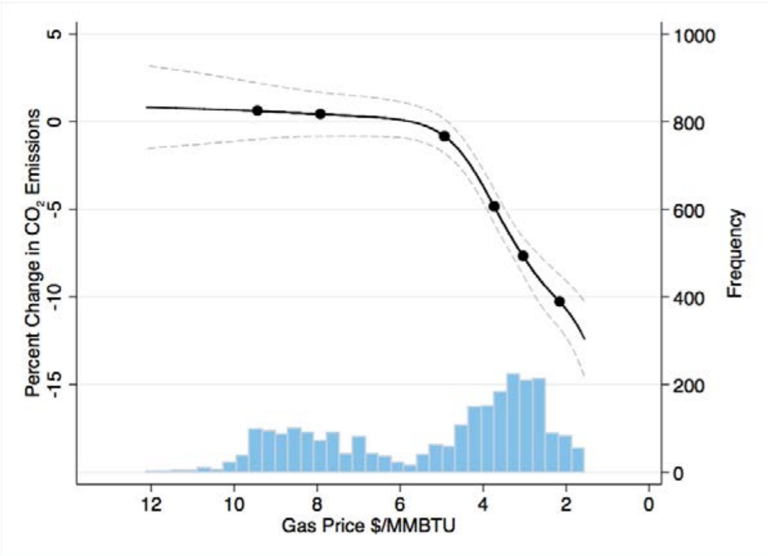

Figure A.5: Imputed CO_2 Response to Carbon Prices

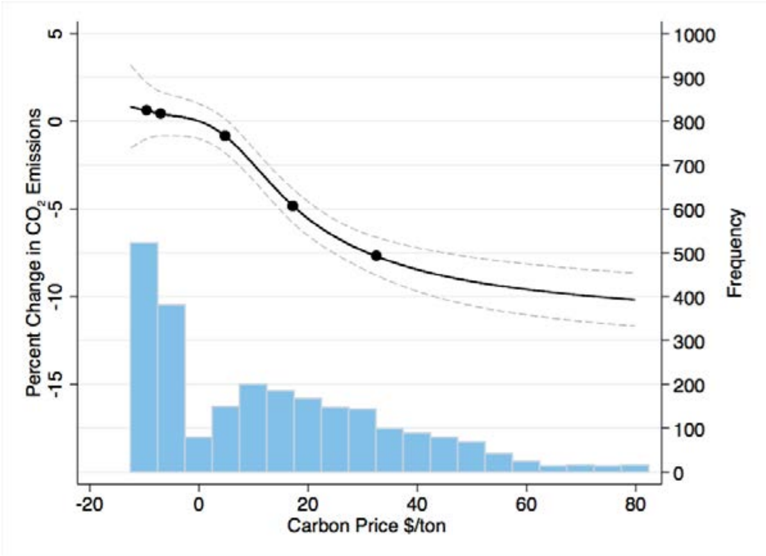

Figure A.6: CO_2 Response for Full Range of Imputed Carbon Prices

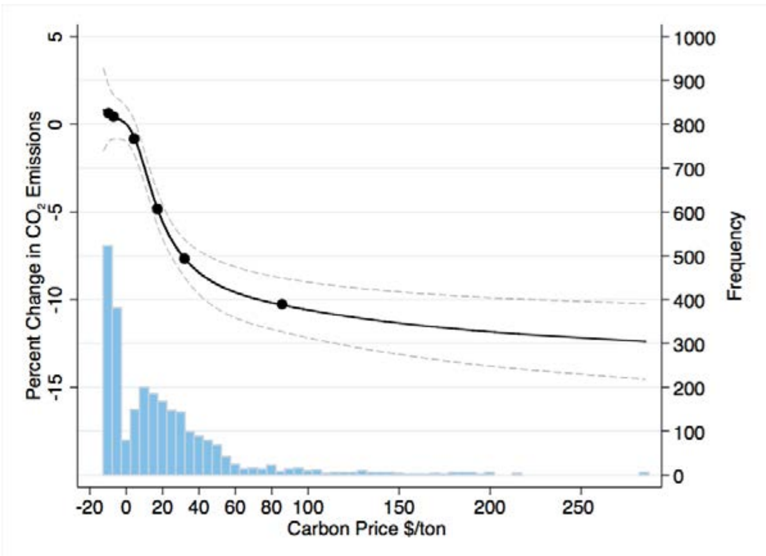

Robustness to Function of Relative Fuel Costs

Figure A.7 estimates the level of emissions in the East as a function of fuel prices using three different functional forms of relative fuel prices. They are (1) our main specification of coal/gas ratio, (2) the inverse ratio of gas/coal, and (3) the price difference of the natural gas price minus the coal price (in \$/mmBTU). All three ratios show very similar mappings of gas prices to emissions, holding coal prices fixed.

Figure A.7: Robustness to Functional Form of Fuel Costs

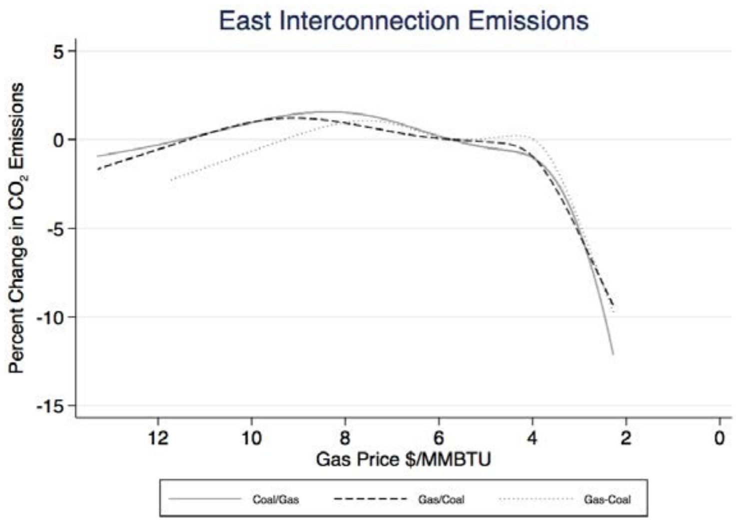

Robustness to the Number of Knots

The main results use six knots in the cubic splines of several variables, including the cost ratio. Figure A.8 shows how the predicted emissions in the East change with the number of knots in comparison to Figure 5a. We see that for four and five knots, the results are virtually identical to the model with six knots and lie completely within the 95% confidence of the six knot specification. For the model with only three knots, the emissions response is overstated at high gas prices, where there are few observations. This model also smoothes over the sharp drop in emissions around $5/mmBTU. The model with seven knots is more sensitive to noise in the data that leads to non-monotonicity in the response curve. It is also almost entirely within the 95% confidence interval of the six knot specification. Overall, the results are qualitatively and statistically insensitive to the number of knots used to form the cubic spline.

Figure A.8: Robustness to Number of Knots

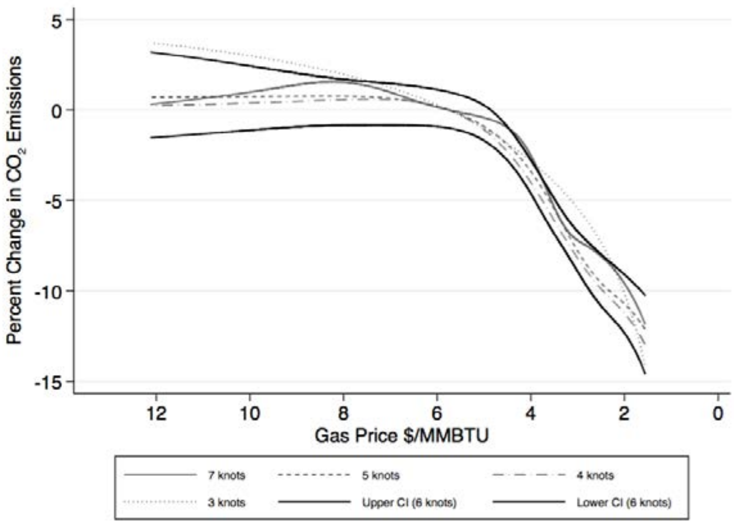

Electricity Demand

As noted in section 7.5, the analysis does not account for the change in electricity demand that may occur when equilibrium electricity prices increase under a carbon tax. Although we cannot impute the demand response to higher counterfactual electricity prices, we can understand how a price on carbon would affect emissions if demand were lower. We do this by splitting the sample by the median daily demand in each month of the sample. That is, we define high demand days in a month to be those days whose demand is higher than the median demand in that month. Low demand days are those below the median. We do this by month of sample to avoid selection on seasonality or time trend. That is, we avoid comparing only winter months to only summer months or the beginning of the sample to the end of the sample. In a month, some days will be happen to be higher demand due to weather and other shocks to demand. This limits the difference in demand between the high and low demand groups, but ensures a comparable sample. The average difference between the high and low demand sample is about 10% of demand as shown in table A.2. This can be seen in the kernel densities of the two regimes. Figure A.9 shows the density of demand has a similar shape but shifts in the high and low demand regimes. Although there is a shift in the distribution of demand, the distribution of gas/coal cost ratios is almost identical in the two regimes as shown in figure A.10. We use the same specification used for our main results to estimate the response of emissions to variation in the cost ration in both the high demand and the low demand sample. The results, shown in figure A.11, demonstrate that there is very little difference between the high and low demand samples. Only in ERCOT at very low gas prices is there a divergence. As we would expect, low demand implies a higher response to relative prices.

Table A.2: Average Daily Demand (GWH) in High and Low Demand Samples

	Low	High	% Difference
East	7053	7855	10.2
ERCOT	810	922	12.2
West	1758	1910	8.0

Figure A.9: Distribution of Demand by Demand Regime: Eastern Interconnection

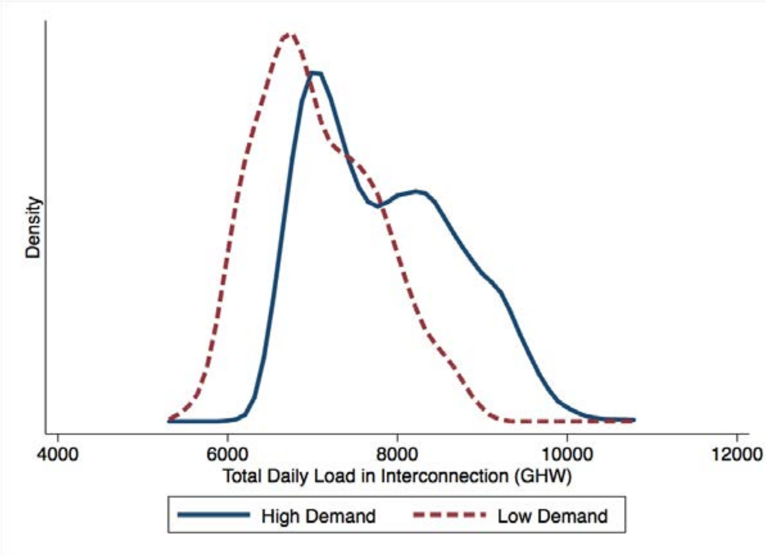

Figure A.10: Distribution of Gas/Coal Ratio by Demand Regime: Eastern Interconnection

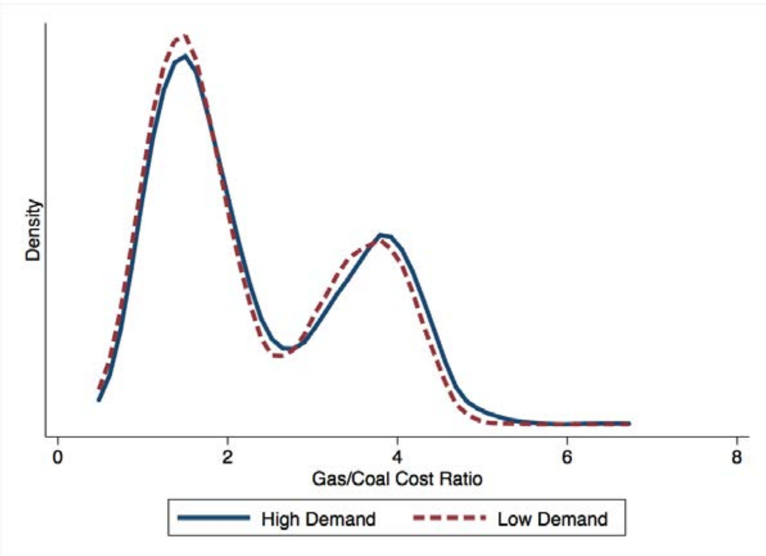

Figure A.11: CO_2 Response in High and Low Demand Periods

(d) Eastern Interconnection

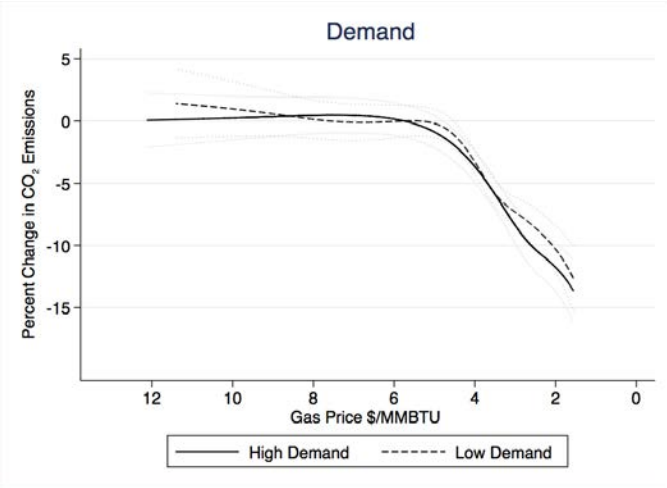

(e) ERCOT Interconnection

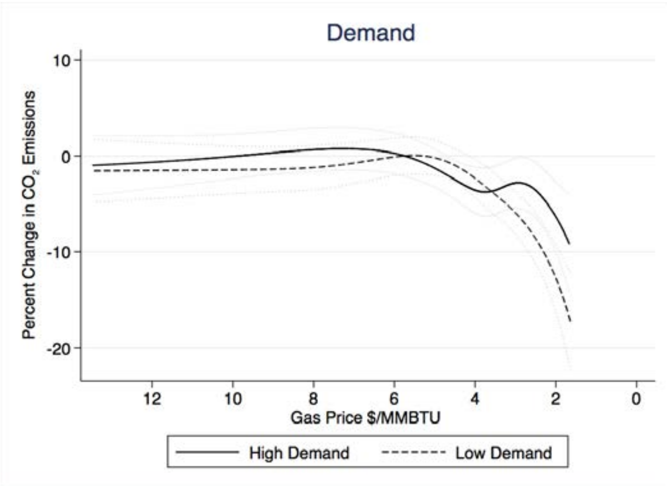

(f) Western Interconnection

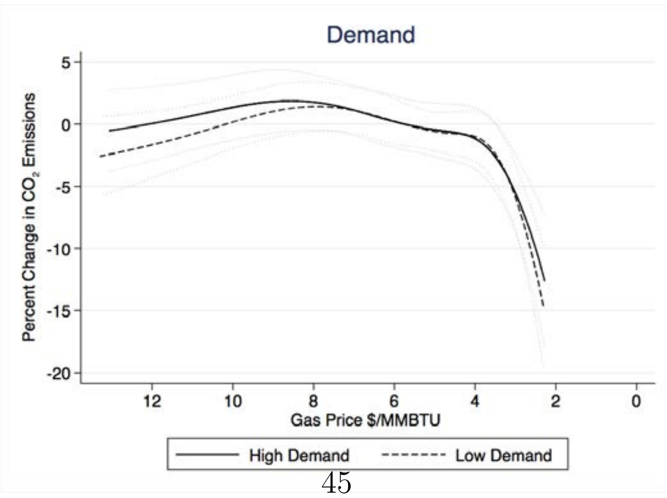

Co-pollutants

This section complements the analysis shown in Figures 8 and 9 of the text by examining how the effects of SO_2 and NO_x are disbursed spatially. The different subregions of the US are defined in the web appendix. Figure A.12 shows how a price of $20 per ton of carbon dioxide would affect CO_2, SO_2, and NO_x emissions in each region. The method is described in the text.

We see from the figure that the emissions response varies regionally. The subregions SRVC (the North Carolina, South Carolina, and Virginia region), RFCE (the New Jersey, Maryland, Delaware, and eastern Pennsylvania region), and NEWE (New England) show the largest reductions of about ten percent for CO2.

However, this does not lead to the largest percent reductions in local pollutants. The co-benefits (in percentage terms) are largest in SRVC and SRSO (the Alabama and Georgia region) for NO_x, and NEWE, SRVC, and CAMX (California) for SO_2. Note that California has very little SO_2. The regional responses of three pollutants are positively correlated across pollutants. However, they do show very different patterns.

Figure A.12: Emissions Response to Carbon Price by Subregion