The Effect of Unique Abatement Options in the U.S. SO_2 Emissions Trading Program and Inferences about the EPA Clean Power Plan

Kevin Griest

March, 2016

Introduction:

Tradable permit markets have long been pushed for by economists, claiming that market based incentives lead to more efficient abatement measures. However, until fairly recently with the Clean Air Act's Acid Rain Program, there had been little implementation of this technique in environmental protection policies. Now, tradable emission permits are gaining attention with policy makers around the world due to the success of the large scale implementation of this policy in the U.S. Acid Rain Program. The Kyoto Protocol came into effect in 2005 which uses an international emissions trading system to reduce global greenhouse gas emissions (GHGs) of developed countries, while California also implemented a cap and trade program for emissions under AB-32 starting in November 2012. These examples of recently developed cap and trade systems along with numerous others have provided evidence that this policy instrument is an economically efficient method of attaining desired environmental regulation.

Other policy tools that have been commonly used in the past instead of tradable permits include standards and taxes. The underlying idea of almost all environmental regulation is that the true cost of production is higher than is represented by the private market because of unaccounted for externalities. Thus it is the government's responsibility to protect the injured parties by internalizing the externality through taxes or some other policy. Under an emissions standard, all power plants are usually required to meet certain physical restrictions which is efficient when all the targeted units have similar characteristics, but this often places overly strict requirements on less technologically advanced units when there is firm heterogeneity. A tax's impact on the other hand is relative to the individual firm's characteristics, making it a more attractive measure for addressing markets with large variation in firm technology. However, as Zhao (2003) shows, assuming rational expectations and uncertainty about future permit prices and abatement costs, tradable permit markets reduce a firm's incentive to invest in permanent abatement technology less than a fixed emissions tax. This means that in theory, tradable permit markets will result in a more efficient populations of firms relative to emissions standards, and that under uncertainty they also provide the strongest incentive to invest in permanent abatement technology. Although the economic theory has been suggesting this for years it wasn't until the U.S. Acid Rain Program was implemented that a large scale market was able to be studied empirically.

The market based system implemented in the U.S. SO2 emissions trading program appears to have been extremely effective in reducing the costs of complying with EPA SO2 regulation. While initial estimates of permit prices ranged from 250 to 1500 dollars per ton of SO2, actual prices stayed below 200 dollars per ton for most of the compliance period, indicating that firms were able to achieve the necessary emissions reductions at a much lower than expected cost. Schmalensee (1998) also estimated that this program saved between 225 and 275 million dollars per year compared to a classic emissions standard, demonstrating its effectiveness relative to other policy options.

However, a critical difference between SO2 and CO2 policy is the ability to install scrubbers at existing plants to control SO2 emissions, a technology that cannot be used to reduce CO2. This paper will look at the U.S. SO2 emissions trading market under the Acid Rain Program, and estimate the impact SO2 scrubbing technology and low sulfur coal had on the ease of compliance. This will provide insight to how potential CO2 markets under the EPA Clean Power Plan may perform considering the lack of CO2 scrubbing technology. Specifically, we will run a time series OLS regression with EPA and EIA data to determine how the number of operating Flue Gas Desulfurization units and average sulfur content of coal used for electricity generation contributed to the reduction in SO2 emission rates among regulated US coal fired power plants. Because of the ARIMA nature of each individual time series, we will take into account auto-correlated residuals and run several procedures to estimate the true relationship between the variables. Once we have an accurate model of the integrated relationships between these variables, we can back out the coefficients and estimate how coal-fired power generation would have been affected had scrubbing not been available, giving us insight to the importance of scrubbers in reducing the cost of Acid Rain Program compliance.

Background:

The EPA Acid Rain Program under Title IV of the Clean Air Act was enacted in 1995 in order to reduce the SO2 and NOx emissions from electric generating units that were contributing to the production of acid rain. This program set a permanent cap on SO2 emissions with interim targets between 1995 and 2009 and the final target of 8.95 million tons set for 2010. Because of the heterogeneity of targeted power plants, the EPA established a SO2 emissions permit market where the majority of the allowances were grandfathered based on historic data, and a few were held for auction in order to give an indication for the proper market price. Under the system, plants were required to have one permit for every ton of SO2 emitted during the year and unused permits could be banked for future years. Between 1995 and 1999 only the worst emitting plants labeled "Table A" units were regulated, but the EPA also allowed the owners of the Table A units to include other less dirty plants as "substitution" units in order to shift the load to a more efficient plant and reduce the cost of abatement. However, Phase II of the program came into effect in 2000, requiring almost all U.S. power plants to significantly reduce SO2 emissions. Looking back, the ARP has been considered extremely effective, with the realized costs of permits much lower than predicted and substantial over-compliance in Phase I. This indicates plants were able to lower emissions relatively cheaply and that despite uncertainty about future allowance prices and abatement costs, firms were strongly incentivized to adopt permanent abatement technology in the form of sulfur dioxide scrubbers. The success of this program has been analyzed extensively by economists and explanations for the observed over compliance and low permit costs further substantiate economic theory on the efficiency of market based incentives.

Schmalensee (1998), presents two possible explanations for the over compliance seen in 1995 and 1996. Both explanations rely on the fact that in the early 90s experts predicted SO2 permit prices to be anywhere from 250 to 700 dollars per ton of SO2, with the expectation that prices in Phase II of the program would be substantially higher than in Phase1. The first reason offered for the over compliance seen in the early years of the program is that because of the expected increase in future permit prices, firms took advantage of the banking and substitution provisions by decreasing the utilization of dirty plants early on to save allowances for future periods. Under the assumption the price of permits increases at the rate of interest, firms are incentivized to abate more than required in order to save the lower costing Phase I permits for use in Phase II. However, while the data does show a surprisingly high level of banked permits in these periods, the plants under regulation of Phase I actually increased their heat input in these years which suggests that other mechanisms were significantly contributing to abatement.

The second explanation Schmalensee (1998) provides is widely regarded as the main cause of the overall success of the program and has been the subject of much more analysis. This is that railroad deregulation and increased competition in the rail sector dramatically reduced coal transportation costs, which often accounts for the majority of power plant fuel costs. Because of this, the relative price of surface mined, low sulfur coal from the Powder River Basin in Wyoming decreased substantially compared to the high sulfur coal mined in the mid and eastern regions of the United States. This allowed power plants across the country to switch to low sulfur fuel at almost no cost and resulted in substantial SO2 emissions reductions. However, Kunce (2005) shows that the extensive fuel switching to PRB coal was not necessarily the direct result of abatement incentives provided by the ARP, but most likely a natural transition due to the decrease in rail transport prices. Although the fuel switching may not have been a mechanism of the emissions trading market, it still had a large impact on the program's success.

Because this nearly cost free abatement was not widely forecasted, many firms overinvested in scrubbers for Phase I plants. A number of papers support this idea, as the amount of allowance banking seen in these two years far surpassed analysts' predictions. Compounding on this, is the incentive for plants to install scrubbers early to maximize their return since the lifespan of an electric generating unit (EGU) is finite. Because the transition to low sulfur coal would likely have occurred whether Title IV of the Clean Air had passed or not, it is of interest to look specifically at the impact scrubbers had on the success of the program. Schmalensee (1998) estimates that had there been no fuel switching to PRB coal, the average cost of abatement seen in 1995 would have increased from 187 dollars per ton to 210. We will contribute to this by estimating the long run impacts of scrubbing and fuel's sulfur content on emission rates over time. This will allow us to estimate whether or not fuel switching would have been adequate to meet EPA caps or if scrubbing was an essential component.

The most recent U.S. legislation encouraging the use of tradable permit markets is the EPA Clean Power Plan which was announced August 3, 2015. This plan enforces a strict emissions standard for three categories of power plants: New, Modified, and Reconstructed; but also includes CO_2 performance goals for each state's mix of existing EGUs. The policy is split into these two pieces in order to manage the heterogeneity of existing plants while also ensuring the more homogeneous new plants are equipped with the most efficient emissions technology. The two types of existing EGUs targeted under the performance goal are fossil fuel-fired steam plans (coal or oil) and natural gas combine cycle plants. These performance goals can be met primarily through two avenues: a rate based goal for each state's total pounds of SO_2 emitted per MWh generated, and a mass based goal in total tons of SO_2 emitted. There are a number of interim targets between 2022 and 2029, with the final goal of a 32% reduction in national CO_2 emissions set for 2030. Although the EPA determines these goals based on each state's mix of EGUs, they give each state power over implementing their own system to meet these goals and even allow multiple states to work together. States can employ a variety of methods to comply but because they are allowed to combine efforts with other states as well, emissions trading markets may be very attractive. This would allow states with more flexibility in reducing CO_2 emissions to sell unused permits to states with more restrictive plant technologies such as West Virginia, whose power generation is comprised of more than 95 percent coal. If the results of the Acid Rain Program are any indication to the potential performance of these hypothetical CO_2 emissions markets, states will be strongly incentivized to adopt the market based system rather than a command and control strategy as has been seen in the past.

However, unlike sulfur dioxide there are no established methods of direct CO_2 abatement other than improving the heat efficiency of the plant, switching from coal to natural gas which has a lower emissions rate, or switching to emissions free sources such as wind or hydroelectric. This makes the investment in permanent CO_2 abatement measures much more complicated because there is limited capacity for renewable energy and relatively small NG reserves. Also the capital costs required for these abatement options are much higher than the cost of installing SO2 scrubbers. Furthermore, unlike sulfur the amount of CO2 emitted per million Btu does not vary significantly across different sources of coal. According to the EPA there is only a 4.5% difference between Bituminous and Sub-bituminous CO2 emission rates: the two types of coal primarily used for fuel at EGUs. Thus the only option for abatement that does not require intensive capital investment is heat rate improvements.

A generating unit's heat rate measures the amount of energy used to generate 1 kilowatt-hour of electricity. In 2015, the EIA released their "Analysis of Heat Rate Improvement Potential at Coal-Fired Power Plants" which used the National Energy Modeling System (NEMS) to statistically model potential heat rate improvements and costs for "32 configurations of existing coal-fired plants". This analysis estimated heat rate potentials for U.S. plants to be between 2% and 5.3% with an average potential of 4% improvement. They also estimate that the capital costs required for these improvements would average $300,000 per MW and that average fixed O&M costs were approximately $6,000 per MW-yr. However, lower fuel costs resulting from these improvements would offset a large portion of the upgrade cost. Unfortunately, a one percent heat rate improvement only results in about a 1 percent CO2 emission reduction as well, which means even if power plants took advantage of all heat rate improvements, they would only reduce emissions an estimated 4% which is substantially less than the 32% reduction the Clean Power Plan requires by 2030. A number of other sources have analyzed this issue extensively and come to slightly differing opinions on efficiency potential and heat rate reduction costs, but none claim that there exists enough potential improvement to meet the EPA CO2 cap without implementing other abatement measures as well.

One other potential CO2 abatement technology is Carbon Capture and Storage. CCS is a fairly new technology that has the potential to substantially reduce CO2 emissions from existing plants without fuel switching, but it is currently extremely expensive because it requires the capture, compression, and long term storage of CO2. Rubin (2007) provides an in-depth analysis of the costs and benefits of CCS technology including the fact that these systems have negative impacts on the heat rates of plants, reducing overall efficiency and increasing fuel costs. Despite its current issues, CCS technology has the potential to have a huge impact similar to scrubbers if costs can be reduced in the future.

Our paper will address these differences between the SO2 and CO2 markets by estimating the impact SO2 scrubbers and low sulfur coal had on the success of the ARP between 1995 and 2006. By identifying how each abatement option effected SO2 emissions on its own, we can estimate what the emissions market would have looked like without them, and thus make inferences about how CO2 markets with similar options may behave. We will also be able to use our results to check previous findings on the relative impacts of scrubbing and sulfur content.

Data and Methodology:

Because the power plants cannot use scrubbers to reduce CO2 and most coal has similar CO2 content, we want to estimate the impact both these abatement options had on the sulfur market in order to make inferences on how similar CO2 markets may behave. To do this, we used plant level data from EIA form 423 which contains information on fuel characteristics including monthly cost of purchase and delivery, quantities delivered, and sulfur content of coal delivered to EGUs. We also took data from the EPA database for clean air programs on the SO2 and CO2 emissions levels, gross load, and number of operating SO2 scrubbers for coal fired power plants under the Acid Rain Program. Unfortunately, due to issues identifying specific boilers across different data sources we were unable to use the data to do a cross sectional analysis, so instead we aggregated it into national monthly data and ran a time series cointegration analysis. A summary and graphs of the data used for the cointegration analysis can be seen in table 1 and figure 1 respectively, where Y is SO2 emissions rate, X1 is number of scrubbers, X2 is the average sulfur content, and .diff indicates the first difference. The discontinuity in both the SO2 rate and number of scrubbers, seen in figure 1, is due to the inclusion of phase II plant data in the EPA data set starting in 1997.

Table 1: Descriptive Statistics

	SO2 Rate	SO2 Rate diff	Scrubbers	Scrubbers diff	Sulfur Content	Content diff
Min	.004270	(.002323)	47	(4.000)	1.050	(.0700)
1st Qu.	.004915	(.000099)	197	0.000	1.110	(.0200)
Mean	.005214	(.000020)	210	0.000	1.160	.0000
Median	.005857	.000031	191	1.406	1.206	(.0017)
3rd Qu	.006455	.000067	232	0.000	1.310	.0100
Max	.008966	.000474	252	143.000	1.410	.0700

( ) indicates negative value. SO2 Rate is tons of SO2 per MWh of load. Scrubbers is the number of operating scrubbers. Sulfur content is average sulfur content of coal used for electricity generation as a percentage of weight.

Figure 1: Time Series Plots

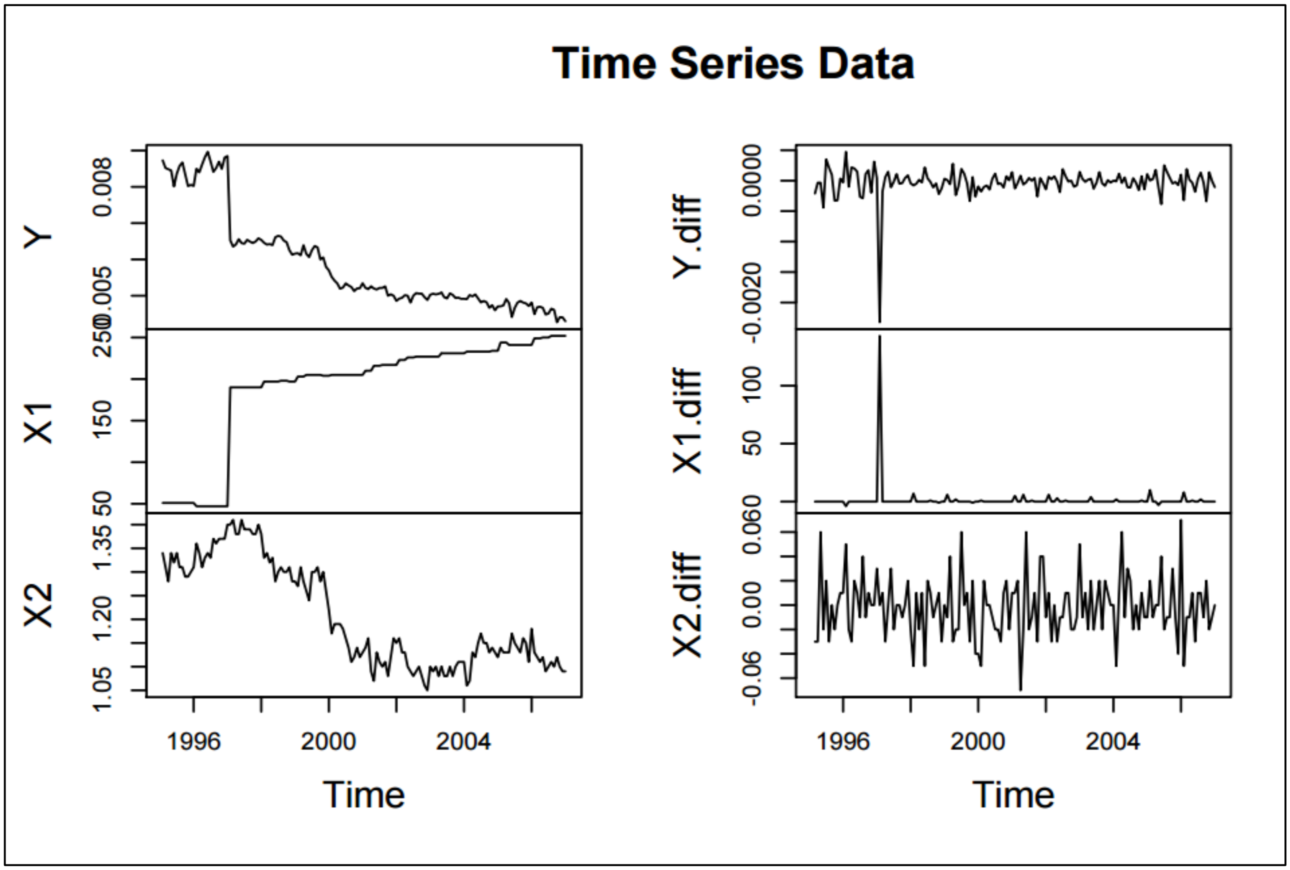

A major issue with running regressions on economic time series data is that the variables are often I(1) series. This means that there is either a deterministic trend or a random walk stochastic process in the individual variable's time series. The problem with this is that randomly generated stochastic processes like these are often highly correlated with each other despite no underlying relationship. This is known as spurious regression and has been demonstrated through both simulating random series and analyzing real data that should in no way have a relationship. In order to address this, we will use Engle and Granger's two step method of cointegration, which allows us to identify with greater confidence the underlying long-run equilibrium as well as the short-run behavior of the series. The general idea of the model is that if the each original series is I(1), indicating that it achieves stationarity after only one difference, and there is a true relationship between the variables, then there exists a linear combination of these variables that achieves stationarity. This would indicate that the series do not vary too far from each other at any given time and that the static relationship estimated by OLS regression holds through time. Furthermore, we can analyze short-run changes in the variables and their interaction with the cointregration relation's error, to test how the series behave when they are not at equilibrium. This is commonly known as an Error Correction Model and its residuals should appear to be white noise. Intuitively, we hope these models indicate a negative relationship between the SO2 rate and number of scrubbers and a positive relationship between the sulfur content of fuel used and SO2 emissions rate.

Applying this method to our data, we first analyze the average SO2 emissions rate, number of operating SO2 scrubbers, and average sulfur content of fuel used by coal fired power plants to check the order of integration using two unit-root tests for robustness: Augmented Dickey Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS). The null hypothesis of the ADF test is that there exists a unit root vs. the alternative that the series is stationary, while the null hypothesis of the KPSS test is that the series is stationary and the alternative is that there exists a unit root. The results of these tests when applied to the original series and then the first difference, as seen in table 2, indicate that all three series are I(1) variables.

Table 2: Unit root tests on individual series

Test	ADF	ADF	KPSS	KPSS
Series	Level	First Difference	Level	First Difference
SO2 Rate	-1.908	-5.548 *	4.101*	0.106
Scrubbers	-2.114	-5.037*	3.387*	0.150
Sulfur Content	-1.455	-5.642*	3.961*	0.071

ADF: Augmented Dickey-Fuller Test Statistic. KPSS: Kwiatkowski–Phillips–Schmidt–Shin Test Statistic. Null of ADF is presence of unit root. Null of KPSS is stationarity. * indicates rejection of the null hypothesis at the 5% significance level.

This indicates that a C(1,1) cointegration model can be appropriately applied to test for underlying relationships. We also looked at the autocorrelation and partial-autocorrelation functions (acf and pacf) of each series and its first difference to better identify the behavior of each series. The acf measures the correlation between two data points of a series at different lag periods and the pacf is an estimate of these correlations after removing the effect of all other lag periods' correlations. Autoregressive, integrated, and moving average processes, or any combination of these can be represented as an ARIMA(p,d,q) model, where p, d, and q represent the order of each process. After taking a first difference, the acf and pacf of the average SO2 emissions rate appear to resemble white noise, implying that the variable is a simple ARIMA(0,1,0) process. We can double check this by running Durbin Watson and Ljung-Box tests which affirm the first difference results in white noise. Repeating this for the number of scrubbers and average sulfur content shows that the number of scrubbers also follow an ARIMA(0,1,0) model while the average sulfur content seems to be an ARIMA(0,1,2) process. Knowing these variables are I(1), we can now run an OLS regression and analyze the residuals to determine whether there is a cointegrated relationship. As seen in figure 2, the linear combination of emissions rate, number of scrubbers, and sulfur content estimated by OLS regression (the residual time series) appears to be stationary with a mean of 0.

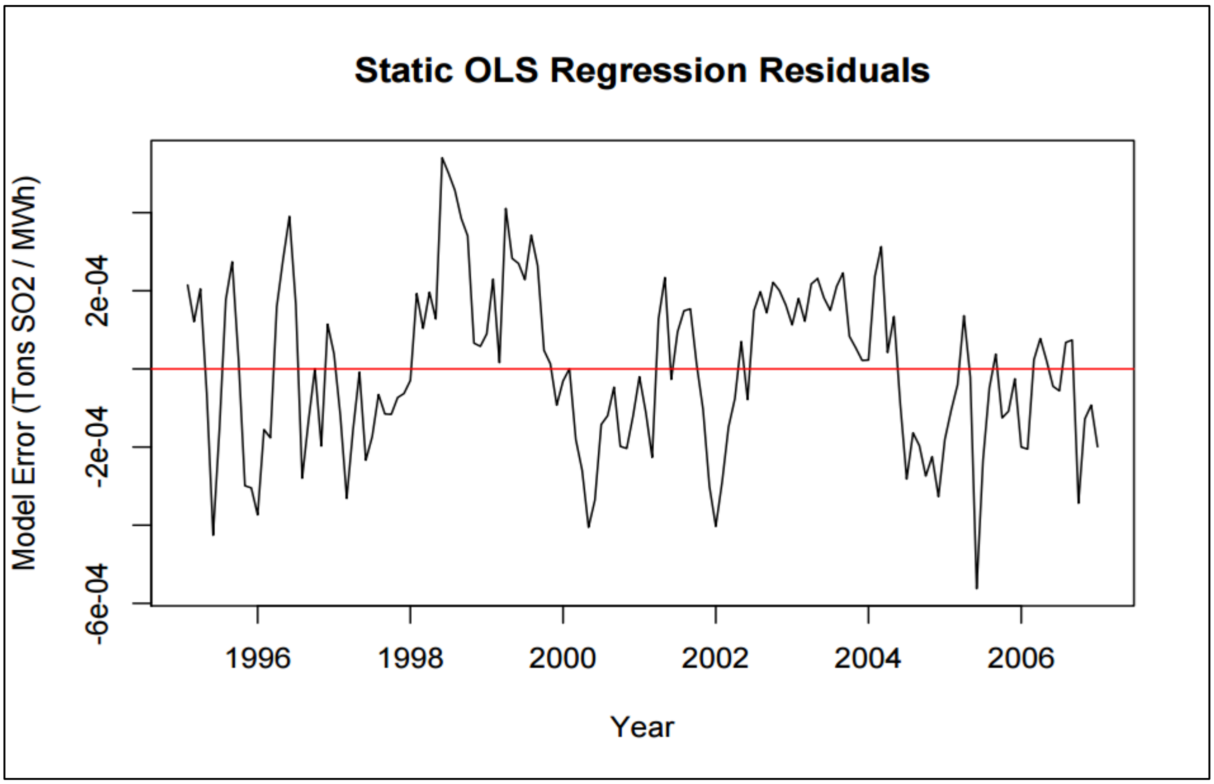

Figure 2: Time series of OLS residuals

- Indicates the zero line. Errors appear to be evenly spread around 0.

Testing this with ADF and KPSS tests confirms that the residuals are a stationary series. However, looking at the acf and pacf, we can see that there appears to be a AR(1) process in the errors. The Durbin Watson and Ljung-Box tests also confirm that the errors are significantly different from white noise, but all that is needed for cointegration is stationarity so we are satisfied with these results. Our regression results in the estimated model;

$$so2\ rate = c1 + b1*(number\ of\ scrubbers) + b2*(sulfur\ content) \quad (1).$$

where SO2 rate is tons of SO2 per MWh, and sulfur content is percentage of sulfur by weight. This represents the long-run equilibrium relationship between the variables and the signs of the coefficients seen in table 3 are what we should expect to see intuitively.

Table 3: Static OLS Regression

Variable	Estimate	t statistic
Intercept	.00405572	13.01***
# Scrubbers	(.00001514)	(42.16) ***
Sulfur Content	.00390378	18.06 ***

( ) indicates negative value. *** indicates rejection of the null hypothesis at the .1% significance level. The adjusted R-squared of this static OLS Regression is .9756.

The statistically significant negative coefficient on scrubbing indicates that an increase in the number of scrubbers reduces power plants' SO2 emission rates and the positive coefficient on sulfur content shows that the drop in sulfur content due to the increased utilization of Powder River Basin coal also significantly contributed to the lowering of power plant SO2 emissions.

Continuing our analysis, we can now estimate an error correction model. Following the Engle-Granger 2 step process, we will do this by running the OLS regression:

$$so2rate.diff(t) = c2 + a1*so2rate.diff(t-1) + a2*scrubbers.diff(t-1) + a3*sulfurcontent.diff(t-1) + r*E(t-1) \quad (2.)$$

The error term is the stationary cointegrated relationship;

$$E(t\text{-}1) = c1 + so2rate(t\text{-}1) - so2rate\text{-}hat(t\text{-}1)$$

$$= c1 + so2rate(t\text{-}1) - b1*scrubbers(t\text{-}1) - b2*sulfurcontent(t\text{-}1) \quad (3).$$

where b1 and b2 are the coefficients from equation *(1)*. The interpretation of this equation is that the change in SO2 emissions rate is a function of the last period's changes in emissions rate, number of scrubbers, and sulfur content, with an adjustment for the last periods deviation from equilibrium. The parameter *r* can therefore be interpreted as the speed at which the series returns to equilibrium after shocks pull it away. Thus a large value of *r* would indicate that the series has a strong inclination to return to the equilibrium relationship estimated in equation *(1)*. We expect *r* to be negative and less than one, because that would indicate that the series is adjusting for the previous periods disequilibrium. After running this regression, we found that the last period's change in SO2 emissions rate did not significantly contribute so we dropped this term from in order to achieve a more parsimonious model.

Table 4: VECM Results

Variable	Estimate	t statistic
Intercept	(.000006139)	(.531)
Scrubbers.diff(t-1)	(.000016025)	(16.792)***
Sulfurcontent.diff(t-1)	.001134415	2.417*
E(t-1)	(.250738637)	(4.456)***

( ) indicates negative value. * , ** , and *** indicate statistical significance at alpha = 5% , 1% , and 0.1% respectively. The adjusted R-squared of this VECM is .6816.

The estimated coefficients and output from R in table 4 above show that all variables in the final model are statistically significant at an alpha of .05 and the adjusted R-squared of the model is .6816. Furthermore, analyzing the errors reveals that they resemble white noise according to both the Durbin Watson and Ljung-Box tests. This verifies the assumption that our residuals are independent. The estimated coefficients on this model are also in line with our expectations, with a negative coefficient on the change in scrubbers and a positive coefficient on the change in sulfur content. As hoped our parameter *r* is also negative with a magnitude of less than 1, indicating the tendency to return to e equilibrium. Figure 3 depicts the fit of this model to the changes in so2 emissions rates.

Figure 3: VECM fit to SO2 emissions rate first difference

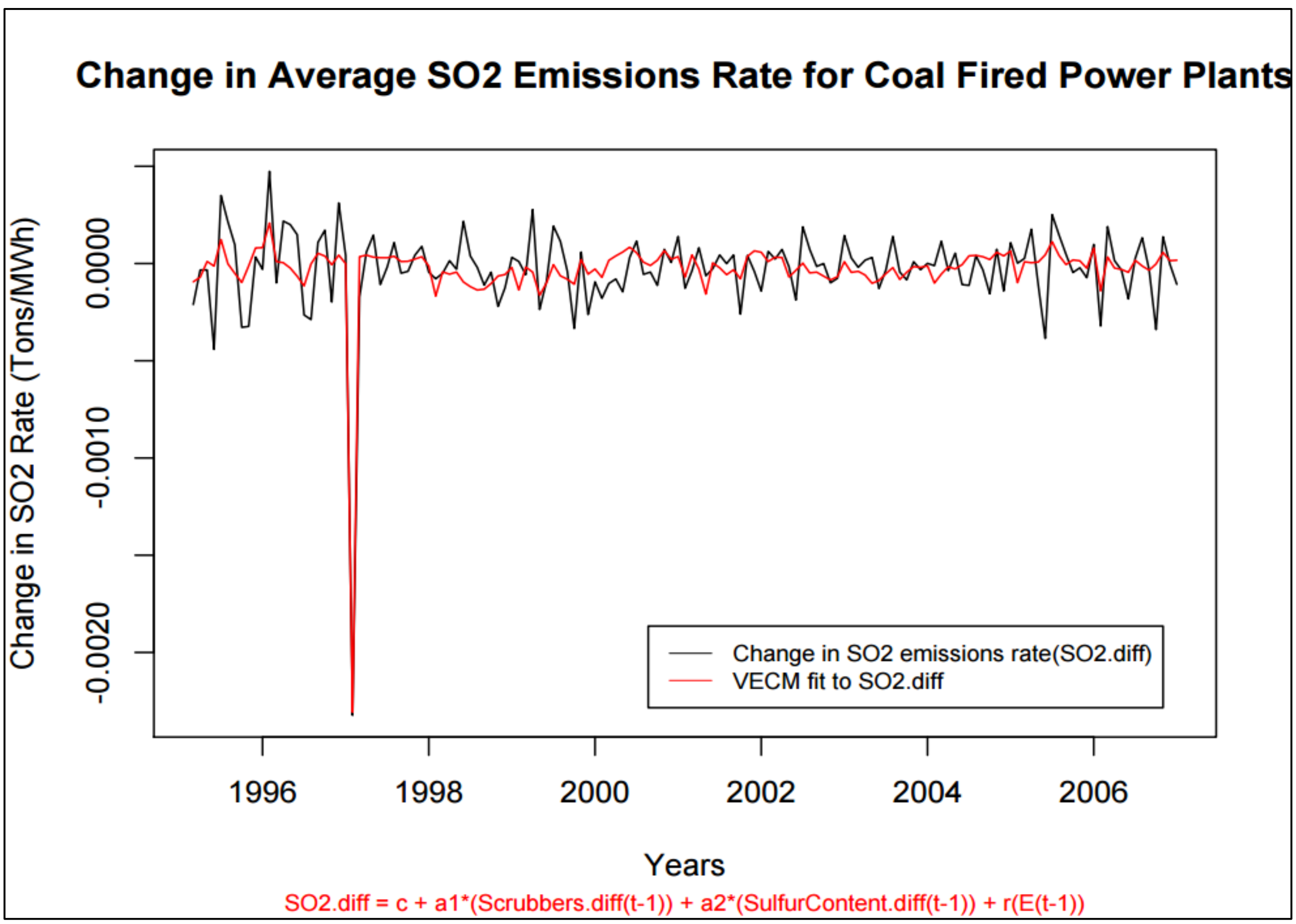

Results:

The results of our cointegration model suggest that in both the short-run and long-run, scrubbers have the largest impact on power plant's SO2 emission rates. Although the sulfur content of coal used by the power plants also had a significant impact, it was not as influential as scrubbing as shown by the t-values of our significance tests. Figure 4 below illustrates our estimation of SO2 emissions rates both with and without scrubbing and low sulfur coal. The blue line shows what we would expect SO2 emission rates to look like if there were no operational scrubbers after controlling for the effect of low sulfur coal.

Figure 4: Long-run Estimate of SO2 emissions without scrubbing technology or low sulfur

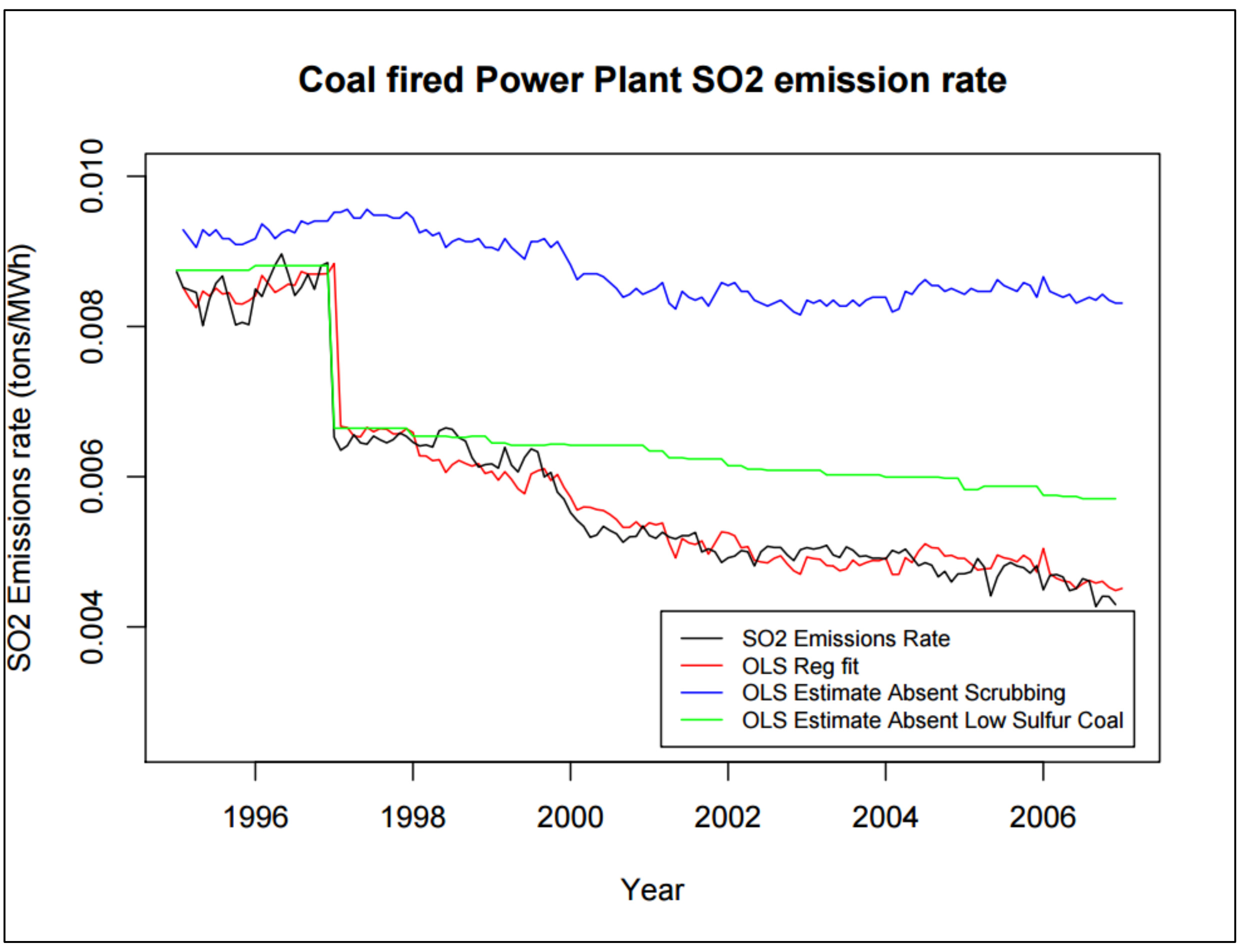

coal

The green line on the other hand shows what we predict SO2 emission rates would have been if the sulfur content of coal stayed at a level close to what was seen in 1995 and prior, after accounting for the impact of scrubbing. Clearly the average emission rate was much higher when we took away scrubbing than when we eliminated low sulfur coal.

Converting the average emission rate to total emissions, we find that without installing scrubbers, the power plants under regulation would not have achieved the target levels between 1997 and 2005. In fact, we estimate that in order to meet the same emissions levels seen between 1995 and 2006 without resorting to other abatement options, coal fired power plants would have had to reduce their average yearly generation by approximately 37%, translating to an average of

Figure 5

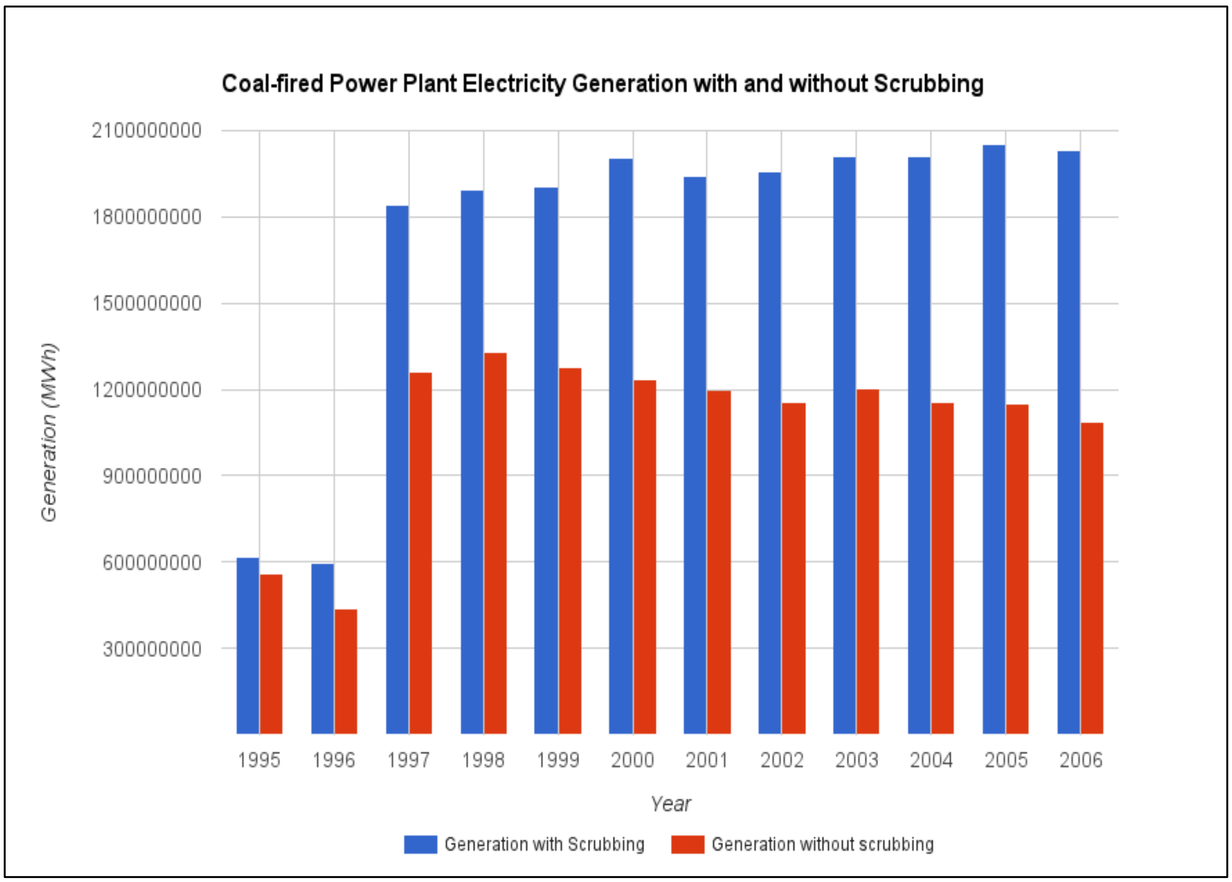

over 600 million MWh in reductions per year. Figure 5 shows the observed coal fired power plants' total yearly generation with scrubbers and the estimated yearly generation needed to reach the same emissions level absent scrubbers. Considering total revenue from U.S. electricity sales was $326 billion in 2006, without substituting other energy sources for the lost production of coal fired plants, power companies would lose almost $45 billion in revenue per year assuming the same average price of electricity. If this were the only option available, electricity prices would instead skyrocket due to the substantial decrease in supply (illustrated in figure 6) and firms would not lose as much revenue as estimated above. However, in reality firms can substitute cleaner energy sources for the lost coal generation. Unfortunately, the capital requirements and other costs associated with switching to cleaner energy sources are substantially higher than scrubbing because they involve building entirely new plants, so the cost of compliance would be substantially higher as well. However, the marginal costs of cleaner sources are actually lower than they are for coal plants because fuel is the primary source of marginal cost and natural gas prices have historically been lower than coal prices. Also, renewable energy plants have almost no marginal costs associated with them because their sources of power: the sun, wind, and water are all free. This suggests that although the initial cost of compliance will be substantially higher in a market without scrubbing technology, the long run price of electricity should actually decrease the most under an emissions trading system because tradable permit markets incentivize permanent abatement investment more than any other policy. Considering CO2 abatement options are severely limited, we believe

Figure 6

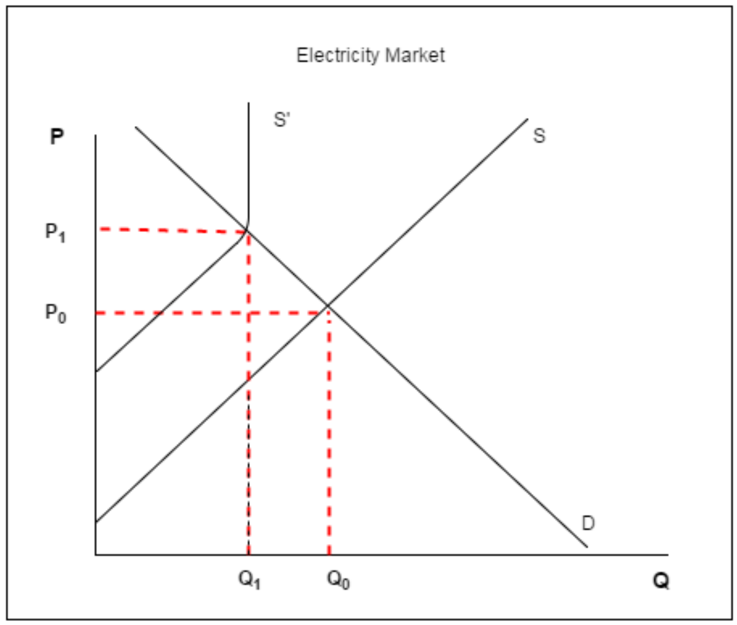

that the cost of complying with the Clean Power Plan through an emissions market will be substantially higher than it was for the Acid Rain Program, but it will also result in a much larger shift to permanent abatement in the form of cleaner energy sources.

References:

Boutabba, Mohamed Amine, Olivier Beaumais, and Sandrine Lardic. "Permit Price Dynamics in the U.S. SO2 Trading Program: A Cointegration Approach." *Energy Economics* 34.3 (2012): 714-22.

"Cap and Trade: Acid Rain Program Results." *EPA*. Environmental Protection Agency. <http://www3.epa.gov/captrade/documents/ctresults.pdf>.

DeLurgio, Stephen A. *Forecasting Principles and Applications*. Boston: Irwin/McGraw-Hill, 1998. Print.

Engle, Robert F., and C.W.J. Granger. “Co-Integration and Error Correction: Representation, Estimating, and Testing.” *Econometrica* 55.2 (1987): 251-276.

"FACT SHEET: Overview of the Clean Power Plan." *EPA*. Environmental Protection Agency, 6 Aug. 2015. <https://www.epa.gov/cleanpowerplan/fact-sheet-overview-clean-power-plan>.

Hanley, Nick, Jason F. Shogren, and Ben White. *Environmental Economics in Theory and Practice*. Houndmills: Macmillan, 1997. Print

Joskow, Paul L. "Restructuring, Competition and Regulatory Reform in the U.S. Electricity Sector." *Journal of Economic Perspectives* 11.3 (1997): 119-38.

Kunce, Mitch, Steve Hamilton, and Shelby Gerking. “Marketable Permits, Low-Sulfur Coal, and the Behavior of Railroads.” (2005). < http://are.berkeley.edu/documents/seminar/Coal.pdf>.

Rubin, Edward S., and Chao Chen. "Cost and Performance of Fossil Fuel Power Plants with CO2 Capture and Storage." *Energy Policy* 35.9 (2007): 4444-454.

Schmalensee, Richard, Paul L. Joskow, A. Denny Ellerman, Juan Pablo Montero, and Elizabeth M. Bailey. "An Interim Evaluation of Sulfur Dioxide Emissions Trading." *Journal of Economic Perspectives* 12.3 (1998): 53-68.

Springer, Urs, and Matthew Varilek. "Estimating the Price of Tradable Permits for Greenhouse Gas Emissions in 2008-2012." *Energy Policy* 32 (2004): 611-21.

Sueyoshi, Toshiyuki, Mika Goto, and Takahiro Ueno. "Performance Analysis of US Coal-fired Power Plants by Measuring Three DEA Efficiencies." *Energy Policy* 38.4 (2010): 1675-1688.

"U.S. Energy Information Administration - EIA - Independent Statistics and Analysis." *Power Plant Emissions of Sulfur Dioxide and Nitrogen Oxides Continue to Decline in 2012*. U.S. Energy Information Administration, 27 Feb. 2013. Web. 15 Jan. 2016.

Zhao, Jinhua. "Irreversible Abatement Investment under Cost Uncertainties: Tradable Emission Permits and Emissions Charges." *Journal of Public Economics* 87 (2003): 2765-2789.