

Di Maria, Corrado; Lange, Ian A.; van der Werf, Edwin

Working Paper

Should we be worried about the green paradox? Announcement effects of the acid rain program

CESifo Working Paper, No. 3829

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Di Maria, Corrado; Lange, Ian A.; van der Werf, Edwin (2012) : Should we be worried about the green paradox? Announcement effects of the acid rain program, CESifo Working Paper, No. 3829, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/58343>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Should we be Worried about the Green Paradox?
Announcement Effects of the Acid Rain Program

Corrado Di Maria
Ian Lange
Edwin van der Werf

CESIFO WORKING PAPER NO. 3829
CATEGORY 10: ENERGY AND CLIMATE ECONOMICS
MAY 2012

PRESENTED AT CESIFO AREA CONFERENCE ON ENERGY AND CLIMATE ECONOMICS, OCTOBER 2011

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

Should we be Worried about the Green Paradox? Announcement Effects of the Acid Rain Program

Abstract

This paper presents the first empirical test of the green paradox hypothesis, according to which well-intended but imperfectly implemented policies may lead to detrimental environmental outcomes due to supply side responses. We use the introduction of the Acid Rain Program in the U.S. as a case study. The theory predicts that owners of coal deposits, expecting future sales to decline, would supply more of their resource between the announcement of the Acid Rain Program and its implementation; moreover, the incentive to increase supply would be stronger for owners of high-sulfur coal. This would, all else equal, induce an increase in sulfur dioxide emissions. Using data on prices, heat input and sulfur content of coal delivered to U.S. power plants, we find strong evidence of a price decrease, some indication that the amount of coal used might have increased, and no evidence that the announcement of the Acid Rain Program led the use of higher sulfur coal. Overall, our evidence suggests that while the mechanism indicated by the theory might be at work, market conditions and concurrent regulation prevented a green paradox from arising. These results have implications for the design of climate policies.

JEL-Code: Q310, Q380, Q530, Q540, Q580.

Keywords: Green Paradox, implementation lags, announcement effects, climate policy, acid rain policy.

Corrado Di Maria
University of Birmingham
Birmingham B15 2TT / United Kingdom
c.dimaria@bham.ac.uk

Ian Lange
University of Stirling
Stirling FK9 4LA / United Kingdom
i.a.lange@stir.ac.uk

Edwin van der Werf
Wageningen University
Wageningen 6700 EW / The Netherlands
edwin.vanderwerf@wur.nl

This version: May 21, 2012

We thank Kathy Baylis, Don Fullerton, Stephan Heblich, Beat Hintermann, Emiliya Lazarova, Mirko Moro, conference participants at the 2011 CESifo Area Conference on Energy and Climate Economics, the 2012 Royal Economic Society Conference, and seminar participants at University of Birmingham, University of Groningen, University of Illinois at Urbana-Champaign, University of Stirling, and Wageningen University for useful comments and discussions.

1 Introduction

Like any other policy aimed at correcting externalities, environmental policy may have unintended detrimental effects if it is not optimally designed to take into account behavioral changes on the part of the regulated agents. A classic example of this type of problems is the introduction of vintage differentiated regulation, whereby older vintages of a given technology are subject to laxer (if any) environmental standards. Under such circumstances, the incentives to invest in more efficient – but also more stringently regulated – technology are greatly reduced. As a consequence, dirtier production processes are locked-in, perpetuating higher than necessary emissions levels (Nelson et al., 1993; Stavins, 2006). A similar situation emerges when regulators want to control fishing efforts mandating maximum levels for the use of specific inputs (e.g. fishing time, number of vessel, type of fishing gear, etc.). Rationally, fishermen react by substituting away from the regulated inputs into permissible ones, and the excessive pressure on the fish populations doesn't abate.

The notion that regulated agents may respond to environmental policy differently than anticipated by policy-makers has recently received new attention, thanks to the thought-provoking contributions by Hans-Werner Sinn (Sinn, 2008, 2012). Sinn argues that most climate policies currently implemented not only fail to provide a solution to the problem of increasing greenhouse gas emissions, but actually aggravate the problem by providing perverse incentives to the owners of stocks of fossil fuels, an outcome he called a 'green paradox'. The crux of Sinn's argument is that resource owners rationally change their behaviour in response to the introduction of environmental policy: as the regulation threatens to reduce future profits, resource owners may modify their extraction plans to increase near-term supply, which increases emissions. When policy makers fail to take into account this type of rational response to the policy shock, the realized effect of the policy on the path of emissions may not be the desired one.

Sinn's seminal contribution has spawned a rich theoretical literature, which discusses several mechanisms that might lead to a green paradox (see e.g. Eichner and Pethig, 2011; Gerlagh, 2011; Grafton et al., 2011; Di Maria et al., 2012; Fischer and Salant, 2012; Hoel, 2012; Smulders et al., 2012; Van der Ploeg and Withagen, 2012).¹ Overall, this literature emphasizes reactions on the supply side, whereby producers shift extraction forward in time, leading to a reduction in current fossil fuel prices. Making fossil fuels cheaper in the short run may have serious consequences as demand for polluting inputs increases, with a consequent raise in short-run emissions.² Until now, however, no empirical investigation has addressed the crucial question whether the effects suggested by Sinn (2008, 2012) and others give, in fact, cause for concern, or they are likely to be tempered by other features of actual markets. This paper fills this gap by presenting what, to our knowledge, is the first empirical test of the green paradox hypothesis.

Using data on coal deliveries to U.S. power plants, we study changes in the price, quantity and quality of the coal in the period between the signing into law of the 1990 Clean Air Act Amendments (CAAA, Public Law 101-549), and the implementation of the Acid Rain Program (ARP, regulated by Title IV of the 1990 CAAA) in 1995. The introduction of the ARP acted as a signal to owners of stocks of coal that it would be harder to sell their

¹Van der Werf and Di Maria (2012) provide a comprehensive overview of this literature.

²Smulders et al. (2012) represent an exception in this literature as they show that a green paradox may also arise in the absence of scarcity. In their model, which completely abstracts from the existence of exhaustible resources, a green paradox arises as agents, anticipating future policy changes, start investing in physical capital ahead of the implementation of the policy. Given the complementarity between capital and pollution, this leads to an increase in emissions as a consequence of the expectation of future policy.

product from 1995 onwards: as it put a nation-wide limit on sulfur-dioxide (SO_2) emissions, the future prospects for coal were consequently restricted.³ According to the green paradox hypothesis sketched above, this would have given mine owners the incentive to increase their supply ahead of the implementation of the Program.⁴ The first testable implication that we derive from the green paradox literature, therefore, refers to the fact that we should observe a fall in the price of coal following the announcement of the 1990 CAAA. A reduction of the price of coal would make coal-fired generation more competitive, moving coal-burning plants down the merit order, and lead to an increase in coal demand. Thus, the second implication of the green paradox hypothesis is that coal-fired utilities would increase their coal input (measured in energy units) over the period 1991-1994. For a given emissions rate (pounds of SO_2 per million Btu of heat input), this would imply an increase in SO_2 emissions. Finally, according to the theory, the incentive to move extraction forward in time would be stronger, the higher the sulfur content of the coal. Indeed, pricing sulfur would make high-sulfur coal more expensive to burn, inducing utilities to substitute towards lower-sulfur coal once the regulation came into effect. High-sulfur mines therefore stand to lose more from the introduction of the ARP, and would have stronger incentives to expand their supply ahead of the upcoming implementation of the cap in January 1995. The third and last testable implication of the theory is, therefore, that the sulfur intensity of coal-fired utilities increased in the interim of the regulation. For a given level of heat input, this would result in a higher level of SO_2 emissions.⁵

We find strong evidence for the hypothesis that prices fell after the announcement of the ARP, as predicted by the green paradox theory. However, this drop in prices did not fully translate into higher emissions. Over all plants in the sample, no statistical increase in heat input is found, but plants that were sufficiently flexible on the coal market are found to have increased their heat input. Regarding the sulfur intensity of coal, we find no evidence that plants switched to dirtier coal. Rather, we find evidence of the opposite, as electricity generators seem to have operated on the spot market to purchase cleaner coal. Moreover, our evidence suggests that firms operating in states where regulators required pre-approval of compliance plans for the new policy seem to have played it safe, and reduced rather than increased the sulfur intensity of their coal.

Our focus on the ARP as a case study for the green paradox is driven by obvious data availability motivations. There are, however, several reasons, both theoretical and empirical, why our results provide general lessons in the green paradox debate. While it is true that the green paradox literature is generally framed within the context of climate change, the relative theoretical analyses do not typically model the process of pollution accumulation. As such, the green paradox literature applies to any type of polluting exhaustible resource and is not restricted to greenhouse gas emissions. Thus, SO_2 policy falls within the relevant range of policies. More importantly, the ARP exhibits strong similarities with the type of climate policy options currently implemented or planned, and can thus be used as a useful acid test for the green paradox hypothesis. First, both climate and SO_2 policies aim to regulate the future consumption of fossil fuels and hence are expected to affect the supply behaviour of coal, oil and gas producers. Second, both types of policies focus on large emitters,

³This is because the combustion of coal, which contains a variable percentage of sulfur, implies the generation of SO_2 as a by-product. Compliance with a cap on SO_2 emissions can be achieved either by reducing the amount of coal burned, the sulfur content of the coal, or by adopting appropriate abatement technologies, e.g. flue gas desulfurization units (scrubbers).

⁴Note that coal-fired power plants consume more than 90% of all coal mined in the U.S. (U.S. EIA, 2011b).

⁵Di Maria et al. (2012) provide a complete theoretical treatment of these effects in the context of a model of exhaustible resource extraction à la Hotelling (1931), extended to allow for differences in pollution intensity across multiple resources.

markedly on electricity generators. Third, the ARP has been the market-based template most of the existing and planned cap and trade schemes have been modelled on (notably the European Union Emissions Trading Scheme, the Regional Greenhouse Gas Initiative, and the Western Climate Initiative). Fourth, the compliance options for the scheme participants are limited, and very similar under both SO₂ and CO₂ regulation: market participants may buy allowances, switch to less polluting fuels⁶, or adopt end-of-pipe abatement technologies (flue-gas desulfurization units in the case of SO₂, carbon capture and sequestration units for CO₂). Finally, the implementation lags that characterized the introduction of the ARP – 5 years for the oldest and dirtiest power plants (Phase I plants), 10 years for all other plants (Phase II) – are of the same order of magnitude as the lags that are relevant in the context of climate policy: the Kyoto Protocol was signed in late 1997, entered into force in 2005, and its first commitment period started in 2008; the EU ETS was first announced in 2001, had a ‘pilot’ phase in 2005-2007, and started in 2008 (Ellerman et al., 2010). The insights we derive below from the implementation of the ARP may therefore be useful to policymakers contemplating future climate policies, especially given the pressure on developing countries to start curbing their emissions from 2020 onward.

The rest of the paper is organized as follows. Section 2 provides an overview of the U.S. electricity sector in the 1980s-1990s, together with a discussion of the evolution of SO₂ regulation in the sector, from the 1970 CAAA to Title IV of the 1990 CAAA. We conclude the section by focussing on the implications of regulatory design for our empirical endeavour. Section 3 presents our empirical models, section 4 our results. Finally, section 5 discusses our results in the context of the green paradox literature, and draws implications for climate policy.

2 Coal-fired electricity generation in the U.S., and the SO₂ trading program

Historically, coal-fired power plants have supplied more than 50% of all electricity used in the United States (U.S. EIA, 2010). Due to the oil price increases of the mid-1970s, coal-fired generation capacity increased throughout the 1980s. Since 1990, however, significant amounts of gas-fired generation have been added, raising the share of gas to around 15% of total generation capacity (U.S. EIA, 2011a), see Figure 1

Throughout the period covered by our analysis, power plants were economically regulated by the state they were located in, generally with a rate of return regulation so that each plant’s output price was set as some fraction above its costs of production. Plants faced the obligation to meet the state’s electricity demand, and thus had less choice over how much and when to produce than would be the case in a liberalized electricity market. This requirement to produce implied that plants were very concerned with assuring a steady supply of fuel. This concern was heightened for coal-fired power plants as they are often base-load plants in an electricity system. Base-load plants are utilized as close as possible to full capacity at all hours of the day, because of their low marginal cost of production and the higher costs to stop and re-start. Nuclear power plants also tend to represent the base-load of the system, while the more flexible natural gas plants tend to be used at peak demand times.

The concern over fuel supply meant that a large majority of coal transactions occurred under long-term forward contracts between plants and coal mines. The contracts were quite complex with many provisions to protect against the ‘hold-up’ problem. Joskow (1985, 1990)

⁶Differences in SO₂ contents for different grades of coal can be large, while Quick (2010) reports differences in CO₂ emission factors across coal types in excess of 10%.

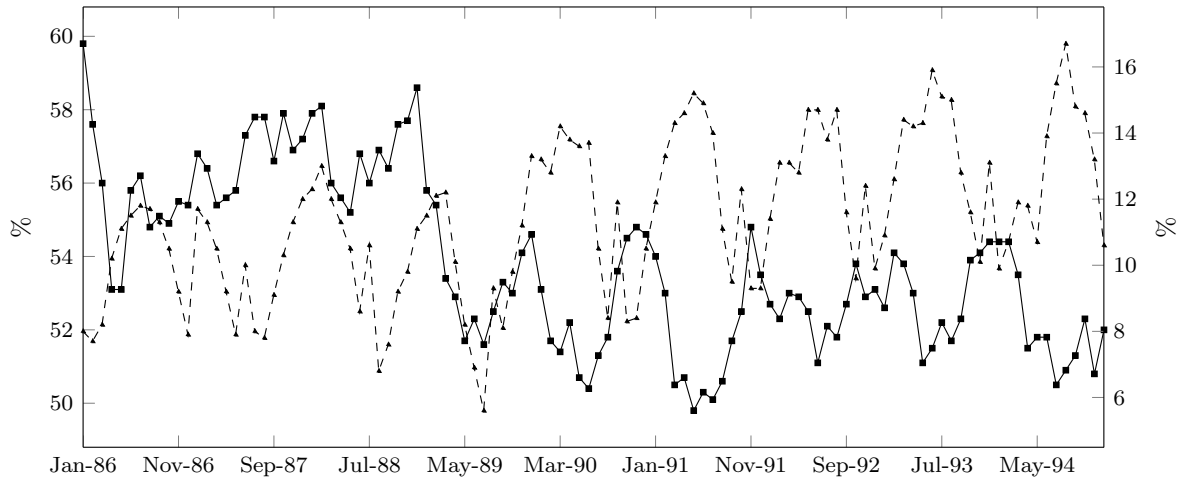


Figure 1: Shares of electricity produced using Coal (solid line, left axis) and Natural Gas (dashed line, right axis) as percentages of total U.S. electricity generation, 1987-1994. Source: U.S. EIA (2010).

has shown that these contracts were largely adhered to even in the face of changes in the spot market coal price, and regulation. The average duration of contracts was about 10 years throughout the 1980s and 1990s, though it was decreasing over time (Lange and Bellas, 2007). The decreased duration was accompanied by an increase in spot market transactions from 10% to 20% of all transactions. The largest increase in spot market activity came from the Western coal region (Kozhevnikova and Lange, 2009). Several commentators have attributed this to the railroad deregulation which began in the mid-1980s: the real prices for shipping coal by rail fell considerably in the late-1980s and 1990s, making Western coal more attractive to utilities in the Mid-West (Ellerman and Montero, 1998; Gerking and Hamilton, 2008). The emergence of a significant spot market in the late 1980s allows us to identify the impact of the introduction of the ARP on coal (spot) market prices and quantities delivered, which would otherwise be very sluggish in the presence of long run contracts. For this reason, we focus our analysis on the period after 1985, when the spot market became deep enough for the price to be considered a meaningful indicator of scarcity (Kozhevnikova and Lange, 2009).

Coal-fired power plants were not only regulated for economic reasons, but also for environmental reasons, as the burning of coal causes the emission of atmospheric pollutants such as SO_2 . U.S. federal regulation of SO_2 emissions from coal-fired plants began with the 1970 CAAA, under which a vintage differentiated emission standard was employed. New plants were subject to the New Source Performance Standards (NSPS), known as NSPS-D, a federal emissions standard of 1.2 pounds of SO_2 per million Btu (lb/mmBtu). The 1977 CAAA tightened restrictions on new plants by expanding the NSPS (known as NSPS-Da) to add a requirement to remove 70-90% of SO_2 post-combustion. The introduction of the NSPS induced owners of power plants to extend the lifetime of existing boilers and resulted in a slower reduction in SO_2 emissions than policymakers had hoped for (Nelson et al., 1993; Stavins, 2006). To fill this gap in the regulation, the Bush Administration introduced provisions to regulate SO_2 emissions via a ‘cap and trade’ program in the summer of 1989. The proposal went through the necessary steps of legislation between 1989 and 1990, being finally signed into law by President G.H.W. Bush on November 15, 1990.

The provisions contained in Title IV of the 1990 CAAA introduced emissions trading in two phases. During Phase I, starting on January 1, 1995, the older, still unregulated boilers were brought under federal regulation. Starting from January 1995, 263 generating units with an

emission rate larger than or equal to 2.5 lb/mmBtu in 1985 were granted emission allowances of about 2.5 lb/mmBtu at baseline 1985-87 fuel use (Ellerman et al., 2000). Each emission allowance would allow its holder to emit one ton of SO₂ in the year of issue or any subsequent year. Phase II, starting on January 1, 2000, covered all units with a capacity of at least 25 MW. Phase II units were to receive allowances at an emission rate of 1.2 lb/mmBtu. The ARP required firms to deliver valid allowances to the U.S. Environmental Protection Agency within thirty days following the end of the calendar year.

The history of SO₂ regulation in the U.S. and the design of the ARP allow us to exploit regulatory differences across U.S. coal-fired power plants to identify the effects of the announcement of the 1990 CAAA on Phase I plants. These plants had been previously unregulated at the federal level, and were generally emitting sulfur dioxide at a much higher rate than other plants. Emissions standards that applied to them at the state level were generous and usually non-binding (Ackerman and Hassler, 1981). NSPS plants, however, were subject to regulations set down in earlier versions of the CAAA. These plants were federally regulated either by an emissions standard or an emissions standard and an implicit technology standard. Given that NSPS plants were subject to binding emissions or technology standards these plants were – contrary to Phase I plants – unable to respond to lower coal prices through an increase in emissions.

In what follows, our aim is not to evaluate the ARP but rather to isolate the effect, if any, of the announcement of the CAAA 1990 on the price of coal, the quantity of coal purchased by power plants and the sulfur intensity of the coal input, ahead of the coming into force of the regulation. In our empirical analysis we argue that, while plants were not assigned randomly into “treatment” and “control” groups, as would be the case in an experimental situation, the existence of two groups of power plants allows for a pseudo-experimental approach. We contend that given the nature of the pre-existing regulation, non-Phase I plants were unlikely to react to changes in the price of coal in search for additional margins (with NSPS-Da plants less likely than NSPS-D plants). Phase I plants, on the contrary, were ideally positioned to benefit from the changes in suppliers’ behaviour, ahead of the implementation of the regulation in 1995. Given that both types of plants operate in otherwise similar economic environments, we can investigate econometrically whether Phase I plants changed their behavior after the announcement of the 1990 CAAA, relative to Phase II plants, and to attribute whatever changes we are able to identify to the announcement of the ARP. More importantly, as will be shown later on while discussing our results, we find strong statistical evidence in support of our approach.

3 Empirical strategy

In this section, we present the empirical models that we employ in the next section to test the three hypotheses discussed in the introduction. The theoretical predictions of the green paradox literature imply that, following the announcement of the ARP, we should observe (1) a drop in the (spot) price of coal, as producers bring extraction forward in time; (2) an increase in the total amount of the resource consumed, due to the reduction in price; and (3) an increase in the use of the resource with the higher pollution content.

In order to verify whether coal prices fell after the announcement of the ARP, we utilize a hedonic price regression similar to those that have been used in past literature (see e.g. Keohane and Busse, 2007; Lange and Bellas, 2007). The model is:

$$p_{j,t} = \alpha_{0,j} + \alpha_{1,n} + \alpha_2 \text{Sulfur}_{j,t} + \alpha_3 \text{Sulfur}_{j,t} \times \text{Interim}_t + \mathbf{x}'_{j,t} \boldsymbol{\alpha} + \varepsilon_{j,t}. \quad (1)$$

The dependent variable is the weighted average real price (per million Btu) of coal delivered to plant j in month t , for deliveries agreed upon on the spot market. As Joskow (1988, 1990) discusses, contract prices do not respond to market conditions as quickly as spot prices due to the price adjustment mechanism in the contract. As a result, this analysis is undertaken for spot market transactions. Both plant- and year-fixed effects are accounted for, and are represented by $\alpha_{0,j}$ and $\alpha_{1,n}$, respectively. The year fixed effects will determine whether prices fell in conjunction with the announcement of the ARP. Sulfur $_{j,t}$ is the weighted average sulfur content of deliveries to plant j in month t and Interim $_t$ is a dummy that has a value equal to one in the period December 1990 - December 1994 and zero otherwise.⁷ Coefficient α_3 will reveal whether the discount for high sulfur coal increased after the passage of the 1990 CAAA. The vector $\mathbf{x}_{j,t}$ indicates a vector of control variables, consisting of the heat and ash content of the coal delivered, variables that control for the region of origin of the coal, a proxy for transportation costs, the coal region mining productivity and the natural gas price. Finally, ε is an IID error term.

We estimate equation (1) using data for the period 1986-1994. Appendix A describes the data in more detail, and provides both the sources of the data and the summary statistics for all variables.

Our second hypothesis states that mandatory Phase I plants should increase the amount of heat consumed after announcement of the ARP. To distinguish between changes that are due to the announcement of the 1990 CAAA and general trends in the industry, a difference-in-difference methodology is utilized. Plants that contain at least one boiler that was mandated to be part of Phase I are the treatment group with different subsets of non-Phase I plants as control groups. First, a model which shows how the pre-and post-announcement trends evolve is estimated. The model is:

$$h_{j,t} = \beta_{0,j} + \beta_{1,n}\delta_n + \beta_{2,n}\delta_n \times \text{PhaseI}_j + \mathbf{x}'_{j,t}\boldsymbol{\beta} + \eta_{j,t}. \quad (2)$$

The dependent variable is the natural log of billion Btu purchased by plant j in month t . Both plant- and year-fixed effects are accounted for, and are represented by $\beta_{0,j}$ and $\beta_{1,n}$, respectively. To test whether pre-announcement trends are equivalent across the two groups, year dummies are interacted with a dummy that is equal to one for mandatory Phase I plants and zero otherwise: $\delta_n \times \text{PhaseI}_j$. If the coefficients of the pre-announcement interaction terms, $\beta_{2,n}$, are not statistically different from zero, this would support the assumption that trends in the heat consumption over time are the same for the treatment and control plants. If the coefficients of the *post*-announcement interaction terms are positive and statistically significant, this would provide support in favour of the green paradox hypothesis, as it would imply that mandatory Phase I plants increased their heat consumption relative to the control group's consumption. The vector $\mathbf{x}_{j,t}$ indicates a vector of control variables, consisting of an index for state-level economic activity, a dummy equal to one for plants that have a scrubber installed, summer and winter dummies, and the real natural gas price. Finally, $\eta_{j,t}$ is an IID error term.

While the difference-in-difference estimation in equation (2) will help reveal the validity of the empirical model, other factors may alter how mandatory Phase I plants respond to the

⁷Various dates can be picked as the announcement date. The announcement of the clean air proposal in the summer of 1989 and the signing into law of the CAAA in November 1990 appear to be most relevant. In addition, the question is how fast spot prices were able to respond to these announcements. To be general, we therefore use year (interaction) dummies in our core regressions throughout the paper, where it can be checked whether prices or quantities changed in 1989 or 1990. Our 'Interim period' dummy variable is equal to one for months after November 1990.

1990 CAAA passage relative to the control group. To explore whether regulatory or market factors are confounding the analysis, a triple difference-in-difference model is estimated. The triple difference model is:

$$h_{j,t} = \gamma_{0,j} + \gamma_{1,n} + \gamma_2 \text{Interim}_t + \gamma_3 \text{Interim}_t \times \text{PhaseI}_j \\ + \gamma_4 \text{Interim}_t \times \text{PhaseI}_j \times \text{Triple}_j + \mathbf{x}'_{j,t} \boldsymbol{\gamma} + \nu_{j,t}. \quad (3)$$

The interaction of the ‘Phase I’ and ‘Interim’ dummy variables is the first difference-in-differences variable: if γ_3 is positive and significant, Phase I plants purchased more heat after announcement, relative to non-Phase I plants. We focus on three possible factors that may have affected the ability of Phase I plants to take advantage of the cheaper coal. In the first place we concentrate our attention on plants that operate in states where no nuclear plants are operating. The presence of nuclear stations would make it more difficult for coal-fired generators to expand their production since nuclear generation is typically cheaper and has higher start-up costs. Thus, when ‘Triple’ is ‘No-nuclear’, we would expect a positive estimate for the γ_4 coefficient. The second aspect we are interested in here is to what extent Phase I plants were constrained by existing long-term contracts on the coal market. Our second ‘Triple’ variable is then ‘High Spot’, a dummy that identifies plants that purchase a large share of their coal on the spot market, rather than rely on long-term contracts. We would expect that plants with a higher degree of exposure to the spot market would be able to buy more coal at lower prices during the interim phase, implying a positive estimate for γ_4 . Finally, we consider whether the fact that some plants were required by their state-level Public Utility Commission (PUC) to submit a compliance plan before the start of Phase I might have limited their ability to increase production. Indeed, if plants had to undergo additional maintenance or install new machinery, for example, to ensure that their plan would function as expected, we might expect a negative estimate for γ_4 when ‘Triple’ is ‘Pre-approval’.

We test our third hypothesis, which states that the sulfur content of coal used increased after announcement, using a difference-in-difference analysis, similar to equation (2). The model, which has as dependent variable the sulfur intensity (in lb/mmBtu) of the coal delivered to the power-plants, is:

$$i_{j,t} = \mu_{0,j} + \mu_{1,n} \delta_n + \mu_{2,n} \delta_n \times \text{PhaseI}_j + \mathbf{x}'_{j,t} \boldsymbol{\mu} + \zeta_{j,t}. \quad (4)$$

Again this will show how the pre-and post-announcement trends evolve for Phase I plants and with different subsets of non-Phase I plants as control groups. Our set of control variables for the sulfur content of coal include a scrubber dummy, rail transportation costs per ton-mile interacted with a dummy equal to one for plants located in a relevant range from the Powder River coal Basin (PRB), and controls for the region of origin of the coal. The region of origin is relevant for sulfur content as coal from the PRB has the lowest sulfur content of the three main regions, while coal from the Interior region has the highest sulfur content and the Appalachian region is in between the other two. In addition, Ellerman and Montero (1998) show that the declining rail transport prices in the 1980s and 1990s changed the economics of coal choice in favor of the coal from the PRB for power plants in a range of 400-1200 miles from the PRB.

The final model is a triple difference-in-difference model similar to equation (3), where ‘Triple’ controls for whether the state the Phase I plant is located in has strict environmental standards, or whether the plant was required to submit a compliance plan before the start of the ARP. Our priors for both controls is that plants facing stringent state regulation or that are required to obtain pre-approval of their compliance plans would find it difficult to switch to dirtier coal (Lile and Burtraw, 1998). As a consequence, we would expect a negative coefficient for the triple difference-in-difference coefficients.

Table 1: Hedonic coal price regressions

Dependent Variable:	Weighted Average Real Price ^(a)	
	(1) Spot Transactions	(2) Spot Transactions for Plants in States with Phase I Plants
Sample:	Coefficient (S.E)	Coefficient (S.E)
1987	-28.26*** (1.28)	-27.80*** (1.31)
1988	-11.55*** (1.87)	-11.39*** (1.94)
1989	-30.17*** (2.62)	-30.51*** (2.55)
1990	-40.31*** (3.39)	-39.17*** (3.77)
1991	-35.36*** (3.42)	-34.49*** (4.11)
1992	-37.49*** (3.52)	-37.68*** (4.24)
1993	-27.13*** (3.69)	-26.51*** (4.27)
1994	-16.87*** (3.85)	-17.13*** (4.32)
Sulfur Content	-3.57*** (0.93)	-3.50*** (0.94)
Sulfur Content × Interim period	-1.43*** (0.52)	-1.36** (0.60)
Observations	19,863	15,799
Plants	367	276

Additional controls for all regressions are: Coal Region Share, Real Transport Costs, Real Natural Gas Price, Ash Content, Heat Content, Year and Plant Dummies. Standard errors corrected for panel serial correlation.

*, **, *** indicate 10%, 5% and 1% statistical significance, respectively.

Time Period is 1986-1994.

(a). Average delivery price, weighted by heat-content for each plant, month.

4 Results

Table 1 shows the results of our estimations of equation (1), which refers to the hypothesis that coal prices fell after announcement of the 1990 CAAA. In the interest of brevity, we suppress the results for the coefficients for additional control variables.⁸ In column (1) we present results for all spot transactions, in column (2) we focus on transactions on the spot market by plants located in states with Phase I plants. Both regressions reveal that coal prices were lower during the period 1990-1992 by about 14 cents per mmBtu, accounting for a drop of 9% relative to the price during the previous three years (1987-1989).⁹ This suggests that the announcement of the 1990 CAAA might indeed have had an impact such as the one

⁸The full output for these regressions, and for all the others presented in the paper is available from the authors upon request.

⁹The t -test for the difference between the average of the year-fixed effects for 1987-89 and the average for 1990-92 shows a statistically significant price difference of 14 cents per mmBtu after announcement (t -stat=6.81).

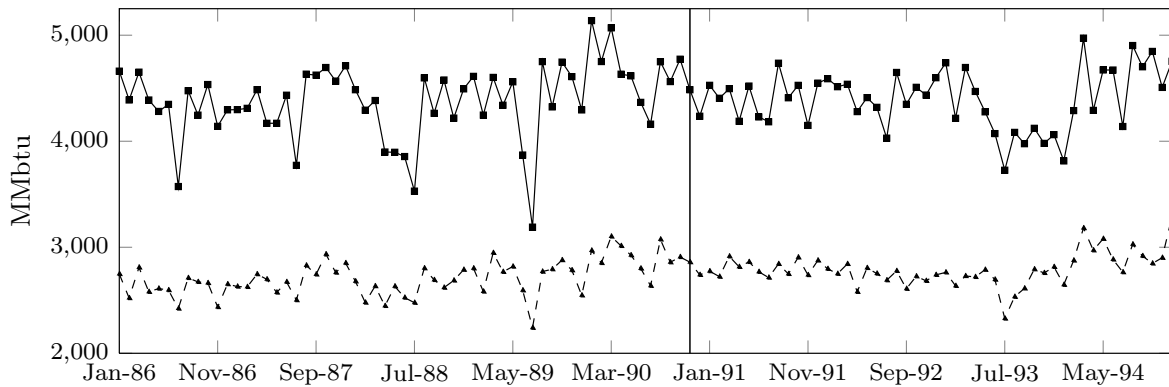


Figure 2: Average heat input per month for Phase I plants (solid line) and for non Phase I plants in States with Phase I plants (dashed line). Source: Authors calculations using FERC 423 data.

suggested by the green paradox hypothesis.¹⁰ The price recovery in 1993-1994 revealed by our regressions might instead be linked to supply disruptions, following the 7-month United Mine Workers' strike in the second half of 1993. As expected, high-sulfur coal is traded at lower prices than low-sulfur coal. The coefficient for the interaction term of sulfur content and a dummy for the period after announcement shows that the sulfur premium increased significantly (by about 40%), after it became clear that high-sulfur coal would become harder to sell in the future. Overall, this evidence supports our first hypothesis, namely that the announcement of the cap on SO_2 emissions appears to have had a depressing effect on coal prices, and a stronger one for high-sulfur coal.

Next, we ask whether these lower coal prices in the interim of the ARP affected the amount of coal burned by U.S. power plants.¹¹ Specifically, we investigate whether coal-fired generators increased their heat input after the announcement of the 1990 CAAA. Table 2 shows the results of a panel estimation of equation (2), focusing on the pre- and post-announcement trends. Column (1) uses the whole sample of non-Phase I power plants as the control group. Our estimates reveal that using this control group is not appropriate, as its pre-announcement trends are different from those of mandatory Phase I plants (the treatment group). This outcome is not surprising as a number of papers have shown that spatial variation is important in the coal market. Almost all mandatory Phase I plants are located east of the Mississippi River, while the low-sulfur coal mining region PRB is located west of the Mississippi. Keohane and Busse (2007) show that railroad firms price discriminated based on plants proximity to the PRB. Joskow (1987) finds that the contract duration is dependent on which coal region the mine is located in (Appalachian, Interior, or Western). As discussed above, Ellerman and Montero (1998), Gerking and Hamilton (2008) and Kozhevnikova and Lange (2009) all find differences in plants decision-making by their geographic location.

In order to provide a better control group, the other two columns restrict the non-Phase I plants to those in states with mandatory Phase I plants. Column (2) shows the results of our second regression model, when the control group is represented by non-Phase I plants operating in states that have at least one Phase I plant. Besides the economic rationale

¹⁰These results correlate to the finding of Kahn and Knittel (2003) that the unveiling of the proposed amendments to the Clean Air Act by President Bush in June of 1989 had a negative impact on coal mining company's stock prices.

¹¹To control for possible State-level serial correlation, we follow Bertrand et al. (2004) and compute standard errors clustering by State.

Table 2: Heat input hypothesis

Dependent Variable:	Natural Log of Total Heat Purchased		
Sample:	All Transactions	All Transactions for Plants in States with Phase I Plants	All Transactions for Phase I Plants or NSPS-Da Plants in States with Phase I Plants
	(1)	(2)	(3)
	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)
1987 × Phase I	-0.02 (0.03)	-0.02 (0.03)	-0.02 (0.04)
1988 × Phase I	-0.07** (0.04)	-0.05 (0.04)	-0.03 (0.05)
1989 × Phase I	-0.09** (0.05)	-0.04 (0.06)	-0.03 (0.07)
1990 × Phase I	-0.03 (0.04)	0.01 (0.05)	0.04 (0.06)
1991 × Phase I	-0.06 (0.05)	-0.01 (0.06)	0.03 (0.06)
1992 × Phase I	-0.05 (0.05)	0.03 (0.06)	0.06 (0.07)
1993 × Phase I	-0.17*** (0.06)	-0.07 (0.06)	-0.05 (0.06)
1994 × Phase I	-0.16*** (0.06)	-0.09 (0.06)	-0.05 (0.07)
Observations	38,315	27,190	22,314
Plants	401	288	237

Controls for all regressions are: State GDP, Scrubber Dummy, Real Natural Gas Price, Summer, Winter, Year and Plant Dummies.

Standard errors corrected for state-level serial correlation.

*, **, *** indicate 10%, 5% and 1% statistical significance, respectively.

Time period is 1986-1994.

previously discussed, the regression results suggest that the choice for this control group is consistent with a proper research design. Indeed, the pre-announcement trend is not statistically different from that of the mandatory Phase I plants. Obviously, the use of these plants as our control group is only justifiable if we can be sure that the announcement of the ARP did not affect non-Phase I plants the way it might have affected Phase I plants. Figure 2 shows the average heat input over time for both mandatory Phase I plants and non-Phase I plants in states with Phase I plants. The average heat input for the control group doesn't exhibit any break following the announcement of the change in legislation. This visual impression is confirmed by a statistical test on the equality of the coefficients of the year dummies over the 1989-1991 period.¹²

The post-announcement interaction variables in column (2) are all statistically insignificant,

¹²The F -test statistic for the joint hypothesis test equals 1.46, with a p -value of 0.26. We also perform a test for the stability of the slope parameters in the period before and after the policy announcement, and we fail to reject the null that the parameters are indeed stable. The complete results are available from the authors upon request.

suggesting that Phase I plants did not increase their heat input relative to the control group. Thus, despite the fall in coal prices documented above, our analysis so far fails to identify any quantity response to the announcement of the ARP.

As a robustness check, column (3) of Table 2 shows the results of the second regression model when the control group is further restricted to only include NSPS-Da plants in states with Phase I plants. NSPS-Da plants are essentially required to have a scrubber in addition to an emissions standard, thus the announcement of the ARP doesn't increase the stringency of the regulation for these plants. The statistically insignificant difference in the pre-announcement trends suggests that such plants also constitute an appropriate control group. Also in this case, the post-announcement interaction variables are all statistically insignificant, providing further evidence that the treated group did not behave differently from the control group.

From our discussion so far, it appears that the announcement of the ARP did not have any impact on heat input. It is still possible, however, that specific factors may restrict power plants' ability to change their behaviour in response to the announcement of the ARP. As mentioned in section 3, we focus on three such factors here. We start by considering the possibility that plants facing the competition of nuclear power stations in providing cheap energy to cover baseload demand might be less able to expand output, despite the availability of cheaper coal. Similarly, plants that obtain a large share of their coal based on pre-existing long-term contracts are unlikely to be able to benefit from cheaper coal on the spot market. Finally, it is plausible that plants that operate in states whose public utility commissions require the pre-approval of detailed compliance plans, would tend to focus on achieving compliance on their own rather than relying on purchasing permits from the market (see e.g. Rose, 1997). This would make them more reluctant to expand output, and switch to more polluting coal and instead would lead them to behave more conservatively. Table 3 shows the results of our triple difference-in-difference analyses.

The results of Table 3 reveal some interesting findings. First of all, mandatory Phase I plants that use the spot market for a large share of their coal purchases are shown to have increased their heat input in the period between the signing into law of the 1990 CAAA and the start of the ARP in January 1995.¹³ These results support the existence of a green paradox for this particular subgroup of mandatory Phase I plants. For high-spot plants, the triple difference-in-difference coefficient is positive and statistically significant at the 5% level for both control groups. We fail to find any differential effect of the announcement in terms of heat input for plants operating in states without nuclear power stations, and in states that required the pre-approval of firms' compliance plans.

Taken together, the results from Tables 2-3 suggest, quite intuitively, that the likelihood of the emergence of a green paradox is reduced when plants rely more heavily on long-term contracts for their fuel supply. Indeed, plants that make only a minimal use of the spot market would naturally be much less likely to be able to benefit from favourable fluctuations in the coal price.

The third testable implication we derive from the green paradox literature refers to the possibility that, as the sulfur premium increases after the announcement of the ARP, more of the – now relatively cheaper – high-sulfur coal would be used by plants, in the interim

¹³The 'High-Spot' dummy used in the regressions in Table 3 assumes the value 1 for all years in which a plant received more than 40% of its coal deliveries from spot market contracts. To check the robustness of our results, we first lowered the threshold to 30%, and then increased it to 50%. Moreover, we focussed on the share of deliveries over the entire period, using the same threshold levels. Finally, we experimented designating 'High-Spot' plants only those whose spot share was higher than 40% (30%, 50%, respectively) in every year of our sample. In all cases the results were qualitatively very similar.

Table 3: Heat input hypothesis, additional results

Dependent Variable:	Natural Log of Total Heat Purchased					
	(1) All Transactions for Plants in States with Phase I Plants	(2) All Transactions for Phase I or NSPS-Da Plants for States with Phase I Plants	(3) All Transactions for Plants in States with Phase I Plants	(4) All Transactions for Phase I or NSPS-Da Plants for States with Phase I Plants	(5) All Transactions for Plants in States with Phase I Plants	(6) All Transactions for Plants in States with Phase I Plants
Sample:	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)
Interim Period	-0.03 (0.03)	-0.03 (0.04)	-0.02 (0.03)	-0.02 (0.04)	-0.03 (0.04)	-0.05 (0.05)
Interim \times Phase I	-0.04 (0.03)	-0.03 (0.03)	-0.05 (0.04)	-0.04 (0.04)	-0.03 (0.04)	0.01 (0.04)
Interim \times Phase I \times No-Nuclear	0.09 (0.12)	0.17 (0.14)				
Interim \times Phase I \times High-Spot			0.13*** (0.04)	0.16*** (0.05)		
Interim \times Phase I \times Compl. Pre-approval					0.04 (0.07)	0.02 (0.05)
Observations	27,190	22,314	27,190	22,314	27,190	22,314
Plants	288	237	288	237	288	237

Controls for all regressions are: State GDP, Scrubber Dummy, Real Natural Gas Price, Summer, Winter, Year and Plant Dummies
Standard errors corrected for state-level serial correlation.

*, **, *** indicate 10%, 5% and 1% statistical significance, respectively.
Time Period is 1986-1994.

Table 4: Intensity hypothesis

Dependent Variable:	Sulfur Intensity			
	(1)	(2)	(3)	(4)
Sample:	Spot Transactions	Spot Transactions for Plants in States with Phase I Plants	Spot Transactions for Phase I or NSPS-Da Plants for States with Phase I Plants	All Transactions
	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)
1987 × Phase I	-0.07 (0.08)	-0.10 (0.08)	-0.09 (0.09)	-0.07 (0.06)
1988 × Phase I	-0.03 (0.11)	-0.05 (0.12)	-0.4 (0.12)	-0.11 (0.07)
1989 × Phase I	0.06 (0.10)	0.03 (0.11)	0.05 (0.12)	-0.09 (0.07)
1990 × Phase I	-0.02 (0.12)	-0.05 (0.13)	-0.07 (0.12)	-0.06 (0.10)
1991 × Phase I	-0.16* (0.09)	-0.15* (0.08)	-0.17* (0.09)	-0.13 (0.11)
1992 × Phase I	-0.18** (0.07)	-0.19** (0.08)	-0.16* (0.08)	-0.10 (0.12)
1993 × Phase I	-0.17** (0.08)	-0.17* (0.09)	-0.10 (0.09)	-0.23** (0.10)
1994 × Phase I	-0.22** (0.11)	-0.21* (0.11)	-0.13 (0.11)	-0.33*** (0.12)
Observations	20,563	16,396	13,804	40,766
Plants	369	277	228	409

Controls for all regressions are: Coal Region Share, Scrubber Dummy; Rail Price, near Powder River Basin, Year and Plant Dummies.

Standard errors corrected for state-level serial correlation.

*, **, *** indicate 10%, 5% and 1% statistical significance, respectively.

Time Period is 1986-1994.

of the regulation. Table 4 presents the results of the estimation of equation (4), testing the hypothesis that the sulfur content of coal did in fact increase after announcement of the future cap on SO₂ emissions. As argued already in section 2, non-Phase I plants already faced multiple regulations in terms of their SO₂ emissions, thus it is difficult to imagine that they would react to the increase in the sulfur premium by switching to dirtier coal and increase their emissions intensity. Our priors that non-Phase I plants constitute a good control are supported by the results in Table 4, which show that the pre-treatment trends are not significantly different across the treated and the control group. Moreover, Figure 3 illustrates that the trend in emissions intensity of the control group didn't change as a result of the announcement of the ARP.

For all control groups utilized, the pre-announcement trends do not differ significantly from zero, supporting the assumption that pre-announcement trends are equivalent between the two groups. The first three columns of Table 4 focus only on deliveries from spot market contracts, since we expect it to be more likely to observe changes in sulfur content for these transactions, relative to deliveries based on long-term contracts. Indeed, we are able to identify changes in the spot data long before any change emerges from the full transaction sample. The changes we identify, however, go in the opposite direction to that hypothesized

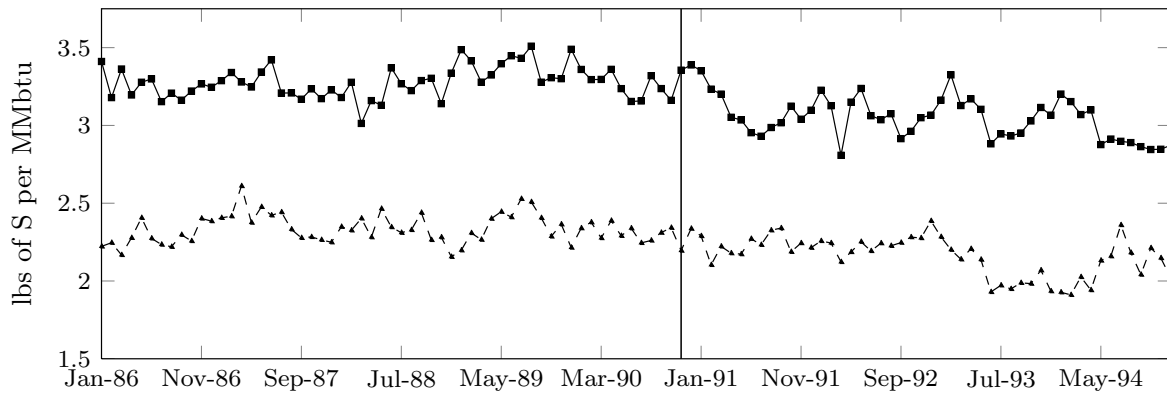


Figure 3: Average monthly sulfur intensity for Phase I plants (solid line) and for non Phase I plants in States with Phase I plants (dashed line). Source: Authors calculations using FERC 423 data.

by proponents of the green paradox theory. The negative and significant coefficients we estimate for the regressions in columns (1)-(3) of Table 4, seem rather to suggest that Phase I plants were actively gearing up to comply with the regulation as early as 1991, possibly experimenting with different types of fuel. Using data on all transactions – see column (4) – reveals that overall mandatory Phase I plants decreased their sulfur intensity relative to non-Phase I plants starting from 1993. This is line with the evidence suggesting that several Phase I plants aggressively renegotiated their long-term contracts in anticipation of the commencement of the ARP, in order to reduce the sulfur-content of the coal they would receive in the future (Kosnik and Lange, 2011).

Also in this case, we resort to a triple difference-in-difference method to control for potential differences amongst mandatory Phase I plants. Once again, we focus on possible differential behaviour by plants that were required to submit compliance plants for pre-approval to their PUC. Columns (1) and (2) of Table 5 show the results of this analysis. Neither the regression using all plants in States with Phase I plants, nor the one using as control the sample restricted to only NSPS-Da plants show any increase in the sulfur intensity for mandatory Phase I plants. As expected, however, the compliance plan requirement seems to have induced reductions in the sulfur intensity of coal inputs by mandatory Phase I plants, already in the interim phase. This is consistent with the view that utilities subject to such scrutiny would tread more cautiously and make sure of compliance well ahead of time. Our last result is that, once we isolate plants operating in states with stringent SO_2 regulation, other Phase I plants can be seen to reduce their sulfur intensity. Once more, this is consistent with early action on the part of previously unregulated Phase I plants to prepare to comply with the requirements of the ARP.

5 Concluding remarks

A recent surge in the literature on suboptimal climate policy has provided several conditions under which such a policy in theory could lead to detrimental environmental outcomes (a ‘green paradox’). In this paper, we have taken the first step to empirically assess this theory.

One of the mechanisms through which environmental policy could lead to an increase in harmful emissions is through the response of suppliers of nonrenewable resources in the presence of an implementation lag for the policy. We have argued that the implementation lag

Table 5: Intensity hypothesis: additional results

Dependent Variable:	Sulfur Intensity			
	(1)	(2)	(3)	(4)
Sample:	Spot Transactions for Plants in States with Phase I Plants	Spot Transactions for Phase I or NSPS-Da Plants for States with Phase I Plants	Spot Transactions for Plants in States with Phase I Plants	Spot Transactions for Phase I or NSPS-Da Plants for States with Phase I Plants
	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)	Coefficient (S.E)
Interim Period	0.05 (0.06)	0.05 (0.06)	0.08 (0.06)	0.04 (0.06)
Interim \times Phase I	-0.09 (0.06)	-0.07 (0.07)	-0.16** (0.07)	-0.11 (0.07)
Interim \times Phase I \times Pre-approval	-0.26*** (0.07)	-0.15** (0.07)		
Interim \times Phase I \times Strict State SO ₂			0.09 (0.10)	0.02 (0.10)
Observations	16,396	13,804	16,396	13,804
Plants	277	228	277	228

Controls for all regressions are: Coal Region Share, Scrubber Dummy; Rail Price, near Powder River Basin, Year and Plant Dummies.

Standard errors corrected for state-level serial correlation.

*, **, *** indicate 10%, 5% and 1% statistical significance, respectively.

Time Period is 1986-1994.

that characterized the SO₂ cap and trade program under the 1990 CAAA has strong parallels with implementation lags in climate policy. We have presented three testable implications regarding the effects of the announcement of the cap. The first hypothesis states that the price of coal would fall after the announcement, as resource owners see the prospects of future sales decline; moreover, this effect should be stronger for poorer quality (in terms of sulfur content) grades of coal. Our second hypothesis is that, in response to this fall in price, coal use by utilities would increase. Our third and final hypothesis is that given the expected increase in the sulfur premium for coal, utilities would find it advantageous to temporarily shift to higher-sulfur coal. Thus, we should be able to observe an increase in the sulfur content of coal delivered after the signing into law of the ARP.

We use data on coal deliveries to U.S. coal-fired power plants to test these hypotheses. We find strong evidence for the first hypothesis. The average monthly price of coal delivered on the spot market dropped by 9% after the announcement of the ARP, and the sulfur premium increased by roughly 40%. However, we are able to identify only a weak effect in terms of our second hypothesis: only coal-fired power plants that were sufficiently flexible on the coal market seem to have increased their heat input in response to the lower price of coal. It stands to reason that plants that rely heavily on long-term contracts for their coal supply would have neither the possibility nor the inclination to purchase cheaper coal via spot-market operations. Furthermore, we find no evidence in support of the hypothesis that the sulfur content of coal would increase. If anything, we are able to show that coal deliveries from spot market contracts have exhibited a marked shift towards cleaner coal as early as 1991. We also find that firms operating in states requiring the pre-submission of compliance plans for the ARP appear to have reduced the sulfur intensity of their coal already in the interim

phase. Both results are likely the result of Phase I plants preparing for the implementation of the ARP ahead of time, and of their relying on self-compliance, rather than on trading.

In conclusion, going back to our title question, should we be worried about the Green Paradox? Although the theoretical literature has found several ways in which suboptimal environmental policies in the presence of nonrenewable resources may induce an increase in emissions, we find mixed evidence in the case of the ARP announcement. We do find a drop in coal prices, in line with the predictions of the green paradox hypothesis, yet it seems that only a small subset of regulated coal-fired power plants have responded to the price signal by increasing their coal use. Moreover, as far as our evidence goes, there seems to have been no shift towards dirtier coal, despite the increase in the sulfur premium. Rather, we find evidence of a moderate shift towards cleaner fuel. Thus, we find no substantive evidence that emissions were higher in the interim of the regulation than they would otherwise have been. The green paradox seems not to have materialized. On the face of this evidence, it seems that the answer to our title's question is a rather resounding "no". We need to consider, however, that important restrictions limited U.S. coal-fired utilities' opportunities to respond to lower coal prices. First of all, the implementation lag might simply have been too short for firms to react. It makes very little economic sense to adjust the productive capacity of a power plant in order to exploit a business advantage spanning only four years. This is even more compelling for the Phase I plants that are the object of our analysis, which are older plants already approaching retirement. Moreover, most of the plants in our sample are large base-load plants, producing at full capacity almost all the time. Their technical ability to produce more might have been severely constrained. Another aspect worth noting here is the pervasiveness of long-term contracts and the limited scale of the spot market, both of which would further limit the possibility to adjust coal quantities, in response to changes in prices. Moreover, the U.S. coal generating sector is subject to a plethora of competing regulations, and local, state and federal rules overlap and interact with each other in a highly complex way (see Rose, 1997). Lile and Burtraw (1998), for example, provide an interesting compendium on how state rules might have biased compliance options for Phase I plants towards capital intensive ones, and promoted self compliance over trading. They also discuss the role of pre-approval of compliance plants, which would favour early compliance over green-paradoxical increases in emissions. Thus, although our results regarding the existence of a green paradox are mixed, they do suggest that fossil fuel prices react to policy announcements as predicted by the theory. Our analysis underlines the importance for the emergence of the green paradox of factors such as the length of the implementation lag, the preexistent and concurrent regulation – environmental or otherwise – of fossil fuel users, and the elasticity of demand for fossil energy. All play a role in determining the way in which fossil fuel users are likely to respond to the drop in input prices.

In terms of the implications of our research for climate policy, we believe that there might still be reasons to be worried. In the presence of unilateral policies by a limited number of countries, the largest emitters on the planet do not appear likely to face stringent climate policies for many years, the demand for energy is increasing at alarming rates and fossil fuels are likely to be highly sought after for the foreseeable future. In these circumstances, fossil fuel users would be in a much better position to exploit a reduction in prices, making the emergence of a green paradox much more than a scientific curiozum.

References

- Ackerman, B., and W. Hassler (1981) *Clean Coal/Dirty Air* (New Haven, CT: Yale University Press)
- Bertrand, M., E. Dufflo, and S. Mullainathan (2004) ‘How much should we trust differences-in-differences estimates?’ *The Quarterly Journal of Economics* 119(1), 249–275
- Di Maria, C., S. Smulders, and E. Van der Werf (2012) ‘Absolute abundance and relative scarcity: environmental policy with implementation lags.’ *Ecological Economics* 74, 104–119
- Eichner, T., and R. Pethig (2011) ‘Carbon leakage, the green paradox and perfect future markets.’ *International Economic Review* 52(3), 767–805
- Ellerman, A. D., and J.-P. Montero (1998) ‘The declining trend in sulfur dioxide emissions: Implications for allowance prices.’ *Journal of Environmental Economics and Management* 36(1), 26–45
- Ellerman, A. D., F. J. Convery, and C. de Perthuis (2010) *Pricing Carbon: the European Union Emissions Trading Scheme* (Cambridge, MA: Cambridge University Press)
- Ellerman, A. D., P. L. Joskow, R. Schmalensee, J.-P. Montero, and E. M. Bailey (2000) *Markets for Clean Air: the U.S. Acid Rain Program* (New York: Cambridge University Press)
- Fischer, C., and S. Salant (2012) ‘Climate policies and intertemporal emissions leakage: quantifying the green paradox.’ Resources for the Future Discussion Paper 12-16
- Gerking, S., and S. F. Hamilton (2008) ‘What explains the increased utilization of Powder River Basin coal in electric power generation?’ *American Journal of Agricultural Economics* 90(4), 933–950
- Gerlagh, R. (2011) ‘Too much oil.’ *CEsifo Economic Studies* 57(1), 79–102
- Grafton, R. Q., T. Kompas, and N. V. Long (2011) ‘Substitution between biofuels and fossil fuels: is there a Green Paradox?’ mimeo
- Hoel, M. (2012) ‘The supply side of CO₂ with country heterogeneity.’ *Scandinavian Journal of Economics* 113(4), 846–865
- Hotelling, H. (1931) ‘The economics of exhaustible resources.’ *Journal of Political Economy* 39(2), 137–175
- Joskow, P. (1985) ‘Vertical integration and long-term contracts: The case of coal-burning electric generating plants.’ *Journal of Law, Economics, & Organization* 1(1), 33–80
- (1987) ‘Contract duration and relationship specific investments: Empirical evidence from coal markets.’ *American Economic Review* 77(1), 168–185
- (1988) ‘Price adjustment in long-term contracts: The case of coal.’ *Journal of Law & Economics* 31(1), 47–83
- (1990) ‘The performance of long-term contracts: Further evidence from coal markets.’ *RAND Journal of Economics* 21(2), 251–274

- Kahn, S., and C. R. Knittel (2003) 'The impact of the Clean Air Act Amendments of 1990 on electric utilities and coal mines: evidence from the stock market.' Center for the Study of Energy Markets Working Paper 118
- Keohane, N., and M. Busse (2007) 'Market effects of environmental regulation: Coal, railroads, and the 1990 Clean Air Act.' *RAND Journal of Economics* 38(4), 1159–1179
- Kosnik, L., and I. Lange (2011) 'Contract renegotiation and rent re-distribution: Who gets raked over the coals?' *Journal of Environmental Economics and Management* 62(2), 155–165
- Kozhevnikova, M., and I. Lange (2009) 'Determinants of contract duration: Further evidence from coal-fired power plants.' *Review of Industrial Organization* 34(3), 217–229
- Lange, I., and A. Bellas (2007) 'The 1990 Clean Air Act and the implicit price of sulfur in coal.' *The B.E. Journal of Economic Analysis & Policy*. Article 41
- Lile, R., and D. Burtraw (1998) 'State-level policies and regulatory guidance for compliance in the early years of the SO₂ emission allowance trading program.' Discussion Paper 98-35, RFF
- Nelson, R. A., T. Tietenberg, and M. R. Donihue (1993) 'Differential environmental regulation: effects on electric utility capital turnover and emissions.' *Review of Economics and Statistics* 75(2), 368–373
- Quick, J. C. (2010) 'Carbon dioxide emission factors for U.S. coal by origin and destination.' *Environmental Science & Technology* 44(7), 2709–2714
- Rose, K. (1997) 'Implementing an emissions trading program in an economically regulated industry: Lessons from the SO₂ trading program.' In *Market Based Approaches to Environmental Policy: Regulatory Innovations at the Fore*, ed. R. F. Kosobud and J. Zimmerman (New York: Van Nostrand Reinhold)
- Sinn, H.-W. (2008) 'Public policies against global warming.' *International Tax and Public Finance* 15(4), 360–394
- (2012) *The Green Paradox: A Supply-Side Approach to Global Warming* (Cambridge, MA: MIT Press)
- Smulders, S., Y. Tsur, and A. Zemel (2012) 'Announcing climate policy: Can the green paradox arise without scarcity?' *Journal of Environmental Economics and Management*, forthcoming
- Stavins, R. N. (2006) 'Vintage-differentiated environmental regulation.' *Stanford Environmental Law Journal* 25(1), 29–63
- U.S. EIA (2004) 'Coal transportation: Rates and trends.' United States Energy Information Administration
- (2005) 'Annual energy review.' United States Energy Information Administration
- (2010) 'Annual energy review.' United States Energy Information Administration
- (2011a) 'Age of electric power generators varies widely.' *Today in Energy*. June 16

- (2011b) ‘Annual coal distribution report.’ United States Energy Information Administration
 - (2011c) ‘Electric power annual 2010.’ United States Energy Information Administration
- Van der Ploeg, F., and C. Withagen (2012) ‘Is there really a Green Paradox?’ *Journal of Environmental Economics and Management*, forthcoming
- Van der Werf, E., and C. Di Maria (2012) ‘Imperfect environmental policy and polluting emissions: the Green Paradox and beyond.’ *International Review of Environmental and Resource Economics* 6(2), 153–194

A Data

Descriptive statistics of all variables are provided in Table A.1.

We use information for the period 1987-1999 for our analysis. This time period starts three years before the 1990 CAAA was passed and ends at the start of Phase II of the ARP. The sample of plants in our main source of data, the U.S. Federal Energy Regulatory Commission (FERC) form FERC-423 ‘Monthly Cost and Quality of Fuels for Electric Plants’, shrinks after 1999 as plants in restructured electricity markets are dropped from the sample by the FERC. The sample begins in 1987 to avoid including the changing structure of the coal market in the pre-policy period. As discussed in Ellerman and Montero (1998) and Kozhevnikova and Lange (2009), a larger share of coal was coming from the Western coal basin and transacted in the spot market by 1987 relative to earlier in the decade.

FERC-423 contains a panel data set with monthly information on the cost and quality of coal deliveries to plants of 50 MW or larger capacity. The dataset records both deliveries made as part of long-term contracts, and deliveries that originate from spot purchases on the market. We obtain data on coal prices, sulfur content, heat content, ash content, region of origin of coal delivered and the share of coal from spot deliveries from this dataset.

Data on the NSPS status of plants were generously provided by Danny Ellerman.

Our proxy for monthly transport costs in cents per million Btu is constructed using the average distance of each plant to each of the three mining regions (from the EIA’s Coal Transportation Rate Database), the rail rates variable ‘Rail Price’ (see below), the size of the delivery in tons, and total heat delivered to the plant from each of the three regions.

Data on state-level output, to control for business cycle effects and industrial activity at the state level, are obtained from the Coincident Economic Activity Index of the Federal Reserve Bank of St. Louis. The monthly index includes nonfarm payroll employment, the unemployment rate, average hours worked in manufacturing and wages and salaries.¹⁴

Data on the gas price are taken from U.S. EIA (2010), are in U.S. dollar cents per thousand cubic feet and are discounted using the PPI for crude energy to 1982 dollars. Data on electricity generation from natural gas are collected from the Annual Energy Review, which is published by the U.S. Energy Information Administration (U.S. EIA, 2005).

‘Scrubber’ is a dummy variable equal to one in each month in which a plant had a flue gas desulfurization unit (SO₂ control equipment, also known as a ‘scrubber’) installed. It is constructed using data from EIA Form 767.

¹⁴State-level GDP data are only available from 1997 onwards.

The summer dummy equals one during the months June, July, August and September and is zero otherwise. The winter dummy equals one during the months December, January, February and March and is zero otherwise.

The rail rates variable 'Rail Price' is created with data from the EIA (U.S. EIA, 2004) on the average rate per ton-mile per year, deflated with the producer price index. The variable 'Close to PRB' includes Arkansas, Illinois, Indiana, Iowa, Kansas, Kentucky, Louisiana, Michigan, Minnesota, Missouri, Nebraska, Oklahoma, Texas, Washington and Wisconsin.

Coal mines are assigned to coal mining regions following Joskow (1987). Regional coal mine productivity data are obtained from U.S. EIA (2010).

Information on whether the state the Phase I plant is located has strict environmental standards and whether the state the Phase I plant is located in required a compliance plan prior to the start of the ARP come from Lile and Burtraw (1998).

Data on whether state has a nuclear plant located in it are obtained from U.S. EIA (2011c).

Table A.1: Descriptive statistics

Sample	(1) Full Sample	(2) Before 1990 CAAA	(3) After 1990 CAAA	(4) Phase I	(5) Non Phase I plants in States with Phase I Plants	(6) NPS-Da Plants in States with Phase I Plants	(7) Diff. (5)-(4)
Variable	Mean (S.D.)	Mean (S.D.)	Mean (S.D.)	Mean (S.D.)	Mean (S.D.)	Mean (S.D.)	
Spot Price (Real U.S.c./MBtu)	171.64 (42.85)	178 (43.38)	164.42 (39.97)	160.58 (33.11)	177.87 (42.49)	181.23 (39.07)	17.04*** (0.65)
Spot Sulfur Intensity (lbs/MBtu)	2.33 (1.47)	2.43 (1.46)	2.21 (1.47)	3.16 (1.42)	2.24 (1.45)	2.29 (1.33)	-0.91*** (0.02)
Total Heat (GBtu)	3429.94 (3462.57)	3359.74 (3407.09)	3517.78 (3528.15)	4388.5 (3649.9)	2747.35 (3177.12)	3909.25 (2364.87)	-1641.15 *** (42.13)
State Output (Coincident Index)	96.78 (8.65)	91.68 (7.29)	103.11 (5.42)	96.77 (8.68)	97.37 (8.45)	97.88 (7.91)	0.58*** (0.1)
Gas Price (U.S.c./,000 cu ft)	318.04 (54.99)	297.68 (64.04)	341.53 (27.39)	317.99 (55.2)	318.06 (54.92)	322.59 (51.35)	0.17 (0.7)
Productivity (ton/employee hour)	5.03 (3.53)	4.39 (3.01)	5.83 (3.95)	3.25 (1.75)	4.42 (3.15)	4.80 (3.58)	1.18*** (0.02)
Rail Price (Real U.S.c./ton mile)	3.09 (0.73)	3.58 (0.66)	2.49 (0.14)	3.10 (0.73)	3.09 (0.73)	3.00 (0.69)	0.01 (0.01)
Plant w/ Scrubber	0.20 (0.40)	0.19 (0.39)	0.21 (0.40)	0.05 (0.22)	0.20 (0.43)	0.75 (0.69)	0.15*** (0.01)
Mandatory Phase I Plants	0.22 (0.41)						
Appalachian Heat Share	0.45 (0.48)	0.46 (0.48)	0.45 (0.48)	0.51 (0.47)	0.53 (0.48)	0.41 (0.46)	0.02*** (0.01)
Interior Heat Share	0.23 (0.41)	0.24 (0.41)	0.21 (0.38)	0.42 (0.46)	0.23 (0.40)	0.27 (0.43)	-0.19*** (0.01)
Western Heat Share	0.30 (0.44)	0.28 (0.44)	0.32 (0.45)	0.06 (0.20)	0.21 (0.39)	0.27 (0.44)	0.15** (0.01)
Import Heat Share	0.01 (0.04)	0.01 (0.03)	0.01 (0.06)	0.01 (0.01)	0.01 (0.05)	0.04 (0.13)	0.04*** (0.01)
No Nuclear State	0.23 (0.42)			0.31 (0.45)			
High Spot Procurement	0.25 (0.23)			0.27 (0.18)			
Pre-Approval	0.17 (0.38)			0.37 (0.48)			
Strict Env. Regulations	0.15 (0.35)			0.08 (0.27)			
Number of Plants	418	409	418	89	329	201	

*, **, *** indicate 10%, 5% and 1% statistical significance, respectively.