DUKE ENVIRONMENTAL AND ENERGY ECONOMICS WORKING PAPER SERIES
organized by the
NICHOLAS INSTITUTE FOR ENVIRONMENTAL POLICY SOLUTIONS
and the
DUKE UNIVERSITY ENERGY INITIATIVE

Do Differentiated Performance Standards Help Coal? CO_2 Policy in the U.S. Electricity Sector

David A. Bielen*

Working Paper EE 14-12
December 2014
Revised August 2015

* Department of Economics, Duke University

Do Differentiated Performance Standards Help Coal? CO_2 Policy in the U.S. Electricity Sector

David A. Bielen[*,†,‡]

August 13, 2015

Abstract

The EPA's Clean Power Plan imposes different CO_2 emission rate standards on coal and natural gas power plants. This paper explores the distributional effects of a "differentiated standard" policy in the electricity sector, compared to a sector-wide single standard policy. Analytically, I find that differentiation leads to increases in coal usage, and ambiguous effects on electricity prices and coal-plant profits. Numerical simulation results suggest the impacts on coal usage are modest, resulting in modest welfare gains to upstream coal interests. They also reveal that coal plants benefit from greater profits at the expense of electricity consumers, who face higher prices.

---

[*]This paper was the central paper of my dissertation at Duke University and is being submitted as an employee at the National Renewable Energy Laboratory (NREL). I'm grateful to my dissertation committee (Billy Pizer, Chris Timmins, Andy Yates, and Richard Newell) for their feedback at all stages of the project. I would also like to extend a special thanks to Jim Bushnell for useful suggestions at multiple stages of the process. Finally, I would like to acknowledge helpful comments from attendees at the Duke University IO/Environmental Lunch Group, the 2014 NBER Environmental and Energy Economics Summer Institute, the 2014 Heartland Conference, and seminars at Wake Forest University, University of California-Davis, and NREL.

[†]Contact: National Renewable Energy Laboratory, 15033 Denver West Parkway, Golden, CO 80401, US; Email: David.Bielen@nrel.gov.

[‡]I have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

While academic economists have traditionally focused on the cost-effectiveness properties of environmental policy instruments, their distributional consequences are receiving increasing attention.[1] With the United States Environmental Protection Agency (EPA) recently finalizing regulations for carbon dioxide (CO_2) emissions from existing power plants in the form of the Clean Power Plan, equity issues have emerged as a central issue in the policy debate. Here, the EPA and ultimately individual states must make choices that will determine the relative burden on the coal industry, coal regions, and downstream coal users. These choices will likely influence the ease or difficulty with which the policies are implemented as well as views about future policies to address climate change.

The options available for influencing distributional impacts depend on the type of policy under consideration. Under a cap-and-trade system, the allocation of permits can be manipulated to address equity concerns. The current regulatory process, however, leans towards rate-based standards that cap emissions per unit of electricity, rather than emissions (Burtraw, Fraas and Richardson, 2012). Trading is still possible and cost-effective, allowing low-mitigation-cost facilities to overcomply, in turn earning and selling credits to high-mitigation-cost facilities that undercomply. Under this type of policy, the analogy to using permit allocation for redistribution would seem to be manipulating the rate standard across facilities, yet the efficacy of this approach is neither demonstrated nor well understood. In our particular context, the question is whether "differentiation" helps coal and its constituents when the standard for coal-fired generation is relaxed while the standard for natural gas generation is tightened. Answering this question and quantifying the associated impacts is the central focus of this paper.

Does differentiation improve outcomes for various coal-oriented stakeholders under a tradable performance standard policy in the U.S. electricity sector? To answer this question, I start with proxies for improved outcomes.

[1]For example, see Bovenberg, Goulder and Gurneyl (2005), Bushnell and Chen (2012), Rausch and Mowers (2014)).

Aggregate coal usage, coal plant profits, and wholesale electricity prices represent the economic impacts on the interests of coal producers and laborers, coal-fired power plants, and electricity distributors and consumers in coal-heavy generation regions, respectively. The analysis begins with a simple analytic model to understand how differentiation affects these outcomes and what is theoretically possible. Where the theory is ambiguous, and in order to provide detailed quantitative estimates, I then turn to a state-of-the-art simulation model of the U.S. wholesale electricity market.

Using the analytic model, I demonstrate how, compared to a single standard policy, differentiation is expected to increase coal usage, but that the directions of price and profit changes depend on two competing effects. The coal usage result is driven by changes in emission reduction behavior: I show that differentiation prioritizes emission reductions from within-fuel switching of generation from dirty to clean (coal) plants, versus between-fuel switching from coal- to natural gas-fired plants. This results in less coal displacement. Differentiation also results in higher prices for compliance credits.[2] As a result, relaxing the standard for coal-fired plants has the counteracting effects of requiring fewer credit purchases but raising the credit prices. I show how these competing cost effects are related to electricity price and plant profit outcomes, and how the net effects depend on the entire system of generating facilities and opportunities for various types of mitigation.

I use the simulation model to generate plausible quantitative impacts of differentiation under a national tradable performance standard policy. Consistent with the analytic results, I find that differentiation increases coal usage through an increase in dirty to clean plant switching within fuel categories. However, the extent is modest ($\approx 1\%$ in coal usage) due to the limited potential of within-fuel switching as a means of compliance. I also find that, on average, differentiation increases electricity prices in almost every region in

[2]This result is a variant of a key result from Bohringer and Lange (2005), which studies benchmarking schemes for permit allocation in a cap-and-trade program. The authors show that if emission rates are identical within fuel categories, then allocating permits proportionally to emissions in a closed system will result in a proportionate rise in the credit price relative to a pure output-based allocation scheme.

the country, even regions that rely heavily on coal-fired generation. Increasing electricity prices help to bolster coal-fired plant profits. Additionally, more than half of the utilized coal plant capacity in the model faces decreasing per unit costs under differentiation. In the extreme, the combination of these two effects results in aggregate coal-fired plant profits that actually exceed the level in the absence of regulation.

Taken together, the results imply that differentiation of a tradable performance standard on the basis of fuel-type does little to aid coal producers and laborers, and hurts electricity consumers in coal-heavy regions even more. Unless the goal is to assuage coal-fired power plants, this type of policy is an ineffective one, at least in a short-run analysis. However, the analytic results suggest that differentiation could be more effective in situations where switching between differentiated facilities (e.g., from coal to gas in this case) is not the overwhelming source of mitigation. For example, this might be the case for programs with regional- rather than fuel-based differentiation, an area that I plan to investigate in future work.

The paper is organized as follows. In Section 2, I describe the policy setting for CO_2 regulation in the electricity sector. Section 3 reviews the existing literature on tradable performance standards and other similar policies. In Sections 4 and 5, I develop a simplified version of the simulation model and use it to examine how differentiation could impact coal usage, coal plant profits, and electricity prices. Section 8 describes the construction of the simulation model, the equilibrium conditions, and the policies for which I report results. Those results are described in Section 7. This section also contains discussion relating the analytic section to the simulation results. Section 8 concludes with a discussion of the implications of the model results.

2 Carbon Policy in the U.S. Electricity Sector

In this section, I describe the policy setting for the EPA's Clean Power Plan (CPP), focusing on recent developments in the executive and legislative branches of the U.S. government. I briefly explain the mechanics of regulating through the Clean Air Act Section 111(d), which is what the EPA has used for the existing source rulemaking. The background provided will provide a more comprehensive motivation for the policy design feature studied in this paper.

Until recently, greenhouse gas (GHG) policy seemed to be most likely introduced through legislative action. The American Clean Energy and Security Act of (2009), also known as the Waxman-Markey Bill, proposed an economy-wide cap-and-trade program for CO_2; it was narrowly approved by the House of Representatives but comparable legislation was never brought to a vote in the Senate. Since then, several clean energy standards have been proposed, including the Clean Energy Standard Act of 2012, known as the Bingaman Proposal.[3] Unlike the Waxman-Markey Bill, the Bingaman Proposal focused solely on the electricity sector. The Proposal also suggested a crediting system for each megawatt-hour (MWh) of electricity produced below a given emissions intensity standard, rather than establishing a cap-and-trade system.[4] Although such a clean energy standard is not dissimilar from what is expected from the EPA process, attempts to establish one on a national basis through legislation proved unable to break through congressional gridlock.

Though attempts to regulate GHGs through new legislation failed, an older piece of legislation, the Clean Air Act (1970), has eventually proven to be a vehicle for executive branch action. At the same time that climate policy was being debated in Congress, there were several interesting developments regarding the EPA's ability and obligation to regulate GHGs. First, the Supreme Court ruled in *Massachusetts v. EPA* (2007) that the EPA needed to articulate a reasonable basis if it was not going to regulate. Two years later, EPA released a formal "endangerment finding," which relied on scientific evidence to establish that GHGs are air pollutants. These two events have since initiated a string of rule-makings as set out by the Clean Air Act. First, the EPA dealt

[3]For a detailed discussion of the merits of a clean energy standard similar to the Bingaman Proposal, see Aldy (2011). Simulation results for various clean energy standards using Resources for the Future's Haiku model are featured in Paul, Palmer and Woerman (2013*b*,*a*).

[4]The intensity standard is denoted in terms of emissions, e.g., metric tons of CO_2, per MWh. Such a standard is also often referred to as an efficiency or performance standard; this paper utilizes the latter.

with the transportation sector by establishing vehicle emission standards. The next step was to address the electricity sector, which is currently the focus of the EPA.

Before addressing existing power plants, the EPA needed to first regulate new power plants, and the result of that process helped to motivate the policy studied here. In its New Source Performance Standards, the EPA established different emission rate targets for natural gas and coal plants. This outcome suggested that a similar distinction could be applied under the existing source regulation. This is exactly what has occurred in the EPA's final rule.

The distinction between natural gas and coal in the CPP manifests in two ways (U.S. Environmental Protection Agency, 2015). First, the EPA specifies technology-specific, i.e., coal and natural gas, targets that states can choose to comply with. Alternatively, states can choose to comply with state-specific targets, which are tied to the ratio of coal to natural gas generation in each state. States with more coal-fired plants received higher, i.e., more tons of CO_2 per MWh, standards. The Clean Air Act is set up so that the EPA sets emission rate regulations, and the individual states construct implementation plans that inform the EPA how the state intends to comply; additionally, states are encouraged to form coalitions and submit joint plans. This suggests a few ways in which differentiated standards could be applied in the implementation phase. For one, individual states or coalitions could choose to treat coal and natural gas differently within their own tradable performance standard policy, using either the technology-specific standards laid out by EPA or other standards that achieve equivalent emission reductions. Also, coalitions could differentiate on the basis of state, using the EPA's specific rates as a benchmark for crediting. It seems quite possible, if not likely, that a variety of compliance pathways will emerge across the U.S. Even though the national policy modeled in this paper is unlikely to arise in the current regulatory framework, the intuition developed here is a useful aid to understanding the impacts of any style of differentiated policy.

My focus on tradable performance standard policies as a means of achieving compliance is supported by the structure of the Clean Air Act regulations.

Because the EPA specifies targets in terms of emission rates rather than levels, tradable standards are a natural candidate for a flexible, market-based implementation policy (Burtraw, Fraas and Richardson, 2012). However, to be clear, the CPP also calculates for each a state a mass-based equivalent (U.S. Environmental Protection Agency, 2015). The extent to which states will choose to exercise that option remains to be seen. As will be demonstrated in the next section, tradable standard policies have become ubiquitous in energy sectors. The working assumption motivating this paper is that, at the very least, they will play a role in regulating existing sources.

3 Literature Review

Despite the emergence of tradable performance standards in a variety of contexts, the vast majority of research on market-based policies has focused on cap-and-trade. This paper contributes to an emerging literature on tradable performance standards, as well as a literature concerned with the distributional impacts of environmental policy. In this section, I demonstrate how my paper fits into those two literatures. I also establish a conceptual link between differentiated performance standard policies, and cap-and-trade policies that allocate permits according to a bench-marking rule.

One of the earliest studies of tradable performance standards is Helfand (1991), which theoretically studies simultaneously the output, profit, and efficiency properties of a variety of emission policies, including a rate-based standard. More recently, Fischer (2001) performed a theoretical comparison of schemes with a revenue rebating feature, including tradable performance standards and cap-and-trade programs with output-based allocations. The author demonstrates how each of these policies results in inefficient outcomes by shifting abatement efforts away from demand contraction and towards emission rate reduction. Fischer (2011) extends the theoretical analysis of this type of policy to markets with imperfect competition, while Fischer (2003) examines the environmental consequences of allowing trading across sectors with rate-based and cap-and-trade regulations. Finally, Fischer and Fox (2007) considers the tax interaction effects of tradable performance standards for both efficiency and distribution of rents.

The work described so far has focused on either general theory or economy-wide policy; other work has focused on the application of tradable performance standard policies in a variety of specific contexts. For example, Holland, Hughes and Knittel (2009) and Lemoine (2014) both analyze tradable performance standards applied to CO_2 emissions from transportation fuels, referred to as low-carbon fuel standard policies, focusing on theoretical efficiency properties. Lemoine (2014) in particular points out that the regulator can achieve higher welfare by choosing each fuel's emission ratings rather than just setting the level of the standard, which is akin to a result we see in the next section that a differentiated standard policy can yield efficiency gains relative to a single tradable standard policy. Tradable performance standards have also been addressed in the context of regulating the energy efficiency of automobiles, or Corporate Average Fuel Economy (CAFE). For instance, Rubin, Leiby and Greene (2008) considers the implications of various market structures in the CAFE compliance credit market on cost-savings, while Ito and Sallee (2014) analyzes the welfare properties of attribute-based standards in the context of Japanese fuel-economy standards. Regarding the latter work, differentiated standards can be considered a special case of attribute-basing in which the attributes are simply the fuel-types or automobile class in the cases of electricity and transporation, respectively. This paper is most similar to Lemoine (2014) and Ito and Sallee (2014) in terms of the type of policy under consideration, but the context is different and my primary focus is on distributional effects rather than cost-effectiveness or efficiency.

Another related paper is Bushnell and Chen (2012), which performs a simulation study of cap-and-trade programs in the western United States. Although much of the authors' focus is on leakage, they also compare a few different permit allocation rules, including one that is "fuel-based." This allocation design is similar to a generic output-based scheme, but the number of permits allocated per megawatt hour is higher for coal-fired plants than for natural gas-fired plants. Citing Bohringer and Lange (2005), the authors note that relative to a pure output-based design, under fuel-based allocation electricity prices rise and many of the benefits to firms with an abundance of coal-fired capacity are limited by rising permit prices. My paper builds on the work of Bushnell and Chen (2012) by investigating a related policy analytically and extending the simulation to the entire country. In addition to the mid-stream electricity generators, I also extend the analysis to coal producers and electricity consumers.

Finally, this paper is related to Bushnell et al. (2015), which also investigates policy design options for existing source regulation of electricity generators under the Clean Air Act. In that work, the authors conduct a game-theoretic analysis of whether states in a coalition would choose to implement mass- or rate-based policies, finding that even though mass policies are more efficient in aggregate, strategic states may choose to implement rate-standards to protect their own interests.

4 A Simple Electricity Market Model

In this section, I develop a framework for analyzing a differentiated standard policy in the U.S. electricity sector. The analytic model presented here is constructed to capture key stylized features of the power market. To build intuition, I demonstrate how the model works in the absence of regulation. I then introduce the performance standard policies to draw conclusions about the impact of differentiation on aggregate coal usage, regional electricity prices, and coal plant profits.

The analytic model is a static representation of an electricity sector with perfect competition in both the product (electricity) and credit markets. Electricity markets are region-specific in the sense that electricity cannot be traded across regions. Each region is endowed with a set of power plants that supply electricity. Individual plants are indexed by region (r), fuel type (f), and unit (i). For simplicity, the analytic model only considers coal (subscript *coal*) and natural gas plants (*gas*) as distinct fuel types.[5] The collection of electricity

[5] Other plant types, such as nuclear and renewables, are typically run at full capacity due to their low marginal costs. They are also not subject to any regulatory costs that would change the plant economics. Therefore, in a short run analysis of emissions policy, the relevant economic features of the policy are captured by the tradeoffs between coal and

generators in each region faces a downward-sloping demand curve represented by the function $P_r(\cdot)$.

Each plant is endowed with a capacity $K^i_{f,r}$, emission rate $\sigma^i_{f,r}$, and unit cost of generation $c^i_{f,r}$. All plant-specific parameters do not vary with generation; for costs, this is equivalent to assuming a perfectly elastic supply of fuel. Additionally, in accordance with the short-run nature of the model, plant capacities, unit costs, and emission rates are held fixed, and no new plants are built.[6] The endogenous variable associated with each plant is its generation level, denoted by $q^i_{f,r}$, which is bounded below by zero and above by the plant capacity. Total plant emissions and costs of generation are linear functions of the generation level, represented by $\sigma^i_{f,r} \cdot q^i_{f,r}$ and $c^i_{f,r} \cdot q^i_{f,r}$, respectively.

4.1 No Regulation

To form a basis for comparison, I present the model first in the absence of regulation, a scenario that I refer to as business-as-usual (BAU). The social welfare function is given by the sum of consumer and producer surplus:

$$W(q^i_{f,r}) = \sum_r \left[\int_0^{\sum_i \sum_f q^i_{f,r}} P_r(u)\mathrm{d}u - \sum_i \sum_f c^i_{f,r} \cdot q^i_{f,r} \right]. \tag{1}$$

The social planner's program is given by

$$\begin{aligned} \underset{\{q^i_{f,r}\}}{\text{maximize}} \quad & W(q^i_{f,r}) \\ \text{subject to} \quad & 0 \leq q^i_{f,r} \leq K^i_{f,r} \quad \forall i, f, r. \end{aligned} \tag{2}$$

Thus, the social planner's objective is to choose plant-specific generation levels so that social welfare is maximized subject to the constraints on plant capacities. Given the structure of the model, the solution to (2) will be equivalent to the competitive market equilibrium.

Alternatively, production decisions can be presented from the perspective

natural gas plants.

[6] Although this paper does not explicitly model the dynamic effects of a differentiated policy, it does provide insight into what they would be.

of the individual price-taking plants. The objective is to choose the generation level that maximizes profits subject to the capacity constraint:

$$\begin{aligned} \underset{q^i_{f,r}}{\text{maximize}} \quad & p_r \cdot q^i_{f,r} - c^i_{f,r} \cdot q^i_{f,r} \\ \text{subject to} \quad & 0 \leq q^i_{f,r} \leq K^i_{f,r}. \end{aligned} \tag{3}$$

Profits are the difference between total revenues and total costs. From the plant perspective, the electricity price is given by a parameter, p_r. The Lagrangian associated with this program is defined as follows:

$$\mathcal{L}(q^i_{f,r}, \lambda^i_{f,r}) \equiv p_r \cdot q^i_{f,r} - c^i_{f,r} \cdot q^i_{f,r} - \lambda^i_{f,r}(q^i_{f,r} - K^i_{f,r}). \tag{4}$$

Each plant faces the following first-order conditions (FOCs):

$$\begin{aligned} & q^i_{f,r} \geq 0, \quad p_r - c^i_{f,r} - \lambda^i_{f,r} \leq 0 && \lambda^i_{f,r} \geq 0, \quad q^i_{f,r} - K^i_{f,r} \leq 0 \\ & q^i_{f,r}(p_r - c^i_{f,r} - \lambda^i_{f,r}) = 0 && \lambda^i_{f,r}(q^i_{f,r} - K^i_{f,r}) = 0. \end{aligned} \tag{5}$$

Market equilibria are defined by the above conditions for each plant, as well as an aggregate resource constraint for each region that ensures that aggregate supply quantity equals the demand quantity.

On the individual plant level, generation decisions can be distilled into three possibilities. First, it is not economical for the plant to generate, and so $q^i_{f,r} = 0$. This will be the outcome if the equilibrium electricity price exceeds the plant's unit cost of generation. Second, the plant generates up to its capacity so that $q^i_{f,r} = K^i_{f,r}$. This will occur is the equilibrium price is greater than the cost of generation, and so the plant will receive a rent ($\lambda^i_{f,r} > 0$) for each unit it generates. Finally, the plant might be on the margin where it's cost of generation sets the equilibrium price. In this case, it will be the case that $0 < q^i_{f,r} < K^i_{f,r}$ and $\lambda^i_{f,r} = 0$.

One possible business-as-usual equilibrium is illustrated in Figure 1. The graphical example demonstrates the outcome for a two-region scenario with two plants in each region. Region 1 is endowed with a low-cost coal plant and a higher-cost gas plant. Region 2 generates its electricity at two coal plants, one which has lower costs of generation than the other. For simplicity, electricity demand is modeled as perfectly inelastic. In both regions, electricity supply is

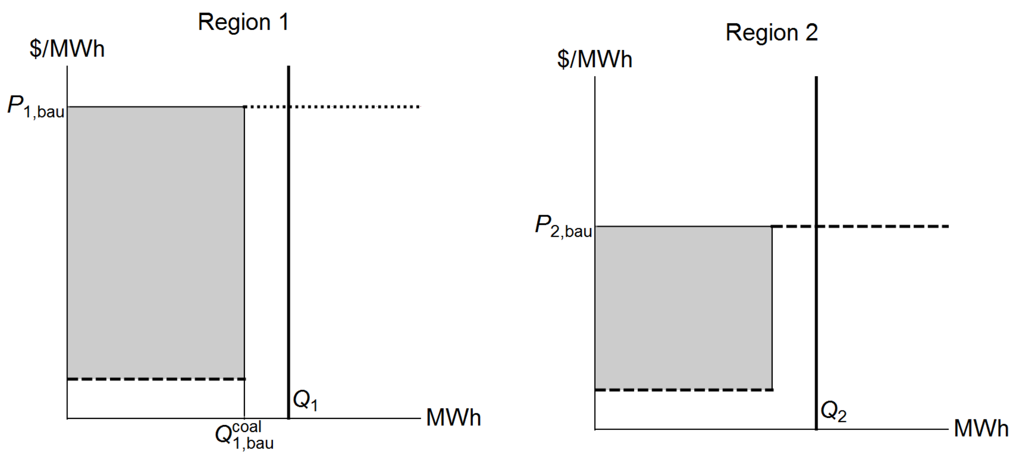

Figure 1: Two-Region Business-As-Usual Electricity Market Equilibrium

an increasing step-function; supply for gas and coal plants are represented by dotted and dashed lines, respectively. To meet electricity demand (represented by the thick, vertical solid lines), units are "dispatched" in order of increasing unit cost. The more expensive plants in each region set the equilibrium price of electricity. The cheaper plants are run at capacity, and accrue profits equal to the shaded areas.

4.2 Tradable Performance Standard Policies

Now suppose that the social planner seeks to reduce emissions through a tradable performance standard policy. Initially, the policy under consideration sets a single standard for all fuel types. Moreover, the performance standard chosen by the regulator must be below the aggregate emission rate in order to change firm behavior. If the quantities $q^{i,bau}_{f,r}$ represent plant generation levels in the absence of regulation and s represents the standard, this means that the inequality

$$\frac{\sum_r \sum_i \sum_f \sigma^i_{f,r} q^{i,bau}_{f,r}}{\sum_r \sum_i \sum_f q^{i,bau}_{f,r}} \leq s \tag{6}$$

must be satisfied. If it is, then I say that the standard is binding.

How will the market comply with the standard? In the electricity sector, I focus on two margins.[7] The first is the coal-to-gas margin, in which coal-fired generation is displaced by gas-fired generation. The second is the within-fuel margin, either coal-to-coal or gas-to-gas, in which generation from a high emission rate plant is displaced by generation by a low emission rate plant within the same fuel category. As I will explain later, these margins are important when comparing a single standard to a differentiated policy.

Implementation of a binding tradable performance standard policy will incorporate a regulatory cost component to the economics of the individual plants. Plants with emission rates below the standard ($\sigma^i_{f,r} < s$) will be suppliers of compliance credits, whereas plants with emission rates above the standard will be demanders. For the individual plant, the profit-maximization problem now has to account for credit sales/purchases:

$$\begin{aligned} \underset{q^i_{f,r}, z^i_{f,r}}{\text{maximize}} \quad & p_r \cdot q^i_{f,r} - c^i_{f,r} \cdot q^i_{f,r} - \tau \cdot z^i_{f,r} \\ \text{subject to} \quad & 0 \leq q^i_{f,r} \leq K^i_{f,r} \\ & (\sigma^i_{f,r} - s) q^i_{f,r} \leq z^i_{f,r}. \end{aligned} \tag{7}$$

Net credit demand is denoted by $z^i_{f,r}$, and the equilibrium price for a credit is τ . If $z^i_{f,r} > 0$, then the plant will demand credits; otherwise, it will supply them. Because there is no benefit to not selling excess credits or purchasing more than is required to comply, the credit constraint will bind for all plants. As a result, credit demand can be substituted out of the objective function, and maximization occurs over a single variable.

[7]Other possible options would be to produce more electricity through gas without adjusting coal-fired generation levels, or to produce less electricity at coal plants while holding gas-fired levels fixed. However, in the short-run, electricity demand is considered to be largely inelastic, which increases the cost associated with these mechanisms. For simplicity, I ignore the contribution to compliance through these channels.

The FOCs for plants under the single standard policy account are

$$
\begin{aligned}
q^i_{f,r} &\geq 0 & \lambda^i_{f,r} &\geq 0 \\
p_r - c^i_{f,r} - \tau(\sigma^i_{f,r} - s) - \lambda^i_{f,r} &\leq 0 & q^i_{f,r} - K^i_{f,r} &\leq 0 \\
q^i_{f,r}(p_r - c^i_{f,r} - \tau(\sigma^i_{f,r} - s) - \lambda^i_{f,r}) &= 0 & \lambda^i_{f,r}(q^i_{f,r} - K^i_{f,r}) &= 0.
\end{aligned}
\tag{8}
$$

The market equilibrium is defined by these conditions, the aggregate resource constraints for each region, and a credit market clearing condition that requires that the number of credits purchased must equal the number of credits sold. Mathematically, this is equivalent to

$$
\sum_r \sum_i \sum_f z^i_{f,r} = \sum_r \sum_i \sum_f (\sigma^i_{f,r} - s) q^i_{f,r} = 0. \tag{9}
$$

4.3 Calculating the Credit Price

The only difference in between the business-as-usual and performance standard conditions from a plant's perspective is the presence of the regulatory cost, $\tau(\sigma^i_{f,r} - s)$. The emission rate and standard are exogenous in this setting, while the credit price is endogenously determined by the model. To see how the credit price is determined, imagine a low stringency policy such that the standard is set just below the business-as-usual aggregate emission rate.

Recalling our discussion of compliance margins, it must be the case that cheaper, higher emission rate generation is displaced by expensive, lower emission rate generation. In order to maintain the same total level of electricity supply in each region, this displacement must occur between plants within the same region. Furthermore, because of the structure of the model with constant unit costs and fixed capacity, it must be the case that the displaced plant is operating at capacity, and the displacing plant is operating below capacity. For a low stringency policy, the level of displacement could be fairly minimal so that both plants are operating below capacity. As implied by the FOCs, these plants will both be at the margin, i.e., the equilibrium price level will be set by their net generation costs. We can use this fact to calculate the credit price:

$$\begin{aligned} & c^i_{f,r} + \tau(\sigma^i_{f,r} - s) = c^{i'}_{f',r} + \tau(\sigma^{i'}_{f',r} - s) \\ \Rightarrow \qquad & \tau = \frac{c^{i'}_{f',r} - c^i_{f,r}}{\sigma^i_{f,r} - \sigma^{i'}_{f',r}}. \end{aligned} \tag{10}$$

Because the credit price is a function of plant-specific parameters, it follows that only two plants will be actively switching – across all regions – for any given level of the standard (unless there are two pairs of plants in different regions who have unit cost equalized by the same credit price). Going forward, I will refer to the plants on the switching margin as the *marginal switching pair*.

To illustrate the single-standard policy graphically, I return to the two-region example. Figure 2 demonstrates how generation, electricity prices, and profits are altered by implementing a performance standard. Consistent with the above discussion, the policy illustrated is of relatively low stringency; in the next section, I illustrate what happens as the standard is tightened. The

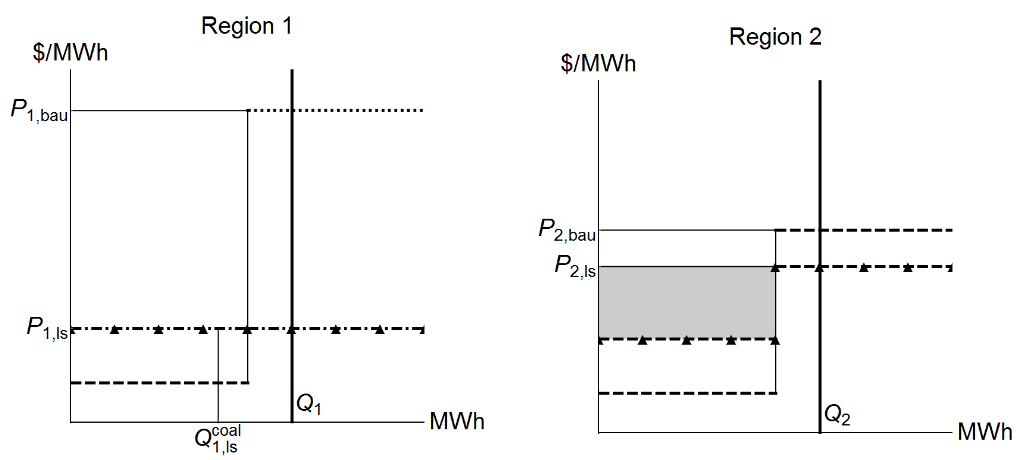

Figure 2: Two-Region Single Standard Equilibrium: Low Stringency

supply curves in each region under the policy are represented by the triangle-studded lines. For comparison, the business-as-usual supply curves are reproduced as in the previous figure. Once again, gas and coal plant supplies are represented by dotted and dashed lines, respectively. A dashed-dotted line is used to represent supply on the margin formed by coal and gas plants with equal costs. Compliance is achieved by switching generation from a high emission rate plant to a low emission rate plant, which is what happens in Region 1. Coal plant generation is displaced by generation from the natural gas plant. The unit costs of these plants are equal, as the coal plant is taxed and the gas plant is subsidized. The amount of coal generation in Region 1 is denoted by $Q_{1,ls}^{coal}$, which is clearly less than the amount under BAU.

In Region 2, the marginal coal plant is subsidized under the standard while the infra-marginal plant is taxed. This has the effect of lowering the electricity price while raising the cost of the low cost plant. As a result, the profits of this plant are reduced significantly.

4.4 Building Emissions Reductions

The policy applied to the two-region example in the previous section featured a single switching pair; however, meaningful aggregate emission reductions for the U.S. electricity sector would require switching among many pairs of plants. In the current setting, this requires tightening the performance standard. In this section, I examine what happens to the credit price as the policy becomes more stringent. The intuition developed here is useful in analyzing the impacts of a differentiated standard.

Using the equilibrium from Figure 2 as a starting point, suppose that the regulator tightens the standard slightly: I refer to this as a medium stringency policy. The result is depicted in Figure 3. The new supply curves are represented by the square-studded lines. In order to comply with the new standard, there is additional switching from the coal plant to the gas plant in Region 1, as depicted by the distance between $Q_{1,ls}^{coal}$ and $Q_{1,ms}^{coal}$. Because the credit price is purely a function of the parameters on the switching margin, it does not vary with the tightening of the standard, provided the switching pair does not change. However, the unit costs do depend on the level of the standard, and raising the standard means increasing unit costs for each plant, which is reflected as an upward shift in the supply curve. Therefore the electricity price increases in both regions, even though the credit price remains unchanged. Moreover, with the marginal switching pair in Region 1, generation in Region

2 does not change.

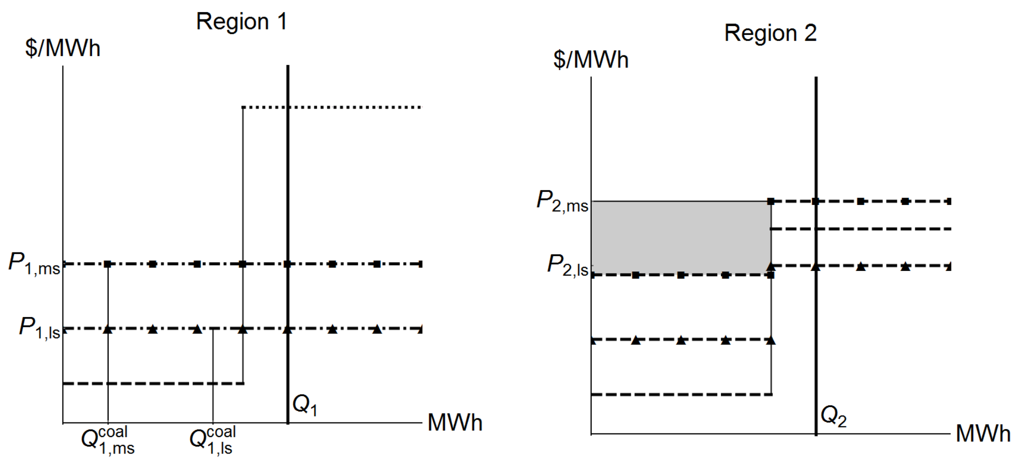

Figure 3: Two-Region Single Standard Equilibrium: Medium Stringency

As the regulator continues to tighten the standard (high stringency), additional switching opportunities are undertaken. In the present example, this occurs when the single gas plant in Region 1 reaches capacity. Figure 4 demonstrates this graphically. The supply curves for the high stringency standard are represented by diamond-studded lines. When the standard has been tightened sufficiently, the dispatch order in Region 1 flips: the coal plant is marginal and sets the electricity price and the gas plant now has the lower unit cost and accrues profits. Region 2 supplies the additional compliance

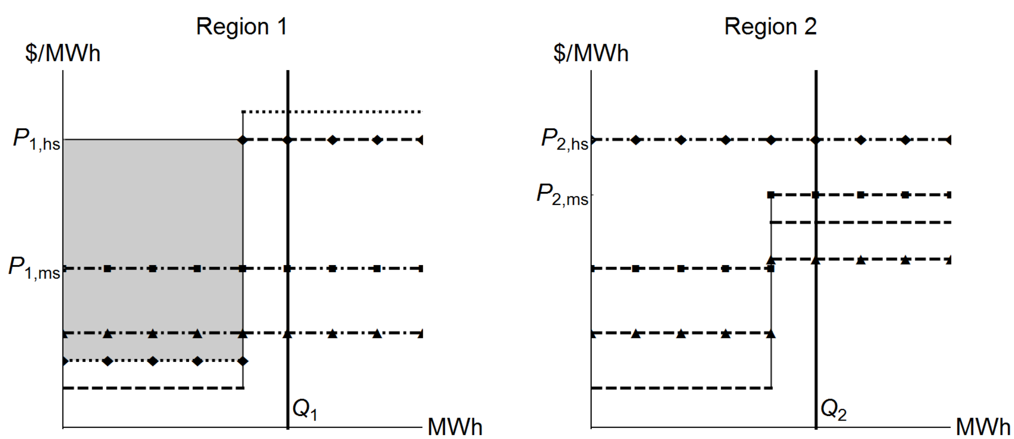

Figure 4: Two-Region Single Standard Equilibrium: High Stringency

margin; the marginal switching pair is now the two coal units in Region 2. In order for this margin to exist, the higher cost coal plant must also have a lower emission rate; otherwise, there is no way to wring further emission rate improvements from this scenario. Additionally, this switching pair will now define the equilibrium credit price, which must be greater than the credit price resulting from the Region 1 switching pair.[8] Therefore, the credit price will jump discretely, and, as a result, the compliance cost schedule will be an increasing step function.

4.5 Differentiating the Standard

In this section, I describe what happens when the social planner differentiates the standard. The insights developed here build off of the developments from the previous section, and form the basis for the analytic results on coal usage, coal profits, and electricity prices provided in the next sub-section.

The impact of differentiated standards on electricity market economics requires an understanding of how they alter the compliance margins. Recall the emission rate constraint given by (9). Unlike in the previous section, however, imagine that the regulator provides coal and natural gas plants with different standards such that $s_{coal} > s_{gas}$. As with the single standard, there is a single marginal switching pair associated with a particular choice of s_{coal} and s_{gas}, and the plants in that pair have equal unit costs. In this case, the new equilibrium credit price when these two plants are providing marginal mitigation can be determined as follows:

$$\begin{aligned} & c^i_{f,r} + \tau(\sigma^i_{f,r} - s_f) = c^{i'}_{f',r} + \tau(\sigma^{i'}_{f',r} - s_{f'}) \\ \Rightarrow \qquad & \tau = \frac{c^{i'}_{f',r} - c^i_{f,r}}{\sigma^i_{f,r} - \sigma^{i'}_{f',r} - (s_f - s_{f'})}. \end{aligned} \tag{11}$$

[8] To see this, I invoke the well-known equivalence between the social planner's optimal outcome and the competitive equilibrium outcome. The equilibrium credit price will be equal to the shadow value on the social planner's credit constraint. Lagrangian duality implies that this credit price is the smallest such price that the emissions constraint is satisfied. Therefore, as the constraint becomes more and more stringent, the credit price must rise.

Compared to the single standard policy, the switching margins have been distorted by the parenthetic expression in the denominator. Note that if the switching margin is between two fuels of the same type, the associated credit price is unchanged by differentiation. However, in the case of a switch between coal and natural gas, the credit price increases unambiguously. The way to think about this is if you line up the switching pairs in the order of credit price they require, all else equal, within-fuel switches have become less costly relative to coal-to-gas switches. Moreover, as the level of differentiation increases, the penalty on coal-to-gas switches rises. In an absolute sense, all switches are at least as costly under differentiation, which cause the equilibrium price to rise for a given level of stringency.[9] In the following sub-section, these credit price and switching behaviors will be illustrated graphically, and then linked to coal usage.

5 Impacts from Analytic Model

In this section, I use the framework developed previously to derive some basic facts about the potential impacts of differentiating the standard on coal usage, coal plant profits, and regional wholesale electricity prices. I also address the impacts on cost-effectiveness. The goal of this section is to demonstrate analytic results where possible and facilitate understanding of the simulation results presented later.

5.1 Coal Usage

Given relatively inelastic electricity demand and the fact that coal-fired power plants typically emit CO_2 at double the rate of natural gas plants, *any* regulation to reduce CO_2 will necessarily result in significant reductions in the amount of coal-fired generation. Accordingly, a tradable performance standard would have a negative effect on the demand for coal, leading to a decrease in coal consumption. In this section, I demonstrate how differentiating the standard will typically lead to a "less negative effect" on coal usage, which the simulation model will then quantify.

[9] This result is reminiscent of the Bohringer and Lange (2005) result for bench-marked allocation in a cap-and-trade program.

The mechanism through which differentiating an otherwise equivalent performance standard might increase coal usage can be seen by comparing within-fuel and coal-to-gas switching margins. Previously, I established that differentiation reduces the cost of within-fuel switches relative to coal-to-gas switches. This suggests that for a given level of emissions reduction, a greater proportion will be met by within-fuel "clean-to-dirty" switches. Such a switch will not change the composition of electricity generation by fuel. Therefore, the extent to which these types of switches are available will determine in part how much less coal generation is displaced by natural gas generation under the differentiated standard, and thus how much more coal is used.

The impact of differentiation on within-fuel switching frequency in a hypothetical setting is illustrated in Figure 5. The graph depicts the marginal compliance cost schedule against the level of emission reductions. Focus first on the thick dotted and solid horizontal lines in the left panel. Each step represents a different switching pair under a single standard policy. The dotted and solid steps correspond to coal-to-gas and coal-to-coal switches, respectively. The vertical line represents the emission reduction target of the policy. Each point along the x-axis has a one-to-one mapping to specific choice of the standard. The marginal compliance cost at the intersection with the emission reduction target will yield the equilibrium credit price. Note that under the single standard policy, two of the three coal-to-coal switches are not undertaken under the target. The thin horizontal lines correspond to the marginal compliance costs under a differentiated standard policy. In the left panel, the pure cost effect is isolated. In mathematical terms, this is the impact of including the difference in the fuel-specific standards from the credit price expression in (11). This causes the costs to increase for all coal-to-gas switches, while the costs for coal-to-coal switches remain the same. In the right panel, the marginal compliance costs are resorted into the new ordering that will occur and the new equilibrium can be observed. As a result, the coal-to-coal switches that previously were not employed are used to achieve the same target. This results in less coal displacement, and therefore greater coal usage under the differentiated policy.

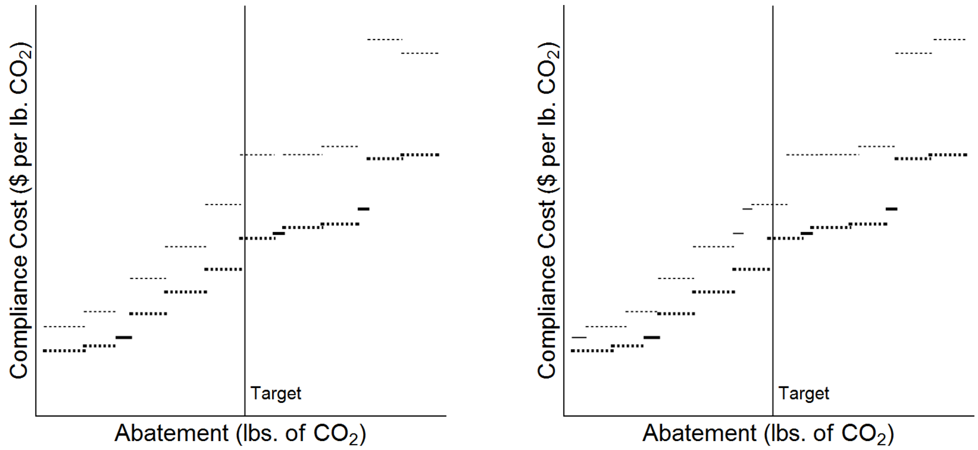

Figure 5: Credit Price as a Function of Emission Reductions

The graphic also demonstrates how the credit price will increase under a differentiated policy. This observation will be very important in understanding the impact of differentiation on plant unit costs, which will ultimately drive the price and coal plant profit results.

5.2 Per Unit Costs

To analyze the cost impacts of a differentiated policy, I investigate what happens at the plant level. Unit generation costs for plant i of fuel-type f in region r are defined as:

$$UC^i_{f,r} = c^i_{f,r} + \tau(\sigma^i_{f,r} - s_f) \tag{12}$$

Now, I explore how these costs change with an incremental increase in the coal standard. Mathematically, this can be seen by treating the credit price and the gas standard as functions of the coal standard. These functions are implicitly defined by the market equilibrium conditions on generation, the credit market constraint, and an equation specifying that emissions reductions are the same across all policies. Taking the derivative of (12) with respect to that standard, the resulting change (for representative coal and gas plants) is

$$\frac{\partial UC^i_{coal,r}}{\partial s_{coal}} = \frac{\partial \tau}{\partial s_{coal}}(\sigma^i_{coal,r} - s_{coal}) - \tau \tag{13a}$$

$$\frac{\partial UC^i_{gas,r}}{\partial s_{coal}} = \frac{\partial s_{gas}}{\partial s_{coal}} \cdot \left(\frac{\partial \tau}{\partial s_{gas}}(\sigma^i_{gas,r} - s_{gas}) - \tau \right) \tag{13b}$$

Note that the assumption of constant unit generation costs means that changes in plant unit costs will depend entirely on the net change in the unit costs of regulation. Also, in order to keep emissions reductions fixed, the gas standard must tighten as the coal standard relaxes, i.e., $\partial s_{gas}/\partial s_{coal} < 0$.

For a "small" degree of differentiation, the net change in unit costs for both coal and gas plants depends on the relative magnitudes of two competing effects. As previously demonstrated, increasing the coal standard while keeping emission reductions fixed will result in an increasing credit price; this also implies that $\partial \tau / \partial s_{gas} < 0$. In isolation, the rising credit price results in larger per unit taxes for coal plants and subsidies for gas plants: this is captured by the first term of each equation in (13). However, an increasing coal standard also reduces the amount of compliance credit purchases for each unit generated at the coal plants, which reduces their per unit tax. For gas plants, the amount of compliance credits created (and hence sold) for each unit generated reduces, which reduces their per unit subsidy. This is captured by the "τ" term in the parenthetical expression for both plant types. Ultimately, the net change in unit cost for a given plant will depend on which of these effects dominates.

Eventually, if the standards are differentiated to a sufficient degree, generation from low emission rate coal plants can become subsidized. Alternatively, generation from high emission rate gas plants can become taxed. For plants in those two categories, the impact of further differentiation on unit costs is unambiguous: they decrease for the coal plants and increase for the gas plants.

Because the expressions in (13) depend on plant-specific emission rates, it is possible, if not likely, that unit costs for some plants within a fuel category will increase while others will decrease. In any case, there is an emission rate in each category for which costs do not change. I refer to these rates, associated with a given level of differentiation, as *iso-cost emission rates*. Mathematically, they are defined for both coal and gas plants by the following equations:

$$\frac{\partial UC_{coal,r}}{\partial s_{coal}} = 0 \Leftrightarrow \sigma^*_{coal}(s_{coal}) = \frac{\tau}{\frac{\partial \tau}{\partial s_{coal}}} + s_{coal}, \tag{14}$$

and

$$\frac{\partial UC_{gas,r}}{\partial s_{coal}} = 0 \Leftrightarrow \sigma^*_{gas}(s_{coal}) = \frac{\tau \frac{\partial s_{gas}}{\partial s_{coal}}}{\frac{\partial \tau}{\partial s_{coal}}} + s_{gas}. \tag{15}$$

Unit costs will be increasing for plant i if its emission rate exceeds the corresponding iso-cost emission rate at a given level of differentiation for its fuel type. Otherwise, unit costs will be decreasing. Note that the iso-cost rates will be complex functions of the universe of plants and mitigation opportunities.

The relationship between the degree of differentiation, iso-cost emission rates, and plant-specific emission rates is demonstrated graphically in Figure 6. The graphic is labeled in the context of coal plants, though the analysis for gas plants is equivalent. The vertical axis measures emission rates, and the horizontal axis measures the coal standard, which serves as a proxy for the level of differentiation. The dashed line represents a hypothetical path for the iso-cost emission rate. Note that this rate is a function of the coal standard, and changes as the level of differentiation increases. Coal plants with high, medium, and low emission rates are represented by the solid horizontal lines labeled $\sigma_{coal,h}$, $\sigma_{coal,m}$, and $\sigma_{coal,l}$, respectively. For the high emission rate plant, unit costs are increasing regardless of the level of the coal standard. The opposite is true for low-emission rate plants. For medium-emission rate plants it is possible that unit costs will be increasing for some levels of the standard and decreasing for others. For those plants, the sign of the net effect on unit costs of a particular level of differentiation (relative to the single-standard) needs to be calculated by integrating the marginal effects over the less differentiated policies.

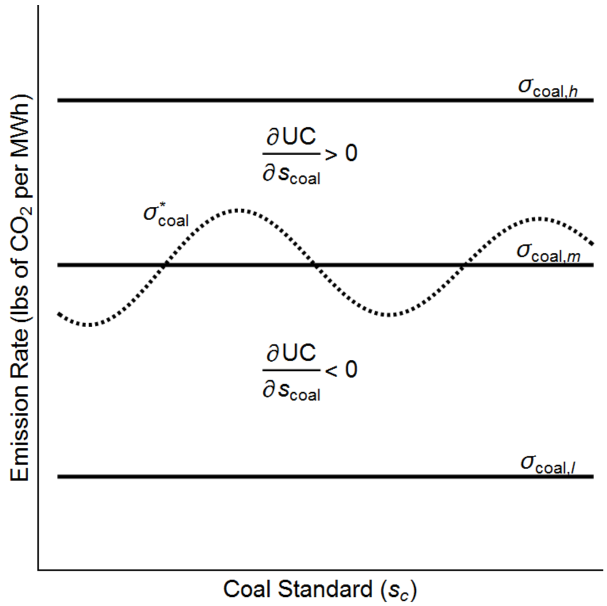

Figure 6: Hypothetical Iso-Cost Emission Rate Curve and Plant Emission Rates

5.3 Regional Prices

A fundamental feature of the U.S. electricity sector is regional segmentation. As a consequence, regulatory policies can have different impacts on different regions of the country. Analytically, the impact of differentiated standard policy on electricity prices in a given region is ambiguous. Furthermore, regional prices need not be monotonic in the degree of differentiation. Ultimately, it will depend largely on the joint distribution of the unit costs and emission rates for both gas and coal plants within the region. As I argue below, in order for prices to not rise in a region, there must be sufficient capacity with low emission rates (conditional on fuel type) to replace relatively low-cost, higher-rate generation.

There are two avenues through which differentiation can alter electricity prices in a region. First, it can influence the unit cost of the marginal plant in the region, as was discussed in the previous section. If the marginal plant has an emission rate below the iso-cost rate, then the wholesale price in the region will decrease, and vice versa for marginal plants above the iso-cost rate.

For small changes in the degree of differentiation, this is the mechanism that would be expected to drive price changes. For larger changes, it is more likely that the marginal plant itself would change. How the dispatch order changes will depend on the changes in the compliance costs of switching opportunities within the region of interest, as well as those same changes across the entire sector.

Figure 7 illustrates how electricity prices might change with further differentiation in a hypothetical region at a particular level of regulation. The vertical axes represent net unit costs of generation under a given performance standard policy, and the horizontal axes represent observed plant emission rates. Unit cost and emission rate pairs for specific plants are represented by the dots, and the horizontal solid line represents the electricity price at the current level of differentiation. This price is determined by the unit cost of the plant at the the margin. The vertical dashed black lines represent three

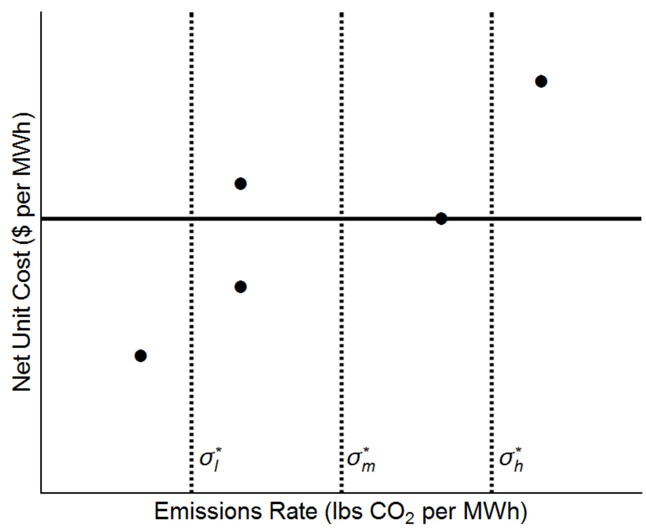

Figure 7: Potential Effects of Differentiation on Prices

possibilities for the iso-cost emission rates, all of which have different implications for the electricity price. For simplicity, assume that these rates remain constant as the level of differentiation increases – this is equivalent to forcing the dashed line in Figure 6 to be horizontal. Under this assumption, both gas and coal plants can be represented on the same graph by pegging their emissions rate to their respective iso-cost rates.[10] For plants with emission rates to the right of the dashed lines, unit costs increase with an increase in differentiation; for plants to the left, unit costs decrease.

Consider first the case in which iso-cost emission rates are depicted by the line labeled σ_h^*. Here all currently employed plants are to the left of the dashed line, meaning that their unit costs are decreasing in the coal standard. Therefore, prices will decrease monotonically with the degree of differentiation, even if the identity of the marginal plant changes. Conversely, suppose that the iso-cost emission rates are depicted by the σ_l^* line. In this setting, all units with decreasing unit costs are producing at capacity, but additional generation is required to meet demand. The additional generation must come from plants with increasing unit costs, and so prices must increase monotonically with the coal standard.

In the first two scenarios, the effect of differentiation on prices in the region was straightforward. Now suppose the σ_m^* line represents the iso-cost emission rate. In this case, a portion of currently employed capacity experiences increasing unit costs, while there is idle capacity with decreasing unit costs. The net effect on prices depends on the magnitude of the rates of change in unit prices, which will determine at what point the low emission rate plant replaces the high rate plant. This ambiguity can result in non-monotonic responses in price to the degree of differentiation. Non-monotonicities can also result from the fact that the dashed lines are likely to shift with the degree of differentiation, which I have assumed away in this diagram.

5.4 Coal Plant Profits

The impact of differentiation on coal plant profits depends on factors described above: net unit costs and regional electricity prices. Compared with the business-as-usual equilibrium, imposing a single standard policy would be expected to have negative consequences for coal plant profits as it will tax

[10] One way to accomplish this transformation would be to subtract $\sigma_{coal}^* - \sigma_{gas}^*$ from the coal plant emission rates. This generates a standardized emission rate for those plants that can be used to plot gas and coal plants on the same axes.

all coal generation without increasing electricity prices appreciably.[11] When the standards are differentiated, however, outcomes are less clear. Drawing on the preceding discussion of unit cost and price effects, I describe how differentiation influences plant profits at the individual unit, regional, and national levels.

Differentiation will tend to enhance the economic position, in terms of profits, of plants with low emission rates while negatively impacting those with higher rates. As demonstrated in Figure 6, coal plants with particularly low emission rates will enjoy growing cost benefits from differentiation. As the degree of differentiation becomes very large, it is conceivable that the plant economics can become even more favorable than in the case of no regulation for the cleanest of facilities. Whether this occurs or not depends also on the price impacts of differentiation. To this point, the actual impacts of differentiation on output prices have been acknowledged as ambiguous. In actuality, the simulation results will point to a nearly universal increase in prices, which will contribute further to the profit advantage of clean coal plants.

The aggregated impacts on coal plant profits both regionally and nationally are difficult to assess analytically, but the preceeding analysis suggests a few general themes. First, heterogeneity for aggregate coal profits across regions could be quite significant. For instance, consider a region with a large endowment of low emission rate coal capacity, but even greater electricity demand. If the resources available to meet that residual demand are high-emission rate natural gas plants, prices will increase while unit coal costs will decrease, leading to substantial profit improvements for coal plants in that region. On the other hand, imagine a region with cheap, high emission rate coal plants and expensive, low emission rate natural gas plants, with the latter on the margin. Under a single standard policy, the coal plants might still be utilized heavily if coal-to-gas switches in other regions are relatively cheaper. As demonstrated previously, differentiation will result in unit cost decreases for

[11]The muted impact of a single standard policy on prices, as opposed to, for instance, a cap-and-trade policy results from the implicit output subsidy established by the policy, as demonstrated by Fischer (2001).

the marginal, low emission rate gas units and increases for the high emission rate coal plants. The result will be decreasing electricity prices and increasing coal plant costs, and so regional coal profits will decrease as a result of differentiation. In a national market with both types of regions, as well as regions between these extremes, it is difficult to assess the aggregate effect on coal profits nationally; doing so will require implementation of a numerical model.

5.5 Cost-Effectiveness

While the primary focus of this paper is on distributional outcomes, I also discuss briefly aggregate cost-effectiveness.[12] It is well-established in the environmental regulation literature that tradable performance standards are second-best instruments for pollution abatement if demand has any price response (see, e.g., Fischer (2001) and Holland, Hughes and Knittel (2009)). This is because the implicit output subsidy that they establish results in less than optimal reductions through conservation; instead, an inefficient proportion of emissions reduction comes from supply-side measures such as fuel-switching. The lack of conservation under a tradable, single standard policy leaves room for improvements through differentiation, which I show here.

To demonstrate how differentiation can improve cost-effectiveness outcomes, I frame the problem in Stackelberg fashion. Taking the market equilibrium as given, the social welfare function is defined as a function of the coal standard, assuming that the gas standard is an implicit function of the coal standard:

$$W(s_{coal}) = \sum_r \left[\int_0^{\sum_i \sum_f q^i_{f,r}(s_{coal})} P_r(u)\mathrm{d}u - \sum_i \sum_f c^i_{f,r} \cdot q^i_{f,r}(s_{coal}) \right]. \tag{16}$$

The question is what happens to welfare as the coal standard is increased by a small amount. Taking the derivative yields

$$W'(s_{coal}) = \sum_r \sum_i \sum_f (P_r(q_r) - c^i_{f,r} - \lambda^i_{f,r}) \cdot \frac{\partial q^i_{f,r}}{\partial s_{coal}}, \tag{17}$$

[12] See Lemoine (2014) for a more thorough treatment.

where $q_r = \sum_i q^i_{gas,r}(s_{coal}) + \sum_i q^i_{coal,r}(s_{coal})$. For small changes in the standard, only the marginal units will respond, meaning that $\partial q^i_{f,r}/\partial s_{coal} = 0$ for all units such that $\lambda^i_{f,r} > 0$. Removing those units and plugging in the first-order conditions for the marginal generators yields

$$\begin{aligned} W'(s_{coal}) &= \sum_r \sum_i \sum_f \tau(\sigma^i_{f,r} - s_f) \cdot \frac{\partial q^i_{f,r}}{\partial s_{coal}} \\ &= \tau\left(\sum_r \sum_i \sum_f \sigma^i_{f,r} \cdot \frac{\partial q^i_{f,r}}{\partial s_{coal}} - \sum_r \sum_i \sum_f s_f \cdot \frac{\partial q^i_{f,r}}{\partial s_{coal}} \right). \end{aligned} \tag{18}$$

Note that the first term in parentheses is the change in emissions in response to a change in the standard. Assuming, as has been the case to this point, that the coal and gas standards are changed in such a way that emissions are held constant, that term drops out of the equation, leaving

$$W'(s_{coal}) = -\tau\left(\sum_r \sum_i \sum_f s_f \cdot \frac{\partial q^i_{f,r}}{\partial s_{coal}} \right). \tag{19}$$

Solving at the single-standard rate such that $s_f = s$ for all f, the equation reduces further to

$$W'(s_{coal}) = -\tau \cdot s\left(\sum_r \sum_i \sum_f \frac{\partial q^i_{f,r}}{\partial s_{coal}} \right). \tag{20}$$

The summation term in the above equation is simply the change in aggregate generation. This implies that the change in welfare due to an incremental increase of the coal standard at the single standard level is a decreasing function of the change in output. Therefore, the initial effect of differentiation away from a uniform standard can yield increases in welfare if and only if it induces an increase in conservation. Accordingly, if demand is assumed to be perfectly inelastic, then aggregate generation is fixed and differentiated policies are less cost-effective than the analogous single standard.

6 Simulation Model

Here I describe the construction of my detailed electricity sector model, and how I use it to simulate a suite of performance standard policies. The simulation model allows me to calibrate the policies to forecasted supply and demand and quantify their impacts. Several data sources, each with relative strengths and weaknesses, are used in the parameterization of the model in an effort to mirror actual conditions in 2017.

The simulation model structure is essentially a more detailed version of the model presented in Section 4; see the Appendix for model formulation details. Electricity can now be transferred across regions (still denoted by r) to meet demand. Due to the non-storable nature of generated electricity, the market is also broken into several time segments, reflecting both the necessity for supply to meet demand instantaneously as well as the seasonal and daily fluctuations in demand. Unlike the analytic model, the simulation model also accounts for electricity transfers across regions, which are subject to transmission losses, wheeling costs, and transmission constraints. Finally, in addition to a capacity constraint within time segments, generation plants also have a seasonal availability constraint to reflect forced outages for maintenance.

6.1 Data Sources and Model Parameterization

In this section, I describe the data and assumptions utilized in my simulation model. The model is constructed to provide a detailed yet transparent exploration into static performance standard policies for CO_2. Overall, the model structure is similar in spirit to the dispatch module of several other electricity market models, such as the U.S. Energy Information Administration's National Energy Modeling System and ICF International's IPM (U.S. Energy Information Administration, 2013).

6.1.1 Electricity Demand

Demand for electricity in the model is adapted from EPA Base Case v.5.13 (U.S. Environmental Protection Agency, 2013). Each region has an hourly demand profile, also referred to as a load curve, that chronologically specifies a demand level of electricity for each hour in the year. For computational reasons, I first sort the load curve (within seasons) from highest to lowest demand, and then I aggregate those profiles into representative time segments. Demand is split into eight time segments within two seasons ("Summer" and "Winter"), meaning that there are 16 total segments. The segments are chosen that there are fewer hours at the high and low ends and more hours in the intermediate sections of the curve. Hourly demand within each segment is summed for each region. The result of this procedure are business-as-usual demand levels for each region and time segment.

For the coal usage, electricity price, and coal plant profit results reported in the next section, I maintain an assumption of inelastic demand for three reasons. First, since I am focusing on a short-run analysis, the price response from consumers is expected to be small. Second, the output subsidy inherent in a performance standard policy greatly dampens the impact of conservation as a compliance mechanism. Finally, the inelastic demand assumption makes it easier to track and quantify the primary emission reduction margins of interest, coal-to-gas switching and within-fuel switching.

For the cost-effectiveness analysis, I report results for the inelastic demand case, as well as two elastic demand cases. Consistent with the short-run analysis, I keep the elasticity parameter small ($\epsilon = -0.1$ and $\epsilon = -0.2$). The elasticity is identical across all regions and time segments, and is combined with the price and quantity levels at business-as-usual to parameterize a linear demand function. For more details, see the Appendix.

6.1.2 Unit-Specific Attributes

Electricity supply is modeled for individual generating units, which I refer to as "plants" for simplicity. This is possible because of the availability of specific data collected by the EIA, EPA, and U.S. Federal Energy Regulatory Commission. The key characteristics of each plant are the region, fuel type, capacity, availability, and variable operations and maintenance (VOM) costs. Fossil plants also have heat rates, which are a measure of efficiency, and emission factors. Below I describe where these data are obtained, and how they are incorporated into the model; see the Appendix for details.

Data for basic plant characteristics come from a variety of sources. The overall universe of generating units comes from SNL Financial. Units that are expected to retire prior to 2017 are omitted from the model, while units that are characterized as being "under construction" with an in-service date prior to 2017 are included. Fuel type and capacity data is also extracted from SNL Financial. The set of units is merged with the EPA's NEEDS v.5.13 database to place each unit in a region; unmatched units are placed into regions based on additional grid interconnection data from SNL Financial (U.S. Environmental Protection Agency, 2013). Finally, unit availability data is extracted from the NEEDS database.

Individual plant costs and plant emissions are assumed to be linear for all levels of generation, and are constructed using heat rate, fuel price, VOM cost, and emission factor data. Heat rates and VOM costs are taken from SNL Financial data whenever possible; otherwise, they are extracted from the NEEDS database. Fuel prices are collected from a combination of EPA Base Case v.5.13 and the 2014 North American Database for EPIS's Aurora software. Emission factors are extracted from SNL Financial when possible; otherwise, an average factor based on the unit's fuel type is used.

6.1.3 Transmission Network

A key feature of the U.S. electricity sector is the interconnectness of the transmission grid. Using data from EPA Base Case v.5.13, I divide the lower 48 U.S. states into 66 model regions. Sixty-one of these regions are power market regions which have a positive demand for electricity. The remaining five regions are power-switching regions for which demand is zero. These regions are used to capture units with the ability to provide electricity to multiple regions. Transmission flows across regions are subject to capacity constraints, wheeling charges, and transmission losses; all of these parameters are supplied by EPA Base Case v.5.13.

7 Simulation Results and Discussion

In this section, I describe the results of several different tradable performance standard policies. The results suggest that, as expected from the analytic model, implementing a tradable standard policy causes coal usage to decline significantly. Relative to a single standard baseline, differentiation leads to a modest increase in coal generation and usage. Under a single standard, electricity prices either decrease or increase depending on the region. Differentiation from the single standard causes prices to rise in almost every region. While single-standard policies result in substantial profit gains for natural gas plants and losses for coal units, differentiation leads to improved profit outcomes for low emission rate units in both categories. Finally, differentiation is slightly more cost-effective for small degrees, but it becomes significantly more costly than the single standard for large degrees.

Using the model described in the previous section, I simulate generation for the U.S. wholesale electricity market under a variety of tradable performance standard policies. I establish two reference cases: a business-as-usual scenario without an emissions constraint and a single standard scenario that achieves a given emission reduction target. The differentiated-standard policies are calibrated to achieve the same target. Since the vast majority of switching behavior is driven by the difference between the coal and gas standards, the various differentiated-standard policies are reported in terms of this parameter, i.e., $\Delta \equiv s_{coal} - s_{gas}$. Therefore, single-standard policies will correspond to the case where $\Delta = 0$.

The model can be simulated for any degree of differentiation and policy stringency; I illustrate a few examples in Figure 8. In the figure, I plot fuel-specific, capacity-weighted kernel density estimates of emission rates for all natural gas and coal units in the model. The red distribution is natural gas, and the blue distribution is coal. The vertical dashed lines represent various performance standards. The black line is placed at the single-standard level for a 5% emissions reduction. For the most part, this line divides the coal and gas distributions. Under this regulation, coal-fired power plants have higher unit costs and gas plants have lower unit costs. The green lines represent the differentiated coal and gas standards when $\Delta = 600$. This captures a substantial level of differentiation where production from the lowest-emitting coal plants is actually subsidized and have lower unit costs compared to BAU. The orange lines represent an even more extreme degree of differentiation ($\Delta = 1200$) for which a significant proportion of the coal units are being subsidized. Correspondingly, a large measure of natural gas units are taxed at this degree. Beyond this degree of differentiation, the emission reduction target cannot be achieved.

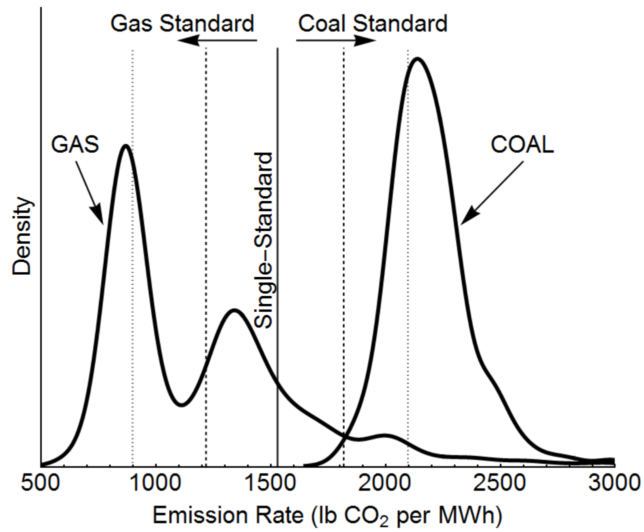

Figure 8: Illustrating Differentiated-Standard Policies

I report results under each policy scenario for coal usage, electricity prices, and coal generator profits. The results are simulated at a fairly specific level of geographic detail, and some aggregation is necessary. The model contains 61 market regions; I aggregate the model regions into 17 "reporting" regions on the basis of the electricity transmission system architecture. In the context of the model, market regions among which large volumes of electricity is transferred tend to be grouped together. The reporting regions are illustrated in Figure 9. I also aggregate across time segments in the main results. This means that, in the case of electricity prices, that the reported value is a weighted average of prices within each time segment and reporting region. Because there are 16 time segments for each market region, if a reporting region consists of 4 market regions, then the given price is the demand-weighted average of $4 \times 16 = 64$ equilibrium prices. My aggregation strategy attempts to capture the most relevant trends in price changes across regions at the time segment-level.

Figure 9: Reporting Region Map

The main results are summarized in the next three subsections. In addition to simply describing the key results, I explain the underlying economic factors that drive them. In doing so, I rely on the intuition developed in the analytic sections.

7.1 Coal Usage

As a measure of the impact on upstream coal interests, most notably coal producers and laborers, I present aggregate results for the change in coal usage as the coal and gas standards are increasingly differentiated. I then explain the simulation outcomes in the context of the analytic discussion of switching margins.

The primary coal usage result is captured in Figure 10. Unlike in the case of electricity prices and generator profits, I emphasize the impact on coal usage at the national rather than the regional level. This is because the origin of the fuel is not necessarily confined to the region in which it is produced. Coal usage has been standardized to their percentages of business-as-usual levels. Both series are plotted on the vertical axis against the corresponding degree of differentiation.

The graph shows that differentiation increases coal usage by a modest

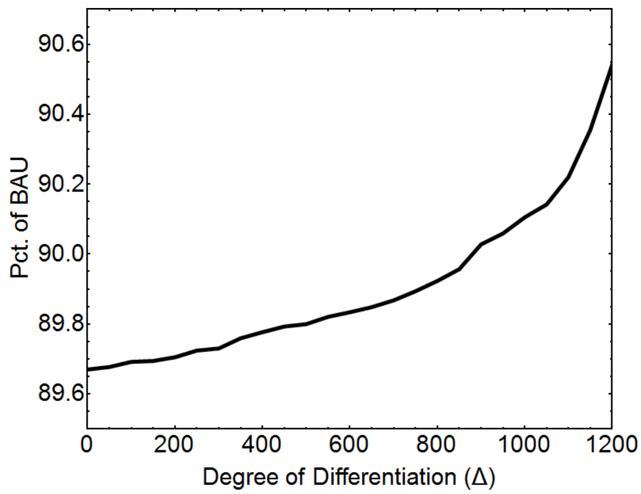

Figure 10: Coal Generation and Usage with Differentiation

amount. As expected, implementing the single standard causes usage to drop; for my particular emissions reduction target, usage is reduced by more than 10% from business-as-usual, which is demonstrated by the level of the curve when $\Delta = 0$. The usage curve increasing and largely convex as the degree of differentiation increases. As a result, substantial increases don't occur until the standard is highly differentiated. However, even for the most extreme differentiation, the increase in usage from the single-standard level is relatively small (less than 1% relative to business-as-usual).

The positive relationship between degree of differentiation and coal usage was predicted in Section 5. Due to rising costs of coal-to-gas switches in response to increased differentiation, it was shown that less coal would be displaced in favor of gas generation for a given emissions target. This is illustrated for the simulation results in Figure 11. In the single standard case, illustrated by the inner-most pie chart, nearly all emission reductions come from coal-to-gas switching. A small slice is the result of coal-to-coal switching.[13] Moving outward, under a moderate level of differentiation ($\Delta = 600$) a small portion of the reductions achieved through coal-to-gas switches is dis-

[13] The "Other" category results from replacing coal generation with dispatchable biomass-fueled generation. I assume that biomass plants, which contribute a very small fraction of U.S. electricity generation, are not covered by the emission regulation.

placed by additional coal-to-coal switching. This is consistent with Figure 10, which shows that coal usages increase at an extremely slow rate up for small and modest levels of differentiation. For extreme differentiation, illustrated by the chart labeled as $\Delta = 1200$, the amount of coal-to-coal switching increases significantly, and some reductions are achieved by gas-to-gas switching. From the aggregate perspective, however, the overwhelming majority of emission reductions are the result of coal-to-gas switches even for the most differentiated policy. This explains why even though coal usage increases with differentiation compared to a single standard, the extent is modest compared to the original decline associated with the single standard.

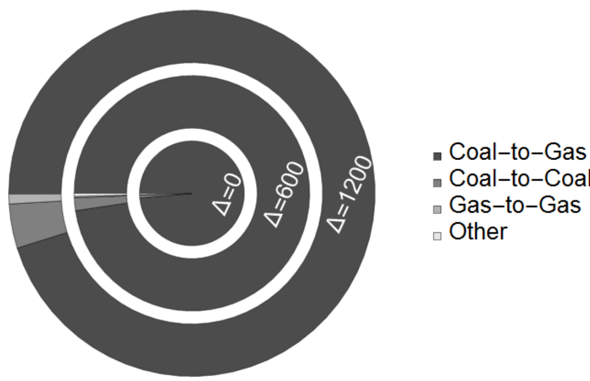

Figure 11: Emission Reductions by Switching Type

7.2 Per Unit Costs

In order to explain the output price and coal plant profit results, I first demonstrate the effect of differentiation on unit costs in the context of the policy simulations. More than half of the utilized coal capacity realizes decreasing costs for almost all differentiation degrees, while most natural gas capacity experiences increasing costs.

As discussed in Section 4, the effect of differentiation on per unit cost for any particular plant depends on the plant emission rate and the behavior of the credit price as well as the relationship between the coal and gas standard. The empirical versions of $\tau(s_{coal})$ and $s_{gas}(s_{coal})$ are used to construct the iso-cost emission rates for the simulation model; the results are plotted in Figure 12.[14] The left panel pertains to coal plants, while the gas results are depicted in the right panel. The paths of the iso-cost emission rates are given by the dashed lines. Close to the single-standard level, the paths fluctuate greatly, but the paths smooth out for greater degrees of differentiation. Roughly speaking, the iso-cost emission rate functions for gas and coal hover near 800 and 2,200 lbs. of CO_2 per MWh.

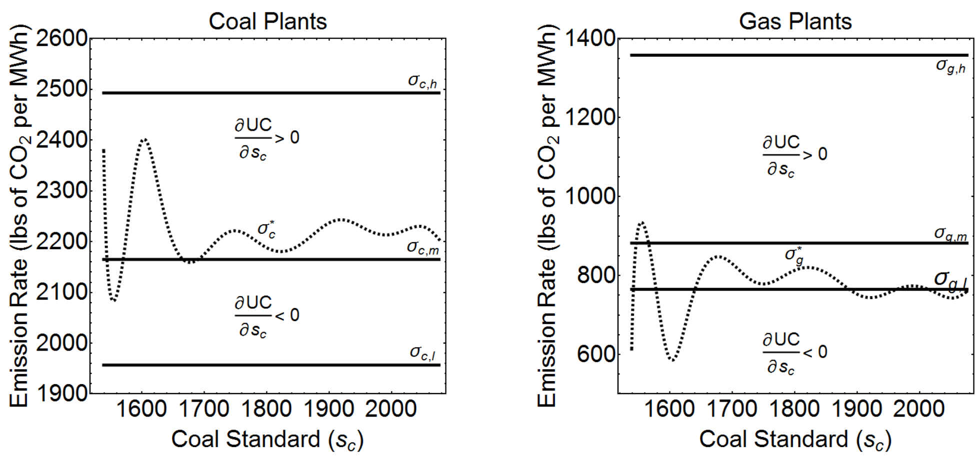

Figure 12: Empirical Iso-Cost Emission Rate Curves

To get a sense for how much fuel-specific capacity is above and below the iso-cost emission rates, the capacity-weighted 5$^{\text{th}}$, 50$^{\text{th}}$, and 95$^{\text{th}}$ emission rate percentiles for both fuels are imposed on the graphs. The 95$^{\text{th}}$ percentile emission rates, represented by the lines labeled $\sigma_{f,h}$, are well above the dashed lines for both coal and gas; costs for plants with emission rates in this neighborhood will be increasing for all differentiation degrees. At the clean end of the distribution, represented by the lines labeled $\sigma_{f,l}$, coal plant costs are decreasing regardless of the coal standard. Contrast this to the gas diagram, which shows that the iso-cost emission rate is largely straddling the 5$^{\text{th}}$ percentile emission rate. This suggests that even for the cleanest gas plants, cost movement is roughly neutral as the standards are increasingly differentiated.

[14]See the Appendix for a graphical representation of the credit price and gas standard as a function of the coal standard.

To get a sense for average behavior, the median emission rate plants are represented by the lines labeled $\sigma_{f,m}$. In the case of coal, costs are decreasing for almost all differentiation degrees. The opposite is true for natural gas. Consequently, on an average basis, this graphic reveals that coal plants benefit more from differentiation than natural gas plants in terms of generation costs.

7.3 Prices

In this section, I report and explain the results for the effect of differentiation on regional wholesale electricity prices. Differentiation causes prices to increase relative to the single-standard price in almost every region, including those that rely heavily on coal for generation. The price increases are the result of increasing unit costs for plants near the margin (both coal and gas) in most regions.

The main results for the policy impacts on electricity prices are given in Table 1. These impacts are reported as percentage deviations in the columns labeled SS, DS1, and DS2, which correspond to single standard, moderately differentiated ($\Delta = 600$), and extremely differentiated policies ($\Delta = 1200$), respectively. For the single standard, the reported percentage deviation in prices is from business-as-usual. For the differentiated standards, the reported percentage deviation in prices is from the single standard. Therefore, adding one of the differentiated standard columns to the single standard column yields the percent change from business-as-usual prices for the differentiated standard policy.

When the single-standard policy is implemented, electricity prices in most regions decrease. In extreme examples, like Texas and the Pacific Northwest, they decrease by 5-6%. This is because electricity prices are set by the unit generation costs of marginal plants, which tend to be gas plants. Under a single standard policy, almost all gas generation is subsidized, resulting in lower unit generation costs. This not universally true, however, as prices increase in other (mostly coal-heavy) regions such as the Midwest-West and Midwest-East. In terms of magnitude, most regions experience price changes, either negative or positive, in the range of 2-4%. For the entire country, aver-

Region	Size TWh	%Coal	BAU $/MWh	SS $\frac{P_{ss}-P_{bau}}{P_{bau}}$	DS1 $\frac{P_{ds}-P_{ss}}{P_{bau}}$	DS2 $\frac{P_{ds}-P_{ss}}{P_{bau}}$
Texas	332	44%	34	-4.9%	0.9%	8.7%
Florida	230	3%	38	-3.1%	0.7%	3.9%
Midwest-North	176	60%	32	-2.4%	0.8%	5.9%
Midwest-West	98	70%	29	4.2%	0.2%	3.4%
Midwest-East	409	64%	33	4.2%	0.3%	3.0%
New England	128	3%	44	-2.7%	0.2%	4.5%
New York	149	4%	45	-2.0%	0.5%	4.3%
Mid-Atlantic	708	47%	38	-1.2%	0.5%	4.2%
Tennessee Valley	231	43%	35	2.7%	0.2%	1.6%
Delta	103	28%	37	-1.8%	0.7%	5.4%
Southeast	256	21%	39	-2.0%	0.6%	0.2%
Carolinas	229	27%	43	3.3%	-0.2%	-3.3%
Heartland	263	64%	33	-3.2%	0.9%	6.7%
California	279	1%	42	-4.8%	0.6%	4.7%
Southwest	197	47%	34	-4.2%	0.8%	5.5%
Rocky Mountains	71	72%	31	0.4%	-0.4%	2.1%
Pacific Northwest	191	17%	32	-6.0%	0.6%	4.6%
United States	**4,048**	**39%**	**37**	**-1.4%**	**0.5%**	**3.8%**

Table 1: Regional Price Impacts of Performance Standard Policies

age prices decrease modestly.

The magnitude of the price impacts of differentiation depend on the degree, but the direction is almost uniformly positive. For the moderately differentiated policy, prices increase from the single-standard level in almost every region; the exceptions are the Carolinas and Rocky Mountain regions. In all cases, the price change is less than a percentage point, including the national average. For the larger degree of differentiation, the price changes are positive in all regions except for the Carolinas. Note that prices in the Rocky Mountain region are not monotonic in the degree of differentiation, a possibility we highlighted in the analytic discussion. In terms of magnitude, the impacts are much larger in almost all regions (the Southeast region is the lone exception). In most cases, the price increases by 3-5%; however, Texas experiences an increase of almost 9% relative to the single-standard level. Moreover, overall price changes under extreme differentiation relative to business-as-usual, which can be determined by summing the BAU and DS2 columns, are in most cases positive with a national average increase of just under 4%.

As demonstrated in the analytic section, the effect of differentiation on the electricity price in a particular region depends on the unit cost behavior of plants near the margin. In most regions, unit costs increase for the marginal plants, regardless of whether they are gas or coal. This occurs because overall per unit costs tend to be correlated with emission rates so higher emission rate plants tend to be on the margin. The lone exception occurs in the Carolinas. I next demonstrate why this region experiences falling prices by contrasting it with a more "typical" region (Texas).

The price impacts of differentiation for the Carolinas (right panel) and Texas (left panel) are demonstrated graphically in Figure 13. The vertical axes represent per unit generation costs at the single standard equilibrium, while the horizontal axes represent emission rates standardized by the fuel-specific iso-emission rates, as was the case in Figure 7.[15] This allows the gas (gray markers) and coal (black markers) plants to be combined on a single graph; the size of the markers reflects the plant capacity, with larger markers representing larger plants. The solid horizontal lines denote the market-clearing electricity price in each region at the single-standard. The vertical dashed lines split the units into two categories: plants to the left have costs that are decreasing with differentiation, while plants to the right have increasing costs. In the Carolinas, all plants near the margin have decreasing unit costs. This occurs because coal plants in the region are relatively efficient, and hence have low emission rates, but fuel costs for importing coal into the region are relatively high. As a result, electricity prices decrease as the standards are differentiated. However, this scenario is atypical. As Table 1 demonstrated, prices generally increase with differentiation. This is what occurs in Texas,

[15]Recall that iso-cost emission rates are a function of the level of the coal standard, or equivalently, the degree of differentiation. To construct this diagram, I use the iso-cost emission rates corresponding to the case where $\Delta = 1200$. For gas and coal, this implies rates of 767 and 2,212 lbs. of CO_2 per MWh. As demonstrated in Figure 12, these rates are not constant for all differentiation degrees, but they serve as a useful proxy for the purposes of illustration.

where most of the units at the margin have increasing costs.

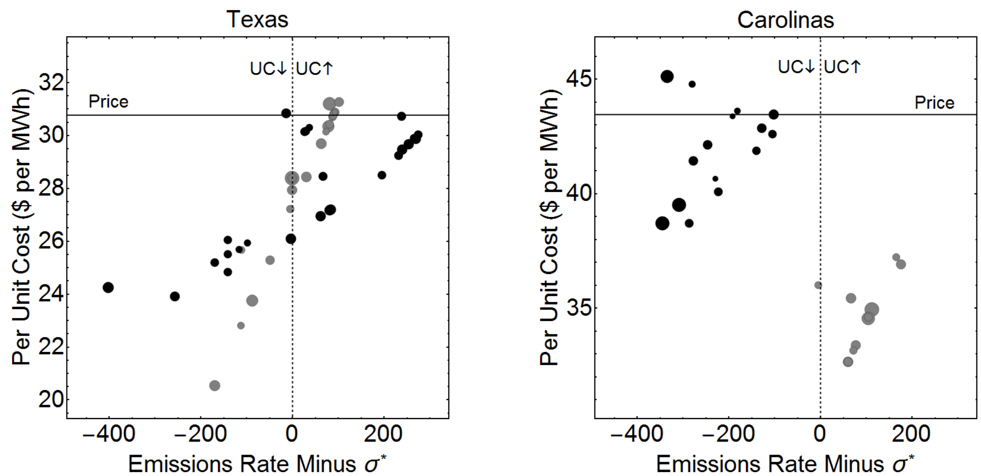

Figure 13: Texas versus Carolinas: Price Impacts of Differentiation

Figure 13 also demonstrates the typical relationship between unit costs and emission rates. Within fuels, these parameters are positively correlated. As a result, plants near the margin are more likely to have costs that increase with differentiation, regardless of whether those plants are coal or gas. This effect should be even more pronounced for higher levels of demand, as the additional plants brought on-line will tend to have greater emission rates. By looking at summer peak demand prices, I demonstrate that this occurs in the context of the simulation model.

The increase in average prices with differentiation was demonstrated in Table 1; I use Figure 14 to convey the influence of high demand time periods on this increase. In the graphic, the horizontal axis denotes the summer peak price difference under extreme differentiation ($\Delta = 1200$) versus the single-standard policy, i.e., $P_{ds} - P_{ss}$. The plot is the density for the summer peak price difference across all regions. The graphic shows that the peak price difference is typically well-above the average price increase under differentiation. This supports the notion that higher emission rate units are employed to meet higher levels of demand. The result is that prices tend to increase with differentiation, regardless of whether coal or gas is on the margin.

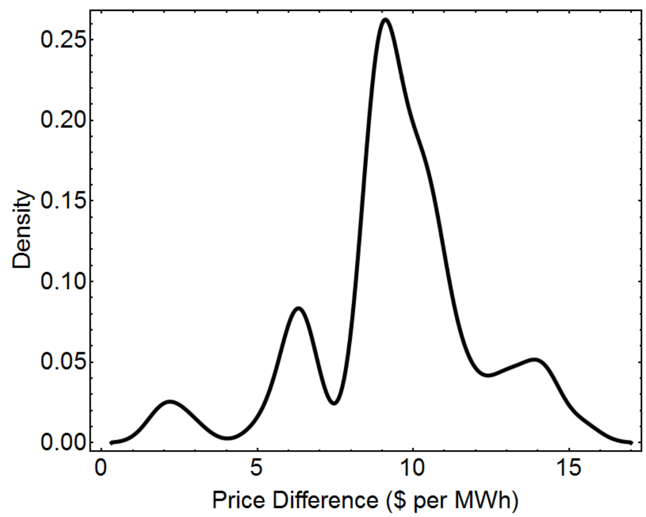

Figure 14: Comparing Summer Peak Prices

7.4 Coal Profits

The previous section documented the adverse consequences of differentiation on electricity prices, a proxy for effects on consumers. In this section, I report the positive effect that differentiation has on coal plant profits. This is due to a combination of higher electricity prices (discussed in Section 7.3) and reduced unit generation costs for coal plants (discussed in Section 7.2).

The main profit results are recorded in Table 2. Unlike the price table, all results in this table are given in levels, including the differences. The units for profits are millions of dollars ($M). Analogous to the reporting for price impacts, the column labeled SS denotes the difference between coal profits under the single-standard policy versus the business-as-usual scenario. The columns labeled DS1 and DS2 report the additional difference between the differentiated standard policies and the single standard policy. The sum of either of these columns and the SS column yields the coal unit profit deviation for the differentiated policy relative to the business-as-usual level.

Not surprisingly, under a single standard policy coal unit profits decrease substantially as expected as coal plant unit costs increase and electricity prices largely decrease. Heavy coal generation regions such as the Mid-Atlantic, Texas, Heartland, and the Southwest experience the largest absolute decreases

Region	Size TWh	%Coal	BAU \$M	SS $\Pi_{ss} - \Pi_{bau}$	DS1 $\Pi_{ds} - \Pi_{ss}$	DS2 $\Pi_{ds} - \Pi_{ss}$
Texas	332	44%	1,101	-573	22	285
Florida	230	3%	24	-5	4	23
Midwest-North	176	60%	663	-290	16	119
Midwest-West	98	70%	648	-73	12	129
Midwest-East	409	64%	1,310	-192	51	452
New England	128	3%	11	-8	2	17
New York	149	4%	23	-13	2	16
Mid-Atlantic	708	47%	2,218	-655	148	1,137
Tennessee Valley	231	43%	552	-91	22	180
Delta	103	28%	117	-65	2	15
Southeast	256	21%	558	-143	16	48
Carolinas	229	27%	229	-9	31	171
Heartland	263	64%	1,602	-544	46	376
California	279	1%	15	-6	1	6
Southwest	197	47%	912	-319	40	282
Rocky Mountains	71	72%	493	-118	2	112
Pacific Northwest	191	17%	244	-116	2	25
United States	**4,048**	**39%**	**10,723**	**-3,218**	**418**	**3,407**

Table 2: Regional Profit Impacts of Performance Standard Policies

in profits. On a percentage basis, coal units in Texas suffer the most, with region-wide profits decreasing by more than half. Compare this, for example, to the relatively mild decrease in profits within the Midwest-East. A large portion of this disparity is due to the disparity in price changes that were reported in Table 1.

Differentiating the standards improves regional coal unit profits across the board, as electricity prices generally rise and coal unit generation costs fall for most capacity. For the lower degree of differentiation, coal unit profits increase in all regions, but by small amounts relative to the lost profits from implementing the single standard. The Carolinas is the only region in which profits increase beyond the single standard level. For a large degree of differentiation, regional coal unit profits increase dramatically in all regions. Relative to business-as-usual profits, aggregate profits in many regions actually increase. On a percentage basis, the largest increase occurs in the Carolinas, where profits are nearly doubled. In Texas and the Midwest-North, however, large profit losses brought on by regulation are not recouped through differentiation. On a nationwide-basis, profits increase beyond business-as-usual levels.

7.5 Cost-Effectiveness

In this section, I report the effect that differentiation has on aggregate cost-effectiveness. Three different cases are considered: one in which demand is perfectly inelastic, and two with relatively small demand elasticities. In both cases, the most extreme level of differentiation ($\Delta = 1200$) increases overall program costs significantly (57 to 79 percent).

The program cost results are displayed in Figure 15. The metric used to evaluate aggregate cost (plotted on the vertical axis) is the additional cost introduced by differentiation relative to the cost imposed by the single standard policy. Mathematically, I calculate

$$Add.\ Cost_{ds} = \frac{Cost_{ds} - Cost_{ss}}{Cost_{ss} - Cost_{bau}}, \tag{21}$$

which provides a unit-free measure of how much cost distortion is introduced by differentiation. The additional cost value is plotted against the level of differentiation under three different assumptions about the price elasticity of demand. The solid line represents the simulation results when demand is perfectly inelastic, as labeled on the figure. The demand elasticity is relaxed as the lines move in a southeast direction across the figure. The dashed line represents a national demand elasticity of $\epsilon = -0.1$, and the dotted line represents a national demand elasticity of $\epsilon = -0.2$

In general, differentiation increases the costs of the performance standard policy, for all versions of the simulation model. As demonstrated in the analytic section, differentiation can lower program costs compared to a single-standard policy, but only when it induces additional conservation. When demand is perfectly inelastic, the conservation margin is eliminated, resulting in necessarily higher costs. In Figure 15, the solid line shows that the simulation results match the theory. Moreover, when demand elasticity is relaxed, the additional induced conservation offsets to some degree the additional costs of differentiation. In fact, although it is difficult to see in the figure, the cost-minimizing version of the policy when demand is elastic is slightly differentiated. However, to keep the model realistic, even the most demand-elastic version of the model is parameterized with very low values of elasticity, and so the improvement induced by differentiation is extremely small.

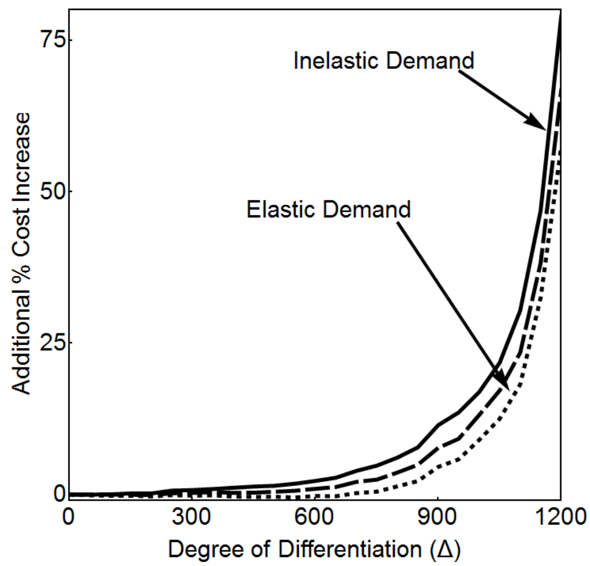

Figure 15: Cost-Effectiveness With Inelastic and Elastic Demand

8 Conclusion

Recent climate change policy discussions, both in the United States and abroad, have underscored the importance of accounting for distributional impacts in proposals. Economists have long touted the virtue of cost-effectiveness in policy, and for good reason, but often equity, legal, and political constraints necessitate deviations from first-best solutions. Ultimately, when it comes to emissions policy, sacrificing a certain degree of cost-effectiveness is preferable to failing to address the issue entirely.

This paper studies a key design feature in a tradable performance standard policy aimed at addressing potential inequities between coal and natural gas interests. By relaxing the coal standard while simultaneously tightening the natural gas standard, I analyze a continuum of differentiated policies aimed at achieving a fixed emission target while mitigating windfall welfare transfers from upstream, midstream, and downstream coal-oriented stakeholders to their natural gas counterparts. I examine the effects of differentiation through both analytic and simulation models with the same essential structure.

In the analytic model, I show that coal producers and laborers benefit from increased coal usage under differentiated policies, while the impacts on electricity consumers and coal-fired power plants through prices and profits, respectively, are ambiguous. The increase in coal usage is the result of a change in the relative cost of two compliance margins: within-fuel switching becomes less expensive relative to coal-to-gas switching under a differentiated policy. The ambiguity of the price and profit outcomes is shown to be the result of increasing compliance credit prices coupled with decreasing credit purchases/sales for coal-/natural gas-fired power plants.

To develop a quantitative sense of the impact of differentiation, I construct and implement a detailed short-run dispatch model of generation in the U.S. wholesale electricity market. The results suggest that differentiation does little to alter outcomes until the distance between the coal and gas standards is significant. However, once the standards are sufficiently differentiated, coal-fired power plants benefit greatly at the expense of electricity consumers. Increasing electricity prices, coupled with decreasing per unit costs for most coal capacity, results in aggregate profits beyond business-as-usual levels. For coal producers, the benefit to differentiation remains modest at any level of differentiation, as they regain less than one percent of the 10 percent market decrease they observed under the single standard policy.

The results suggest that if the purpose of differentiation is to aid coal producers and electricity consumers, then it is a poor mechanism. Going forward, it would be interesting to use the framework developed in this paper to compare to other versions of differentiated policy. For instance, one might consider how the results change when differentiation is made on the basis of states, rather than fuels, as suggested by both the initial proposal and finalized version of the CPP. Differentiation could also be studied outside of the electricity context: for instance, CAFE standards already differentiate on the basis of vehicle class. There is potential for interesting research on differentiated standards in automobile markets, since (1) they already exist; and (2) electricity and automobile markets differ in interesting ways.

References

Aldy, Joseph E. 2011. "Promoting Clean Energy in the American Power Sector." The Hamilton Project Discussion Paper 2011-04.

Bohringer, Christoph, and Andreas Lange. 2005. "Economic Implications of Alternative Allocation Schemes for Emission Allowances." *Scandinavian Journal of Economics*, 107(3): 563–581.

Bovenberg, A. Lans, Lawrence H. Goulder, and Derek J. Gurneyl. 2005. "Efficiency Costs of Meeting Industry-Distributional Constraints Under Environmental Permits and Taxes." *The RAND Jounal of Economics*, 36(4): 951–971.

Burtraw, Dallas, Arthur G. Fraas, and Nathan Richardson. 2012. "Tradable Standards for Clean Air Act Carbon Policy." *Resources for the Future Discussion Paper*, 12(05).

Bushnell, James, and Yihsu Chen. 2012. "Allocation and Leakage in Regional Cap-and-Trade Markets for CO_2." *Resource and Energy Economics*, 34: 647–668.

Bushnell, James B., Stephen P. Holland, Jonathan E. Hughes, and Christopher R. Knittel. 2015. "Strategic Policy Choice in State-Level Regulation: The EPA's Clean Power Plan." *NBER Working Paper Series*, 21259.

Fischer, Carolyn. 2001. "Rebating Environmental Policy Revenues: Output-Based Allocations and Tradable Performance Standards." *Resources for the Future Discussion Paper*, 01(22).

Fischer, Carolyn. 2003. "Combining Rate-Based and Cap-and-Trade Emissions Policies." *Climate Policy*, 3(2): S89–S103.

Fischer, Carolyn. 2011. "Market Power and Output-Based Refunding of Environmental Policy Revenues." *Resource and Energy Economics*, 33: 212–230.

Fischer, Carolyn, and Alan Fox. 2007. "Output-Based Allocation of Emissions Permits for Mitigating Tax and Trade Interactions." *Land Economics*, 83(4): 575–599.

Fischer, Carolyn, and Richard G. Newell. 2008. "Environmental and Technology Policies for Climate Mitigation." *Journal of Environmental Economics and Management*, 55: 142–162.

Fullerton, Don, and Garth Heutel. 2007. "The General Equilibrium Incidence of Environmental Taxes." *Journal of Public Economics*, 91: 571–591.

Helfand, Gloria E. 1991. "Standards versus Standards: The Effects of Different Pollution Restrictions." *The American Economic Review*, 81(3): 622–634.

Holladay, J. Scott, and Jacob LaRiviere. 2013. "The Effect of Abundant Natural Gas on Air Pollution from Electricity Production." Working Paper.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. "Greenhouse Gas Reductions under Low Carbon Fuel Standards?" *American Economic Journal: Economic Policy*, 1(1): 106–146.

Ito, Koichiro, and James M. Sallee. 2014. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards." *NBER Working Paper Series*, 20500.

Lemoine, Derek. 2014. "Escape from Third-Best: Rating Emissions for Intensity Standards." *University of Arizona Department of Economics Working Paper*, 12(03).

Linn, Joshua, Erin Mastrangelo, and Dallas Burtraw. 2014. "Regulating Greenhouse Gases in the Electricity Sector under the Clean Air Act." *Journal of the Association for Environmental and Resource Economics*, 1(1/2): 97–134.

Paul, Anthony, Karen Palmer, and Matt Woerman. 2013*a*. "Analysis of the Bingaman Clean Energy Standard Proposal." *Resource and Energy Economics*, 36: 113–129.

Paul, Anthony, Karen Palmer, and Matt Woerman. 2013*b*. "Modeling a Clean Energy Standard for Electricity: Policy Design Implications for Emissions, Supply, Prices, and Regions." *Energy Economics*, 36: 108–124.

Rausch, Sebastian, and Matthew Mowers. 2014. "Distributional and Efficiency Impacts of Clean and Renewable Energy Standards for Electricity." *Resource and Energy Economics*, 36: 556–585.

Rubin, Jonathan, Paul N. Leiby, and David L. Greene. 2008. "Tradable Fuel Economy Credits: Competition and Oligopoly." *Journal of Environmental Economics and Management*, 58: 315–328.

U.S. Energy Information Administration. 2013. "The Electricity Market Module of the National Energy Modeling System Model Documentation." http://www.eia.gov/forecasts/aeo/nems/documentation/electricity/pdf/m068

U.S. Environmental Protection Agency. 2013. "Documentation for EPA Base Case v.5.13 Using the Integrated Planning Model." http://www.epa.gov/airmarkets/programs/ipm/psmodel.html.

U.S. Environmental Protection Agency. 2014. "Translation of the Clean Power Plan Emission Rate-Based CO_2 Goals to Mass-Based Equivalents." http://www2.epa.gov/sites/production/files/2014-11/documents/20141106tsd-rate-to-mass.pdf.

U.S. Environmental Protection Agency. 2015. "Clean Power Plan Final Rule."

Appendix

Simulation Model Formulation

The social welfare function for the simulation model is the following:

$$W(q^i_{j,r,m}, t_{r,r',m}) = \sum_r \sum_m \Bigg[\int_0^{\sum_i \sum_j q^i_{j,r,m} + \sum_{r'} ((1-\ell_{r',r,m}) t_{r',r,m} - t_{r,r',m})} P_{r,m}(u) \mathrm{d}u \\ - \sum_i \sum_j c^i_{j,r,m} \cdot q^i_{j,r,m} - \sum_{r'} w_{r',r} \cdot t_{r',r,m} \Bigg], \tag{22}$$

where m indexes time segments as described in Section . The social planner's objective function is now maximized over unit-specific generation ($q^i_{j,r,m}$) and the level of electricity transfer ($t_{r',r,m}$) from region r' to region r within each time segment. Transfers between regions are subject to losses and wheeling costs, denoted by $\ell_{r',r,m}$ and $w_{r',r}$, respectively.

Without regulation, the social planner solves the following program:

$$\begin{aligned} \underset{q^i_{j,r,m}, t_{r,r',m}}{\text{maximize}} \quad & W(q^i_{j,r,m}, t_{r,r',m}) \\ \text{subject to} \quad & 0 \leq q^i_{j,r,m} \leq K^i_{j,r,m} && \forall i, j, r, m \\ & 0 \leq \sum_{m \in S} q^i_{j,r,m} \leq A^i_{j,r,S} && \forall i, j, r, S \\ & 0 \leq t_{r,r',m} \leq T_{r,r',m} && \forall r, r', m. \end{aligned} \tag{23}$$

Inter-regional transmission capacity limits are given by $T_{r,r',s}$. In addition to a capacity constraint within time segments, generation units also have a seasonal availability constraint to reflect forced outages for maintenance, where seasons are denoted by S. Availability is specified as a proportion of the sum of capacity over all time segments:

$$A^i_{j,r,S} = a^i_{j,r,S} \sum_{m \in S} K^i_{j,r,m}, \tag{24}$$

where $a^i_{j,r,S} \in (0,1)$.[16]

The program under regulation is analogous to the one described in Section 4. Seeking to meet an emissions target through tradable performance standards, the regulator seeks to maximize the following Lagrangian:

$$\begin{aligned}
&\mathcal{L}(q^i_{j,r,m}, t_{r,r',m}, \lambda^i_{j,r,m}, \mu^i_{j,r,S}, \gamma_{r,r',m}, \tau) \equiv W(q^i_{j,r,m}, t_{r,r',m}) \\
&-\sum_r \sum_m \sum_i \sum_j \lambda^i_{j,r,s}(q^i_{j,r,m} - K^i_{j,r,s}) - \sum_r \sum_S \sum_i \sum_j \mu^i_{j,r,S}\Big(\sum_{m\in S} q^i_{j,r,m} - A^i_{j,r,S}\Big) \\
&-\sum_r \sum_{r'} \sum_m \gamma_{r,r',m}(t_{r,r',m} - T_{r,r',m}) - \tau\Big(\sum_r \sum_m \sum_i \sum_j (\sigma^i_{j,r,m} - s_j) q^i_{j,r,m}\Big).
\end{aligned} \tag{25}$$

Under the specification described in Section 8, the program to be solved has a quadratic objective function with linear constraints.

Model Parameterization

Electricity demand in the model is assumed to be linear. For each region and time segment in the model, I define the function

$$P_{r,m}(q_{r,m}) = a_{r,m} - br, mq_{r,m},$$

where $q_{r,m}$ is the quantity of electricity demand in region r in time segment m. For given demand elasticity $\epsilon < 0$, the demand parameters are defined as

$$a_{r,m} = \frac{\epsilon - 1}{\epsilon} p^{bau}_{r,m}$$

and

$$b_{r,m} = \frac{p^{bau}_{r,m}}{\epsilon q^{bau}_{r,m}},$$

[16]The approach to modeling seasonal availability is similar to the EPA's Base Case v.5.13, which uses ICF International's Integrated Planning Model (IPM) (U.S. Environmental Protection Agency, 2013). Alternatively, models such as the one developed in Bushnell and Chen (2012), specify a similar availability constraint in each time segment. As a result, the capacity constraint becomes superfluous. Conceptually, the seasonal approach reflects the ability of the unit operator to schedule maintenance strategically to minimize lost profits, while the per-segment approach reflects the unexpected nature of some outages.

where $q_{r,m}^{bau}$ and $p_{r,m}^{bau}$ represent equilibrium prices and quantities under business-as-usual.

Generating unit costs per MWh are constructed as VOM costs plus the product of the heat rate (HR) and the fuel price (FP):

$$c_{j,r,s}^{i} \equiv VOM_{j,r}^{i} + HR_{j,r}^{i} \cdot FP_{j,r,s}^{i}. \tag{26}$$

The heat rate is defined in terms of mmBtu (million British thermal units) per MWh, i.e., the amount of primary heat energy required by the plant to produce a unit of electricity. Fuel prices are measured by $ per mmBtu, i.e., the price per unit of heat; they are assumed to be zero for non-thermal units (e.g., nuclear and renewables).

Emission rates, defined in terms of pounds of CO_2 per MWh, are the product of the heat rate of the generating unit and an emissions factor that measures the CO_2 intensity of the associated fuel:

$$\sigma_{j,r}^{i} \equiv EF_{j,r}^{i} \cdot HR_{j,r}^{i}. \tag{27}$$

The emissions factor is defined in terms of pounds of CO_2 per mmBtu. For natural gas generators, this factor is assumed to equal across all units. For coal plants, this factor defers depending on the type of coal that the unit burns.

Empirical Credit Price

The credit price as well as the relationship between the coal and gas standard are equilibrium results from the model simulations; both are displayed graphically in Figure 16, plotted against the coal standard. As expected, the credit price, represented by the solid line, not only increases with the coal standard, but exhibits substantial convexity. For instance, the credit price for the most extreme degree of differentiation in the figure is around six times larger than the value at the single standard. The dashed line demonstrates the relationship between the two standards. As a function of the coal standard, the gas standard is decreasing and slightly concave.

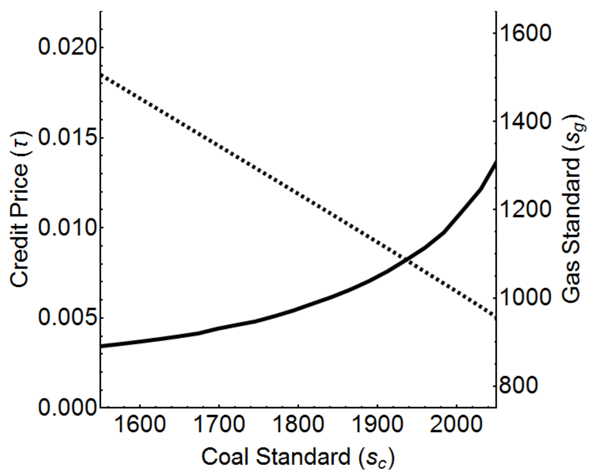

Figure 16: Credit Price and Gas Standard Plots