Designing markets for pollution when damages vary across sources: What are the gains from differentiation?

Meredith Fowlie and Nicholas Muller*

October 31, 2011

Preliminary and incomplete. Comments very welcome.

Abstract

A majority of the air pollution currently regulated under U.S. emissions trading programs is non-uniformly mixed, meaning that health and environmental damages depend on the location and dispersion characteristics of the sources. Most emissions trading programs ignore this fact. Emissions are penalized at a single permit price, regardless of the location of the source. In theory, differentiated policies can be designed to accommodate non-uniformly mixed pollution using emissions penalties that vary with emissions damages. We present a simple framework to illustrate the gains from differentiation in first-best and second-best settings. This serves as foundation for a detailed analysis of the gains from differentiation in a major U.S. emissions trading program. We take two complementary approaches to estimating these gains. Our preferred estimate, which is generated using an econometrically estimated model of firms' compliance choices, is surprisingly small given the extent of the variation in damages across sources. A more standard approach to simulating of policy outcomes, one that assumes strict cost minimization on behalf of all firms, predicts larger benefits. A comparison of the two approaches provides insights into the determinants of the gains from policy differentiation, some of which are inadequately captured by standard policy simulation models.

Keywords: Market-based Policy, NO_x Budget Program, Policy Instrument Choice.

JEL Classifications: Q54, Q53, Q58

*Fowlie: Department of Agricultural Economics, UC Berkeley, and NBER. Muller: Department of Economics and the Environmental Studies Program, Middlebury College. We thank Christian Vossler, Jacob LaRiviere, and seminar participants at Carnegie Mellon, Resources for the Future, University of Pittsburgh, and the University of Tennessee for helpful discussions and comments. We thank Dilek Uz for invaluable research assistance. Muller thanks the United States Environmental Protection Agency for support under award: EPA-OPEI-NCEE-08-02.

1 Introduction

Economists have long advocated for market-based approaches to pollution regulation (Montgomery, 1972; Baumol and Oates, 1988). The past three decades have witnessed large scale experimentation with implementing emissions trading programs in practice. By many measures, this experimentation has been very successful. Targeted emissions reductions have been achieved or exceeded, and it is estimated that total abatement costs have been significantly less than what they would have been in the absence of the trading provisions (Carlson et al. 2000; Stavins, 2005)

In terms of economic efficiency, however, many existing cap-and-trade programs likely fall short of the theoretical ideal. Efficiency requires that marginal abatement costs be set equal to marginal damages at all sources (Baumol and Oates, 1988; Montgomery, 1972). Most policies are currently implemented as spatially uniform, "undifferentiated" trading programs. That is, all regulated emissions are taxed at the equilibrium permit price. By equalizing marginal abatement costs across all sources, an undifferentiated trading regime will minimize the total abatement costs incurred to meet the emissions cap. However, when a pollutant is "non-uniformly mixed" (i.e. health and environmental damages from emissions vary across sources), efficiency cannot achieved by equating marginal abatement costs across sources. Spatial variation in emissions damages across sources thus gives rise to a trade off between minimizing pollution abatement costs and minimizing the damages caused by permitted emissions. This paper investigates these trade-offs in theory and practice.

Market-based policies can, in theory, be designed to account for spatial variation in damages (Montgomery, 1972; Tietenberg, 1980; Muller and Mendelsohn, 2009). Baumol and Oates (1988) use a general equilibrium model to depict optimal pollution taxes in a setting with heterogeneous costs and damages. The optimal tax rate is calibrated to the marginal damage caused by emissions. When damages vary by source, so do the tax rates. Other authors have proposed "differentiated" emissions permit market designs wherein differences in marginal social cost are reflected in different "prices" for emissions (Mendelsohn, 1986; Tietenberg, 1995.

This paper begins by developing a conceptual framework useful for quantifying the benefits from differentiation. As emphasized by Mendelsohn (1986), these benefits depend on both the extent of the variation in damages across sources and the steepness of the marginal abatement costs. We extend the established theoretical literature so as to accommodate two practical considerations that complicate the design and implementation of differentiated policy designs. First, we consider a second-best setting in which a regulator seeks to minimize pollution damages plus abatement cost subject to an exogenously set cap on emissions. Second, we note that the policy maker is rarely, if ever, fully informed. We analyze the design of differentiated policy when there is both uncertainty surrounding estimates of pollution damages and very limited information about abatement costs.

The conceptual framework serves as foundation for a detailed analysis of the gains from policy differentiation in the context of a landmark emissions trading program, The NOx Budget Program limits nitrogen oxide (NOx) emissions from large point sources in the Eastern United States. Previous work has documented considerable variation in the per ton damages from NO_x emissions (Mauzerall et al., 2005; Tong et al., 2006; Levy et al., 2009; Muller, Tong, Mendelsohn, 2009; Muller and Mendelsohn, 2009). In the design stages of this program, policy makers were aware of this heterogeneity and considered imposing restrictions on interregional trading (FR 63(90): 25902). Ultimately, it was decided that the potential benefits from this additional complexity would not justify the costs (US EPA, 1998). The program was therefore implemented as a single jurisdiction, spatially uniform trading program in which all emissions are traded on a one-for-one basis. With the benefit of hindsight, we revisit the decision to forego spatially differentiated NOx trading in favor of the simpler undifferentiated alternative.

Our analysis proceeds in four stages. First, source-specific marginal damage estimates are generated using a stochastic integrated assessment model (APEEP, Muller, Mendelsohn, 2007;2009), AP2. Second, we model firms' compliance choices under the observed (undifferentiated) and counterfactual (differentiated) policy regimes. In the third step, the simulated compliance choices are mapped into source-specific NO_x emissions and abatement costs. This allows us to estimate the effect of policy differentiation on abatement costs and the spatial distribution of the permitted emissions. Finally, the integrated assessment model is re-introduced for the purpose of estimating the aggregate health and environmental impacts of the simulated emissions associated with each policy scenario. With estimates of the abatement costs and damages associated with the observed and counterfactual policy designs in hand, we construct our estimate of the gains from policy differentiation.

This approach has several attractive features. First, the detailed integrated assessment model we use allows us to estimate *source-specific* marginal damages. We find significant variation in these damages; almost half of this variation occurs within (versus between) state. This finding is germane to the unfolding debate about market-based regulation of non-uniformly mixed pollutants because it suggests that multi-state trading zones are a blunt instrument for capturing spatial variation in NOx emissions damages.

Second, we are able to address an important limitation of earlier work that has investigated the benefits of policy differentiation. In the literature, it is standard to assume that firms adhere to strictly cost-minimizing strategies. Emissions market outcomes are simulated using a deterministic, cost-minimization algorithms. Previous authors have noted that this approach fails to capture salient features of the real world decision processes that drive emissions abatement decisions (Krupnick et al. 2000). Given the retrospective nature of our analysis, we are able to observe the compliance choices that regulated firms actually made. Our preferred approach to simulating firms' response to observed and counterfactual policy incentives makes use of an econometrically estimated model of facility-level compliance choices in the NOx Budget Program (Fowlie, 2010).

Finally, our approach lends itself to analyzing the policy implications of an important form of uncertainty. Increasingly, integrated assessment modeling (IAM) is being used to inform policy design and implementation (e.g. Transport Rule RIA). Models used to estimate marginal damages from emissions of criteria pollutants incorporate many imprecisely estimated parameters and stochastic inputs. Estimates of source-specific marginal damages, which form a key aspect of the design of efficient policy, are therefore uncertain. We characterize the joint distribution of the marginal damage parameters associated with the sources in the NBP. We then consider how a policy maker can use this rich information in the design and analysis of differentiated policy.

Our preferred estimate of the gains from policy differentiation, which is generated using the econometrically estimated model of firms' compliance decision, is surprisingly small given the extent of the intersource variation in pollution damages. Expected benefits associated with the mandated emissions reduction increase by only 2 to 5 percent under the differentiated policy vis a vis the undifferentiated benchmark. In contrast, the cost minimization-based simulations predict gains of 15-18 percent. We argue that the latter, more standard approach overestimates the extent to which firms would likely have responded to differentiated policy incentives.

2 Theoretical framework

Consider a group of N firms emitting a non-uniformly mixed pollutant. The extent of the damage caused by these emissions depends not only on the level of emissions, but also how the emissions are distributed across sources.[1] We define abatement cost functions in terms of emissions: $C_i(e_i)$. We assume that $C_i'(e_i) \leq 0 \leq C"_i(e_i)$.

We define pollution damage functions, also in terms of emissions: $D_i(e_i)$. And we assume that $D_i'(e_i) \geq 0$, and $D_i''(e_i) \geq 0$. Further, we capitalize on a series of empirical simulations, which are discussed in appendix A1, to put further structure on the damage function; we assume that the damage function is linear and additively separable. The implication is that the product of the marginal damage times emissions is equal to the area beneath the total damage function. We begin by assuming that policy makers know these marginal damage parameters with certainty prior to implementing the emissions policy.

Suppose that the policy maker's objective is to minimize the total social cost (TSC) associated with emissions of this pollutant:

$$TSC = \sum_{i=1}^{N} (D_i(e_i) + C_i(e_i)) \tag{1}$$

The first component in Eq. (1) measures damages from pollution. The second term measures the costs of reducing emissions levels below unconstrained "business as usual" levels. To minimize total social costs, one differentiates Eq. (1) with respect to source-level emissions. Assuming an interior solution, first-order conditions for a cost-minimization imply:

$$-C_i'(e_i^*) = D_i'(e_i^*) \;\; \forall \; i. \tag{2}$$

Intuitively, marginal costs are set to equal marginal damages across all sources. The * superscript denotes efficient emissions levels. The efficient level of aggregate emissions is thus $E^* = \sum_{i=1}^{N} e_i^*$.

.

[1] In this analysis, we will focus exclusively on the spatial heterogeneity in damages. See Joskow, Martin and Ellerman (CITE) for an analysis of the implications of temporal variation in damages.

2.1 Market-based regulation of non-uniformly mixed pollution

Having characterized the first-best emissions outcome, we next evaluate the performance of alternative market-based policy designs against this benchmark. We will focus exclusively on cap-and-trade programs which impose a cap on total emissions and issue a corresponding number of tradeable emissions permits.[2] We assume that the emissions permits are allocated to participating sources either by auction or a gratis using some allocation rule that does not depend on production decisions going forward. Any free allocation of permits to firm (i) is represented by the initial allocation A_i.

As a point of departure, we assume that the emissions cap is set optimally at E^* and that markets are efficient and free of distortions. We subsequently refer to this as the "first-best" case. To keep the analytics simple and intuitive, we consider a case with only two price taking firms. It is straightforward to generalize the following analysis to $N > 2$, although this complicates notation unnecessarily. Producers are denoted h and l to indicate high and low damage areas, respectively, and in order to clean-up the exposition of our results, the marginal damage for firm (h), $D_h'(e_h)$, is denoted (δ_H), and the marginal damage for firm (l), $D_l'(e_l)$, is denoted (δ_L).

We first consider a market-based policy design that implicitly assumes that marginal damages are equal across all sources. We then contrast this with policy regimes in which the terms of compliance are designed to reflect heterogeneity in damages.

2.2 Undifferentiated Permit Market

Most existing and planned emissions trading programs feature undifferentiated permits: meaning that firms are required to hold a permit to offset each unit of emissions, regardless of where the emissions occur. Trading occurs on a ton-for-ton basis. Building on the formulation in Muller and Mendelsohn (2009), we assume that each firm chooses emissions (e_i), emissions permit purchases (e_{bi}), and permit sales (e_{si}), both valued at the market-determined price (τ), to minimize the costs of complying with this emissions-based trading program.

$$\begin{aligned} \min_{E_i, E_{si}, E_{bij}} \quad & C_i(e_i) + \tau\left(e_{bi} - e_{si} - A_i\right) \\ s.t. \quad & e_i \leq A_i - e_{si} + e_{bij} \\ & e_i, e_{si}, e_{bij} \geq 0, \end{aligned} \tag{3}$$

If we assume an interior solution, cost-minimization implies that marginal abatement costs are set equal across all sources:

$$C_h'(e_h^u) = C_l'(e_l^u) = \tau, \tag{4}$$

where the u superscript denotes the undifferentiated trading equilibrium.

[2] Although we are ultimately concerned about limiting the *damages* associated with pollution exposure, it is far more standard to regulate *emissions*, presumably because imposing a cap on emissions is relatively simple and easy to communicate.

Figure 1 illustrates these first order conditions in the simple two firm case. The width of this figure, measured in units of emissions, is equal to the total quantity of permitted emissions $\overline{E}$. In this first best case, we assume this cap has been optimally set to E^*. At the left origin, all emissions occur at the low damage firm (i.e. $E_l = \overline{E}$) and emissions at the high damage firm are driven to zero ($E_h = 0$). The upward sloping solid line, moving from left to right, represents the marginal abatement costs at the low damage firm $C_l'(e_l)$. At the right origin, the high damage firm emits E^* (i.e. $e_h = \overline{E}$) and the low damage firm emits nothing ($e_l = 0$).The solid line increasing from right to left measures marginal abatement costs at the high damage firm $C_h'(e_h)$.

Equilibrium emissions under the undifferentiated trading regime are given by $\{e_l^u, e_h^u\}$. This equilibrium occurs at the intersection of $C_h'(e_h)$ and $C_l'(e_l)$. This allocation of permitted emissions minimizes the total abatement costs required to meet the emissions cap $\overline{E}^*$. However, this is not the optimal outcome. Total social welfare could be improved by shifting some of the permitted emissions away from the high damage source to the low damage source (Muller and Mendelsohn, 2009). As long as the policy facilitates trading and compliance based on one permit price (τ), and $\delta_h \neq \delta_l$, an undifferentiated market cannot achieve allocative efficiency.

2.3 Differentiated Permit Market

We now consider how this market-based policy design can be modified so as to achieve the socially optimal allocation of permitted emissions. There is a growing literature that examines "differentiated" policies that are designed to reflect variation in pollution damages (Teitenberg, 1995; Farrow et al., 2004; Horan and Shortle, 2005; Muller and Mendelsohn, 2009). To date, work in this area has focused on the construction of trading ratios based on the ratio of marginal damages between each pair of regulated sources. It is straightforward to operationalize these ratios within our simple analytical framework.

Let $\overline{\delta}$ represent the average of the marginal damage across all sources in a trading program. In this simple two firm case, $\overline{\delta} = \frac{\delta_l+\delta_h}{2}$. We construct firm-specific damage ratios r_i, normalizing each firm's marginal damage by the mean damage parameter : $r_i = \frac{\delta_i}{\overline{\delta}}$. To remain in compliance, each firm must hold r_i permits to offset each unit of uncontrolled emissions. Returning to Eq. (3) the compliance constraint is:

$$r_i e_i = A_i - e_{si} + e_{bi}.$$

All else equal, the more damage caused by emissions at a given source, the more permits that source needs to hold to offset its emissions.

The first order conditions with respect to (e_i) for cost minimization in this differentiated regime imply that the ratio of marginal damages will be set equal with the ratios of marginal costs across the two firms:

$$\frac{C_j'(e_j^r)}{C_i'(e_i^r)} = \frac{\delta_j}{\delta_i},\ i \neq j. \tag{5}$$

where the r superscript denotes the equilibrium outcome under a regime that incorporates these trading ratios.[3]

In this first-best setting, (5) delivers the socially optimal allocation of permitted emissions across sources (see Appendix 1). Figure 1 illustrates this result graphically. The broken lines represent the marginal abatement cost schedules scaled by the inverse of the corresponding marginal damage: $C_i'(e_i)\frac{1}{\delta_i}$, $i = l, h$. By (5), the allocation of emissions across these two sources occurs where these broken lines intersect. This allocation of the permitted emissions achieves the optimal trade off between abatement costs and benefits from reduced damages.

This differentiated policy design will welfare-dominate the undifferentiated system if benefits (in the form of avoided damages) exceed the increase in abatement costs. In Figure 1, gross benefits from differentiation are represented by area ABCE The increase in abatement costs is equal to area ACD. The net benefits from differentiation, defined by the shaded areas ABC + CDE, are positive.

To fix ideas, we derive a more general expression for these net benefits. We maintain our assumption that the aggregate damage function is linear and additively separable in source-specific damages and the notation that $\delta_i = D_i'(e_i)$. We now assume the following functional form for the source-specific abatement cost functions: $C_i(e_i) = \alpha_{0i} - \alpha_{1i}e_i + \beta_i e_i^2$. We accommodate heterogeneity in abatement costs by allowing the parameters of this function to vary across sources. The first derivative of $(C_i(e_i))$ serves as a linear approximation to the marginal abatement cost at each source. In practice, these cost curves may be discontinuous as pollution abatement often involves lumpy investments in emissions reducing capital equipment.

Solving for the optimal cap yields the following:

$$\overline{E}^* = \frac{(\alpha_{1h}\beta_l + \alpha_{1l}\beta_h - \delta_h\beta_h - \delta_l\beta_h)}{2\beta_l\beta_h}. \tag{6}$$

Note that this optimal limit depends not only on the damage parameters $\{\delta\}$ and the cost parameters$\{\alpha_1, \beta\}$, but also on the correspondence between these two sets of parameters.

Solving for equilibrium emissions under the emissions-based and differentiated policy designs, respectively, we obtain an expression for the change in source-specific emissions induced by this differentiation:

$$e_h^r - e_h^u = \frac{\delta_l - \delta_h}{2(\beta_h + \beta_l)} \tag{7}$$

$$e_l^r - e_l^u = \frac{\delta_h - \delta_l}{2(\beta_h + \beta_l)} \tag{8}$$

[3] It is worth pointing out that the phrase "trading ratio" is somewhat misleading insofar as these ratios affect emissions trading activities only indirectly via the effect on compliance requirements.

Intuitively, differentiation shifts some share of permitted emissions from the high damage source to the low damage source. The extent of this reallocation depends on the difference in damage parameters and the steepness of the marginal abatement cost curves.

It is then straightforward to derive an expression for the net benefits from differentiation in terms of the model parameters (see Appendix 2 for complete derivation):

$$TSC^u(\delta,\beta) - TSC^r(\delta,\beta) = \frac{(\delta_l - \delta_h)^2}{4(\beta_l + \beta_h)} \geq 0, \tag{9}$$

The e superscript denotes the emissions-based design; r denotes the differentiated design that incorporates ratios. The vectors of damage parameters and abatement cost coefficients are denoted δ and β, respectively.

We can now make three observations based on equation (9):

(1) In a first best setting, an emissions trading program that incorporates differentiated "trading ratios welfare dominates more standard policy designs when damages vary across sources.

(2) *The extent to which differentiation reduces pollution damages (via a reallocation of permitted emissions) is increasing with the variation in damages and decreasing with the slope of the marginal abatement cost functions.*

(3) The net benefit conferred by differentiation (i.e. the value of avoided damages less the increase in abatement costs) is increasing with the variation in damages and decreasing with the slope of the marginal abatement cost functions.

Points (1) through (3) should be intuitive. First, if damages do not vary, there is no advantage to differentiated policy. Accordingly, the more heterogeneous the damages, the greater the benefits from differentiation, all else equal. Finally, if marginal abatement costs are steeply increasing in abatement, it will be relatively more costly to shift emissions from the high damage to the low damage source.

Note that the benefits from differentiation do not depend on the correlation between source-specific abatement costs and source-specific damage parameters. This contrasts with the findings of Mendelsohn (1986) who finds that positive covariance between abatement cost parameters and emissions damages increase the relative effectiveness of differentiated policy designs. Given our maintained assumptions regarding the linearity of the damage function, this relationship disappears.

2.3.1 Exogenously determined emissions constraint

The above analysis is predicated on some strong simplifying assumptions that are likely to be violated in practice. In what follows, we investigate two practical design considerations. This section we explore optimal policy in the presence of an exogenously determined emissions constraint which may differ from the optimal cap. In the next section, we incorporate uncertainty in the marginal damage estimates.

In the theory literature that considers the design and implementation of spatially differentiated emissions policies, it is often assumed that the cap can be optimally set (Muller and Mendelsohn, 2009). In fact, this is unlikely to be a safe assumption. It is often the case that the emissions constraint is (explicitly or implicitly) determined by a superseding authority. The implementing agency must therefore determine the optimal policy design conditional on this cap. Additionally, even if the regulator could set caps optimally, it is not clear that they could do so without firm-specific estimates of marginal social cost and marginal abatement costs.

To examine this more common situation, we now assume that the policy maker seeks to minimize total social costs subject to an emissions constraint:

$$\begin{aligned} \min_{e_h, e_l} TSC &= D_h(e_h) + D_l(e_l) + C_h(e_h) + C_l(e_l) \\ s.t. e_h + e_l &\leq \overline{E} \end{aligned}$$

First order conditions for constrained cost minimization imply:

$$C_h'(e_h) - C_l'(e_l) = \delta_h - \delta_l \tag{10}$$

When $C_i'(e_i) = \delta_i$ across all sources, setting ratios of marginal damages equal to ratios of marginal abatement costs implies that differences in marginal damages are equal to differences in marginal abatement costs. Put differently, equations (2) and Eq. (10) can be simultaneously satisfied. However, if the cap has not been set optimally, the market fails to clear when marginal abatement costs are set equal to marginal damages. In a regime that incorporates first best trading ratios, marginal abatement costs will exceed (fall below) marginal damages if the cap has been set too stringently (loosely). In either case, Eq. (10) is not met if the first-best ratios are used to design policy.

Figure 2 illustrates the case in which the emissions constraint is too stringent. As in Figure 1, the intersection of the broken lines defines the equilibrium allocation of permitted emissions under a policy regime that employs first-best trading ratios $\{e_h^r, e_l^r\}$.At this allocation, the ratio of marginal abatement costs equals the ratio of marginal damages. Because the cap has been set too stringently, the difference in marginal abatement costs exceeds the difference in marginal damages. Contrast this with the allocation of emissions denoted by the superscript ^{w} which satisfies Eq. (10). This allocation minimizes the total social cost associated with the emissions constraint $\overline{E}$. Net benefits vis a vis the emissions-based equilibrium are given by the value of avoided damages (area DEFG) less increased costs (area ABC).

Can the constrained optimal outcome be achieved using a decentralized, market-based policy intervention? The theory literature has investigated the use of "second best" ratios which are intended to minimize the total social costs associated with a given emissions constraint (see, for example, Horan and Shortle, 2005; Muller, 2011). As compared to the first-best ratios introduced above, these second-best ratios are more difficult to construct. For example, to implement the second-best ratios developed by Horan and Shortle (2005) the policy maker must correctly anticipate the shadow value of the imposed emissions constraint.

We introduce an alternative approach to achieving the constrained optimum. Let w_i represent a source-specific difference (or "wedge") between the firm-specific marginal damage and the expected marginal damage across sources: $w_i \equiv \delta_i - \overline{\delta}_i$. As in the emissions-based regime, firms must hold permits to offset uncontrolled emissions in order to remain in compliance. But we now introduce an additional compliance requirement. In addition to holding permits, each firm must pay a compliance payment equal to: $e_i(\delta_i - \overline{\delta}_i)$. This is simply their compliance wedge times their emission level. It reflects either the additional (or subtracted) damage relative to an average firm. The firm's compliance cost minimization problem can now be written:

$$\begin{aligned} \min_{e_i, e_{si}, e_{bij}} \quad & C_i(e_i) + (\tau + w_i)\,(e_{bi} - e_{si} - A_i) \\ s.t. \quad & e_i \leq A_i - e_{si} + e_{bij} \\ & e_i, e_{si}, e_{bij} \geq 0, \end{aligned} \tag{11}$$

This policy design modification drives a wedge between the market clearing permit price and the price paid by firms whose marginal damage exceeds, or falls below, the average damage value. As in the trading ratio regime, relatively high (low) damage firms face a higher (lower) cost of offsetting emissions using permits. This tends to reallocate emissions from sources that cause high marginal damage to sources whose discharges cause lower damages.

The first order conditions under this spatially differentiated regime indicate that the firm's cost minimizing behavior is aligned with the socially optimal outcome in the constrained setting.

$$C_h'(e_h) - C_l'(e_l) = \delta_h - \delta_l. \tag{12}$$

A policy regime that incorporates these wedges will achieve the constrained optimum, even when the cap is not optimally set. The net welfare gains (vis a vis the undifferentiated design) are precisely the same as in the first best case (see Eq. [9]). Note that these gains do not depend on the exogenously set emissions cap. Simply stated:

(4) When the emissions cap is not set optimally (as will often be the case), differentiation based on differences in marginal damages (versus ratios) achieves the optimal allocation of the permitted emissions.

The advantage of using wedges (versus second-best trading ratios) to define compliance requirements is that the information requirements are not as onerous. Only source specific measures of marginal damages are required. A potential disadvantage is that the wedge-based policy is not public revenue neutral. Under the benchmark emissions-based design and policies that incorporate trading ratios, the implementing agency does not pay out or take in funds while administering program compliance. In contrast, under the wedge-based regime, program compliance requires high damage firms to pay $\omega_i e_i$ to the implementing agency, while the agency is required to pay $w_i e_i$ to low damage firms. Depending on how the permitted allocations are allocated across high and low damage firms in equilibrium, the implementing agency may be net long or short after compliance requirements have been satisfied.

2.3.2 Uncertain damages

Another important practical consideration is the uncertainty that pervades the estimation of damages from pollution. The integrated assessment models that are used to estimate marginal damages from emissions of criteria pollutants incorporate many imprecisely estimated parameters and stochastic inputs. Estimates of source-specific marginal damages, which form a key aspect of the design of efficient policy, are therefore uncertain.

Increasingly, integrated assessment modeling (IAM) is being used to inform policy design and implementation (e.g. Transport Rule RIA). We assume that the joint distribution of the marginal damage parameters $f(\delta_l, \delta_h)$ is known ex ante. The policy maker seeks to minimize expected total social costs from emissions subject to the exogenously determined emissions constraint:

$$\begin{aligned} \min_{e_h, e_l} TSC &= C_h(e_h) + C_l(e_l) + \int\int (\delta_h e_h + \delta_l e_l) f(\delta_l, \delta_h) \\ s.t. \quad e_h + e_l &= \overline{E} \end{aligned} \tag{13}$$

Substituting in the constraint, the first order condition for cost minimization yields:

$$\begin{aligned} -(C_h'(e_h) - C_l'(e_l)) &= \int\int (\delta_h - \delta_l) f(\delta_l, \delta_h) \\ -(C_h'(e_h) - C_l'(e_l)) &= E[\delta_h] - E[\delta_l] \end{aligned} \tag{14}$$

The constrained optimum in the presence of uncertainty about the marginal damage parameters equates differences in marginal abatement costs with differences in expected marginal damages. If the w_i parameters are constructed using expected source-specific marginal damages $E[\delta_h]$, $E[\delta_l]$, Eq. (14) will be satisfied in expectation.

We note two important qualifications. First, our assumptions regarding the linear form of the damage function come into play. In the literature that examines the implications of uncertain damages on optimal policy, researchers have argued that the optimal trading ratio between two sources with equal expected damages but varying degrees of uncertainty should penalize the more uncertain damages (e.g. Horan, 2001; Horan and Shortle, 2005;). In our case, linearity eliminates the covariance term that gives rise to this penalty. Our assumed policy objective function is also important. We assume that the regulator seeks to minimize expected social costs. If instead the regulator wants to meet an ambient target probabilistically, varying degrees of uncertainty will matter because otherwise identical firms will have differential marginal effects on the probability the target is violated.

Before looking into how uncertain damages affect the expected benefits from differentiation, it is instructive to consider how net benefits vary across draws from the distribution $f(\delta)$. Consider a policy regime that incorporates wedges defined using $E[\delta_l]$ and $E[\delta_h]$. Let $\delta' = \{\delta_h, \delta_l\}$ denote a particular draw from the joint distribution of marginal damages. We can think of this as the damages that were actually realized. The difference in total social costs under the compliance wedge design (TSC^w), relative to an undifferentiated policy (TSC^u) is given by:

$$TSC^w - TSC^u = \frac{1}{2} \frac{\left(\delta_l' - \delta_h'\right)\left(E[\delta_l] - E[\delta_h]\right)}{(\beta_l + \beta_h)} - \frac{1}{4} \frac{\left(E[\delta_l] - E[\delta_h]\right)^2}{(\beta_l + \beta_h)}. \tag{15}$$

The first argument in Eq. (15) represents the difference in damages across emissions-based and differentiated (second-best) policy regimes. The second argument captures the increase in abatement costs associated with a move to the differentiated policy design. Given uncertain damages, it is possible for the realized net benefits of differentiation to be negative. Consider an extreme case where the realized damage values $\{\delta_h, \delta_l\}$ are negatively correlated with the expected damage values, $E[\delta_l]$ and $E[\delta_h]$. This implies that the source that was expected to be associated with relatively low damages is actually the relatively high damage source. The differentiated policy will incorrectly penalize the low damage source vis a vis the high damage source. Because, in the case of negative covariances, the $E[\delta_l]$ and $E[\delta_h]$, are effectively misleading the regulator, society is better off to maintain a policy that treats all sources as if their damages at the margin are equal. The uniform, emissions-based regime welfare dominates the (misguided) differentiated regime.

The expected net welfare gain associated with the spatially differentiated policy is obtained by integrating over the entire distribution of damages (see Appendix X):

$$E[TSC^w(f(\delta), \beta) - TSC^e(f(\delta), \beta)] = \frac{(\delta_h - \delta_l)^2}{4(\beta_h + \beta_l)} \left(2cov(\delta_l, \delta_h) - var(\delta_l) - var(\delta_h)\right)$$

Having derived Eq. (**??**), we can make some additional observations regarding the determinants of the benefits from differentiation under uncertainty:

(5) Benefits from differentiation are decreasing with the variance of source-specific damage distributions.

(6) Benefits from differentiation are increasing with the covariance in damages across sources.

Intuitively, if the damage parameters are precisely estimated, the policy maker can design the policy to more accurately reflect the ex post realized damages. If there is strong positive correlation in damage realizations across sources, the average damage parameters which are used to define the terms of compliance will be more positively correlated (in expectation) with the ex post realized damages.

3 The NOx Budget Program

The NOx Budget Program (NBP) is a market-based emissions trading program created to reduce the regional transport of NOx emissions in the eastern United States. The program establishes a region-wide cap on emissions of NOx from large stationary sources in twenty eastern states during ozone season (May-September). The NBP was primarily designed to help Northeastern and Mid-Atlantic states attain Federal ozone standards. When the NBP was promulgated, significant portions of the Northeast, Mid-Atlantic, and parts of the Midwest were failing to meet Federal standards (Ozone Transport Assessment Group (OTAG), 1997).

Although the precise contribution of individual sources to the non-attainment problems in this region was difficult to estimate at the time of the rulemaking, there was plenty of evidence to suggest that marginal damages varied significantly across sources. The EPA received over 50 responses when, during the planning stages of the NOx SIP Call, it solicited comments on whether the program should incorporate trading ratios or other restrictions on interregional trading in order to reflect the significant differential effects of NOx emissions across states(FR 63(90): 25902). Most commentors supported unrestricted trading and expressed concerns that "discounts or other adjustments or restrictions would unnecessarily complicate the trading program, and therefore reduce its effectiveness" (FR 63(207): 57460). These comments and some accompanying analysis (US EPA, 1998a) led regulators to design a single jurisdiction, undifferentiated trading program. There are no spatial restrictions on trading within the program. All emissions are treated symmetrically for compliance purposes.[4]

In what follows, we revisit the decision to forego a differentiated policy design in favor of the simpler, undifferentiated alternative. Our analysis will focus exclusively on the coal-fired generating units in the program. Although gas- and oil-fired generators and other industrial point sources are also included in the NBP, coal-fired units represent approximately 94 percent of the NOx emissions regulated under the program and at least 94 percent of the NOx emissions reductions over the first five years (U.S. EPA, 2005; US E.P.A. 2008). Natural gas and oil-fueled plants tend to have much lower uncontrolled NOx emissions rates.[5] By exempting these units from our analysis, we are implicitly assuming that operating decisions at these units would not be differentially affected under an emissions-based design and the counterfactual policy designs we consider. Future versions of the paper will test this assumption explicitly.

4 Estimating the benefits from differentiation

We use the theory model presented in section 2 as a conceptual framework for a detailed analysis of the benefits from differentiation in the context of the NOx Budget Program. This applied analysis proceeds in five steps, each of which are described in detail in the following five subsections.

4.1 Estimating damages from pollution

[4] The US EPA has also investigated the potential to use weather and atmospheric chemistry forecasts to vary the NOx permit price over time (US EPA, 2007).

[5] Whereas the average pre-retrofit NOx emissions rate among coal plants exceeded 5.5 lbs/MWh, average NOx emissions rates among marginal electricity producers are estimated to range between 0.3 to 2.2 lbs NOx/MWh (NEISO, 2006; Keith et al., 2003).

NOx emissions affect health and environmental outcomes through two main pathways: ozone formation and particulate matter formation.[6] Specifically, emitted NO_x interacts with ambient ammonia to form ammonium nitrate, a constituent of ambient $PM_{2.5}$. And NO_x also forms tropospheric O_3 through a series of chemical reactions (Seinfeld, Pandis, 1998). Both $PM_{2.5}$ and O_3 are criteria air pollutants regulated under Title I of the Clean Air Act. As such, exposures to these two pollutants have been shown to have a number of adverse effects on human health and welfare. Prior research has shown that the majority of damages due to exposures to both $PM_{2.5}$ and O_3 are premature mortalities and increased rates of illness (USEPA, 1999; Muller and Mendelsohn, 2007;2009). Exposure to elevated concentrations of either pollutant has been linked to significant human health and ecosystem damages (see, for example, Brunekreef and Holgate, 2002; WHO, 2003).

The extent to which NO_x emissions react with precursors to form ozone or particulate matter depends upon prevailing meteorological conditions, pre-existing precursor emissions concentrations, and other factors that vary across time and space. Furthermore, the health impacts associated with a change in ozone and/or particulate matter at a particular location will depend on the human and non-human populations at that location. For these reasons, the damage caused by a given quantity of NOx emissions will depend significantly on the spatial distribution of the emissions.

4.1.1 Source-specific damage parameters

The source-specific δ_i parameters capture the effect of an incremental change in NOx emissions at source i on health and environmental impacts across the airshed. We use a stochastic integrated assessment model, AP2, to estimate these source-specific damage parameters (Muller, 2011).

The AP2 model is comprised of six modules; emissions, air quality modeling, concentrations, exposures, physical effects, and monetary damages. The emissions data used in AP2 is provided by the US EPA's National Emission Inventory for 2005 (US EPA, 2009). These data encompass emissions of NO_x, $PM_{2.5}$, sulfur dioxide (SO_2), volatile organic compounds (VOCs), and ammonia (NH_3). AP2 attributes these data to both the appropriate source location and source type. Specifically, AP2 models emissions from 656 individual point sources (mostly large EGUs). Emissions from the remaining point sources are decomposed according to height of emissions and the county in which the source is located. For ground-level emissions (these are produced by cars, residences, and small commercial facilities) AP2 attributes these discharges to the county in which they are reported (by US EPA) to occur.

The approach to air quality modeling used in AP2 relies on the Gaussian Plume model (Turner, 1994). This approach uses a reduced form statistical model to capture the processes that connect emissions (e) to concentrations (C). The relationship between emissions of nitrogen oxides released at source j and the concentration of pollutant s (ozone or particulate matter) at receptor point r is captured by $C_{sri}(e_i)$. The predicted pollutant concentrations generated using the AP2 model have been tested against the predictions made by a more advanced air quality model (see the appendix in Muller, 2011). The agreement between the county-level surfaces produced by the two models is quite strong.

[6] NOx emissions also contribute to acid rain in some mountain regions, and exacerbate eutrophication problems.

AP2 then connects ambient concentrations to physical impacts using peer-reviewed dose-response functions. Let β^{s}_{kp} represent the dose response coefficient which captures the effect of an incremental change in the concentrations of pollutant s on health outcome k in population cohort (p). In order to model impacts of exposure to $PM_{2.5}$ on adult mortality rates, this analysis uses the findings reported in Pope et al., (2002). The impact of $PM_{2.5}$ exposure on infant mortality rates is modeled using the results from Woodruff et al., (2006). For O_3, we use the findings from Bell et al., (2004). In addition, this analysis includes the impact of exposure to $PM_{2.5}$ on incidence rates of chronic bronchitis (Abbey et al., 1995).

The final modeling step in connecting emissions to damages translates the physical effects predicted by the dose-response functions into monetary terms. Let α_k represent the valuation coefficient that is used to translate the health outcome k into dollar terms. We rely on valuation methodologies used in the prior literature. In order to value the risk of premature mortalities due to pollution exposure, we employ the Value of a Statistical Life (VSL) method. (See Viscusi and Aldy, 2004 for a summary of this literature.) In particular, we employ a VSL of approximately $6 million; this value, which is used by US EPA, results from a meta-analysis of nearly 30 studies that compute VSLs using both stated and revealed preference methods. Further, each case of chronic bronchitis is valued at approximately $300 thousand which is also the value used by US EPA.

The marginal ($/ton) damage for NO_x for the 632 coal-fired EGUs are estimated using the marginal damage algorithm used in Muller (2011) which is based on the routine developed in Muller and Mendelsohn (2007; 2009) . This algorithm includes the following steps. First, baseline emissions are constructed from detailed emissions data collected by the US EPA in the years immediately preceding the introduction of the NOx Budget Program. These emissions reflect the NO_x controls required for all sources in non-attainment areas. AP2 computes total national damages associated with these baseline levels of NOx emissions. Next, one ton of NO_x is added to baseline emissions at a particular EGU. AP2 is then re-run. Concentrations, exposures, physical effects, and damages are recomputed. Since the only difference between the baseline run and the "add-one-ton" run is the additional ton of NO_x, the change in damages is strictly attributable to the added ton. This design is then repeated over all of the EGUs encompassed by the NBP.

The marginal damage calculation in the context of statistical uncertainty involves the following steps. First AP2 makes a random draw (denoted the t^{th} draw) from the input distributions. Next, AP2 computes m^{th} realization for emissions, concentrations, exposures, physical effects, and damages based on the realized draw from each input distribution. AP2 then adds one ton of NO_x that has been added to source (i). Again, AP2 tabulates concentrations, exposures, physical effects, and damages conditional on the added ton of NO_x at source (i), (using the same t^{th} realization from the input distributions) . AP2 computes the difference between damages with baseline emissions and after adding the ton of NO_x to (i). This is repeated 4,999 times to estimate the empirical distribution of marginal damages for NO_x emitted from facility (i). This process is then repeated for each EGU in the analysis.

Equation [16] provides a very parsimonious description of the marginal damage estimates used in our analysis:

$$\delta_j = \sum_{ri}\sum_{k}\sum_{s} \alpha_k \beta_{sp}^{s} P_{ri} \frac{dC_{srj}(e_j)}{de_j}. \tag{16}$$

Given the stochastic nature of AP2, the parameters of the atmospheric model, the population estimates P_{ri} , the dose response parameters β_{sj}^{s} and valuation parameters α are treated as being uncertain. Even the emissions levels at individual sources cannot be predicted with certainty. These multiple sources of uncertainty beget significant uncertainty in the marginal damage estimates. As described above, the AP2 model is used to compute δ_{jt}, where t indexes draws. This exercise yields an empirical distribution for each δ_j parameter.

The extent to which marginal damage estimates vary across draws is striking. Figure 3 summarizes the distribution of a single marginal damage parameter. This source, a single coal-fired electricity generating unit in Ohio, was chosen because the variance and skewness of the corresponding empirical distribution are very close to the median values across all units. The point estimate, or expected value, of the damage caused by an incremental change in emissions at this source is \$1496/ton NOx. The standard deviation is \$1796/ton. Muller (2011) finds that most of this within source variation stems from uncertainty in the air quality modeling component, adult mortality dose-response parameter estimates, and mortality valuation parameters. The skewness of the distribution stems from the multiplicative nature of the process that links emissions to damages.

Figure 4 illustrates the extent to which the expected values of source specific damage parameters $E[\delta_i]$ vary across sources. The average parameter value (averaged across all sources) is \$1711/ton of NOx. In the subsequent discussion, we classify any source with estimated damages exceeding (falling below) \$1711/ton NOx as "high" ("low") damage.

Notably, a significant amount of the inter-source variation (approximately 45 percent) occurs within (versus between) states. This suggests that a zonal trading regime that employs state-level trading ratios (and permits one-for-one trading within states) is a fairly blunt policy tool to capture heterogeneity in emissions damages.

For five of the 632 units in our data, we find that the expected value of the marginal damage parameter δ is negative. This suggests that a decrease in NOx emissions at these sources leads to increased overall damages. This result is driven by the complex, non-linear photochemical reactions that transform NO_x and VOCs into ozone. Daily ozone concentrations are non-linear and monotonic functions of NOx and the ratio of volatile organic compounds (VOCs) and NOx. At sufficiently low ratios, the conversion of NOx to ozone is limited by the availability of VOCs. In these VOC limited conditions, reductions of NOx can increase peak ozone levels until the system transitions out of a VOC-limited state (Seinfeld and Pandis, 1998).

4.1.2 Parameterizing a damage-differentiated policy

The unit-specific damage parameter estimates summarized by Figure 4 can be used to define the compliance requirements imposed in a differentiated emissions trading program. Eq. (14) implies that the constrained optimum is obtained by equating differences in marginal abatement costs with differences in expected marginal damages. To construct these source-specific "wedges", we subtract the average damage parameter ($1711) from the source-specific expected damage measures $E[\delta_i]$ Figure 5A plots these wedges versus the expected damage parameters. Relatively "high damage" units are required to pay an amount that exceeds the market clearing permit price for each unit of emissions, whereas relatively "low damage" units pay less than the permit price per ton. We assume that incentivizing pollution at facilities with negative damage parameter estimates would be politically unpopular. Instead, we exempt any units with wedges that exceed the average marginal damage value.from the differentiated policy regime.

We will also simulate outcomes under differentiated policy designs that use the first best trading ratios to define the terms of compliance. To construct these ratios, the source-specific damage measure δ_i is divided by the mean value $\bar{\delta}$. Figure 5B plots these ratios as a function of the estimated damage parameters. Relatively "high damage" units are required to hold $r_i > 1$ permit per ton of emissions under the spatially-differentiated trading counterfactual, whereas relatively "low damage" units are required to hold $r_i < 1$ permit per ton.

Both of the differentiated policy designs we consider are predicated on the assumption that the aggregate damage function $D(e)$ is linear and additively separable in source specific damages. Appendix 5 discusses the extent to which the simulation data support this assumption.

4.2 Estimating source-specific NOx abatement costs

The NBP mandated a dramatic reduction in average NOx emissions rates.[7] In the period between when the rule was upheld by the US Court of Appeals (March 2000) and the deadline for full compliance (May 2004), firms had to make costly decisions about how to comply with this new regulation. We will assume perfect compliance on behalf of all units. In fact, compliance has been close to 100 percent for the duration of the program (US EPA, 2008).

To comply, firms can do one or more of the following: purchase permits to offset emissions exceeding their allocation, install NOx control equipment, or reduce production at dirtier plants during ozone season. For the coal-fired units in our analysis, we rule out reduction in ozone season output as a compliance strategy and assume that firm-level production and aggregate output are exogenously determined and independent of the environmental compliance choice. Coal-fired units are typically inframarginal due to their relatively low fuel operating costs. Consistent with this observation, Fowlie (2010) finds that the introduction of the NBP reduced profit margins at operating units, but not production levels.

We also rule out the possibility that the emissions regulation would cause an electricity generating to exit the market prematurely. The coal plants in our analysis are long-lived. The average retirement age of a coal plant is 49 years(). We do observe a small number of coal-fired boilers retiring during the study period. These are units with decades of service stretching as far back as the end of World War II. We

[7] Pre-retrofit emissions rates at affected coal plants were, on average, three and a half times higher than the emissions rate on which the aggregate cap was based (0.15 lbs NOx/mmbtu).

omit these units from our analysis, equivalent to assuming that retirement status is exogenous to the policy design.

The specific NOx control options available to a given unit vary across units of different vintages and boiler types. In general, the more capital intensive the compliance option, the greater the emissions reductions. Compliance options that incorporate Selective Catalytic Reduction (SCR) technology, a very capital intensive post-combustion control, can reduce emissions by up to ninety percent. NOx emissions rates can be reduced by thirty-five percent through the adoption of Selective Non-Catalytic Reduction Technology (SNCR). Pre-combustion control technologies such as low NOx burners (LNB) or combustion modifications (CM) require much smaller upfront investments and can reduce emissions by fifteen to fifty percent, depending on a boiler's technical specifications and operating characteristics. Some of these technology options are physically additive. For example, SCR technologies can be combined with pre-combustion control technologies to deliver even greater emissions reductions.

Three factors that are likely to significantly influence a manager's choice of environmental compliance strategy are the up-front capital costs K, the anticipated variable operating costs V, and the expected emissions rate m. The capital costs, variable operating costs, and emissions reduction efficiencies associated with different compliance alternatives vary significantly, both across NOx control technologies and across generating units with different technical characteristics. We do not directly observe the variable compliance costs and fixed capital costs or the post-retrofit emissions rates that plant managers anticipated when making their decisions. We can, however, generate detailed, unit-specific engineering estimates of these variables using detailed unit-level and plant-level data.

In the late 1990s, to help generators prepare to comply with market-based NOx regulations, the Electric Power Research Institute[8] developed software to identify all major NOx control options (including combinations of control technologies) available to coal-fired boilers, conditional on unit and plant level characteristics. The software has been used not only by plant managers, but also by regulators to evaluate proposed compliance costs for the utilities they regulate (Himes, 2004; Musatti, 2004; Srivastava, 2004). This software was used to generate the unit-specific cost estimates used in this analysis (EPRI, 1999b). This cost estimation exercise is described in detail in Fowlie (2010).

Table 1 presents summary statistics for unit-level operating characteristics that significantly determine NOx emissions levels. To construct this table, units are classified as either "high damage" (above average) or "low damage" (below average) units. This damage classification is described in detail in section 3.1. Overall, these unit-level characteristics are very similarly distributed similarly across the two groups.

Table 2 presents means and standard deviations of the capital and variable costs (estimated at the unit level) for the most commonly adopted NOx control technologies. It is important to emphasize that these are estimates of the ex ante expected (versus ex post realized) abatement costs.

[8]The Electric Power Research Institute (EPRI) is an organization that was created and is funded by public and private electric utilities to conduct electricity industry relevant R&D.

4.3 Simulating facility-level compliance decisions

The next step in the analysis is to simulate firms' compliance choices under differentiated and undifferentiated policy regimes. We do this in two complementary ways.

4.3.1 Cost-minimization algorithm

In the ex ante analysis that informs the design and implementation of market-based emissions regulations, it is standard to assume strict cost minimization on behalf of all firms (IPM, ISIS). Our first approach to simulating firms' compliance decisions under the observed and counterfactual policy designs is based on this assumption. To simulate firms' compliance decisions, we use a simple algorithm to find the combination of NOx control options that minimizes the cost of meeting the emissions cap.

Let $j = 1...J_i$ index the NOx control technology options available to the ith electricity generating unit. Let K_{ij} represent the engineering cost estimates of required capital investments specific to unit i and technology j; V_{ij} is the corresponding variable operating cost estimate (per kWh) and m_{ij} represents the corresponding post-retrofit emissions rate. Let e_{i0} represent the pre-retrofit emissions rate; this is the amount of NOx the ith unit emits per kWh of electricity generated if it installs no new pollution controls.

In the baseline, undifferentiated policy regime, we calculate the ex ante expected annual compliance cost associated with unit i and compliance strategy j as follows:

$$\begin{aligned}\min_j \quad &: \quad C_{ij} = v_{ij}Q_i + l_iK_{ij} \\ v_{ij} \quad &= \quad (V_{ij} + \tau m_{ij})Q_i.\end{aligned}$$

Capital investments K_{ij} are converted to annual costs using a levelized annual cost factor l_i. Unit-specific time horizons are constructed by subtracting the unit age from the assumed life span (50 years).[9] This discount rate we use, 5.34 percent, was derived from financial data for electric utilities [10] Expected annual compliance costs v_{ij} are obtained by multiplying estimated per kWh variable compliance costs by expected seasonal production Q_i. Historic electricity production during the ozone season, Q_i, is used to proxy for expected ozone season production.[11]

To simulate outcomes under the differentiated policy that incorporates first-best ratios, variable compliance costs are redefined as $(V_{ij} + \tau r_i m_{ij})Q_i$. In the differentiated regime that incorporates wedges, the variable compliance costs are $(V_{ij} + \tau m_{ij} + w_i m_{ij})Q_i$.

[9]Note that, in addition to treating the retirement decision as exogenous, we are not attributing any costs to NOx reductions from the new plants replacing these units once they retire. This is equivalent to assuming that the new capacity investment will comply with new source standards, and that the cap will cease to bind as these new plants make up a larger share of the fleet.

[10]This rate is used in the US EPA Integrated Planning Modeling of investments in environmental retrofits.For more complete documentation of this model and its assumptions, see http://www.epa.gov/airmarkt/progsregs/epa-ipm/.

[11]Anecdotal evidence suggests that managers used past summer production levels to estimate future production (EPRI, 1999a). We adopt this approach and use the historical average of a unit's past summer production levels ($\bar{Q}_n$) to proxy for expected ozone season production.

The process for simulating permit market clearing is as follows. The cap is set equal to the seasonal NOx emissions associated with the compliance choices we actually observe (597,000 tons).[12] Beginning with an initial permit price τ^0, we find the compliance option j^0 at each electricity generating unit that minimizes (). The ozone season NOx emissions associated with these choices are summed across units. If these aggregate emissions exceed (fall below) the cap, the permit price is incrementally increased (reduced) and the process is repeated until the aggregate emissions constraint is just satisfied.[13]

Table 3 summarizes the compliance choices associated with the simulated equilibrium under the observed (undifferentiated) policy regime. This table shows that the cost minimization algorithm poorly predicts the compliance choices that firms actually made. In particular, the model overestimates the share of less capital intensive combustion modifications and underestimates the share of capital intensive SCR retrofits.

We are not the first to find that observed outcomes in an emissions permit market depart markedly from those consistent with strict cost minimization (see, for example, Carlson et al., 2000). If cost minimization algorithms fail to capture real world distortions and idiosyncracies, they will likely fail to accurately predict firms' compliance decisions in a counterfactual policy regime. For this reason, we pursue a second (preferred) approach to modeling firms' compliance decisions.

4.3.2 An econometric model of the compliance decision

Given the retrospective nature of the analysis, we can actually observe the compliance choices that firms in the NBP actually made. Fowlie (2010) estimatea an econometric model of the compliance choices made by plant managers in the NBP. Let j index the compliance strategy alternatives that are technologically feasible at unit i, $j,\ j = 1...J_i$. The decision maker at unit i is assumed to choose the compliance strategy that minimizes the unobserved latent value C_{ij} :

$$\begin{aligned} C_{ij} &= \alpha_j + \beta_m^v v_{ij} + \beta_m^K K_{ij} + \beta^{KA} K_{ij} \cdot Age_{ij} + \varepsilon_{ij}, \\ where \quad v_{ij} &= (V_{ij} + \tau m_{ij}) Q_i \end{aligned} \tag{17}$$

The deterministic component of C_{ij} is a weighted sum of expected annual compliance costs v_{ij}, the expected capital costs K_{ij} associated with initial retrofit and technology installation, and a constant term α_j that varies across technology types. The technology fixed effects are intended to capture average biases for or against particular types of NOx control equipment. An interaction term between capital costs and demeaned plant age is included in the model because older plants can be expected to weigh

[12] The estimated emissions associated with observed compliance choices exceed the emissions levels that were actually observed. In 2004, the first year of full compliance, NOx emissions from coal units were 564,000 tons. Emissions levels dropped below 500,000 tons in later years (US EPA, 2007). One possible explanation for this discrepancy is that many units that made no capital investment in abatement equipment were able to make extensive small-scale improvements to reduce emissions intensity. Linn (2008) estimates that 10-15 percent of emissions reductions were the product of these small process changes and modifications.

[13] If this iterative procedure arrives at a point where it is vascillating around the cap, the price that delivers the quantity of emissions just below the cap is chosen to be the equilibrium price. Equilibrium emissions are calculated and the simulation stops.

capital costs more heavily as they have less time to recover these costs. Expected annual compliance costs are obtained by multiplying estimated per kWh variable costs by expected seasonal production Q_i. We maintain the assumption that expected seasonal electricity production (Q_n) is independent of the compliance strategy being evaluated.

With some additional assumptions, this model can be implemented empirically as a random-coefficients logit (RCL) model. More specifically, the ε_{nj} are assumed to be *iid* extreme value and independent of the covariates in the model. The variable cost coefficient (β^v) and the capital cost coefficient (β^K) are allowed to vary randomly in the population according to a bivariate normal distribution, thereby accommodating any unobserved heterogeneity in responses to changes in compliance costs.[14] The econometric model is estimated separately for units serving restructured wholesale electricity markets versus publicly owned units and units subject to cost-of-service regulation. A more detailed description of the econometric specification and estimation results can be found in Fowlie (2010).

An electricity generating facility or "plant" can consist of several physically independent generating units, each comprising of a boiler (or boilers) and a generator. The 632 boilers in our data represent 221 power plants. The same plant managers make compliance decisions for all boilers at a given plant. To accommodate correlation across choices made by the same plant managers, the β_m coefficients are allowed to vary across managers according to the density $f(\beta|b, \Omega)$, but are assumed to be constant across choices made by the same manager. Estimates of the parameters of the distribution of β^v and β^K in the population of managers can be combined with information about observed choices in order to make inferences about where in the population distribution a particular decision maker most likely lies (Allenby and Rossi, 1999; Revelt and Train, 2000; Train, 2003). Using the means of these plant-manager specific distributions, versus the population means, should improve our ability to simulate the choices that these plant managers would have made in counterfactual policy scenarios.

Table 4 summarizes the the parameter estimates that define the policy simulations. The top panel reports the estimated technology specific fixed effects. These are all negative, suggesting that the average plant manager was biased against emissions abatement technology retrofits (vis a vis the compliance option that relies exclusively on purchasing permits).

The bottom panel reports the means of the manager-specific distributions of the two cost coefficients (β^K and β^v).The ratio $(\beta^K + \beta^{KA} Age) : \beta^v$ is of particular interest as it can be interpreted as a measure of how a plant manager trades off fixed capital costs (i.e. investments in NOx control equipment) and variable compliance costs (including the cost of holding permits to offset uncontrolled emissions each year). On average, this willingness to pay to reduce variable annual compliance costs is well below the levelized annual cost factor used to parameterize (17).

With this econometrically estimated compliance choice model in hand, our approach to simulating permit market outcomes is mechanically very similar to cost minimization exercise described above. We begin

[14] It is common in the literature to assume that cost coefficients are lognormally distributed, so as to ensure the a priori expected negative domain for the distribution (with costs entering the model as negative numbers). Model specifications that assumed a log-normal distribution for cost coefficient failed to converge.

with an initial permit price τ^0. Given this permit price, the coefficient vector β_m, and the characteristics of the choices available to the unit, choice probabilities are approximated. Unit-level emissions are calculated as the product of these choice probabilities and the corresponding ozone season emissions (measured in lbs of NOx). Emissions are summed across units. If the total quantity of emissions exceeds (is less than) the cap, τ is increased (decreased) by $0.01. This process iterates until aggregate emissions just equals the cap.[15] The resulting allocation of permitted emissions defines the equilibrium outcome.

The bottom panel of Table 3 summarizes the equilibrium choice probabilities under the observed (undifferentiated) policy regime . The simulation correctly predicts approximately 80 percent of compliance decisions.[16] Whereas the cost minimization model predicts that the emissions cap would be met with hundreds of relatively small investments in pre-combustion controls and modifications, the econometric model correctly predicts that a significant portion of the mandated emissions reductions is achieved using more capital intensive compliance strategies. One possible explanation for the apparent over-investment in more capital intensive compliance options could be the regulatory incentives faced by public power authorities and plants operating under rate-of-return regulation (Fowlie, 2010; Sotkiewitz, 2008),

To simulate outcomes under the counterfactual, differentiated trading regimes, variable compliance costs are redefined so as to incorporate source specific ratios r_i or wedges w_i. This assumes that the fundamental structure of the firm-level compliance decisions we model (and in particular, the relative weighting of capital costs and variable operating costs) would not change under a differentiated regime.

4.4 Estimating abatement costs

Abatement costs are calculated by summing the levelized annual capital costs and variable operating costs associated with the simulated compliance choices. When the econometric model is used to simulate policy outcomes, abatement costs are constructed by multiplying choice probabilities with the corresponding capital and variable operating costs. To estimate the cost of differentiation, the total abatement costs associated with the equilibrium outcome under the undifferentiated policy is subtracted from the abatement costs associated with the counterfactual policy outcome.

The unit-specific, technology specific cost estimates used in our analysis capture ex ante expected costs. These numbers capture was know in 2001 about financing, installing, and operating different NOx abatement equipment at different kinds of coal-fired boilers. This is precisely the information we want to use when modeling compliance decisions. However, these are not the ideal data to use when estimating the *ex post realized* costs. Our cost numbers do not capture financial variables (such as credit ratings) that could

[15] If this iterative procedure arrives at a point where it is vascillating around the cap, the price that delivers the quantity of emissions just below the cap is chosen to be the equilibrium price. Equilibrium emissions are calculated and the simulation stops.

[16] Where the model has the most difficulty is in correctly distinguishing between relatively low cost combustion modifications and those units that report no retrofit. It should be noted that the line between these two classes of compliance decisions is not always clear. Linn (2008) estimates that 10 to 15 percent of emissions reductions were the result of small-scale changes that do not require any significant capital investments. We are unable to observe minor process changes in the data as these are not explicitly reported. This will lead us to overestimate emissions at some units that report no compliance retrofits.

significantly affect the cost of financing large capital investments in abatement equipment. Moreover, ex ante expected costs may inaccurately reflect the costs that were actually realized (or would have been realized) if and when a plant manager chose to implement a particular compliance strategy.

There is some anecdotal evidence to suggest that ex post realized NOx abatement costs may have been lower than expected. There is also evidence to suggest that plant managers found low cost ways to reduce NOx emissions rates that did not involve significant investment or modification (Linn, 2006). Consequently, our estimates of the costs of complying with the observed, undifferentiated policy are likely overestimates. Absent additional information on which of our ex ante cost estimates are more or less accurate, we cannot say whether our estimates of the incremental costs of differentiation are biased up or down.

4.5 Simulated damages

AP2 is also used to quantify the change in damages due to the various policy scenarios explored in the study. In this context, rather than systematically perturbing NO_x emissions one source at-a-time, NO_x emissions change simultaneously at many of the regulated EGUs in response to the different modeled policies. Here, a vector of NO_x emissions corresponding to the output from the econometric cost model is processed by AP2. The resulting damages associated with both O_3 and $PM_{2.5}$ exposure are computed. It is important to note that for each policy scenario, total NO_x emissions are held fixed. What varies is the allocation of emissions across the regulated EGUs. Therefore, any difference in damages found to occur between the policy scenarios is attributable to the spatial redistribution of emissions (rather than a change in the overall stringency of the policies).

The damage simulations make use of the Monte Carlo algorithm built into AP2 (see above). However, the algorithm is adapted to accept simultaneous emission changes at many EGUs, rather than the marginal damage procedure which is executed offline to produce the marginal damage estimates used in setting the trading ratios. Specifically, each emissions vector (corresponding to a particular policy scenario) is processed through the Monte Carlo machinery. This yields a distribution of estimates of the damages stemming from each policy design. The simulations feature all input parameters (emissions, transfer coefficients in the stochastic air quality model, population, dose-response, and valuation) as random variables consisting of 5,000 possible realizations. Conditional on a policy emissions vector, one realization is selected from each input distribution, and the total exposures, physical effects, and monetary damages are computed. This is repeated 5,000 times for each policy scenario. And, to ensure a rigorous comparison of outcomes from each policy scenario, the resultant damages are compared for the same draws from the input distribution.

As an additional robustness check for our assumption in section 2. regarding the additively separable and linear NO_x damage function, we also compute total damages by simply multiplying marginal damages times emissions at each facility. This is repeated 5,000 times since the marginal damages also consist of 5,000 estimates for each EGU. In comparing the estimated welfare impacts across policy scenarios using the margin-times emission approach against the simulations that allow emissions to change simultaneously, we are careful to align the 5,000 draws such that the same "states-of-the-world" are compared to one another.

5 Results and synthesis

This section is divided into three subsections. The first characterizes the simulated outcomes (including the allocation of permitted emissions across sources, abatement costs incurred, and total damages avoided) under the observed, undifferentiated policy regime. The second section provides estimates the benefits and costs of differentiation relative to this benchmark. The third examines the uncertainty in these estimates.

5.1 Benchmark case: Undifferentiated trading

Table 5 summarizes simulated outcomes under the observed, undifferentiated regime. Results generated using the cost minimization algorithm and econometrically estimated model are reported in column (1) and (2), respectively. These outcomes serve as a benchmark for the subsequent analysis of the counterfactual, differentiated policy regimes.

The differences between these two sets of results are quite striking. First, note that the NOx permit price required to incentivize the mandated level of abatement is more than twice as high when the econometric model is used to simulate facility-level compliance choices. To understand why, recall that the econometric model captures both negative biases against specific NOx control technologies (in the negative technology fixed effects) and a reluctance on the part of many firms to take on large capital investments in exchange for reducing annual variable compliance costs (in the relative size of the β_k and β_v coefficients). The combined effect of these two factors, which are not captured in the cost minimization model, is a much smaller response (in terms of a firms' choice of emissions level) to a given change in the permit price. The higher permit price in column (2) reflects a higher marginal abatement cost (as perceived by the firms making these compliance decisions) as compared to the simulations that mechanistically minimize our assumed, ex ante expected costs.

The suite of abatement technologies that are used to meet the emissions cap look quite different across the two sets of simulations (see Table 4). Any deviation from the choices predicted using the cost minimization model will, by definition, increase our measure of the overall (ex ante expected) cost of meeting the emissions cap. Table 5 shows how these differences in compliance choices translate into significant differences in estimated compliance costs. Levelized annual compliance costs are estimated to be almost 50 percent higher when the econometric model of the compliance choice is used.

Differences in the suite of abatement technologies deployed to achieve the mandated emissions reductions also imply differences in the allocation of permitted emissions across sources (and differences in the aggregate estimates of avoided damages). Facility-level data collected following the introduction of the NBP indicate that approximately 38 percent of permitted NOx emissions occurred at sources with higher than average damage parameters. The cost minimization model over-predicts the share of emissions occurring at these "high damage" sources (42 percent). The econometric model allocates 39 percent of permitted emissions to high damage firms. Consequently, the estimate of the benefits (in terms of avoided damages) achieved by the NOx Budget Program is higher when the empirical model is used to simulate outcomes.

Our estimate of the net benefits (i.e. the value of avoided damages less abatement costs) conferred by the NOx Budget Program is $383,000,000 per year. To the extent that our ex ante cost estimates overestimate the ex post realized costs, this net benefit estimate will be conservative. The estimate obtained using the model that assumes strict cost minimization is $573,000,000 per year.

5.2 Policy counterfactuals

We consider two differentiated policy designs: one based on the "wedges" introduced above, and one that incorporates first-best ratios.

Figure 6 provides a graphical summary of the extent to which policy differentiation shifts emissions from high to low damage sources. The vertical axes in these figures measure the change in emissions (in percentage terms) moving from an undifferentiated to a differentiated policy. A positive percentage change indicates that the emissions level chosen in the undifferentiated regime exceeds the emissions level chosen in the differentiated regime. Equation (7) suggests these changes should depend on the degree of heterogeneity in the source specific damage parameters (which is significant in the NBP) and the steepness of the marginal abatement cost curves.

The horizontal axis measures source-specific marginal damage estimates. The circular markers correspond to the wedge-based differentiated policy. The triangle markers correspond to the differentiated policy that incorporates first best ratios. Each marker corresponds to a different electricity generating unit. Intuitively, differentiation increases (decreases) emissions at units with below-average (above-average) damage parameters.

The left panel summarizes the emissions changes implied by the cost minimization model. Emissions changes are fairly significant, with some low (high) damage units increasing (reducing) emissions by more than 50 percent under the differentiated policy (relative to the undifferentiated regime). Taken together, policy differentiation moves an estimated 14 to 15 percent of permitted emissions from high damage sources to low damage sources.

A comparison of the right and left panels of Figure 6 introduces an important result. The simulations generated using the econometric model predict a much smaller reallocation of emissions as compared to the simulations that are based on strict cost minimization. The differentiated policy that incorporates wedges shifts only 2 percent of emissions from high to lower damage sources. To understand the reason for this relatively anemic response of emissions to policy differentiation in the left panel, it is important to remember what is held constant across these two simulation approaches, and what is allowed to differ. The emissions cap, the damage parameters and raw cost inputs (i.e. the assumed capital and variable costs K_{ij} and v_{ij}) used to parameterize both sets of simulations are identical. The basic structure of the underlying cost models are also very similar. What differs is the relative weighting of capital and variable abatement costs. The econometrically estimated equation (17) penalizes capital intensive abatement options more heavily as compared to equation (17). Loosely speaking, this implies that marginal abatement costs are more steeply sloped (because capital intensive abatement options are associated with significantly greater emissions intensity reductions). By Observation (2), steeper abatement costs imply a smaller response to policy differentiation, all else equal.

Table 6 provides a more detailed summary of the results from the simulations that assume strict cost minimization. All costs and benefits are expressed relative to the undifferentiated policy baseline. The first thing to note is that the simulated equilibrium permit prices under the differentiated regimes ($0.90-$0.95/lb NOx) are remarkably close to the average damage parameter ($0.86/lb NOx). Appendix 1 shows that the equilibrium permit price will equal the average damage parameter when the cap is optimally set. Thus, conditional our estimated marginal damage parameters δ and source-specific NOx abatement cost assumptions (i.e. the K_{ij} and v_{ij}), the emissions cap imposed under the NBP was very close to optimal.

Intuitively, because the emissions constraint imposed in these cost minimization simulations is very close to the optimal, the advantage of using wedges over first-best ratios (in terms of efficiency) is very small. Note that the wedge-based regime allocates a slightly larger quantity of emissions to high damage sources as compared to the ratio-based regime. Thus, the costs of differentiation (in terms of increased abatement costs) under the wedge-based regime are slightly lower.

Table 6 also summarizes the estimated net benefits from policy differentiation. These are positive (consistent with observation (1)) and very similar across the ratio-based and wedge-based regimes. The expected net benefits amount to approximately $90 M per year. In percentage terms, policy differentiation increases the net benefits conferred by the trading program by 14 percent.

Table 6 suggests that the efficiency gains from using wedges versus ratios are quite small. It is also important to consider the distributional implications of this design choice. Under all policy scenarios, the majority of permitted emissions occur at low damage sources. Consequently, under the differentiated policy that incorporates wedges, $\sum w_i e_i < 0$. The revenues paid to relatively low damage sources exceeds the wedge-based revenues collected from relatively high damage sources by an estimated $150M. Although this transfer from the implementing agency to low damage sources is of no consequence for our welfare calculations, it could impact the political palatability of wedges versus ratios.

Table 7 summarizes the results of the simulations that incorporate the econometric model of firms' compliance choices. All costs and benefits are expressed relative to the undifferentiated policy baseline. Because facility-level compliance decisions are less responsive to incremental changes in the cost of uncontrolled emissions (as compared to the model that assumes strict cost minimization), simulations using the econometric model predict a much smaller impact of policy differentiation on costs and benefits. The wedge-based regime increases expected program benefits (in terms of avoided emissions) by only 2 percent. The ratio-based policy increases the expected benefits of the program by an estimated 6 percent.

Perhaps the most surprising result reported in Table 6 is that the expected net benefits from differentiation appear *larger* under the design that incorporates ratios (versus wedges). This runs counter to the intuitive arguments presented in section 2 which suggest that wedges should be preferred to first-best ratios when the cap is not set optimally. How can this be?

Figure 7 provides an intuitive, albeit highly stylized, illustration of why the ratio-based regime welfare dominates the wedge-based regime in these simulations. This figure extends the simple graphical analysis conducted using Figures 1 and 2. The main difference is that we now distinguish between "true" social costs of abatement (i.e. the opportunity cost of the financing, equipment, and variable inputs to abatement activities).and the abatement costs that are perceived by the firms.

The solid lines in Figure 7 represent linear approximations to the source-specific, "true" marginal abatement cost curves. These are intended to represent, very simply, the ex ante expected cost components we use to estimate the abatement costs associated with a given set of simulated compliance choices.

The econometric model we use to simulate firms' compliance choices assumes that marginal abatement costs perceived by firms deviate systematically from these true cost curves. In particular, the econometric model assigns a cost penalty to specific NOx control technologies and, for most managers, penalizes capital intensive options relatively more heavily. For now, we assume that pre-existing distortions or imperfections in the product market drive an artificial wedge between the true costs of abatement (as perceived by society) and the abatement costs perceived by firms. The broken lines in Figure 7 represent linear approximations to these "perceived" abatement cost functions. These functions are more steeply sloped, reflecting the fact that emissions are much less responsive to changes in the permit price when compliance choices are simulated using the econometric model.

The intersection of the broken lines denotes the undifferentiated trading equilibrium. Note that this allocation of permitted emissions lies to the right of the equilibrium that would be observed if firms perceptions and society's perceptions were exactly aligned. This is consistent with our finding that aggregate damages under undifferentiated trading are slightly lower when the econometric model is used to simulate compliance decisions (see Table 5).

For ease of exposition, lines denoting the true marginal damages have been removed from Figure 7. What is shown explicitly is the difference between source-specific marginal damages ($\delta_h - \delta_l$). By Eq. (10), the optimal allocation of permitted emissions, conditional on the imposed cap $\overline{E}$, occurs at the point where the difference in true marginal abatement costs equals the difference in these damage parameters. The constrained optimal allocation is denoted $\{e_h^*, e_l^*\}$. When the wedge-based policy is introduced, however, firms equate *perceived* marginal abatement costs with $(\tau + w_i)$. The equilibrium outcome is $\{e_h^w, e_l^w\}$; the constrained optimum is not obtained.

To compensate for the differences in the true (from a social perspective) and perceived (from the firm's perspective) abatement cost functions, the fully informed regulator should introduce a policy that exaggerates the differences in emissions penalties across sources with different damage parameters.[17] Note

[17] From a practical standpoint, it seems very unlikely that a policy maker would have access to reliable information regarding how firms' and society's abatement costs will diverge during the policy design stage.

that the first-best ratios will exaggerate these differences vis a vis wedges when the emissions cap is too stringent. This explains why the net benefits of differentiation can be larger under the ratio-based regime, as compared to the wedge-based regime. The use of ratios, versus wedges, helps to mitigate the pre-existing distortions in firms' perceived costs of abatement.

Figure 7 provides a clean and intuitive illustration of why our estimated net benefits under the ratio policy exceed net benefits under the wedge-based policy. But this figure, and our approach to measuring the efficiency gains from policy differentiation more generally, is predicated on the very strong assumption that our ex ante expected cost estimates accurately capture the costs that were (or would have been) incurred. Although there are reasons to think that firms' compliance choices were influenced by pre-existing distortions and idiosyncratic errors, it seems unlikely that all of the differences between the abatement cost estimates we construct, and those consistent with observed choices, can be attributed to artificial distortion and inefficiency. For these reasons, our cost estimates should be interpreted carefully.

5.3 Uncertainty

Thus far, our discussion of the simulation results has focused exclusively on the point estimates of the gains from policy differentiation. In this section, we characterize the uncertainty surrounding these estimates. What we are analyzing amounts to a policy lottery in which probabilities can be attached to a range of possible benefit outcomes. The simulated effects of policy are uncertain because the specific parameters of the model are uncertain. Figures 8A and 8B illustrate the distribution of our estimates of the gains from policy differentiation (i.e. the reduction in damages from the permitted emissions vis a vis the undifferentiated benchmark). Each distribution summarizes 5000 realizations of benefits, each corresponding to a different vector of source-specific damage parameters δ^r.

Observations (5) and (6) provide some intuition for what drives the significant variation in these benefit realizations. Note that the covariance between the realized damage vector δ^r and the vector of source-specific average parameters that defines the terms of compliance varies across realizations. Realized benefits from differentiation are increasing with this covariance. In cases where undifferentiated trading welfare dominates differentiated trading (i.e. realizations lying to the left of zero), the realized damages are negatively correlated with the source-specific average damage parameters. In these instances, the differentiated policy designs perversely shift damages from low to high damage sources.

In each figure, the vertical black line represents the estimated cost of differentiation (relative to the undifferentiated, benchmark regime). The area of the distribution to the right of this line is equal to the probability that the net benefits from policy differentiation are positive. We choose to summarize our results in this way in order to clearly separate the estimated benefits of differentiation from the estimated costs. Taking our estimates of costs as given, the expected net benefits of policy differentiation are positive across all scenarios we consider.

6 Conclusion

How should market-based emissions regulations, and cap and trade programs in particular, be designed and implemented when damages from emissions vary significantly across sources? To shed light on this question, we first introduce a conceptual framework that is useful for analyzing the efficiency gains from policy differentiation. We extend the theoretical work in this area in order to consider key factors that complicate real-world implementation of differentiated policy. These include jurisdictional constraints, limited information about abatement costs, and uncertainty about how to value damages from pollution. We examine how differentiated policy designs can be modified so as to accommodate these constraints and limitations.

The conceptual framework serves as foundation for an applied analysis of the gains from policy differentiation. In particular, we consider the landmark NOx Budged Program. Prior research has shown that the damages due to NO_x emissions vary considerably according to where the emission occurs (source location). The policy design that is currently in place fails to appropriately capture this heterogeneity. In 2008, a federal district court vacated the rule that was to succeed the NOx Budget Program due to policy's failure to adequately accommodate spatial variation in damages.[18] Since that time, debates surrounding this program should be redesigned have become increasingly contentious.

The empirical analysis first estimates marginal damages for each of the 632 boilers regulated under the NBP. These damage estimates are used to parameterize counterfactual, spatially differentiated designs. An econometric model of the compliance decisions made by firms subject to the NBP is used to simulate outcomes under both the observed and two counterfactual policy designs (Fowlie, 2010). The corresponding aggregate abatement costs and environmental damage are then tabulated and compared. Importantly, the total emission levels between these two policy designs is held fixed; only the spatial distribution of emissions change as a function of the imposition of the trading ratios.

Our preferred estimates of the gains from policy differentiation are smaller than expected given the extent to which damages vary across sources in the program. We estimate net gains from policy differentiation on the order of $20 to $50 million per year. The benefits (in terms of reduced damages from emissions) increase by only 2 to 6 percent. We contrast these findings with a more stylized policy simulation model that assumes strict cost minimization on behalf of all firms. This optimization-based approach to simulating policy outcomes predicts larger gains from differentiation, with benefits on the order of 16-18 percent. We argue that the more stylized model likely overestimates how firm-level compliance decisions would respond to differentiated policy incentives.

Finally, we are able to account for parameter uncertainty in the modeling of damages from pollution. We explore the implications of this uncertainty for both policy design and policy analysis. Uncertainty in the damage parameters begets significant uncertainty in our estimated gains from differentiation.

[18] The court found that the CAIR regulation "does not prohibit polluting sources within an upwind state from preventing attainment of National ambient air quality standards in downwind states." *State of North Carolina v. Environmental Protection Agency, No. 05-1244,* slip op. (2008), District of Columbia Court of Appeals.

References

[1] Abbey, D.E., B.E. Ostro, F.Peterson, R.J. Burchette, 1995. "Chronic Respiratory Symptoms Associated with Estimated Long-Term Ambient Concentrations of Fine Particulates Less than 2.5 Microns in Aerodynamic Diameter and other Air Pollutants"Journal Expo Analysis Environmental Epidemiology. 5:2., pp 137-159.

[2] Banzhaf, H. Spencer, Dallas Burtraw. Karen Palmer. 2004. "Efficient emission fees in the US electricity sector." Resource and Energy Economics. 26: 317-341.

[3] Banzhaf,H.S., B. A. Chupp, "Heterogeneous Harm vs. Spatial Spillovers: Environmental Federalism and US Air Pollution, NBER Working Paper No. 15666, January 2010

[4] Baumol, William J. and Wallace E. Oates. 1988. *The Theory of Environmental Policy.* 2nd ed. Cambridge: Cambridge University Press.

[5] Bell, Michelle L., Adrian McDermott, Scott L. Zeger, Jonathan M. Samet, and Francesca Domenici. 2004. "Ozone and Short-Term Mortality in 95 US Urban Communities, 1987-2000." *Journal of the American Medical Association,* 292(19):2372-2378.

[6] Byun, Daewon, W., and Kenneth L. Schere. 2006. "Review of the Governing Equations, Computational Algorithms, and Other Components of the Models-3 Community Multiscale Air Quality (CMAQ) Modeling System." *Applied Mechanics Reviews,* 59(2):51-77.

[7] Carlson, Carlson, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer. 2000. "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?," *Journal of Political Economy*, University of Chicago Press, vol. 108(6): 1292-1326.

[8] Dales, John H. 1968. *Pollution, Property & Prices: An Essay in Policy-making and Economics.* Toronto: University of Toronto Press.

[9] Farrow, R. Scott, Martin T. Schultz, Pinar Celikkol, and George L. Van Houtven. 2005. "Pollution Trading in Water Quality Limited Areas: Use of Benefits Assessment and cost effective Trading Ratios." *Land Economics,* 81(2):191-205.

[10] Hoel, M., L. Karp, 2000. Taxes and Quotas for a Stock Pollutant with Multiplicative Uncertainty. UC Berkeley: Department of Agricultural and Resource Economics, UCB. CUDARE Working Paper 870. Retrieved from: http://escholarship.org/uc/item/9v86p5s7

[11] Krupnick, Alan, Virginia McConnell, Terrell Stoessell, Matthew Cannon, and Michael Batz. 2000. "Cost-Effective NOx Control in the Eastern United States," Discussion Papers dp-00-18, Resources For the Future.

[12] Klaassen, Ger A.J., Finn R. Forsund, and Markus Amann. 1994. "Emissions Trading in Europe with an Exchange Rate." *Environmental and Resource Economics,* 4:305-330.

[13] Lippman, Morton, and Richard B. Schlessinger. 2000. "Toxicological Bases for the Setting of Health-Related Air Pollution Standards." *Annual Review of Public Health*, 21: 309-333.

[14] Linn, Joshua. 2008. "Technological Modifications in the Nitrogen Oxides Tradable Permit Program." *The Energy Journal* , 29(3):153-176.

[15] Mauzerall, Denise, Babar Sultan, Namsoug Kim, and David F. Bradford. 2005. "NO_x Emissions from Large Point Sources: Variability in Ozone Production, Resulting Health Damages and Economic Costs." *Atmospheric Environment,* 39: 2851-2866.

[16] Mendelsohn, Robert O. 1986. "Regulating Heterogeneous Emissions." *Journal of Environmental Economics and Management*,13:301- 312.

[17] Montgomery, W. David. 1972. "Markets in Licenses and Efficient Pollution Control Programs." *Journal of Economic Theory,* 5:395-418.

[18] Muller, Nicholas Z. and Robert O. Mendelsohn. 2007. "Measuring the Damages from Air Pollution in the United States." *Journal of Environmental Economics and Management*, 54(1):1-14.

[19] Muller, Nicholas Z. and Robert O. Mendelsohn. 2009. "Efficient Pollution Control: Getting the Prices Right." *American Economic Review* 99(5):1714-1739.

[20] Muller, Nicholas Z. 2011 "Linking Policy to Statistical Uncertainty in Air Pollution Damages." The B.E. Press Journal of Economic Analysis & Policy. Vol. 11 : Iss. 1 (Contributions), Article 32.

[21] Pope, C. Arden, Richard T. Burnett, Michael J. Thun, Eugenia E. Calle, Daniel Krewski, Kazuhiko Ito, and George D. Thurston. 2002. "Lung Cancer, Cardiopulmonary Mortality, and Long-Term Exposure to Fine Particulate Air Pollution." *Journal of the American Medical Association,* 287(9):1132-1141.

[22] Seinfeld, J.H., S.N. Pandis. 1998 "Atmospheric Chemistry and Physics." John Wiley & Sons, Inc., NY, NY, USA.

[23] Stavins, Robert N., 2005. "Lessons Learned from SO2 Allowance Trading," *Choices*, Agricultural and Applied Economics Association, vol. 20(1).

[24] Tong, Daniel Q., Nicholas Z. Muller, Denise L. Mauzerall, and Robert O. Mendelsohn. 2006. "Integrated Assessment of the Spatial Variability of Ozone Impacts from Emissions of Nitrogen Oxides," *Environmental Science and Technology,* 40(5):1395-1400.

[25] Turner, D.B., 1994. "Workbook of Atmospheric Dispersion Estimates - An Introduction to Dispersion Modeling." 2nd Ed. Lewis Publishers, Ann Arbor, MI, USA.

[26] US Environmental Protection Agency (USEPA). 2006. *National Emissions Inventory (NEI): 2002.* Washington, DC: Office of Air Quality Planning and Standards, Emissions Inventory Group; Emissions, Monitoring, and Analysis Division.

[27] US Environmental Protection Agency. 2003. *Documentation Supplement for EPA Modeling Applications (V.2.1.6) Using the Integrated Planning Model.* Office of Air and Radiation. Washington D.C.

[28] US Environmental Protection Agency. 1999. *Regulatory Impact Analysis for the Final Section 126 Petition Rule*, Office of Air and Radiation, Washington DC.

[29] US Environmental Protection Agency. 1998a. *Regulatory Impact Analysis for the NBP*, Office of Air and Radiation, Washington DC.

[30] US Environmental Protection Agency (USEPA). 1999. *The Benefits and Costs of the Clean Air Act: 1990 - 2010.* EPA Report to Congress. EPA 410-R-99-001, Washington, DC, Office of Air and Radiation, Office of Policy.

[31] Viscusi, W. Kip and Joseph E. Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World." *Journal of Risk and Uncertainty,* 27(1): 5-76.

[32] Woodruff, Tracey J., Jennifer D. Parker, and Kenneth C. Schoendorf. 2006. "Fine Particulate Matter ($PM_{2.5}$) Air Pollution and Selected Causes of Postneonatal Infant Mortality in California" *Environmental Health Perspectives,* 114(5):786-790.

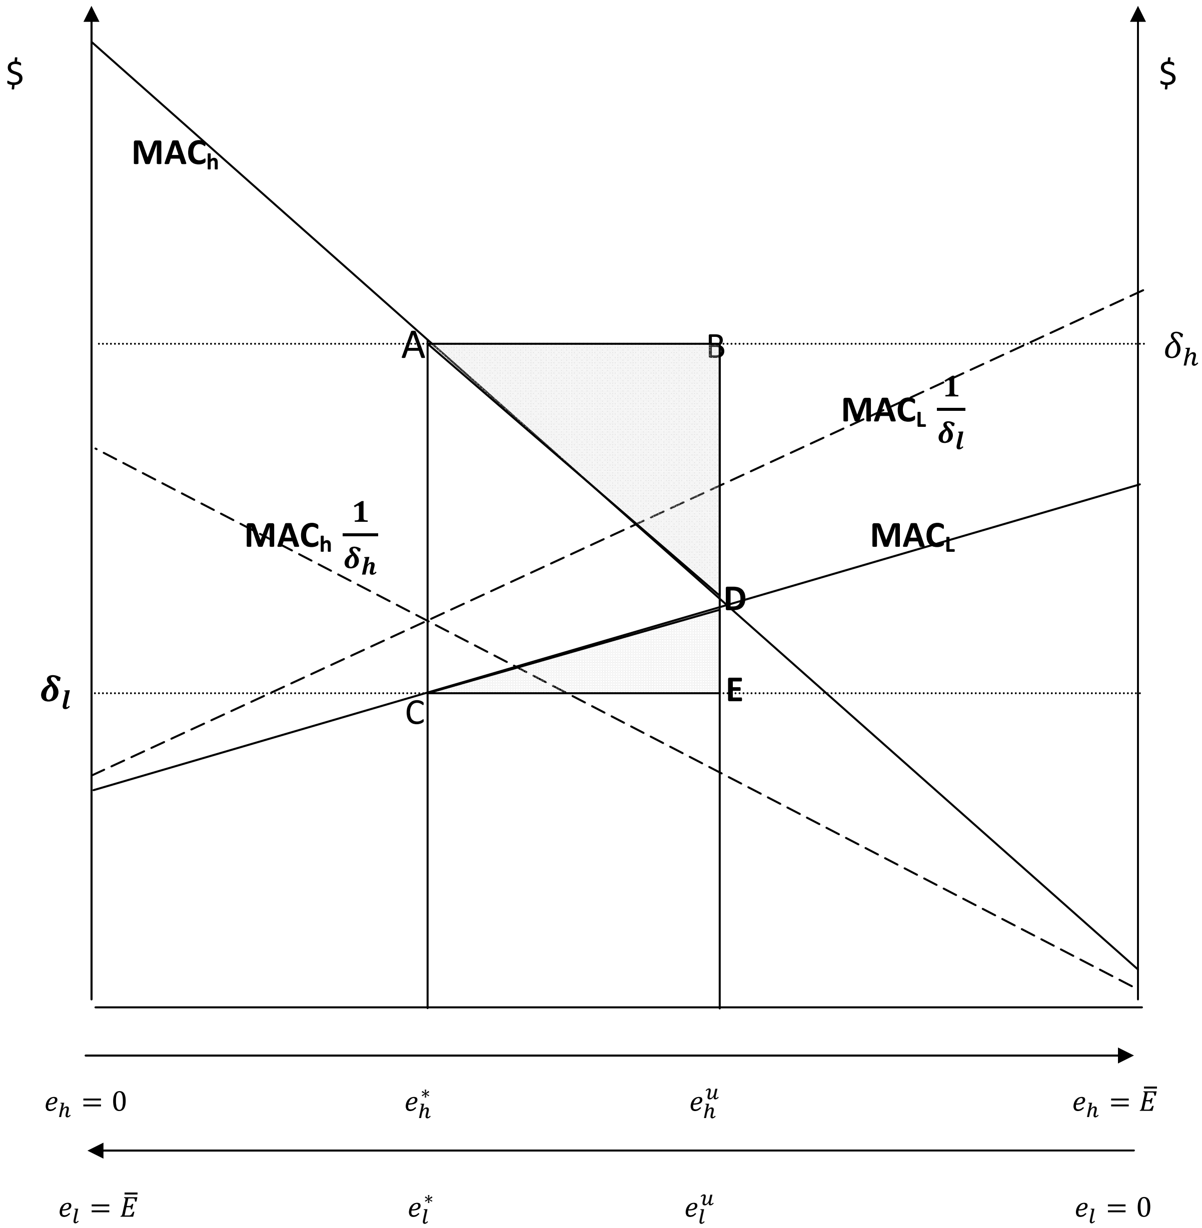

Figure 1: Emissions permit market outcomes under differentiated and undifferentiated policies: Optimal emissions constraint

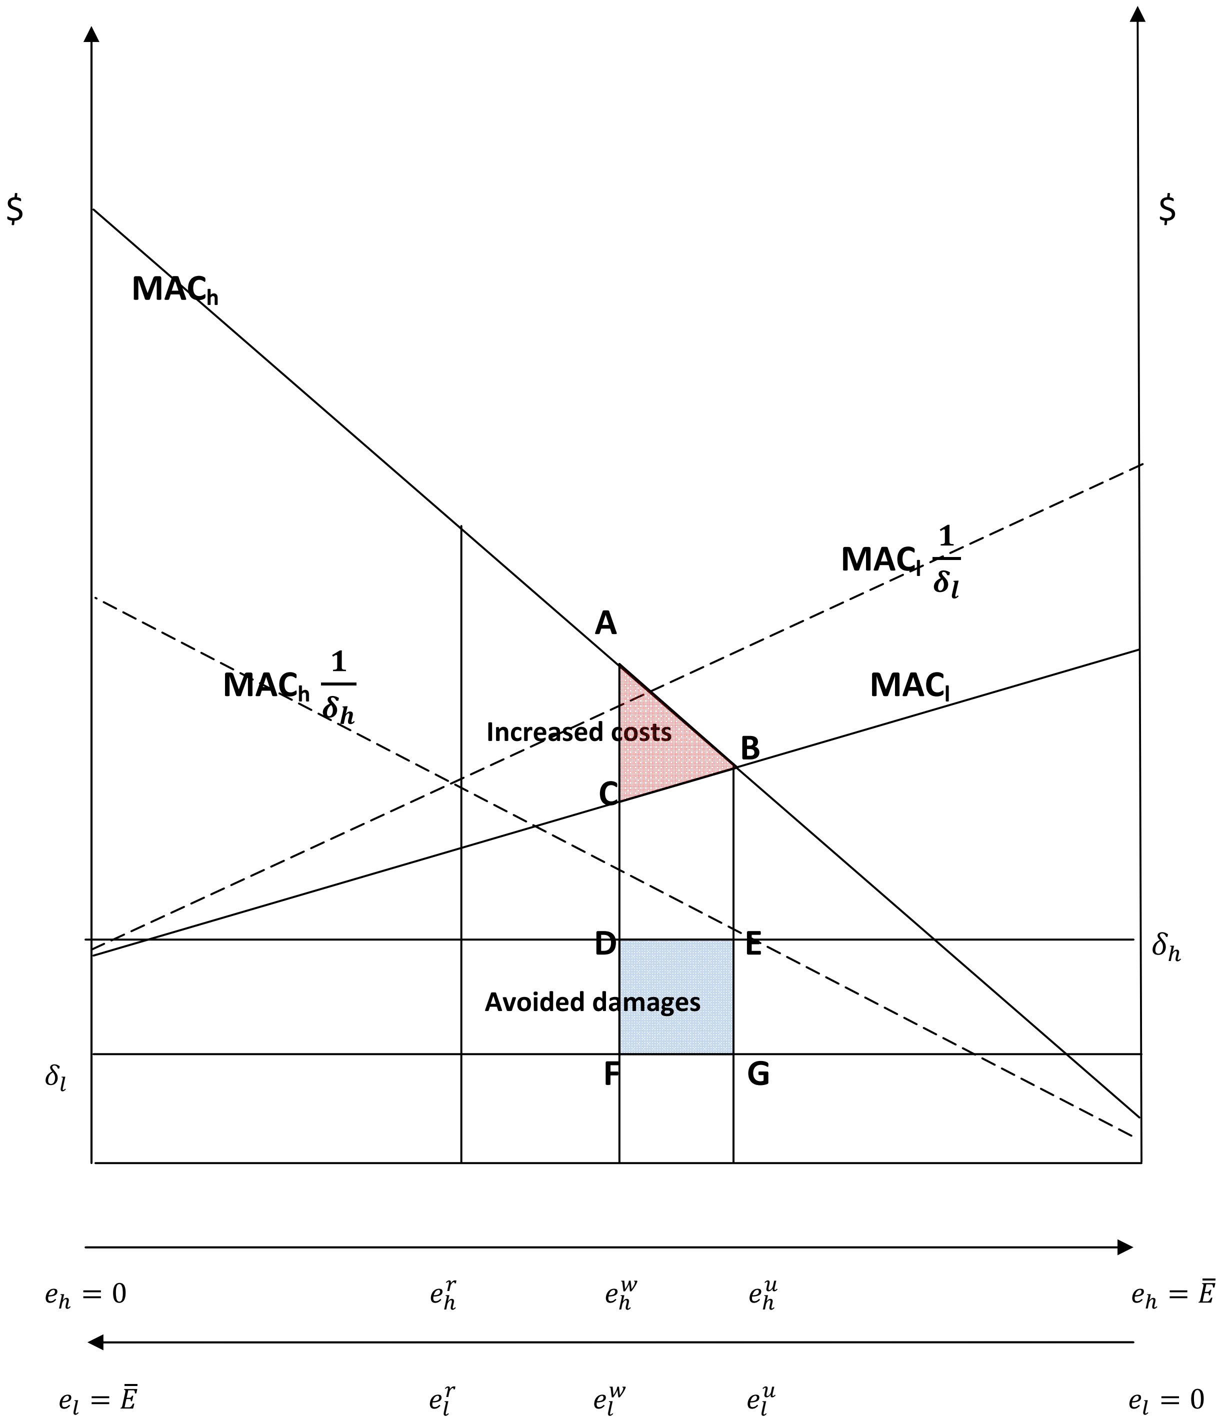

Figure 2: Emissions permit market outcomes under differentiated and undifferentiated policies: Sub-optimal emissions constraint

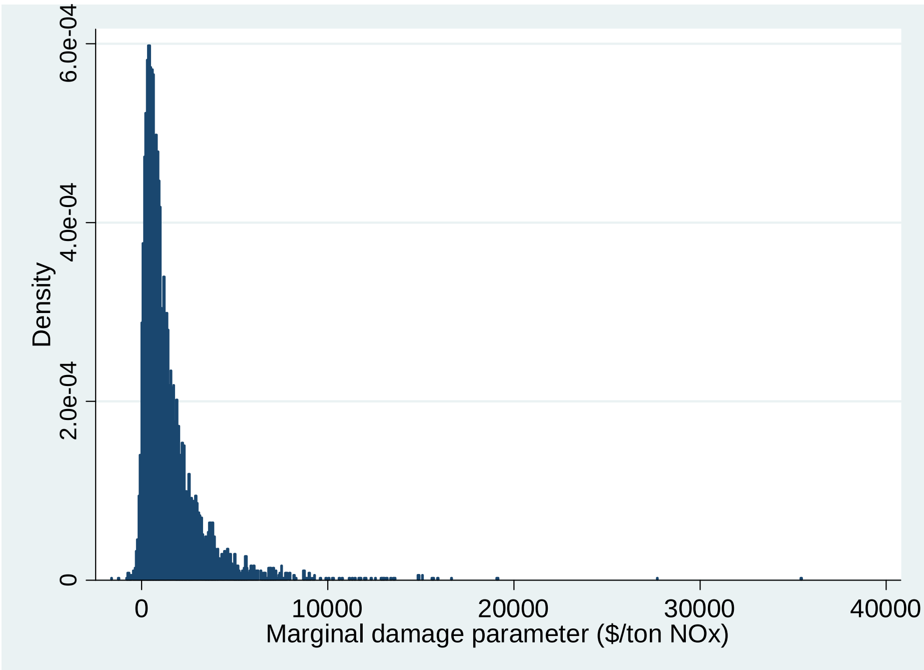

Figure 3 : Representative distribution of source-specific marginal damages

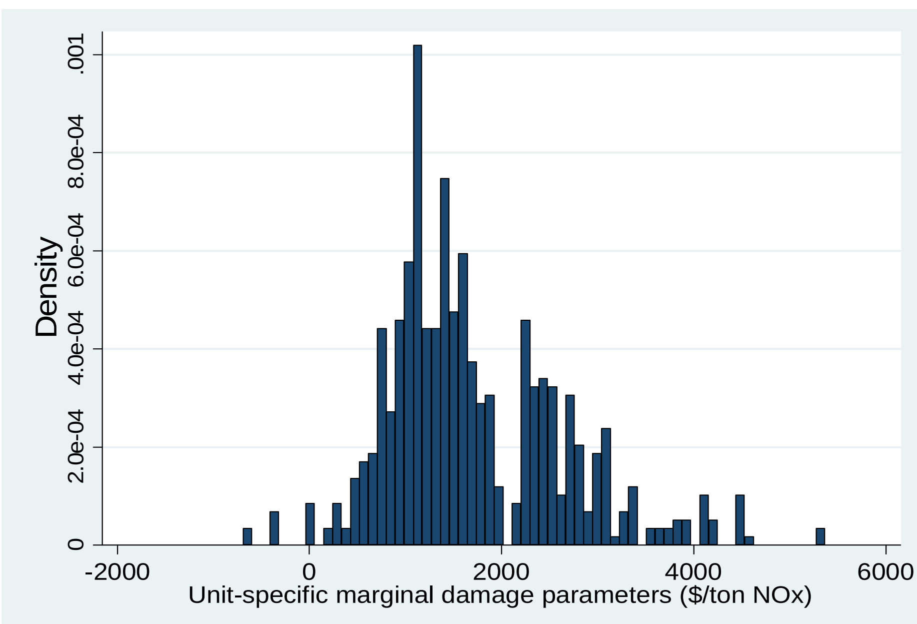

Figure 4: Histogram of damage parameter point estimates

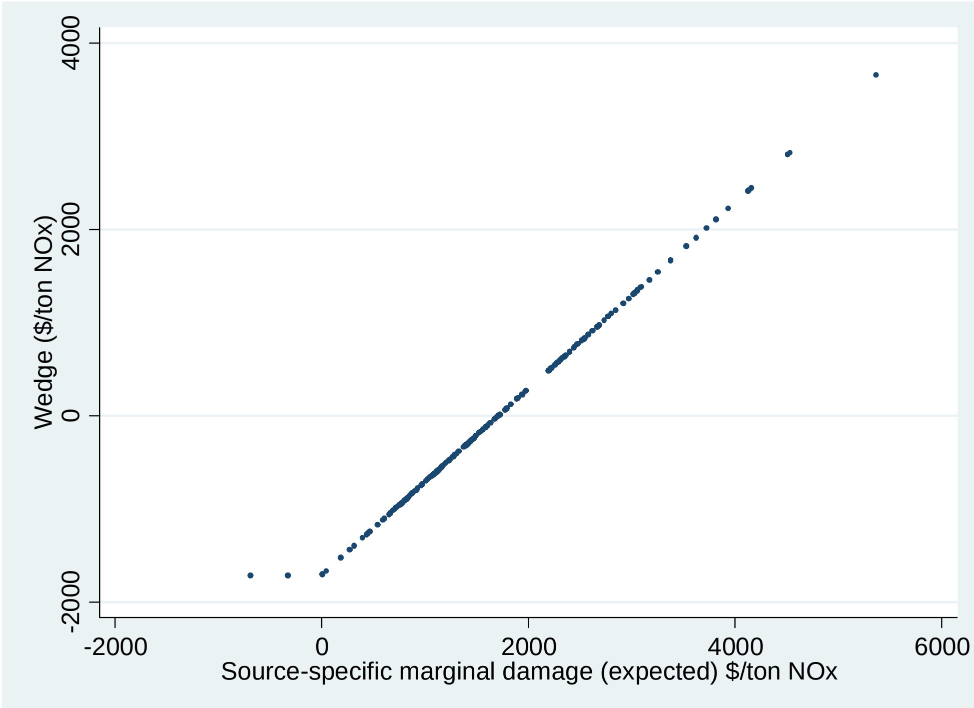

Figure 5A: Trading wedges used in counterfactual policy simulations

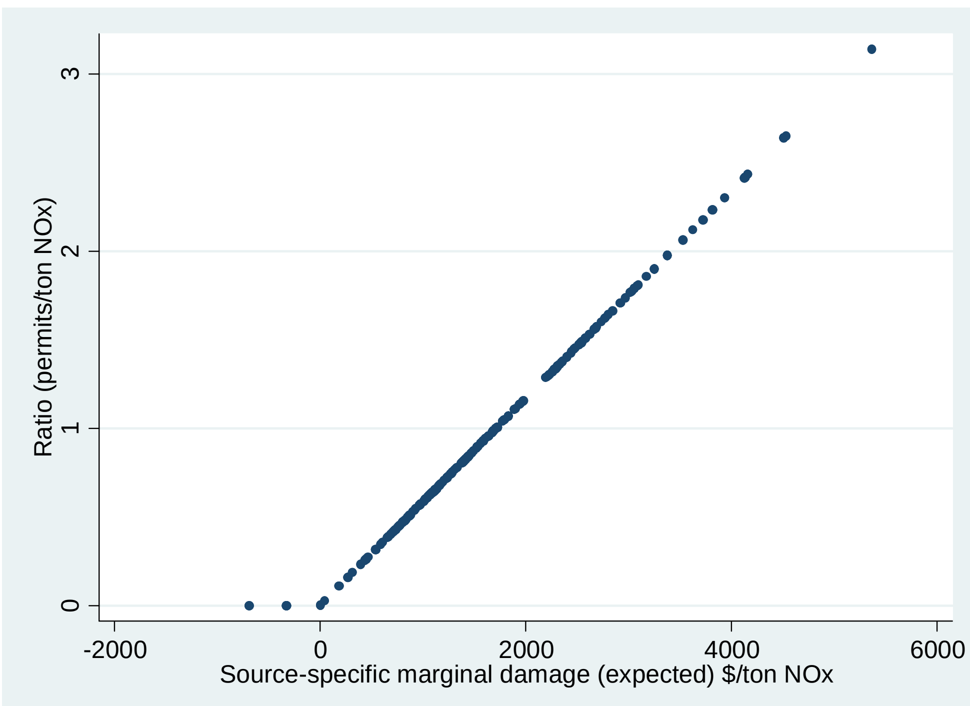

Figure 5B: Trading ratios used in counterfactual policy simulations

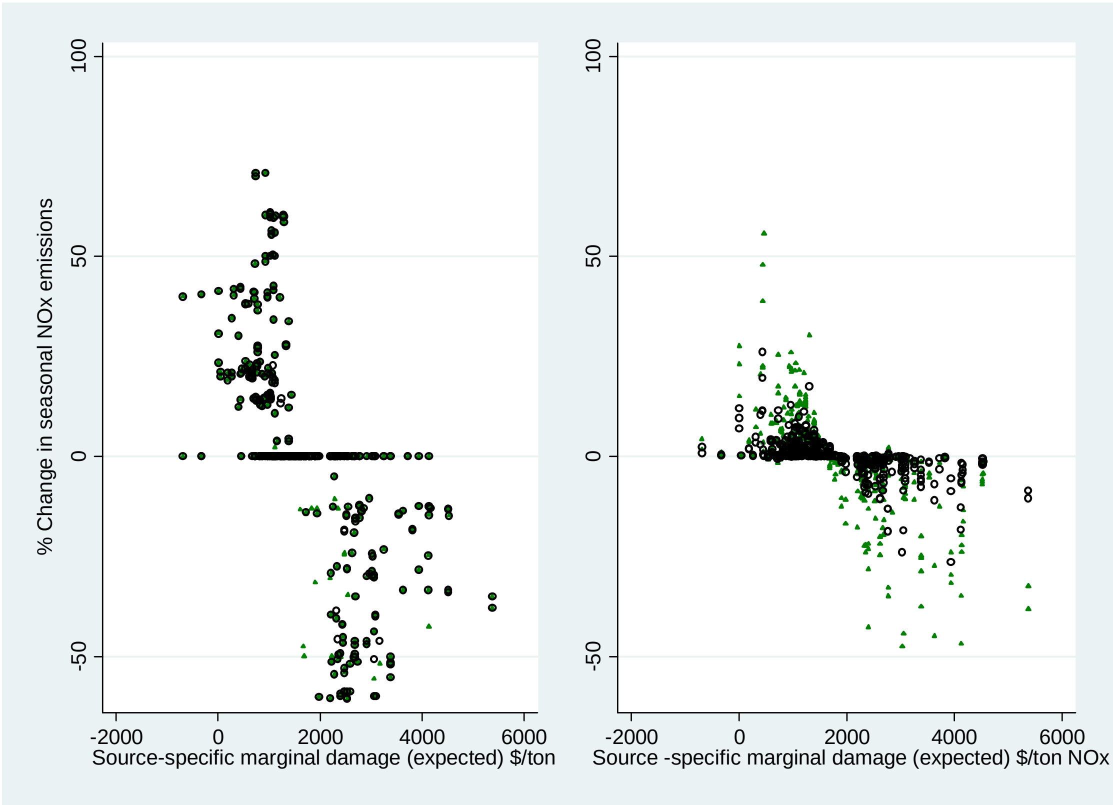

Figure 6: Reallocation of permitted emissions under differentiated NOx permit trading

Notes: The left panel summarizes emissions simulated using the cost minimization algorithm. The right panel summarizes emissions simulated using the econometrically estimated choice model. The vertical axis measures percent changes in simulated ozone season emissions in the observed, undifferentiated regime versus the simulated emissions under the counterfactual, differentiated regimes. The black circles denote emissions changes under the wedge-based design. The small triangles denote emissions changes under the ratio-based design.

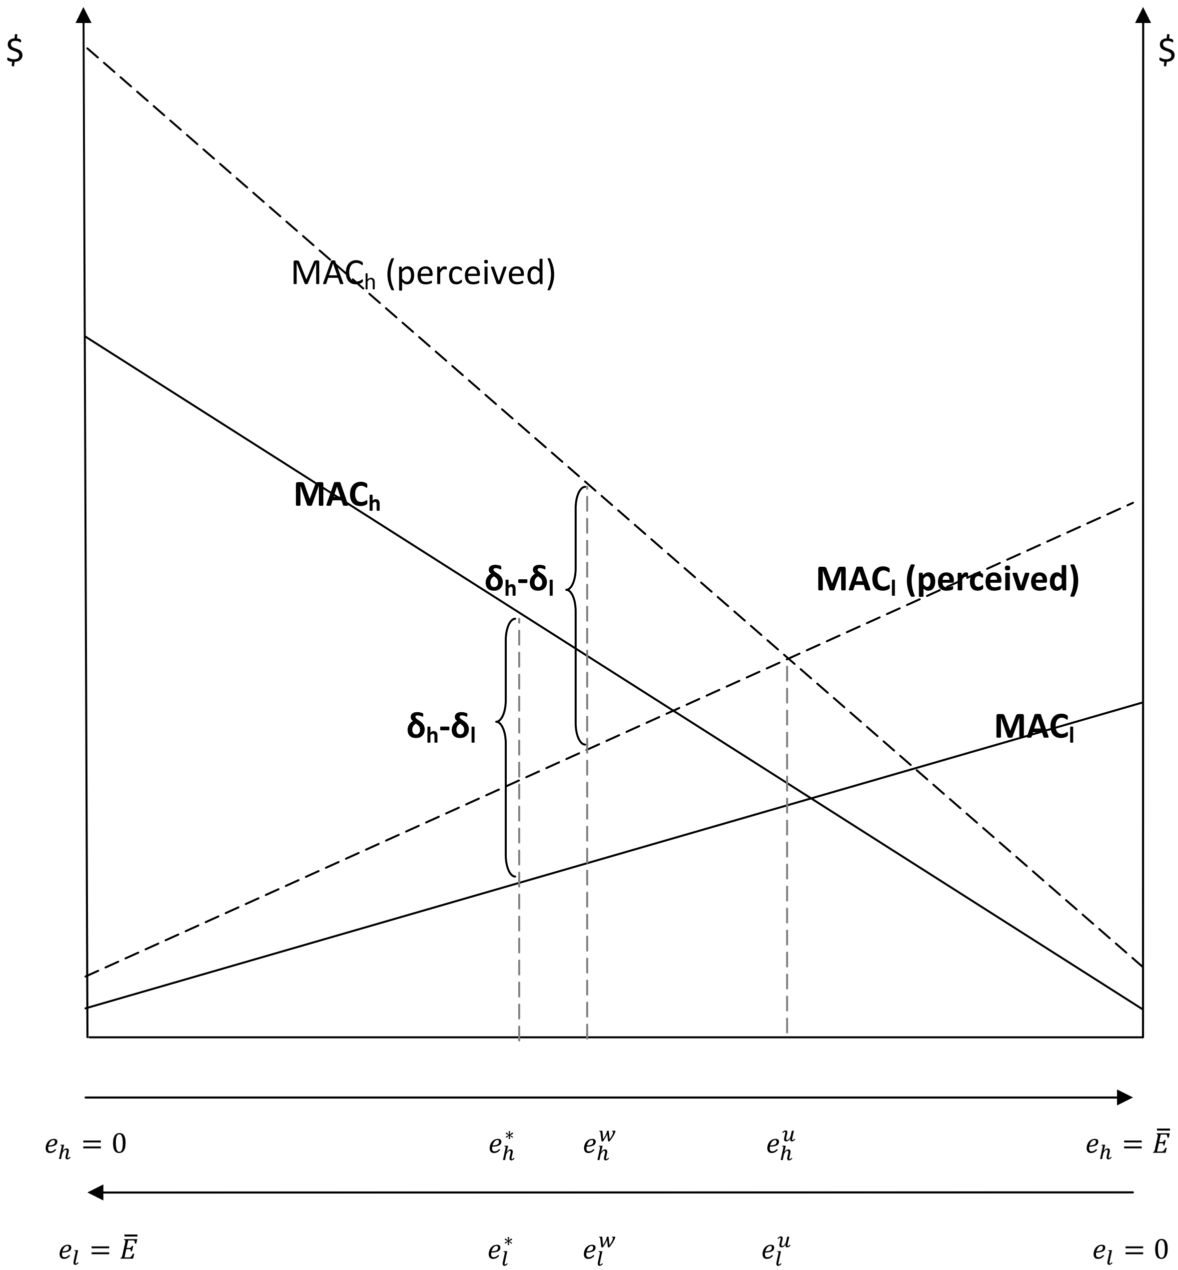

Figure 7: Emissions permit market outcomes under emissions-based and damage-based policies: Sub-optimal emissions constraint

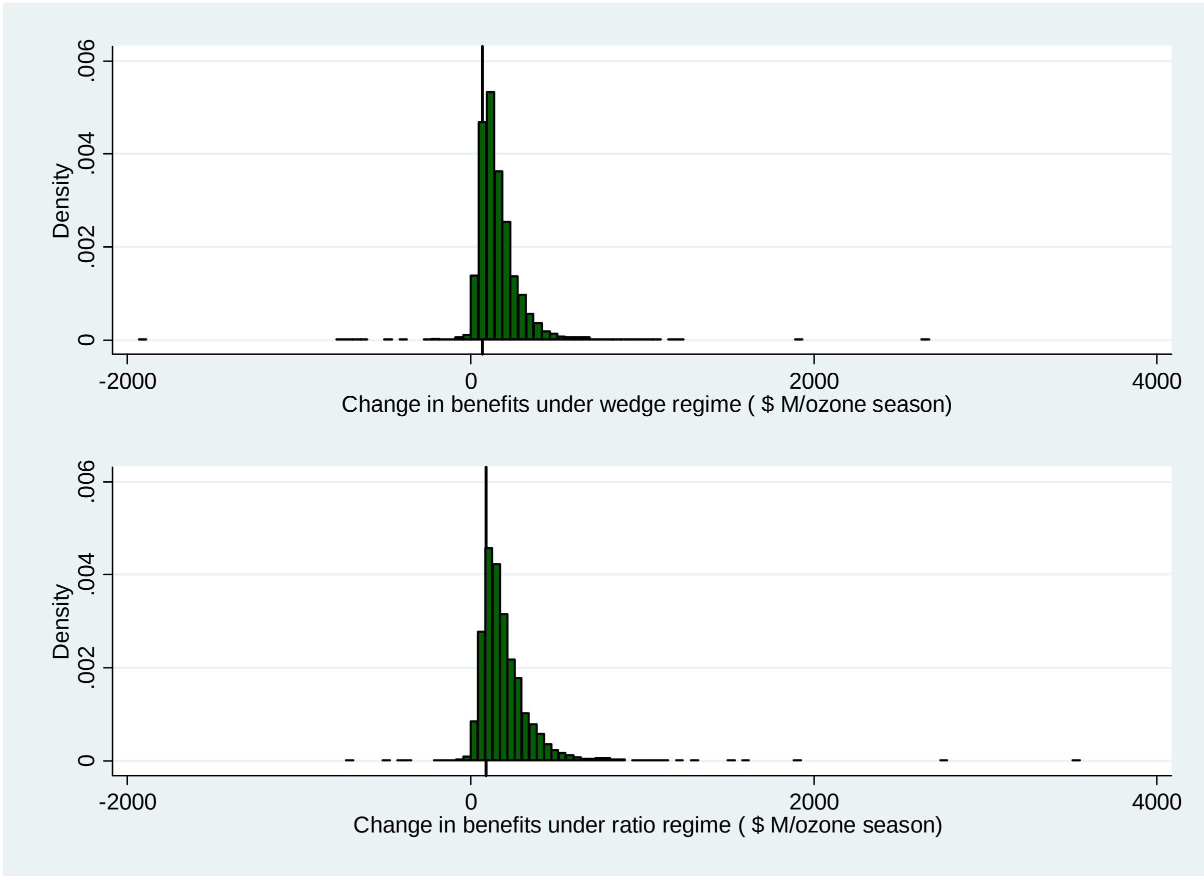

Figure 8A : Benefits of policy differentiation (cost minimization model of compliance decisions).

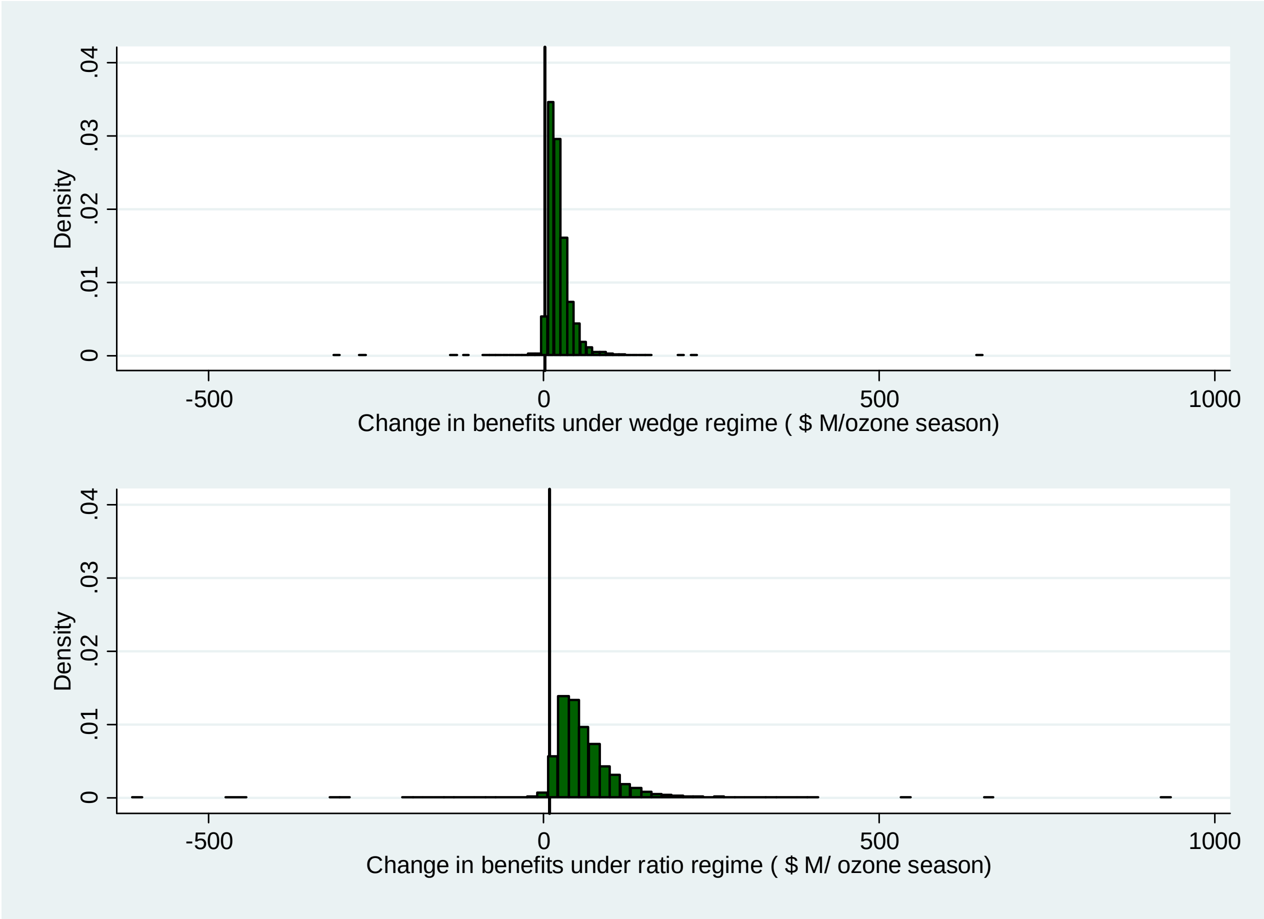

Figure 8B : Benefits of policy differentiation (econometric model of compliance decision)

Table 1 : Unit-level summary statistic

Variable	High damage	Low damage
# Units	241	391
Capacity (MW)	255.61 (234.52)	281.64 (259.84)
Pre-retrofit NOX emissions rate (lbs NOx/mmbtu)	0.55 (0.25)	0.50 (0.20)
Boiler age (years)	35.80 (10.51)	36.59 (11.53)
Summer capacity factor	65.03 (15.22)	66.07 (15.07)
Ozone season production (MWh)	780,000 (683,000)	794,000 (678,000)

Notes: This table summarizes the operating characteristics of 632 coal-fired generating units regulated under the NOx Budget Trading Program. Standard deviations are in parentheses. “High damage” units are those with above average damage parameter point estimates. “Low damage” units are those with below average damage parameter point estimates.

Table 2: Compliance Cost Summary Statistics for Commonly Selected Control Technologies

NOx control technology	Capital cost ($/kW)		Variable cost (cents/kWh)	
	High damage	Low damage	High damage	Low damage
Combustion modification	6.13 (10.64)	8.12 (17.74)	1.00 (0.40)	1.00 (0.39)
Low NOx burners	17.45 (19.94)	21.98 (28.47)	0.68 (0.18)	0.65 (0.14)
SNCR	7.01 (10.09)	8.93 (11.66)	0.98 (0.37)	1.01 (0.41)
SCR	70.94 (127.99)	80.40 (155.01)	0.55 (0.19)	0.52 (0.29)

Notes: This table summarizes the ex ante predicted NOx control costs for 632 coal-fired generating units regulated under the NOx Budget Trading Program. Standard deviations are in parentheses. "High damage" units are those with above average damage parameter point estimates. "Low damage" units are those with below average damage parameter point estimates. Costs were estimated using proprietary software developed by EPRI. See text for details.

Table 3 : Observed, predicted, and correctly predicted compliance choices

Compliance choice	SCR	SNCR	Low NOx burners	Combustion Modifications	No retrofit	Total
Observed choices	187	42	53	58	292	632
Cost minimization model						
Predicted adoption rate	65	79	258	184	46	632
Correctly predicted	48	5	52	11	3	152 (24%)
Econometric model						
Predicted choices	190	16	33	22	371	632
Correctly predicted	172	8	23	18	279	500 (79%)

Notes: This table summarizes predicted and observed compliance choices for the 632 electricity generating units included in the study.

Table 4 : Econometrically estimated coefficients of the compliance choice model

Technology specific constants	High damage units	Low damage units
Post-combustion controls	-2.21 (1.66)	-3.06 (1.34)
Low NOx burners	-2.06 (0.53)	-2.33 (0.43)
Combustion modifications	-1.89 (0.85)	-2.32 (0.69)
Age* capital cost interaction	-0.17 (0.07)	-0.13 (0.06)
Manager-specific coefficients		
Annual compliance cost ($1.000,000)	-1.08 (0.81)	-0.99 (0.59)
Capital cost ($1.000,000)	-0.45 (0.43)	-0.28 (0.33)
# Units	383	269

Notes: Only point estimates are used to parameterize the simulation model. This table reports average coefficient values (averaged across facilities). Standard deviations are in parentheses. For a more detailed discussion of these econometric estimates, see Fowlie (2010).

Table 5 : Simulated outcomes under undifferentiated policy

Model of compliance choice	**Cost minimization (1)**	**Econometric (2)**
Permit price ($/lb NOx)	$0.85	$2.23
Levelized annual abatement costs ($ M)	$488	$692
Annual benefits ($M) (monetized avoided damages)	$1,060 ($884)	$1,075 ($1,030)
Annual net benefits ($M)	$573 ($884)	$383 ($1,030)
% permitted emissions occurring at high damage sources	42.4%	38.9%
% abatement costs incurred by high damage sources	31.4%	38.7%

Notes: This table summarizes the results from simulating investment in NOx abatement and the associated ozone-season emissions under the observed, undifferentiated trading regime. Standard deviations are in parentheses. "High damage" units are those with above average damage parameter point estimates. "Low damage" units are those with below average damage parameter point estimates.

Table 5 : Simulated outcomes under differentiated policy – Cost minimization model

Differentiated policy	Wedges (1)	Ratios (2)
Permit price ($/lb NOx)	$0.90	$0.95
Change in levelized annual abatement costs ($ M)	$67.5	$106
Change in annual benefits ($M) (monetized avoided damages)	$157 ($116)	$194 ($139)
Change in annual net benefits ($M)	$89.8 ($116)	$88.5 ($139)
% permitted emissions occurring at high damage sources	28.1%	27.2%
% abatement costs incurred by high damage sources	61.3%	62.4%

Notes: This table summarizes the results from simulating investment in NOx abatement and the associated ozone-season emissions under two alternative differentiated policy designs. Standard deviations are in parentheses. “High damage” units are those with above average damage parameter point estimates. “Low damage” units are those with below average damage parameter point estimates.

Table 6 : Simulated outcomes under differentiated policy – Econometric model

Differentiated policy	Wedges (1)	Ratios (2)
Permit price ($/lb NOx)	$2.20	$2.33
Change in levelized annual abatement costs ($ M)	$1.6	$8.4
Change in levelized annual abatement costs ($M)		
Change in annual benefits ($M) (monetized avoided damages)	$22.5 ($21.5)	$60.0 ($59.2)
Change in annual net benefits ($M)	$20.9 ($21.5)	$51.6 ($59.2)
% permitted emissions occurring at high damage sources	37.1%	34.2%
% abatement costs incurred by high damage sources	41.5%	45.9%

Notes: This table summarizes the results from simulating investment in NOx abatement and the associated ozone-season emissions under two alternative differentiated policy designs. Standard deviations are in parentheses. "High damage" units are those with above average damage parameter point estimates. "Low damage" units are those with below average damage parameter point estimates.