Designing markets for pollution when damages vary across sources : Evidence from the NO_x Budget Program

Meredith Fowlie* and Nicholas Muller+
*UC Berkeley and NBER
+Department of Economics
Environmental Studies Program
Middlebury College
email: nmuller@middlebury.edu

December 17^{th}, 2010

Abstract

Designing markets for pollution when damages vary across sources :
Evidence from the NO_x Budget Program.

Existing and planned emissions trading programs are almost exclusively "emissions-based", meaning that a permit can be used to offset a ton of pollution, regardless of where in the program region the ton is emitted. Designing programs in this way presumes that the health and environmental damages resulting from the permitted emissions are independent of where in the regulated region the emissions occur. A growing body of scientific evidence indicates that this is not the case for nitrogen oxides (NOx). When marginal damages from incremental emissions reductions vary significantly across sources, there is the potential to significantly improve the efficiency of permit market outcomes by using facility or region-specific marginal damage estimates to determine the terms of permit trading. We estimate the efficiency gains from "damage-based" trading in the context of a major NOx emissions trading program. We find that, under the damage-based trading regime, levelized annual abatement costs increase by an estimated $12 M (i.e. less than 2 percent). However, damages associated with permitted emissions decrease by approximately $62 M annually. The net benefits under the policy that incorporates spatially differentiated trading increase by 17%, or almost $50M annually.

Keywords: Market-based Policy, NO_x Budget Program, Policy Instrument Choice.
JEL Classifications: Q54, Q53, Q58

Acknowledgements: Muller wishes to thank the United States Environmental Protection Agency for support under award: EPA-OPEI-NCEE-08-02.

Economists have long advocated for market-based approaches to pollution regulation (Montgomery, 1972; Baulmol 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 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 allocative efficiency, however, most existing cap-and-trade programs fall short of the theoretical ideal. This is because most policies feature spatially uniform emissions permit trading. That is, all sources in an emissions trading program are permitted to trade allowances with all other sources at an effective one-to-one (i.e. ton-for-ton) exchange rate. By equalizing marginal abatement costs across all sources, a one-for-one trading regime will minimize the total abatement costs incurred to meet the emissions cap. However, spatially uniform permit trading will fall short of allocative efficiency when the impact of emissions - the health and environmental harm - varies across regulated sources.

Allocative efficiency requires that marginal abatement costs be set equal to marginal damages across all sources (Baumol and Oates, 1987; Montgomery, 1972). For well-mixed pollutants such as CO_2, abatement cost minimization and allocative efficiency can be achieved simultaneously since the damage caused by emissions does not vary by source (Hoel, Karp, 2000). When a pollutant is "non-uniformly mixed" (i.e. damages from emissions vary across sources), allocative efficiency cannot no longer be achieved by equating marginal abatement costs across sources.

Spatial variation in 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, and the inefficiencies that arise when a non-uniformly mixed pollutant is regulated using a policy that ignores spatial variation in damages from pollution. In an applied exercise, we estimate how outcomes under a landmark emissions trading program, the NOx Budget Program (NBP), would have differed had the program incorporated spatially differentiated trading.

Market-based policies can, in theory, be designed to account for spatial variation in damages (Montgomery, 1972; Tietenberg, 1980). Baumol and Oates (1987) 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 characterized efficient quantity-based instruments. The key to realizing allocative efficiency is the calibration of permit exchange rates to the ratio of the marginal damage of emissions for each pair of regulated sources (Klaassen, Forsund, Amann 1994; Farrow et al. 2005; Muller and Mendelsohn, 2009).

Heterogeneity in pollution damages adds complexity to both the policy design and to the modeling that informs implementation and ex post evaluation. This complexity can increase administrative costs and reduce political palatability. Thus far, regulators have concluded that the benefits from spatially differentiated trading do not appear to justify the added complexity. Existing emissions trading programs have adopted ton-for-ton trading of non-uniformly mixed pollution.[1]

The focus of this paper is the landmark NBP, a large regional emissions trading program affecting 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 the NBP, 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.

In this paper, with the benefit of hindsight, we revisit the decision to forego spatially differentiated (or so-called "damage-based") NOx trading in favor of the simpler emissions-based alternatives. The paper begins with a conceptually straightforward model that characterizes the welfare implications of moving from an emissions-based permit trading regime to one that incorporates spatially-differentiated, damage-based permit trading. We first consider a stylized, "first best" setting that is free of constraints, distortions, or market failures. We then extend the analysis

[1] Some programs do incorporate some measures to address spatial variation in damages. In principle, the Acid Rain Program prohibits trades that lead to exceedence of NAAQS. Southern California's Reclaim program limits NOx permit trading between coastal and inland areas.

to accommodate some pre-existing distortions and institutional constraints that are particularly relevant to the policy setting we analyze.

Having laid down the theoretical foundations, we turn to the applied policy analysis. In order to estimate the welfare consequences of implementing an emissions-based, versus damage-based, NOx trading program, we simulate firms' compliance decisions under both the observed and counterfactual policy designs. Our analysis proceeds in four stages. First, source-specific marginal damage estimates are generated using a stochastic version of the Air Pollution Emission Experiments and Policy analysis model (APEEP, Muller, Mendelsohn, 2007;2009), AP2. Second, an econometric model of the compliance decisions made by firms subject to the NBP is used to simulate investment in NOx abatement and the associated ozone season NOx reductions under the existing (i.e. emissions-based) policy and a series of counterfactual (i.e. damage-based) designs (Fowlie, 2010). In the third step, these source-specific NO_x emissions are processed through AP2 in order to estimate the aggregate health and environmental impacts under each policy scenario. Finally, the analysis is extended to consider the welfare implications of political constraints, pre-existing distortions in the regulated (i.e. electricity) industry, and uncertainty.

We define the net benefits of a given policy to be equal to the monetized benefits associated with the mandated emissions reduction (i.e. the avoided damages) less the costs incurred to reduce NOx emissions. Estimated net-benefits increase by 17 percent (or approximately \$50 M, annually) when spatial variation in damages is accounted for. This estimate is somewhat conservative in that we assume policy makers would face political constraints that would limit their ability to implement the optimal policy design. If we remove these constraints, cost savings exceed 30 percent.

These findings are germane to the unfolding debate about market-based regulation of non-uniformly mixed pollutants. As policy makers work to design the next generation of emissions trading programs, this debate has reached a fever pitch. When a spatially uniform cap-and-trade system for regulating mercury emissions was proposed in 2004, it attracted a record number of rulemaking comments. Critics were adamant that emissions trading was inappropriate for a toxic, non-uniformly mixed pollutant like mercury.[2] The courts ultimately invalidated the rule. In 2008,

[2] Several studies indicated that CAMR would create local hot spots of mercury pollution, disproportionately impacting some communites.

a federal district court vacated the Clean Air Interstate Rule and the associated regional NO_x trading program, in large part due to policy's failure to adequately accommodate spatial variation in damages.[3]

The paper also contributes to a growing literature that compares spatially differentiated and spatially uniform emissions trading in a variety of policy contexts.[4] Most relevant to this study is work that has been done to analyze spatially differentiated NOx trading in the Eastern United States. This small literature is comprised of ex ante analyses of zonal trading regimes (i.e. market designs that limit or prohibit trading between multi-state trading zones) (ICF Kaiser, 1996; Krupnick et al. 2000; US EPA, 1998). In general, researchers have found the differences in damages across multi-state trading zones to be relatively small. Consequently, it has been estimated that of the welfare gains from spatially differentiated NOx permit trading are negligible.

Our findings contradict this earlier work. To understand why, it is instructive to highlight two distinguishing features of this study. First, we use a detailed integrated assessment model to estimate *source-specific* marginal damages. We find significant variation in these damages; almost half of this variation occurring within (versus between) state. This suggests that multi-state trading zones are a very blunt instrument to capture spatial variation in NOx emissions damages. Second, in all previous work, emissions market outcomes are simulated using a deterministic, cost-minimization algorithms. Krupnick et al. (2000) acknowledge a limitation of this approach, noting that they can make "no claim that optimizations of the kinds described here reflect emissions trading or other particular policies". These optimization models fail to capture salient features of the real world decision processes that drive emissions abatement decisions, and thus market outcomes. In contrast, in our ex post analysis, we are able to use an econometric model to estimate the compliance choices that these plant managers most likely would have made had the NOx emissions market been designed to reflect spatial heterogeneity in marginal damages from pollution.

http://www.ens-newswire.com/ens/feb2005/2005-02-07-10.html

[3]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.

[4]Policies analyzed in prior work include spatially differentiated groundwater permit trading in Nebraska (Kuwayama and Brozovic, 2010), particulate matter trading in Santiago Chile (O'Ryan, 2000) and the monumental U.S. Acid Rain Program ((Kete, 1992; Muller and Mendelsohn, 2009; Chupp, Banzhaf, 2010)

The paper proceeds as follows. Section 1 introduces the theoretical framework and derives some basic theoretical results. Section 2 provides background on the NOx Budget Program. Section 3 introduces the applied analysis and presents the results. Section 4 concludes.

1 Welfare implications of spatially differentiated emissions trading

In this section, we introduce the theoretical foundations underlying our applied policy analysis. A simple theory model is used to characterize the welfare implications of moving from an emissions-based permit trading regime to one that incorporates spatially-differentiated permit trading. We first consider a stylized, "first best" setting. We then extend the analysis to accommodate some pre-existing distortions and institutional constraints that could potentially affect outcomes in the NOx emissions trading program we analyze.

1.1 Theory model

Consider an industry comprised of N firms producing a homogenous good. Industrial production generates harmful pollution. This pollution is non-uniformly mixed, meaning that the extent of the damage caused by emissions depends not only on the level of emissions, but also how the permitted emissions are distributed across sources.[5]

Let E_i^0 denote the baseline emissions at firm i; this is the level of emissions we would observe absent any regulatory constraint. The firm can reduce emissions below E_i by investing in abatement a_i. Firm-level emissions are thus $E_i = E_i^0 - a_i$.We define abatement cost functions in terms of emissions: $C_i(E_i)$. We assume that $C_i'(E_i) \leq 0, C"_i(E_i) \geq 0$. We accommodate heterogeneity in abatement costs by allowing the parameters of the abatement function to vary across firms.

Damages from emissions also vary across facilities. We define firm-specific damage functions $D_i(E_i)$. We make several simplifying assumptions regarding the structure of these damages. First, we assume that aggregate damages D are additively separable: $D = \sum_{i=1}^{N} D_i(E_i)$. Second, we assume firm-level damages are linear in emissions: $D_i(E_i) = k_i + \delta_i E_i$. Finally, we assume the the

[5] 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.

parameters of the firm-specific damage functions are known with certainty. In subsequent sections we investigate the implications of relaxing these assumptions.

The policy designs we consider are, in many respects, standard "cap-and-trade" programs. An emissions cap $\overline{E}$ limits the total quantity of permitted emissions. A corresponding number of tradable permits are allocated. We assume that emissions permits are allocated either by auction or a gratis using some allocation rule that does not depend on production decisions going forward (such as grandfathering). Any free allocation of permits to firm i is represented by the initial allocation A_i.

All of the policy designs we consider are "emissions equivalent", meaning that emissions constraint $\overline{E}$ is held constant across the policy scenarios. Although we are ultimately concerned about limiting the *damages* associated with pollution exposure, in all existing and planned cap-and-trade programs, the cap is defined in terms of *emissions*. This is presumably because imposing a cap on emissions is relatively simple and easy to communicate.

To comply with the regulation, firms must hold permits to offset their uncontrolled emissions. We assume that facilities comply with the regulation either by holding emissions permits, investing in emissions abatement, or some combination of these two strategies. We rule out reduction in output as a compliance strategy and assume that firm-level production and aggregate output are exogenously determined and independent of the environmental compliance choice. This assumption is appropriate for the policy context we consider.[6]

We use the following total social cost measure (TSC) to evaluate equilibrium outcomes under the alternative policy scenarios:

[6] Our analysis focuses exclusively on coal plants who account for the vast majority of NOx emissions regulated under the NBP. These coal plants are are typically inframarginal due to their relatively low fuel operating costs. The introduction of the cap-and-trade program therefore reduced profit margins but not capacity factors at these units (Fowlie, 2010). Price-setting units (typically natural gas or oil-fueled plants) represent a very small fraction of the NOx emissions regulated under the NBP and tend to have much lower uncontrolled NOx emissions rates. 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). If compliance costs incurred at marginal units were passed directly through to consumers, retail electricity prices would be unlikely to increase by more than one percent on average. Given the small magnitude of this price change, and given the inelasticity of electricity demand, we make the simplifying assumption that total production levels are unaffected by the introduction of the NOx emissions trading program.

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

Any policy-induced change in social welfare will be captured by changes in (1) (because production and consumption levels are exogenous).

1.2 First-best Outcome.

To keep the analytics simple and intuitive, we consider a case with only two price taking firms. Producers are denoted h and l to indicate high and low damage areas respectively. For each firm we define a marginal damage parameter $\delta_i = D_i'(E_i)$ which captures the damages (measured in dollars) caused by an incremental change in E_i. We assume, $\delta_l < \delta_h$.

We first consider a setting in which marginal damages are known with certainty and the cap is optimally set. Aside from the emissions externality, markets are efficient and free of distortions. We refer to this as the "first-best" policy setting.

In this first best setting, the policy maker's objective is to coordinate investment in pollution abatement, and thus emissions, so as to minimize (1):

$$\min_{E_h, E_l} TSC = \delta_h E_h + \delta_l E_l + C_h(E_h) + C_l(E_l). \tag{2}$$

First order conditions with respect to (E_i) for cost minimization imply:

$$\begin{aligned} C_i'(E_i^*) q_i &= \delta_i \quad i = h, l. & (3) \\ &\rightarrow \quad C_h'(E_h^*) \frac{1}{\delta_h} = C_l'(E_l^*) \frac{1}{\delta_l}. & (4) \end{aligned}$$

The efficient level of aggregate emissions is denoted $E^* = E_h^* + E_l^*$; the * superscript denotes efficient levels of emissions. To minimize the costs of meeting the emissions constraint, marginal costs are set to equal marginal damages at all sources.

Figure 1 illustrates these first order conditions graphically. The width of this figure, measured in units of emissions, is equal to the total quantity of permitted emissions $\overline{E}$. We assume the cap has been set optimally: $\overline{E} = E^*$. 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 $C_h'(E_h)$.

The broken lines in Figure 1 represent the marginal abatement cost schedules scaled by the inverse of the corresponding damage parameter: $C_i'(e_i)\frac{1}{\delta_i}$, $i = l, h$. By [4], the allocation of emissions across these two sources occurs where these broken lines intersect. The shaded region represents damages associated with this allocation of emissions (i.e. $\delta_l E_l^* + \delta_h E_h^*$). To see why this outcome is optimal, note that allocating more of the permitted emissions to the low damage firm would reduce welfare because the net increase in abatement costs would exceed the net decrease in damages. Allocating more of the permitted emissions to the high damage firm would reduce welfare because the associated increase in damages would exceed the incremental reduction in abatement costs.

1.3 Emissions-based trading

Having characterized the optimal allocation of pollution abatement activity (and thus emissions), we now compare equilibrium outcomes under two market-based policy designs against this first best benchmark. We first analyze outcomes under an "emissions-based" market design.

In an emissions-based trading regime, we assume that each firm chooses emissions (E_i), emissions permit purchases (E_{bi}), and permit sales (E_{si}) to minimize the total cost of complying with the regulation:

$$\min_{E_i, E_{si}, E_{bij}} \quad C_i(E_i) + \tau^e (E_{bi} - E_{si} - A_i) \tag{5}$$
$$s.t. \quad E_i \leq A_i - E_{si} + E_{bij}$$
$$E_i, E_{si}, E_{bij} \geq 0,$$

where τ^e represent the equilibrium permit price.

The Lagrangian for the firm's cost minimization problem is:

$$L_i = C_i(E_i) + \tau (E_{bi} - E_{si} - A_i) + \lambda_i (E_i - A_i + E_{si} - E_{bi}). \tag{6}$$

The well known first order conditions with respect to (E_i) imply:

$$C_i'(E_i^e) \geq \lambda_i, \quad i = h, l, \tag{7}$$

The superscript $(^e)$ denotes the the emissions-based trading equilibrium. If we assume an interior solution, the first order condition for optimal purchasing (E_{bi}) and selling (E_{si}) of allowances is:

$$\tau = \lambda_i \tag{8}$$

Taken together, these first order conditions imply that marginal abatement costs are set equal across all sources in an emissions based trading system:

$$C_h'(E_h^e) = C_l'(E_l^e) = \tau. \tag{9}$$

In Figure 1, equilibrium emissions under the emissions-equivalent emissions-based trading regime are given by $\{E_l^e, E_h^e\}$. This equilibrium occurs at the intersection of $C_h'(E_h)$ and $C_l'(E_l)$. This outcome is not optimal; welfare could be improved by shifting more of the permitted emissions away from the high damage source to the low damage source.

1.4 Exposure-based trading

We now consider a policy design that is identical in every respect to the emissions-based trading regime, except that compliance requirements are now defined in terms of relative damages. Let $\overline{\delta}$ represent the average of the damage coefficients across regulated firms. 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 damage coefficient by the mean damage parameter $\overline{\delta} : r_i = \frac{\delta_i}{\overline{\delta}}$. For example, if emissions from firm h cause twice as much damage as emissions from firm l, firm h is required to hold twice as many permits as firm l to offset a unit of pollution.

Firms are assumed to choose their compliance strategy so as to minimize total compliance costs:

$$\min_{E_i, E_{si}, E_{bij}} \quad C_i(E_i) + \tau(E_{bi} - E_{si} - A_i). \tag{10}$$

$$s.t.\ r_i E_i = A_i - E_{si} + E_{bi}.$$

$$E_i, E_{si}, E_{bij} \geq 0$$

The supporting Lagrangian is:

$$L_i = C_i(E_i) + \tau(E_{bi} - E_{si} - A_i) + \gamma_i(r_i E_i - A_i + E_{si} - E_{bi}).$$

The first order conditions with respect to (E_i) for cost minimization imply:

$$C_i'(E_i^d) \geq r_i \gamma_i \quad i = h, l. \tag{11}$$

where the d superscript denotes an equilibrium outcome under damage-based trading.

The first order conditions with respect to the purchase and sale of permits are, equivalently:

$$\tau = \gamma_i \tag{12}$$

Rearranging these first order conditions yields:

$$C_i'(E_i^d) = \tau r_i. \tag{13}$$

In equilibrium, the ratio of marginal abatement costs is equated to the ratio of marginal damages of firms' emissions. Appendix 1 shows that the equilibrium permit price in this damage-based trading regime will equal $\overline{\delta}$ if the cap is optimally set at $\overline{E} = E^*$. The equilibrium outcome under damage-based trading will be the socially efficient outcome:

$$C_i'(E_i^d) = \delta_i \quad i = h, l. \tag{14}$$

.

1.5 Welfare comparisons of emissions-based and exposure-based trading.

The primary objective of this paper is to investigate the welfare consequences of moving from an emissions-based NO_x trading program design to the damage-based alternative. To more concretely motivate this estimation exercise, we derive analytical expressions for the costs and benefits associated with this policy design change. We use a quadratic approximation of the firm-level abatement cost function defined in terms of emissions $C_i(E_i)$:

$$C_i(E_i) = \alpha_{0i} - \alpha_{1i}E_i + \alpha_{21}E_i^2, \tag{15}$$

and a linear approximation of the firm-specific damage function for emissions $D_i(E_i)$:

$$D_i(E_i) = \delta_{0i} + \delta_{1i}E_i. \tag{16}$$

Appendix 1 derives an expression for the optimal cap as a function of the damage parameters δ and the abatement cost function parameters α:

$$\overline{E}^*(\boldsymbol{\alpha}, \boldsymbol{\delta}) = \frac{(\alpha_{1h}\alpha_{2l} + \alpha_{1l}\alpha_{2h} - \delta_h\alpha_{2l} - \delta_l\alpha_{2h})}{2\alpha_{2l}\alpha_{2h}}. \tag{17}$$

This is the level of emissions that minimizes [1]. All else equal, the more damaging the emissions, the more stringent the cap. In contrast, the more costly it is to reduce emissions, the more lax the optimal cap.

Once the optimal emissions cap has been identified and introduced into the model, it is straightforward to derive expressions for equilibrium emissions outcomes under the emissions-based and damage-based policy designs, respectively. Subtracting emissions under an emissions-based regulatory regime E_i^e from emissions under an damage-based trading regime E_i^d yields:

$$E_i^d(\alpha,\delta) - E_i^e(\alpha,\delta) = \frac{\delta_j - \delta_i}{2\left(\alpha_{2i} + \alpha_{2j}\right)}, i \neq j, i, j = l, h. \tag{18}$$

When relative damages are accounted for, a larger share of the permitted emissions occurs at the low damage source.

Expositionally, in the applied analysis, it will be useful to decompose the welfare implications of moving from an emissions-based regime to exposure based trading into two parts:

1. *Changes in aggregate investment costs*:

   Moving from an emissions-based to damage-based emissions trading design will increase overall emissions abatement costs. An analytical expression for this abatement cost increase, conditional on the assumptions of the model, is derived in Appendix 2:

$$C^d(\alpha,\delta) - C^e(\alpha,\delta) = \frac{(\delta_l - \delta_h)^2}{4(\alpha_{2h} + \alpha_{2l})} \geq 0 \tag{19}$$

   The more heterogeneous the damages, the more significant the reallocation of permitted emissions under damage-based trading, the greater the increase in industry-wide abatement costs.

   In Figure 1, moving from the emissions-based to the damage-based emissions trading regime results in an increase in abatement costs incurred at the high damage firm (represented by the area $ACED$). Abatement costs at the low damage firm are reduced by an amount defined by the triangular area CDE. Taken together, abatement costs increase by an amount equal to the triangular area ADC.

2. *Changes in damages caused by permitted emissions*: Although total emissions are held equal across the market designs we consider, total damages will differ.The analytical expression for this change in damages (also derived in Appendix 2):

$$D^d(\alpha,\delta) - D^e(\alpha,\delta) = -\frac{(\delta_l - \delta_h)^2}{2\,(\alpha_{2h} + \alpha_{2l})} \leq 0, \tag{20}$$

   In Figure 1, moving from an emissions-based to an exposure-based design decreases damages by an amount equal to the rectangular area $ABCE$.

Net welfare impacts of spatially differentiating emissions permit trading are obtained by subtracting the costs from the benefits. In the context of our analytical model, the net welfare impacts are positive and given by:

$$\begin{aligned} \Delta TSC(\alpha,\delta) &= \frac{(\delta_l - \delta_h)^2}{2\,(\alpha_{2h} + \alpha_{2l})} - \frac{(\delta_l - \delta_h)^2}{4(\alpha_{2h} + \alpha_{2l})} && (21) \\ &= \frac{(\delta_l - \delta_h)^2}{4(\alpha_{2h} + \alpha_{2l})} \geq 0 && (22) \end{aligned}$$

This simple analytical expression serves to highlight the factors that determine the relative cost effectiveness of spatially differentiated trading (as compared to more standard, emissions-based designs). First, welfare gains from spatially differentiated trading are increasing with the difference in damage parameters $\delta_h - \delta_l$.The greater the variation in damages across sources, the larger the benefits associated with spatially differentiated trading, all else equal. Trivially, if damages do not vary across sources, emissions-based trading and damage-based trading will be one in the same.

The abatement cost structure is also important. Welfare gains associated with damage-based trading are decreasing with the slope of the marginal abatement cost curves. Intuitively, the more steeply sloped the marginal abatement cost curves, the more costly it is to incrementally reduce damages by reallocating emissions from the relatively high damage source to the relatively low damage source.

1.6 Damage-based emissions trading in a second best setting

The preceding section demonstrates that, in a first best setting, damage-based emissions trading welfare dominates emissions-based emissions trading when damages vary across sources. In practice, the welfare implications of moving from an emissions-based to a damage-based program design may be more nuanced and ambiguous. Pre-existing market failures, jurisdictional limitations, and other distortions can complicate the application of this theory in practice.

In this working paper, we consider three factors that impact the welfare implications of spatially differentiated permit trading in the institutional setting we investigate. We first consider political constraints that limit the regulators' ability to optimally define the damage based trading ratios r, or the cap $\overline{E}$. We then consider how pre-existing regulatory distortions in the polluting industry can affect the relative efficiency properties of emissions and exposure-based permit trading. In future work, we will extend this analysis to allow for non-linear damage functions. We will also evaluate how the uncertainty inherent in marginal damage estimates should best be incorporated into permit market design.

1.6.1 *Incorrect damage parameters*

In practice, the ratios that define the terms of compliance may deviate from the true damage parameters.For example, political constraints may limit the ability of regulators to define trading ratios that reflect the true range of marginal damages. Let θ_i represent the difference between the true damage parameter at firm i, δ_i and the parameter that is used to define the terms of compliance. In the presence of these distortions, the welfare impacts of transitioning from an emissions-based regime to a damage-based regime are given by:[7]

[7] Appendix 1 provides a detailed derivation.

$$\Delta TSC(\alpha, \delta, \theta) = \frac{(\delta_l - \delta_h)^2}{4(\alpha_{2h} + \alpha_{2l})} - \frac{(\alpha_{2h} + \alpha_{2l})(\theta_l - \theta_h)^2 \delta_h^2 \delta_l^2}{4(\theta_h \delta_h \alpha_{2l} + \theta_l \delta_l \alpha_{2h})^2} \tag{23}$$

If $\theta_l = \theta_h$, the source-specific distortions cancel out and compliance ratios are unaffected. In this special case, incorrectly defined damage parameters have no impact on market outcomes. However, if $\theta_l \neq \theta_h$,the compliance ratios will not accurately reflect the true degree of heterogeneity in damages; the equilibrium outcome under damage-based trading will no longer be efficient. In the analysis that follows, we will consider one particular political constraint that would likely result in this kind of distortion.

1.6.2 *Exogenously determined emissions cap*

For a variety of reasons, it may be difficult in practice to set the emissions cap at the optimal level E^*. Political, jurisdictional, or other implementation constraints may result in an emissions cap that is too stringent- or not stringent enough. When the cap is not set optimally, the equilibrium permit price under damage-based trading will no longer equal the average damage parameter $\overline{\delta}$.

Let Δ measure the difference between the emissions cap and the optimal level of industry emissions: $\overline{E} = E^* + \Delta$. The equilibrium permit price under the damage-based trading regime is:[8]

$$\tau(\alpha, \delta, \Delta) = \frac{\delta_h + \delta_l}{2} \frac{(\delta_l \alpha_{2h} + \delta_h \alpha_{2l} - 2\Delta \alpha_{2h} \alpha_{2l})}{(\delta_h \alpha_{2l} + \delta_l \alpha_{2h})} \tag{24}$$

If $\Delta \neq 0$, damage-based emissions trading will fail to coordinate investment in emissions abatement efficiently. The permit price will be too high if the cap is set too stringently (i.e. $\Delta < 0$) and too low if the cap is too lax (i.e. $\Delta > 0$).

Figure 2 helps to illustrate the qualitative implications of an emissions cap that is too stringent. This figure is identical to Figure 1, except that we now assume smaller damage parameters for each source, such that $\overline{E} < E^*(\alpha, \delta)$. By (11), marginal abatement costs are set equal to the product of the permit price and the damage-based compliance ratio:

[8] See Appendix 2 for the derivation.

$$C_i'(E_i^d) = \tau\frac{\delta_i}{\overline{\delta}} = \delta_i\frac{(\delta_l\alpha_{2h} + \delta_h\alpha_{2l} - 2\Delta\alpha_{2h}\alpha_{2l})}{(\delta_h\alpha_{2l} + \delta_l\alpha_{2h})} > \delta_i \tag{25}$$

This implies that $C_i'(E_i^d) > \delta_i$ for $i = l, h$.

In Figure 2, industry-wide abatement costs under damage-based permit trading exceed costs under emissions-based trading by an amount equal to the area of the shaded triangle. The shaded rectangle represents the benefits (in the form of avoided emissions damages) associated with spatially differentiated trading using the new damage parameters. In this case, the benefits from spatially differentiated permit trading no longer exceed the costs such that the emissions-based trading regime now welfare dominates the damage-based regime. In contrast, if $\Delta > 0$ and the cap is not sufficiently stringent, the welfare gains from spatially-differentiated trading increase relative to the benchmark first-best case.

1.6.3 *Pre-existing abatement cost distortions in the product market*

Pre-existing distortions in the product market can also affect the efficiency properties of permit market outcomes. For example, the electricity generating units in the NO_x emissions trading program face different economic regulatory incentives in their respective electricity markets. Some forms of electricity market regulation have the potential to distort environmental compliance choices (Fowlie, 2010).

Figure 3 helps to illustrate a case in point. Here we assume that the high damage firm is subject to economic regulation that subsidizes investment in emissions abatement. This may occur in regulated electricity markets when producers are guaranteed to earn a positive rate of return on capital investments in pollution abatement equipment. The marginal abatement cost schedule denoted MAC_h reflects the true abatement costs, whereas the abatement cost curve denoted MAC'_h reflects the abatement costs as perceived by the firm. This regulatory distortion moves the equilibrium outcome under emissions-based trading closer to the optimum. In the damage-based trading regime, marginal abatement costs exceed damages at the high damage firm, whereas the reverse is true at the low damage firm. As a result, these pre-existing regulatory incentives distort environmental compliance incentives in a way that reduces the net benefits from spatially differentiated permit trading are reduced. Of course, pre-existing distortions can work in the other direction if they effectively reduce costs as perceived by relatively low damage sources.

1.6.4 *Damage function mispecification*

In this working paper, we will assume that the effect of an incremental change in emissions at source i on health and environmental outcomes can be adequately captured by the scalar δ_i. More precisely, we assume that aggregate damages D are linear and additively separable in terms of source-specific emissions E_i. These damage parameters, and corresponding trading ratios, are estimated once; they are not permitted to adjust as investments are made in NOx control equipment and emissions rates of other sources are reduced.

In future work, we will evaluate the plausibility of the assumptions we make about the damage function. In particular, we will investigate the extent to which periodically updating trading ratios could improve the efficiency of damage-based permit trading.

1.6.5 *Risk and uncertainty in damage measures*

Another maintained assumption in this working paper is that the damage parameters δ are known with certainty. In fact, estimation of these parameters is predicated upon a complex series of assumptions and approximations. For expositional clarity, we define three main different sources of uncertainty that can complicate the estimation of these marginal damage parameters.

The first is air quality modeling uncertainty. In order to estimate source-specific damage parameters, we must first model how changes in emissions levels at one point in the airshed affects pollution concentrations at all other points in the airshed. The pollution formation, transport, and deposition processes that determine how a change in emissions at one location affects pollution concentrations and exposure at other points in space and time are complex. This complexity begets uncertainty about how these processes should best be represented in a modeling framework.

The second source of variability stems from the data used to estimate these models. Pollutant formation and transport can be a highly stochastic process that depends fundamentally on weather patterns, meteorological conditions, pre-existing precursor concentrations, and other hard-to-predict phenomena.

The third source of uncertainty is introduced when pollution concentrations are converted into monetized damages. This requires a series of additional assumptions about how changes in pollution concentrations at a given location affect health and environmental outcomes, and about how much these outcomes are worth.

In this working paper, we will use point estimates of the source-specific marginal damage coefficients to define the terms of trade. In future work. we will explore the implications of the three aforementioned sources of uncertainty surrounding these estimates.

2 Empirical application: the NOx Budget Program

The NOx Budget Program (NBP) is an emissions trading program that limits emissions of NOx from large stationary sources in nineteen eastern states. The NBP was primarily designed to help Northeastern and Mid-Atlantic states attain Federal ozone standards. Prior to the introduction of this program, large point sources in the region were subject to a prescriptive standard that required the installation of low NOx burners. These standards proved insufficient. 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 and associated damages in this region was difficult to estimate precisely, 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, together with a simulation exercise which indicated that imposing spatial constraints on trading would not significantly affect the location of emissions (US EPA, 1998a), led regulators to design a single jurisdiction trading program in which all emissions are traded on a one-for-one basis.

In our analysis, we use data collected from 632 coal-fired generating units that are regulated under the NOx Budget Program. Although gas- and oil-fired generators and other industrial point sources are also included in the NBP, these coal-fired units represent over 90 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). 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.

3 Ex post analysis of spatially differentiated NOx trading

The primary objective of our applied analysis is to examine the implications of spatially differentiated NOx trading in a landmark emissions trading program. Section 1 laid the conceptual foundations for this exercise. The theoretical model also helps to highlight the essential inputs and outputs of our analysis. Inputs include the source-specific demand parameters δ, the unit specific abatement cost schedules $C_i(E_i)$, and the decision rule that dictates how firms in an emissions trading program make their abatement investment decisions. Key outputs include the relative impact of spatially differentiated NOx trading on aggregate abatement costs, analogous to $C^x(\alpha, \delta) - C^e(\alpha, \delta)$ in (19); and the relative impact of spatially differentiated NOx trading on aggregate damages, analogous to $D^x(\alpha, \delta) - D^e(\alpha, \delta)$ in (20). In what follows, we describe these inputs and outputs in detail.

3.1 Source-specific damage parameters

NOx emissions affect health and environmental outcomes through two main pathways: ozone formation and particulate matter formation.[9] 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. In sum, the damage caused by a given quantity of NOx emissions will depend on the spatial distribution of the emissions.[10]

The source-specific δ_i parameters capture how an incremental change in NOx emissions at source i affects damages across the airshed through changing concentrations of (and thus exposure to) ozone and particulate matter. In order to connect emissions of NO_x to concentrations of both $PM_{2.5}$ and O_3 and to the resultant physical impacts and damages, this paper employs AP2, a stochastic version of the APEEP model which has been used in prior research (Muller and Mendelsohn, 2007;2009).

Figure 9 provides a diagram of the AP2 model. AP2 is a standard integrated assessment model in its overall structure. The model is comprised of six modules; emissions, air quality modeling, concentrations, exposures, physical effects, and monetary damages. The emissions data used in

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

[10] The NOx Budget Program does not explicitly account for spatial variation in marginal damages from emissions; a permit can be used to offset a unit of NOx emissions, regardless of where in the program region the unit is emitted.

AP2 is provided by the USEPA'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 USEPA) to occur.

The approach to air quality modeling used in AP2 relies on the Gaussian Plume model (Turner, 1994). This tack uses a reduced form statistical model to capture the processes that connect emissions to concentrations. The predictions from the AP2 model have been tested against the predictions made by a more advanced air quality model (see Muller and Mendelsohn, 2007). The agreement between the county-level surfaces produced by the two models is quite strong. AP2 then connects ambient concentrations to physical impacts using peer-reviewed dose-response functions. 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., 1993). The final modeling step in connecting emissions to damages is expressing the physical effects predicted by the dose-response functions in monetary terms. To do this, 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 USEPA, results from a meta-analysis of nearly 30 studies that compute VSLs. Further, each case of chronic bronchitis is valued at approximately $300 thousand which is also the value used by USEPA.

The marginal ($/ton) damage for NO_x for the 632 coal-fired electricity generating units are estimated using the marginal damage algorithm 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 the 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.[11]

Figure 4 summarizes the unit specific point estimates of marginal damages from NOx emissions. The average parameter value is \$2180/ton of NOx emitted during ozone season. These parameters vary significantly across the sources in the program. Notably, a significant amount of this 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) would be a fairly blunt policy instrument.

For five of the 632 units in our data, we find that the estimated damage parameters are negative. This suggests that a decrease in NOx emissions at these sources leads to increased overall damages. This seemingly counterintuitive 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.

With these unit-specific damage parameter estimates in hand, it is straightforward to construct the compliance ratios. For each source, we divide the source-specific damage measure δ_i with the mean value $\overline{\delta}$. Figure 5 plots the compliance ratios we will use in our primary damage-based policy counterfactual as a function of the estimated damage parameters. Relatively "high damage" units

[11] In this working paper, all of our analysis is based on the point estimates of these marginal damage parameters. Future work will investigate the implications of the uncertainty surrounding these estimates.

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. Note that there are no negative trading ratios. We assume that incentivize pollution at facilities with negative damage parameter estimates would be politically unpopular. We thus assign the five units with negative marginal damage estimates a compliance ratio of zero, equivalent to dropping them from the program under the damage-based regime.

3.2 NOx abatement costs

The NBP mandated a dramatic reduction in average NOx emissions rates.[12] 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. To comply, firms can do one or more of the following: purchase permits to offset emissions exceeding their allocation, install one or more NOx control technologies, or reduce production at dirtier plants during ozone season.

Two factors that are likely to significantly influence a manager's choice of environmental compliance strategy are the up-front capital costs and anticipated variable compliance costs (*i.e.* compliance costs incurred per unit of electricity produced). 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. The specific NOx control options available to a given unit also vary across units of different vintages and boiler types. Compliance options that incorporate Selective Catalytic Reduction (SCR) technology 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) can reduce emissions by fifteen to fifty percent, depending on a boiler's technical specifications and operating characteristics.

[12] 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).

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[13] developed software to generate cost estimates for all major NOx control options 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 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. On average, capital costs are somewhat higher among units located in low damage areas.

3.3 Firm-level compliance decisions

In order to simulate firms' response to different emissions trading program designs, we use an empirical model developed by Fowlie (2010). The basic structure of this discrete choice model is as follows. The manager of unit i $(i = 1..632)$ faces a choice among J_i compliance strategy alternatives (indexed by $j,\ j = 1...J_i$). Plant managers are assumed to choose the compliance strategy that minimizes the unobserved latent value C_{ij}. 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 :

$$\begin{aligned} C_{ij} &= \alpha_j + \beta_i^v v_{ij} + \beta_i^K K_{ij} + \beta^{KA} K_{ij} \cdot Age_{ij} + \varepsilon_{ij}, \qquad (26) \\ where \quad v_{ij} &= (V_{ij} + \tau r_i m_{ij}) Q_i \end{aligned}$$

[13] 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.

An interaction term between capital costs and demeaned plant age is included in the model because older plants can be expected to weigh capital costs more heavily as they have less time to recover these costs. The variable cost (per kWh) of operating the control technology is V_{ij}. The variable cost associated with offsetting emissions with permits is equal to the product of the permit price τ, the compliance ratio r_i and the post-retrofit emissions rate m_{ij}. In the observed emissions-based policy regime, $r_i = 1$ for all units.

Expected average 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.[14]

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.[15]

The model is estimated separately for units serving restructured wholesale electricity markets versus publicly owned units and units subject to cost-of-service regulation (see Fowlie (2010)). The estimated parameters of the random coefficient distributions are then combined with information about observed choices in order to make inferences about where in the population distribution a particular plant manager most likely lies. A more detailed description of the model specification, estimation results, and the derivation of manager-specific coefficient distributions can be found in Fowlie (2010).

We use the RCL coefficient estimates, together with the implied manager-specific distributions of the β^K and β^v parameters, to simulate outcomes under the observed emissions-based policy design and the counterfactual, damage-based designs. The simulations proceed as follows:

[14] 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.

[15] 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.

1. A policy scenario (i.e. the permit market design to be analyzed) is defined in terms of the compliance ratios vector r and the emissions cap E.

2. For each manager, for each random parameter, Z random draws from the appropriate manager-specific density are taken. Let b_{mz} represent the z^{th} draw from the distribution of coefficients associated with manager m.

3. Beginning with repetition $z = 1$, simulation of the market clearing permit price τ_z and emissions begins by setting the permit price equal to 0.

4. For each unit, choice probabilities are approximated for all available compliance choices conditional on the prevailing permit price τ, the coefficient vector $\mathbf{b}_{mz}$, the trading ratio vector r and the unit-specific choice set characteristics. Managers are assumed to choose the compliance strategy with the highest estimated probability.

5. Ozone season emissions (measured in lbs of NOx) and engineering estimates of compliance costs associated with the predicted choices are calculated and summed across units.

6. If the total quantity of emissions equals the cap, τ is the equilibrium price and the simulation stops. If the total quantity of emissions exceeds (is less than) the cap, τ is increased (decreased) by \$0.01.

7. Steps 4-6 are repeated until an equilibrium is reached.[16]

8. Steps 3-7 are repeated Z times for each policy scenario.

Under the baseline policy scenario, $r = 1$ for all units and the emissions cap E is set equal to the emissions implied by the observed compliance choices. To simulate outcomes under damage based trading, the estimated, source specific compliance ratios are used to define the terms of compliance. This effectively increases (decreases) the costs of offsetting uncontrolled emissions using permits for a relatively high (low) damage firm.

[16] 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.

This approach assumes that the structure of these firm-level compliance decisions, and in particular, the relative weighting of capital costs and variable operating costs, would not change if the terms of compliance were damage-based versus emissions-based. We believe this to be a very reasonable assumption. Once differences in variable operating costs across the two regimes are accounted for, we see no reason why a plant manager would take a fundamentally different approach to compliance.

Figures 6 and 7 offer a graphical summary of our simulation results. Each point in these scatterplots represents a different electricity generating unit. The horizontal axis measures the source-specific damage parameters we introduced in section 3.1. These are adjusted slightly; the five negative marginal damage estimates are set to zero. The horizontal axis in Figure 6 measures the percentage change in unit-level emissions, moving from spatially uniform trading to the spatially differentiated trading regime. Intuitively, units with relatively low (high) damage parameters will increase (reduce) emissions under a damage-based regime. Note that the emissions levels at many units are unaffected by the policy design change. Given the discrete nature of the NOx abatement choice, a unit's cost minimizing choice of emissions abatement (and thus emissions) need not be affected by a change in the cost of holding permits to offset emissions.

Figure 8 conducts a similar exercise using data on simulated levelized annual abatement costs. The horizontal axis measures the difference in simulated levelized annual abatement costs across the emissions equivalent, emissions-based and damage-based trading regimes. Intuitively, investment in abatement is lower (higher) in the spatially differentiated trading regime at the units with relatively low (high) damage parameters.

Tables 3 and 4 summarize these simulation results in more detail. Table 3 summarizes the unit-level simulation results. For each unit, simulated emissions and abatement costs are summed within policy scenario across 50 repetitions. Table 3 summarizes the unit-specific averages. Average unit-level emissions reductions are approximately the same across these emissions equivalent policy designs. Slight differences (which disappear when numbers are expressed in millions of lbs) are due to the fact that the emissions cap is rarely met exactly. Note that whereas the average change in emissions across regimes is close zero (because emissions increases at relatively low damage firms balance out the emissions decreases at relatively high damage firms), levelized annual abatement costs increase by $0.02 M on average.

Table 4 reports simulation results aggregated by damage category and averaged across the Z simulation repetitions. The high damage category includes the 241 units with damage parameters that exceed $2180/ton. The low damage category includes the remaining 391 units with below average damage parameters. Transitioning from spatially uniform to spatially differentiated trading shifts approximately 6% of the permitted emissions (or 72 M lbs NOx/year) away from these high damage facilities and into regions where the emissions do less damage. This shift will unambiguously lower the damages caused by the permitted emissions. But this reallocation comes at a cost. As emissions levels increase among low damage units, abatement costs dall by approximately $50M/year in levelized annual costs. However, abatement costs in the high damage area increase by approximately $63 M per year. In sum, estimated levelized annual abatement costs under spatially differentiated trading are approximately $12 higher than simulated abatement costs under the observed, spatially uniform trading regime.

3.4 Estimating damages cause by permitted emissions

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).

Table 5 summarizes the simulation results. The first column reports the simulated equilibrium permit price, the levelized annual abatement costs, and the value of the avoided damages under the observed policy regime. The average permit price (averaged across simulation repetitions) is \$2.44/lb NOx. The observed average permit price that prevailed over the time period in which these compliance decisions were being made was \$2.25, within one standard deviation of the mean simulated price. The estimated benefits accruing from this policy (i.e. approximately \$1B per year in avoided damages from NOx emissions) exceeds the estimated levelized annual costs of \$707 M. Net welfare gains under the implemented policy design are thus estimated to be \$294M/year.

The second column of table 5 reports the results from the counterfactual, spatially differentiated trading regime. Levelized annual abatement costs increase by an estimated \$12 M (less than 2 percent). Marginal abatement costs also increase, the average simulated permit price increases to \$2.75/lb. Damages associated with permitted emissions are lower under the spatially differentiated design; the estimated benefits increase by approximately \$62 M/year, or 6 percent. Taken together, the net benefits under the policy that incorporates spatially differentiated trading increase by 17%, or almost \$50M/year.

3.5 Subsidizing "welfare-improving" pollution

For a very small subset of units, our point estimates of the damage parameters are negative. If we maintain our assumption that source-specific marginal damages are constant, a negative marginal damage parameter implies that a source should be paid to pollute. Thus far, we have assumed that this would not be feasible politically. But we are interested in estimating the cost of this constraint, conditional on the assumptions of the model.

We rerun our simulations of damage-based trading using a vector of unadjusted compliance ratios. That is, we no longer set the negative trading ratios to zero, so that emissions at sources with negative marginal damage parameter estimates are effectively subsidized. The third column of Table 5 reports these additional results. Estimated annual benefits (in terms of avoided damages) increase by almost \$16 M per year. These benefits derive from shifting some of the permitted emissions to units where emissions increases are predicted to decrease rates of ozone formation. Levelized annual abatement costs are not significantly impacted, suggesting that the units that offset the increased emissions at the units with negative marginal damages have very similar marginal abatement costs.

3.6 Setting the optimal cap

In section 2, we noted that political, jurisdictional, or other implementation constraints may result in an emissions cap that is too stringent- or not stringent enough. In this section, we compare the emissions constraint imposed in the NBP cap with the "optimal" cap implied by our estimated marginal damage parameters, abatement costs, and econometric model of firm-level compliance decisions.

In theory, assuming away other distortions or imperfections, the equilibrium permit price τ^d in the damage-based trading simulations should equal the average damage parameter $\bar{\delta}$.if the emissions cap is set optimally. Averaging across our estimated damage parameters yields a mean value of \$2180/ton. The equilibrium permit price in our damage-based trading simulations is \$5500/ton. By (24), this suggests that $\Delta < 0$ and the emissions cap imposed in the observed regulatory regime is too stringent. Figure 2 helps to illustrate how the welfare gains associated with damage-based trading will be undermined if the cap is set too stringently.

In order to investigate the practical implications of this apparent discrepancy, we first identify the aggregate level of emissions that result when the permit price in a damage-based trading regime is set at \$2180/ton NOx. Call this E^{**}.We compare simulated outcomes under emissions and damage-based trading in regimes that impose the less stringent cap E^{**}. Table 6 reports the results. Annual net benefits from the policy increase by approximately \$15 M as compared to an identical damage-based policy design that imposes the observed cap. And the benefits associated with damage-based trading, vis a vis emissions-based trading, increase to 32 percent.

Figure 6 helps to illustrate an important caveat with respect to these simulation results. The figure plots a time series of the NOx permit price. Early on, before the NBP took effect, permit prices were very close to our simulated price. This is the time period in which the vast majority of compliance decisions were being made. As we should expect, the permit price reflects the marginal ex ante expected cost of meeting the emissions cap. However, as the program got underway, several factors contributed to a tumbling of the permit price. In particular, firms discovered that reducing

NOx emissions was not as expensive as initially anticipated (Linn, 2008). Figure 6 illustrates how the permit price stabilized very close to our estimated average damage parameter. This implies that, conditional on our estimated marginal damages, the emissions cap imposed in the NBP appears to be close to optimal after all.

In sum, our NOx control cost estimates reflect the NOx abatement cost information and *ex ante* expectations that informed compliance decisions in this cap-and-trade program. These cost estimates proved to be too high. Using these ex ante expected costs is appropriate when simulating compliance decisions made in preparation for the NBP. However, using these cost estimates to evaluate the overall costs of the program likely results in an over-estimate of compliance costs, and thus an underestimate the net benefits of the NOx Budget Program. Unfortunately, comprehensive data on what compliance with the NBP actually cost are unavailable.

3.7 Pre-existing regulatory distortions

Finally, we consider the implications of pre-existing distortions in the industry in which the majority of the emissions reductions mandated by the NBP occur. The recent wave of electricity industry restructuring in the United States has resulted in significant inter-state variation in electricity industry economic regulation. Thus, in addition to having different production and abatement costs, generators in the NOx Budget Program face very different economic regulation and investment incentives. In particular, rate-base regulated plants are guaranteed to earn a rate of return on prudent investments in pollution abatement equipment, whereas plants operating in restructured electricity markets are offered no such assurances.

Averch and Johnson (1962) illustrate how, under certain conditions, regulated firms earning a positive rate of return on capital investment will find it profitable to invest more heavily in capital equipment than is consistent with cost minimization. Fowlie (2010) finds that economic regulation in the electricity industry has substantively impacted how electricity generating units chose to comply with the NOx Budget Program. In this paper, we investigate the extent to which these regulatory distortions affect the welfare gains from spatially differentiated permit trading.

On average, estimated marginal damages are higher among generating units that are subject to cost-of-service regulation.[17] Figure 3 provides a very stylized representation of how a pre-existing regulation that effectively lowers the perceived abatement costs among high damage sources can reduce the overall net benefits associated with spatially differentiated permit trading.

The policy setting we analyze is not so straightforward: there are hundreds of sources associated with a range of marginal damages, abatement cost functions are discontinuous, effects of the economic regulation are heterogeneous, and there are other distortions and imperfections at work. To assess the implications of economic regulation in the electricity industry on actual policy outcomes, we modify our simulations slightly. The econometric model is estimated separately for units in restructured electricity markets and rate regulated electricity markets, respectively. We use the coefficient estimates obtained using data from restructured electricity markets to parameterize the simulation model for all units. That is, we simulate firm responses to both emissions-based and damage-based emissions regulation in a counterfactual scenario in which all units operate in restructured electricity markets.

Table 7 reports the results. In the first column, our main results are reproduced as a basis for comparison. The second column summarizes results from the policy counterfactual in which all units are assumed to operate in restructured electricity markets. Removing the economic regulation increases the average simulated permit price considerably. A higher permit price is required to incentivize the mandated emissions reductions. More importantly, the estimated net benefits of spatially differentiated trading are reduced. Intuitively, this is because the pre-existing regulatory distortion is pushing the emissions-based equilibrium outcome closer to the optimal outcome, whereas it drives a wedge between the damage-based equilibrium outcome and the optimum.

4 Conclusion

This analysis explores the welfare implications of a series of alternative policy designs for the NO_x Budget Program (NBP), a regional cap-and-trade policy that manages emissions of NO_x produced by large industrial point sources in the Eastern U.S. The current program permits trading pollution

[17] The average damage parameter among units supplying restructured wholesale electricity markets is \$1758/ton NOx (standard deviation \$1217/ton). The average damage parameter among regulated units or units that are publicly owned and operated is \$2529/ton NOx (standard deviation \$1111).

allowances on a ton-for-ton basis across all regulated sources. The paper examines damage-based trading designs which establish exchange rates between firms calibrated to the relative damages caused by their emissions.

The motivation for this exercise is that prior research has shown that the damages due to NO_x emissions vary considerably according to where the emission occurs (source location). And since the NBP permits ton-for-ton trading, the current cost-effective regulatory design fails to appropriately capture this heterogeneity. The empirical analysis first computes the efficient trading ratios between each of the 632 boilers regulated under the NBP; the ratios of each pair of sources' marginal damage for NO_x emissions. Firms' responses to these trading ratios are modeled using an econometric model of the compliance decisions made by firms subject to the NBP (Fowlie, 2010). The outcome of firms' choices (and subsequent emission levels) is modeled using an integrated assessment model (Muller, Mendelsohn, 2007;2009). Aggregate abatement costs and environmental damage are then tabulated and compared to the extant cost-effective design. 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.

We find that, under the damage-based trading regime, levelized annual abatement costs increase by an estimated \$12 M (which is less than 2 percent). Marginal abatement costs also increase, the average simulated permit price increases to \$2.75/lb NO_x from \$2.44/lb. NO_x. Intuitively, damages associated with permitted emissions are lower under the spatially differentiated design. The damages due to NO_x emissions decrease by approximately \$62 M/year (6 percent). Taken together, the net benefits under the policy that incorporates spatially differentiated trading increase by 17%, or almost \$50M annually. It is important to note again that this welfare improvement is *not* due to a change in the total amount of emissions. Rather the policy featuring trading ratios calibrated to firms' marginal damages results in a spatial reallocation of emissions. This shift moves emissions from high-damage sources to sources that cause less damage per ton NO_x emitted. The policy recognizes that damages are not uniformly distributed across the regulated sources which the current cost-effective program fails to capture. And, as the results indicate, although this spatial reallocation of emissions generates larger total abatement costs, the reduction in damages outweighs this increase in costs by a substantial margin.

Additional analysis investigates the role of pre-existing distortions and political constraints in determining the magnitude of the returns to spatially differentiated trading. Removing a political constraint that limits regulators' ability to subsidize emissions in cases where additional emissions might actually reduce overall damages increases the relative benefits of damage-based trading by an estimated $15 M per year. Removing a pre-existing regulatory distortion in the electricity market would work in the opposite direction, reducing relative benefits by almost $20M per year.

These results are accompanied by some important caveats. First, in this working paper, we ignore the uncertainty surrounding damage parameter estimation; we take point estimates of the source-specific marginal damage parameters as given. Second, our source-specific estimates of NOx control costs capture expectations at the time that investments in NOx controls were being made. Anecdotal evidence suggests that ex ante expected control costs exceeded the costs that were actually realized. This would imply that our estimated net benefits of the NBP are conservative. Finally, we make no attempt to estimate the additional costs that could be associated with defining compliance in terms of damages. Our estimated net benefits of damage-based trading do not reflect the costs of designing and implementing a more complex permit market design.

References

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

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

[3] 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.

[4] 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.

[5] 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.

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

[7] 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.

[8] 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

[9] 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.

[10] 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.

[11] 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.

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

[13] 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.

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

[15] 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.

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

[17] 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.

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

[19] 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.

[20] 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.

[21] 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.

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

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

[24] 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.

[25] 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.

[26] 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.

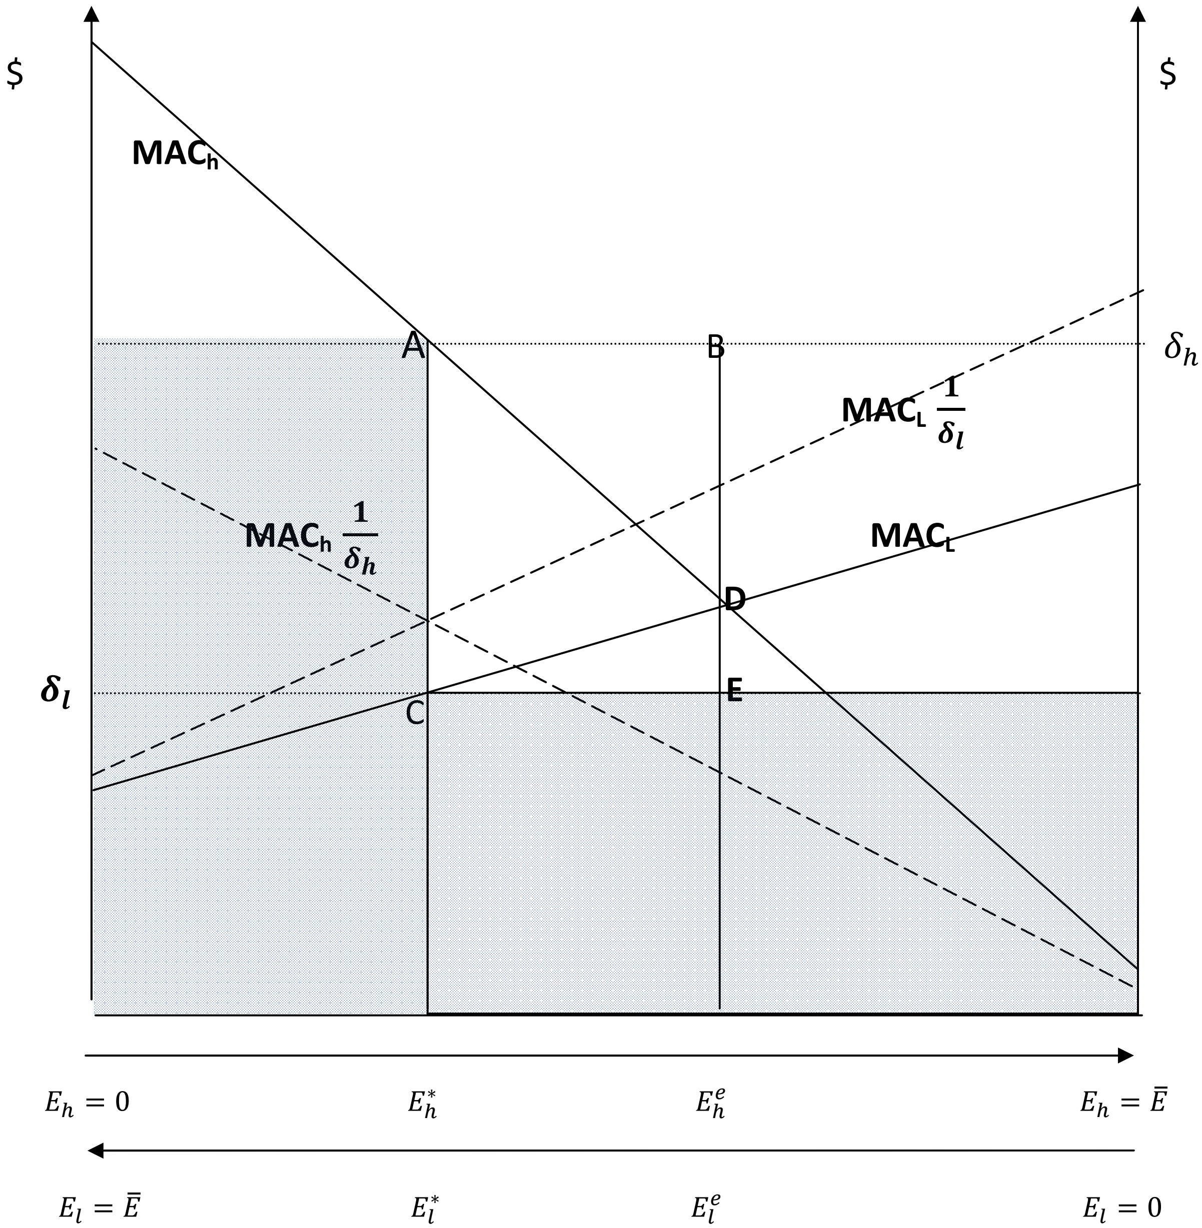

Figure 1: Emissions permit market outcomes under emissions-based and damage-based policies: First best setting

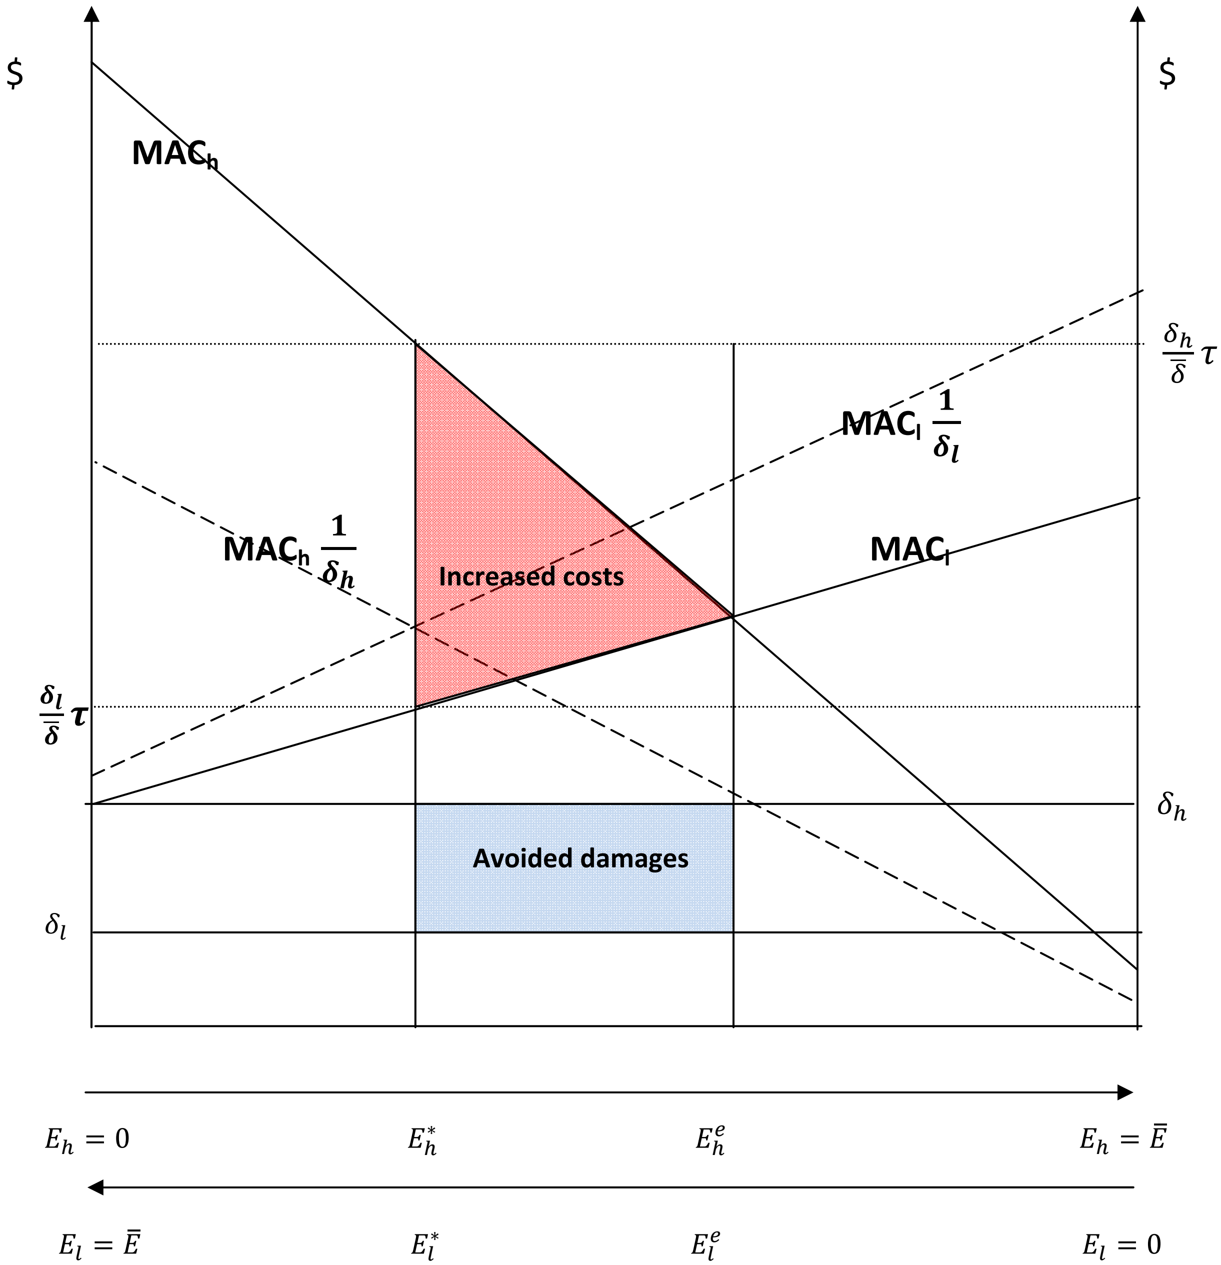

Figure 2: Emissions permit market outcomes under emissions-based and damage-based policies: Sub-optimal emissions constraint

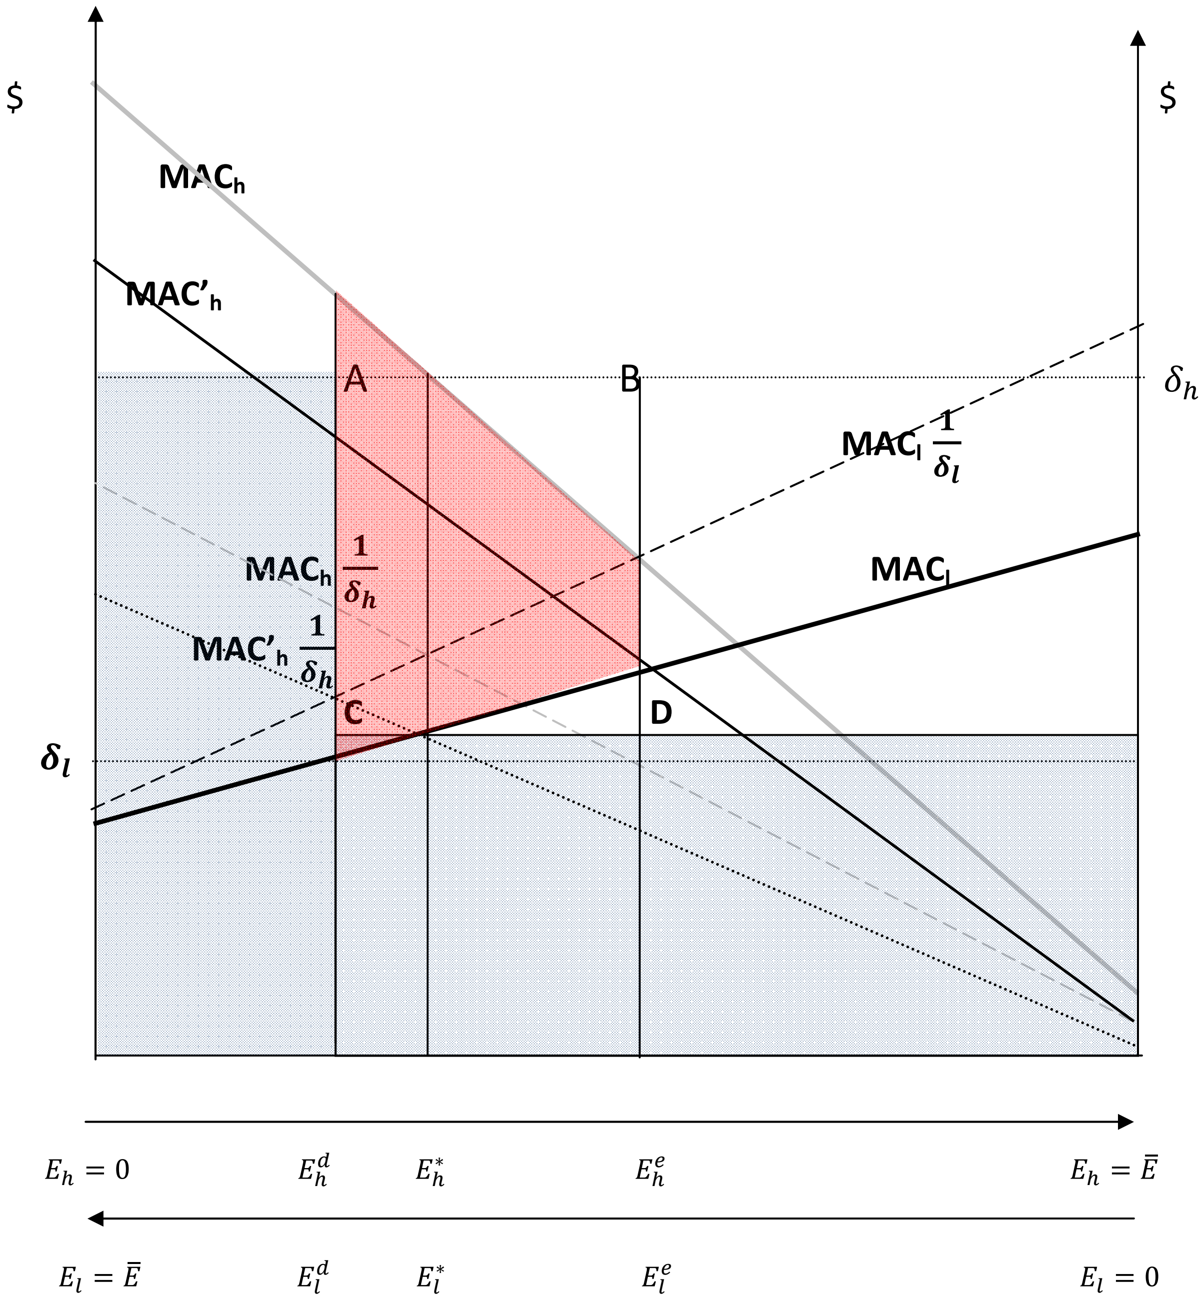

Figure 3: Emissions permit market outcomes under emissions-based and damage-based policies: Discounted costs for the high damage firm

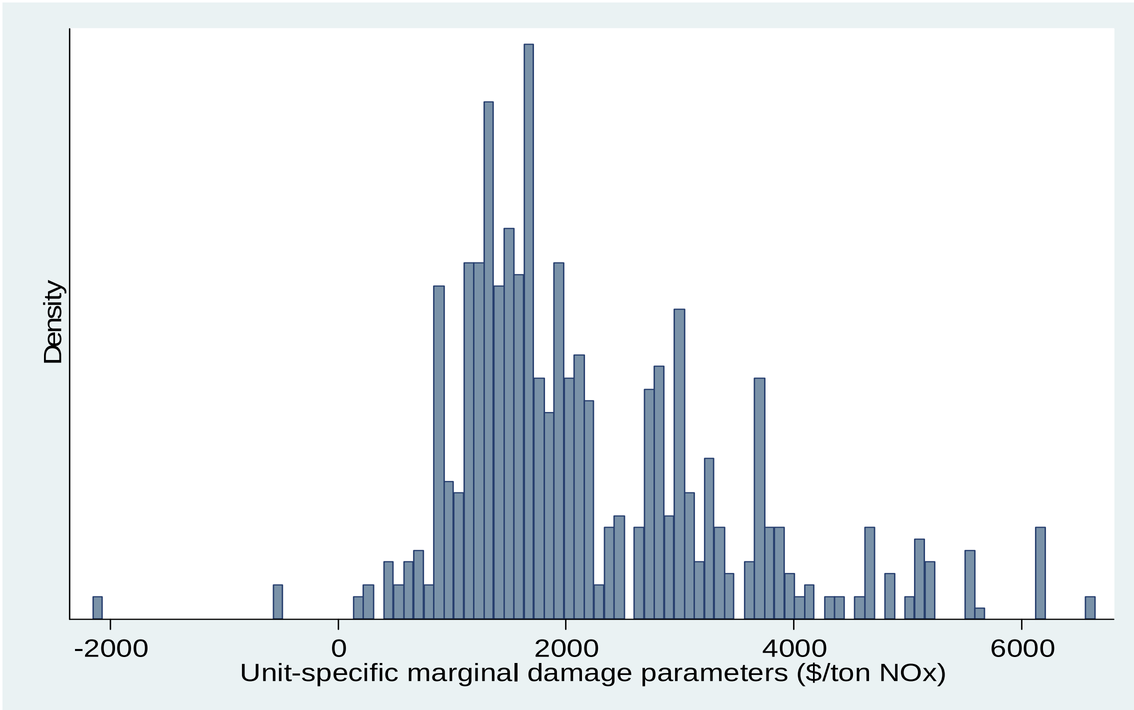

Figure 4: Histogram of unit-specific damage parameters estimated using the APEEP model

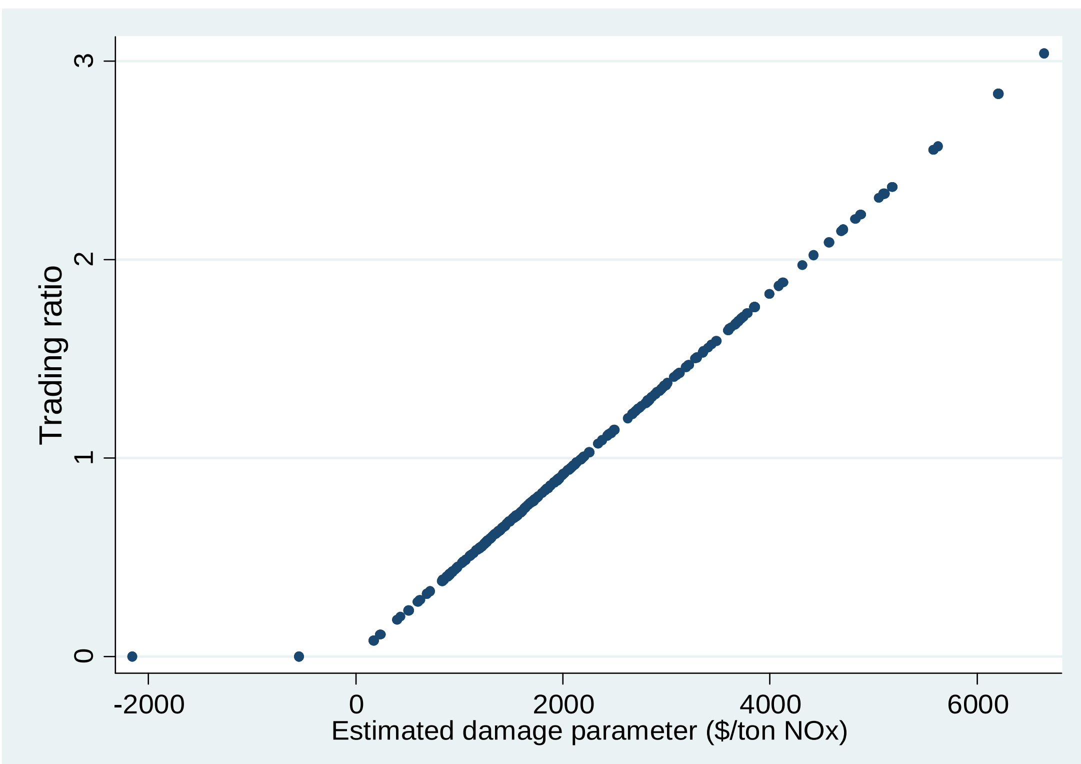

Figure 5: Trading ratios used in damage-based emissions trading simulations

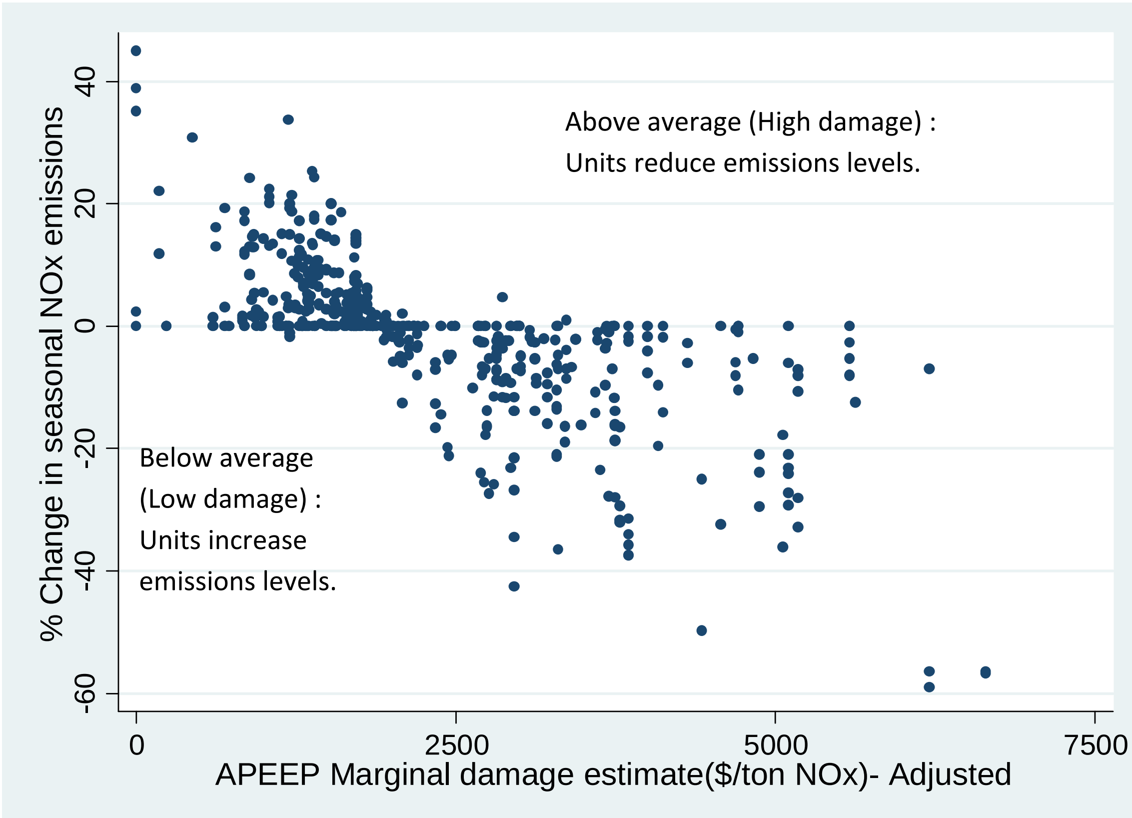

Figure 6: Source-specific changes in emissions induced by spatially differentiated NOx permit trading

Notes: Each point represents a different electricity generating unit. The horizontal axis measures the facility level damage estimates. "Adjusted" implies that any positive damage estimates have been set to zero, (see Figure 5). The vertical axis measures percent changes in simulated ozone season emissions in the observed, emissions based case less simulated emissions under the counterfactual exposure-based trading.

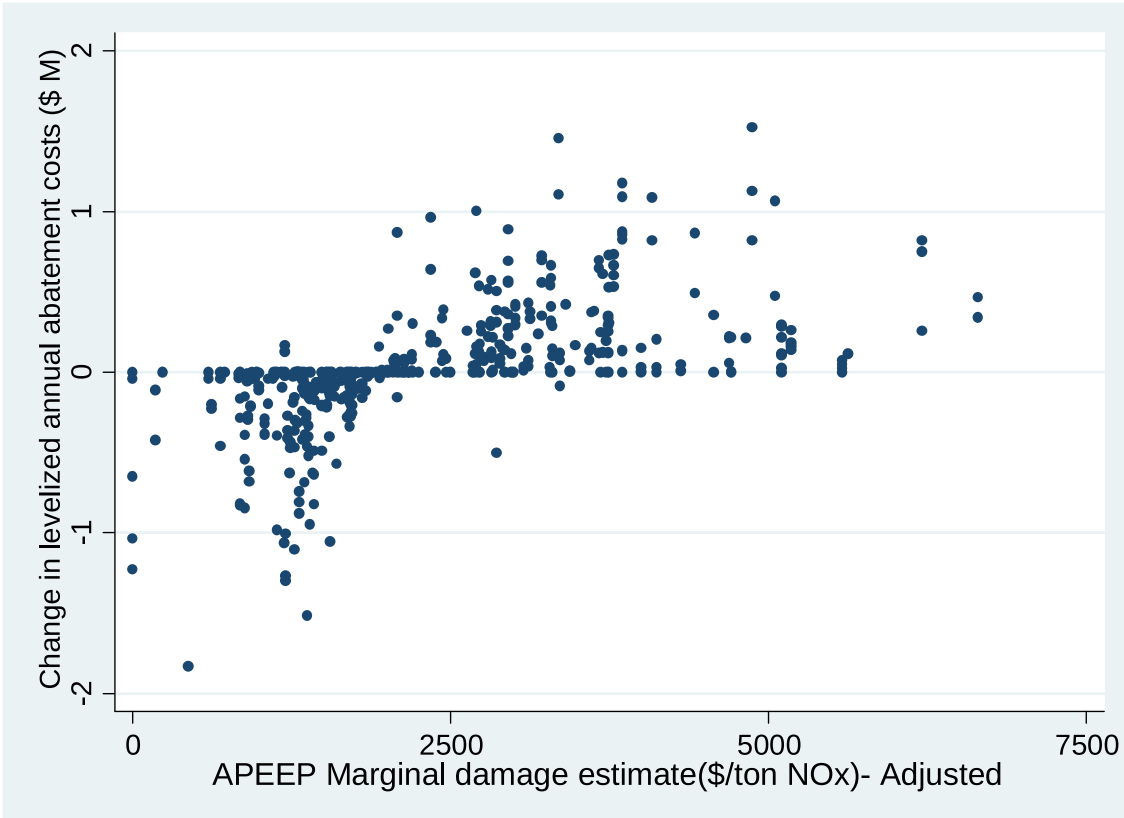

Figure 7: Source-specific changes in investment in emissions abatement induced by spatially differentiated NOx permit trading

Notes: Each point represents a different electricity generating unit. The horizontal axis measures the facility level damage estimates. "Adjusted" implies that any positive damage estimates have been set to zero, (see Figure 5). The vertical axis measures changes in simulated investments in pollution abatement in the observed, emissions based case less simulated emissions under the counterfactual exposure-based trading.

Figure 8: Observed and simulated NOx permit prices

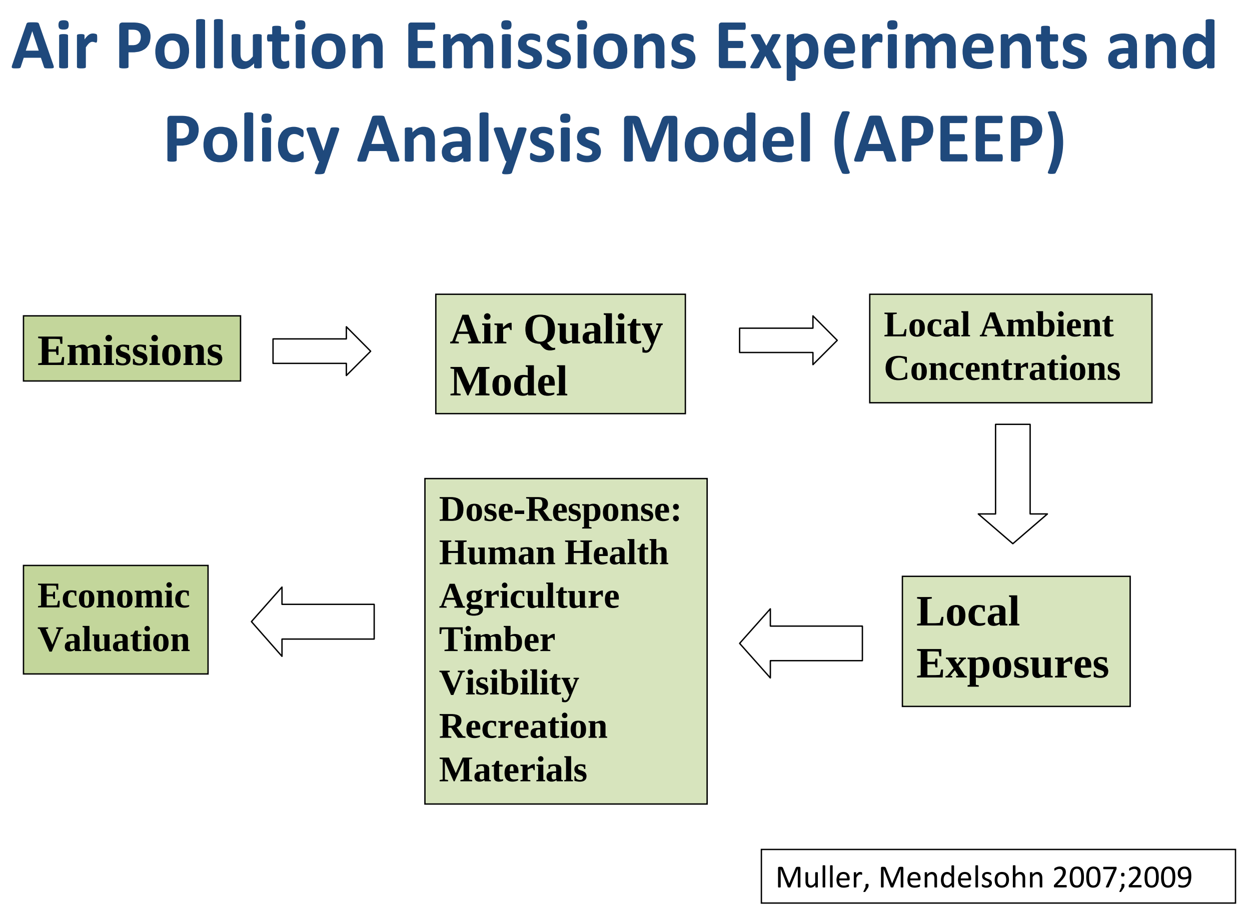

Figure 9. AP2 Model Structure

Table 1 : Unit-level summary statistics

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)
Low NOx burners with overfire air	31.30 (74.15)	26.63 (44.62)	0.64 (0.15)	0.64 (0.19)
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: Simulated unit-level emissions and investment in abatement costs

Outcome	Mean (Standard deviation)
Reductions under emissions-based trading relative to unregulated benchmark (millions of lbs)	1.9 (3.7)
Reductions under exposure-based trading relative to baseline (millions of lbs)	1.9 (3.7)
Change in emissions across regimes (absolute value)	239,015.90 (405,901)
Change in emissions across regimes (millions of lbs)	0.0 (0.5)
Levelized annual abatement cost under emissions-based trading ($M)	$1.12 ($1.87)
Levelized annual abatement cost damage based trading ($M)	$1.14 ($1.84)
Average change in costs across regimes ($M)	$0.02 ($0.34)

Notes: This table summarizes the results from simulating investment in NOx abatement and ozone-season emissions under the observed emissions-based trading regime and the counterfactual damage-based trading regime. The numbers in this table summarize how unit-level outcomes for the 632 coal-fired generating units in the data set. 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 4: Simulated emissions and investment in abatement by damage classification

		Emissions-based trading	Exposure-based trading (beneficial emissions omitted)
High damage facilities (above average)	Emissions reductions (Million lbs NOx)	492.8 (22.8)	564.3 (18.7)
	Percent of permitted emissions	38% (2%)	32% (1%)
	Levelized annual abatement costs ($M)	$258.72 ($19.08)	$321.36 ($14.51)
Low damage facilities (below average)	Emissions reductions (Million lbs NOx)	719.3 (23.1)	648.8 (19.8)
	Percent of permitted emissions	62% (1%)	68% (1%)
	Levelized annual abatement costs ($M)	$448.28 ($16.92)	$398.05 ($15.53)

Notes: This table summarizes the results from simulating investment in NOx abatement and ozone-season emissions under the observed emissions-based trading regime and the counterfactual damage-based trading regime. The numbers in this table summarize outcomes aggregated by damage classification. "High damage" units are those with above average damage parameter point estimates. "Low damage" units are those with below average damage parameter point estimates. These results are summarized across simulation repetitions, within policy scenarios. Standard deviations are in parentheses

Table 5: Simulated emissions and abatement investment by damage classification: Restricted and unrestricted damage-based trading

	Emissions-based trading	Exposure-based trading (beneficial damages omitted)	Exposure-based trading beneficial damages rewarded
Price ($/lb NOx)	$2.44 ($0.32)	$2.75 ($0.36)	$2.74 ($0.36)
Increase in levelized annual abatement cost ($M)	$707.00 ($12.49)	$719.41 ($11.42)	$719.40 ($11.61)
Reduction in damages from emissions ($1M)	$1,000.68 ($18.70)	$1,062.52 ($18.70)	$1,078.49 ($13.95)
Increase in costs vis a vis emissions-based trading ($M)	- -	$12.41 ($9.17)	$12.40 ($9.44)
Increase in benefits vis a vis emissions-based trading ($M)	- -	$61.84 ($13.58)	$77.81 ($13.51)
Welfare gain vis a vis emissions-based trading ($M)	- -	$49.43 ($15.28)	$65.41 (15.39)
Relative welfare gain (% terms)	- -	17% (6%)	23% 7%

Table 6 : Simulated emissions and abatement investment by damage classification: Optimal cap

	Emissions-based trading	Exposure-based trading Beneficial damages omitted
Price ($/lb NOx)	$1.00 ($0.09)	$1.09 ($0.36)
Increase in levelized annual abatement cost ($M)	$427.47 ($15.48)	$428.86 ($15.21)
Reduction in damages from emissions ($1M)	$637.21 ($19.87)	$703.92 ($20.14)
Increase in costs vis a vis emissions-based trading ($M)	- -	$1.39 ($12.76)
Increase in benefits vis a vis emissions-based trading ($M)	- -	$66.72 ($12.67)
Welfare gain vis a vis emissions-based trading ($M)	- -	$65.32 ($16.65)
Relative welfare gain (% terms)	- -	32% (11%)

Table 7: Simulated emissions and abatement investment by damage classification: Symmetric economic regulation

	Damage-based trading Observed regulatory regime	Damage-based trading Counterfactual economic regulatory regime
Price ($/lb NOx)	$2.75 ($0.36)	$7.63 ($3.55)
Increase in levelized annual abatement cost ($M)	$719.41 ($11.42)	$773.52 ($43.50)
Reduction in damages from emissions ($1M)	$1,062.52 ($18.70)	$1,048.10 ($31.05)
Increase in costs under damage-based trading ($M)	$12.41 ($9.17)	$11.20 ($11.51)
Increase in benefits under damage-based trading ($M)	$61.84 ($13.58)	$41.62 ($16.07)
Welfare gain relative to emissions-based trading ($M)	$49.43 ($15.28)	$30.42 ($17.12)
Relative welfare gain (% terms)	17% (6%)	14% (11%)