



EI @ Haas WP 237R

**Market-based Emissions Regulation When Damages
Vary Across Sources: What are the Gains from
Differentiation?**

Meredith Fowlie and Nicholas Muller
February 2013

**Revised version published in
*Journal of the Association of
Environmental and Resource Economists*
6(3), 593-632, May 2019.**

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board.

© 2013 by Meredith Fowlie and Nicholas Muller. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

<http://ei.haas.berkeley.edu>

Market-based emissions regulation when damages vary across sources: What are the gains from differentiation?

Meredith Fowlie and Nicholas Muller*

February 5, 2013

Abstract

Much of the air pollution currently regulated under U.S. emissions trading programs is non-uniformly mixed, meaning that health and environmental damages depend on the location and dispersion characteristics of the sources. Existing policy regimes ignore this fact. Emissions are penalized at a single permit price, regardless of the location of the source. In theory, differentiated policies can be designed to accommodate non-uniformly mixed pollution using emissions penalties that vary with emissions damages. Under perfect certainty, damage-based policy differentiation is unambiguously welfare improving. In the presence of uncertainty about damages and abatement costs, differentiated policies may not welfare dominate simpler, undifferentiated designs. Using rich data from a major U.S. emissions trading program, we estimate the welfare impacts of policy differentiation. Surprisingly, we find that differentiated emissions trading results in welfare *loss* as compared to the undifferentiated trading regime that was implemented. This result manifests because ex post realized abatement costs appear to have exceeded expectations. We further show that, in this context, a differentiated price-based policy welfare dominates the differentiated quantity-based alternative.

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

JEL Classifications: Q54, Q53, Q58

*Fowlie: Department of Agricultural Economics, UC Berkeley, and NBER. Muller: Department of Economics and the Environmental Studies Program, Middlebury College and NBER. We thank Chris Costello, Larry Goulder, Michael Greenstone, Jacob LaRiviere, Steven Puller, Mar Reguant, Leo Simon, Christian Vossler, and seminar participants at Carnegie Mellon, Duke, Resources for the Future, Stanford University, the University of California Berkeley, University of Pittsburgh, the University of Tennessee, the US EPA, and Yale University for helpful discussions and comments. We thank Dilek Uz for invaluable research assistance. Fowlie thanks the Energy Institute at Haas for support under a research contract from the California Energy Commission. Muller thanks the United States Environmental Protection Agency for support under award: EPA-OPEI-NCEE-08-02.

1 Introduction

Economists have long advocated for market-based approaches to pollution regulation (Montgomery, 1972; Baumol and Oates, 1988). The past three decades have witnessed large scale experimentation with market-based policy designs. By many measures, this experimentation has been very successful. Targeted emissions reductions have been achieved or exceeded, and it is estimated that total abatement costs have been significantly less than what they would have been in the absence of the trading provisions (Fowlie et al., 2012; Keohane, 2007; Stavins, 2005).

In terms of economic efficiency, however, many existing policies fall short of the theoretical ideal. Efficiency requires that the marginal cost of pollution reduction be set equal to the marginal damage caused by emissions. If pollution is "non-uniformly mixed" (i.e. health and environmental damages from pollution depend on the location of the source), efficiency requires that the marginal costs of pollution abatement should vary across sources according to the degree of damage caused (Montgomery, 1972; Tietenberg, 2006). To achieve this, policy incentives should reflect differences in marginal damages across sources. However, the majority of existing and planned market-based emissions regulations are implemented as spatially uniform, "undifferentiated" policies; all regulated emissions are penalized at the same tax rate or permit price.

Some of the most pernicious air quality problems in the United States involve non-uniformly mixed pollutants. Examples include nitrogen oxides (NO_x) and sulfur dioxide (SO₂), which are two criteria pollutants currently subject to undifferentiated market-based regulation. In 2008, the Federal Court of Appeals ruled that this regulatory approach fails to adequately accommodate spatial transport of pollution and the associated variation in damages across sources.¹ Since that time, debates surrounding market-based regulation of non-uniformly mixed pollutants have escalated. Our aim in this paper is to investigate the welfare implications of regulating non-uniformly mixed pollutants using undifferentiated - versus differentiated- policy instruments.

In theory, market-based policies can be designed to accommodate non-uniformly mixed pollution. Baumol and Oates (1988) use a general equilibrium model to depict optimal pollution taxes in a setting with heterogeneous costs and damages. The optimal tax rate is calibrated to the marginal damage caused by emissions. When damages vary by source, so do the tax rates. Others have proposed so-called "differentiated" emissions permit market designs wherein differences in marginal damages are reflected in different compliance requirements and incentives (Mendelsohn, 1986; Tietenberg, 1995; Muller and Mendelsohn, 2009).

A simple theoretical model provides a foundation for our analysis. We begin by considering the canonical case in which a regulator equipped with perfect information seeks to minimize

¹On July 11, 2008, in *North Carolina v. EPA*, the U.S. Court of Appeals for the D.C. Circuit vacated the Clean Air Interstate Rule (CAIR) which was intended to provide a cost-effective, market-oriented approach to regulating a non-uniformly mixed pollutant (subject to compliance feasibility). *State of North Carolina v. Environmental Protection Agency*, No. 05-1244, slip op. (2008), District of Columbia Court of Appeals.

pollution damages plus abatement costs. In this setting, moving from an undifferentiated to a differentiated policy design will be unambiguously welfare enhancing so long as policy parameters are set optimally and damages differ across sources. As in Mendelsohn (1986), we show that the gains from policy differentiation will depend on both the extent of the variation in damages across sources and the slopes of the marginal abatement cost functions.

Although the perfect information case provides a useful jumping-off point, it ignores an important feature of real-world policy settings. Information imperfections have played a significant role in constraining the design and implementation of market-based emissions policies, and non-uniformly mixed pollution regulation in particular. We thus extend the baseline model to more accurately capture the information constraints that complicate policy implementation in practice. We consider both uncertainty regarding source-specific damages and imperfect information about abatement costs.

Building on this theoretical foundation, we examine the welfare implications of policy differentiation in the context of a landmark emissions trading program. The NOx Budget Program (NBP) imposed a binding cap on nitrogen oxide emissions from large point sources in the Eastern United States. In the design stages, policy makers considered imposing damage-based restrictions on interregional trading.² Ex ante policy simulations projected only nominal gains from policy differentiation (Krupnick et al. 2000, US EPA, 1998). Consequently, the program was implemented as a single jurisdiction, spatially uniform trading program.

This paper revisits the decision to forego spatially differentiated NOx trading in favor of the simpler undifferentiated alternative. We begin by assessing the benefits and ex post realized costs incurred under the NOx Budget Program as implemented. We then examine how these welfare impacts would have manifested differently under damage-differentiated policy regimes. This retrospective look at both factual and counterfactual policy regimes is motivated by the following three observations:

First, integrated assessment models now provide much richer estimates of the variation in damages from pollution. Recent advances in data collection and computer-analytics have improved modeling of pollution formation, fate, and transport. Advances in epidemiology and related fields have improved our understanding of exposure pathways and dose-response relationships. In principle, these advances allow for more informed and data-driven policy design and implementation (Esty, 2004). We use a state-of-the-art stochastic integrated assessment model, AP2, to estimate marginal damages for each facility in the NBP.³ In contrast to the pollution damage estimates that informed the design of the NOx Budget Program, we document significantly more spatial variation in damages across regulated sources. All else equal, this suggests larger gains from policy differentiation.

²See the Federal Register 63(90), Monday, May 11, 1998, page 25902.

³AP2 is the stochastic version of the APEEP model (Muller, Mendelsohn, 2009).

Second, policy analysts have duly noted that standard approaches to ex ante modeling may fail to capture key features of the real world decision processes that drive emissions abatement decisions (Krupnick et al. 2000). Deterministic, compliance cost-minimization algorithms, together with engineering estimates of the capital and operating costs, are typically used to inform policy design and implementation. If firms' environmental compliance choices deviate significantly from the predictions of these models, anticipated policy outcomes may not materialize. With the benefit of hindsight, we estimate an econometric model of firms' compliance choices. This model is used to simulate firms' response to counterfactual policy incentives and to approximate ex post realized abatement costs. Importantly, we find substantive differences between ex post observed compliance choices and those predicted by engineering cost minimization.

The third observation pertains to the policy relevance of the issues we address. In the years since the NOx Budget Program was implemented, debates about how to regulate non-uniformly mixed pollution have escalated. In court proceedings, the need for more sophisticated analysis has been identified.⁴ Recent improvements in data quality and analysis pave the way for a more sophisticated regulatory response to non-uniformly mixed pollution problems. We bring state-of-the-art integrated assessment modeling of emissions damages, together with econometric modeling of firms' responses to market-based emissions regulation incentives, to this analysis of differentiated policy alternatives.

Our paper contributes to an important empirical literature examining the benefits and costs incurred under market-based emissions regulation. Papers in this literature generally fall into one of two categories: ex ante analysis of likely outcomes under proposed regulations (e.g. Fullerton, 1997; Krupnick et al., 2000); and ex post evaluation of realized outcomes (e.g. Deschenes et al., 2012; Joskow et al., 1998; Fowlie, 2010). Policy makers have stressed the importance of identifying differences between the assumptions underlying ex ante analysis and ex post realized outcomes in order to inform and improve future ex ante cost analysis (US EPA, 2012). This importance notwithstanding, very little work has been done to systematically compare ex ante modeling and ex post analysis.⁵ This paper puts ex ante policy simulation and ex post evaluation on the same footing. We document discrepancies between observed compliance choices and standard assumptions of ex ante analysis. Although these discrepancies imply higher-than-expected abatement costs, we still find that the benefits from the program exceed costs by an estimated \$264 M per year.⁶

⁴See *State of North Carolina v. Environmental Protection Agency, No. 05-1244*, slip op. (2008), District of Columbia Court of Appeals.

⁵Carlson et al. (2000) is an important exception.

⁶Our measure of avoided damages accounts for adverse effects on human health, reduced yields of agricultural crops and timber, reductions in visibility, enhanced depreciation of man-made materials, and damages due to lost recreation services. Notably, we do not account for reductions in pharmaceutical purchases and other defensive expenditures. Deschenes et al. (2012) find that improvements in air quality achieved under the NBP reduced drug expenditures by approximately \$900 million annually. Our estimated annual benefits are therefore conservative.

Having demonstrated the positive welfare gains conferred by undifferentiated emissions trading, we then ask whether these gains would have been larger under differentiated trading. Adopting standard assumptions with respect to valuation parameters and cost minimization modeling, we find that damage-differentiated emissions trading between sources would likely have *reduced* welfare *vis a vis* undifferentiated trading. This surprising result stands in stark contrast to recent work which suggests significant welfare gains from policy differentiation in the context of the Acid Rain Program (Muller and Mendelsohn, 2009). Unlike previous studies, we use an econometrically estimated model of firms' compliance decision to simulate firm-level responses to observed and counterfactual policy incentives.⁷ When ex post realized abatement costs are larger than expected, the costs of policy differentiation can exceed the benefits.

Finally, our work contributes to the classic literature that investigates the relative efficiency of price and quantity-based policies under uncertainty. Echoing Weitzman (1974), we find that price and quantity-based regimes are associated with distinctly different outcomes when abatement costs are uncertain and policy parameters must be fixed before uncertainty is resolved. In the particular case we analyze, policy differentiation increases welfare in the price-based regime, but reduces welfare in the quantity-based regime. Also consistent with the earlier literature (e.g. Baumol and Oates, 1988; Fishelson, 1976) we find that the expected social loss (*vis a vis* the ex post optimum) is the same under differentiated price and quantity-based policy regimes when uncertainty is limited to randomness in the parameters of the marginal benefit function.

The remainder of this paper is organized as follows. Section 2 introduces a simple model useful for analyzing the welfare implications of damage-based policy differentiation. Section 3 provides an overview of the NOx Budget Trading Program. Section 4 describes how we construct our estimates of the potential gains from policy differentiation in the NBP. Section 5 summarizes our results. Concluding remarks are offered in Section 6.

2 Theoretical framework

The theoretical framework developed in this section serves three purposes. First, the model is used to derive the differentiated and undifferentiated policy designs that are the focus of the more detailed, applied analysis to follow. Second, we derive an intuitive expression for the gains from policy differentiation which will aid in the interpretation of the policy simulation results. Finally, the modeling framework is used to explore the implications of uncertainty regarding both abatement costs and damages.

⁷There are some additional distinguishing characteristics of our paper *vis a vis* prior work looking at the acid rain program. We employ stochastic (versus deterministic) marginal damage estimates which allows us to characterize uncertainty surrounding our estimates. Also, nitrogen oxide behaves differently than sulfur dioxide once emitted. This has implications for the degree to which damages vary across sources.

We suppose that a firm's abatement cost function is a quadratic function of the level of abatement. In most cases, including the policy setting that is of primary interest here, a quadratic functional form provides a reasonable approximation to the true form.⁸ We transform this quadratic abatement function so that costs are expressed as a function of emissions: $C_i(e_i) = \alpha_{0i} - \alpha_{1i}e_i + \beta_i e_i^2$ (see Appendix 1). We assume that $C'_i(e_i) \leq 0 \leq C''_i(e_i)$. We further assume that the regulator has estimates of the marginal cost function parameters $\{\alpha_0, \alpha_1, \beta\}$ for each source.⁹

We capitalize on a series of empirical simulations (discussed in Appendix 3) to impose structure on the function that maps emissions to aggregate damages $D(e_1, \dots, e_N)$. These damages depend not only on the aggregate level of emissions, but also on how the emissions are distributed across sources.¹⁰ We assume that the damage function is linear in source-specific emissions and additively separable.

We define source-specific pollution damage functions in terms of emissions: $D_i(e_i)$. For each source, we define a marginal damage parameter $D'_i(e_i) \equiv \delta_i$. The product of the marginal damage δ_i times the corresponding unit-level emissions e_i yields damages caused by emissions at unit i . Summing across these yields the total damages: $D(e) = \sum_i \delta_i e_i$. We suppose that the policy maker's objective is to minimize the total social cost TSC associated with emissions of this pollutant:

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

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

$$-C'_i(e_i^*) = \delta_i \quad \forall i. \quad (2)$$

Intuitively, marginal costs are set to equal marginal damages across all sources. The (*) superscript denotes efficient emissions levels. Solving for e_i^* conditional on the functional form assumptions outlined above yields an expression for the optimal level of emissions at source i :

⁸In the context of the NOx Budget Program, source-specific marginal abatement cost functions are more accurately summarized as step functions. For ease of exposition, we assume a quadratic functional form in the theory model. But we release this assumption in the policy simulations.

⁹This assumption regarding the regulator's information endowments reflects how the US EPA (for example) has set aggregate quantity limits using firm-specific cost estimates in the course of their regulatory decision-making processes.

¹⁰In this analysis, we will focus exclusively on the spatial heterogeneity in damages. See Joskow, Martin and Ellerman (2007) for an analysis of the implications of temporal variation in damages.

$$e_i^* = \frac{\delta_i - \alpha_{1i}}{2\beta_i}. \quad (3)$$

Having characterized the first-best emissions outcome, we can evaluate the performance of alternative market-based policy designs against this benchmark.

2.1 Market-based regulation of non-uniformly mixed pollution under perfect certainty

We are interested in evaluating the performance of both quantity-based and price-based policy instruments. In the United States, emissions trading programs that place a limit on the quantity of permitted emissions are used to regulate non-uniformly mixed point source pollution at both the federal and regional level.¹¹ In Europe, emissions taxes are more commonplace.¹²

Price and quantity-based policy instruments can be implemented as "differentiated" or "undifferentiated" programs. Undifferentiated policies effectively treat emissions as if they cause the same degree of environmental damage on a per-ton basis. In contrast, under a differentiated policy regime, the cost of offsetting a unit of pollution varies according to the degree of damage caused.

We develop a simple two-firm model to motivate the four policy designs that are the focus of this paper. The two firms differ in terms of the emissions damages they cause. Emissions at firm (H) cause relatively high damages; emissions at firm (L) cause relatively low damages. As a jumping off point, we begin by assuming that policy makers know the source specific marginal damage parameters and marginal cost parameters for both firms with certainty. Given this complete information, we derive welfare maximizing policy parameters.

2.1.1 The differentiated tax

To derive the differentiated tax structure that minimizes total social costs as defined by (1), we begin by assuming that both the low and high damage firms seek to minimize the sum of their abatement costs and tax payments. Using the low damage firm as an example, the firm's cost minimization problem is:

$$\min_{e_i} : TC_L = \alpha_{0L} - \alpha_{1L}e_L + \beta_L e_L^2 + \tau_L e_L, \quad (4)$$

where TC denotes the total cost of complying with the program (i.e. abatement costs plus tax liabilities) and τ_L denotes the tax for the low-damage firm. First order conditions for a minimum

¹¹Examples include the Federal Acid Rain Program, the NOx Budget Program, and the Regional Clean Air Incentives Market in Southern California.

¹²Examples include France's tax on NOx and volatile organic compounds, and Sweden's tax on NOx (Millock, Nauges, and Sterner, 2004).

imply that the firm sets marginal abatement cost equal to the tax rate it faces (see Appendix 1.2). The corresponding cost-minimizing emissions level is: $e_L^D = \frac{\alpha_L - \tau_L}{2\beta_L}$. For the high damage firm, the cost-minimizing emission level is $e_H^D = \frac{\alpha_H - \tau_H}{2\beta_H}$. The welfare maximizing tax rates can be found by evaluating TSC at $\{e_L^D, e_H^D\}$ and then differentiating TSC sequentially with respect to (τ_L, τ_H) . This yields the well-known result that the source-specific tax rates should be set equal to the corresponding source-specific marginal damages (see Appendix section 1.2).

2.1.2 The undifferentiated tax

In real-world policy contexts, it may not be possible to levy a tax that varies across sources according to the marginal external cost. Consequently, undifferentiated taxes are far more common. We consider the case in which the regulator is constrained to tax all firms at the same rate.

We maintain the assumption that firms seek to minimize the sum of their abatement costs and tax payments. Compliance cost minimizing firms set their marginal abatement costs equal to the uniform tax rate. The corresponding cost-minimizing emissions levels e_L^U and e_H^U are derived in Appendix 1.3. We evaluate the social welfare function TSC at $\{e_L^U, e_H^U\}$ and then differentiate TSC with respect to τ . Setting $\frac{\partial TSC}{\partial \tau}$ equal to zero and solving for τ yields an expression for the social cost minimizing uniform tax $\tau^* = \frac{(\beta_H \delta_L + \delta_H \beta_L)}{\beta_H + \beta_L}$. This is analogous to the well-known weighted average formula of Diamond (1973) defining the optimal uniform tax rate on activities that generate non-uniform external costs. In our case, the optimal undifferentiated tax is a weighted average marginal damage. Damages associated with less steeply sloped marginal abatement cost curves are weighted more heavily.

It is straightforward to show that, conditional on our maintained functional form assumptions, aggregate emissions are the same under both the first-best differentiated and second-best undifferentiated tax policies: $E^* = \frac{(\alpha_L \beta_H + \alpha_H \beta_L - \beta_L \delta_H - \beta_H \delta_L)}{2\beta_L \beta_H}$. Appendix 1.3 demonstrates this emissions equivalence result.

2.1.3 Undifferentiated emissions trading

In the emissions trading programs we consider, a fixed quantity of tradeable emissions permits is allocated to participating sources either by auction or a gratis using some allocation rule that does not depend on production decisions going forward. Any free allocation of permits to firm i is represented by the initial allocation A_i , $i = \{L, H\}$. Let $\{A_{sij}, A_{bij}\}$ represent permits sold by firm i to firm j and bought from firm i by firm j at permit price τ . In the context of an undifferentiated

permit system firm L faces the following problem:

$$\begin{aligned} \min_{e_j} TC_L &= \alpha_{0L} - \alpha_{1L}e_L + \beta_L e_L^2 + \tau(A_{bLH} - A_{sLH}) \\ s.t. \quad e_L &\leq A_L - A_{sLH} + A_{bLH} \end{aligned} \tag{5}$$

Intuitively, a cost-minimizing firm will equate emissions abatement costs on the margin with the prevailing permit price τ (see Appendix 1.4):

$$-\alpha_{1L} + 2\beta_L e_L = \tau.$$

Having characterized the firms' response, we consider the regulators' constrained optimization problem. Conditional on implementing an undifferentiated emissions trading program, the regulator chooses the emissions cap that minimizes total social costs as defined in (1). In Appendix 1.4, we demonstrate that welfare maximizing emissions cap is precisely equal to the aggregate emissions observed under the undifferentiated (and differentiated) tax regimes.

2.1.4 Differentiated emissions trading

We now consider how the undifferentiated emissions trading program design can be modified so as to achieve the socially optimal allocation of emissions defined by (3). There is a growing literature that examines damage-differentiated, quantity-based policies (Teitenberg, 1995; Farrow et al., 2004; Horan and Shortle, 2005; Muller and Mendelsohn, 2009). To date, work in this area has focused primarily on using the ratio of marginal damages between each pair of regulated sources to define the terms of compliance.¹³ It is straightforward to operationalize this form of policy differentiation within our simple analytical framework.

Let $\bar{\delta}$ represent the average of the marginal damage across all sources in a trading program. In our simple two firm case, $\bar{\delta} = \frac{\delta_L + \delta_H}{2}$. We construct firm-specific ratios r_i , normalizing each firm's marginal damage by the mean damage parameter. For example, the ratio for the low-damage firm is $r_L = \frac{\delta_L}{\bar{\delta}}$. To remain in compliance, this firm must hold r_L permits to offset each unit of uncontrolled emissions. Under this form of policy differentiation, the firm's compliance constraint is slightly modified with respect to (5). The compliance cost minimization problem becomes:

$$\begin{aligned} \min_{e_j} TC_L &= \alpha_{0L} - \alpha_{1L}e_L + \beta_L e_L^2 + \tau(A_{bLH} - A_{sLH}) \\ s.t. \quad r_L e_L &\leq A_L - A_{sLH} + A_{bLH} \end{aligned} \tag{6}$$

¹³It is worth pointing out that the phrase "trading ratio" is somewhat misleading insofar as these ratios affect emissions trading activities only indirectly via the effect on compliance requirements.

In Appendix 1.5, we show that a compliance cost minimizing firm will reduce emissions until the cost of abatement on the margin equals the prevailing permit price weighted by the firm-specific trading ratio r . In equilibrium, the ratio of marginal costs between the two firms is: $\frac{-\alpha_{1L} + 2\beta_L e_L}{-\alpha_{1H} + 2\beta_H e_H} = \frac{-\lambda_L r_L}{-\lambda_H r_H}$, where λ represents the shadow value of the firm's regulatory compliance constraint. Since $\tau = \lambda_L = \lambda_H$, we have:

$$\frac{\alpha_{1L} - 2\beta_L e_L}{\delta_L} = \frac{\alpha_{1H} - 2\beta_H e_H}{\delta_H} \quad (7)$$

Note that this is congruent with the ratio of first-best emission levels.

Conditional on implementing a differentiated trading program described above, the regulator must choose how many permits to distribute. In Appendix 1.5, we derive an expression for the permit allocation that minimizes total social costs. Given this permit allocation, the optimal outcome is achieved, and the aggregate emissions are equivalent to those observed under the other policy regimes.

2.1.5 Welfare maximizing policy parameters

We have thus far derived the following optimal tax rates (differentiated and undifferentiated)

$$\tau_i^* = \delta_i \quad (8a)$$

$$\tau^* = \frac{(\beta_H \delta_L + \delta_H \beta_L)}{\beta_H + \beta_L} \quad (8b)$$

and the following optimal permit allocations (differentiated and undifferentiated):

$$A^{U*} = \frac{\beta_H \alpha_L - \beta_H \delta_L - \delta_H \beta_L + \alpha_L \beta_L}{2\beta_H \beta_L} \quad (9)$$

$$A^{D*} = \frac{\beta_H \alpha_L \delta_L + \delta_H \beta_L \alpha_H - \beta_H \delta_L^2 - \delta_H^2 \beta_L}{\beta_H \beta_L (\delta_H + \delta_L)}. \quad (10)$$

Each of these is chosen to maximize social welfare given the assumed structure of the corresponding policy regime. We can now make a simple observation:

(R1) Accurate information regarding the source-specific cost parameters is required to efficiently implement three of the four market-based policies we consider. The differentiated tax is the one exception.

Intuitively, when source-specific damages are linear in emissions, information about abatement costs is not needed to define the optimal tax rates τ_i^* .

2.2 Welfare Gains from Differentiation: No Uncertainty.

Although policy differentiation has no impact on aggregate emissions, it does affect how emissions are distributed across sources in equilibrium.¹⁴ In both price and quantity-based regimes, policy differentiation reduces emissions at the source that causes relatively high damage source and increases emissions at the relatively low damage source:

$$\begin{aligned} e_L^U - e_L^D &= \frac{\delta_L - \delta_H}{2(\beta_H + \beta_L)} < 0 \\ e_H^U - e_H^D &= \frac{\delta_H - \delta_L}{2(\beta_H + \beta_L)} > 0 \end{aligned}$$

This reallocation of emissions has potentially significant welfare implications. We first illustrate these welfare changes graphically. We then derive an analytical expression for the gains from policy differentiation which will prove useful when interpreting our more detailed policy simulation results.

Figure 1 illustrates the simple two firm case. The width of this figure, measured in units of emissions, is equal to the optimal aggregate emissions level E^* . Given the emissions equivalence result demonstrated in the previous section, we hold total emissions constant at E^* across the differentiated and undifferentiated policy regimes. At the left origin, all emissions occur at the low damage firm 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. $e_H = E^*$) 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 $C'_h(e_h)$.¹⁵

Equilibrium emissions under an undifferentiated policy regime are given by $\{e_l^u, e_h^u\}$. This equilibrium occurs at the intersection of $C'_h(e_h)$ and $C'_l(e_l)$ which is congruent with the first-order condition for cost-minimization under both the undifferentiated tax and the undifferentiated emissions trading regime described above. This allocation of permitted emissions minimizes the total abatement costs required to meet the emissions cap. However, this is not the optimal outcome. Total social welfare could be improved by shifting some of the permitted emissions away from the high damage source to the low damage source.

Equilibrium emissions under either the differentiated tax or the differentiated trading regime are given by $\{e_L^*, e_H^*\}$. The broken lines represent the marginal abatement cost schedules scaled

¹⁴This emissions equivalence does depend on the linearity of the marginal abatement cost curves. If marginal abatement cost curves are only approximately linear, the equivalence of optimal emissions across the differentiated and undifferentiated tax case will only approximately hold. This point will be made clearly in the subsequent applied analysis.

¹⁵In Figure 1, we assume that the relatively high damage firm also faces relatively higher costs of abatement. If the reverse were true, the optimal cap E^* would change, but it would still be the case that too much of the permitted emissions would be allocated to the high damage firm under the undifferentiated policy.

by the inverse of the corresponding marginal damage: $C'_i(e_i)\frac{1}{\delta_i}$, $i = l, h$. By (7), the allocation of emissions under damage-differentiated trading occurs where these broken lines intersect. This allocation of the permitted emissions achieves the optimal trade off between abatement costs and benefits from reduced damages.

The benefits from policy differentiation (in the form of reductions in damages vis a vis the undifferentiated equilibrium) are represented by area ABCE. The increase in abatement costs is equal to area ACD. The net benefits from differentiation are thus defined by the shaded areas $ABD + CDE$. The net gains from policy differentiation are unambiguously positive in this setting.

Using the functional forms from above, it is straightforward to derive an expression for the net benefits from differentiation in terms of the model parameters (see Appendix 2):

$$TSC^U(\delta, \beta) - TSC^D(\delta, \beta) = \frac{(\delta_L - \delta_H)^2}{4(\beta_l + \beta_h)} \geq 0, \quad (11)$$

Note that the welfare gains from differentiation are equivalent for taxes and permits under complete information. We can now make the following two observations based on equation (11):

(R2) The extent to which policy differentiation reduces pollution damages (via a reallocation of permitted emissions) and increases net welfare gains is increasing with the variation in damages across sources.

If damages do not vary across sources, there is no advantage to differentiated policy. Accordingly, the more heterogeneous the damages, the greater the benefits from differentiation, all else equal.

(R3) The extent to which policy differentiation reduces pollution damages (via a reallocation of permitted emissions) and increases net welfare gains is decreasing with the slope of the marginal abatement cost functions.

Intuitively, if marginal abatement costs are steeply increasing in abatement, it will be relatively more costly to shift emissions from the high damage to the low damage source.

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

2.3 Uncertainty

Thus far, we have analyzed market-based regulation of non-uniformly mixed pollution when source-specific marginal damage parameters and source-specific abatement costs are known with certainty.

In fact, real-world policy contexts are characterized by considerable uncertainty regarding both costs and damages from pollution. We begin this section by focusing exclusively on uncertainty surrounding the source-specific damage parameter estimates. We then introduce uncertainty about abatement costs.

2.3.1 Design implications of damage uncertainty

As noted above, recent advances in data collection and computer supported analytics have enhanced our ability to measure and analyze emission fate and transport, exposure pathways, and dose-response relationships (Esty, 2004). State-of-the-art integrated assessment modeling can now provide policy makers with rich information regarding source specific damages. In principle, this facilitates a more informed and data-driven approach to policy design and implementation. However, estimates of source-specific marginal damages are still highly uncertain. Much of the variation in damage estimates is driven either by the stochastic nature of the inputs (e.g. meteorological conditions) or parameter uncertainty that is unlikely to be resolved over the duration of the program (e.g. uncertainty regarding dose response or mortality valuation parameters).

We extend the model to more accurately represent the information that can, in principle, be used to inform policy design. We assume that the joint distribution of the marginal damage parameters $f(\delta_L, \delta_H)$ is known ex ante. This is a reasonable assumption given the kind of modeling tools that implementing agencies currently have at their disposal. We continue to assume that a risk neutral policy maker seeks to minimize the expected social costs of pollution. We assume that policy design parameters are chosen at the outset of the program and are fixed for the duration.

For now, we maintain our assumption that the regulator has complete information regarding abatement costs. While this is clearly a restrictive assumption, we invoke it in order to derive an expression that isolates the affect of damage uncertainty, and damage uncertainty alone, on the gains from policy differentiation.

Maintaining the simple two-firm set-up as above, the policy maker seeks to minimize expected total social costs from emissions:

$$\min_{e_H, e_L} TSC = C_H(e_H) + C_L(e_L) + \int \int (\delta_H e_H + \delta_L e_L) f(\delta_L, \delta_H) \quad (12)$$

The first order condition for cost minimization with respect to $\{e_L, e_H\}$ yields:

$$\begin{aligned} -C'_H(e_H) &= E[\delta_h] \\ -C'_l(e_l) &= E[\delta_l] \end{aligned} \quad (13)$$

In the presence of uncertainty about marginal damage parameters, setting marginal abatement

costs equal to expected marginal damages maximizes expected welfare.

We note some important qualifications. First, our assumption regarding the linear form of the damage function is important here. In the literature that examines the implications of uncertain damages on optimal policy, researchers have argued that the optimal trading ratio between two sources with equal expected damages but varying degrees of uncertainty should penalize the more uncertain damages (e.g. Horan and Shortle, 2005). In our case, linearity in damages eliminates the covariance term that gives rise to this penalty.

Our assumed policy objective function is also important. We assume that the regulator seeks to minimize the expected social costs of pollution. If instead the regulator wants to meet an ambient target probabilistically, varying degrees of uncertainty will matter because otherwise identical firms will have differential marginal effects on the probability the target is violated.

Conditional on these assumptions, we can derive the optimal policy parameters under damage uncertainty. The differentiated tax (8a) and undifferentiated tax (8b) should be defined in terms of expected marginal damage parameters $E[\delta_L]$ and $E[\delta_H]$. Similarly, compliance ratios r_i should be defined in terms of expected marginal damages.

2.3.2 The gains from differentiation under damage uncertainty

Uncertainty regarding source specific damage parameters implies uncertain gains from policy differentiation. To illustrate how this uncertainty manifests, it is instructive to consider the realized gains from differentiation conditional on a particular realization of $\{\delta_L, \delta_H\}$.

Let $\delta' = \{\delta'_H, \delta'_L\}$ denote a draw from the joint distribution of marginal damages $f(\delta_L, \delta_H)$. Given this realization of damages, the difference in total social costs under the differentiated design (TSC^D), relative to an undifferentiated policy (TSC^U) is derived in Appendix 2.1. This expression reduces to:¹⁶

$$TSC^U(\delta'; \delta) - TSC^D(\delta'; \delta) = \frac{1}{2} \frac{(\delta'_L - \delta'_H)(E[\delta_L] - E[\delta_H])}{(\beta_L + \beta_H)} - \frac{1}{4} \frac{(E[\delta_L] - E[\delta_H])^2}{(\beta_L + \beta_H)}. \quad (14)$$

From this expression we see that the net gains from policy differentiation reduce to (11) if source-specific damages manifest as expected. Realized net benefits from policy differentiation will be greater than expected if the realized difference in marginal damages across sources is greater than expected. Finally, if the ranking of sources with respect to damages is the reverse of what was expected, the first argument in (14) turns negative. In other words, policy differentiation can actually reduce welfare when marginal damage parameters are uncertain.

¹⁶We maintain our assumption of certain abatement costs here; the gains are identical for quantity-based and price-based instruments.

2.3.3 Cost uncertainty

As compared to uncertainty about damages, cost uncertainty manifests somewhat differently in the model. Integrated assessment modeling provides policy makers with a fairly rich characterization of the uncertainty surrounding damage estimates. In contrast, the methodologies used to estimate the engineering costs of emissions control options (i.e. the capital equipment expense, the site preparation costs for the application, and annual operating and maintenance costs) typically yield a single point estimate. In this section, we investigate how gains from differentiation are impacted when ex post realized costs deviate from the point estimates that informed policy design and implementation.

It is beyond the scope of this paper to present a fully general analysis of how abatement cost uncertainty affects realized gains from policy differentiation. Instead, we present a fairly stylized treatment of the problem in order to develop the intuition that will be essential to the interpretation of our policy simulation results. For expositional clarity, we shut down any uncertainty regarding damages in order to focus exclusively on costs. We confine our attention to uncertainty regarding the slopes of the firm-specific marginal cost functions as this will prove to be particularly relevant in our setting.

Beginning with the emissions tax regimes, we contrast outcomes under the differentiated and undifferentiated tax. We assume that the policy maker has point estimates of the cost parameters β_{H0} and β_{L0} which she uses to define these tax rates. Let $\beta_{L1} = \beta_{L0} + \Delta_L$ and $\beta_{H1} = \beta_{H0} + \Delta_H$ denote the ex post realized cost parameters. Given the realization of costs (β_1), we compare total social costs incurred under the first-best differentiated tax policy (TSC^D) against the second-best undifferentiated tax policy (TSC^U). Appendix 2.2 derives the following expression:

$$TSC^U(\beta_1; \beta_0) - TSC^D(\beta_1; \beta_0) = \frac{(\delta_L - \delta_H)^2 (\beta_{H1}\beta_{L0}^2 + \beta_{L1}\beta_{H0}^2)}{4\beta_{L0}\beta_{H0}\beta_{L0} + \beta_{H0}^2} \geq 0. \quad (15)$$

While expression (15) is difficult to parse, it is clearly non-negative. If realized costs equal expected costs, the gains from policy differentiation reduce to (11). Moreover, we show in Appendix 2.2 that the partial derivative of (15) with respect to Δ_L or Δ_H is negative.

(R5): For the case of emission taxes, the gains from differentiation are non-negative in the presence of cost uncertainty and certain damages.

This result makes intuitive sense given (R1). Recall that accurate cost information is needed to define the optimal undifferentiated tax rate but not the differentiated tax rates. A differentiated tax thus confers advantages associated via increased flexibility (i.e. the tax scales with the damage caused) and the information required ex ante to deliver ex post optimal results.

For the case of tradable permits, an analogous exercise yields an even more complex expression. One key implication is that, under a quantity-based policy regime, the welfare implications

of policy differentiation are ambiguous. In other words, policy differentiation can be welfare reducing. Under abatement cost uncertainty, the regulator is unlikely to be able to set the optimal aggregate cap, ex ante. If the imposed cap proves to be too tight or too loose given ex post realized abatement costs, the trading ratios (r_L, r_H) will no longer yield the first best allocation of emissions across sources.

To see why this is the case, consider the following. By Eq. (3), the optimal emission levels are: $e_L^D = \frac{\alpha_{1L} - \delta_L}{2\beta_L}$ and $e_H^D = \frac{\alpha_{1H} - \delta_H}{2\beta_H}$. If the cap has not been set optimally, the relative allocation of emissions across high and low damage sources that maximizes welfare *subject to the imposed cap* is no longer defined by (7). However, the differentiated emissions trading regime is hard-wired to allocate emissions across sources such that (7) is achieved. In extreme cases, it may be that setting marginal abatement costs equal across sources (the defining characteristic of the undifferentiated market) will welfare dominate the equilibrium characterized by Eq. (7).

3 The NOx Budget Program

The NOx Budget Program (NBP) is a market-based emissions trading program designed to reduce aggregate NOx emissions and the regional transport of NOx emissions in the eastern United States. Over the period 2003-2008, the program established a region-wide cap on emissions of NOx from large stationary sources in twenty eastern states during ozone season (May-September). The NBP was primarily designed to help Northeastern and Mid-Atlantic states attain Federal ambient ozone standards. When the NBP was promulgated, significant portions of the Northeast, Mid-Atlantic, and parts of the Midwest were failing to meet Federal ambient standards (Ozone Transport Assessment Group (OTAG), 1997).

Although the precise contribution of individual sources to the non-attainment problems in this region was difficult to estimate at the time of the rulemaking, there was plenty of evidence to suggest that marginal damages varied significantly across sources. The EPA received over 50 responses when 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 commentators supported unrestricted trading and expressed concerns that “discounts or other adjustments or restrictions would unnecessarily complicate the trading program, and therefore reduce its effectiveness” (FR 63(207): 57460). These comments and accompanying analysis (US EPA, 1998) led regulators to design a single jurisdiction, undifferentiated trading program. There are no spatial restrictions on trading within the program. All emissions are treated symmetrically for compliance purposes.

In 2008, a federal district court vacated the Clean Air Interstate Rule (CAIR) that was to replace the NOx Budget Program citing the policy’s failure to adequately accommodate regional

transport of pollution (among other factors).¹⁷ Since that time, debates surrounding the market-based regulation of non-uniformly mixed criteria pollutants have become increasingly heated. In the interest of informing policy designs going forward, we revisit the decision to forego a differentiated NBP design in favor of the simpler, undifferentiated alternative.

Our analysis will focus exclusively on the coal-fired generating units in the program. Although other industrial point sources are also included in the NBP, coal-fired electricity generating units represent approximately 94 percent of the NOx emissions regulated under the program and more than 94 percent of the NOx emissions reductions over the first five years (U.S. EPA, 2005; US EPA 2008). Data limitations require that we exempt these small emitters from our analysis. Because these sources account for such a small share of emissions and emissions abatement, this omissions should have very negligible impacts on our results.¹⁸

4 Estimating the welfare impacts of policy differentiation

Figure 2 provides a diagrammatic summary of the policy simulations we conduct. Taken together, we conduct eight sets of counterfactual policy simulations. These can be classified in terms of regulatory regime (prices, quantities, no regulation); our approach to modeling firms' compliance choices (cost minimization algorithm or econometric model); and whether policy incentives are differentiated or undifferentiated.

The analysis proceeds in several steps:

1. Estimate the marginal damage parameters δ for each source in the NOx Budget Program.
2. Construct engineering estimates of source-specific, technology-specific abatement costs and emissions reduction efficiencies. These estimates are constructed using information about costs and technology operating characteristics that was available at the time of the rule making.
3. Use the information obtained in steps (1) and (2) to define the ex ante optimal emissions cap (in the quantity-based policy regimes) and the ex ante optimal emissions tax (in the price-based regimes).

¹⁷The 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.

¹⁸Natural gas and oil-fueled electricity generators subject to the NOx Budget Program are among the units we omit from our analysis. These units account for less than 5 percent of emissions regulated under the program. The policy design changes we consider could marginally affect how these units are dispatched. But these marginal changes would presumably have very small impacts on overall operating costs and aggregate damages. The reason is that these units 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 lbs NOx/MWh (NEISO, 2006; Keith et al., 2003).

4. Simulate firms' compliance choices under both the observed and counterfactual policy regimes.
5. Estimate the total damages associated with the simulated NO_x emissions under observed and counterfactual policy regimes.
6. Estimate the total abatement costs associated with simulated compliance choices.
7. Compare the net benefits across differentiated and undifferentiated policy designs to obtain an estimate of the gains from policy differentiation.

The following subsections describe each step in more detail.

4.1 Estimating source-specific damages from pollution

NO_x emissions affect health and environmental outcomes through two main pathways: ozone formation and particulate matter formation.¹⁹ 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 O₃ through a series of chemical reactions (Seinfeld, Pandis, 1998). Prior research has shown that the majority of damages due to exposures to both PM_{2.5} and O₃ are premature mortalities and increased rates of illness (US EPA, 1999; WHO, 2003; Muller and Mendelsohn, 2009).

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 and 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 age and spatial distribution of the human populations at that location. For these reasons, the damage caused by a given quantity of NO_x emissions will depend significantly on the spatial distribution of the emissions.

In what follows, we characterize both variability and uncertainty in NO_x emissions damages in detail. With respect to the former, we estimate the extent of the variation in marginal damage estimates across sources in the NO_x Budget Program. With respect to the latter, it is important to emphasize that our uncertainty analysis is not comprehensive. We formally quantify the parameter uncertainty inherent in source-specific damage estimates. But we make no attempt to capture modeling uncertainty.

¹⁹NO_x emissions also contribute to acid rain in some mountain regions, and exacerbate eutrophication problems. Neither of these effects are modeled in this analysis.

4.1.1 Source-specific damage parameters

The source-specific marginal damage parameters (δ_i) capture the estimated effect of an incremental change in NOx emissions at source i on health and environmental impacts across the airshed. We use a stochastic integrated assessment model, AP2, to estimate these damage parameters. The AP2 model is comprised of six modules; emissions, air quality modeling, concentrations, exposures, physical effects, and monetary damages. The interested reader can find a detailed description of the model in (Muller, 2011).

The air quality model uses a reduced form approach to mapping emissions (e) into pollution concentrations (C). More precisely, the relationship between emissions of nitrogen oxides released at source i and the concentration of pollutant s (ozone or particulate matter) at receptor point r is modeled through a source-receptor matrix in which each cell in the matrix, $C_{sri}(e_i)$, denotes the incremental contribution of emissions from (i) to concentrations in location (r).²⁰ AP2 then maps ambient concentrations into physical impacts using dose-response functions. Let ϕ_{kp}^s represent the dose response coefficient which captures the effect of an incremental change in the concentrations of pollutant s on health outcome k in population cohort p .²¹

The final modeling step in AP2 translates the physical effects predicted by the dose-response functions into monetary terms. Let ν_k represent the valuation coefficient that is used to translate the health outcome k into dollar terms. We rely on valuation methodologies used in the prior literature. In particular, we employ a VSL of approximately \$6 million.²² It is important to note that our monetized damage estimates do not incorporate defensive investments or costly actions that individuals take to protect themselves from harmful effects of pollution. Deschenes et al. (2012) demonstrate that a failure to account for changes in these defensive expenditures can result in a significant undervaluation of the benefits from NOx emissions reductions.

The marginal damage parameters for the 632 coal-fired EGUs regulated by the NBP are estimated using the marginal damage algorithm used in Muller (2011). First, baseline emissions data collected by the US EPA in the years immediately preceding the introduction of the NBP are used by AP2 to compute total national damages. Next, one ton of NO_x is added to baseline emissions at a particular EGU. AP2 is then re-run. Concentrations, exposures, physical effects, and damages are recomputed. Since the only difference between the baseline run and the "add-

²⁰The source-receptor matrix is based on the Gaussian Plume model (Turner, 1994). The predicted pollutant concentrations generated using the AP2 model have been tested against the predictions made by a more sophisticated and detailed air quality model (see the appendix in Muller, 2011). The agreement between the county-level surfaces produced by the two models is quite strong.

²¹In 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 O₃, we use the findings from Bell et al., (2004). In addition, this analysis includes the impact of exposure to PM_{2.5} on incidence rates of chronic bronchitis (Abbey et al., 1995).

²²This value, which is used by US EPA, results from a meta-analysis of nearly 30 studies that compute VSLs using both stated and revealed preference methods.

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.

4.1.2 Damage parameter uncertainty

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

$$\delta_i = \sum_{ri} \sum_k \sum_s \nu_k \phi_{kp}^s P_{ri} \frac{dC_{si}(e_i)}{de_i}. \quad (16)$$

The parameters of the air quality model, the population estimates P_{ri} , the dose response parameters ϕ_{kp}^s and valuation parameters ν_k are all uncertain. Even the emissions levels at individual sources cannot be predicted with absolute certainty.

We construct an empirical distribution for each δ_i parameter. We first make a random draw (denoted m) from the distributions of each of the parameters in (16). Next, we use the AP2 model to compute emissions, concentrations, exposures, physical effects, and damages based on the draw m from each input distribution. We then add one ton of NO_x to source i and recompute concentrations, exposures, physical effects, and damages (using the same m^{th} realization from the input distributions). AP2 computes the difference between damages with baseline emissions and after adding the ton of NO_x to emissions at source (i). This is denoted δ_{im} . This process is repeated 5000 times using 5000 different, independent draws from the parameter distributions. The result is an empirical distribution of marginal damages for NO_x emitted from facility i . This process is repeated for each EGU in the analysis (using the same set of m independent draws).

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

Figure 4 illustrates the extent to which the expected values of source specific damage parameters $E[\delta_i]$ vary across sources. The average parameter value (averaged across all sources) is \$1734/ton of NO_x . In the subsequent discussion, we classify any source with estimated damages exceeding (falling below) \$1734/ton NO_x as "high" ("low") damage. Notably, a significant amount of the inter-source variation (approximately 45 percent) occurs within (versus between)

states. Previous analyses of spatially differentiated NOx trading considered multi-state zonal approaches to policy differentiation (Krupnick et al., 2000; USEPA, 1998)²³. Our source-specific damage estimates suggests that differentiating policy incentives at the state or region-level would ignore a significant portion of the spatial variation in damages.

These source-specific damage parameters are used to define the tax rates and compliance ratios in the counterfactual policy regimes we consider. For six of the 632 units in our data, we find that the expected value of the marginal damage parameter δ is negative. This suggests that a decrease in NOx emissions at these sources leads to increased overall damages. This result is driven by the complex, non-linear photochemical reactions that transform NO_x and VOCs into ozone.²⁴ We assume that incentivizing pollution at facilities with negative damage parameter estimates would be politically unpopular. Instead, we exempt any units with negative expected damage parameters.

4.2 Engineering estimates of NOx abatement costs

The NBP mandated a dramatic reduction in average NOx emissions rates.²⁵ 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 NOx control equipment, or reduce production at dirtier plants during ozone season.²⁶ For the coal-fired units in our analysis, we rule out reduction in ozone season output as a compliance strategy and assume that firm-level production and aggregate output are exogenously determined and independent of the environmental compliance choice. Coal-fired units are typically inframarginal due to their relatively low fuel operating costs. Consistent with this observation, Fowlie (2010) finds that the introduction of the NBP reduced profit margins at operating units, but not production levels.

We take as given the population of sources in the NOx Budget Program. That is, we rule out the possibility that one or more of the counterfactual policy designs we consider would cause an electricity generating unit to exit the market prematurely.²⁷ The coal plants in our analysis

²³For example, policy makers considered dividing the regulated region into two or three subregions in an effort to make a distinction among the States that may contribute the most to the ozone transport problem and those where the wind patterns may be less likely to affect air quality in the other states.

²⁴Daily ozone concentrations are non-linear and monotonic functions of NOx and the ratio of volatile organic compounds (VOCs) and NOx. At sufficiently low ratios, the conversion of NOx to ozone is limited by the availability of VOCs. In these VOC limited conditions, reductions of NOx can increase peak ozone levels until the system transitions out of a VOC-limited state (Seinfeld and Pandis, 1998).

²⁵Pre-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).

²⁶We assume perfect compliance on behalf of all units. In fact, compliance has been close to 100 percent for the duration of the program (US EPA, 2008).

²⁷Note that, in addition to treating the retirement decision as exogenous, we are not attributing any costs to NOx

are long-lived. The average retirement age of a coal plant in the United States is 49 years. We do observe a small number of coal-fired boilers retiring during the study period. These are units with decades of service stretching as far back as the end of World War II. We assume that these retirement decisions are unaffected by the policy design.

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

In the late 1990s, to help generators prepare to comply with market-based NOx regulations, the Electric Power Research Institute²⁸ developed software to identify all major NOx control options (including combinations of control technologies) available to coal-fired boilers, conditional on unit and plant level characteristics. The software has been used not only by plant managers, but also by regulators to evaluate proposed compliance costs for the utilities they regulate. This software was used to generate the boiler-specific cost estimates used in this analysis (EPRI, 1999). This cost estimation exercise is described in detail in Fowle (2010). Importantly, the boiler-level and plant-level parameters (including boiler technology type, plant vintage, plant capacity) that generate variation in these cost and emissions estimates are plausibly exogenous to the compliance choices we are interested in modeling.

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. Overall, these unit-level characteristics are similarly distributed similarly across these high and low damage groups.

4.3 Simulating facility-level compliance decisions

We use two alternative approaches to modeling firm-level compliance decisions when simulating policy outcomes. We first specify a simple cost minimization model of firms' compliance choices. We calibrate key parameters to match standard policy simulation models. This model is used to approximate the ex ante expectations of a well informed policy maker.

reductions from the new plants replacing these units once they retire. This is equivalent to assuming that operating characteristics of newly constructed units will be independent of the policy design choices we are analyzing. This assumption is reasonable given the stringency of new source standards.

²⁸The 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.

Previous studies have noted that cost minimization algorithms provide a crude and possibly inaccurate approximation of real-world decision-making (Fowlie, 2010; Krupnick et al. 2000). If firm’s environmental compliance choices deviate significantly from cost-minimization algorithms, standard policy simulation models will inaccurately predict the gains from differentiation. Our second approach replaces the cost minimization algorithm with an econometrically estimated model of the compliance choice. These represent our best estimate of the outcomes that would actually materialize.

In both cases, a firm’s compliance choice is modeled as a static, one-shot decision regarding what - if any- abatement equipment to install. This static approach is appropriate because all firms had to make a choice about how to comply with the program during a fairly short period of time. The details of the program were finalized in 2000. The aggregate emissions cap was imposed (and full compliance was required) in 2004.

4.3.1 Cost-minimization algorithm

We use a simple algorithm to find the combination of NOx control options that minimizes the levelized annual cost of complying with the policy designs we consider. Let $j = 1 \dots J_i$ index the NOx control technology options available to the i^{th} electricity generating unit. Let K_{ij} represent the engineering cost estimate of the required capital investment specific to unit i and technology j . Let V_{ij} represent the corresponding variable operating cost estimate (per kWh). Let m_{ij} represent the corresponding post-retrofit emissions rate. Let e_{i0} represent the pre-retrofit emissions rate (i.e. the amount of NOx the i^{th} unit emits per kWh of electricity generated if it installs no new pollution controls).

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

$$\begin{aligned} \min_j \quad & C_{ij} = v_{ij}Q_i + l_iK_{ij} & (17) \\ \text{where } v_{ij} \quad & = (V_{ij} + \tau m_{ij})Q_i. \end{aligned}$$

Capital investments K_{ij} are converted to annual costs using a levelized annual cost factor l_i . We assume a discount rate of 5.5 percent. This rate, which was derived from financial data for electric utilities, is used in the modeling of investments in environmental retrofits conducted by the US EPA.²⁹ Unit-specific investment time horizons are constructed by subtracting the unit age from the assumed life span (55 years).³⁰

²⁹For more complete documentation of this model and its assumptions, see <http://www.epa.gov/airmarkets/progsregs/epa-ipm/docs/v410/Chapter8.pdf>

³⁰The median age of coal plants operating in the United States is 44. We assume a retirement age of 55. If units are older than 55 years, we assume these units will not invest in pollution abatement equipment. However, we also adopt standard EPA assumptions regarding the lifespan of the NOx control technology (30 years).

Expected annual operating costs v_{ij} are obtained by multiplying estimated per kWh operating costs by expected seasonal production Q_i . Historic electricity production during the ozone season, Q_i , is used to proxy for expected ozone season production.³¹ To estimate the annual variable compliance cost v_{ij} , the technology operating costs are added to the expected (per kWh) costs of holding permits to offset uncontrolled emissions.

We use the model to simulate firms' compliance choices under the observed and counterfactual policy regimes. Given an emissions cap and an initial guess for the equilibrium permit price τ^0 , we find the compliance option j^0 at each electricity generating unit that minimizes compliance costs. The ozone season NOx emissions associated with these choices are summed across units. If these aggregate emissions exceed (fall below) the cap, the permit price is incrementally increased (reduced) and the process is repeated until the aggregate emissions constraint is just satisfied.³²

When simulating outcomes under the observed policy regime, the cap is set equal to the seasonal NOx emissions associated with observed compliance choices (658,609 short tons).³³ It is straightforward to modify the model to represent the counterfactual policy regimes we analyze. In the differentiated trading regime that incorporates compliance ratios, variable compliance costs are redefined as $(V_{ij} + \tau r_i m_{ij})Q_i$, where $r_i = E[\delta_i] / (\frac{1}{N} \sum E[\delta_i])$. In an undifferentiated tax regime, τ equals the uniform tax rate per unit of emissions. To simulate outcomes under the differentiated tax, the τ parameter (introduced in section 2.1.2) is replaced with the source-specific expected damage parameter δ_i .

The top panel in Table 2 summarizes the compliance choices that we observe in the NOx Budget Program. A majority of units chose to rely on the permit market exclusively for compliance (the "no retrofit" option). A majority of the mandated emissions reductions were achieved using highly capital intensive selective catalytic reduction (SCR) technologies. The middle panel summarizes the compliance choices simulated using the cost minimization algorithm. This cost minimization model poorly predicts the compliance choices that firms actually made; only 24% of choices are correctly predicted. In particular, the model overestimates the share of less capital intensive combustion modifications and underestimates the share of capital intensive SCR retrofits.

³¹Anecdotal evidence suggests that managers used past summer production levels to estimate future production (EPRI, 1999). 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.

³²If 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.

³³The estimated emissions associated with observed compliance choices exceed the emissions levels that were actually observed. In 2004, the first year of full compliance, NOx emissions from coal units were 605,000 short tons. Emissions from coal plants dropped even lower in later years (US EPA, 2007). One possible explanation for the discrepancy between our emissions estimates (based on engineering estimates of the emissions associated with alternative compliance options) and observed emissions is that many units that made no capital investment in abatement equipment were able to make extensive small-scale improvements to reduce emissions intensity. Linn (2008) estimates that 10-15 percent of emissions reductions were the product of these small process changes and modifications.

In sum, we find that compliance choices observed in the NOx Budget Program depart markedly from those predicted by the cost minimizing algorithm. If the cost minimization-based policy simulations incorrectly predict how firms respond to the observed (undifferentiated) policy, these models should not be relied upon to accurately simulate firms' response to differentiated policy incentives. For this reason, we pursue a second approach to modeling firms' compliance decisions which is designed to more accurately capture the real world distortions and idiosyncrasies that determine firms' environmental compliance choices.

4.3.2 An econometric model of the compliance decision

Fowle (2010) estimates an econometric model of the compliance choices made by plant managers in the NBP. We use this model to simulate the compliance decisions that plant managers under both the observed and counterfactual policy regimes. We argue that this choice model can be used to generate more realistic estimates of firms' response to observed and counterfactual policy incentives (as compared to simulations based on the calibrated cost minimization model).

The structure of this choice model is very similar to (17). The primary difference is that parameters are econometrically estimated versus calibrated to match standard policy simulation models. The decision maker at unit i is assumed to choose the compliance strategy that minimizes the unobserved latent value C_{ij} :

$$C_{ij} = \alpha_j + \beta_m^v v_{ij} + \beta_m^K K_{ij} + \beta^{KA} K_{ij} \cdot Age_{ij} + \varepsilon_{ij}, \quad (18)$$

where $v_{ij} = (V_{ij} + \tau m_{ij})Q_i$

The deterministic component of C_{ij} is a weighted sum of expected annual compliance costs v_{ij} , the expected capital costs K_{ij} associated with initial retrofit and technology installation, and a constant term α_j that varies across technology types. The technology fixed effects are intended to capture average biases for or against particular types of NOx control equipment. An interaction term between capital costs and demeaned plant age is included in the model because older plants can be expected to weigh capital costs more heavily as they have less time to recover these costs. 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.³⁴

Expected annual compliance costs are obtained by multiplying estimated per kWh variable costs by expected seasonal production Q_i . We use the average permit price that prevailed during the period prior to the NBP compliance deadline (\$4500/ton NOx) as a proxy for what managers'

³⁴It 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.

expected cost of offsetting uncontrolled emissions. The econometric model is estimated separately for units serving restructured wholesale electricity markets versus publicly owned units and units subject to cost-of-service regulation. A more detailed description of the econometric specification and estimation results can be found in Fowlie (2010).

Estimates of the parameters of the distribution of β^v and β^K in the population can be combined with information about observed choices in order to make inferences about where in the population distribution a particular decision maker lies (Train, 2003). We use the means of these plant-manager specific distributions to parameterize our policy simulation model. This should improve our ability to simulate the choices that these plant managers would most likely have made in counterfactual policy scenarios.

Table 3 summarizes the choice model parameter estimates. The top panel reports the estimated technology specific fixed effects. These are all negative, suggesting that the average plant manager was biased against emissions abatement technology retrofits (vis a vis the compliance option that relies exclusively on purchasing permits). The bottom panel reports the means of the manager-specific distributions of the two cost coefficients (β^K and β^v). The ratio $(\beta^K + \beta^{KA} Age) : \beta^v$ can be interpreted as a measure of how a plant manager trades off fixed capital costs (i.e. investments in NOx control equipment) and variable compliance costs (including the cost of holding permits to offset uncontrolled emissions each year). Point estimates of this ratio (computed using the estimated means of the manager specific conditional distributions) are 0.48 and 0.21 among managers of deregulated and regulated units, respectively. As compared to the cost minimization model, these econometric estimates imply that plant managers were more strongly biased against more capital intensive compliance options.³⁵

The bottom panel of Table 2 summarizes the equilibrium choice probabilities under the observed (undifferentiated) policy regime. The econometric model correctly predicts that a significant portion of the mandated emissions reductions is achieved using more capital intensive compliance strategies.³⁶

With this econometrically estimated compliance choice model in hand, our approach to simulating permit market outcomes is mechanically very similar to simulation exercise described in section 4.3.1.³⁷ A key assumption is that the fundamental structure of the firm-level compliance decisions we model would be invariant to the policy regime changes we consider. This seems very plausible. We see no reason why managers willingness to trade off annual operating costs and up-

³⁵The high discount rate among plants serving restructured electricity markets likely reflects significant credit rating downgrades which affected several firms during the time period in which plant managers were having to make their compliance decision.

³⁶One possible explanation for the apparent over-investment in more capital intensive compliance options could be the regulatory incentives faced by public power authorities and plants operating under rate-of-return regulation (Fowlie, 2010).

³⁷One difference is that the simulated emissions are defined probabilistically as $e_i = \sum_j P_{ij} e_{ij}$, where P_{ij} is the predicted probability of unit i choosing compliance strategy j .

front capital investment, and/or managers' preferences for or against particular pollution control technologies, should be impacted by policy differentiation.

4.4 Welfare maximizing policy parameters

We are interested in evaluating how alternative approaches to designing market-based policy, namely the differentiated tax and trading programs described in Section 2.1, perform vis a vis more standard, undifferentiated designs. When parameterizing the policy simulations, we approach the policy design problem from the perspective of a well informed, welfare maximizing regulator. We assume the regulator has access to source-specific damage estimates summarized in section 4.1.1, the engineering estimates of abatement costs summarized in section 4.2, and a standard policy simulation model such as that described in section 4.3.1. Using this ex ante available information, we identify the policy parameters (i.e. the emissions taxes or emissions quantity limits) that maximize ex ante expected welfare.³⁸

In the theoretical model that featured only two firms with linear marginal abatement cost curves, it was straightforward to solve for the optimal policy parameters analytically. In the applied analysis that involves hundreds of sources with non-linear abatement cost functions, we solve for the optimal taxes and emissions caps using simulation. In other words, we simulate outcomes under a range of policy parameter values in order to identify the welfare maximizing value. Integrating the fully stochastic integrated assessment model into this policy optimization exercise would be highly resource intensive. When simulating outcomes for the purpose of identifying the ex ante optimal tax or emissions cap, we assume that the policy maker uses the same linear, additively separable approximation that was assumed in section 2 to map emissions into monetized damages.

4.4.1 Ex ante optimal tax rates

By Eq. (8b), the optimal undifferentiated (or uniform) emissions tax τ is equal to the average of the source-specific damage parameters weighted by the slopes of the corresponding marginal abatement cost curves. Let $e_i^*(\tau)$ denote the emissions-level at source i that minimizes compliance costs given the tax rate $\tau_i = \tau$. The uniform tax rate that minimizes $\sum_{i=1}^N D_i(e_i^*(\tau)) + C_i(e_i^*(\tau))$ is \$1630/ton NOx.³⁹ The aggregate, ozone season emissions that correspond with this undifferentiated tax amount to 726,293 tons.

³⁸Integrating the fully stochastic integrated assessment model into this policy optimization exercise would be highly resource intensive. We therefore assume that the policy maker uses the same linear, additively separable approximation that was assumed in section 2 to map emissions into monetized damages.

³⁹Note that this value falls below the average marginal damage value of \$1734/ton. When we approximate the slope of the source-specific marginal abatement cost curves in the neighborhood of the tax, we find that these curves are less steeply sloped on average (or emissions are locally less responsive to changes in the tax rate) among relatively low damage (versus high damage) sources.

The optimal differentiated tax is more straightforward to define (see Eq. (8a)). Compliance cost minimizing compliance choices under this differentiated tax regime yield ozone-season emissions of 746,644 tons of NO_x. Thus, the emissions equivalence result derived in section 2.1 does not hold exactly. Strict emissions equivalence requires linear marginal abatement cost curves. The marginal abatement cost curves in this optimization routine are non-linear step functions.

4.4.2 Ex ante optimal emissions cap

To define the welfare maximizing emissions cap in the undifferentiated trading regime, we simulate cost minimizing compliance behavior in a trading regime characterized by ton-for-ton emissions trading. We compute the total social cost (i.e. abatement costs plus pollution damages) for a range of emissions cap values. Intuitively, we find that the welfare maximizing emissions cap is equal to the level of emissions that corresponds to the optimal undifferentiated tax: 726,293 tons. The corresponding permit price is \$1630/ton NO_x. Note that this emissions cap is close to, but somewhat less stringent than, the cap imposed under the NBP.

Turning to differentiated emissions trading, we confine our attention to the differentiated trading design introduced in section 2. Each source is required to hold $r_i = \delta_i/\bar{\delta}$ permits per unit of uncontrolled emissions. The regulator chooses the total number of tradable permits to distribute. The total quantity of emissions associated with a given permit allocation will depend on which sources hold the permits. We simulate permit market clearing for a range of permit supply quantities. Total social costs are minimized when the aggregate permit supply is constrained to equal 642,158 tons. Intuitively, the simulated emissions that correspond to this permit allocation sum to 746,644 tons of NO_x (i.e. the quantity of emissions that corresponds to the differentiated tax regime).

4.5 Pollution damages

Once the compliance choices associated with a particular policy design have been simulated, the corresponding vector of emissions is processed through the stochastic AP2 model. At this stage of the analysis, rather than systematically perturbing NO_x emissions one source at-a-time, NO_x emissions change simultaneously at many of the regulated EGUs (relative to the baseline case). These pollution damage simulations feature all input parameters (emissions, transfer coefficients in the stochastic air quality model, population, dose-response, and valuation) as random variables. As we explain above, we simulate damages using 5,000 possible realizations of the model parameters. Importantly, when comparing damages across policy scenarios, resultant damages are compared for the same draws from the input distribution.

4.6 Estimating the costs of compliance

Information on the actual expenditures of plants subject to the regulation is not publicly available. Consistent with our approach to simulating policy outcomes, we take two different approaches to estimating these abatement costs. The first assumes that ex ante engineering estimates are a reasonable proxy for the costs that would actually materialize. The second approach finds the abatement costs that are most consistent with observed choices.

4.6.1 Engineering estimates of compliance costs

The most direct approach to estimating abatement costs employs the same cost assumptions and accounting parameters that were used to parameterize equation (17). The levelized annual cost of compliance under policy regime r is defined to be:

$$LAC_r^{CM} = \sum_i V_{ir}Q_i + l_iK_{ir}, \quad (19)$$

where V_{ir} and K_{ir} are the boiler-specific, technology-specific cost estimates associated with the compliance option chosen by firm i under policy regime r .

When the econometric model is used to simulate compliance choice probabilities, this cost estimate is defined to be:

$$LAC_r^{EST} = \sum_i \sum_j P_{irj}(V_{ir}Q_i + l_iK_{ir}),$$

where P_{irj} denotes the simulated choice probability associated with unit i and choice j in regime r .

This approach provides an estimate of ex ante expected abatement costs. However, these costs are inconsistent with observed compliance choices (and thus the choices simulated using the econometrically estimated choice model).

4.6.2 Compliance costs consistent with the compliance choice model

The second approach derives the abatement costs that are most consistent with simulated compliance choices. Note that both the cost minimization algorithm and the econometrically estimated model of the firm's compliance choice can be used to simulate the quantity of NOx a source would emit at a given price per ton of NOx. Repeating this exercise for a range of NOx prices (starting at \$0/ton) traces out a marginal abatement cost (MAC) curve. Integrating under this curve yields the estimate of the levelized annual abatement cost that rationalizes the simulated compliance choices.

Figure 5 provides a graphical illustration. The horizontal axis measures aggregate emissions abatement (in millions of tons of NOx per ozone season). The vertical axis measures marginal abatement costs (in \$/ton). The vertical line corresponds to the emissions cap imposed under the NBP. The lower MAC curve is generated using the model that assumes strict cost minimization and undifferentiated permit trading. Integrating under this lower marginal abatement cost curve up to the level of abatement required to satisfy the NBP cap yields a levelized annual cost estimate that is exactly equal to the sum of (19) across simulated compliance choices. The more steeply sloped MAC curve in figure 5 is generated using the econometric model of the compliance choice (under undifferentiated trading). Integrating under this marginal abatement cost curve obtains an estimate of the aggregate cost of compliance that is most consistent with the compliance costs as perceived by the regulated firms.⁴⁰

The difference in the two MAC curves depicted in Figure 5 is striking. If the difference between these two numbers is attributable to distortions and idiosyncrasies that drive a wedge between private and social costs (such as status quo bias or price discrimination on the part of pollution control equipment manufactures), then the lower curve provides a more appropriate cost measure for our purposes. However, a more likely explanation is that the engineering cost estimates fail to capture all the costs of financing and implementing investments in pollution abatement equipment. This would imply that the cost estimates that rationalize the econometrically estimated choice model should serve as our preferred abatement cost estimate.

5 Results and synthesis

This section is divided into three sub-sections. The first summarizes outcomes under the NOx Budget Program as implemented. The second investigates the counterfactual policy scenarios in the interest of estimating net gains from policy-differentiation. The third sub-section examines the implications of damage and cost uncertainty.

Recall that we generate two sets of policy simulations for each policy scenario we consider. Simulations based on the calibrated cost minimization algorithm are intended to approximate ex ante expectations of a well informed policy maker. Simulations which incorporate the econometrically estimated choice model represent our best estimate of the outcomes that would actually materialize under a given set of policy incentives. The key difference between these two approaches is that the latter incorporates information about abatement costs that is revealed by firms' ex post observed compliance choices. It important to note that ex post observed information does nothing

⁴⁰Note that, by construction, these two alternative approaches to estimating compliance costs yield equivalent results when the cost minimization model is used to simulate compliance choices. Intuitively, the very same cost estimates define both the compliance cost minimization routine and the ex post computation of costs. When the econometric model is used, these two approaches yields very different results.

to resolve damage uncertainty. Thus, information regarding damages does not vary across the two simulation approaches.

5.1 Welfare impacts of the undifferentiated NOx Budget Program

Table 4 summarizes the simulated outcomes under an undifferentiated policy regime that is calibrated to match the NOx Budget Program. In these simulations, the emissions cap imposed is set equal to the seasonal NOx emissions associated with observed compliance choices (658,609 tons). Ex ante expected outcomes are reported in column (1). Estimates which incorporate ex post revealed information regarding abatement costs and firms' response to policy incentives are reported in column (2).

Simulated equilibrium permit prices appear in the top row of this table. Because the marginal abatement cost curves that rationalize observed choices are steeper than the MAC curves generated using the cost minimization algorithm, the permit market clearing price in the second column is significantly higher.

The percentage of permits used to offset emissions at relatively high damage sources, reported in the third row of the table, provides a crude measure of how permitted emissions are distributed across sources. Facility-level data collected following the introduction of the NBP indicate that approximately 38 percent of permitted NOx emissions occurred at sources with higher than average damage parameters. The cost minimization model over-predicts the share of emissions occurring at these high damage sources (41 percent). The econometric model correctly allocates 38 percent of permitted emissions to relatively high damage firms.⁴¹

The second panel in Table 4 estimates the abatement costs incurred to comply with the NOx Budget program. The difference in the levelized annual cost estimates across the two columns highlights the extent to which observed compliance choices depart from the predictions of the cost minimization-based model. In column (1), our two approaches to cost estimation yield equivalent results by construction. In the second column, the cost estimates which are most consistent with the econometric choice model are somewhat higher than those constructed using engineering estimates.

The third panel of table 4 summarizes the simulated impacts of the policy on aggregate damages. Annual benefits (i.e. the monetized value of avoided damages vis a vis the baseline) are generated using the stochastic integrated assessment model. The table reports the average damage impacts (averaged across all 5000 realizations) and the 5th and 95th percentile estimates (in parentheses).⁴²

⁴¹Note that we are using the econometrically estimated model to predict outcomes that were used to estimate the model. This is intended to illustrate the descriptive, versus predictive, power of the model.

⁴²Strictly speaking, these should not be interpreted as confidence intervals. Recall that our analytical approach captures uncertainty about key parameters (such as dose-response parameters and the value of statistical life)

Subtracting our preferred abatement cost estimate from avoided damage values yields an estimate of the annual net benefits of the NOx Budget Program. Our estimate of ex ante expected net benefits is \$546 million per year (\$US 2000). This estimate falls within the range of net benefits that were projected in the design stages of the program, although comparisons with other studies are confounded by differences in modeling assumptions.⁴³ Our preferred, albeit conservative, point estimate of annual net benefits that were actually realized is \$264 million per year.⁴⁴

5.2 Policy counterfactuals

In this subsection, we investigate the likely welfare implications of policy differentiation. For expositional reasons, we begin with a discussion of price-based policy designs.

5.2.1 Differentiated and undifferentiated emissions taxes

Simulated outcomes under both the differentiated and undifferentiated tax regimes are summarized numerically in Table 5 and graphically in Figure 6. The left panel of Figure 6 summarizes outcomes at high damage sources (i.e. sources with marginal damage parameters that exceed the average value of \$1733/ton). The right panel summarizes outcomes aggregated across relatively low damage sources.

The downward sloping lines in Figure 6 connect the dots - literally - defined by simulated seasonal emissions (reported in the top panel of Table 5) and monetary policy incentives (the emissions tax measured in \$/ton NOx). The lower curve in each graph corresponds to ex ante expected abatement costs. The higher curves correspond to abatement costs as perceived by firms.

The horizontal lines correspond to differentiated and undifferentiated tax rates. The labeled intersection points define equilibrium outcomes under the differentiated (D) and undifferentiated (U) tax regimes among high (H) and low (L) damage units, respectively. The 0 and 1 subscripts denote the ex ante expected and ex post revealed costs, respectively.⁴⁵

and the stochastic nature of the meteorological factors that shape pollution formation, transport, and deposition. Conceptually, these numbers summarize the outcomes of a policy "lottery" in which probabilities can be attached to a range of possible emissions damage outcomes. But the set of outcomes represented in our analysis is not comprehensive. For example, we ignore modeling uncertainty.

⁴³The Regulatory Impact Assessment that informed the implementation of the program estimated annual net benefits ranging from -\$560 to \$2510 (\$US 1990) (USEPA, 1998). Burtraw et al (2003)(Burtraw, Bharvirkar, and McGuinness) projected annual net benefits of \$440 million (\$1997) per year.

⁴⁴It is also worth reiterating that our benefits estimates do not account for reduced pharmaceutical purchases and other defensive expenditures. In a recent study that emphasizes health-related impacts of the NBP, Deschenes et al. (2012) find that monetized benefits from avoided defensive expenditures are approximately equal to the monetized value of reduced mortality rates. Because our analysis fails to account for defensive health expenditures; our benefits estimate should be interpreted as a lower bound.

⁴⁵For example, the point labeled D_{H0} denotes aggregate emissions and the marginal cost of abatement among high damage sources (H) under the differentiated tax regime (D) when ex ante expected abatement costs are assumed (0). Simulated emissions among high damage sources amount to 209,000 tons per year (or 28 percent of 746,600 tons as reported in Table 6).

Both the table and the figure show that the cost minimization model predicts lower levels of emissions (or larger tax-induced emissions reductions) as compared to the simulations that incorporate the econometrically estimated choice model. This makes intuitive sense. The more steeply sloped the abatement cost curves, the lower the level of abatement activity induced by a given emissions tax.

To assess the gains from policy differentiation, we compare the net benefits of the policy (i.e. the avoided damages less abatement costs) across the undifferentiated and differentiated tax regimes. In Figure 6, these net gains are represented by area A (i.e. the increase in avoided damages less additional abatement costs among high damage sources) plus area B (reductions in abatement costs less the increase in damages among low damage sources). Taken together, these ex ante expected gains exceed \$60M per year (see Table 5).

The simulations that make use of ex post revealed information regarding compliance choices provide a more credible estimate of how firms would actually respond under these counterfactual tax regimes. These net gains are represented by areas C+D in Figure 6. Gains from policy differentiation are smaller than expected (\$47 M annually). R5 helps provide some intuition for this result. Because the ex post realized abatement cost curve is steeper than expected, gains from policy differentiation fall short of expectation.

Finally, note that the differentiated tax achieves the efficient outcome in expectation, regardless of whether ex ante expectations regarding abatement costs prove accurate. This is illustrated graphically; realized marginal abatement costs are set equal to source-specific expected marginal damages. In contrast, the policy maker requires accurate cost information in order to define the optimal undifferentiated tax ($R1$). Intuitively, the weights used to define this uniform tax may inaccurately capture the ex post realized relative abatement costs.

5.2.2 Differentiated and undifferentiated emissions trading

Table 6 summarizes the results from simulating outcomes under undifferentiated and differentiated emissions trading.⁴⁶ These results are summarized graphically by Figure 7. Note that the left two columns in Table 6 are identical to the corresponding columns in Table 5. Likewise, the equilibrium outcomes D_{H0} , U_{H0} , D_{L0} , and U_{L0} are identical to the emissions tax case. Intuitively, when the policy maker's ex ante information regarding abatement costs is correct, the choice between price-based and quantity-based policy instruments has no bearing on emissions outcomes in equilibrium.

⁴⁶In an earlier working paper, we compared outcomes under the NOx Budget Program as implemented against a differentiated program that delivered the same quantity of emissions. This comparison was problematic insofar as neither policy was optimal conditional on ex ante available information. Here, we elect to compare ex ante optimal differentiated and undifferentiated trading regimes. As noted above, an emissions cap of 726 thousand tons per ozone season minimizes total social costs under an undifferentiated trading regime, conditional on ex ante available information regarding both damages and costs. This is somewhat less stringent than the cap imposed in the NOx Budget Program. (659 thousand tons).

As in the tax case, the differentiated trading program welfare dominates undifferentiated trading. These net gains are represented graphically as the sum of the two shaded triangles in Figure 7 (identical to areas A and B in Figure 6).

When the econometric model is used to simulate compliance decisions under a quantity-based regime, a very different picture emerges. Recall that the policymaker charged with setting the emissions cap does so based on expectations over both source-specific damages and source-specific abatement costs. In contrast to the price-based regime, accurate abatement cost information is required to define optimal differentiated and undifferentiated quantity-based policy parameters (see *R1*).

Because ex post realized abatement costs are higher than expected; the ex ante optimal emissions cap proves to be too restrictive ex post. This has welfare implications for both undifferentiated and differentiated trading.

Under the undifferentiated regime, the permit price must rise well above the level of the undifferentiated tax of \$1630/ton to clear the market. Marginal abatement costs are set to \$3008/ton across sources. In other words, the ex ante optimal emissions cap (which proves to be too stringent) incentivizes significantly more emissions abatement as compared to the ex ante optimal tax. The benefits from this additional abatement do not justify the costs. Consequently, the undifferentiated price-based regime welfare dominates the undifferentiated quantity-based regime. Net benefits are \$353 million under the cap-and-trade program (see Table 6) versus \$378 million under the undifferentiated tax (see Table 5).

Under the differentiated trading regime, the price required to clear the market is more than double the average marginal damage value. This leads to a significant and asymmetric distortion in policy incentives across sources. Each source must pay $\frac{\delta_i}{\bar{\delta}}\tau$ to offset a unit of emissions. When the market clears at a price of $\bar{\delta}$, firms have an incentive to set marginal abatement costs equal to the value of the damage caused by a unit of emissions. When the market clears at a price that exceeds average damages, firms are incentivized to invest too much in abatement activities. And this distortion is proportional to the source-specific damage parameter δ_i .

This distortion is illustrated graphically in Figure 7. To construct this figure, we use the marginal abatement cost curves defined in Figure 6 by the intersection of the emissions tax rates and equilibrium emissions (for high and low damage sources respectively). We add to this graph the equilibrium permit prices and corresponding emissions levels under differentiated and undifferentiated trading. This figure shows that the marginal abatement costs exceed marginal damages under both differentiated and undifferentiated trading. The distortion is particularly dramatic under the differentiated trading regime among high damage sources.

Importantly, we find that policy differentiation leads to a *reduction* in welfare for tradable permits if we assume that the costs that rationalize observed choices represent the true costs incurred. In the right panel of Figure 7, the hatched area represents abatement costs savings less the value

of increased damages. Under the differentiated policy regime, emissions levels among relatively low damage sources increase vis a vis the undifferentiated regime, and abatement costs fall (as expected). In the left panel, the hatched area represents the increase in abatement costs among relatively high damage sources less the value of avoided damages. The increase in abatement costs induced by policy differentiation easily exceeds the value of the additional avoided damages because the market clearing permit price significantly exceeds the average marginal damage value. The distortion created by the overly stringent cap (and the associated permit price that significantly exceeds \$1733/ton NO_x) is so extreme that the costs of policy differentiation (in the form of increased abatement costs) exceed the benefits.

We note two important qualifications with respect to this surprising finding. First, this is not a general result. If regulators had over-estimated abatement costs, then the differentiated policy would welfare dominate undifferentiated emissions trading (all else equal). Hence, the relative welfare performance of undifferentiated versus differentiated quantity-based policies is a function of how (not just if) the regulator mis-judges abatement costs. This result is similar to that reported by Muller (2012) and it prohibits a general result pertaining to the welfare implications of policy differentiation under cost and damage uncertainty for tradable permits.

Second, while the Monte Carlo exercises executed in this analysis capture the statistical uncertainty associated with particular parameters, we do not undertake a sensitivity analysis over different parameter distributions in AP2. For example, in order to value mortality risks, we employ results from USEPA's meta analysis which estimates a mean \$6 million VSL and an associated standard deviation of roughly \$3 million (see USEPA, 1999;2010). An alternative would be to use the results from Mrozek and Taylor (2000) who estimate a mean VSL of approximately \$2 million. Using a different distribution for the VSL would shift our marginal damage values for all power plants. As prior authors note, both the value attributed to mortality risks and the dose-response function governing PM_{2.5} exposure and adult mortality impacts have pronounced impacts on the level of the marginal damage estimates. While our use of the \$6 million VSL is a standard choice, many researchers use a different PM_{2.5}- mortality dose-response function than that used herein. In particular, the most commonly used alternative function (Laden et al., 2006) yields (mean) damages that are on the order of two-to-three times larger than the function employed in the present paper. Note that if we adopted this parameter value and the associated distribution, the resulting shift in expected marginal damage would approximately equalize the inter-source average NO_x marginal damage to the market clearing permit price.

5.3 Damage uncertainty

Thus far, our discussion has focused almost exclusively on the expected values of avoided damages under the policy scenarios we consider. However, our analytical approach allows us to examine how

parameter uncertainty, and the stochastic nature of meteorological factors, generate uncertainty about the welfare impacts of different policy alternatives. As explained above, we simulate damages and associated policy outcomes for 5000 states of the world (each one characterized by draws from the distributions of the damage model parameters). In addition to the mean values, we report the 5th and 95th percentile values of simulated damage outcomes under the policy regimes we analyze.

Tables 5 through 7 show how the range of estimates of the net welfare impacts of the NOx Budget Program is quite large. In states of the world characterized by larger marginal damage values, net benefits are much larger (in excess of \$1.7 B). In states of the world characterized by much lower damage values, the costs of complying with the program (which do not vary across damage realizations) exceed the benefits.

There is also significant variation in our estimates of the welfare gains from policy differentiation across the 5000 parameter draws. One of the factors driving this variation is the degree to which marginal damages across sources. Recall from (R3) that gains from policy differentiation are increasing with the degree of intersource variation in damages, all else equal. Another factor generating variation in the simulated gains from differentiation is the degree of correlation between expected damages and realized damages. As noted in section 3.2, if ex ante information regarding the ranking of sources in terms of marginal damages is incorrect, the parameters that define the terms of compliance will incorrectly incentivize additional abatement activities among relatively low damage sources (and vice versa).

Finally, it is interesting to note that the variance in the simulated gains from policy differentiation is larger under the quantity-based (versus price-based) policy regime. Intuitively, this is because the degree to which the imposed cap is too stringent varies across the 5000 realizations. In some of the cases we consider, realized damages exceed expected damages to such an extent that the imposed cap is too lax. We find that policy differentiation is welfare improving in 25 percent of simulated cases. Notably, in almost 10 percent of cases, gains from policy differentiation in the quantity-based regime exceed gains under the price-based policy.

6 Conclusion

How should market-based emissions regulations be designed and implemented when damages from emissions vary significantly across sources? To shed light on this question, we develop a conceptual framework to analyze the welfare implications of policy differentiation. We consider differentiated taxes and a form of differentiated emissions trading that has been emphasized in the theory literature. Differentiated policy designs welfare dominate undifferentiated designs under perfect information. Once uncertainty is introduced into the model, the welfare implications of policy differentiation become more ambiguous.

The theory model provides the foundation for an applied analysis of the gains from policy differentiation. We consider the landmark NOx Budget Program (NBP). Although emissions damages vary considerably across sources, for the sake of simplicity, the program features undifferentiated emissions trading. This begs the question: could welfare have been significantly increased under a differentiated policy regime?

There are several important findings. The first comes out of a direct comparison of the econometrically estimated model of firms' compliance decisions and a more stylized cost minimization model calibrated to match a deterministic simulation model used to inform policy-making. Consistent with earlier work by Carlson et al. (2000), we find substantive differences between the observed compliance choices and those predicted by the cost minimization algorithm. In our case, abatement costs that rationalize observed choices exceed ex ante engineering estimates by a significant margin. These higher-than-expected compliance costs notwithstanding, we estimate that the gains from the NOx Budget Program as implemented easily exceed the abatement costs incurred.

Having found that the undifferentiated emissions trading program delivered welfare gains, we then ask whether these gains would have been smaller or larger under a differentiated trading regime. More precisely, we compare differentiated and undifferentiated policies designed to maximize ex ante expected welfare. We find that a differentiated emissions trading regime would likely have reduced welfare vis a vis the undifferentiated policy that was implemented. This is due to the fact that the imposed cap was apparently too stringent given ex post realized costs. In contrast, gains from policy differentiation would likely have been positive under a tax-based regime.

It is worth emphasizing that our findings run counter to the general expectation that policy differentiation will be unambiguously welfare improving. Such expectations were formed under the assumption of perfect information and thus overlooked the possibility that the ex ante optimal cap may prove too stringent ex post (e.g. Mendelsohn, 1986). We show that negative welfare impacts are not only possible, but likely in the particular case of a large scale NOx trading program.

Of course, our empirical results speak only to the welfare implications of policy differentiation in the NOx Budget Program. This policy context need not be representative of other programs. For example, the additively separable damage function and underestimation of abatement costs ex ante need not apply in general. That said, uncertainty regarding emissions damage and abatement costs is ubiquitous. The ideas and findings presented here demonstrate the potential importance of this uncertainty in determining the welfare consequences of policy differentiation across a range of policy designs.

References

- Abbey, D., B. Ostro, F. Peterson, and R. Burchette. "Chronic respiratory symptoms associated with estimated long-term ambient concentrations of fine particulates less than 2.5 microns in aerodynamic diameter and other air pollutants". *Journal Expo Analysis Environmental Epidemiology* 5(2), pp.137–159.
- Baumol, W. and W. Oates (1988). *The Theory of Environmental Policy*. Cambridge Univ Pr.
- Bell, M. L., A. McDermott, S. Zeger, J. Samet, and F. Domenici (2004). "Ozone and short-term mortality in 95 US urban communities, 1987-2000.". *Journal of the American Medical Association* 17.
- Burtraw, D., R. Bharvirkar, and M. McGuinness. "Uncertainty and the net benefits of emissions reductions of Nitrogen Oxides from electricity generation". *Land Economics* 79(3).
- Carlson, C., D. Burtraw, and K. Palmer (2000). "Sulfur dioxide control by electric utilities: What are the gains from trade". *Journal of Political Economy*, 1293–1326.
- Deschenes, O., M. Greenstone, and J. Shapiro (2012). "Defensive investments and the demand for air quality: Evidence from the nox budget program and ozone reductions". Technical report, National Bureau of Economic Research.
- Diamond, P. A. (1973). "Consumption externalities and imperfect corrective pricing". *Bell Journal of Economics* 4(2).
- Electric Power Research Institute (1999). UMBRELLA: Software for Assessing NOx Control Technology Combinations, Version 1.0. Technical report.
- Esty, D. C. (2004). "Environmental protection in the information age". *New York University Law Review*.
- Farrow, R., M. Schultz, P. Celikkol, and G. 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.
- Fishelson, G. (1976). "Emission control policies under uncertainty". *Journal of Environmental Economics and Management* 3(3), 189–197.
- Fowle, M. (2010). "Emissions trading, electricity restructuring, and investment in pollution abatement". *American Economic Review* 100(3), 837–69.
- Fowle, M., S. P. Holland, and E. T. Mansur (2012, April). "What do emissions markets deliver and to whom? evidence from southern california's nox trading program". *American Economic Review* 102(2), 965–93.

- Fullerton, D., S. P. McDermott, and J. P. Caulkins (1997). Sulfur dioxide compliance of a regulated utility. *Journal of Environmental Economics and Management* 34(1), 32 – 53.
- Horan, R. D. and J. S. Shortle (2005). "When two wrongs make a right: Second-best point-nonpoint trading ratios". *American Journal of Agricultural Economics* 87(2), 340–352.
- Joskow, P. L., R. Schmalensee, and E. M. Bailey (1998). The market for sulfur dioxide emissions. *The American Economic Review* 88(4), pp. 669–685.
- Keith, G., D. White, and B. Biewald (2003). "OTC emission reduction workbook 2.1: Description and users manual". Technical report, Synapse Energy Economics.
- Keohane, N. (2007). "Cost savings from allowance trading in the 1990 clean air act: estimates from a choice-based model". In C. E. Kolstad and J. Freeman (Eds.), *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience*. New York: Oxford University Press.
- Krupnick, A., V. McConnell, and R. for the Future (2000). *Cost-effective NOx Control in the Eastern United States*. Resources for the Future.
- Laden, F., J. S. F. S. D. D. (2006). "Reduction in fine particulate air pollution and mortality: Extended follow-up of the harvard six cities study". *American Journal of Respiratory and Critical Care Medicine* 173, pp. 667–672.
- Linn, J. (2008). "Technological modifications in the nitrogen oxides tradable permit program". *The Energy journal* 29(3), 153–176.
- Martin, K. C., P. L. Joskow, and A. D. Ellerman (2007). "Time and location differentiated nox control in competitive electricity markets using cap-and-trade mechanisms". Working Papers 0704, Massachusetts Institute of Technology, Center for Energy and Environmental Policy Research.
- Mendelsohn, R. (1986). Regulating heterogeneous emissions. *Journal of Environmental Economics and Management* 13(4), 301–312.
- Millock, K., C. Nauges, and T. N. S. Sterner (2004). "Environmental taxes: A comparison of French and Swedish experience from taxes on industrial air pollution". *CESifo DICE Report* 2(1), 30–34.
- Montgomery, W. (1972). "Markets in licenses and efficient pollution control programs". *Journal of economic theory* 5(3), 395–418.
- Muller, N. "Linking policy to statistical uncertainty in air pollution damages". *Berkeley Electronic Press Journal of Economic Analysis and Policy* 11(1).
- Muller, N. "The design of optimal climate policy with air pollution co-benefits". *Resource and Energy Economics* 34, pp. 696–722.

- Muller, N. (2011). "Linking policy to statistical uncertainty in air pollution damages". *The B. E. Journal of Economic Analysis & Policy* 11(1), 32.
- Muller, N. and R. Mendelsohn (2009). "Efficient pollution regulation: Getting the prices right". *The American Economic Review* 99(5), 1714–1739.
- New England ISO (2006). 2006 New England marginal emissions rate analysis. Technical report, ISO New England System Planning Department.
- OTAG Air Quality Analysis Workgroup (1997, June). Telling the OTAG ozone story with data. Technical report.
- Pope, C., R. Burnett, M. Thun, and et al (2002). "Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution". *JAMA* 287(9), 1132–1141.
- Seinfeld, J. and S. Pandis (2006). *"Atmospheric Chemistry and Physics: From Air Pollution to Climate Change"* (2 ed.). Wiley-Interscience.
- Stavins, R. (2005). "Lessons learned from SO₂ allowance trading". *Choices* 20(1), 53–58.
- Tietenberg, T. (1995). "Tradeable permits for pollution control when emission location matters: What have we learned?". *Environmental and Resource Economics* 5(2).
- Tietenberg, T. (2006). *Emissions Trading-Principles and Practice*. Resources for the Future Pres.
- Train, K. (2003). *Discrete choice methods with simulation*. Cambridge university press.
- Turner, D. (2006). *"Workbook of Atmospheric Dispersion Estimates: An Introduction to Dispersion Modeling"*. CRC.
- United States Environment Protection Agency (1998a). "Analyzing electric power generation under the CAAA".
- United States Environment Protection Agency (1998b). "Regulatory impact analysis for the NO_x SIP Call, FIP, and Section 126 Petitions".
- United States Environmental Protection Agency (2005). NO_x budget program : 2005 compliance and environmental results.
- United States Environmental Protection Agency (2008). "NO_x Budget Trading Program 2008 Progress Reports".
- United States Office of Environmental Protection. "The benefits and costs of the Clean Air Act: 1990 - 2010.
- US Environmental Protection Agency Science Advisory Board (2012). SAB review of EPA's retrospective cost study of the costs of epa regulations: An interim report of five case studies. Technical report.

Weitzman, M. (1974). "Prices vs. quantities". *The Review of Economic Studies*, 477–491.

Woodruff, T., J. Parker, and K. Schoendorf (2006). "Fine particulate matter (pm2.5) air pollution and selected causes of postneonatal infant mortality in California". *Environmental Health Perspectives* 114(5), 786.

World Health Organization (2003). Health aspects of air pollution with particulate matter, ozone and nitrogen dioxide. Technical report.

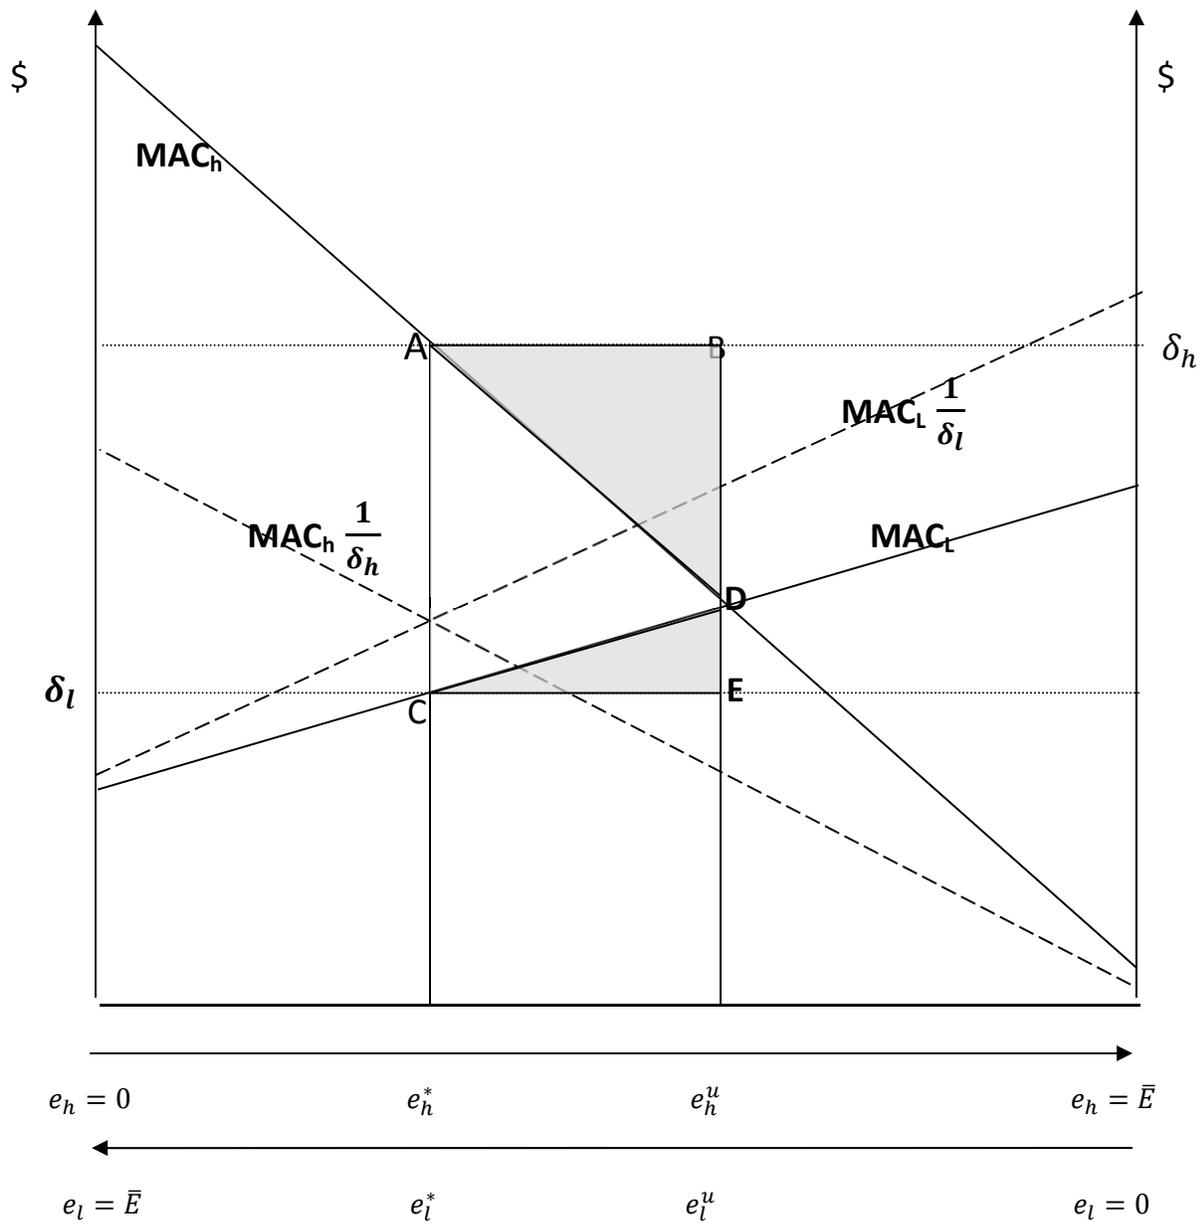


Figure 1: Emissions permit market outcomes under differentiated and undifferentiated policies: Optimal emissions constraint

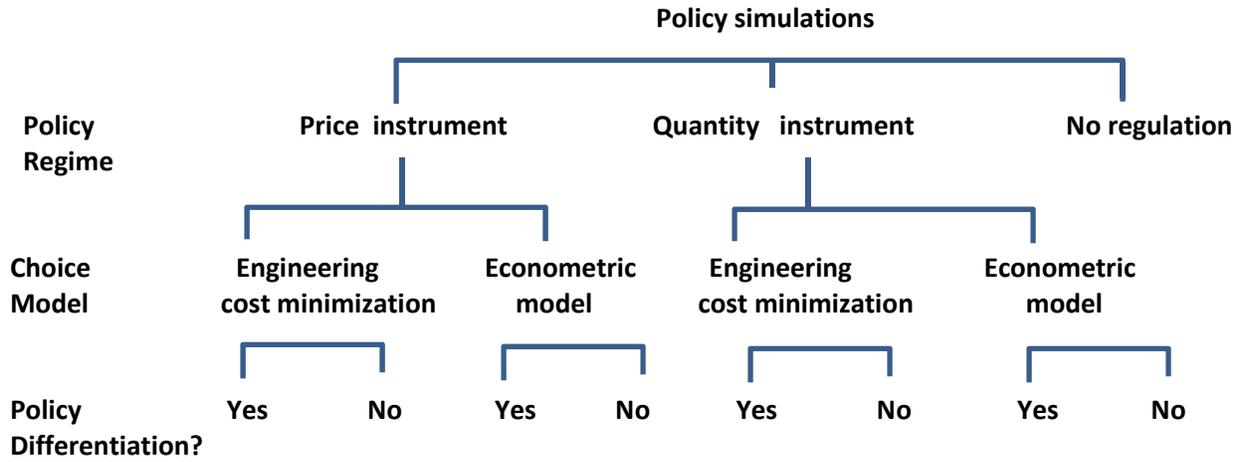


Figure 2: Policy simulation road map

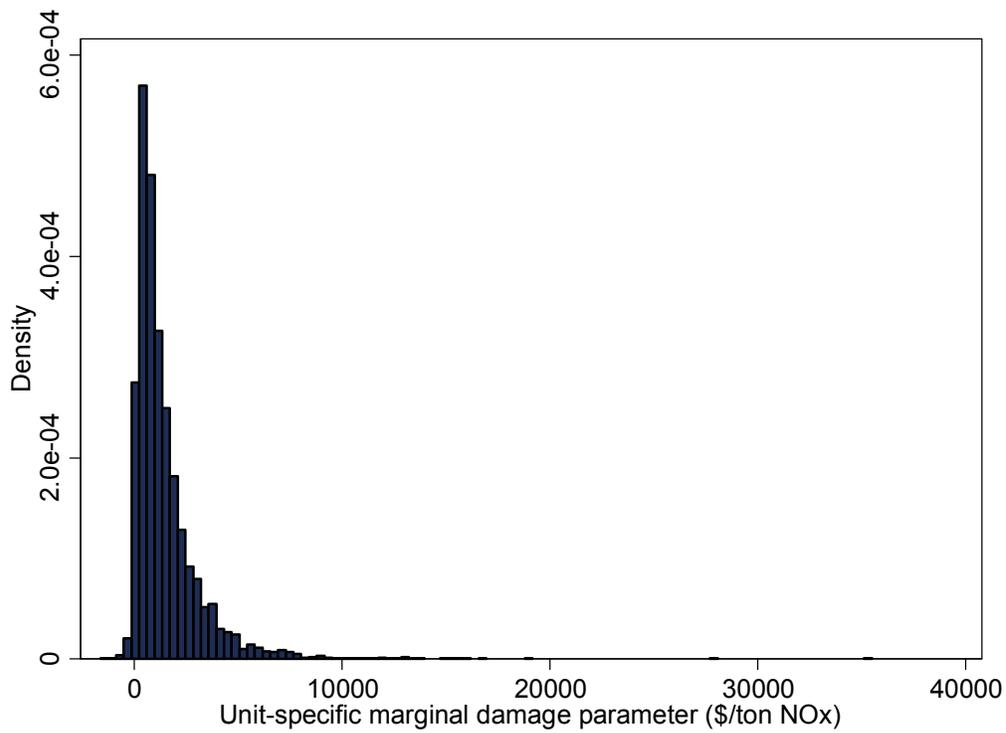


Figure 3 : Within source variation in marginal damage parameter estimates

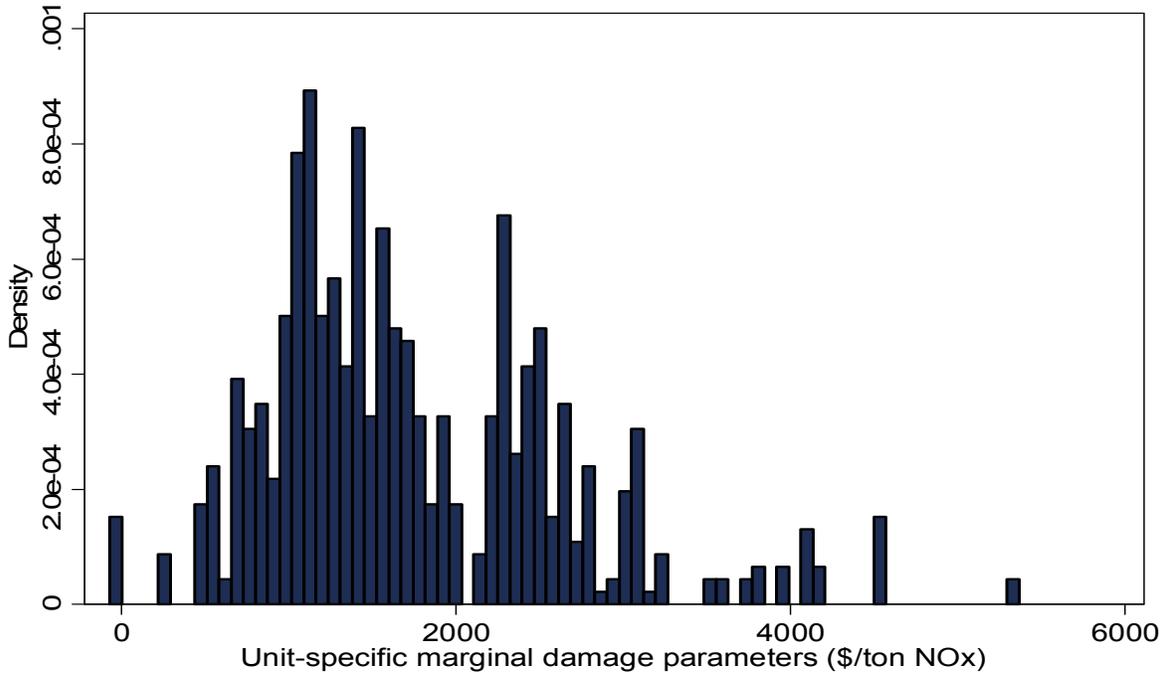


Figure 4: Between source variation in point estimates of marginal damage values

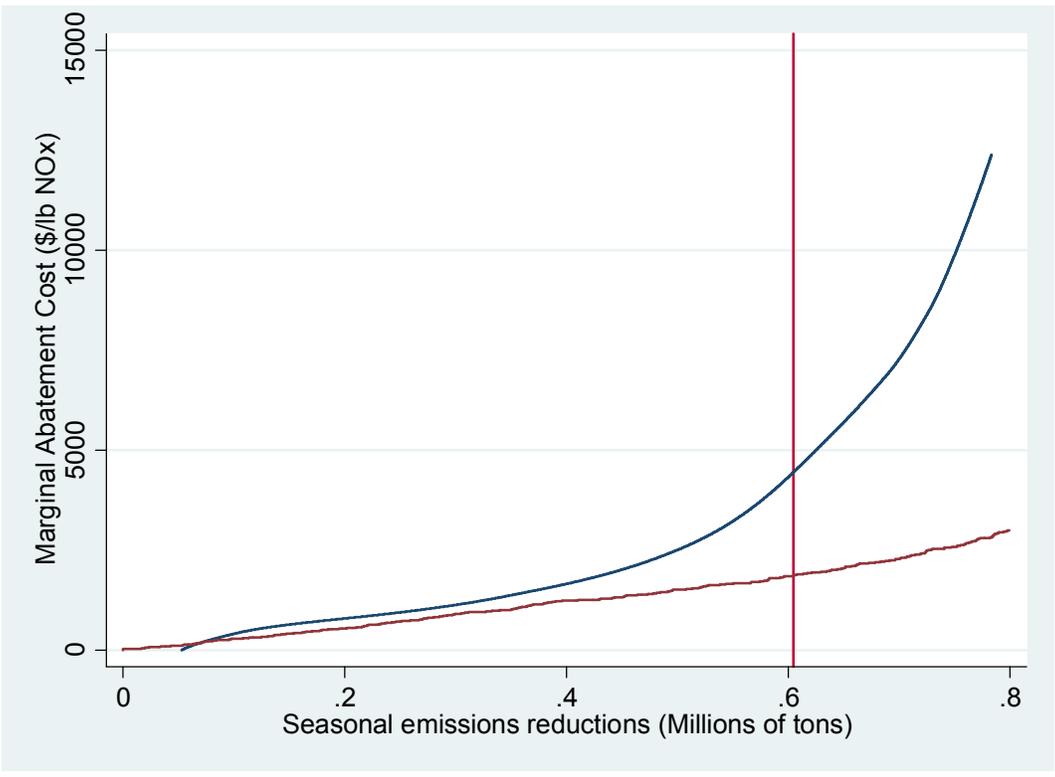


Figure 5 : Aggregate marginal abatement cost curves generated using alternative models of the facility-level compliance choice

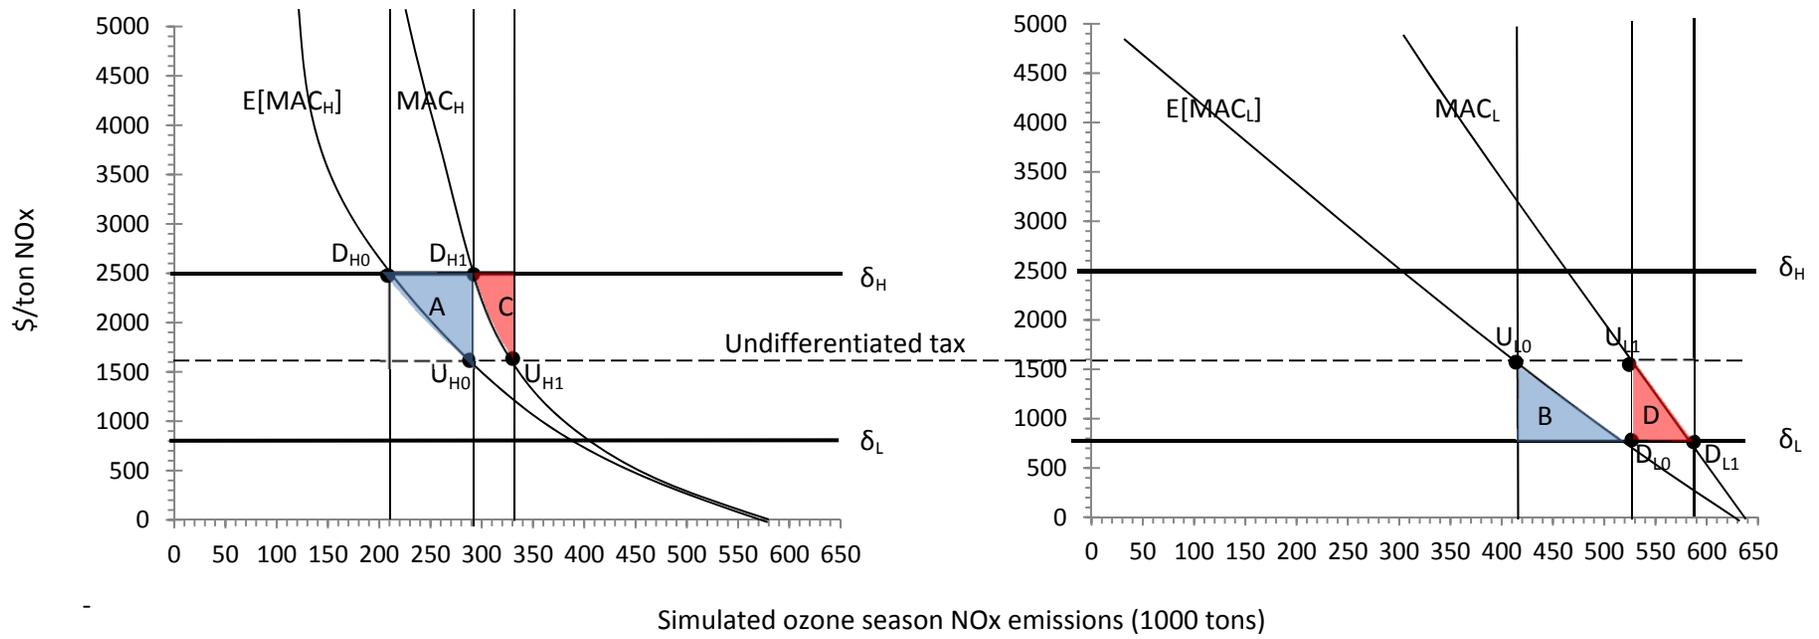


Figure 6 : Gains from policy differentiation under the tax-based regime

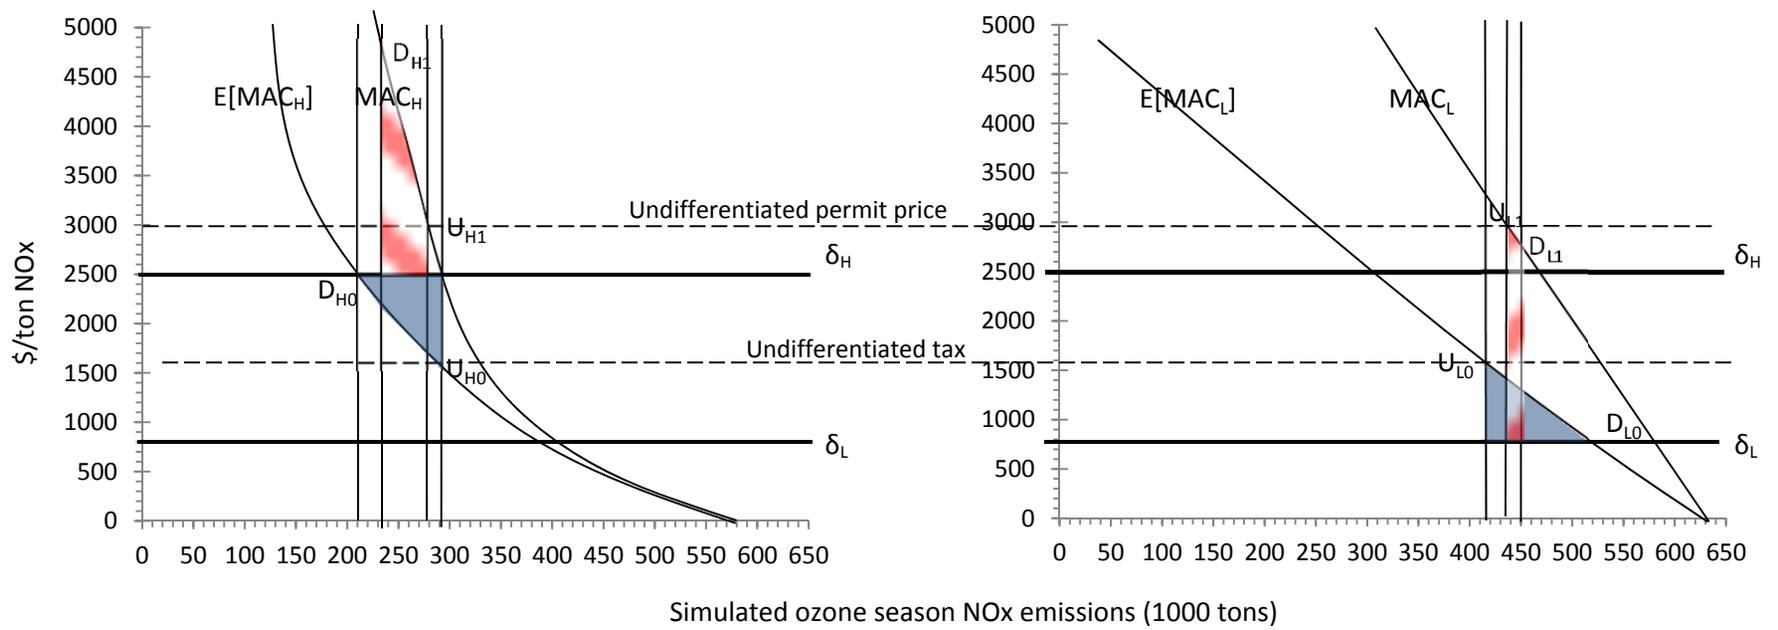


Figure 7 : Gains from policy differentiation under the quantity-based regime

Table 1 : Unit-level summary statistic

Variable	High damage	Low damage
# Units	248	384
Capacity (MW)	255.3 (243.3)	282.1 (255.0)
Pre-retrofit NOX emissions rate (lbs NOx/mmbtu)	0.55 (0.24)	0.50 (0.20)
Boiler age (years)	35.9 (10.5)	36.5 (11.6)
Summer capacity factor	65.0 (14.9)	66.1 (15.2)
Ozone season production (MWh)	687,868 (725,920)	779,450 (767758)
Average damage parameter (\$/ton NOx)	2621 (692)	1161 (370)

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 : Observed, predicted, and correctly predicted compliance choices

Compliance choice	SCR	SNCR	Low NOx burners	Combustion Modifications	No retrofit	Total
Observed choices	187	42	53	58	292	632
Cost minimization model						
Predicted adoption rate	62	80	228	159	103	632
Correctly predicted	47	6	48	9	76	186 (29%)
Econometric model						
Predicted choices	179	15	35	21	382	632
Correctly predicted	166	7	22	18	284	497 (79%)

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

Table 3 : Econometrically estimated coefficients of the compliance choice model

Technology specific constants	High damage units	Low damage units
Post-combustion controls	-3.06 (1.35)	-2.21 (1.61)
Low NOx burners	-2.33 (0.43)	-2.06 (0.53)
Combustion modifications	-2.32 (0.69)	-1.89 (0.85)
Age* capital cost interaction	-0.13 (0.06)	-0.17 (0.07)
Manager-specific coefficients		
Annual compliance cost (\$1,000,000)	-0.99 (0.58)	-1.09 (0.82)
Capital cost (\$1,000,000)	-0.28 (0.33)	-0.45 (0.43)
# Units	248	384

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

Table 4 : Simulated outcomes under observed policy regime

Model of compliance choice	Cost minimization (1)	Econometric (2)
Permit price (\$/ton NO_x)	\$1,864	\$4,460
Ozone season emissions (thousand tons NO_x)	658.5	658.6
% permitted emissions occurring at high damage sources	41%	38%
Levelized annual costs (\$M) (engineering estimate)	\$461	\$712
Levelized annual costs (\$M) (derived from choice model)	\$461	\$733
Annual benefits (\$M) (monetized avoided damages)	\$1,009 (\$175-\$2444)	\$997 (\$173-\$2436)
Annual net benefits (\$M) (Preferred measure)	\$546 (-\$282, \$1979)	\$264 (-\$558, \$1710)

Notes: This table summarizes the results from simulating investment in NO_x abatement and the associated ozone-season emissions under the observed, undifferentiated trading regime. 95 percent confidence intervals are in parentheses. “High damage” units are those with above average damage parameter point estimates. “Low damage” units are those with below average damage parameter point estimates.

Table 5 : Simulated gains from policy differentiation– Emissions tax

Model of compliance choice	Cost minimization		Econometric	
	Undifferentiated	Differentiated	Undifferentiated	Differentiated
Permit price (\$/ton NOx)	\$1630		\$1630	
Ozone season emissions (thousand tons NOx)	723.6	746.6	868	881
% emissions occurring at high damage sources	40%	28%	38%	33%
Levelized annual costs (\$M)	\$361	\$402	\$259	\$272
Avoided annual damages (\$M)	\$883 (\$158, \$2134)	\$988 (\$182, \$2370)	\$637 (\$104, \$1588)	\$697 (\$125, \$1698)
Net benefits (\$M)	\$522 (-\$203, \$1773)	\$586 (\$-220, \$1969)	\$378 (-\$155, \$1329)	\$426 (-\$147, \$1426)
Effects of policy differentiation				
Change in annual costs (\$M)	\$41		\$13	
Change in avoided annual damages (\$M)	-\$105 (\$7, \$256)		-\$60 (\$6, \$143)	
Net gains from differentiation (annual in \$M)	\$64 (-\$34, \$215)		\$47 (-\$7, \$130)	

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

Table 6 : Simulated gains from policy differentiation– Cap-and-trade program

Model of compliance choice	Cost minimization		Econometric	
	Undifferentiated (1)	Differentiated (2)	Undifferentiated (3)	Differentiated (4)
Permit price (\$/ton NOx)	\$1630	\$1732	\$3008	\$3937
Permits allocated (000)	726.3	642.2	726.3	642.2
Ozone season emissions (000 tons NOx)	726.3	746.6	726.2	698.2
% emissions occurring at high damage sources	40%	28%	38%	33%
Levelized annual costs (\$M)	\$361	\$402	\$521	\$674
Avoided annual damages (\$M)	\$883 (\$158, \$2134)	\$988 (\$182, \$2370)	\$878 (\$152, \$2155)	\$990 (\$180, \$2405)
Net benefits (\$M)	\$522 (-\$203, \$1773)	\$586 (-\$220, \$1969)	\$353 (-\$373, \$1629)	\$316 (-\$495, \$1731)
Effects of policy differentiation				
Change in annual costs (\$M)	\$41.0		\$149	
Change in avoided annual damages (\$M)	-\$105 (\$7, \$256)		-\$112 (\$6, \$267)	
Net gains from differentiation (annual in \$M)	\$64 (-\$34, \$215)		-\$37 (-\$130, \$117)	

Notes: This table summarizes the results from simulating investment in NOx abatement and the associated ozone-season emissions under counterfactual policy designs. “High damage” units are those with above average damage parameter point estimates. “Low damage” units are those with below average damage parameter point estimates. The 5th percentile and 95th percentile values are reported in parentheses.