Optimal Climate Policy with Air Pollution Co-Benefits.

Nicholas Z. Muller*

* Department of Economics
Environmental Studies Program
Middlebury College
303 College Street
Middlebury, VT 05753
phone: (802) 443 5918
email: nmuller@middlebury.edu

November 5^{th}, 2010

Abstract

Optimal Climate Policy with Air Pollution Co-Benefits.

This paper develops a model of an optimal regulatory program for greenhouse gas (GHGs) emissions that accomodates the benefits due to reductions of co-pollutants. The model extends the usual treatment of GHGs to include a spatially-variant damage term dictated by the source-specific emission rate of the following co-pollutants: sulfur dioxide (SO_2), nitrogen oxides (NO_x), volatile organic compounds (VOC), and fine particulate matter ($PM_{2.5}$). Employing per ton damage estimates for the co-pollutants produced by an integrated assessment model, co-pollutant damage estimates per ton carbon dioxide equivalent (CO_2e) are developed for over 10,000 sources of GHGs in the lower 48 states including both transportation sources and electric power generation. For coal-fired electric power generation and diesel-powered vehicles, the co-pollutant damages are larger in magnitude than recent peer-reviewed estimates of the marginal damage for GHGs, independent of co-pollutant emissions, although this result is sensitive to damage modeling assumptions in the integrated assessment model. This implies that the practical impact of GHG emissions from these sources is significantly greater than what is reported in current research. Further, the co-pollutant damage per ton CO_2e varies considerably across source types and source location. Exclusive of the damage due to GHG emissions themselves, bundled co-pollutant damages range from -\$200/ton CO_2e to \$1,000/ton CO_2e.The analytical model shows that adopting a policy that encompasses the spatially-variant co-pollutant damage may not be welfare improving relative to the standard approach to managing GHGs. The aggregate welfare impact depends on whether the aggregate limit on GHGs is optimal, too lenient, or excessively strict. If the limit on emissions is too strict, a policy that relfects the co-pollutant damage reduces net welfare, relative to a policy that effectively ignores co-pollutants. If the emission cap is too lenient, then the co-pollutant policy improves welfare. This result has important implications for GHG policy in the United States; although co-pollutant benefits of abating GHGs have been shown to be significant in magnitude, tailoring climate policy to reflect these source-specific co-benefits is not necessarily socially beneficial.

Keywords: Climate change, local air pollution, policy instrument choice.
JEL Classifications: Q54, Q53, Q58

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

1 Introduction

Much has been written on the economics of climate change. The earliest work discussed the design of optimal control policies using models that integrate climate impacts and growth (Nordhaus, 1982;1992). Economists have also explored the measurement of damages and adaptive behavior, (Mendelsohn and Seo, 2008). More recently, economists have focused on the characterization of uncertainty in measuring climate impacts (Weitzman, 2009). Throughout this literature, the standard treatment of modeling greenhouse gases (GHGs) is to treat the damages caused by emissions as independent of source location and specifications. From this perspective; one unit of carbon dioxide (CO_2), or any other GHG, imposes the same externality cost regardless if it is emitted from a diesel-powered truck in Kenya or from a natural gas-fired power plant in California. In light of this property of GHGs, extant policies and policy proposals employ cap-and-trade programs that feature one permit price across all regulated firms. And, if the dynamically tightening aggregate emission targets are set appropriately, having a uniform permit price is sufficient to attain dynamic efficiency since the shadow price for a ton of GHG abated is independent of the source doing the abatement.

This paper's contribution to a crowded literature is in modeling the case where this is not an accurate depiction of the impact of GHG emissions. The point here is not that some new properties of GHGs have been discovered that suggests their impact is location dependent. Rather, the difference is in recognizing that when sources elect to abate GHGs, other pollutants are often also reduced in tandem with GHGs. That is, a large portion of the GHGs that are emitted are due to the burning of fossil fuels. And when fossil fuels are burned, in addition to GHGs, many other pollutants are generated. For example, when coal is burned, sulfur dioxide (SO_2), nitrogen oxides (NO_x), and particulate matter (PM) are produced. Combustion of natural gas produces primarily NO_x. Diesel fuel yields copious quantities of SO_2. Each of these actions also yield GHGs. Since these emissions are produced together - and since abating CO_2 by certain methods implies curtailing these emissions in tandem - the impact of these emissions depends both on the GHGs produced and the specific mix of co-pollutants that are co-generated. The result of this is that the effective impact of GHG emissions may vary according to the different mixes of co-pollutants that are associated with particular production processes and consumption habits.

An additional dimension to this problem increases the likelihood of heterogeneity in damages; even though burning the same fuels yields equivalent mixtures of both GHGs and co-pollutants, the impact per unit of local co-pollutants (such as those listed above) varies considerably according to where they are emitted. Recent research has shown that SO_2 and $PM_{2.5}$ emissions vary in terms of their impact per ton by more than 100-times depending on whether the emission occurs in a large city or a rural area (Muller and Mendelsohn, 2009). Returning to the above example, a ton of CO_2 abated by driving a hybrid vehicle in rural Kenya will generate vastly different benefits than will driving that vehicle in Nairobi, Kenya because avoided exposures to the co-pollutants that would have been produced if a conventional fuel vehicle were driven are many times greater in Nairobi.

The bottom line is this; unless GHGs are abated using a technology that specifically targets GHGs, rather than switching to low-carbon fuels for instance, the avoided damage due to GHG abatement will likely vary according to where abatement is conducted. So if carbon capture and storage (CCS) is widely adopted, then this point of paying close attention to the bundling of abatement of several pollutants becomes moot because CCS isolates and abates just carbon. However, if economies elect to meet carbon caps by switching from coal to natural gas in producing electricity, then this point matters very much. And this is far more likely to occur on a large scale given that alternative abatement strategies - such as fuel-switching - are much less expensive at the margin than CCS.

Bundled or coupled pollutants have important implications for the design of optimal GHG abatement policies. Without the spatially-sensitive component (the co-pollutant component), optimal policy equates marginal abatement costs to the marginal impact of GHG emissions. And since this impact is independent of the emission source, costs are equated across all sources at the margin. The appropriate instruments in this case are either a globally harmonized carbon tax, or a cap-and-trade system with one permit price. However, if the co-benefit of abatement is addressed, when marginal costs are equated to marginal damages, they will differ across sources. In this context, allocatively efficient market-based environmental policies must feature source-specific emission (Pigouvian) taxes, or a system of damage-weighted exchange rates between regulated firms (Baumol and Oates, 1988; Farrow et al., 2004; Muller, Mendelsohn, 2009). The central point for policy design is the following; the usual approach to managing (or proposing to manage) GHGs using uniform-price cap-and-trade policies cannot obtain allocative efficiency when the impact of co-pollutants are recognized.

This paper sets up a model to explore the welfare impact of switching from an underlying belief about GHGs that their impact is spatially uniform to one that acknowledges co-pollutant benefits of abatement. This problem is addressed in the context of an arbitrary cap on GHG emissions set through a political process that bears no resemblance to either what economists would consider an optimal aggregate limit or what climate scientists would necessarily think of as a safe limit. It is arbitrary. The results suggest that incorporating the co-pollutant damages into policy may not be welfare-improving in the context of an arbitrary cap. The intuition is the following; despite the fact that this switch would reduce the deadweight loss due to ignoring co-pollutants, the welfare impact of the interaction between the aggregate cap and the source-specific co-pollutant damage may overwhelm the benefit of capturing the co-pollutant damage if the the cap is inefficiently strict. More specifically, if the aggregate emission cap is too strict a net welfare reduction due to adopting a policy that reflects spatially-variant co-pollutant damages is possible. If the cap is too lenient, the shift in policy cannot decrease welfare.

The implication of this finding for policymakers is that it is likely to make more sense to address spatial variation in co-benefits if GHG policy begins with modest emission caps. Multi-pollutant rules and legislation have recieved significant attention in the U.S. recently, so this finding may be of use in guiding that discussion. If policymakers elect to pursue stringent emission cuts quickly, the better path forward is to structure policy without explicit attention to the variation in co-pollutant benefits.

The empirical section of the paper computes bundled pollutant damages on a per ton CO_2e basis for over 10,000 GHG sources in the contiguous U.S. The results indicate that for coal-fired electric power generation and diesel-powered vehicles, the co-pollutant damages are in excess of recent peer-reviewed estimates of the per-ton damage of GHGs, independent of co-pollutant emissions (Tol, 2008). This implies that the actual benefit due to abating GHG emissions from these sources is considerably greater than what is reported in current research. Further, the co-pollutant damage per ton CO_2e varies considerably across source types and source location. Co-pollutant damages range from -\$200/ton[1] CO_2e to \$1,000/ton CO_2e. These source-specific co-pollutant damages are then used to simulate the welfare impact due to adopting a policy that reflects co-benefits. As the analytical model predicts, the degree and direction of the inefficiency of the aggregate emission constraint determines the welfare impact of the policy shift. For exceedingly lenient emission caps, adopting the optimal co-benefit policy is welfare improving. When the aggregate emission cap is excessively stringent, the policy shift from the standard approach to managing GHGs to the co-benefit policy reduces welfare.

Other authors have explored co-pollutant benefits of climate change policy. Bollen et al. (2009) employ a macro-economic model to examine how co-benefits are affected by the scale of GHG abatement policies and the distribution of co-benefits across multiple countries and regions. Bollen et al. (2009) also gather results from a collection of studies that focus on co-benefits. In particular, Bussolo and O'connor (2001) explore the co-benefits of GHG policies in India. Garbaccio et al. (1999) and O'Connor et al (2003) report the co-benefits due to GHG abatement policies in China. Cao, Ho, and Jorgensen (2008) also focus on estimating the co-benefits of GHG policies in China. Bollen et al. (2009) also summarize many of the behaviorial factors which are likely to impact the co-benefits due to GHG policy such as fuel price elasticities, the complementarities in abatement of both GHGs and local pollutants in Europe, and impacts of forests and agriculture. An underlying, common theme throughout these extant papers is focusing on measuring aggregate co-benefits of various GHG policies. The present paper's exclusive focus on the U.S., on computing source-specific

[1] Although the vast majority of sources generate positive co-pollutant damages, 120 sources produce negative co-pollutant impacts. This stems from the ability of NO_x emissions in dense urban areas to reduce O_3 damages. For a complete discussion of this phenomenon see Tong et al., 2006, or Muller et al., 2009.

co-benefits/ton CO_2e and on efficient co-pollutant policy distinguishes it from this earlier literature.

The remainder of the paper is organized as follows. Section 2 depicts the analytical model and section 3 describes the empirical model used in this study. Section 4 presents the empirical results while section 5 concludes with discussion and policy implications.

2 Analytical Model

Employing a quadratic model of the benefits (avoided damages) and costs of controlling emissions, the analysis begins with a traditional characterization (without recognition of co-benefits) of the benefit of abating GHGs in (1).

$$B_{jt} = \left(\beta_{0t} + \beta_{1t}E_{jt} + \beta_{2t}E_{jt}^2\right)(1+r+\gamma)^{-t} \tag{1}$$

where: r = discount rate

γ = GHG decay parameter

It is assumed that $\left(\frac{\partial B_{jt}}{\partial E_{jt}}\right) < 0$, and $\left(\frac{\partial^2 B_{jt}}{\partial E_{jt}^2}\right) < 0$. In (1) B_{jt} represents the benefits of firm (j) abating GHGs in time (t), which are also the avoided damages of GHG emissions, defined comprehensively to encompass both market and non-market impacts. (E_{jt}) denotes emissions abatement. Following the intercept, the first two terms on the right-hand side of (1) reflect a second-order approximation of the impact of firm (j) emissions. This specification is intended to capture the link between emissions abatement in time period (t) and subsequent benefits that are a function of the accumulated stock of GHGs[2]. Hence, the vector of (β_{kt}), for $k = 1, 2, 3$, terms reflect the different stock levels occuring in different time periods; avoided emissions have different degrees of benefit as the GHG stock grows or declines through time. The effect of emissions abatement from all other sources (N) at time (t): $E_c = \sum_{i=1}^{N} E_{it}$, is subsumed into the ($\beta_{kt}$) vector. Similarly, the effect of GHGs generated by all firms (N) in the infinite past up to the current time period is modeled through the stock effect and also subsumed into (β_{kt}). And finally, the impact that GHG discharges produced by all (N) firms beginning next period (t = 1) and extending into the infinite future will also impact the stock in future time periods (GHGs tend to be long-lived). The future abatement (emission) trajectories of all other firms will impact resulting future benefits of abatement by partially dictating the future stock of GHGs. The effect of future stock levels on current abatement benefits is also reflected in

[2] An alternative approach to specifying (B_{jt}) would include variables that explicitly denote the impact of abatement conducted by other firms in the current time period, the past, and the future. However, the results derived from the analytical model do not depend on these aspects of the specification of (B_{jt}).

the time-variant (β_{kt}).The entire benefit function is multiplied by the $(1+r+\gamma)^{-t}$ term where (r) is the discount rate and (γ) represents the natural decay rate of GHGs in the upper atmosphere[3].

Note that the (β_{kt}) vector does not vary by emission source. This reflects the standard treatment of modeling the damages of GHGs in which the damages caused by emissions are independent of source location and specifications. In order to model the impact of spatially-variant co-pollutant emissions that are inextricably linked to GHG emissions, equation (1) must be modified slightly as shown in (2).

$$B_{jt} = \left(\beta_{0t} + (\beta_{1t} + \theta_{jt})E_{jt} + \beta_{2t}E_{jt}^2\right)(1+r+\gamma)^{-t} \tag{2}$$

The change is focused on the β_1 term and the extension of model (1) essentially boils down to adding a source-specific, time-period-specific parameter (θ_{jt}) to allow for differences in the impact of firm (j) emissions in the current time period. The terms that reflect the impact of accumulated previous emissions, and expected future emissions, are unaltered in this specification because the co-pollutants are not long lived. Hence, their prior (future) emissions have little or no bearing on this apsect of the impact of emissions in the current period. Since this analysis models source-specific damage functions, channeling the (θ_{jt}) component through (E_{jt}) linearly amounts to changing the vertical intercept of the marginal benefit function: $\left(\frac{\partial B_{jt}}{\partial E_{jt}}\right)$.

Two questions on this choice of modeling (θ_{jt}) arise. Why is the change forced through the linear emission term? And why do past or future emissions not matter for the impact of co-pollutants? Regarding the first question; the GHG abatement function is likely to be nearly constant with respect to E_{jt}; β_2 is likely to be quite small (Hoel, Karp, 2000). This is based on the assumption that each firm (j) produces a sufficiently small quantity of emissions, relative to the combined output of all (N) sources, which for GHGs is all sources around the globe, such that changes to E_{jt} do not move the aggregate emission level at time (t), $\sum_{i=1}^{N} E_{it}$, up or down along the benefit (damage) function. Now, bringing in the co-pollutants, in order for the co-pollutant damage to impact β_2, the underlying damage function for the co-pollutants must also be non-linear in E_{jt}. Recent empirical evidence suggests that, for individual sources (even large industrial facilities), the marginal damage function for these co-pollutants are nearly flat, implying that total damages are linear[4]. Therefore, the impact of co-pollutants is to vertically shift the marginal benefit function for GHGs on a source-specific and time period-specific basis.

As for the second question, since local pollutants do not persist or accumulate in the atmosphere across multiple time periods, modeling their impact in a way that does not depend on the stock of GHGs (either

[3] The empirical value for γ for CO_2 is somewhere in the neighborhood of 0.0001 (IPCC, 2007). For other GHGs such as methane (CH_4) $\gamma = 0.08$. And for nitrous oxide (N_2O), $\gamma = 0.009$.

[4] The reader is encouraged to review Muller and Mendelsohn (2009) for a discussion of this phenomenon.

the current stock due to past emissions or the expected future stock) is consistent with how these substances behave in the atmosphere[5].

Next, the abatement cost function is introduced in equation (3). The cost function is expected to have first and second partial derivatives with respect to emissions abatement that are positive: , $\left(\frac{\partial C_{jt}}{\partial E_{jt}}\right) > 0, \left(\frac{\partial^2 C_{jt}}{\partial E_{jt}^2}\right) > 0$.

$$C_{jt} = \left(\alpha_{0t} + (\alpha_{1t} + \delta_{jt})E_{jt} + \alpha_{2t}E_{jt}^2\right)(1+r)^{-t} \tag{3}$$

The $(\alpha_1 + \delta_{jt})$ term allows for variation in abatement costs by source (j) and across time periods (t). This heterogeneity reflects source-specific abatement opportunities. Certain firms may adopt energy-saving technologies to reduce carbon emissions. Others may have to switch to low-carbon fuels or fuel blends. Some industries or sectors may pursue more expensive abatement strategies such as capturing and storing carbon before it is emitted. Equation (3) allows for these differences through the (δ_{jt}) term.

The next step in the analysis is to explore whether and how the inclusion of source-specific parameters (θ_{jt}) into the benefit function impacts the optimal emission path. To accomplish this, the difference, in terms of efficient emissions (E^*_{jt}) characterized by the benefit functions expressed in equations (1) and (2) is determined.

2.1 Allocative Efficiency

In this setting, the regulator's goal is to maximize the net benefits of pollution control as shown in (4).

$$\max_{E_{jt}} \; B_{jt} - C_{jt} \tag{4}$$

The well-know first-order condition with respect to (E_{jt}) is that marginal benefits are equated to marginal abatement costs in present value terms. Note that (5) reflects the regulator's use of the benefit function in (1) which does not allow for source-specific effects of emissions.

$$(\beta_{1t} + 2\beta_{2t}E_{jt})(1+r+\gamma)^{-t} = (\alpha_{1t} + \delta_{jt} + 2\alpha_{2t}E_{jt})(1+r)^{-t} \tag{5}$$

[5] This paper assumes that both the first and second central moments of (θ_{jt}) are well-defined: $\mathrm{E}(\theta_{jt}) = \bar{\theta}$, $\mathrm{E}(\theta_{jt}^2) = \sigma^2$. Although it is hypothesized that $(\theta_{jt} \geq 0)$, the empirical section of the paper finds that approximately 1% of the sources covered in the paper have $(\theta_{jt} < 0)$. Becuase of the small number of such anamolous sources, their impact on the welfare implications of policy outcomes is imperceptible.

Rearranging to solve for (E_{jt}) displays the socially-optimal level of abatement for firm (j) in terms of the marginal benefit and marginal cost functions. Increasing benefits of abatement increase the optimal abatement level. Increasing marginal costs drive down the optimal level. Inelastic benefit and cost functions reduce optimal abatement; the elasticities are captured by ($2\alpha_2$) and ($2\beta_2$). To clarify the analytical results, γ is assumed to be equal to zero. As mentioned above, for CO_2 γ is approximately equal to 0.0001, which, given that γ always appears additively with the discount rate (r), assuming γ is equal to zero will not have a discernable impact when the empirical simulations are conducted.

$$E^*_{jt1} = \frac{\beta_{1t}\,(r+1)^t - ((\alpha_{1t} + \delta_{jt})\,(r+\gamma+1)^t)}{2(\alpha_{2t}\,(r+\gamma+1)^t - \beta_{2t}\,(r+1)^t)} \rightarrow \frac{\beta_{1t} - \alpha_{1t} - \delta_{jt}}{2\,(\alpha_{2t} - \beta_{2t})} \text{ if } \gamma \simeq 0. \tag{6}$$

Running through the same exercise but now employing the benefit function in (2) results in the following characterization of optimal levels for firm (j).

$$E^*_{jt2} = \frac{\beta_{1t} + \theta_{jt} - \alpha_{1t} - \delta_{jt}}{2\,(\alpha_{2t} - \beta_{2t})} \tag{7}$$

The difference between (E^*_{jt1}) and (E^*_{jt2}) reduces to:

$$E^*_{jt2} - E^*_{jt1} = \frac{\theta_{jt}}{2(\alpha_{2t} - \beta_{2t})} \tag{8}$$

Expression (8) indicates that the change to emissions abatement from the case where the source-specific term is ignored (using benefit function 1) to where it is included (employing benefit function 2) is an increasing function of the source-specific co-pollutant damage. The greater the realization of (θ_{jt}) the greater the abatement when (θ_{jt}) is accounted for. Of course, this change is affected by the elasticities (slopes) in the benefit and cost functions.

The welfare implication of this difference in emissions, for firm (j) emissions at time (t), is determined by evaluating the net benefits (using benefit function (2)) at (E^*_{jt1}) and (E^*_{jt2}) and computing the resulting difference. After considerable algebra this works out to be:

$$\Delta W_{jt} = \frac{\theta^2_{jt}}{4(\alpha_{2t} - \beta_{2t})(1+r)^t} \tag{9}$$

Essentially equation (9) represents the usual deadweight loss triangle, in present value terms, attenuated by two-times the slope terms corresponding to B_{jt} and C_{jt}. Since benefit function (1) ignores (θ_{jt}) then the improvement due to adopting a spatial model (like benefit function 2) that encompasses it increases as (θ_{jt}) increases, irrespective of its sign. Equation (9) is always greater than or equal to zero. Unless the underlying benefit and cost curves are perfectly inelastic, welfare always improves if policies adopt benefit function (2).

2.2 The Emission Cap.

The above program indentifies optimal emission paths for benefit functions (1) and (2) while aggregate emissions are unconstrained. This section of the analysis explores optimal emission controls in the presence of an aggregate cap on emissions ($\bar{E}$) that is set exogenously. Setting ($\bar{E}$) exogenously simply means that the regulator that has efficiency as their objective does not control the levers in calibrating the aggregate cap. Some other agent is responsible for this part of the policy process. Often a legislature establishes aggregate goals while an environmental regulator is charged with implementing policy geared to achieve the goal implied by the cap. This situation is captured by the problem shown in (10). The regulator seeks to maximize net benefits (as with the problem in equation 4) but now their ability to engineer the efficient outcome is constrained by the exogenous cap. They pursue constrained efficiency.

$$\max_{E_{jt}} \ (B_{jt} - C_{jt}) \qquad s.t. \ \bar{E}_t \leq \sum_{i=1}^{N} E_{it} \tag{10}$$

The present value Lagrangian expression corresponding to this constrained optimization problem is shown in (11) employing benefit function (2).

$$\begin{aligned} L_{jt} \quad = \quad & \left(\beta_{0t} + (\beta_{1t} + \theta_{jt})E_{jt} + \beta_{2t}E_{jt}^2\right)(1 + r + \gamma)^{-t} - ... \\ & ... - (\alpha_{0t} + (\alpha_{1t} + \delta_{jt})E_{jt} + \alpha_{2t}E_{jt}^2)(1 + r)^{-t} + \lambda_t(\bar{E}_t - \sum_{i=1}^{N} E_{it})(1 + r)^{-t} \end{aligned} \tag{11}$$

The first-order condition with respect to (E_{jt}) is shown in (12).

$$(\beta_1 + \theta_{jt} + 2\beta_2 E)\,(1 + r + \gamma)^{-t} - (\alpha_1 + \delta_{jt} + 2\alpha_2 E)(1 + r)^{-t} = \lambda_t (1 + r)^{-t} \tag{12}$$

The resulting constrained-efficient emission levels for benefit functions (1) and (2), which are denoted (E^C_{jt1}) and (E^C_{jt2}), are very similar to (E^*_{jt1}) and (E^*_{jt2}) shown above with the sole difference being the appearance of λ_t, the Lagrange multiplier, in the numerator[6]. Note that if the aggregate cap ($\bar{E}_t$) corresponds to the optimal aggregate level of emissions, $\left(\sum_{i=1}^{N} E^*_{it}\right)$, then $\lambda_t = 0$, and $\left(\frac{\partial B_{jt}}{\partial E_{jt}}\right) = \left(\frac{\partial C_{jt}}{\partial E_{jt}}\right)$. More interesting, and more likely given the often arbitrary (or a least politically-based) manner in which aggregate emission limits are set, is the case where $\lambda_t \neq 0$. In this situation, with an exogenously calibrated, sub-optimal cap, the difference in terms of quantities (E^C_{jt2}- E^C_{jt1}) according to whether benefit function (1) or (2) is used in (10) is exactly the same as shown in (8) because the Lagrange multiplier appears additively in the numerator of both (E^C_{jt1}) and (E^C_{jt2}).

However, the welfare impact of adopting a policy that results in (E^C_{jt1}) rather than (E^C_{jt2}) differs from the case where $\lambda = 0$. This is shown in (13).

$$\Delta W_{jt} = \frac{\theta^2_{jt} + 2\lambda_t \theta_{jt}}{4(\alpha_{2t} - \beta_{2t})(1 + r)^t} \tag{13}$$

The expression in (13) shows that the welfare impact consists of two components, both in present value terms. First, the welfare improvement of (E^C_{jt2}) relative to (E^C_{jt1}) is an increasing second-order function of the source-specific co-pollutant term (θ_{jt}). This is just the typical welfare triangle shown above in (9). Second, the welfare change is also driven by the wedge between $\left(\frac{\partial B_{jt}}{\partial E_{jt}}\right)$ and $\left(\frac{\partial C_{jt}}{\partial E_{jt}}\right)$, which is λ_t. This "cap inefficiency" is the extent to which regulators cannot induce firms to equate $\left(\frac{\partial B_{jt}}{\partial E_{jt}}\right)$ to $\left(\frac{\partial C_{jt}}{\partial E_{jt}}\right)$ due to the arbitrary (sub-optimal) nature of the aggregate emission constraint.

Note that the sign of the overall welfare change is dictated by the $2\lambda_t \theta_{jt}$ term. If λ_t and θ_{jt} are of opposite signs, then ΔW_{jt} may be negative. Employing benefit function (2) in the constrained-efficient optimization problem may reduce welfare even though it captures the co-pollutant damage. This is somewhat suprising and it is the primary result derived from the analytical model. One would think that if marginal co-benefits vary by emission location, then using a source-specific benefit function would improve - in a welfare sense - the outcome of policy. Expression (13) suggests that this is not necessarily so.

[6] $E^C_{jt1} = \left(\frac{\beta_1 - \lambda_t + \delta_{jt} + \alpha_{1t}}{2(\alpha_{2t} - \beta_{2t})}\right)$
$E^C_{jt1} = \left(\frac{\beta_1 + \theta_{jt} - \lambda_t + \delta_{jt} + \alpha_{1t}}{2(\alpha_{2t} - \beta_{2t})}\right)$

The key in determining whether a reallocation from (E^C_{jt1}) to (E^C_{jt2}) results in a net welfare improvement involves the relative magnitudes and the signs of λ_t, and θ_{jt}; the net marginal benefit, or cap inefficiency, and the source-specific co-pollutant damage term, respectively. First, clearly if λ_t and θ_{jt} are the same sign, then the welfare change is positive; reallocating emissions from (E^C_{jt1}) to (E^C_{jt2}) is beneficial. This outcome is shown in figure (1) for the case where both λ_t and θ_{jt} are positive. The aggregate cap is too lenient (not enough abatement). Employing benefit function (1), MB_1 in figure 1, emissions are E_1, and the resulting deadweight loss is equal to the area (A+B+C). When benefit function (2), MB_2, is employed, emissions adjust to E_2 and the remaining deadweight loss is A. Clearly, welfare improves by (B+C).

However, if $\lambda < 0$, and $\theta_{jt} > 0$, then the direction of the welfare change depends on their relative magnitudes. If $\lambda_t = \frac{1}{2}\theta_{jt}$, in absolute value then there is no welfare change as the $2\lambda_t\theta_{jt}$ term is equal to θ^2_{jt}. If $\lambda_t > \frac{1}{2}\theta_{jt}$, in absolute value, then there is a reduction in welfare as the cap inefficiency dominates the θ^2_{jt} term. Figure 2 depicts the case where $\lambda_t < 0, \theta_{jt} > 0$,and $\lambda_t > \frac{1}{2}\theta_{jt}$. With benefit function (1) emissions begin at E_1 with an associated deadweight loss of B. The adjustment to E_2, stemming from the use of benefit function (2) results in a new deadweight loss of A which is clearly in excess of B. Welfare declines. One can also see from this figure that if $\lambda_t = \frac{1}{2}\theta_{jt}$, A and B would be exactly offsetting and welfare would not change.

From this discussion some general themes emerge. First, if the emission cap is too lax ($\lambda_t > 0$), then the reallocation from (E^C_{jt1}) to (E^C_{jt2}) will be (socially) beneficial since the vast majority of firms have positive realizations of (θ_{jt})[7] and the $2\lambda_t\theta_{jt}$ term will reinforce the θ^2_{jt} term. Hence, aggregate welfare will increase. Similarly, if the emission cap is too stringent ($\lambda_t < 0$), then the reallocation will be (socially) beneficial for firms with realizations of (θ_{jt}) such that: $\theta_{jt} < 0$, and $\theta_{jt} > 2\lambda_t$. However, since most firms have $\theta_{jt} > 0$, when ($\lambda_t < 0$) aggregate welfare will likely decrease due to a move from a policy based on benefit function 1 to a program derived from benefit function 2.

2.3 Policy Design

The design of market-based policy instruments geared to optimally manage GHGs in the context of co-benefits has interesting qualities when viewed in a dynamic sense. First, if the aggregate emission constraint is optimal, a system of emission taxes (a carbon tax) would need to reflect the source-specific marginal benefits of abatement. Sources would face different tax rates. Similarly, if a system of pollution allowances were adopted optimal policy would employ a series of exchange rates (trading ratios) between regulated sources calibrated to the marginal benefits of abatement. While this property of efficient pollution control has long been known, the point that is worth emphasizing here is how the various components of the marginal

[7] This statement will be substantiated in the empricial section of the paper. Suffice it to say that most emissions are co-pollutants are harmful.

benefit of GHG abatement may change through time. The literature that focuses on measuring the damages due to GHG emissions points out that, on a per unit basis, damages from GHG discharges increase into the future. That is, an emission released in the near future will have a smaller impact than an equivalent emission released in twenty years (in present value terms). This increase is driven by the accumulation of GHGs in the atmosphere; as the stock increases, the marginal damage (benefit) due to additional emissions (abatement) increases.

The co-benefit term is also likely to increase but for entirely different reasons. Since co-benefits associated with local air pollutants are driven by adverse impacts on human health, population growth into the future will cause co-benefits per ton to increase. Further, income growth is also likely drive co-benefits upward; greater levels of income are associated with (for example) greater willingness to pay to avoid mortality risks.

The implications of these considerations for optimal GHG policy depends on the relative rates of growth projected to occur in the co-benefit and the GHG components of the marginal benefit function. That is, if the stock of GHGs increases quickly, and the marginal benefit of avoided climate impacts increases quickly as well (more quickly than the co-benefit component), then the degree of heterogeneity across sources in the total marginal benefit will decline. Since the benefit of avoided climate impacts is uniform across sources, if this aspect of the benefits of GHG abatement dominates the co-benefit component, then marginal benefits of abatement converge across sources. If the policy instrument chosen to manage GHGs is an emission tax, then, in this case, the relative tax rates across firms approaches unity as the direct benefits of climate stabilization dominate local co-benefits. Similarly, the optimal trading ratios would tend toward unity.

On the other hand, if population and income growth cause the co-benefit aspect of GHG abatement benefits to overwhelm the benefit of climate stabilization, then the marginal benefit of abatement will tend to diverge across firms. (Aside from rapid population or income grwoth this may occur if climate mitigation strategies are undertaken such that the adverse effects of warming are non-marginally reduced.) The implication for policy in this case would be source-specific tax rates that remain different across sources into the future. And for quantity-based instruments, the optimal allowance exchange rates would remain different than unity for sources with different marginal benefits of abatement.

3 Empirical Model

In order to estimate (θ_{jt}), this paper employs estimates of the damage per ton for SO_2, NO_x, volatile organic compounds (VOCs), and fine particulate matter ($PM_{2.5}$) that were produced by the Air Pollution Emission Experiments and Policy Analysis model (APEEP), (Muller, Mendelsohn, 2007;2009). APEEP is a standard integrated assessment model in its overall structure. However, it provides source-specific marginal damage estimates for the above pollutants. These spatially-variant damages are then used to compute the co-pollutant impacts per ton of CO_2e emitted for a broad range of sources using the methods outlined below. The modeling assumptions employed in APEEP that are critical in determining the magnitude of the marginal damages are standard in the air pollution damage measurement literature (USEPA, 1999). These include modeling the impact of $PM_{2.5}$ on adult mortality rates using the results from Pope et al., (2002), a mortality risk valuation of approximately \$600 per $\left(\frac{1}{10{,}000}\right)$ risk increment, and applying this value uniformly across populations of all ages. Further, the impact of exposures to tropospheric ozone (O_3) on mortality rates is modeled using the results from Bell et al., (2004). APEEP tracks chronic morbidity states including bronchitis and asthma employing the results from Abbey et al., (1995) and McDonnel et al., (1993). In a sensitivity analysis, marginal damages are computed using the age-variant value for mortality risks reported in Aldy and Viscusi (2007). For a complete description of the APEEP model see Muller and Mendelsohn (2007).

To estimate the co-pollutant benefits per ton CO_2e, one must determine the amounts of SO_2, NO_x, VOCs, and $PM_{2.5}$ that are co-emitted with CO_2. This varies according to inputs (fuel choice) and abatement technology. Transportation sources are broadly decomposed into vehicles that employ diesel fuel and those that burn gasoline. The emission rates for both the local pollutants and CO_2 are assumed to be equal to the current regulatory standards for the local pollutants[8]. Table 1 depicts the emission rates for transportation sources. Note that these are expressed relative to one ton of CO_2e. Thus, for each ton of CO_2e that is emitted from a vehicle that uses diesel fuel, 2.6×10^{-3} tons of SO_2 are emitted. For electric power generators (EGUs) distinct co-pollutant emission factors are computed for individual coal-fired, natural gas-fired, and oil-fired generation units. Table 1 indicates that the average emission rate for coal-fired power generation is 6.1×10^{-3} tons of SO_2 emitted for each ton of CO_2e. Similarly, on average, 1.9 x 10^{-3} tons of NO_x are emitted for each ton of CO_2e produced when burning coal. Natural gas-powered facilities generate significantly smaller quantities of both SO_2 and NO_x than do coal-fired plants (per ton CO_2 emitted). Oil-based generators produce SO_2 emissions that are in between coal and natural gas facilties, and NO_x emission rates that are greater than for coal-fired units. For all fuel types emissions of $PM_{2.5}$ and VOCs are considerably smaller than SO_2 and NO_x. It is also important to note the heterogeneity in emissions rates for different types of coal and for different varieties of oil. For instance, some generators burn bituminous coal while others rely on sub-bituminous or lignite. In order to estimate one emission rate for all coal-fired generators in table 1, a capacity-weighted average emission factor was computed. This approach was also employed for oil-based generation units since there are many different types of petroleum products that are used to produce

[8] This suggests that as the older vehicles in the fleet retire, these emission factors become increasing more accurate.

electricity[9].

The per-ton CO_2 emissions factors for source (j) at time (t) of co-pollutant species (s), denoted (CE_{jts}), are computed using the following formula which relies on plant-specific emission data from USEPA's eGrid system (USEPA, 2009).

$$CE_{jts} = \frac{E_{jts}}{MWH_{jt}} \times \frac{MWH_{jt}}{CO_{2jt}} \tag{14}$$

In (14), MWH_{jt} reflects net generation expressed in megawatt-hours for source (j) at time (t), E_{jts} represents emissions of co-pollutant species (s) in tons, and CO_{2jt} corresponds to emissions of CO_2 produced by source (j) at time (t) also in tons. The co-pollutant damages are then estimated by multiplying this fractional ton by the (\$/ton) marginal damage estimate derived from APEEP (MD_{jts}) as shown in (15). Each co-pollutant species (s) is emitted in different quantities relative to GHG emissions based on its content in the fuel burned to produce GHGs and on abatement technology. Coal contains high amounts of SO_2, for example. Adding together the estimated impact across co-pollutants generates a per ton CO_2e co-pollutant damage. So, with the notation from section 2, the co-pollutant benefit (avoided damage) of abating GHGs and therefore abating the co-emissions of pollutant species (s) by source (j) at time (t) on a per-ton CO_2 basis is denoted $\hat{\theta}_{jt}$.

$$\hat{\theta}_{jt} = \sum_{s=1}^{4} (CE_{jts} \times MD_{jts}) \tag{15}$$

An important distinction among the co-pollutant damages, is the type, or specifications, of the source (j). The subscript (j) in equation (15) captures not just location, but also different source types. Damages for the co-pollutants vary considerably according to whether the emission occurs at the ground (as from a vehicle), or high in the air - like from a tall smokestack (Muller, Mendelsohn, 2009; Muller, Tong, Mendelsohn, 2009). As such, distinct per-ton damage estimates for every EGU in the contiguous 48 states produced by APEEP are used for emissions from electric power generators. And ground-level per-ton damage estimates that capture the damage per ton for every county in the coterminous 48 states from APEEP are used for emissions from vehicles.

The advantage of using APEEP is that it provides source and location-specific damage estimates. Thus, the APEEP estimates are able to capture the source-specific variability in the co-pollutant damage that is

[9] The weighted averaging technique was only employed in order to present an emission rate for coal without listing all of the different coal types in the tables. For the co-pollutant damage estimation procedures, the actual source-specific emission rates are calculated, which is not based on the weighted averaging method.

the essence of the (θ_{jt}) term in the analytical model. However, it is important to note that APEEP reports the damages per ton for local pollutants, by source location and type. As noted above, these vary quite a bit across space (Muller, Mendelsohn, 2009). The analysis herein identifies another degree of variation for EGUs; some EGUs may employ abatement devices, such as scrubbers that remove SO_2 or selective catalytic reducers that decrease emissions of NO_x, while others may elect to comply with extant policies using cleaner fuels. Hence, emission rates per ton CO_2e also vary based on these differences. Neither of these factors affect the marginal damage for the local pollutants themselves (SO_2 or NO_x, for example). Hence, the marginal co-pollutant damages reflect both spatial and technological variability.

4 Results

Table 2 provides summary statistics for the estimated $\left(\hat{\theta}_{jts}\right)$ – the co-pollutant damages expressed on a per-ton CO_2e basis. Beginning with transportation sources, the largest damage is due to emissions of SO_2 from diesel-powered vehicles. For each ton of CO_2e, the SO_2 emitted along with the CO_2e causes \$40.4 in damages. The remaining damages are \$6.64 for NO_x, \$5.69 for $PM_{2.5}$, and \$1.62 for VOCs. Note that the damages due to SO_2 are zero for gasoline vehicles; emissions of SO_2 are nearly zero. Also note that the damages are the same for the other three pollutants across vehicle fuel types. This stems from the assumption that emission rates of these co-pollutants are equal to the regulatory standard which are set equally for both diesel and gasoline vehicles (DOE, 2009).

Turning to electric power generation, the greatest co-pollutant damage per ton corresponds to SO_2 emitted when coal is combusted. The damage is \$70.18 per ton CO_2e. NO_xdamage from coal-fired electric generating units (EGUs) is \$3.02/ton CO_2e while for $PM_{2.5}$ the damage is \$2.38. For oil-fired EGUs, SO_2 emissions cause damages of \$34.76 per ton CO_2e, NO_x discharges cause \$3.11 in damages, and $PM_{2.5}$ yields \$1.09 worth of harm for each ton of CO_2e emitted. The impacts from natural gas fired EGUs are considerably smaller for NO_x and especially SO_2. Specifically, SO_2 co-pollutant damages are \$1.05 while for NO_x the damages are \$0.92. $PM_{2.5}$ damages are \$0.18. For all three fuels, VOC damages are around \$0.02 per ton CO_2e.

The bottom row in table 2 displays the total co-pollutant damage due to emitting one ton of CO_2e from each of the five different source categories covered in this analysis. These are the sum across the four co-pollutants - these are the empirical versions of (15). Coal-based power generation yields the greatest co-pollutant damage, which is driven by the SO_2 component, of \$75.59 per ton CO_2e. Diesel vehicles yield the next largest co-pollutant damage of \$54.31 per ton CO_2e. Oil-fired EGUs produce combined co-pollutant damages of \$38.98/ton CO_2e. Gasoline vehicles and natural gas EGUs generate much smaller co-pollutant damages.

A useful exercise is to compare these co-pollutant damages to current estimates of the harm caused by CO_2e emissions. This provides one with a sense of the degree to which optimal GHG abatement paths may change if co-pollutant damages are taken into account. In a recent meta-analysis, Tol (2008) reports that the mean social cost of carbon (SCC) across over 200 studies ranges between \$88/ton and \$127/ton carbon. Converting this to CO_2e produces a damage of between \$24.00/ton and \$34.64/ton CO_2e. The co-pollutant damage caused by coal-fired EGUs, oil-fired EGUs and diesel vehicles lies above this range. This means that the co-pollutant impacts are greater than recent estimate of the climate impacts of CO_2e emissions. Hence, the practical welfare impact of a CO_2e emission from these source types is at least double what the literature has reported when co-pollutant impacts are accounted for. This simply must have significant implications for GHG policy.

The take-home message is this; the effective damage from GHG emissions from these source types would increase dramatically if co-pollutant impacts are counted relative to the case where co-pollutant impacts are ignored. This means that climate policy reflecting the practical impact of GHGs (with bundled co-pollutant damages) is likely to be significantly more stringent than climate policy focusing just on GHG impacts - independent of the bundled pollutants.

Table 2 shows that the co-pollutant impact varies considerably according to source types (coal-fired EGU compared to gasoline vehicles, for example). The next step in the analysis is to recognize that the co-pollutant effects are not homogenous across source locations. Figures 3 and 4 present maps of the combined co-pollutant damages per ton CO_2e. Figure 3 focuses on coal-fired EGUs; the map indicates that the facilities which produce the greatest co-pollutant damage are in the Ohio River Valley stretching from Western Indiana into Pennsylvania. Many of these plants yield damages of greater than \$200/ton CO_2e. It is also notable that facilities in close proximity to one another may have quite different co-pollutant damages. This stems from heterogeneity in input fuels and in abatement devices that control local air pollution emissions. Figure 3 also indicates that most of the coal plants in the Western U.S. have relatively low co-pollutant damages.

Figure 4 maps the impact of co-pollutants emitted by diesel-powered vehicles. Like figure 3, there is a clear west-to-east gradient of increasing co-pollutant damages. Since the emission rates are constant across space[10] this heterogeneity is strictly due to the marginal damage estimates produced by APEEP (Muller, Mendelsohn, 2007;2009). The greatest co-pollutant damages are concentrated in large urban areas.

Table 3 summarizes the distributions of the co-pollutant damages across space for each source type. Again beginning with diesel vehicles, the 5^{th} percentile source generates damages of nearly \$20/ton CO_2e, the median source causes damages of nearly \$50/ton CO_2e, and the 95^{th} percentile source yields harm equivalent

[10] Recall that this portion of the analysis assumes that vehicles emit up to the current emission standards.

to \$110/ton CO_2e. The distribution of damages due to gasoline-powered vehicles shows considerably less dispersion than for diesel vehicles with an interquartile range of just \$9 and a median value of \$12/ton CO_2e. Moving next to the distributions corresponding to EGUs, the 5^{th} percentile source from the distribution of coal-fired facilities generates damages of \$10/ton, the median source causes damages of \$60/ton and the 95^{th} percentile source produces damages equal to \$250/ton. The spread for co-pollutant damages from natural gas facilities is much less than for coal-based generation and the median for this EGU type is \$1.82/ton CO_2e. However, the distribution for oil-fired plants is also quite disperse; the 5^{th} percentile source generates damage of \$2.5/ton, the median yields damage of \$41/ton, while the 95^{th} percentile source produces harm of over \$85/ton. The distributions for diesel vehicles, and coal and oil EGUs are similar in two respects. First, each contains a long right tail indicating the presence of outliers in the distribution. And second, these outliers reflect emissions in densely populated metropolitan areas.

Across all of the source types and locations the maximum co-pollutant damage/ton is approximately \$1,000 due to ground-level emissions produced by diesel vehicles in Los Angeles County, California. In fact, figures 3 and 4 show the pattern that the largest co-pollutant impacts are concentrated in the largest urban and metropolitan areas. In stark contrast are the co-benefits of abating GHGs in rural areas. These co-benefits are as much as 10-times smaller than the co-pollutant impacts in urban areas. Given the relative dominance of health-related impacts of the co-pollutants measured by APEEP, this finding is intuitive; the greatest human exposure per unit of co-pollutant emissions occurs in and near cities.

Table 4 reflects the results of the sensitivity analysis. Employing the age-variant value for mortality risks derived from Aldy and Viscusi (2008) reduces the co-pollutant damage considerably. The reason that this occurs is that most of the co-pollutant induced mortalities occur among the elderly; the impact of exposures to local pollution is proportional to baseline mortality rates which are highest among elderly populations. The mortality risk valuations reported in Aldy and Viscusi (2008) peak at middle age and decline thereafter. Thus mortality risk impacts due to emissions of co-pollutants are are attributed a smaller value under this set of assumptions than when a uniform value is applied. The impact across source types is an approximate 30% reduction in the total co-pollutant damage. Specifically, the total co-pollutant impact for coal-fired EGUs drops from \$75/ton CO_2e to \$53/ton. The co-pollutant damage for diesel vehicles likewise decreases from \$54 to \$39/ton CO_2e. Employing these alternative modeling assumptions, the co-pollutant damage estimates still exceed the damage estimates reported by Tol (2008).

4.1 Policy Simulations

The policy simulations explore the welfare impact of moving from a GHG abatement policy that ignores co-pollutant benefits to a policy that encompasses the spatially-heterogeneous co-pollutant effects. To do this, equation (13) is estimated empirically, and it is decomposed into its two component parts. Table 5 reports empirical estimates of $\sum_{i=1}^{N}\left(\frac{\theta_{it}^2}{4(1+r)^t(\alpha_2-\beta_2)}\right)$ which characterizes the welfare impact of this policy shift due to capturing the heterogeneity in the co-benefits of abatement. The analysis explores multiple values for the discount rate and the elasticities of the marginal benefit and marginal cost functions with respect to (E_{jt}): $2\alpha_2$, $2\beta_2$. Table 5 indictates that the projected welfare improvement due to capturing spatial heterogeneity in the co-pollutant damages ranges from \$18 million to \$90 million, per annum. This amount is considerably smaller in magnitude than the benefit of capturing spatial heterogeneity in damages for the SO_2 cap-and-trade program that was found by and reported in Muller and Mendelsohn (2009). Specifically, those authors found that the analogous welfare improvement in that program was between \$300 million and \$1 billion, annually.

This range is dictated largely by the elasticities on the marginal benefit and marginal cost functions. Note that the simulations assume a value of t = 3. This corresponds to the year 2012 (if we assume that the analysis was conducted in 2009) which is the first year of emission reductions of GHGs under the proposed legislation before the U.S. Congress. As such, the discount rate has a modest impact on the results. (If the analysis looked at the impact of emission reductions further into the future, of course, the discount rate becomes much more important.)

To complete the welfare analysis, figure 5 displays the impact of the aggregate emission cap as reflected by $\sum_{i=1}^{N}\left(\frac{2\lambda_t\theta_{jt}}{4(1+r)^t(\alpha_2-\beta_2)}\right)$ from equation (13). Specifically, figure 5 displays the relationship between the aggregate welfare impact and the magnitude of (λ_t). This figure, which employs r = 0.03, t = 3, and $2\alpha_2$, $2\beta_2 = 0.5$, shows empirically the relationship that was posited in section 3; aggregate welfare is strongly dependent on the magnitude and sign of (λ_t). Note that when ($\lambda_t = 0$), the change in total welfare is approximately \$18 million, which supports the results shown in table 5. Figure 5 also shows that the degree to which the aggregate emission constraint is sub-optimal plays a critical role in determining the aggregate welfare change. For $\lambda_t = 500$, or -500, the total welfare impact is $\pm$\$75 million. These extreme values of (λ_t) reflect aggregate emission caps that are grossly inefficient. This figure emphasizes the key result from the analytical model; if the aggregate emission constraint is highly sub-optimal (meaning λ_t is large in absolute value) it only makes sense to pursue the spatial co-pollutant policy if the cap is too lenient ($\lambda_t > 0$). This is shown in figure 5 where for ($\lambda_t > 0$), the spatial co-benefit policy is welfare-improving. In contrast, if the emission cap is too strict adopting the spatially-variant co-benefit policy will cause aggregate welfare to decline.

For the empirical (θ_{jt}) in this paper, the transition from an aggregate welfare loss to a welfare gain occurs when ($\lambda_t \simeq -55$). Note that in section 3, the rule was identified that determined the sign of equation (13) - if λ_t and θ_{jt} are of different signs, the welfare change shifts from negative to positive where $\lambda_t \simeq \frac{1}{2}\theta_{jt}$. Of course, if $\lambda_t > 0$, welfare always improves. That is reflected in figure 5. However, for $\lambda_t < 0$, if $\lambda_t > \frac{1}{2}\theta_{jt}$, in absolute value, then there is a reduction in welfare, and if $\lambda_t < \frac{1}{2}\theta_{jt}$, in absolute value, then there is an increase in welfare. In light of this, one would expect to observe the aggregate welfare change shifting from negative to positive where $\lambda_t \simeq \frac{1}{2}\bar{\theta}$, where $\bar{\theta}$ equals the arithmetic mean of (θ). The empirical estimate of ($\bar{\theta}$) is approximately 40, and yet the transition from a welfare loss to a welfare gain occurs when ($\lambda_t \simeq -55$) which is clearly larger (absolute value) than $\left(\frac{1}{2}\bar{\theta} = 20\right)$[11]. The discrepancy between the rule derived from the analytical model and the observed threshhold value for λ_t is due to the right-skewed nature of the distributions of (θ). Specifically, the large positive outliers in (θ) dominate the effect of (λ_t) because (θ) enters equation (13) as a second-order term, while (λ_t) enters (13) linearly. Hence, moving from $\lambda_t = -500$ to $\lambda_t = 500$ the transition from negative to positive welfare impacts occurs at larger negative values of λ_t than $-\frac{1}{2}\theta_{jt}$.

The interplay between the two components of equation (13) is shown in figure 6. This figure displays the welfare impact across all of the 10,000 sources in this study - each with its own realization of (θ_{jt}) - for the case where ($\lambda_t = -55$). The total welfare impact, depicted by (+), tracks closely to the independent benefit of capturing variation in the co-pollutant damage, shown by (o). In contrast, the impact of (λ_t), marked by x's, begin positive for the few sources with $\theta_{jt} < 0$, and then slowly (linearly) declines as the (θ_{jt}) increase. This occurs because ($\lambda_t = -55$) in this simulation. Although the vast majority of sources generate very small welfare impacts - as their emissions adjust to a policy that reflects spatially-variant co-benefits - the aggregate welfare impact is dominated by no greater than 10 sources whose realizations of (θ_i) are large. In fact, for this simulation, over 9,000 sources have a net negative welfare impact. The aggregate welfare impact is swamped by the sources that are positive outliers in the distribution of (θ_{jt}). The important point for policy design is that it is very important to capture the welfare improvement associated with efficient levels of co-pollutant emissions from these outlier sources because they drive the aggregate welfare impact.

5 Discussion

The literature on the economics of climate change is extensive. Throughout this literature, the impact of GHGs has been modeled as being independent of source type and location. This paper takes a different

[11] This would imply that the switch from negative to positive welfare impacts would occur at $\lambda_t = -20$.

tack by exploring the impact of co-pollutant that are coupled with GHGs emissions. Since the co-pollutants measured in this paper have per-ton damages that vary by source location, if GHGs are effectively bundled with these other pollutants, then the practical impact of GHGs varies by location as well. The questions explored by this paper are two-fold. First, what are the magnitudes and the spatial distribution of these co-pollutant impacts? Second, if regulators design an optimal policy reflecting these co-pollutant impacts, will this tend to improve social welfare relative to the case where these impacts are ignored?

In answering the first line of inquiry, this paper uses an integrated assessment model to compute the damages due to SO_2, NO_x, VOCs, and $PM_{2.5}$ that are co-emitted with CO_2. This study reports the damages of these pollutants per ton of CO_2e emitted. The magnitude of these ancilliary impacts are large. The total co-pollutant impact due to an emission of CO_2 from a coal-fired power plant is \$75/ton CO_2 and the co-pollutant damage due to CO_2 emissions from a diesel-powered vehicle is \$55/ton CO_2. Emissions from natural gas power plants and gasoline-powered vehicles generate co-pollutant damages that are considerably smaller. An important point of comparison is between these co-pollutant damage estimates and the damage caused by CO_2 itself that are reported in the literature. For example Tol (2008) reports CO_2 damages that range between \$24/ton and \$35/ton CO_2e. Hence, the findings from this study indicate co-pollutant damages are comparable in magnitude, and for some sources in excess of, the primary impacts from CO_2.

In pursuing the second question, this analysis finds that the welfare impact of implementing a spatially-tailored policy that reflects source-specific co-pollutant damages may not improve aggregate welfare. This occurs because such a policy is considered in the context of an arbitrary (sub-optimal) aggregate emission cap and the inefficiency due to the arbitrary cap may overwhelm the benefit due to adopting spatially-nuanced policy. This is an important result because it acts as a cautionary note to policymakers in the following sense; although the co-pollutant benefits of GHG reductions are significant in magnitude, adopting a policy that attempts to manage emissions from different sources based on their degree of harm, either through source-specific taxes or a cap and trade program using exchange rates, may not be welfare improving.

Table 1: Co-pollutant Emission Factors: (tons/ton CO_2e)

Pollutant	Transportation		Power Generation		
	Diesel	Gasoline	Coal*	Natural Gas	Oil*
CO_2	1	1	1	1	1
SO_2	0.0026	0	0.0061	0.0001	0.0032
NO_x	0.0021	0.0021	0.0019	0.0002	0.0045
$PM_{2.5}$	0.0005	0.0005	0.0002	0.00001	0.0001
VOC	0.0002	0.0002	0.00001	0.00002	0.00001

Emission rates for SO_2 and NO_x for coal and oil in table 1
reflect capacity-weighted average across types of coal and oil.

Table 2: Summary Statistics for Co-pollutant Damages: ($/ton CO_2e)

Pollutant	Transportation		Power Generation		
	Diesel	Gasoline	Coal	Natural Gas	Oil
SO_2	40.36 (33.11)	0	70.18 (84.51)	1.05 (5.67)	34.76 (106.44)
NO_x	6.64 (11.17)	6.64 (11.17)	3.02 (4.20)	0.92 (9.86)	3.11 (27.01)
$PM_{2.5}$	5.69 (8.18)	5.69 (8.18)	2.38 (1.69)	0.18 (0.27)	1.09 (1.24)
VOC	1.62 (2.32)	1.62 (2.32)	0.01 (0.01)	0.02 (0.03)	0.02 (0.03)
Total	54.31 (39.59)	13.95 (10.53)	75.59 (86.56)	2.17 (11.70)	38.98 (108.21)
n	3,110	3,110	1,106	2,034	800

Mean, (standard deviation).
Damages reported in table 1 reflect capacity-weighted average across types of coal and oil.

Table 3: Distribution of Co-pollutant Damages by Source and Fuel Type ($/ton CO_2e)

Percentile	Transportation		Power Generation		
	Diesel	Gasoline	Coal	Natural Gas	Oil
5^{th}	19.20	5.26	10.23	-0.40	2.45
25^{th}	31.21	8.57	28.69	0.71	13.9
50^{th}	48.60	11.76	60.83	1.82	41.00
75^{th}	64.85	16.46	115.52	3.96	65.17
95^{th}	112.52	30.81	252.00	34.08	85.29

Table 4: Co-pollutant Damages with Age-Variant VSL ($/ton CO_2e)

Pollutant	Transportation		Power Generation		
	Diesel	Gasoline	Coal	Natural Gas	Oil
SO_2	28.56 (23.44)	0	49.4 (57.03)	0.74 (3.88)	24.43 (62.83)
NO_x	5.63 (3.64)	5.63 (3.64)	2.17 (2.80)	1.03 (6.85)	2.86 (18.44)
$PM_{2.5}$	4.03 (5.80)	4.03 (5.80)	1.67 (0.31)	0.12 (0.27)	0.76 (1.24)
VOC	1.15 (1.65)	1.15 (1.65)	0.01 (0.00)	0.01 (0.03)	0.02 (0.03)
Total	39.36 (30.40)	10.80 (8.52)	53.24 (58.25)	1.90 (8.11)	28.07 (64.27)
n	3,110	3,110	1,106	2,034	800

Mean, (standard deviation).

Damages reported in table 4 reflect capacity-weighted average across types of coal and oil.

Table 5: Welfare Impact of
Capturing Spatial Heterogeneity ($ x 10^6)

(r)	$2\alpha_2$	$2\beta_2$	ΔW
0.03	0.1	0.1	90.5
0.03	0.25	0.25	36.2
0.03	0.5	0.5	18.1
0.03	0.1	0.5	30.2
0.05	0.25	0.25	34.2

t = 3.

$\Delta W = \left(\sum_{i=1}^{N} \frac{\theta_{it}^2}{4(1+r)^t(\alpha_2 - \beta_2)}\right)$ from equation (13).

Figure 1: Welfare Impact of Benefit Function 2: $\lambda, \theta > 0$.

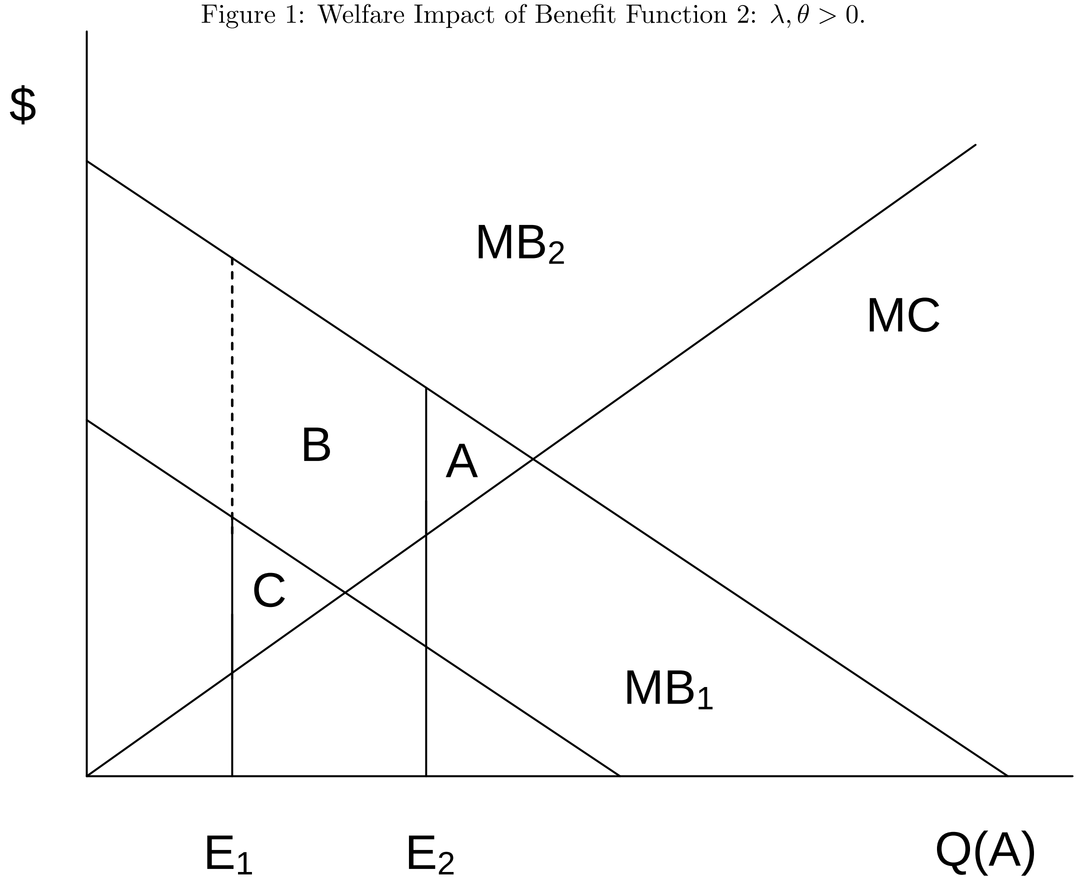

Figure 2: Welfare Impact of Benefit Function 2: $\lambda < 0, \theta > 0, \lambda > \frac{1}{2}\theta$.

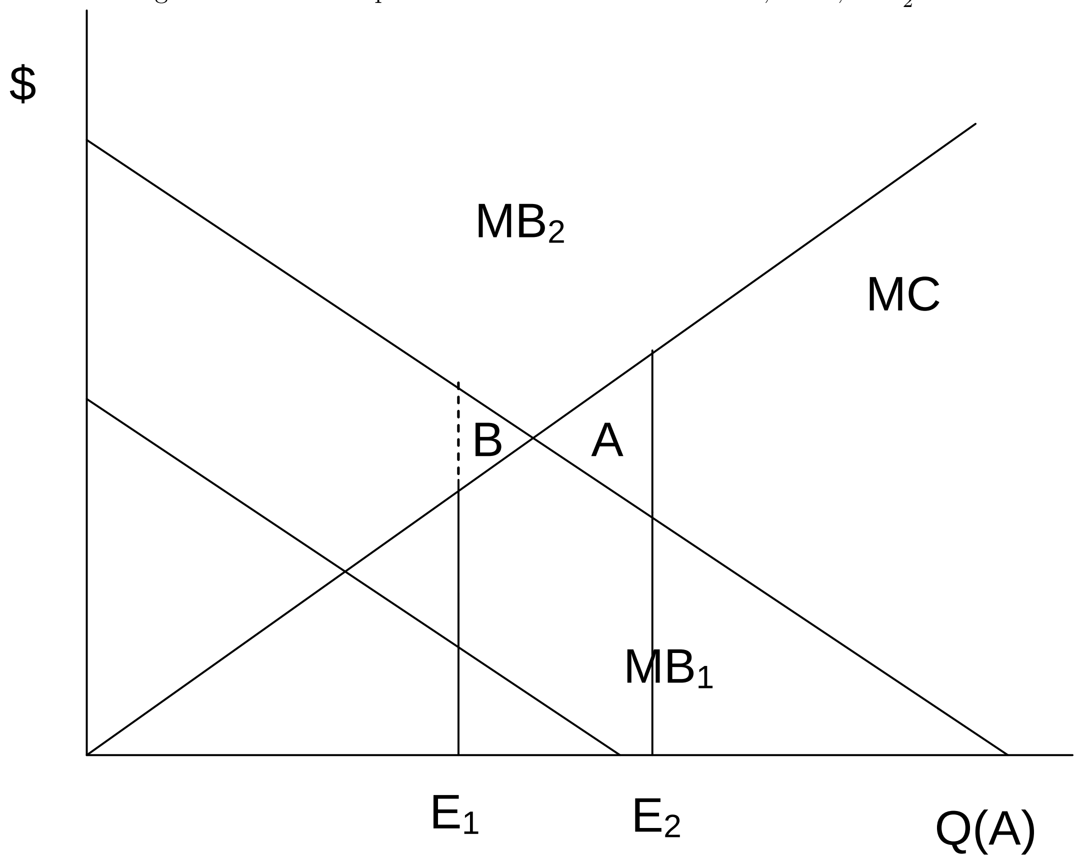

Figure 3: Co-Pollutant Damage Coal-fired EGUs ($\$$/ton CO_2).

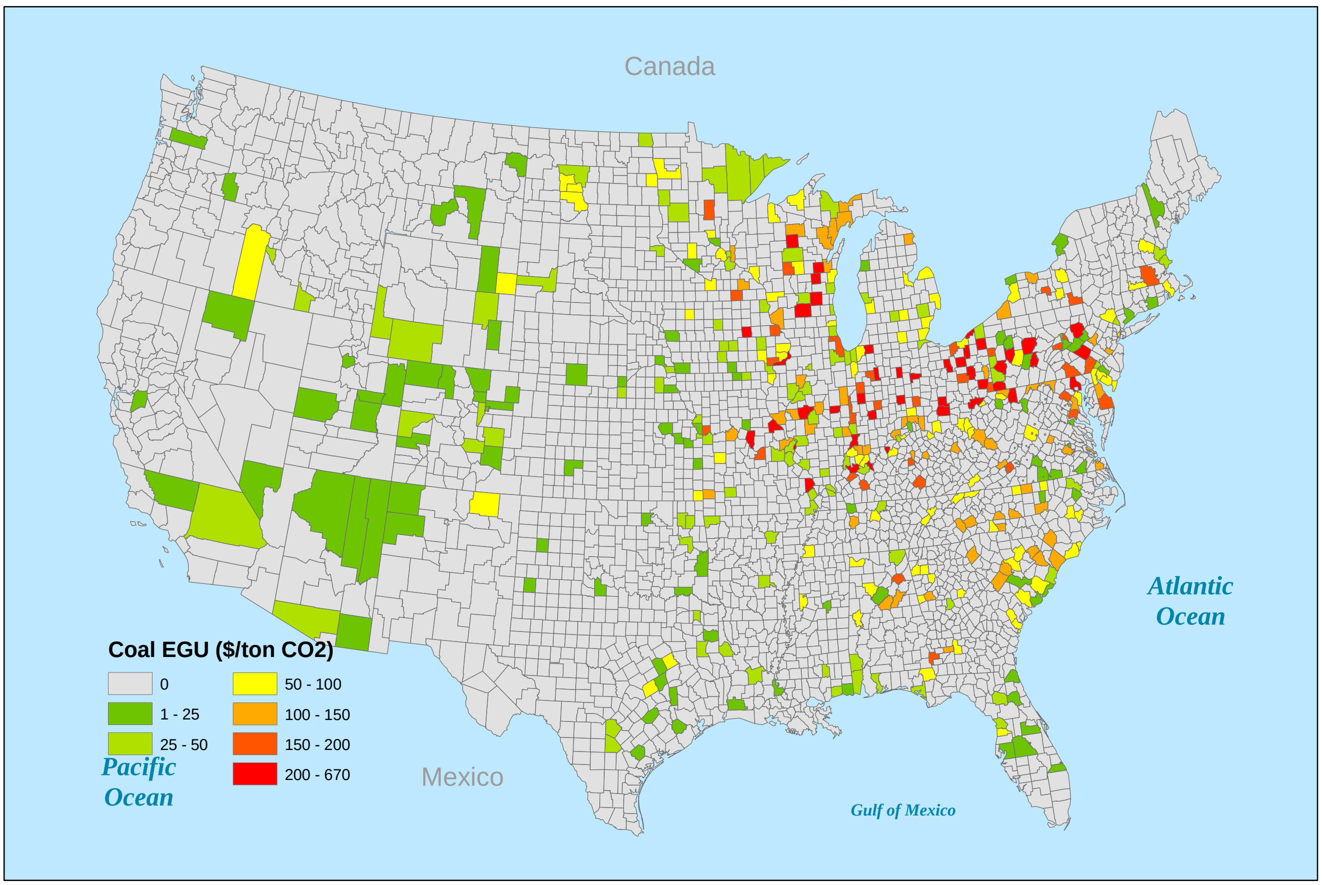

Figure 4: Co-Pollutant Damage Diesel Vehicles (\$/ton CO_2).

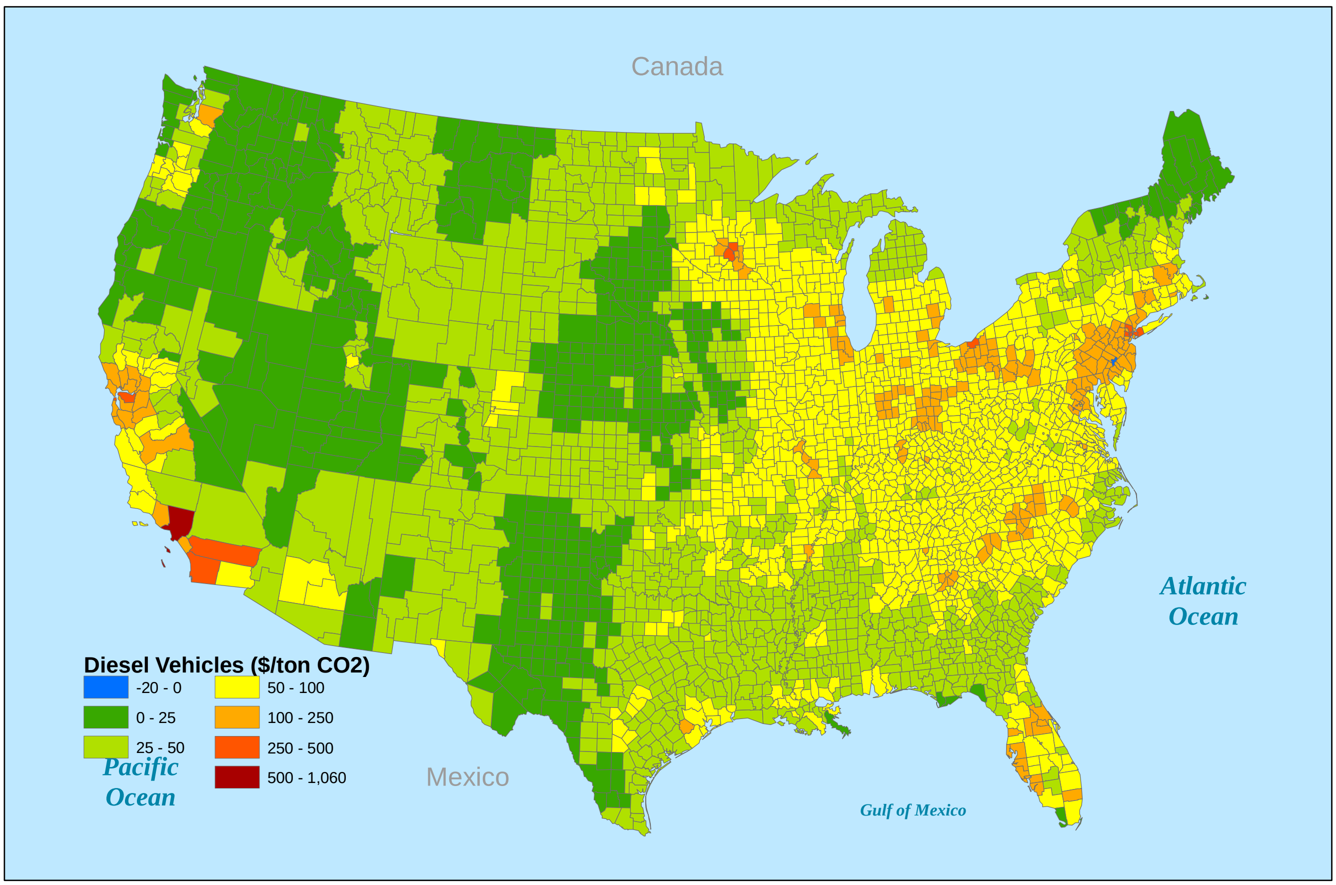

Figure 5: Welfare Change as a Function of λ.

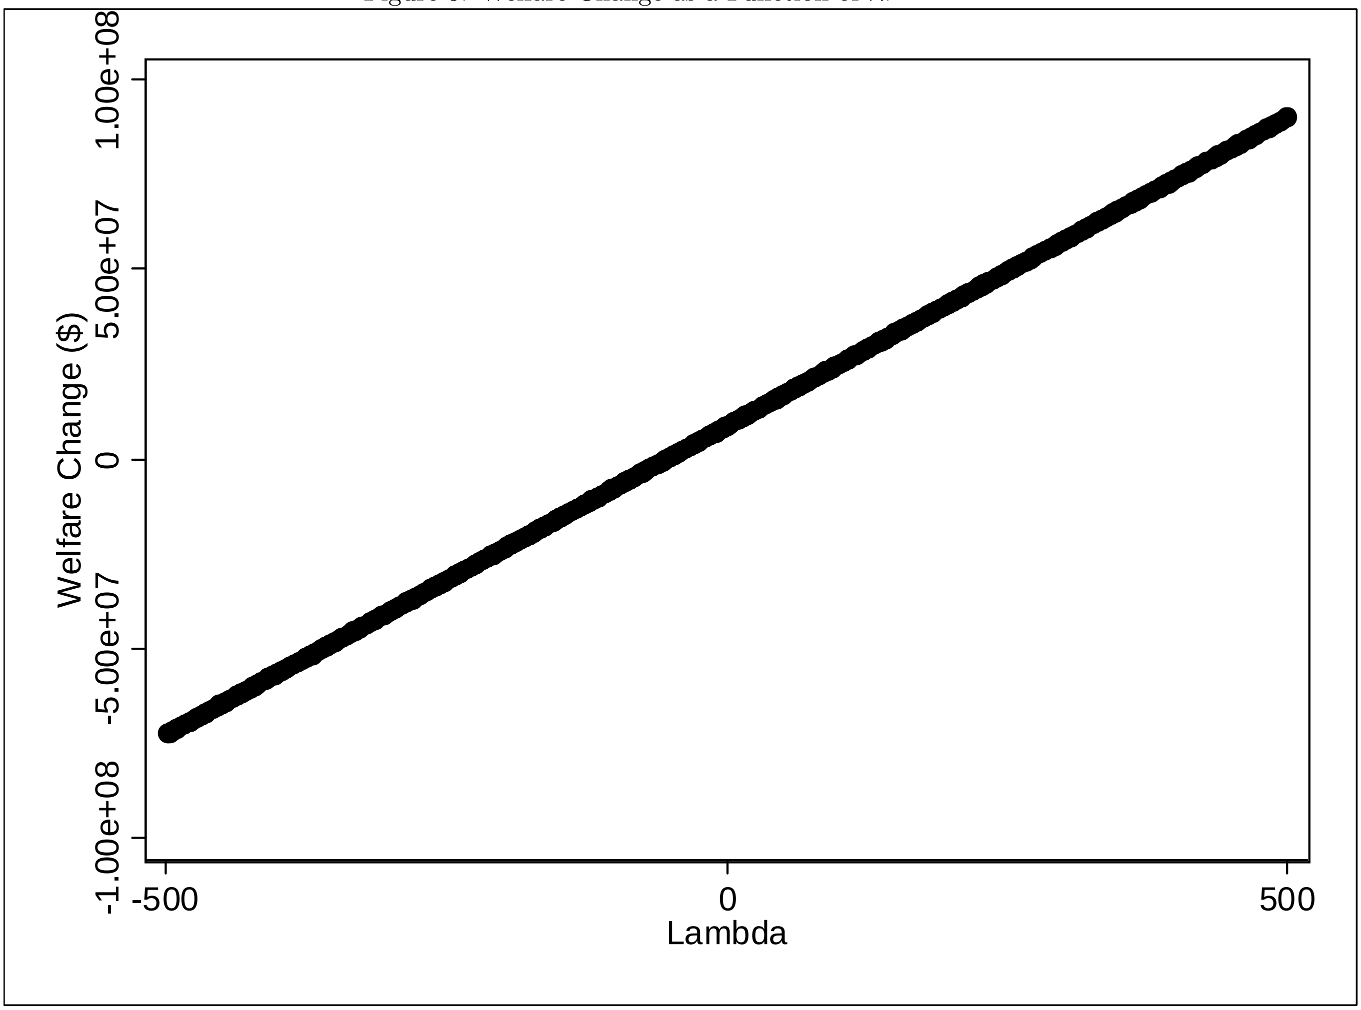

Figure 6: Welfare Change as a Function of (θ):$\lambda = -55$.

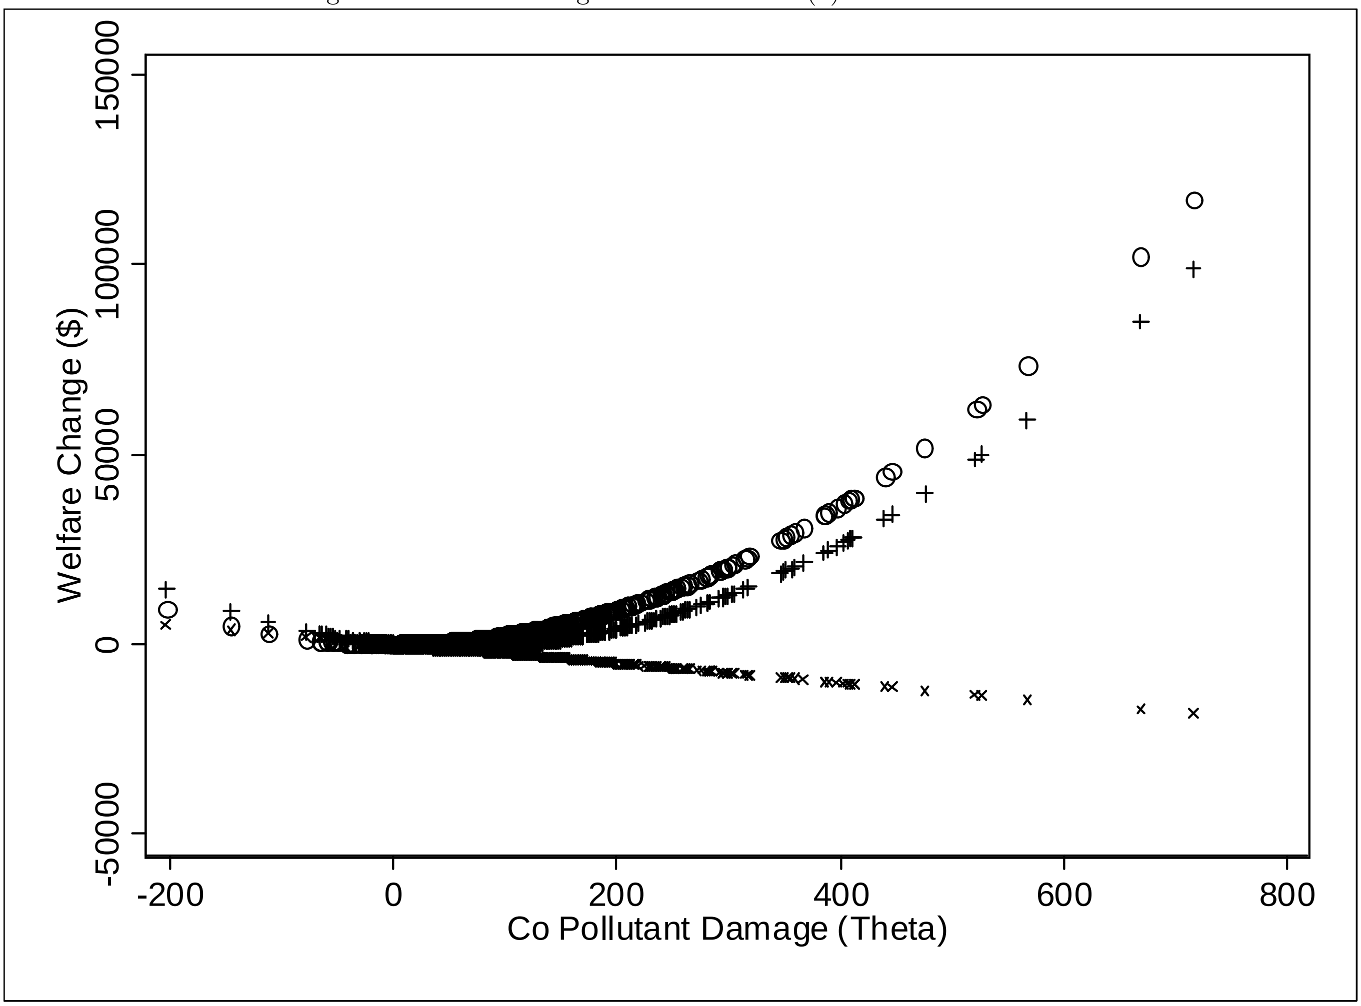

References

[1] Bollen, J., B.Guay, S. Jamet, J. Corfee-Morlot. 2009. "Co-Benefits of Climate Change Mitigation Policies: Literature Review and New Results", OECD Economics Department Working Papers, No. 693, OECD publishing, © OECD.

[2] Bussolo, M., and D. O'Connor (2001), Clearing the Air in India: the Economics of Climate Policy with Ancillary Benefits. Responding to Local and Global Environmental Challenges, OECD Development Centre, Paris: 61 pp.

[3] Cao, J., M.S. Ho, D.W. Jorgensen. 2008. "Co-Benefits of Greenhouse Gas Mitigation Policies in China: An Integrated Top-Down and Bottom-Up Modeling Analysis." Environment for Development Discussion Paper Series. EfD DP 08-10.

[4] Garbaccio, R.F., M.S. Ho and D.W. Jorgenson,1999. "Controlling Carbon Emissions in China", *Environment and Development Economics*, 4(4): 493-518.

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

[6] Krupnick, A., D. Burtraw, A. Markandya, 2000. "The Ancilary Benefits and Costs of Climate Change Mitigation: A Conceptual Framework. Ancillary Benefits and Costs of Greenhouse Gas Mitigation", OECD. Paris, OECD: pp.53-94.

[7] Mendelsohn, R.O. N.Seo. 2008. "Measuring Impacts and Adaptation to Climate Change: A Structural Ricardian Model of African Livestock Management" Agricultural Economics 38: 150-165.

[8] Muller, N.Z., R.O. Mendelsohn. 2009. "Efficient Pollution Regulation: Getting the Prices Right." *The American Economic Review.* 99(5): 1714 - 1739.

[9] Muller, N.Z., D.Q. Tong, R.O. Mendelsohn. 2009. "Regulating NO_x and SO_2 in Atlanta." *The B.E. Journal of Economic Analysis & Policy: Vol. 9 : Iss. 2 (Contributions), Article 3.*

[10] Nordhaus, W.D. 1982. "How Fast Should We Graze the Global Commons?", *American Economic Review.* 72(2): 242-246.

[11] Nordhaus, W.D. 1992. "An Optimal Transition Path for Controlling Greenhouse Gases." *Science.* 258(5086):1315.

[12] O'Connor, D., F. Zhai, K. Aunan, T. Berntsen and H. Vennemo, 2003. "Agricultural and Human Health Impacts of Climate Policy in China: A General Equilibrium Analysis with Special Reference to Guandong." Responding to Local and Global Environmental Challenges, O.D. Centre, Paris, OECD.

[13] Smith,K. R., E. Haigler. 2008. "Co-Benefits of Climate Mitigation and Health Protection in Energy Systems: Scoping Methods." *Annual Review of Public Health*, 29:11-25. (doi: 10.1146/annurev.publhealth.29.020907.090759)

[14] Tol, R.S.J. 2008. "The Social Cost of Carbon: Trends, Outliers, and Catastrophes." Economics: The Open-Access, Open-Assessment E-Journal, Vol. 2, 2008-25.

[15] USEPA, 2007. "eGrid Year 2005 Summary Tables." http://www.epa.gov/cleanenergy/energy-resources/egrid/index.html

[16] Weitzman, M.L. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *The Review of Economics and Statistics.* 91(1):1-19.