NBER WORKING PAPER SERIES

DISTRIBUTIONAL WEIGHTS IN ECONOMIC ANALYSIS.

Robert W. Hahn
Nicholas Z. Muller

Working Paper 31475
http://www.nber.org/papers/w31475

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2023, revised August 2024

The authors thank seminar participants at Georgetown University, Matthew Adler, Scott Farrow, Joel Franklin, Al McGartland, Brian Prest, and Nathan Hendren for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Distributional Weights in Economic Analysis.
Robert W. Hahn and Nicholas Z. Muller
NBER Working Paper No. 31475
July 2023, revised August 2024
JEL No. H41,Q51,Q52,Q58

ABSTRACT

The use of distributional weights in economic analysis is receiving increasing attention in both research and policy circles. This paper examines the extent to which distributional weights affect economic analysis of public good provision. First, we present a modelling framework with distributional weights that allows for marginal benefits and costs of the public good to differ across groups and regions. We show how the provision of a pure public good and the cost of providing that good varies with distributional weights different from unity. Second, we analyze distributional weights in conjunction with the value of a statistical life (VSL). We compare the use of an average VSL with group-specific VSLs adjusted for income. We show when using an average VSL will increase or decrease optimal public goods provision relative to group-specific VSLs for given distributional weights. We also identify conditions under which a low-income group prefers using an average VSL to group-specific VSLs. Third, we argue there is a "dictator paradox" that arises using conventional distributional weights that have been suggested by the U.S. government, where policy outcomes are essentially driven by a small number of very low-income individuals. Fourth, we apply our framework to a policy proposal to lower the ambient standard for fine particulate matter. We find that the level of aggregation for data analysis has a large impact on the magnitude and sign of net benefits. Moreover, the policy could fail a benefit-cost test when distributional weights are applied, but pass a traditional benefit-cost test.

Robert W. Hahn
University of Oxford
and Carnegie-Mellon University
robert.hahn@smithschool.ox.ac.uk

Nicholas Z. Muller
Department of Engineering, and Public Policy
Tepper School of Business
Carnegie Mellon University
4215 Tepper Quad
5000 Forbes Avenue
Pittsburgh, PA 15213
and NBER
nicholas.muller74@gmail.com

1. Introduction

The analysis of inequality in the distribution of economic resources across groups and nations has become an important research area in economics (Piketty and Saez, 2003; 2006; Piketty, Saez, and Zucman, 2018; Piketty, 2022). Concerns about how best to assess distributional outcomes from public policies, such as those related to race and income, are also receiving more attention by policy makers (The White House, 2021). One way to accomplish this is to conduct an economic impact analysis. This typically entails analyzing and reporting the benefits and costs that accrue to different groups in a population affected by a policy (USEPA, 2014). Another approach suggested by economists to measure the impact of policies on different groups is to introduce distributional weights (Adler, 2016). Conceptually, this approach acknowledges that a dollar (or other currency) of benefits *or costs* does not have the same value to all members of society. The distributional weights embody these differences in value. The weights are typically based on the diminishing marginal utility of income, such that benefits and costs are given higher weights when incurred by lower income groups. As evidence of the pressing policy-relevance of this topic, the U.S. Office of Management and Budget recently suggested that federal benefit-cost analyses could use distributional weights that are based on income (The White House, 2023, Sunstein, 2024).

This paper extends the literature on distributional weights with the goal of informing the practice of policy evaluation. It makes four contributions. First, we present a tractable model with distributional weights that allows for marginal benefits and costs to differ across groups and regions. The first order conditions for maximizing welfare are derived and compared with the case where distributional weights are unity, which is the typical case used in benefit-cost analysis. We show that Samuelson's analysis of pure public goods is a special case in our framework (Samuelson, 1954), as are other settings in which marginal benefits and costs may differ by region, but where distributional weights are unity (Muller and Mendelsohn, 2009). We also show how the provision of a pure public good and the cost of providing that good varies with distributional weights different from unity. We suggest this could have implications for the design of market mechanisms to internalize externalities.

Second, we analyze the use of distributional weights in conjunction with the value of a statistical life (VSL). The VSL is a key parameter used in analyses of government policy that calculate the costs and benefits of policies where mortality risks are affected (Robinson, 2007; USEPA, 1999; 2011; Viscusi, 2018). This value measures the tradeoff between fatality risk and income, and it has been shown to vary systematically with income levels (Viscusi, 2017). We compare the use of an average VSL, often used in the design of government policy, with the use of group-specific VSLs adjusted for income. We show how to determine when using an average VSL will increase or decrease optimal public goods provision relative to the case of group-specific VSLs for a given set of distributional weights. Much of the motivation for distributional analyses stems from concern about whether policies reflect the preferences of low-income groups. Motivated by the frequent use of average VSLs in practice, we identify conditions under which a low-income group would prefer using an average VSL to group-specific VSLs in a setting in which their mortality risks are valued less than the mortality risks faced by a high-income group. Our result depends critically on the fraction of the costs that the low-income group bears in the provision of the public good. This analysis of the VSL is also applicable to any willingness-to-pay based determinant of the benefits or costs of public good provision, provided the willingness-to-pay varies with race, income, or other drivers of equity concerns. More generally, our analysis highlights the importance of considering costs as well as benefits in attempting to address equity concerns. We show that when using distributional weights that differ across groups, the allocation of costs will generally affect total welfare. We suggest that this may have implications for how decisionmakers and analysts frame equity issues, including concerns about environmental equity. In particular, proponents of environmental equity or justice may want to consider benefits *and* costs to specific groups within the population, as opposed to simply focusing on differences in environmental quality (Cecot and Hahn, 2023).

Third, we use our theoretical framework to introduce what we call the "dictator paradox" which argues that the application of conventional distributional weights recommended by governments and researchers may yield policy outcomes that are driven by a very small number of very low-income individuals. This occurs because as household income approaches zero, the associated distributional weight approaches infinity. We refer to this as a paradox because distributional analysis is motivated by equity concerns. Yet, a small number of low-income individuals essentially become dictators of policy outcomes. We suggest that this paradox has implications for a wide range of policy issues where distributional weights may be used. Our focus on extreme weights has immediate policy relevance. Recently, the Federal Emergency Management Agency (FEMA) issued guidance on the use of distributional weights in certain areas where it conducts BCA (FEMA, 2024). In light of the extreme distributional weights that occur in the tails of the income distribution, FEMA recommends truncating weights. For example, all benefits and costs accruing to households under the 5^{th} percentile are weighted as if their income were equal to households *at the 5^{th} percentile.* While this is one approach to resolving the dictator paradox, it raises serious concerns about how federal agencies choose to evaluate welfare among the lowest income households.

Fourth, we use the theoretical framework to estimate the benefits and costs associated with a recent proposal by the U.S. Environmental Protection Agency (USEPA) to lower the National Ambient Air Quality Standard for fine particulate matter ($PM_{2.5}$) from 12 ug/m^3 to 10 ug/m^3. We do so with and without distributional weights. We consider cases in which weights are only applied to benefits and cases where weights are applied to both benefits and costs. Using raw household level income data and households grouped according to income percentiles, we test how the level of aggregation affects our conclusions on aggregate benefits and net benefits. This is especially important when using distributional weights because there is no theoretical guidance on group assignment and very low-income households have very large distributional weights which may disproportionally influence the outcome of policy analysis.

Our empirical policy analysis yields three key conclusions. First, using household level income data along with conventionally defined weights, we find that weighted net benefits go to a small fraction of the population. Sixty-six percent of the estimated net benefits accrue to the poorest 0.01 percent of households. This occurs in part because these very low-income households have distributional weights that are very large relative to the average income, and their per-capita net benefits are also higher than the average. Second, increasing the level of aggregation can increase or decrease net benefits when weights are applied. In our example, explained in more detail in section 4, this is due to changes in risk levels, changes in how risks are valued, and changes in the distributional weights that are applied to different groups. We also note that the level of aggregation can actually change the sign of net benefits for the case in which distributional weights and income-adjusted VSLs are used. Our findings raise an important and broadly relevant issue for policy analysis: if the level of aggregation can have a dramatic effect on net benefits with distributional weights, policy makers at the very least need to be aware of this fact. Third, assigning costs a weight of one, but using distributional weights for benefits yields net benefits that are higher than when distributional weights are applied to both benefits and costs. The reason relates to the dictator paradox. A small number of low-income individuals have their very high weight on cost reduced from some number substantially greater than one to one, and this drives the overall increase in net benefits.

The theory we present builds on three key areas of research. First, we extend a model by Montgomery (1972), which used a general equilibrium framework to assess the cost-effectiveness of different policy instruments. Montgomery demonstrated that, in principle, pollution licenses could be cost-effective. His model includes a mapping between emission sources and environmental quality receptors, which we also use in developing our theory. In contrast to Montgomery, we introduce distributional weights and focus on social welfare maximization rather than cost-effectiveness.

A second literature that we build on highlights the importance of distributional weights. Adler (2016) summarizes many of the theoretical contributions in this area. Early research that discussed distributional weights includes Meade (1955), Dasgupta and Pearce (1972) and Dasgupta, Sen, and Marglin (1972) and Harberger (1978). The UK Green Book explains how distributional weights could be used in analyzing government projects, programs, and regulations. It notes that, “in principle, each monetary cost and benefit should be weighted according to the relative prosperity of those receiving the benefit or bearing the cost.” (HM Treasury, 2003, p. 92).

A third literature we address attempts to link discussions of different VSLs with welfare. Closest to the spirit of our analysis of the VSL is Banzhaf (2011; 2023). Banzhaf compares the use of an average VSL with a group-specific VSL using unit distributional weights. He finds that the preferences for specific policies may shift depending on which set of weights are used. Hemel (2022) analyzes a regulation that would require new vehicles to have rearview cameras. He examines the welfare implications of using different VSLs for low and high-income groups, as well as the possibility of using different distributional weights. Sunstein (2023) also analyzes the welfare implications of using a VSL and argues that knowing the benefit and cost incidence is important. In contrast to these authors, we present a formal model of optimizing choice, we apply distributional weights to both benefits and costs, and we allow the distributional weights and the benefit and cost functions to vary. Our flexible modeling approach facilitates an analysis of how consumer choices would vary with different distributional weights and different VSL choices.

Our empirical analysis relies on methods that have been applied in both the environmental economics literature and in federal policy analyses. The link between exposure to air pollution, elevated mortality risk, and monetary benefits used here was applied by Muller and Mendelsohn (2009) to evaluate an efficient policy design for fossil fuel-fired power plants. Muller, Mendelsohn, and Nordhaus (2011) applied these tools to attribute damages from air pollution to individual industries and sectors in the U.S. economy. Tschofen et al., (2019) applied these methods to assess sectoral level air pollution damages. Regulatory Impact Analyses (RIAs) conducted by the U.S. Environmental Protection Agency for fine particulate matter (the subject of our application) employ the same approach to benefit estimation as we invoke here (USEPA, 2022).

The paper proceeds as follows: Section 2 presents the theoretical framework and discusses how it relates to global and local public goods. Section 3 provides an application to the VSL. It explores the impact of weighting individual VSLs and whether low-income groups and individuals would prefer using an average VSL to their group-specific VSL. Section 4 presents our empirical analysis of a proposal to regulate fine particulate matter. Finally, section 5 concludes and identifies areas for future research.

2. Theoretical model

We first present a general model of the benefits and costs of public goods provision with distributional weights. We consider the case of local public goods and pure public goods with and without distributional weights (Muller and Mendelsohn, 2009; Samuelson, 1954). Weights and the level of the public good consumed are allowed to vary by group. A group may refer to an income group, a racial group, or any other group. Typically, a group, or groups, will be associated with a region that experiences some kind of externality, such as pollution. For purposes of the theory, groups are assumed to be mutually exclusive and exhaustive.

In this formulation, a planner is assumed to maximize the difference between the weighted sum of benefits and the weighted sum of costs to different groups of consumers. We first consider the weighted benefits to groups and then the weighted costs to groups. There are assumed to be m groups, indexed by j. Group j is assigned a distributional weight, w_j, which is assumed to be constant and positive.[1] In the standard benefit-cost

[1] The distributional weight needs to be greater than or equal to zero. One economic interpretation for this weight could be the product of the marginal utility of income for group *i* times the weight that the decision maker puts on an extra unit of utility for group *i*. Note that in the interest of simplicity we assume that the distributional weight is constant. To the extent that weights vary with income, and income changes with the introduction of the public good, this would induce variation in the weights that could be taken into account.

evaluation, the distributional weights are set to one for all groups. Each group has a benefit function $B_j(q_j)$ where q_j represents the level of the public good consumed by group j. The public good can be local or global (e.g., environmental quality). Marginal benefit functions are assumed to be positive and non-increasing ($B_j' > 0, B_j'' \leq 0$ for all groups). The weighted benefits for a group j are given by $w_j B_j(q_j)$. Summing over all groups gives aggregate weighted benefits of a given level of the public good: $\sum_{j=1}^{m} w_j B_j(q_j)$.

To identify the weighted costs to different groups, we need to introduce the costs to firms and translate that into costs to groups of consumers in different regions. For concreteness, assume n firms produce a single type of emissions as a byproduct of their production.[2] The cost function for emissions reductions for firm i is given by $C_i(r_i)$, where r_i represents emission reductions. The cost function could be associated with a multi-product or single product firm.[3] The marginal cost of emission reductions is assumed to be greater than zero and non-decreasing in r_i, ($C_i' > 0, C_i'' \geq 0$ for all i). When a firm reduces emissions by r_i, the impact of its reduction on group j is given by $r_i h_{ij}$. Therefore, h_{ij} is a transfer coefficient that defines the impact of one unit of a reduction in emissions from firm i on environmental quality for group j.

A critical Issue in the subsequent analysis is the economic incidence of costs, or how the cost of emission reductions by firms is allocated across different consumer groups. This factor is often not given adequate attention in policy discussions pertaining to distributional outcomes (Banzhaf, 2023; Hemel, 2022). We define p_{ij} as the fraction of firm i abatement costs borne by group j. Furthermore, we assume these abatement costs are fully allocated across groups, so that $\sum_{j=1}^{m} p_{ij} = 1$ for $i = 1, \ldots, n$.[4]

We can now derive the costs of a particular policy. The unweighted cost to group j is $\sum_{i=1}^{n} p_{ij} C_i(r_i)$, which is the fraction of firm i's cost incurred by group j multiplied by firm i's cost and summed across all polluting firms. In order to derive the total weighted cost, we aggregate over all firms and apply the groups weights, yielding: $\sum_{j=1}^{m} w_j \sum_{i=1}^{n} p_{ij} C_i(r_i)$.

Subtracting costs from benefits and simplifying yields the weighted sum of net benefits over groups:

$$\sum_{j=1}^{m} w_j \{B_j(q_j) - \sum_{i=1}^{n} p_{ij} C_i(r_i)\} \quad (1)$$

[2] This assumption can be relaxed to allow for multiple types of emissions, and the same analytical framework can be applied.
[3] Individual firm costs are presumed to represent the loss in firm surplus from reducing emissions. Similarly, the benefit function for a group is presumed to represent the gain in total surplus for the group.
[4] In order to focus on the issue of cost allocation across groups, we do not consider the inefficiencies associated with raising revenue from various groups. See, for example, Mirrlees (1971) and more recently Hendren and Sprung-Keyser (2020).

The planner in this model chooses the level of emissions reduction, r_i, for each firm, to maximize this weighted sum. The public good for group j can be connected to the level of emission reduction by noting that $q_j = \sum_{i=1}^{n} r_i h_{ij}$. This says that the abatement level for group j is determined by the emission reductions by each firm i, multiplied by the corresponding transfer coefficient h_{ij}, and summed over all firms. Substitution into (1) yields the following maximization problem:

$$\underset{r_1,\dots,r_n}{\text{Maximize}} \sum_{j=1}^{m} w_j \{B_j \left(\sum_{k=1}^{n} r_k h_{kj}\right) - \sum_{i=1}^{n} p_{ij} C_i (r_i)\} \tag{2}$$

for $i = 1, \dots, n$.[5]

Differentiating the preceding expression with respect to r_i yields:

$$\sum_{j=1}^{m} w_j \{B_j{}' h_{ij} - p_{ij} C_i{}' (r_i)\} = 0 \tag{3}$$

for $i = 1, \dots, n$.

Equation (3) says the weighted sum of marginal net benefits must equal zero (from reducing a unit of emissions at the i th firm). Alternatively, the weighted sum of marginal benefits over j groups equals the weighted sum of marginal costs borne by different groups (again from reducing a unit of emissions at the i th firm).[6]

2.1 Local public goods

Consider the local public goods case where abatement and distributional weights can differ across groups. This is a generalization of Muller and Mendelsohn (2009) and Mendelsohn (1986), who implicitly assume that distributional weights are one for each group. Even in the case where distributional weights are one, the marginal cost of emissions reduction will, in general, *not* be equal across all firms because the h_{ij} vector for a given firm i can differ across firms. In that sense, the optimal plan for reducing emissions may not be cost effective. This is because implementing the cost-effective plan would not maximize net benefits to society. Intuitively, when benefits of emission reductions can vary by group, it may no longer be socially optimal to achieve these reductions in a manner that simply minimizes the cost of achieving a given emissions reduction target. We will explore the issue of cost minimization in more depth below in the context of pure public goods.

2.2 Pure public goods

In this framework, the pure public good case (or global public good case) can be modeled as an arbitrary number of groups that experience the same level of abatement.

[5] We use the index k in the summation for the benefit function to make the differentiation with respect to r_i transparent when applying the chain rule.

[6] In what follows, we assume that the second order condition for an interior maximum are satisfied, which will be the case given our assumptions about marginal benefit and marginal cost curves.

Samuelson (1954) considered the case where $w_j = 1$ for all j (*i.e.*, weights are unity across groups) and $h_{ij} = 1$ for all i, j (*i.e.*, transfer coefficients are set equal to 1). The level of abatement is given by $q = \sum_{i=1}^{n} r_i$ and it is the same for all groups.

Rewriting (3) with these assumptions yields $\sum_{j=1}^{m}\{B_j{}'(q) - p_{ij}C_i{}'(r_i)\} = 0$ for $i = 1, \ldots, n$. Simplification yields:

$\sum_{j=1}^{m} B_j{}'(q) = C_i'(r_i)$ for $i = 1, \ldots, n.$[7] (4)

We will refer to this as the "unit weight" case which assumes that distributional weights are unity.[8] The first order conditions defined by (4) say at the optimum, the sum of marginal benefits for all groups, which is a function of total abatement, equals the marginal cost of abatement for each firm. It is a version of the Samuelson first order condition for a pure public good.[9] Note that because marginal costs of abatement are equated across firms, the optimal level of abatement defined by (4) can be represented by a point on the aggregate marginal cost curve, defined as C'.

The generalization with arbitrary welfare weights is given by:

$\sum_{j=1}^{m} w_j\{B_j{}'(q) - p_{ij}C_i{}'(r_i)\} = 0$ for $i = 1, \ldots, n.$ (5)

We will refer to this as the "general" case because weights can vary, whereas in the original Samuelson formulation, the weights are (implicitly) restricted to unity. Equation (5) says that a weighted sum of marginal benefit curves equals a weighted sum of marginal cost curves.

If weights are not equal, marginal costs will generally differ across firms at the optimum and it may not be optimal for the planner to achieve a cost minimizing reduction of emissions. At the optimum, the weighted marginal benefits, $\sum_{j=1}^{m} w_j B_j{}'(q)$, will be the same across firms. However, the marginal costs will not, in general, be the same because the weight on marginal costs can differ across firms. Thus, at the optimum, there is no guarantee that marginal costs will be equated across firms. A similar point was presented in an earlier paper by Chichilnisky and Heal (1994), which explored optimal abatement across countries.

This has implications for market design. The cost-minimizing solution in terms of emissions will not, in general, be the optimal solution because of the different weights on the cost function.[10] In general, the planner may need to choose between a design that

[7] The derivation relies on the assumption that costs are fully allocated, so that $\sum_{j=1}^{m} p_{ij} = 1$ *for* $i = 1, \ldots, n$.

[8] The weights only need to be constant and positive to get the same result here in terms of public good provision.

[9] Graphically, the aggregate marginal cost curve is the horizontal aggregation of the marginal cost functions for all firms.

[10] Consider an example with two groups and two firms: $w_1 = 1$, $w_2 = 2$, $p_{11} = 1$, $p_{12} = 0$, $p_{21} = 0$, and $p_{21} = 1$. The groups have different weights and group 1 pays for firm 1's abatement and group 2 pays for firm 2's abatement. In this case the weight on $C_1{}'$ is 1*1 = 1, and on $C_2{}'$ is 2*1 = 2. The point is that the weighted

could achieve reductions in emissions in a cost-effective manner and a design that would take account of the impact of emissions on overall social welfare that is implied by the distributional weights. The introduction of distributional weights may, thus, complicate policy design.

2.2.1 Distributional weights and provision of pure public goods

In general, the optimal amount of the public good will not be the same when different distributional weights are used across groups (compared to when the weights are the same). We wish to understand how the provision of the public good (*e.g.*, abatement) will vary when distributional weights vary for different consumer groups. We will assume that each group pays a fixed fraction of total abatement costs, so group j pays p_j.

Using this assumption, we can solve (5) for the marginal cost to obtain:

$$C_i' = \frac{\sum_{j=1}^{m} w_j B_j'}{\sum_{k=1}^{m} w_k p_k} \text{ for } i = 1, \dots, n. \qquad (6)$$

Note that, with the simplifying assumption regarding cost shares, neither the numerator nor denominator in equation (6) depend on i. Thus, C_i' are equal for all i, which means that the marginal cost of reducing emissions is equal across all firms, and hence equal to a point on the aggregate marginal cost curve, C'. This point is worth noting. Efficiency requires that firms have the same marginal cost of emissions reductions even in the case where distributional weights differ from unity.

The intuition for the result in equation (6) is straightforward. Suppose the planner does not minimize costs. Then she could increase welfare by minimizing costs. By reducing overall costs for a given level of abatement, the planner raises the overall weighted welfare for each group, and thus increases the value of the welfare function. Thus, the planner must select the cost minimizing level of abatement to maximize welfare. This cost minimization implies that marginal costs will be equal across firms, given the assumptions of the model. The result about the independence of efficiency and distribution has been noted by many scholars in different contexts (Coase, 1960; Bergstrom and Cornes, 1983 and Montgomery, 1972).

While marginal costs will be equated even in the presence of non-uniform weights in this example, the level of abatement can change with the introduction of these weights. Equation (6) also says that whether abatement goes up or down when incorporating nonunitary weights depends on whether a *weighted average* of the sum of the marginal benefit curves goes up or down compared to the unit weight case. This direction of change does *not* depend on the marginal cost function. Rather it depends on how the marginal

marginal benefits across the two firms are the same, but the weights on the marginal cost functions differ, yielding a result that would not be cost-effective in terms of emissions. If the groups were associated with different regions, that would mean that the two regions had different marginal costs of abatement, which is one motivation for implementing cap-and-trade regimes for climate change (Stavins, 2022).

benefit functions adjusted by their weights compare (for the case of unit weights and the general case) and it depends on the cost shares.[11]

Another way of interpreting the equilibrium condition represented by equation (6) is to note that $C' \sum_{k=1}^{m} w_k p_k = \sum_{j=1}^{m} w_j B_j'$. Viewed in this way the equilibrium is determined by the intersection of a weighted marginal cost curve, where the weight is $\sum_{k=1}^{m} w_k p_k$, and a weighted sum of marginal benefit curves, where the weight on the j th curve is w_j . Both interpretations of the equilibrium can be illustrated graphically.[12]

Leveraging the equilibrium condition in (6), we wish to compare the optimal provision of the pure public good when distributional weights are arbitrary with optimal provision when distributional weights are set to one. Let q_w^* be the optimal abatement with arbitrary distributional weights, w_j; and q^* be the optimal abatement with unit weights.

Suppose

$$\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \frac{\sum_{j=1}^{m} B_j'(q_w^*)}{\sum_{k=1}^{m} p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$$

where the equality follows because $\sum_{k=1}^{m} p_k = 1$. The above expression says the weighted marginal benefit at q_w^* exceeds the unit weight marginal benefit at q_w^*. Reducing q from q_w^* will decrease C' (if the marginal cost curve is upward sloping) or increase the sum of marginal benefits (if that curve is downward sloping). Thus, q must be reduced from q_w^* to satisfy the unit weights first order condition, and it follows that $q_w^* > q^*$.

This yields the following result:

<u>Proposition 1</u>: Assume the aggregate marginal cost curve is upward sloping and/or the sum of marginal benefits curve for all groups is downward sloping.[13] Then:

$$q_w^* \,(>)(=)(<) q^* \; iff \; \frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} (>)(=)(<) \sum_{j=1}^{m} B_k'(q_w^*) \tag{7}$$

Proof: See appendix.

This proposition shows that the distributional weights and the cost shares are both crucial for determining whether abatement goes up or down with the introduction of non-unit distributional weights. Further, this comparison does not depend on the marginal cost function.

---

[11] The marginal benefit function in the unit weight case is the vertical summation of the marginal benefit functions across groups, following Samuelson. In the weighted case, the vertical summation must be weighted in accord with the weights on the right-hand side of equation (6).

[12] Figure 1 in the next section illustrates this graphical approach using an application to the VSL.

[13] We also assume a unique interior equilibrium exists, which will be satisfied for reasonable choices of the model parameters, given the assumed form of the marginal benefit and marginal cost functions.

3. Model application to the VSL

In this section we demonstrate the policy relevance of the framework we developed in the previous section by applying it to the VSL.[14] The following analysis of the VSL is, in fact, applicable to any willingness-to-pay (WTP) based determinant of the benefits of public good provision, provided the WTP varies with race, income, or other drivers of equity concerns. Examples include ecosystem services, endangered species preservation, or morbidity risk reductions.

This section begins by presenting a simplified version of the pure public good model to explore how the government's use of an average VSL affects overall welfare and the welfare of specific groups. We then demonstrate how application of an income elasticity to impute individual WTP interacts with distributional weights in applied policy analysis. We conclude by introducing the Dictator Paradox, which argues that policy analysis using distributional weights will often be driven by a very small number of households or individuals with very low income.

We consider whether the use of an average VSL may be *preferred* by a group that we refer to as low-income, even when use of the average VSL is not preferred by society. This result depends on the share of the costs for the public good paid by the low-income group. The intuition is that the application of the average VSL to the mortality risk incurred by the low-income group attributes a higher value to such risks than their group-specific VSL. Provided costs borne by the low-income group are sufficiently small, this group enjoys higher net benefits. We underscore the importance of economic cost incidence in driving this result.

Because government agencies often use a VSL that does not differ by income, it is instructive to consider how this might affect welfare. To model this, we consider two population groups: group 1 is the low-income group and group 2 is the high-income group. Each group has a constant marginal benefit curve which incorporate different VSLs. For group 1, the marginal benefit curve is given by $B_1' = n_1 c$, where n_1 represents the number of people in group 1 and $c > 0$ represents the per capita marginal benefits for that group associated with each unit of the public good. For group 2, the marginal benefit curve is given by $B_2' = n_2 d$, where n_2 represents the number of people in group 2 and $d > c$ represents the per capita marginal benefits for group 2 associated with each unit of the public good. The high-income group thus has higher per capita marginal benefits than the low-income group for each unit of the public good that is provided. The motivation for this specification lies in the positive VSL-income elasticity (Viscusi and Aldy, 2003). The marginal cost curve is given by a linear upward sloping curve through the origin, which implies $C'(q) = bq$ and $b > 0$. We assume for simplicity there is a single firm.[15] The cost share for group 1 is given by $n_1 p_1 = p$ and for group 2 is given by $n_2 p_2 = 1 - p$, so that

[14] Note that we choose to focus here on the VSL because of its importance in federal benefit-cost analyses and because the VSL varies systematically with income. Our analysis in this section is sufficiently general to be applied to any such parameter.
[15] This assumption is easily relaxed to allow for multiple firms.

costs are fully allocated. In this case p_1 and p_2 can be interpreted as the per capita cost share for a member of group 1 and group 2, respectively.

We can link this formally to the use of VSLs by assuming there is a constant relationship between provision of the public good, q, and the number of statistical lives saved in each group.[16] The benefits for a unit reduction in emissions can be expressed as the product of some constant and the group-specific VSL. As above, the distributional weights are assumed to be positive. In this context, the social planner maximizes social welfare as shown below:[17]

$$\underset{q}{\text{Maximize}}\ w_1\left(n_1cq - p\frac{bq^2}{2}\right) + w_2\left(n_2dq - (1-p)\frac{bq^2}{2}\right) \quad (8)$$

We first note that, conditional on q, net benefits are declining in the cost share borne by the low-income group (denoted p), since $w_2 < w_1$. In contrast, with unit weights, cost shares do not affect net benefits since this amounts to reallocating equally-weighted cost between groups (holding total cost fixed). We demonstrate the importance of this point in section 4.2.3 below.

We consider the first order conditions for a maximum. Differentiating (8) with respect to q yields:

$$w_1(n_1c - pbq) + w_2(n_2d - (1-p)bq) = 0$$

This first order condition is a linear equation in q. We define the solution to this problem as q_t^*, where "t" denotes using "true," group-specific values for the VSL for both groups.[18] Solving for q_t^* yields:

$$q_t^* = \frac{w_1n_1c + w_2n_2d}{w_1p\,b + w_2(1-p)b} \quad (9)$$

Note that q_t^* is a function of both the weighted sum of marginal benefits and costs.

We next define the average per capita marginal benefit as $\frac{c+d}{2}$, which is a constant proportion of the average VSL. If we substitute the average VSL for the true per capita VSL for both groups, then the first order conditions for a maximum become:

$$w_1\left(n_1\left(\frac{c+d}{2}\right) - pbq\right) + w_2\left(n_2\left(\frac{c+d}{2}\right) - (1-p)bq\right) = 0$$

---

[16] The benefits functions also allow for a formulation in which the VSLs differ by group, and the risk per unit of the public good differs by group.

[17] This formulation assumes total costs are $C(q) = \frac{bq^2}{2}, B_1(q) = cq$ and $B_2(q) = dq$. This follows from integrating the marginal benefit functions and the marginal cost function from 0 to q.

[18] We use the word "true" to connote a calculation based on the actual willingness to pay of different groups.

The first-order condition using the average VSL is also a linear equation in q. We define the solution to this problem as q_a^*, where the "a" denotes using average values for the VSL for both groups. Solving for q_a^* yields:

$$q_a^* = \frac{(w_1 n_1 + w_2 n_2)(\frac{c+d}{2})}{w_1 p\, b + w_2 (1-p) b} \quad (10)$$

The denominators in equation (9) and equation (10) are the same and positive, but the numerators will, in general be, different. Under the special case when the distributional weights multiplied by population are equal (*i.e.*, $w_1 n_1 = w_2 n_2$), then $q_t^* = q_a^*$. To see this result let $\bar{w} = w_1 n_1 = w_2 n_2$. Then the numerator in (9) is $\bar{w}c + \bar{w}d$ and the numerator in (10) is $\bar{w}\left(\frac{c+d}{2}\right) + \bar{w}\left(\frac{c+d}{2}\right)$. These terms are clearly equal.[19] If, however, weighted marginal benefits are higher (lower) using the true value versus the average, then $q_t^* > (<) q_a^*$.

In the more general case where $n_1 \neq n_2$, we have the following proposition:

<u>Proposition 2</u>: $q_t^* (<)(=)(>) q_a^*$ iff $w_1 n_1 (>)(=)(<) w_2 n_2$.[20]

Proof: See appendix.

This proposition says that for the case of linear marginal benefit and marginal cost curves, the distributional weights and the size of the groups determine how the optimal public good provision changes with the introduction of an average VSL. Note that these results do not depend on the fraction of costs borne by group 1. As can be seen from the expression for q_t^* and q_a^*, these quantities are affected by changes in p, but the relationship between these two values is not.

In summary, this model provides a transparent way of determining the welfare implications of substituting an average VSL for a true VSL for low-income and high-income groups. We have assumed in this particular analysis that the distributional weights are group-specific. However, the model also enables analyses comparing the case with distributional weights set to unity and an average VSL with distributional weights set to their group-specific value with true VSLs.

3.1 Low-income group preferences and the average VSL

In this subsection, we raise the possibility that certain population groups may prefer policy calibrated to the population average VSL rather than their true VSL. We do so both because policymakers often employ the average VSL in policy design and evaluation and

---

[19]In the example where $n_1 = n_2 = 1$, the term for weighed marginal benefits is $\frac{w_1 c + w_2 d}{w_1 p + w_2 (1-p)}$ in the true case and $\frac{(w_1 + w_2)(\frac{c+d}{2})}{w_1 p + w_2 (1-p)}$ in the average VSL case. Viewing the numerator of those two terms reveals they will be equal if and only if $w_1 = w_2$. In this case they yield the same optimal level of the public good.

[20] We use the assumptions on the linear marginal benefit function and cost functions defined above.

because this exercise highlights the importance of economic cost incidence in determining welfare outcomes. In particular, we show that there may be situations in which the low-group group prefers an average VSL to the true VSL. This situation can arise, for example, when a low-income group shoulders a sufficiently low fraction of the costs. In this case, that group may prefer an average VSL that results in a higher level of abatement to using a true VSL that results in a lower level of abatement. In contrast, with a sufficiently high-cost share, the high-income group may prefer a lower level of abatement. We present this example below along with a graphical illustration.

We begin by defining the optimal level of the public good from group 1's point of view, (q^*). We then construct an example in which $q_t^* < q_a^* < q^*$. If this is the case, we will show that group 1 prefers q_a^* to q_t^*, that society prefers $q_t^* < q_a^*$, and that group two prefers q_t^* to q_a^*. For simplicity, assume the low-income group and the high-income group each consist of one agent so that $n_1 = n_2 = 1$.[21]

Group 1's maximization problem is:

$$\underset{q}{\text{Maximize}}\ w_1\left(n_1 cq - p\frac{bq^2}{2}\right)$$

Differentiating with respect to q yields $c - pbq = 0$ for the first order condition, which yields $q^* = \frac{c}{pb}$. From (10), we have $q_a^* = \frac{(w_1+w_2)[\frac{1}{2}(c+d)]}{w_1 p\, b + w_2(1-p)b}$. We wish to see if there is a $q_a^* < q^*$. From the formula for q^* and q_a^*, it follows that

$$q_a^* < q^* \text{ iff } \frac{(w_1+w_2)[\frac{1}{2}(c+d)]}{w_1 p + w_2(1-p)} < \frac{c}{p}. \qquad (11)$$

The key insight is that for p sufficiently small in (11), the right-hand side approaches infinity, while the left-hand side remains finite. The numerator stays constant and the denominator tends to w_2 as $p \to 0$. This shows that for p sufficiently small, we have $q_a^* < q^*$. To establish $q_t^* < q_a^*$, we assume $w_1 > w_2$ and $n_1 = n_2 = 1$, and apply Proposition 2.[22]

The intuition for this result is straightforward. Think of the limiting case in which p is zero, so that group 1 does not pay for the public good. In this case, group 1 prefers an infinite amount of the public good. Thus, it is possible to move group 1's preferred level of the public good beyond q_a^* for p sufficiently small. Stated another way, as the public good becomes cheaper from the standpoint of group 1, it wishes to have more of it.

[21] The example can easily be extended to an arbitrary number of members in each group, but we assume one member in each group to avoid notational clutter.

[22] More generally, for $q_t^* < q_a^*$ to hold, we require conditions related to both the weights and the population group sizes. First, we require that $w_1 > w_2$ which, if the weights reflect group-specific marginal utility of income, should hold for concave utility functions. Second, if $n_1 >= n_2$, then $q_t^* < q_a^*$. This is likely to hold writ large in societies with right-skewed income distributions. Finally, the only case in which $q_t^* > q_a^*$ occurs if $n_2/n_1 > w_1/w_2$.

We are now in a position to demonstrate the three claims noted above. First, under the conditions laid out above, group 1 prefers q_a^* to q_t^* because q_a^* is closer in distance to q^*, and group 1's objective function is a parabola that achieves its maximum at q^*. Second, society, or the social planner, prefers q_t^* to q_a^* for the same reason—that is, society's objective function is a parabola that achieves its maximum at q_t^* and $q_t^* < q_a^*$ by construction. Finally, at q_a^*, we know that group 1's welfare is higher than at q_t^*. Because total social welfare is lower at q_a^* than at q_t^*, this implies group 2's welfare is lower at q_a^* than at q_t^*. This demonstrates the third claim.

A graphical representation of this result is shown in Figure 1. Figure 1a shows the optimum values for the average and true levels of abatement when group 1 and group 2 are both included. The weighted marginal benefit function for the average VSL case is given by MB_a and the weighted marginal benefit function for the true VSL case is given by MB_t (with MB_t < MB_a). The weighted marginal cost curve is given by MC_w. The weights are given by the first order condition for a pure public good for the linear functions defined in this VSL example.[23]

The optimum for the true value case, q_t^*, occurs at the intersection of MB_t and MC_w; and the optimum for the average VSL value case, q_a^*, occurs at the intersection of MB_a and MC. Note that $q_t^* < q_a^*$. The social loss in moving from q_t^* to q_a^* is given by triangle A.[24]

Figure 1b shows why group 1 is better off with the optimum, q_a^*, which uses the average VSL, rather than the true optimum, q_t^*, which uses the true VSL The optimum value for group 1, q^* is defined by the intersection of its weighted marginal cost curve MC_1, and its weighted marginal benefit curve MB_1. Note that that $q_t^* < q_a^* < q^*$. The increase in weighted welfare for group 1 in moving from q_t^* to q_a^* is given by trapezoid B.[25]

This framework for comparing an average VSL with a true VSL (or willingness to pay) can be extended in a number of different ways. To illustrate, we consider three extensions. First, we consider the case of a pure public good in which the low-income group and the high-income group face different risk reductions per unit of abatement. This may stem from differential baseline risks that interact with pollution exposure (Spiller et al., 2021). Suppose, for the sake of illustration, that the risk reduction per person for the low-income group from a unit reduction in abatement were higher than the risk reduction per person for the high-income group. In this case, we can no longer assume that the marginal benefit per person for a unit of abatement was higher for the high-income group than the low-

[23] The specific values for the objective function are $w_1 = 2$, $w_2 = .5$, $b = 1$, $c = 1$, $d = 2$, $n_1 = 1$, $n_2 = 1$, and $p = .15$.

[24] This social loss is measured in weighted dollars because we are using the distributional weights.

[25] Sunstein (2023) explores the case when a low-income group could be better off, but does not use a formal model. He argues that costs are important, which this model supports. He also argues that a subsidy would likely make this group better off. His subsidy appears to be an in-kind transfer used to provide more abatement. If the subsidy reduced the price of the public good for the low-income group by reducing p instead, then Proposition 2 would apply.

income group (*i.e.*, $d > c$). If, in fact, it were the case that $d < c$, then the signs in Proposition 2 would be reversed, but the framework we have developed would still apply.

A second extension is to consider the case in which the transfer coefficients are different for the low-income group and the high-income group. This could arise if the two groups live in different neighborhoods. This is the case of local public goods. The analysis in this case is very similar to the first case presented in the preceding paragraph, except risk now varies by income class per unit of emission reduction because of the location of the groups rather than differences in health status. The same type of analysis would apply and would be dependent on whether the marginal benefit per person per unit of emission reduction would be higher for the low-income group than the high-income group.

Finally, consider a case in which we have two local public goods, one for the low-income group and one for the high-income group, both of which are provided separately by different firms (so the costs and benefits are separate). In this case we make the following two observations. First, if both groups are paying the full costs, each will always weakly prefer the true value of the VSL to the average value. The reason is that costs are fully internalized in these cases, so a group cannot do better, and may do worse, if it does not use its true benefit function (or VSL in this case). Second, if one group, say the low-income group, is not paying its full cost, then it may prefer an average VSL to a true VSL. The reasoning is similar to the reasoning provided in Proposition 2. As the price the group faces for the public good declines (*e.g.*, measured in terms of its cost share), the group will want more of the public good to increase its welfare. Fundamentally, this is because the costs to that group are not fully internalized.

3.2 Distributional Weights and The VSL Income Elasticity

We next address the comparison of average and individual, or true, VSLs from the perspective of empirical policy analysis. In an applied setting, it is essential to recognize that *individual* VSLs are not observable. Applied policy analysis uses sample or population average VSLs (USEPA, 1999; 2011). Estimation of individual weighted VSLs involves a two-step procedure. First, a sample or population average VSL is used in conjunction with an income elasticity to impute an individual VSL (Viscusi and Masterman, 2017). Second, distributional weights are applied to the imputed individual VSL to arrive at a weighted VSL. The analysis in this section demonstrates that this procedure yields an effective weight that differs from conventionally defined distributional weights that are based on knowing the true VSL (Acland and Greenberg, 2023, The White House, 2023).

We begin by estimating the VSL for individual i, denoted V_i. We use the VSL-income elasticity (ε), the average VSL ($\bar{V}$), individual $i's$ income (I_i), and the median income (I_m) in the following formula:

$$V_i = \left(\frac{I_i}{I_m}\right)^{\varepsilon} \bar{V}.^{26} \tag{12}$$

From Kaplow (2005) we can decompose ε as

$$\varepsilon = \eta + \rho \tag{13}$$

where

$\eta =$ the elasticity of utility with respect to consumption
$\rho =$ the coefficient of relative risk aversion or the elasticity of marginal utility with respect to consumption.

Kaplow (2005) assumes that the utility function is concave in consumption and that utility depends only on consumption and expenditures that reduce the individual's probability of death. The decomposition in (13) therefore invokes this assumption.

Using equation (13), we can rewrite equation (12) as:

$$V_i = \left(\frac{I_i}{I_m}\right)^{(\eta+\rho)} \bar{V} \tag{14}$$

Next, we apply a conventionally defined welfare weight, which adjusts V_i for the marginal utility of income for individual i relative to that of the median income (Acland and Greenberg, 2023). We further assume a constant relative risk aversion (CRRA) utility function for individual i. Employing the CRRA functional form[27] yields the following weight:

$$\varpi_i = \left(\frac{I_i}{I_m}\right)^{(-\rho)} \tag{15}$$

The second step in the estimation of individual, weighted VSLs is shown in (16). The multiplication of the imputed individual VSL by the welfare weight counteracts the effect of ρ in deriving the individual unweighted VSL in (14).

$$V_i\varpi_i = \{\left(\frac{I_i}{I_m}\right)^{(-\rho)} \left(\frac{I_i}{I_m}\right)^{(\eta+\rho)} \bar{V}\} = \left(\frac{I_i}{I_m}\right)^{(\eta)} \bar{V} \tag{16}$$

Equation (16) indicates that the weighted V_i depends on relative income and η.
Hence, the weighted *and income adjusted* VSL is no longer a function of the elasticity of the marginal utility of income. Rather it is a function of the rate at which total utility changes as consumption changes. This result obtains from the fact that the income adjustment in (12) and (14) already accounts for ρ. Thus, the welfare weight removes this effect and the resulting coefficient on the average VSL depends only on η. The expression in (16) differs

[26] In what follows, we will call V_i the "group-specific" or "true" VSL and apply it to households or larger groups based on their income.
[27] Here, we assume that income is allocated only to consumption, so we can write utility as a function of income.

from typical characterizations of welfare weights (Acland and Greenberg, 2023). And, as stated above, this two-step procedure is required because individual VSLs are not observable. The implication of (16) is that the weight depends on income levels both through the $\left(\frac{I_i}{I_m}\right)$ term and through η, since $\eta = \frac{I(1-\rho)}{I-I^{\rho}}$, under the assumption of CRRA utility. In contrast, ϖ_i only varies through the $\left(\frac{I_i}{I_m}\right)$ term.

While the literature commonly uses the income elasticity to estimate group-specific VSLs, and the literature commonly uses weights such as those in equation (15), prior research does not appear to have identified the relationship shown in (16).

3.3 Distributional Weights and the "Dictator Paradox"

As Hahn (2023) pointed out, a conventional assumption about the functional form of distributional weights can lead to peculiar policy outcomes. If we assume welfare weights are defined by equation (15), then as income for individual (i) approaches zero, $\varpi_i \to \infty$. Assuming there is only one such individual, this individual's preferences would dominate all others in term of weights, and thus they become a "dictator." This is paradoxical because, in an effort to increase the weight placed on low-income individuals to promote equity or justice, the distributional weights can create a kind of dictatorship, or at least a "tyranny of the minority." That is, a small group of people at the lower end of the income distribution could dominate policy choices.

To illustrate this idea, consider the planner's problem in equation (8), where there are two groups with different VSLs. Assume for simplicity that each group has one individual and one group has zero income and the other has non-zero income, so that $n_1 = n_2 = 1$, and $w_1 \to \infty$.

In this case, we can show that the optimal level of public good provision is $q_t^* = \frac{c}{pb}$.[28] This is also group 1's preferred outcome (not surprisingly since it is the "dictator"), but not group 2's.

The conclusion is that policy outcomes converge to those preferred by individuals with near-zero income levels because of their near infinite distributional weights. Hence, the very low-income individual effectively becomes a dictator, driving policy outcomes without regard for the preferences of other people. From an applied perspective, policy analysts are more likely to encounter this phenomenon when outcomes are tracked at the

[28] Recall that $q_t^* = \frac{w_1 n_1 c + w_2 n_2 d}{w_1 p\, b + w_2 (1-p) b}$. Using the simplifying assumptions, this yields $q_t^* = \frac{w_1 c}{w_1 p\, b + w_2(1-p)b} + \frac{w_2 d}{w_1 p\, b + w_2(1-p)b}$. Taking the limit as $w_1 \to \infty$ yields $q_t^* = \frac{c}{pb}$. To identify group 1's preferred outcome, call it q_1^*, differentiate the first term in (8) and set it to zero, which yields $q_1^* = \frac{c}{pb}$. Using the same reasoning and notation for group 2 yields $q_2^* = \frac{d}{(1-p)b}$. Also, note that when group 1 bears no cost ($p = 0$), an infinite amount of the public good would be supplied. A similar analysis applies when an average VSL is used because the structure of the problem is the same. In this case $q_a^* = \frac{c+d}{2pb}$.

individual level, rather than for large social groups. We explore the importance of aggregation in our empirical analysis in section 4.

3.4 Modelling framework strengths and limitations

The modeling framework developed here is quite flexible, and allows for many formulations that allow operationalization of various concepts of equity. For example, an important issue for policy makers is how to address "environmental equity." A full treatment of this issue is beyond the scope of this paper, but we wish to make two points. Recent efforts by the federal government to address environmental justice considerations focus exclusively on benefits (The White House, 2021). However, our work demonstrates the importance of considering costs borne by the low-income group as well as benefits. First, as the previous example shows, consideration of costs has fundamental implications for group welfare related to the use of an average or true VSL. Second, the model demonstrates that not considering cost could result in corner solutions where the welfare maximizing allocation that results from including costs is not achieved. For example, without considering costs, the optimum quantity demanded by group one in the example would be infinite. Our joint inclusion of distributional weights *and* cost shares brings these two insights to light in a manner that facilitates operationalizing a broader definition of equity in economic analysis.

The model also has some limitations as well. We consider two here. The first relates to the fact that it is developed in a partial setting, and a second relates to the inclusion of uncertainty. One limitation of the model is that it relies on a partial equilibrium framework that does not consider tradeoffs with other goods and leisure. In principle, this could be added to our model, but would add complexity. A second limitation is that prices are assumed to be given, and do not change with the level of the public good that is provided. While a general equilibrium framing may be more realistic for some public goods problems, we decided to opt for model simplicity using a partial equilibrium setting that is applied in many real-world applications.

A second issue relates to the inclusion of uncertainty. The basic problem is that the decision maker may have very limited information over the key parameters in the model. Uncertainty in benefits and costs has been analyzed by several scholars (Weitzman, 1974), and we do not have much to add to that discussion because we are not considering different policy instruments. Uncertainty in other key parameters, such as the distributional weights and the cost allocation shares, could be quite important. Under plausible independence assumptions, one can show that uncertainty over the weights and the cost shares would still lead to using a framework in which expected net benefits are maximized. The basic insight in this case would be a familiar one. As uncertainty in parameter estimates increase, the value of obtaining better (or perfect) information would tend to increase.

4. Application

We next demonstrate how the theoretical framework developed above could be applied to a real-world policy proposal. We analyze a proposed change to the regulation of fine particulate matter (or $PM_{2.5}$), a common air pollutant. We focus on fine particulate matter for three reasons. First, the data rich context enables empirical estimation of both benefits and costs. Second, air pollution generally, and $PM_{2.5}$ specifically, has large, and economically significant, consequences for social welfare in the U.S. economy (Muller, Mendelsohn, and Nordhaus, 2011; Tschofen et al., 2019). Third, the USEPA is currently considering a policy to lower the ambient limits for $PM_{2.5}$ (USEPA, 2023c). This work could help inform how regulators frame policy analysis. We begin by discussing our data sources and empirical methods before presenting the simulation design and empirical results.

4.1 Data and Methods

The empirical analysis employs data gathered from the following sources. Ambient readings of $PM_{2.5}$ are from the USEPA AQS databases (USEPA, 2023a). These data are from the year 2021. Individual and household level income data are obtained from the IPUMS datasets, which are derived from the U.S. Census Bureau's American Community Survey (IPUMS, 2023). The dataset used in our subsequent calculations comprise the intersection of the IPUMS and the AQS samples for the year 2021. Also important to the calculation of mortality risk from exposure to ambient $PM_{2.5}$ are data on baseline mortality risk which are provided by the Centers for Disease Control and Prevention's (CDC) Wonder databases (CDC, 2023). These data are county-by-age averages for 2021.

In order to calculate the monetary damages from exposure to ambient $PM_{2.5}$, we employ the methods used by the USEPA in their benefit-cost analyses for air pollution policy (USEPA, 1999; 2011). The mortality risk incurred by age-group (a), in county (j), in year (t) attributable to $PM_{2.5}$ exposure is shown in (17).

$$M_{a,j,t} = Pop_{a,j,t}\gamma_{a,j,t}\left(1 - \frac{1}{exp(\beta\Delta PM_{j,t})}\right) \quad (17)$$

where: $Pop_{a,j,t}$ = population of age-group (a), in county (j), year (t).

$\gamma_{a,j,t}$ = baseline mortality risk of age-group (a), in county (j), year (t).

β = statistically-estimated parameter from Krewski et al., (2009).

$\Delta PM_{j,t}$ = change in ambient $PM_{2.5}$ in county (j), year (t).

The functional form in (17), specifically the parenthetical term, is based on the epidemiological dose-response modeling in Krewski et al., (2009). This function is used by the USEPA and we invoke it here to enhance the policy relevance of our application. Monetary damage is the product of $M_{a,i,t}$ and the VSL, which we begin by applying uniformly across all exposed populations using the USEPA's preferred VSL. Specifically, the USEPA employs a VSL of $7.4 million expressed in year-2006 USD (denoted VSL_{2006} in (18) below) which we update to the 2021 currency year and 2021 median income levels (USEPA, 2023b).

The adjustments of the VSL for inflation uses the CPI (BLS, 2023) and for changes in real median income are shown in (18):

$$\bar{V}_t = (VSL_{2006})(CPI_{t,2006})\left(\frac{I_{t,m}}{I_{2006,m}}\right)^{\epsilon} \quad (18)$$

where: $I_{t,m}$ = real median income in year (t).
$I_{2006,m}$ = median income in 2006.
ε = VSL-income elasticity, set to 0.7.
$CPI_{t,2006}$ = CPI for year (t) with base year 2006.

This calculation yields the following uniform VSL for 2021: ($7.4 million)(1.32)(1.078) = $10.53 million.

One objective of the empirical analysis is to explore the influence of data aggregation on the benefit and cost estimates. We consider three "levels" of aggregation for the IPUMS income data: households, four income groups, and two income groups. Our sample consists of about 242,000 households. Thus, with four income groups, each group contains about 60,000 households. With two income groups, there are 121,000 households in each group. We adjust (18) using income data at each resolution for 2021 (IPUMS, 2023). For example, to estimate unweighted, income-adjusted household VSLs, we employ the approach in equation (19):

$$V_{j,t} = \left(\frac{I_{j,i,t}}{I_{m,t}}\right)^{\varepsilon} \bar{V}_t \quad (19)$$

where the $I_{j,i,t}$ term in the numerator reflects reported income for household (i), in county (j), in year (t).

The distributional weights are calculated using the functional form in (15), which is commonly used by researchers and policymakers (Acland and Greenberg, 2023; The White House, 2023).

$$\omega_{j,i,t} = \left(\frac{I_{j,i,t}}{I_{m,t}}\right)^{-\rho} \quad (20)$$

where: ρ = coefficient of relative risk aversion, set to 1.4.
$I_{j,t}$ = income for individual (j), time (t)
$I_{m,t}$ = median income time (t).

By multiplying (19) and (20) the weighted, individual VSLs embody the effective weights shown in (16). For group-specific weights with more than one individual, the numerator in (19) is replaced with average income for the group.

To estimate the cost of attaining the proposed 10 ug/m^3 NAAQS for $PM_{2.5}$, we employ existing cost estimates derived from prior benefit-cost analyses for $PM_{2.5}$ (USEPA, 1999). Specifically, USEPA (1999) reports the annual compliance costs associated with abatement of PM to meet the PM NAAQS in the years 2000 and 2010. These cost estimates and additional related figures are reported in Table A1. To estimate the cost per ug/m^3 reduced, we calculate the national average change in ambient $PM_{2.5}$ from 2010 to 2011 (USEPA, 2023a) and then divide total reported cost by the change in $PM_{2.5}$. Finally, we divide through by total population to derive a per-capita cost per ug/m^3 reduction. This figure, for costs reported in 2010, is \$120/(ug/m^3)/person. (See Table A1 in the appendix.) Household costs are calculated by multiplying the per-capita cost times the number of persons per household, as reported in IPUMS (2023), and by the $PM_{2.5}$ improvement (the necessary reduction in $PM_{2.5}$ to reach 10 ug/m^3).

To capture statistical uncertainty associated with the key parameters used in our empirical calculations we use a Monte Carlo simulation. These exercises include variation in the VSL (USEPA, 2017), the premature mortality-$PM_{2.5}$ concentration response function (USEPA, 2023e), and the per-capita abatement cost estimates. The uncertain parameters are varied simultaneously in each of the 1,000 iterations of the Monte Carlo simulation. We assume that the realizations of these parameters are uncorrelated.

4.2 Empirical Simulation Design and Results

We design and execute the following empirical simulation to apply the framework developed in sections 2 and 3. Currently, the primary National Ambient Air Quality Standards (NAAQS) for $PM_{2.5}$ are set to 12 ug/m^3 (USEPA, 2023d). Recently, the USEPA proposed tightening the NAAQS to between 9 and 10 ug/m^3 (USEPA, 2023c). We leverage this proposed policy change to examine the implications of distributional weights for benefit-cost analysis. Using annual average $PM_{2.5}$ readings by county from the AQS databases (USEPA, 2023a), we calculate the benefits and costs associated with attaining an average $PM_{2.5}$ level of 10 ug/m^3 in any county that exceeded this level in 2021. Thus, we evaluate (17) and (18) using the $\Delta PM_{i,t}$ calculated as the difference between observed readings and 10 ug/m^3 provided the difference is non-negative.

In 2021, there were 100 monitored counties in the coterminous U.S. for which annual average $PM_{2.5}$ levels exceeded 10 ug/m^3. This set of counties falls to 51 when we merge the AQS data with the IPUMS databases. Among these 51 counties, the average difference between observed $PM_{2.5}$ levels and the proposed 10 ug/m^3 standard was 1.32 ug/m^3. Figure A1 in the appendix shows the distribution of these exceedances. While most counties exhibit a difference under 2 ug/m^3, 12 counties show a difference greater than 2 ug/m^3 with a maximum difference of over 7 ug/m^3.

4.2.1 Total Benefits

Table 1 reports the benefits of achieving the 10 ug/m^3 NAAQS. The three rows represent different levels of aggregation of the income data. The first row employs household-level income data. The second row assigns households to four groups: below the 25th percentile of income; between the 25th percentile and the median; between the median and the 75th percentile; and above the 75th percentile. To calculate the group-specific weights using (20), we employ within-group average income relative to the national median income. The third row considers two groups, one below the median one above the median. The columns consider different assumptions about the VSL and the distributional weights. The first column uses a uniform VSL with unit weights; the second column uses group-specific VSLs with unit weights (19); and the third column uses group-specific VSLs with distributional weights, the product of (19) and (20).

To understand the patterns observed in the table, we begin by considering the first row of the table which uses household-level data. We first compare column 2 with column 1 and then column 3 with column 2. The first column uses USEPA's uniform VSL of $10.5 million with unit weights.[29] In this case, total annual benefits are estimated to be $441 million. In the second column of Table 1, we convert the uniform VSL to group-specific VSLs and still apply unit weights. These benefits are $468 million at the household level of aggregation, a 6% increase over the uniform VSL case. The change in benefits from the uniform VSL to the household specific VSL with unit weights depends on two offsetting factors. First, absolute reductions in ambient $PM_{2.5}$ are weakly negatively correlated with income (see Figure A2 in the appendix). Hence, poorer households, on average, would experience greater improvements in air quality under the 10 ug/m^3 NAAQS. All else equal, this will decrease total benefits since the unweighted household VSLs are lower for low-income households. Countering this effect is the fact that the distribution of household-specific VSLs is right-skewed (see Figure A3 in the appendix). Thus, the smaller pollution reductions experienced by high-income households are valued much more than when using the uniform VSL, which would tend to increase benefits. The fact that benefits increase in the second column relative to the first suggests that the latter effect dominates.

Now, consider the change in the first row from column 2 to column 3. We see that household-specific distributional weights have a pronounced effect on the benefit estimates. Total benefits at the household level of aggregation increase by 60% relative to the unweighted case (compare $750 million to $468 million). The significant increase in total benefits with distributional weights arises because households with very low incomes are assigned very large distributional weights *and* they tend to incur larger improvements in air quality. (See figures A4 and A6 in the appendix.) These effects reinforce each other.

We now turn to patterns across rows within a given column. The following three patterns emerge. First, with a uniform VSL, aggregation does not affect benefits (column 1). Second, using group-specific VSLs with unit weights (column 2), total benefits increase as the data are more aggregated. Third, using the group-specific VSLs with distributional weights (column 3), total benefits decrease with aggregation. Mechanisms for these patterns are discussed below.

In the first column, the level of aggregation has no effect on the total benefit estimates (ignoring rounding error), since the VSL is applied uniformly across all people. The second column indicates that the when using unit weights and group-specific VSLs, total benefits increase as the level of aggregation increases. In the four-group case, total benefits increase by 3 percent relative to when using household income data (compare

[29] The use of a uniform VSL embodies implicit distributional weights (Hemel, 2022).

$481 million to $468 million). The increase in benefits with aggregation of the income data can be explained by two factors. First, recall that low-income households tend to incur the largest improvements in air quality. Second, averaging income by groups increases the VSL for very low households relative to the case that uses household income data. Hence, the value attributed to the largest air quality changes and mortality risk reductions rises. The result is an increase in total benefits. The same reasoning applies in moving from four to two income groups. Total benefits increase by about 4 percent from $481 million to $499 million.

The third column shows that aggregation reduces total weighted benefits compared to when household level data are used. Comparing the first and second rows in column 3 shows that benefits fall from $750 million to $587 million (a 22 percent reduction). Further aggregating households into two groups yields benefits of $503 million which is a 33 percent decrease relative to the first row. The reason for this is that averaging incomes within groups smooths extreme outliers at the bottom of the income distribution. This eliminates the very large outliers in the distributional weights and the weighted VSLs calculated using household income data.

4.2.2. Net Benefits

Table 2 combines total costs and total benefits to report the net benefits associated with meeting the proposed NAAQS. Table A1 in the appendix reports the cost estimates. USEPA (1999) reports that in 2010, expenditures on PM NAAQS attainment were $5.1 billion. AQS monitoring data indicates $PM_{2.5}$ levels fell by an average of 0.13 ug/m^3 between 2010 and 2011. Recognizing that the 2010 population was 332 million, the per capita cost/(ug/m^3) amounts to about $120/(ug/m^3)/person. This cost estimate is used together with our estimated $PM_{2.5}$ reductions to meet the 10 ug/m^3 standard and population data from IPUMS (2023) for the counties covered in this analysis. Total unweighted costs as incurred by the sampled population are $74 million. We explore different approaches to allocating costs to households and the implications for net benefits in section 4.2.3 below.

Several results emerge from Table 2 that relate to how costs are weighted. The first column in Table 2 parallels the result for the first column in Table 1, except net benefits are lower than total benefits by $74 million, reflecting the fact that costs have unit weights. (See Figure A7 for the full Monte Carlo results.) The same reasoning holds when comparing the second column of Table 2 with the second column of Table 1. Again, net benefits are reduced by $74 million.

The results are more interesting in column 3 which applies weights to both benefits and costs. In the first row, using the household level income data, net benefits are *negative.* Figure A8 shows the Monte Carlo results. The large negative value occurs because the distributional weights for very low-income households are very large and both benefits *and* costs are weighted. As such, if the unweighted benefits per household are smaller than the costs for these low-income households, *weighted* net benefits may be sufficiently negative to outweigh the positive net benefits accruing to other households. In fact, we find that about two-thirds of the total -$11.6 billion in net benefits (-$7 billion) arise from just 5 households with near-zero income. This is a clear example of the dictator paradox discussed in section 3. As household income falls to zero, the distributional weights become very large, which heavily influences policy outcomes.

The reason for the heavy influence on outcomes is shown in Figure 2. The figure plots the cumulative sum of the distributional weights normalized by the total number of households: $\left(\frac{1}{242k}\sum_{i=1}^{242k} w_i\right)$. This metric depicts the cumulative value of a total of $1 in either benefits or costs accruing equally to households across the income distribution. The figure illustrates two points. First a weighted dollar equally distributed has 185 times the utility of an unweighted dollar equally distributed across the population (see the top line in Figure 2). Second, it shows that the vast majority of the weight is accounted for by a very small fraction of the population. For example, the 1st percentile of households account for 98% of the weight. As Hahn (2023) has noted, this could provide a justification for simply transferring money from higher income to lower income individuals with large increases in welfare. In addition, the figure serves as a further illustration of the dictator paradox. When distributional weights are used to evaluate policy, the sign and the magnitude of net benefits are entirely driven by a *very* small number of households. This point is likely to be quite general and not dependent on the policy application. It simply requires that a small number of households have very low incomes and that the functional form of distributional weights used here is relevant.[30] The paradoxical nature of this result suggests why policy makers may want to exercise extreme caution if they choose to apply such distributional weights.

Column 4 reports net benefits with weighted benefits and unweighted costs. Compared with column 3, which applies weighted costs, net benefits are always higher. This is consistent with the results shown in Figure 2, which indicate that weighting tends to accentuate economic impacts (weighting costs in this case) compared with the case of unit weights. Perhaps of greater interest is that when costs are not weighted, net benefits switch signs when using household-level data (going from -$11,600 million in the case when costs have group-specific weights to $676 million in the case when costs have unit weights). That the sign of the net benefits can hinge on whether both benefits and costs are weighted provides empirical support for the claim made in section 2 above regarding the importance of weighting both.

A comparison of the net benefits across rows in the third column in Table 2 highlights the importance of aggregation when using distributional weights that differ from unity. Aggregating households into four income groups yields net benefits that are positive and smaller (in absolute value) by a factor of 27 relative to the household level data (compare $436 million to -$11.6 billion). This change in net benefits occurs because the extreme outliers in the left-tail of the income distribution are smoothed by allocating households into groups and averaging incomes. Further aggregating households into two groups

[30] The measure we use is based on income. An alternative might be to consider consumption, which might yield less dramatic results than those obtained here (e.g., if consumption were somewhat higher at the very low end of the income scale). We use income because that variable is more commonly associated with distributional weights in the literature and because of data availability.

reduces total net benefits from $436 million to $400 million. Our results suggest that when using distributional weights, the degree to which households are aggregated can materially affect a policymaker's conclusions.

4.2.3 Impact of Changing Cost Allocation on Net Benefits

Not very much is known about the incidence of costs for many environmental policies. Here, we consider a range of outcomes to highlight the potential importance of cost allocation on both costs and net benefits. We will show that cost allocation matters for net benefits, and that the results are driven largely by the lowest-income households.

Table 3 reports four different approaches to cost allocation: 1. Equal per capita cost allocation (the base case considered above); 2. costs are allocated in direct proportion to a household's share of total income; 3. costs are a logarithmic function of household income; and 4. costs are allocated proportional to each household's income share, subject to the constraint that each household must pay at least $20. We consider only the household level of aggregation here to help isolate the impact of cost allocation on costs and net benefits.

The results in Table 3 can be explained by one key factor: the policy cost share for the lowest 1 percent of households in the sample. Column 1 in Table 3 reports the percent of total costs borne by the bottom 1 percent of the households when distributional weights are applied. It shows that when the bottom 1 percent of households account for a very large share of the costs in weighted terms, net benefits are negative. When the bottom 1 percent of the households account for more than 90 percent of the total costs net benefits are negative (-$11,600 million in the equal per capita allocation case, and -$156 million in the minimum cost case). This result follows because these costs are relatively heavily weighted in welfare terms, which makes costs very large. Second there is a large variation in costs and net benefits across the 4 different allocation schemes as shown in columns two and three. For example, net benefits range from -$11,600 million to $688 million depending on the allocation scheme. Third, there is an inverse relationship between the percent of total cost paid by the bottom 1 percent of households and net benefits. The relationship is non-linear. These results are driven by the relatively very large welfare weights associated with households at the bottom end of the income distribution as shown in figure 2.

In the fourth row of Table 3, we assume that the minimum cost faced by low-income households is $20/year. To explore the sensitivity of net benefits to this assumption we vary this minimum cost from zero to $100. Figure A9 plots weighted net benefits against the various minimum household policy costs[31]. With zero minimum cost, net benefits align with those reported in first row of Table 3. As the minimum cost rises, net benefits fall. Net benefits become negative when the minimum cost faced by households reaches just $17, annually. Thus, if the $PM_{2.5}$ NAAQS revision imposes even very small costs on low-income households, weighting benefits *and* costs is likely to change the sign of net

[31] We tabulate weighted net benefits for each assumed minimum cost. In effect, adding a minimum cost simply adds an intercept to the function that allocates cost to households proportional to income.

benefits relative to an analysis using unit weights. The sensitivity of both the sign and the magnitude of the net benefit estimates to the cost shares borne by a very small number of very low-income households is further evidence of the dictator paradox.

5. Conclusion

There is increased interest among researchers and policy makers in measuring the impact of policies on different groups within a population. One approach that is gaining traction in U.S. policy is to introduce distributional weights into a conventional BCA framework. By doing so, the analysis recognizes that a dollar of benefits or costs does not have the same value to all individuals or households affected by a policy. This approach to policy analysis introduces a kind of equity analysis that is different than traditional approaches that consider whether it may be possible for the winners from a policy to compensate the losers and for society as a whole to still be better off (Kaldor, 1939, and Hicks, 1939). In this paper we develop a general model that shows how the optimal provision of a public good varies with the introduction of distributional weights. We show that the level of public good provision may increase, decrease or remain the same compared with the unit weight case. In addition, we identify conditions under which public good provision will (or will not) be produced at minimum cost.

Next, we illustrate how the framework can be used to shed light on an important policy issue: the welfare impacts of using a single average VSL versus multiple, group-specific, VSLs that may depend on a parameter, such as income. We show how to determine when using an average VSL will increase or decrease optimal public goods provision relative to the case of group-specific VSLs. We also identify conditions under which a low-income group would prefer using an average VSL to group-specific VSLs in which their mortality risks are valued less than the mortality risks faced by a high-income group.

Using our general framework with a set of distributional weights suggested by the U.S. government and researchers, we derive the dictator paradox, which states that a very small number of low-income households may account for a disproportionate share of the weight policymakers apply in a benefit-cost analysis. We emphasize the paradoxical nature of this outcome. The use of distributional weights is motivated by equity considerations. Yet, this approach may result in policy outcomes essentially dictated by a very small number of individuals or households. This outcome is similar to, but not the same as, the Rawlsian maximin principle which places all weight on the worst off individual. The dictator paradox differs in that some positive distributional weight is placed on all members of society, not just the lowest income household.

We apply the formal model to evaluate a recent proposal by the USEPA to tighten the National Ambient Air Quality Standard for fine particulate matter. The empirical analysis highlights several important aspects of using distributional weights in applied policy analysis. First, how policymakers elect to aggregate household income data has a marked effect on the magnitude of net benefit estimates.[32]

Second, the empirical analysis underscores the importance of weighting *both* benefits and costs. Using household level income data, we report that weighted net benefits are sizable *and negative*. This result occurs because weights calculated using household income data are extremely large for very low-income households and, for these outlier observations, per capita costs exceed benefits. In contrast, when only benefits are weighted, net benefits are positive. Furthermore, the negative net benefit estimate only occurs when using household level data. Once households are aggregated into larger groups, and incomes are averaged within groups, net benefits are positive. In essence, this occurs because very low incomes are averaged away and the distributional weights are considerably smaller for very low- income households. This result emphasizes that the level of data aggregation interacts with whether benefits and costs are both weighted.

While we have illustrated the impacts of using distributional weights on an important local environmental problem, there are several important global public goods problems that could be analyzed through this lens. Examples include climate change and pandemics (e.g., Anthoff and Tol, 2010 and Budolfson et al., 2021). We argue that it would be useful to conduct BCA in these two global scale contexts and to compare the results obtained using distributional weights and unit weights using the framework developed here.

Because of the challenges we identify associated with the use of distributional weights, we wish to be clear that we are *not* endorsing the use of this approach in policy analysis. Given our results related to the impact of aggregation on welfare, and the dictator paradox, we think more research is needed to assess the desirability of using distributional weights in the economic analysis of policies. This research might focus on how different types of weights affect numerical benefit and cost estimates, and the sense in which aggregation matters. We conjecture that the kinds of aggregation issues that we identify here will persist in many other settings.

Policy makers should be aware of the issues associated with distributional weights before using them in benefit-cost analysis. FEMA's guidance for dealing with extreme-valued weights suggests agencies have begun to grapple with these issues, but in ways that appear to be arbitrary (FEMA, 2024). Still, if policy makers decide to employ weights in spite of the challenges we highlight, there is a strong case to be made for applying weights to *both* benefits and costs because this could have a big impact on the kinds of policies that are identified as welfare-improving. Furthermore, there is no obvious reason why costs and benefits for a particular group should be treated differently.

Alternatives to distributional weights that highlight the distributional impacts of policies should also be considered. When large amounts of government expenditures are involved, a framework that can explicitly address these issues is the marginal value of public funds (Hendren and Sprung Keyser, 2020). This approach can be used to highlight

[32] OMB notes that choices about the level of aggregation should be "tailored to the context," but does not appear to consider how the level of aggregation would affect net benefits (OMB, 2023, p. 63).

the welfare impacts of expenditures on different groups as well as on government revenues (Finkelstein and Hendren, 2020). Similar to disaggregated benefit-cost analysis, this approach allows decision makers to decide if they think a policy is worthwhile based on the welfare weights they wish to assign to different groups.

Robert W. Hahn
University of Oxford
Carnegie Mellon University

Nicholas Z. Muller
Carnegie Mellon University
NBER

References

Acland, D., and D.H. Greenberg. 2023. "The Elasticity of Marginal Utility of Income for Distributional Weighting and Social Discounting: A Meta-Analysis." *Journal of Benefit-Cost Analysis.* 1-20.

Adler, Matthew. 2016. "Benefit–cost analysis and distributional weights: an overview." *Review of Environmental Economics and Policy*. 10: 264-285.

Anthoff, David, and Richard SJ Tol. 2010. "On international equity weights and national decision making on climate change." *Journal of Environmental Economics and Management* 60(1): 14-20.

Banzhaf, H. Spencer. 2023. "Distribution and Disputation: Net Benefits, Equity, and Public Decision Making." Available at SSRN: https://ssrn.com/abstract=4449328.

Banzhaf, H. Spencer. 2011. "Regulatory impact analysis of environmental justice effects." *J. Land Use & Envtl. L.* 27(1): 1-30.

Bergstrom, Theodore C., and Richard C. Cornes. 1983. "Independence of allocative efficiency from distribution in the theory of public goods." *Econometrica*. 1753-1765.

Budolfson, Mark B., et al. 2021. "Utilitarian benchmarks for emissions and pledges promote equity, climate and development." *Nature Climate Change* 11(10): 827-833.

Cecot, Caroline, and Robert W. Hahn. 2024. "Incorporating equity and justice concerns in regulation." *Regulation & Governance* 18(1): 99-120.

Chichilnisky, Graciela, and Geoffrey Heal. 1994. "Who should abate carbon emissions?: An international viewpoint." *Economics Letters* 44(4): 443-449.

Coase, Ronald. 1960. "The Problem of Social Cost." *Journal of Law and Economics*, 3: 1–44.

Dasgupta, Ajit K., and D. W. Pearce. 1972. *Cost-benefit analysis: Theory and practice.* New York: Harper and Row.

Dasgupta, Partha, Amartya Kumar Sen, and Stephen A. Marglin. 1972. *Guidelines for project evaluation.* United Nations. New York, NY, USA.

Finkelstein, Amy, and Nathaniel Hendren. 2020. "Welfare analysis meets causal inference." *Journal of Economic Perspectives* 34 (4): 146-167.

Hahn, Robert W. 2023. "Reforming regulation with an eye toward equity." *Science* 380(6648): 899-901.

Harberger, Arnold C. 1978. “On the Use of Distributional Weights in Social Cost-Benefit Analysis.” *Journal of Political Economy.* 86(2) Part 2: Research in Taxation: S87 – SS120.

Hemel, Daniel. 2022. "Regulation and Redistribution with Lives in the Balance." *U. Chi. L. Rev.* 89: 649-734.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A unified welfare analysis of government policies." *The Quarterly Journal of Economics* 135(3): 1209-1318.

Hicks, John. 1939. "The Foundations of Welfare Economics." *The Economic Journal*. 49(196): 696–712.

HM Treasury. 2003. The Green Book: Central Government Guidance on Appraisal and Evaluation, Stationery Office.

Kaldor, Nicholas. 1939. "Welfare propositions of economics and interpersonal comparisons of utility." *The Economic Journal.* 49(195): 549-552.

Meade, J. E. 1955. *Trade and welfare: Mathematical supplement.* London: Oxford University Press.

Mendelsohn, Robert. 1986. "Regulating heterogeneous emissions." *Journal of Environmental Economics and Management* 13(4): 301-312.

Montgomery, W. David. 1972. "Markets in licenses and efficient pollution control programs." *Journal of Economic Theory* 5(3): 395-418.

Muller, Nicholas Z., and Robert O. Mendelsohn. 2009. "Efficient pollution regulation: getting the prices right." *American Economic Review* 99(5): 1714-1739.

Muller, Nicholas Z., Robert O. Mendelsohn, William D. Nordhaus. 2011. “Environmental Accounting: Methods with an Application to the United States Economy.” *American Economic Review.* 101(5): 1649 – 1675.

Piketty, Thomas, E. Saez. 2003. “Income Inequality in the United States: 1913 – 1998.” *Quarterly Journal of Economics.* 118(1): 1 – 39.

Piketty, Thomas, E. Saez. 2006. “The evolution of top incomes: a historical and international perspective.” *American Economic Review.* 96(2): 200-205.

Piketty, Thomas, E. Saez, G. Zucman. 2018. “Distributional National Accounts: Methods and Estimates for the United States.” *Quarterly Journal of Economics.* 133(2): 553-609.

Piketty, Thomas, 2022. *A brief history of equality*. Harvard University Press. Cambridge, MA, USA.

Rawls, J. 1971. *A theory of justice.* Belknap Press/Harvard University Press.

Robinson, Lisa A. 2007. "Policy Monitor: How US Government Agencies Value Mortality Risk Reductions." *Review of Environmental Economics and Policy.* 1(2): 283-299.

Samuelson, Paul A. 1954. "The pure theory of public expenditure." *The Review of Economics and Statistics.* 36(4): 387-389.

Spiller, Elisheba, Jeremy Proville, Ananya Roy, Nicholas Z. Muller. 2021. "Mortality Risk from $PM_{2.5}$: A Comparison of Modeling Approaches to Identify Disparities Across Ethnic Groups in Policy Outcomes." *Environmental Health Perspectives*. 129(12): 1 – 12.

Stavins, Robert N. 2022. "The relative merits of carbon pricing instruments: Taxes versus trading." *Review of Environmental Economics and Policy* 16(1): 62-82.

Sunstein, Cass R. 2023. "Inequality and the Value of a Statistical Life." *Journal of Benefit-Cost Analysis* 14(1): 1-7.

Sunstein, Cass R. 2024. "The Economic Constitution of the United States." *Journal of Economic Perspectives* 38(2): 25-42.

United States Environmental Protection Agency. 1999. *The Benefits and Costs of the Clean Air Act: 1990--2010.* EPA Report to Congress. EPA 410-R-99-001, Office of Air and Radiation, Office of Policy, Washington, D.C.

United States Environmental Protection Agency. 2011. *The Benefits and Costs of the Clean Air Act: 1990--2020.* Final Report. Office of Air and Radiation, Office of Policy, Washington, D.C.

United States Environmental Protection Agency. 2017. “Appendix B. Mortality Risk Valuation Estimates.” https://www.epa.gov/sites/default/files/2017-09/documents/ee-0568-22.pdf

United States Environmental Protection Agency, 2022. Regulatory Impact Analysis for the Proposed Reconsideration of the National Ambient Air Quality Standards for Particulate Matter. EPA-452/P-22-001. Office of Air Quality Planning and Standards.

Health and Environmental Impacts Division. Research Triangle Park, NC, USA. https://www.epa.gov/system/files/documents/2023-01/naaqs-pm_ria_proposed_2022-12.pdf

United States Environmental Protection Agency, 2023a. https://aqs.epa.gov/aqsweb/airdata/download_files.html#Raw

United States Environmental Protection Agency, 2023b. https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatvalue

United States Environmental Protection Agency, 2023c. https://www.epa.gov/pm-pollution/national-ambient-air-quality-standards-naaqs-pm

United States Environmental Protection Agency, 2023d. https://www.epa.gov/criteria-air-pollutants/naaqs-table

United States Environmental Protection Agency, 2023e. “Estimating $PM_{2.5}$ and Ozone-Attributable Health Benefits.” Technical Support Document (TSD) for the 2022 PM NAAQS Reconsideration Proposal RIA.” EPA-HQ-OAR-2019-0587. USEPA OAR. Research Triangle Park, North Carolina, USA.

Viscusi, W. Kip. 2018. "Pricing lives." In *Pricing Lives*. Princeton University Press, Princeton, NJ, USA.

Viscusi, W. Kip, and Joseph E. Aldy. 2003. “The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World.” *Journal of Risk and Uncertainty.* 24: 5 – 76.

Viscusi, W. Kip, and Clayton J. Masterman. 2017. “Income Elasticities and Global Values of a Statistical Life.” *Journal of Benefit Cost Analysis*. 8(2): 226 – 250.

Weitzman, Martin L. 1974. "Prices vs. quantities." *The Review of Economic Studies.* 41(4): 477-491.

The White House. 2021. "Executive order on advancing racial equity and support for underserved communities through the federal government." Executive Order No. 13985.

The White House. 2023, “Circular No. A-4.” (2023), Office of Management and Budget. November 9th.

Figure 1: Social welfare loss and low-income group gain with average VSL

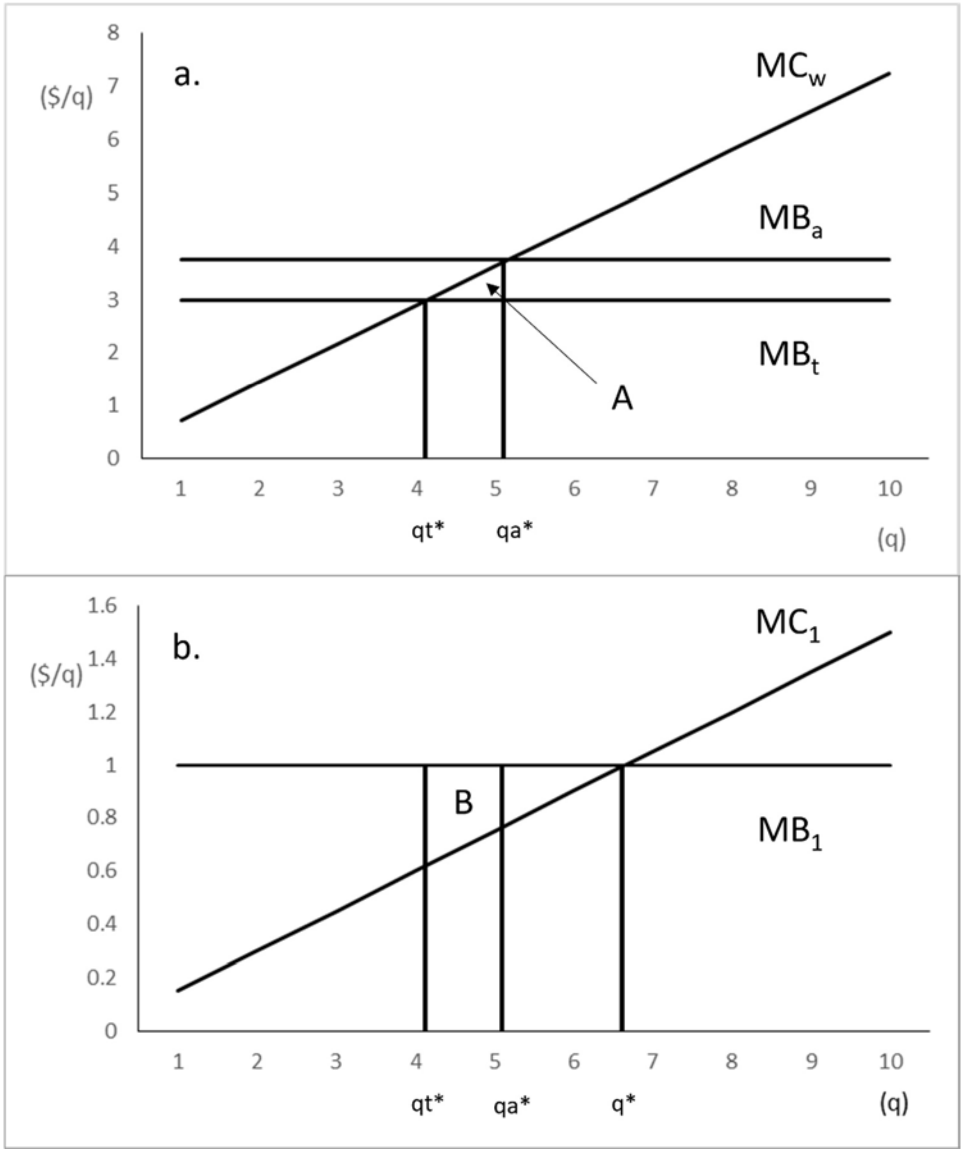

Notes: MC_w = weighted marginal cost; MB_a = weighted marginal benefit using average VSL; MB_t = weighted marginal benefit using the true VSL; MC_1 = low-income group marginal cost; MB_1 = low-income group marginal benefit; A = social welfare loss associated with q_a^* rather than q_t^*; and B = welfare gain to low-income group associated with q_a^* rather than q_t^*. q refers to quantity of the public good.

Figure 2: Distributional Weights and Household Income

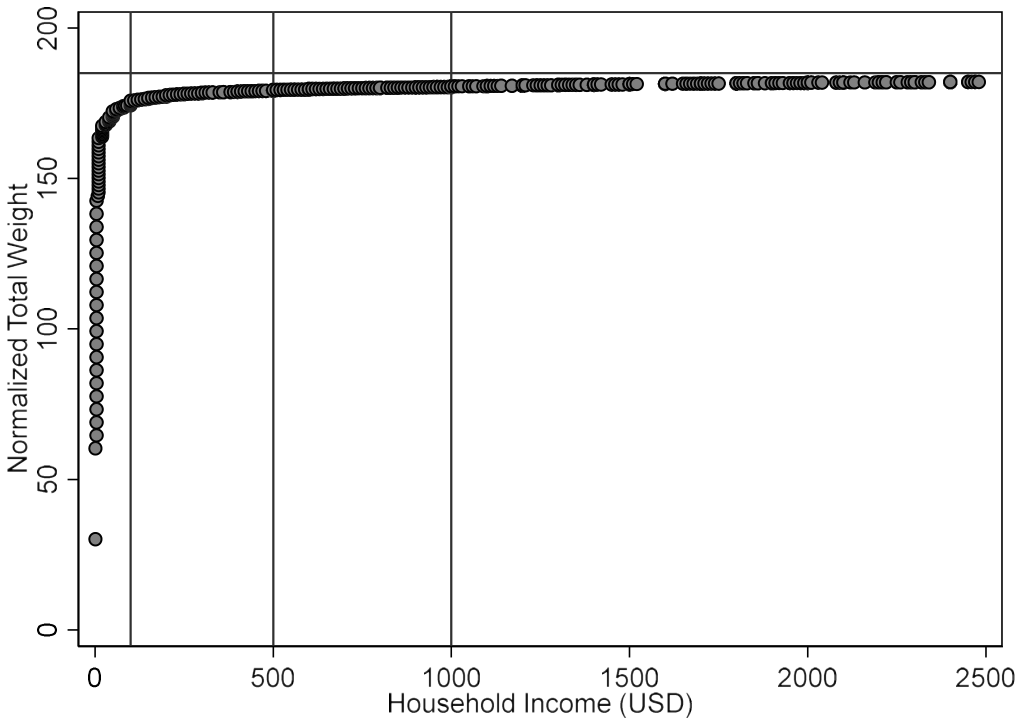

Notes: Vertical axis defined as: $\frac{1}{N}\sum_{i=1}^{N} w_i$. N = 242,630. This depicts the cumulative value of $1 in either benefits or cost across the households included in the study. Figure 2 includes the lowest 1% of households by income. Vertical lines at $100, $500, and $1,000. Horizontal line at $185 is the maximum value of the normalized total weight.

Table 1: Benefits of Achieving $PM_{2.5}$ NAAQS of 10 ug/m^3

Unit of Observation	Uniform VSL Unit Weights	Group-Specific VSL Unit Weights	Group-Specific VSL Distributional Weights
Household	441 (97)	468 (102)	750 (164)
Four Groups	440 (97)	481 (105)	587 (129)
Two Groups	441 (97)	499 (109)	503 (110)

Notes: Benefits are in millions of 2021 dollars. Values in parentheses are standard deviations across 1,000 draws from Monte Carlo simulation. Column 1 uses USEPA's uniform VSL and unit weights. Column 2 converts the uniform VSL to group-specific values using equation (18), applying unit weights. Column 3 weights the group-specific VSLs used in column 2 using equation (19). There are approximately 242,630 households. The four group case features roughly 60,000 households per group. The two group case features roughly 120,000 households per group.

Table 2: Net Benefits Achieving $PM_{2.5}$ NAAQS of 10 ug/m^3.

Unit of Observation	Uniform VSL Unit Weights	Group-Specific VSL Unit Weights	Group-Specific VSL Distributional Weights	Group-Specific VSL Distributional Weights Unit Weight Costs
Household	367 (321)	394 (323)	-11,600 (51,400)	676 (348)
Four Groups	368 (321)	406 (327)	436 (638)	513 (343)
Two Groups	367 (321)	425 (327)	400 (446)	430 (332)

Notes: Net benefits are in millions of 2021 dollars. Values in parentheses are standard deviations across 1,000 draws from Monte Carlo simulation. Column 1 uses USEPA's uniform VSL and unit weights. Column 2 converts the uniform VSL to group-specific values using equation (18), applying unit weights. Column 3 weights the group-specific VSLs used in column 2 using equation (19). Column 4 weights the VSLs used in column 2 using equation (19) but costs have unit weights. There are approximately 242,630 households. The four group case features roughly 60,000 households per group. The two group case features roughly 120,000 households per group. Costs are allocated on a per-capita basis.

Table 3: Total Costs and Net Benefits Under Different Cost Assumptions.

Cost Allocation	% Total Cost With Distributional Weights For Bottom 1% of Households	Total Cost With Distributional Weights	Net Benefits With Distributional Weights
Equal per Capita	98.1	12,400 (51,400)	-11,600 (51,400)
Proportional	5.8	62 (259)	688 (307)
Logarithmic	40.6	371 (485)	378 (513)
Minimum Cost	91.7	896 (261)	-146 (309)

Notes: Costs and net benefits are in millions of 2021 dollars. Values in parentheses are standard deviations across 1,000 draws from Monte Carlo simulation. Column 1 percentage of total cost borne by lowest 1 percent of households with distributional weights. Column 2 reports the total cost using weights in equation (19). Column 3 shows net benefits using weights in equation (19). Equal per capita cost allocation attributes cost to households with equal per capita cost times reported number of persons in each household. Proportional cost allocation attributes cost proportional to household share of total income. Logarithmic cost allocation attributes cost based on the natural log of household income. Minimum cost allocation attributes cost proportional to household share of total income, but no household incurs cost less than $20.

Appendix

This appendix contains the proofs of Proposition 1 and Proposition 2 as well as supplementary tables and figures.

Proposition 1: Assume the aggregate marginal curve is upward sloping and/or the sum of marginal benefits curve for all groups is downward sloping.

Then:

$$q_w^* \,(>)(=)(<) q^* \; iff \; \frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} (>)(=)(<) \sum_{j=1}^{m} B_k'(q_w^*) \quad (17)$$

Proof: Consider the case of equality. Assume that $q_w^* = q^*$. From equation (4), $\sum_{j=1}^{m} B_j'(q^*) = C'(q^*)$, and from equation (6), $C'(q_w^*) = \frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k}$. Because $q_w^* = q^*$, it follows that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$. Now, assume that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$. We wish to show that this implies $q_w^* = q^*$. q_w^* is optimal in the general case. But it is also optimal in the unit weight case because $C'(q_w^*) = \frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$. Thus, $q_w^* = q^*$. This shows $q_w^* = q^*$ iff $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$.

Consider the greater-than case and assume that $q_w^* > q^*$. We wish to show that this implies that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \sum_{j=1}^{m} B_j'(q_w^*)$. We use a proof by contradiction. That is, assume that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} \leq \sum_{j=1}^{m} B_j'(q_w^*)$ and arrive at a contradiction. Consider the cases of equal-to and less-than separately. We know that if $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$ then $q_w^* = q^*$, which violates the assumption that $q_w^* > q^*$. Now, consider the case where we assume $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} < \sum_{j=1}^{m} B_j'(q_w^*)$. Given q_w^* is optimal for the general case, $C'(q_w^*) = \frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k}$. But this implies under this assumption that $C'(q_w^*) < \sum_{j=1}^{m} B_j'(q_w^*)$. This inequality implies that q must be increased from q_w^* to satisfy the unit weight case first order condition, which means that $q^* > q_w^*$. But this contradicts the assumption that $q_w^* > q^*$. This shows that $q_w^* > q^*$ implies $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \sum_{j=1}^{m} B_j'(q_w^*)$.

Now, consider the greater-than case and assume that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \sum_{j=1}^{m} B_j'(q_w^*)$. We wish to show this implies $q_w^* > q^*$. We use a proof by contradiction. That is, assume that $q_w^* \leq q^*$ and arrive at a contradiction. Consider the cases of equal-to and less-than separately. We know that if $q_w^* = q^*$, $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} = \sum_{j=1}^{m} B_j'(q_w^*)$ which violates the assumption that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \sum_{j=1}^{m} B_j'(q_w^*)$.

Consider the case where $q_w^* < q^*$. Given q_w^* is optimal for the general case, $C'(q_w^*) = \frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k}$ which implies, under the assumption that $q_w^* < q^*$, that $C'(q_w^*) < \sum_{j=1}^{m} B_j'(q_w^*)$,[33] which in turn implies that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} < \sum_{j=1}^{m} B_j'(q_w^*)$. This contradicts our assumption. This shows that $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \sum_{j=1}^{m} B_j'(q_w^*)$ implies $q_w^* > q^*$. Taking the results on the less-than case together, we have shown that $q_w^* > q^*$ iff $\frac{\sum_{j=1}^{m} w_j B_j'(q_w^*)}{\sum_{k=1}^{m} w_k p_k} > \sum_{j=1}^{m} B_j'(q_w^*)$.

The same proof holds, *mutatis mutandis*, for the less-than case.

<u>Proposition 2</u>: $q_t^*(<)(=)(>)q_a^*$ iff $w_1 n_1 (>)(=)(<) w_2 n_2$, with the assumption on the linear marginal benefit function and cost functions defined above.

<u>Proof</u>: Because the denominator is positive and equal, we form the difference between the numerators in equation (9) and equation (10). This difference is $(<)(=)(>)0$ iff $(q_t^* - q_a^*)\ (<)(=)(>)0$.

This yields: $(q_t^* - q_a^*)\ (<)(=)(>)0$

iff

$$[(w_1 n_1 c + w_2 n_2 d) - (w_1 n_1 + w_2 n_2)(\frac{c+d}{2})]\ (<)(=)(>)0$$

iff

$$(w_1 n_1 c + w_2 n_2 d) - (w_1 n_1 + w_2 n_2)(\frac{c+d}{2})\ (<)(=)(>)0$$

iff

$$w_1 n_1 (c - (\frac{c+d}{2})) + w_2 n_2 (d - (\frac{c+d}{2}))\ (<)(=)(>)0$$

iff

$$.5[w_1 n_1 (c - d) + w_2 n_2 (d - c)]\ (<)(=)(>)0$$

iff

[33] This follows directly from the assumption that the marginal benefit function is non-decreasing and the marginal cost function is increasing.

$w_1 n_1 (c - d) + w_2 n_2 (d - c)\ (<)(=)(>) 0$

Letting $k = d - c$, and substituting yields

iff
$k\ (w_2 n_2 - w_1 n_1)\ \ (<)(=)(>)\ 0$

iff
$w_2 n_2 - w_1 n_1\ \ (<)(=)(>)\ 0$

iff $w_1 n_1 (>)(=)(<) w_2 n_2$

Each of the if-and-only-if steps represent standard algebraic operations. This proves Proposition 2.

Supplementary Tables and Figures

Figure A1: Distribution of $PM_{2.5}$ Exceedances.

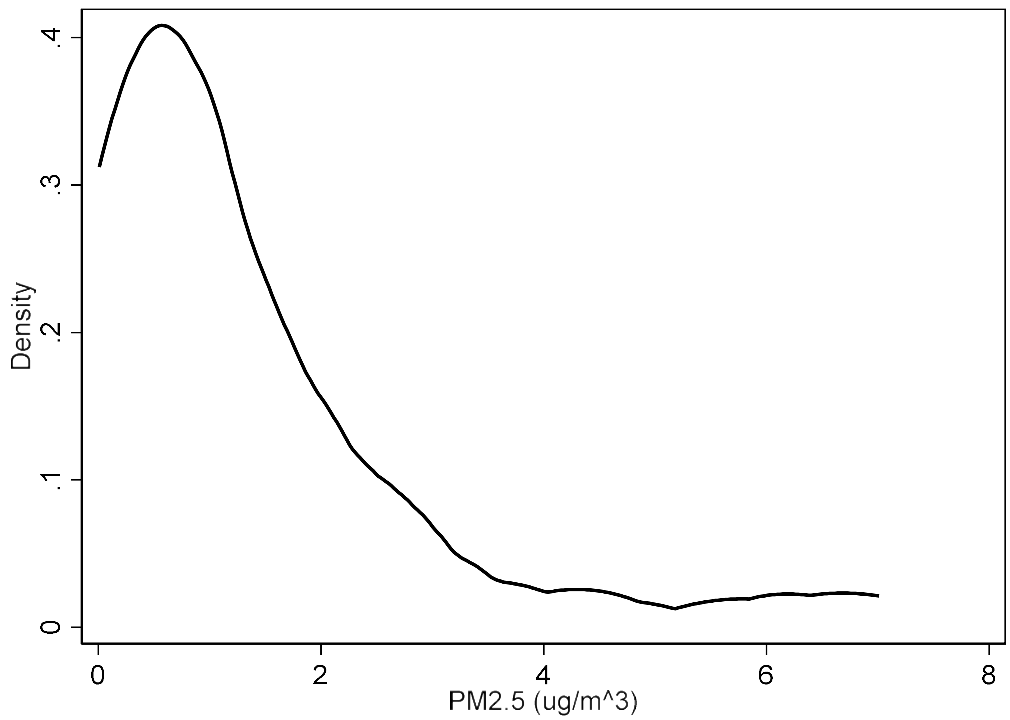

Notes: Figure A1 shows the probability density function of the difference between annual average county level $PM_{2.5}$ readings and the proposed NAAQS of 10 ug/m^3 for all counties with observed readings at 10 ug/m^3 or greater.

Figure A2: $PM_{2.5}$ Exceedances and Household Income

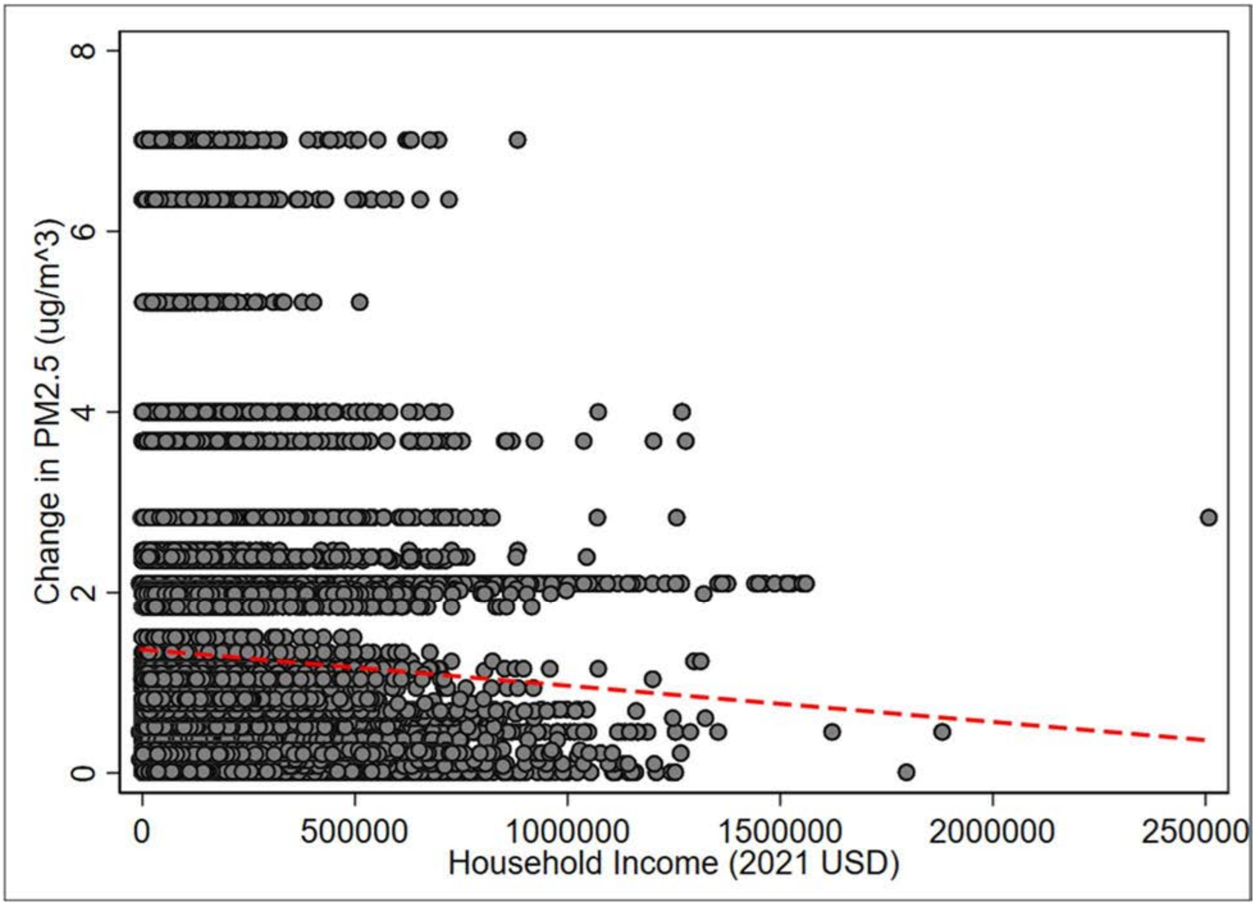

Notes: Vertical axis shows the difference between annual average county level $PM_{2.5}$ readings and the proposed NAAQS of 10 ug/m^3 for all counties with observed readings at 10 ug/m^3 or greater. Each dot represents one household, matched to county-average $PM_{2.5}$ exceedance. N = 242,630. The dashed red line is linear fit of $PM_{2.5}$ reductions to the 10 ug/m^3 NAAQS on household income.
The regression model assumes the following form, with income in $1,000.

$$ln(\Delta PM_{2.5,j}) = \alpha_0 + \alpha_1 Income_j + \varepsilon_j$$

$$\widehat{\alpha_1} = -0.001^{***}$$

Figure A3: Household Income Adjusted VSL with Unit Weights

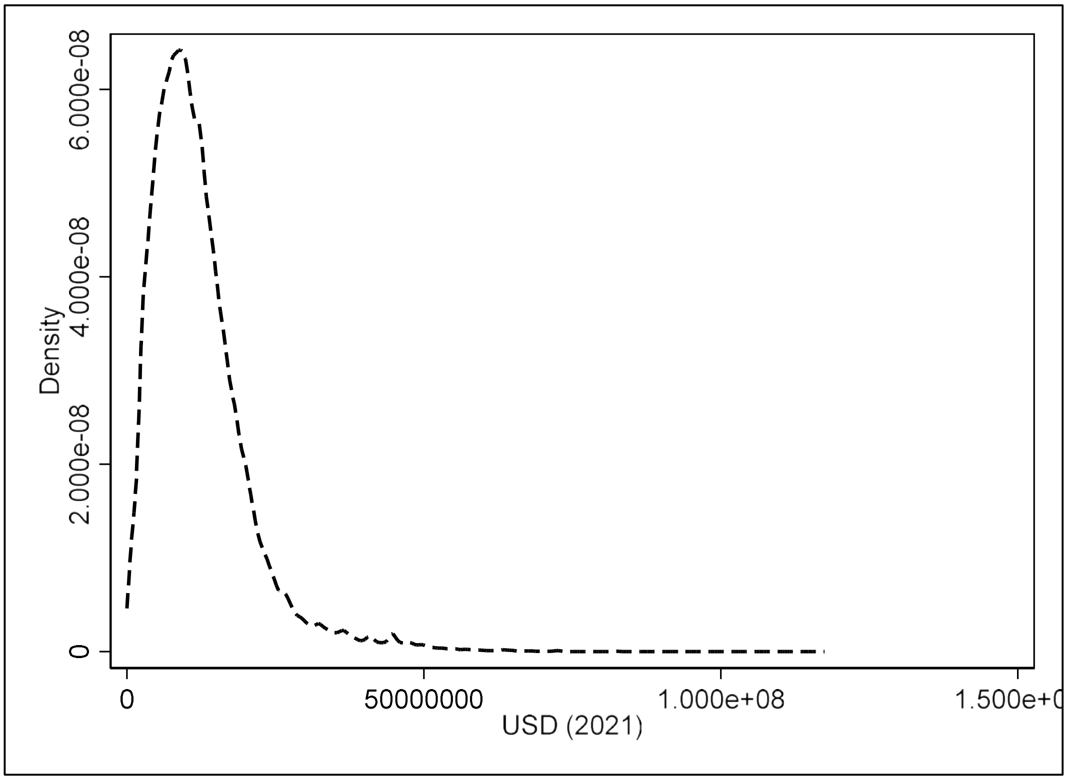

Notes: Figure A3 shows the probability density function of the income adjusted VSLs with unit weights.

Figure A4: Household Income Adjusted VSL with Distributional Weights

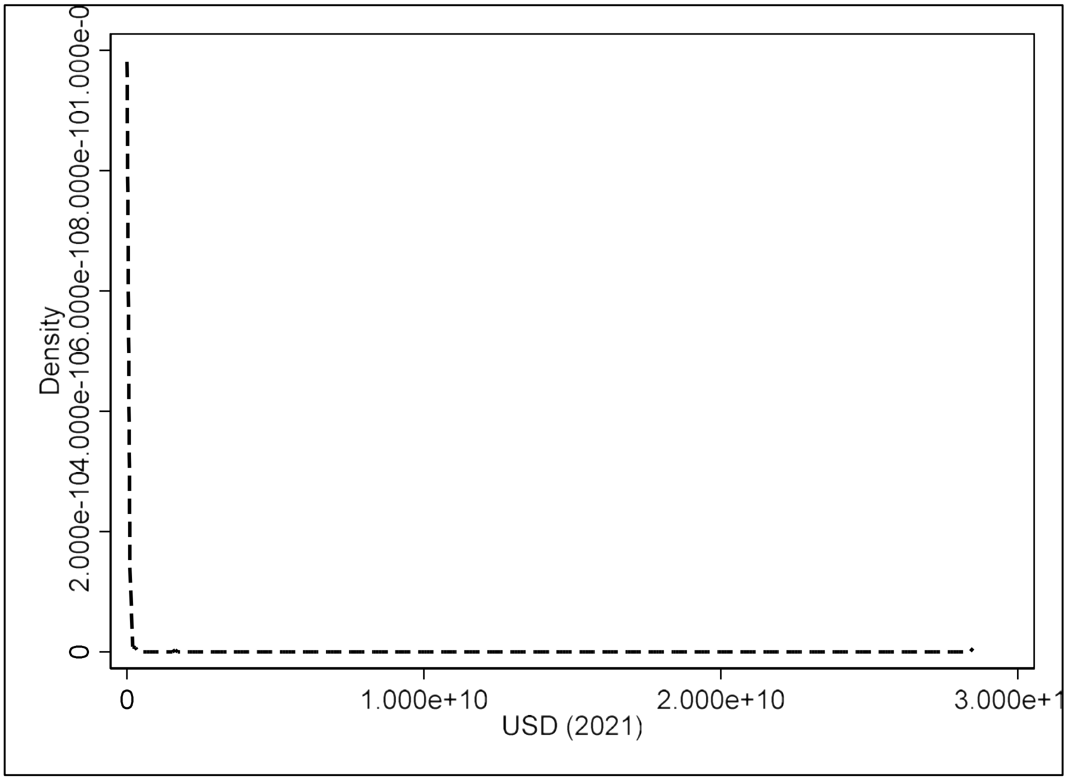

Notes: Figure A3 shows the probability density function of the income adjusted VSLs with distributional weights.

Figure A5: Household Income Adjusted VSL with Unit Weights and Household Income

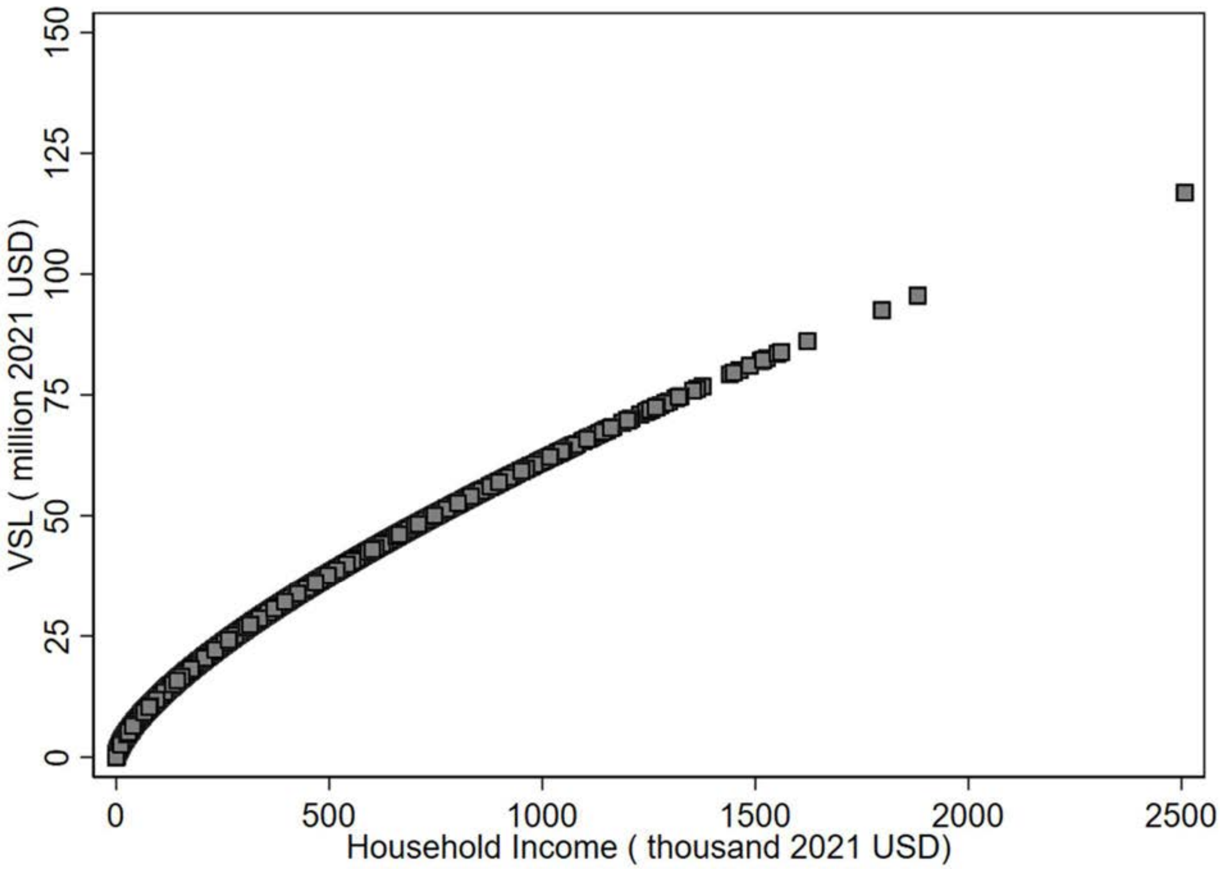

Figure A6: Household Income Adjusted VSL with Distributional Weights and Income

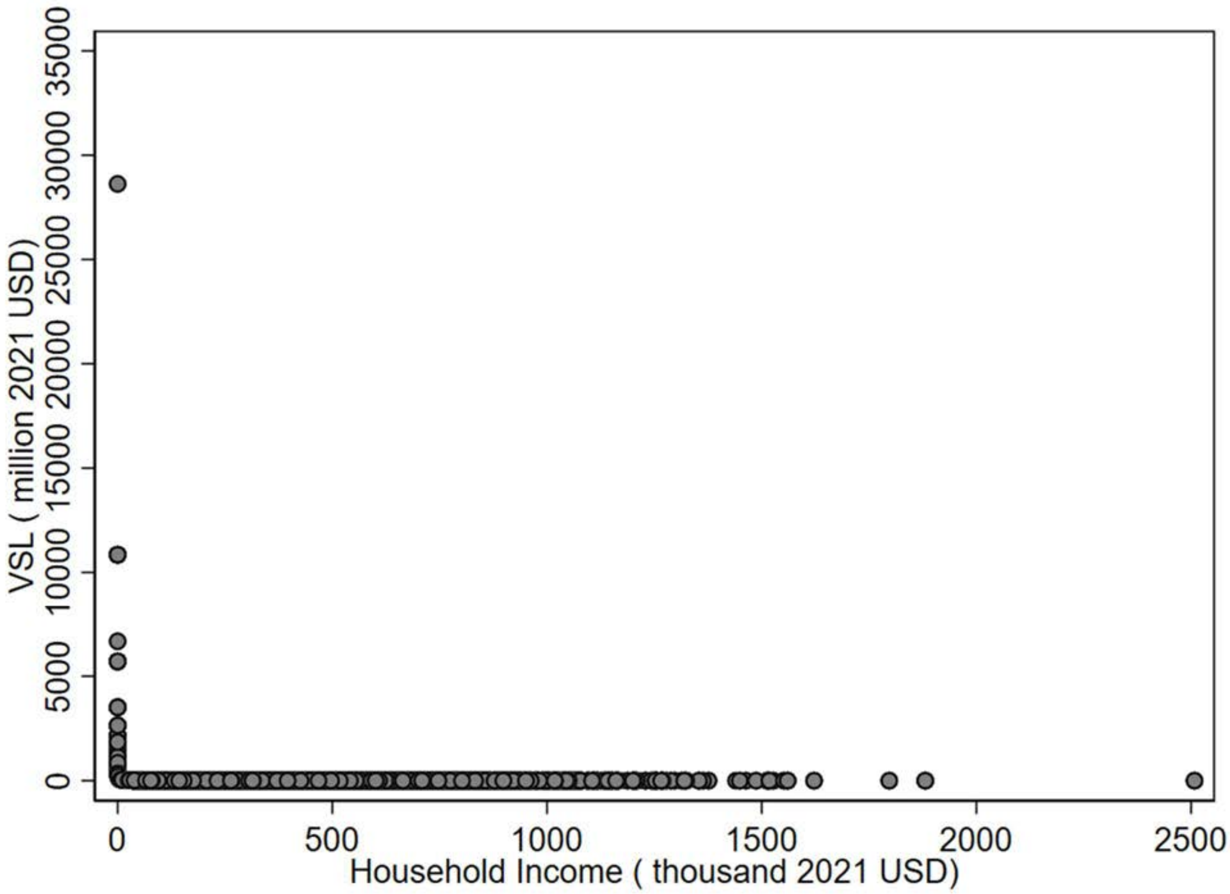

Figure A7: Monte Carlo Simulation Results: Net Benefits with Unit Weights

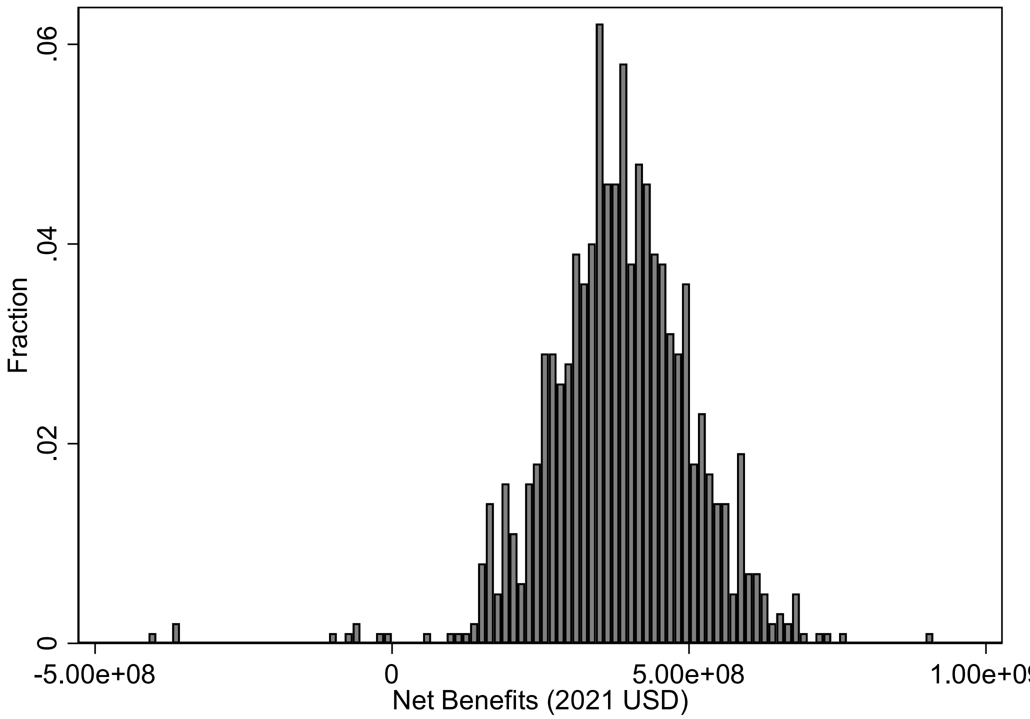

Notes: Figure A7 is a histogram showing the net benefits with unit weights. The Monte Carlo simulation consists of 1,000 iterations. The mean net benefit value is $394 million and the standard deviation is $323 million as reported in the top row of table 2.

Figure A8: Monte Carlo Simulation Results: Net Benefits with Distributional Weights

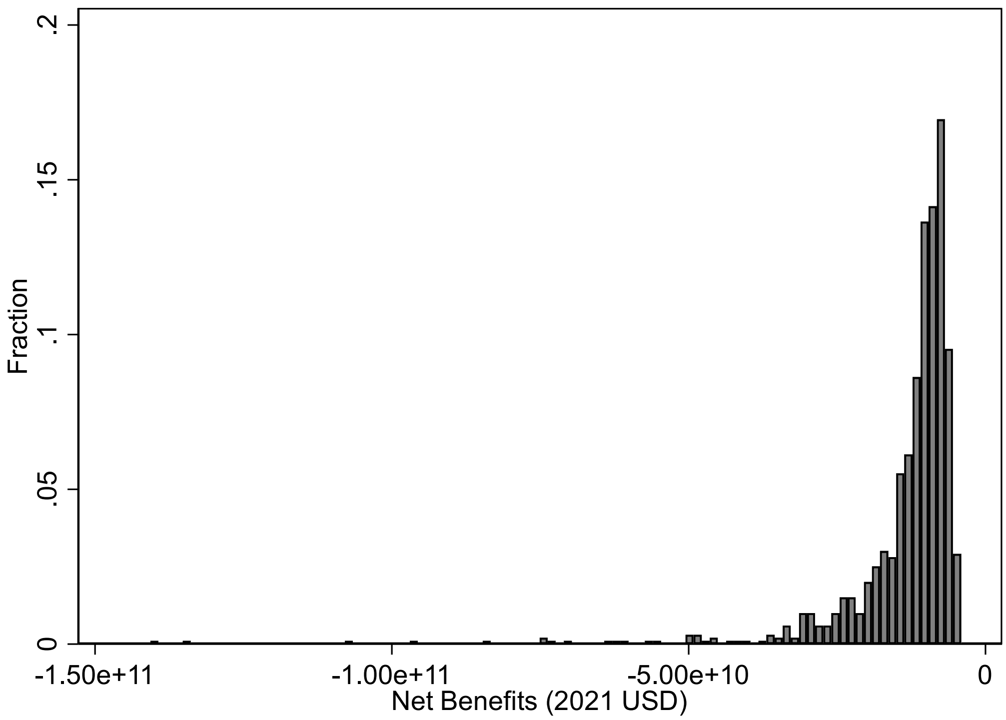

Notes: Figure A8 is a histogram showing the net benefits with distributional weights. The Monte Carlo simulation consists of 1,000 iterations. The mean net benefit value is -$11.6 billion and the standard deviation is $51.4 billion as reported in the top row of table 2.

Figure A9: Net Benefits with Distributional Weights and Cost for Lowest Income Household

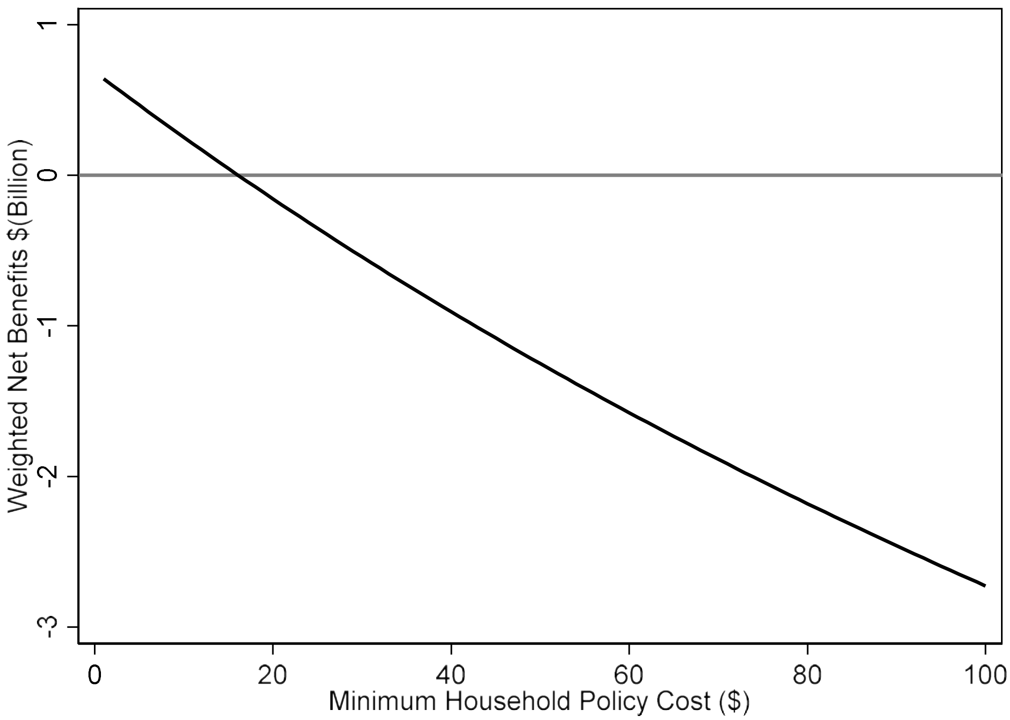

Notes: Minimum household policy cost is the cost for the lowest income household. Total policy costs are calculated as proportional to the share of total income across households. For households with near zero income, policy costs approach zero. In the simulation that produces figure A9, household costs cannot fall below the values shown on the horizontal axis. For each minimum cost, total weighted benefits and costs are calculated, and the net is reported on the vertical axis.