

NBER WORKING PAPER SERIES

ON THE USES OF BENEFIT-COST  
REASONING IN CHOOSING POLICY  
TOWARD GLOBAL CLIMATE CHANGE

David F. Bradford

Working Paper 5920

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
February 1997

Prepared as a response to the letter of invitation for the workshop on "Discounting in Intergenerational Decision making," sponsored by Resources for the Future (RFF) and the Energy Modeling Forum and held at RFF, Washington, DC, November 14-15, 1996. The operative portions of the letter have been attached as an appendix to the paper. In the spirit of the exercise, I have omitted references to the literature. I would, however, like to acknowledge a particular debt to Thomas C. Schelling's paper, "Intergenerational Discounting," Energy Policy, (XXIII-4/5) 1995, pp. 395-401. This paper is part of NBER's research program in Public Economics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

© 1997 by David F. Bradford. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

On the Uses of Benefit-Cost Reasoning in  
Choosing Policy Toward Global Climate Change  
David F. Bradford  
NBER Working Paper No. 5920  
February 1997  
JEL No. D61  
Public Economics

### **ABSTRACT**

In the debate about the correct discount rate to use in evaluating policy with regard to climate change, which covers the entire world and extends for centuries, the conditions for deploying benefit-cost analysis are often overlooked. Where (a) income distributional effects of policies are large and (b) one cannot take for granted compensating adjustment in other policy instruments affecting distribution, simple aggregation of gains and losses is unlikely to provide a convincing basis for action, as an ethical matter, or predictor of policy, as a political matter.

David F. Bradford  
Woodrow Wilson School  
Princeton University  
Princeton, NJ 08544-1013  
and NBER  
BRADFORD@PRINCETON.EDU

# ON THE USES OF BENEFIT-COST REASONING IN CHOOSING POLICY TOWARD GLOBAL CLIMATE CHANGE

David F. Bradford\*

If economists could manage to get themselves thought of as humble, competent people, on a level with dentists, that would be splendid. (John Maynard Keynes, 1931)

To understand discounting, and the application of benefit-cost reasoning more generally, in the context of setting policy toward global climate change, it may be helpful to start with an easier case. In the 1980s Voluntary Export Restraint (VER) arrangements applied to trade in automobiles between Japan and the United States. It is estimated that the effect was to raise the prices of Japanese automobiles sold in the United States on average about \$1,500 above what they would otherwise have been. Presumably, as a result of the policy, the price in the United States of automobiles manufactured in countries other than Japan, including the United States, was also somewhat higher than otherwise. Consequently, owing to the VER, U.S. buyers of automobiles suffered a loss of some significant amount, relative to the status quo. Those involved in the automobile industry, including owners of United States-based companies and their

---

\* Professor of Economics and Public Affairs, Princeton University. Prepared as a response to the letter of invitation for the workshop on "Discounting in Intergenerational Decision making," sponsored by Resources for the Future (RFF) and the Energy Modeling Forum and held at RFF, Washington DC, November 14-15, 1996. The operative portions of the letter have been attached as an appendix to the paper. In the spirit of the exercise, I have omitted references to the literature. I would, however, like to acknowledge a particular debt to Thomas C. Schelling's paper, "Intergenerational Discounting," Energy Policy, (XXIII-4/5) 1995, pp. 395-401.

employees as well as owners of Japanese companies and their employees, gained an amount roughly comparable to what U.S. consumers lost. The gains to the gainers presumably fell short of the loss to the losers, however. This is because production took place where it was more costly and because there was distortion in the composition of the output of automobiles.

Was the VER a good policy? What would benefit-cost analysis have to say? The bottom line of a benefit-cost analysis would be a measure of the net gain in the aggregate, encompassing all affected parties. This total would compare the outcome under the policy with the outcome (counterfactual) under the status quo. An important question is, outcome for whom? Let us suppose the analysis is global. Everyone's benefits and costs count.

As I have suggested, the net benefit of the policy in this case was probably negative from a global perspective (and certainly so from a national perspective). In the language of policy analysis, the policy failed the benefit-cost test and resulted in a "welfare loss." A hard-line advocate of benefit-cost as a decision criterion would say that the policy therefore should not have been implemented. I believe there is a reasonable case made for adopting a benefit-cost test as a decision criterion, or at least a presumptive criterion. But the strict logic of what we learn from the benefit-cost exercise does not allow us to draw any conclusion about whether a policy that passes the test *should* be implemented. Nor does it tell us whether a policy that fails the test *should* be rejected. For somewhat similar reasons, passing a benefit cost test is neither necessary nor sufficient to predict policy will be accepted in the actual political process.

The reason is well known, but sometimes forgotten. The simple aggregate of benefits and costs does not tell us anything about who enjoys the benefits or bears the costs. In the case of the VER, beneficiaries were stakeholders in the automobile industry, especially in the United States and Japan. The losers were U.S. consumers. Economic science does not tell us whether the

resulting redistribution of wealth – from U.S. consumers to stakeholders in the automobile industry – was good or bad.

The fact that the aggregate net benefit of the policy was negative does, however, contain some useful information. It tells us that there was an alternative policy to the VER that would have been better for everyone. That is, a team of economists could have dreamed up an alternative way to deliver the positive benefits to the auto industry's stakeholders at a lower cost to U.S. consumers. Whether the required alternative policy would have been in accord with the U.S. Constitution, or acceptable to a majority of members of the House of Representatives, or explicable to David Brinkley, are questions on which economics is silent. (There is a pretty strong presumption among economists -- call it optimism -- that if a way could have been found that satisfied such political constraints and that did give a better deal to everyone, it would have been adopted.)

Notice that to undertake this analysis, no information is needed about who ought to get more or less, who is more or less deserving. Nor is it other than a technical question what prices should be used to evaluate the changes in consumption of the various people with an interest in the outcomes – auto workers and auto buyers, for example. These are matters about which there is plenty of room to argue – does the degree of monopoly in the automobile or other markets affect the way one should value changes in the outcomes? – but the issues are not at a deep philosophical level.

To go further, to say that a policy is a good or bad thing, requires us to weigh in the scales the gains to the gainers and the losses to the losers. In the utilitarian tradition, I am inclined to apply as a rough metric the sort of “declining marginal social utility of consumption” that is the usual justification for policies that distribute from richer to poorer. I try, however, to keep

straight that this is my metric, or maybe even the dominant metric in my group of professional colleagues, not one based in value-free science.

There are, however, other ways to deal with the distributional aspect of policies that would permit us to go further on a more scientific basis. First, one might argue that routinely basing decisions on the aggregate net benefits will produce, on average for everyone, a better result than case-by-case decisions in a political setting. The free enterprise economy chooses projects on essentially benefit-cost grounds. The fact that competitors will lose if I bring out a superior mousetrap does not affect my decision to do so. It would not be easy to write down a convincing model of recurring decision situations to prove rigorously that obeying the benefit-cost criterion produced a better result on average than alternative methods. But the success of competitive systems in raising living standards, by comparison with other systems, suggests the idea is not to be dismissed out of hand either.

A stronger presumption for the benefit-cost standard obtains if distributional concerns can be summed up in terms of something like “earning power” or “ability to pay” and a tax and transfer system is in place that aims to choose the best feasible outcome in terms of this measure. If “you,” the decision-maker, are also in charge of the tax-transfer system, to a first approximation you will want to make allocative decisions, such as whether to invoke VER, on the basis of benefit-cost calculus. (Again, there are many technical caveats, but I believe the basic thrust of the proposition is correct as a matter of logic, not deeper philosophy.) An alternative interpretation is descriptive, rather than prescriptive. The political process may be thought of as maintaining some sort of equilibrium in the distribution between richer and poorer. Here richer and poorer is measured by the criteria of the tax-transfer system (e.g., income as defined by the

tax code). In this setting you will tend to be correct in predicting that policies with positive net benefits will be adopted.

Notice, however, that the condition of a redistributive tax-transfer system based on the criterion of who ought to have the greater weight in the decision did not apply in the VER case. The gainers from the policy were not identified by their income but rather by their connection to an industry (and geographical region). Since we may presume that the gains to the Japanese companies and workers did not get much weight in the calculus of U.S. decision makers, we must guess that the relative weights on the U.S. residents in different locations were unrelated to income as defined by the tax and welfare systems. So benefit-cost reasoning leads to a bad prediction of the outcome. The outcome chosen by taking the criterion of maximizing net benefits in the aggregate would also have been wrong from an ethical point of view, if one accepts the political equilibrium as expressing a social judgment of equity.

Let us apply the lessons of this simple example to the very different case of global climate change. The starting point is some notion of the baseline policy path. This is not a wholly obvious concept. Does the baseline path incorporate any sort of corrective adjustment in response to changing circumstances? It is usual to invoke a business-as-usual path that is a more or less well specified version of doing what we are now doing. Alternatives are then considered, typically achieving reduced emissions of greenhouse gases by one or another set of mechanisms. An estimate is then developed of the value of period-by-period benefits, which may be due to bad things avoided, such as flooding, species destruction, etc. These are netted against the costs, which may be due to greater sacrifice to production of things that take energy or foregone good effects of climate change, such as higher agricultural productivity in northern regions. The thing evaluated in any given period will, incidentally, not be a fixed outcome, but rather a probability

distribution of outcomes, depending on resolution of the many sources of uncertainty that inevitably attend projections into the future.

This is a formidable research problem (which is not to suggest that pulling apart the VER is child's play). To focus on the matter of discounting, let us suppose, however, it has been satisfactorily solved. We have, for each time point, the equivalent in dollars to the impact of the policy, relative to the status quo, for everybody around at the time. How should we aggregate these quantities for our benefit-cost analysis? Neglecting technical quibbles that I believe are not likely to upset the story very much, we should do the same thing we did in the VER case. There we used the going prices for food, shelter, automobiles, etc., to evaluate the changes in consumption of the various interested parties. Here we would use the market prices for assets with the same profiles of returns to price out the quantities obtaining at different times.

One issue that is usually raised in connection with discounting is dealing with risk. I have attempted to integrate this issue into my description of an ideal benefit-cost analysis. The thing to be evaluated is the impact of a policy on the probability distribution of outcomes experienced by people in the future. If the impact is like a risky security, it should be evaluated as such. It is, however, important to remember that the fact that a security has a highly variable outcome does not reduce its value if its variability is inversely related to other risky outcomes. We call such a security an insurance policy.

It is an understatement to say that we can have little confidence in predictions about economic affairs many decades in the future. In making our best guesses about the impact of climate policy on the distant future, however, it is the correlation with the distribution in the absence of policy that counts. Interventions to moderate climate change will work to pay off more when the outcome would otherwise be worse. Therefore, considerations of risk are likely to



affect the calculation of net benefits in a way similar to reducing the discount rate used in the analysis.

When I say that our best guess of asset market prices “should” be used in the analysis, the force of the “should” is logical, not ethical. That is, by using market prices we obtain a result that allows us to answer the question: Could we combine implementation of the policy in question (for example, cutting emissions of carbon dioxide by switching to nuclear fuel) with redistribution so as to make everyone better off? As in the case of the VER, this is, strictly speaking, the question that economics can claim to answer. As in the case of the VER, there is no necessary implication that a policy that fails this test is “bad” and should not be adopted. Someone might claim that, the way life really works, the policy will affect the distribution of well-being. The induced changes in wealth will not be adjusted by the tax-transfer function of government, or, at least, will not be adjusted in the dimension of particular concern to this observer. The beneficiaries in the process will be particularly worthy and the losers less deserving, so a project that fails the test is a good thing on balance. This apparently described the VER for the U.S. Congress.

There seems to be a presumption in connection with global climate change that policies undertaken today will benefit future generations at the cost of consumption sacrificed by present generations. That is, distributional consequences are inextricably bound to the allocational issue of controlling greenhouse emissions. This connection is, however, not logically necessary. Suppose there were adequate means of adjusting the outcomes for their distributional effects. Then the same presumption in favor of the benefit-cost criterion in making decisions would apply to choice of policies toward global warming as in other areas.

What would such a redistributional procedure involve? Consider the content of a benefit estimate. A typical specification would be the amount of money the affected person would pay to obtain the impact in question, rather than experience the business as usual world. (If the mitigating policy were taken as the baseline, the benefit estimate would be the amount one would require to receive to be compensated for not obtaining the impact in question.) For example, a population that would have to move to a higher location, or build dikes, or suffer higher risks of loss of life due to severe storms, would be prepared to pay some amount for an insurance policy that would mitigate these effects. The redistribution required to offset the adoption of a policy of mitigation would be a transfer away from this population toward people who would be negatively affected by the policy – presumably those living today who would sacrifice as a result of, say, a carbon tax. “Transfer away” would take the form of reducing the inheritance of the group reaping the benefit of the policy.

Apart from its apparent mean-spiritedness, this description suggests the part of the puzzle that is lacking to make benefit-cost analysis a likely sufficient guide to action or likely predictor of political acceptance of mitigating policies. It may be reasonable to think that there is a systematic redistributive mechanism in place within the United States, both at a given time and over time. There is, however, clearly no such mechanism connecting results at great geographical distance, either now or over time.

In the case of the VER, the aggregate of benefits net of costs proved a bad predictor of the collective choice. Taking the Congressional judgment as the measure of social welfare, it also proved a bad indicator of the socially better choice. For much the same reason, the aggregate net benefit of climate change, whether correctly discounted or not, is unlikely to provide enough

information for either predicting or guiding choices. What must be added is information about the distribution of gains and losses, over time and across space.

By doing so, we would draw attention to the potential for developing compensating changes that might be necessary to make acceptable politically a policy that would otherwise be vetoed by a losing group. We could, for example, run larger budget deficits in order to reduce the inheritance of future generations. We might be induced to sympathize with a population distant in time and place if we were convinced they would be very poor and that they could be helped by our sacrifice today.

My remarks about discounting presuppose that it has been sensible or feasible to translate the impacts of policies into monetary terms. It seems to me, this is the toughest part. The hard questions are how to value the possible annihilation of species or the flooding of areas of human settlement. It is not likely that we can be completely successful assigning money values to such impacts. To be sure, at some level, we make such implicit tradeoffs every day. Decisions about public safety, national defense, intervention in foreign wars or natural disasters, and any number of other matters of high and low policy require balancing of money costs and nonmonetary benefits. Benefit-cost analysis can at best help us put such matters in perspective. An example would be comparing the costs of saving a statistical life by mandating seat belt use and regulating the cleanup of abandoned waste dumps.

One has to be impressed with the uncertainty that exists about the shape of human civilization in two or three hundred years. What would we have said in 1775 if we happened to live in the British colonies on the east coast of North America about what life would look like in the same place in 1996? What can we now say about what life will look like here, or in the

Middle East, or sub-Saharan Africa in 2150? With this vagueness about the distant future in mind, let me conclude with a brief response to the four questions posed to us by the conveners:

- We have no real choice but to treat projects whose effects will be spread out over hundreds of years as longer versions of projects having impacts over a mere forty years.
- To a first approximation, the appropriate way to discount probability distributions of consumption for purposes of determining a policy's potential for making everybody better off is to use market discount rates appropriate for the distributions in questions.
- Because mechanisms do not exist for adjusting the distribution of gains and losses more or less automatically, to assess a policy requires information about gains and losses of specific groups or cohorts through time. The most practical way to describe such distributive effects is in terms of the value of consumption, discounted to the present. These are the magnitudes we need to know to assess the alternative distributions of benefits and costs that are available. The effects cannot be reliably added together, however.
- If its logic is well understood, benefit-cost analysis should be very helpful in decision making on even such complex issues as global climate change. The most serious problem is evaluating hard-to-value things. Many policy alternatives, such as the choice of instrument to control greenhouse gas emissions do not raise such difficult questions.

Appendix

Excerpt from the August 15, 1996 letter from John P. Weyant and Paul R. Portney to the invited contributors to the workshop on “Discounting in Intergenerational Decision making,” sponsored by Resources for the Future (RFF) and the Energy Modeling Forum and held at RFF, Washington DC, November 14-15, 1996:

We are hoping you will be willing to prepare a short paper (no more than 10 double-spaced pages) and present it at the workshop. Unlike most other conferences and workshops, however, we will be asking all the participants to address the same set of questions in their papers, and the discussion following each presentation will focus at least in part on the degree to which the views just presented are in accordance or divergence with those of other participants. The questions we hope you will address (ideally, in a way that will be accessible to non-economists) include the following:

- Should projects whose effects will be spread out over hundreds of years or more be treated simply as “longer versions” of projects whose principal effects extend no more than, say, thirty or forty years?
- If the answer to the above question is “yes”, what is the appropriate way to determine the discount rate to be used?
- If projects with significant intergenerational effects are to be evaluated differently, how should this be done? Should benefits and costs in the distant future not be discounted at all? At a different rate?
- Perhaps more fundamentally, is it appropriate to use benefit-cost analysis at all in decision making on such issues as climate change, disposal of high-level nuclear wastes, etc.?