

CHOOSING PRICE OR QUANTITY CONTROLS FOR GREENHOUSE GASES

William Pizer

Climate Issues Brief No. 17

July 1999

Resources for the Future
1616 P Street, NW
Washington, DC 20036
Telephone 202-328-5000
Fax 202-939-3460

Climate Issue Briefs are short reports produced as part of RFF's Climate Economics and Policy Program to provide topical, timely information and analysis to a broad non-technical audience.

The preparation of these briefs is funded in part by The G. Unger Vetlesen Foundation.

© 1999 Resources for the Future. All rights reserved. Reproduction of this paper is permitted with attribution and copyright citation.

Choosing Price or Quantity Controls for Greenhouse Gases

William Pizer, Resources for the Future

I. Introduction

Much of the debate surrounding climate change has centered on verifying the threat of climate change and deciding the magnitude of an appropriate response. After years of negotiation, this effort led to the 1997 signing of the Kyoto Protocol, a binding commitment by industrialized countries to reduce their emissions of carbon dioxide to slightly below 1990 recorded levels. Without approving or disapproving of the response effort embodied in the Kyoto Protocol, I believe that an important element has been ignored. Namely, should we specify our response to climate change in terms of a quantitative target?

The appeal of a quantitative target is obvious. A commitment to a particular emissions level provides a straightforward measure of environmental progress as well as compliance. Commitment to an emissions tax, for example, offers neither a guarantee that emissions will be limited to a certain level nor an obvious way to measure a country's compliance (when other taxes and subsidies already exist). Yet, it is precisely this concern which points to an important observation.

Quantity targets guarantee a fixed level of emissions. Emission taxes guarantee a fixed financial incentive to reduce emissions. Both can be set at either aggressive or modest levels. Aside from the appeal of the known and verifiable emissions levels that quantity targets can ensure, might there be other important differences between price and quantity controls?

Economists would say yes. With uncertain outcomes and policies that are fixed for many years, it is important to carefully consider both the costs and benefits of alternate price and quantity controls in order to judge which is best. My own analysis of the two approaches indicates that price-based greenhouse gas (GHG) controls are much more desirable than quantity targets, taking into account both the potential long-term damages of climate change and the costs of GHG control. This can be argued on the basis of both theory and numerical simulations. Based on the latter, I find that price mechanisms produce expected net gains *five times higher* than even the most favorably designed quantity target.

To explain this conclusion, I first characterize the differences between price and quantity controls for GHGs. I then present both theoretical and empirical evidence that price-based controls are preferable to quantity targets based on these differences. Finally, I discuss how price controls can be implemented *without* a general carbon tax. This last point is particularly salient for the United States, where taxes are generally unpopular. The "safety valve," as it is often called, involves a cap-and-trade GHG system accompanied by a specified fee or penalty for emissions beyond the initial cap.

II. How do quantity- and price-based mechanisms work?

A quantity mechanism—usually referred to as a permit or cap-and-trade system—works by first requiring individuals to obtain a permit for each ton of carbon dioxide they emit, and then limiting the number of permits to a fixed level.¹ This permit requirement could be imposed on the individuals who actually release carbon dioxide into the atmosphere by burning coal, petroleum products, or natural gas. However, unlike emissions of conventional pollutants which depend on a variety of other factors, carbon dioxide emissions can be determined very accurately by the volume of fuel being used. Rather than requiring *users* of fossil fuels to obtain permits, we could therefore require *producers* to obtain the same permits. This has the advantage of involving far fewer individuals in the regulatory process, thereby reducing both monitoring and enforcement costs (see the papers by Carolyn Fischer, Suzi Kerr and Michael Toman in Further Readings). This type of system has been used with considerable success in the United States to regulate both sulfur dioxide and lead.

A key element in a permit system is that individuals are free to buy and sell existing permits in an effort to obtain the lowest cost of compliance for themselves, in turn leading to the lowest cost of compliance for society. In particular, when individuals observe a market price for permits, those that can reduce emissions more cheaply will do so in order to either sell excess permits or avoid having to buy additional ones. Similarly, those who face higher reduction costs will avoid reductions by either buying permits or keeping those they already possess. In this way, total emissions will exactly equal the number of permits while only the cheapest reductions are undertaken.

A price mechanism—usually referred to as a carbon tax or emissions fee—requires the payment of a fixed fee for every ton of CO₂ emitted. Like the permit system, this fee could be levied upstream on fossil fuel producers or downstream on fossil fuel consumers. Either way, we associate a positive cost with emissions of CO₂ and create a fixed monetary incentive to reduce emissions. Such price-based systems have been used in Europe to regulate a wide range of pollutants (although the focus is usually revenue generation rather than substantial emissions reductions).

Like a tradable permit system, price mechanisms are cost-effective. Only those emitters who can reduce emissions at a cost below the fixed fee or tax will choose to do so. Since only the cheapest reductions are undertaken, we are guaranteed that the resulting emission level is obtained at the lowest possible cost.

The important distinction between these two systems is how they adjust when costs change unexpectedly. A quantity or permit system adjusts by allowing the permit price to rise or fall while holding the emissions level constant. A price or tax system adjusts by allowing the level of total emissions to rise or fall while holding the price associated with

¹ Here and throughout this brief, we discuss policies designed to limit carbon dioxide emissions from fossil fuel sources. These emissions constitute the bulk of GHG emissions and are the general focus of most policy discussions. Regardless, the arguments made in this context apply equally well to the regulation of GHG emissions more broadly defined.

emissions constant. Ignoring uncertainty and assuming we know the costs of controlling CO₂, both policies can be used with the same results. Consider the following example:

Suppose we *know* that with a comprehensive domestic CO₂ trading system in place in the United States by the year 2010, a permit volume of 1.2 gigatons of carbon equivalent emissions (GtC) will lead to a \$100 permit price per ton of carbon. (1998 US emissions of carbon from fossil fuels are estimated at 1.5 GtC.) In other words, faced with a price incentive of \$100 per ton to reduce emissions, regulated firms in the United States will find ways to reduce emissions to 1.2 GtC. Therefore, the same outcome can be obtained by imposing a \$100 per ton carbon tax.

III. Uncertainty about costs


In reality, we have only a vague idea about the permit price that would occur with emissions of 1.2 GtC or any other emission target. There are three reasons why such costs are hard to pin down. First, little evidence exists concerning reduction costs. There are no recent examples of carbon reductions on a substantial scale from which to base estimates. In the 1970's, energy prices doubled and encouraged increased energy efficiency, but these events occurred both in a context of considerable uncertainty about the future and alongside many other confounding factors (such as increased environmental regulation). Alternatively, engineering studies provide a bottom up approach to estimating costs. However, comparisons of past engineering forecasts to actual implementation costs suggest that they are inaccurate at best (see work by Winston Harrington and Richard Morgenstern under Further Readings).

A second source of uncertainty arises because we need to forecast compliance costs in the *future*. This involves difficult predictions about the evolution of new technologies. Proponents of aggressive policy argue that reductions will be cheap as new low-carbon or carbon-free energy technologies become available. Proponents of more modest policies argue that these are unproven, pie-in-the-sky technologies that may never be practical.

Finally, it is impossible to know how uncontrolled emission levels will change in the future. That is, to achieve 1990 emission levels in 2010, it is unclear whether reductions of 5, 25 or even 50% will be necessary. The Intergovernmental Panel on Climate Change (IPCC), the international agency charged with studying climate change, gives a range of six possible global emission scenarios in the year 2010 that include a low of 9 GtC and a high of 13 GtC. My own simulations suggest a broader possible range, from 7 to 18 GtC.

The low end of both ranges reflects the possibility that population and economic growth may slow in the future and the energy intensity of production may fall. The high end reflects the opposite possibility, that growth is high and energy intensity rises. Figure 1 shows the distribution of uncontrolled emissions arising from my simulations of one thousand possible outcomes in 2010 alongside the six IPCC scenarios. (For details on the modeling, see paper by Pizer in Further Readings.)

Figure 1: Distribution of Emissions in 2010



In summary, there are two important reasons why we have only vague ideas about the cost of alternative emission targets. First, there is little historic evidence on costs. Second, as we examine policies ten or more years in the future, it is unclear how both baseline emissions and available technologies will change between now and then. From the preceding figure, global emissions could be anywhere from 7 to 18 GtC in 2010. The cost associated with a 8.5 GtC (1990 level) target will be uncertain both because the necessary reduction is uncertain—somewhere between zero and 10 GtC—and because even knowing the reduction level, costs are difficult to estimate.


IV. Effects of price and quantity controls with cost uncertainty

When the cost of a particular emission target is uncertain, price and quantity controls will have distinctly different consequences for the actual level of emissions as well as the overall cost of a climate policy. Even if both policies are designed to deliver the same results under a best guess scenario, they will necessarily behave differently when control costs deviate from this best guess. These differences arise because a price policy provides a fixed \$/ton incentive regardless of the emission level, while a quantity policy generates whatever incentive is necessary in order to strictly limit emissions to a specified level.

Figure 2 illustrates these differences by showing the emission consequences in 2010 associated with two policies that are roughly equivalent under a best-guess scenario: a quantity target of 8.5 GtC and a carbon tax of \$80 per ton. Using the same one thousand emission scenarios shown in Figure 1, simulations are used to calculate the effect of these two policies for each outcome. With a carbon tax, the left panel indicates that emissions

are below 8.5 GtC in over 75% of the outcomes. In other words, on average the carbon tax achieves *more* reductions than a quantity target of 8.5 GtC. Sometimes, the reductions are much more: note that emissions may be as low as 3 GtC. Yet, the carbon tax fails to *guarantee* that emissions will always be below any particular threshold.

Figure 2: Effect of Price and Quantity Controls on Emissions in 2010

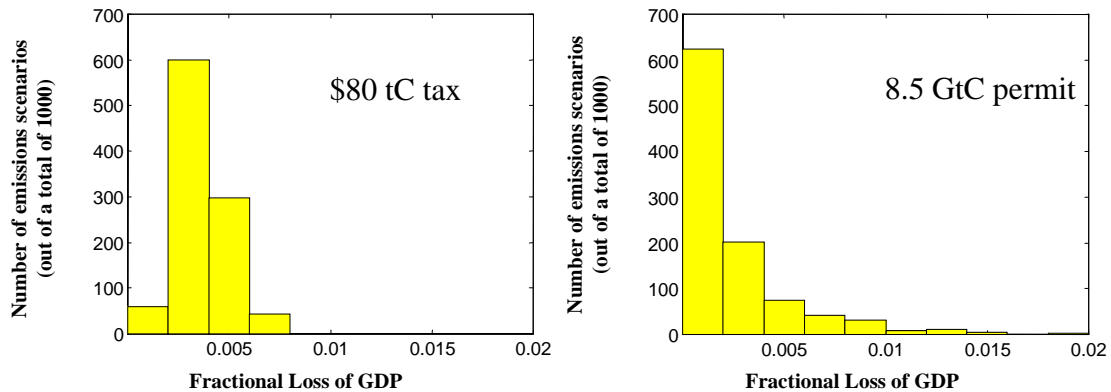


The quantity target, in contrast, never results in emission levels above 8.5 GtC. Since some emission outcomes in the absence of controls were rather high, on the order of 18 GtC, we would expect that the cost of this policy could be quite high. At the other extreme, the quantity policy could be costless if uncontrolled emissions are unexpectedly low.

This suggests that the cost associated with quantity controls will be high or low depending on future reduction costs as well as the future level of uncontrolled emissions. In contrast, price controls create a fixed incentive to reduce each ton of carbon dioxide regardless of the uncontrolled emission level. Therefore, costs under a carbon tax should fluctuate much less than costs under a quantity control.

With this distinction in mind, Figure 3 shows the estimated cost consequences of both policies. The range of costs associated with the quantity target is quite wide as we suspected. The estimates extend from zero to 2.2% of global gross domestic product (GDP). That is almost four times higher than the highest cost outcome under the carbon tax. In fact, the cost associated with emission reductions under a carbon tax are concentrated entirely in the range 0.2% to 0.6% of GDP. Since the carbon tax always applies the same per ton incentive to reduce emissions, the cost outcomes are more narrowly distributed than those occurring under a quantity target.

Figure 3: Distribution of 2010 Costs associated with Price and Quantity Controls



V. Choosing between price and quantity controls

So far the discussion has been limited to the different emission and cost consequences of alternative price and quantity controls. Choosing between them, as well as choosing the appropriate stringency of either policy, requires making judgments about climate change consequences as well as control costs. In order to help us understand when one policy instrument will likely be preferred to the other, it is useful to consider two extreme cases.

First, imagine that there is a known climate change *threshold*. When carbon dioxide emissions are below this threshold, the consequences are negligible. Above this threshold, however, damages are potentially catastrophic. For example, research suggests that the process by which carbon dioxide is absorbed at the surface of the oceans and circulated downward could change dramatically under certain circumstances (see article by Broecker listed in Further Readings). If we further believe that these changes will have severe consequences *and* we can identify a safe emission threshold for avoiding them, then quantity controls seem preferable. Quantity controls can be used to avoid crossing the threshold and, in this case, large expenditures in order to meet the target are justified by the dire consequences of missing it.

Now, imagine instead that every ton of carbon dioxide emitted causes the same incremental amount of damage. These damages might be very high or low, but the key is that each ton of emissions is just as bad as the next. Such a scenario is also plausible, as indicated by a survey of experts including both natural and social scientists who do research on global warming. Their beliefs suggest that the damage caused by each ton of emitted CO₂ may be quite high but that there is no threshold: damages are essentially proportional to emissions. Each additional ton is equally damaging, whether it is the first ton emitted or the last (Tim Roughgarden and Steven Schneider discuss this survey, originally conducted by William Nordhaus; see Further Readings for both references).

In this case, it makes sense to use a price instrument. Specifically, a carbon tax equal to the damage per ton of CO₂ will lead to exactly the right balance between the cost of

reducing emissions and the resulting benefits of less global warming. Every time a firm decides to emit CO₂, it will be confronted with an added financial burden equal to the resulting damage. This will lead to reduction efforts as well as investments in new technology that are commensurate with the alternative of climate change damage. In this scenario, little emphasis is placed on reaching a particular emission target because there is no obvious quantity target to choose. This argument applies even if we are uncertain about the magnitude of climate damage per unit of CO₂.

VI. Arguments for Price Policies

Given this characterization of circumstances under which alternative price and quantity mechanisms are preferred, we can now make the argument for price controls. This argument hinges on two basic points. The first point is that climate change consequences generally depend on the *stock* of greenhouse gases in the atmosphere, rather than annual emissions. Greenhouse gases emitted today may remain in the atmosphere for hundreds of years. It is not the level of annual emissions that matters for climate change, but rather the total amount of carbon dioxide and other greenhouse gases that have accumulated in the atmosphere. The second point is that while scientists continue to argue over a wide range of climate change consequences, few advocate an immediate halt to further emission. For example, the most aggressive stabilization target discussed by the IPCC is 450 ppm (roughly 1035 GtC), a level that we will not reach before 2030 even in the absence of emission controls (see the Technical Summary provided in the IPCC report listed in Further Readings).

If only the stock of atmospheric GHGs matters for climate change, and if experts agree that the stock will grow at least in the immediate future, there is virtually no rationale for quantity controls (for further discussion see my paper with Richard Newell in Further Readings). The fact that only the stock matters should first draw our attention away from short-term quantity controls for emissions and toward long-term quantity controls for the *stock*. It cannot matter whether a ton of CO₂ is emitted this year, next year or ten years in the future if all we care about is the total amount in the atmosphere. Taking the next step and presuming that the stock will grow over the next few decades, this suggests that there is some room to rearrange emissions over time and that a short-term quantity control on emissions is unnecessary.

Quantity controls derive their desirability from situations where strict limits are important, when dire consequences occur beyond a certain threshold. Such policies trade off lower expected costs in favor of strict control of emissions in all possible outcomes. However, under the assumption that it is acceptable to allow the stock of greenhouse gases to grow in the interim, there is no advantage to such strict control. We give up the flexible response of price controls without the benefit of an avoided catastrophe.

Even for those who believe the consequences of global warming will be dire and that current emission targets are not aggressive enough, price policies are still better. An aggressive policy designed to *eventually* stabilize the stock does not demand a strict limit on emissions before stabilization becomes necessary. Additional emissions this year are

no worse than emissions next year. Why not abate more when costs are low and less when costs are high—exactly the outcome under a price mechanism? When we eventually move closer to a point where the stock must be stabilized, a switch to quantity controls will be appropriate.

In addition to these theoretical arguments, one can also turn to integrated assessment models for support. To this end, I have constructed an integrated model of the world economy and climate based on the DICE model developed by William Nordhaus. In contrast to the DICE model, I simultaneously incorporate uncertainty about everything from growth in population and energy efficiency to the cost of emission reductions, to the sensitivity of the environment to atmospheric CO₂ and the damages arising from global warming.

The results of these simulations indicate the price-based mechanisms can generate overall economic gains (expected benefits minus expected costs) that are *five times higher* than even the most prudent quantity-based mechanism. These results are robust. Even allowing for catastrophic damages beyond three degrees centigrade of warming, price mechanisms continue to perform better. This robustness can be explained in two ways. First, the catastrophe, if it exists, lies in the future. Before we reach that point, it is desirable to have some flexibility in emission reductions. Specifically, one will want to delay those reductions if the costs are unexpectedly high in the short run, provided those reductions can be obtained more cheaply in the future but before the catastrophe.

Second, unlike the earlier, stylized description where climate consequences depend directly on CO₂ concentrations, in this model damages instead depend on temperature change. In reality, damages probably depend on an even more complex climatic response. Either way, the link between CO₂ emissions, concentrations, temperature change and other climatic effects are not precisely known. Therefore, a quantity control on *emissions* is not equivalent a quantity control on *climate change*. Both price and quantity controls will lead to uncertain climate consequences. Therefore the advantage of the quantity control—namely its ability to avoid with certainty the threat of climate catastrophe—is substantially weakened.

VII. Combined price and quantity mechanisms

Even if a carbon tax is preferable to a cap-and-trade approach in terms of social costs and benefits, this policy obviously faces steep political opposition in the United States. Businesses oppose carbon taxes because of the transfer of revenue to the government. Under a permit system there is a hope that some, if not all, permits would be given away for free. Environmental groups oppose carbon taxes for an entirely different reason: they are unsatisfied with the prospect that a carbon tax, unlike a permit system, fails to guarantee a particular emission level. Such antagonism from both sides of the debate makes it unlikely that a carbon tax will become part of the US response to the Kyoto protocol.

However, the advantages of a carbon tax can be achieved without the baggage accompanying an actual tax. In particular, a combined mechanism—often referred to as a hybrid or “safety-valve”—can obtain the economic advantages of a tax while preserving at least some of the political advantages of a permit system (other concerns about the revenue aspects of different policies have been discussed by Ian Parry; see Further Readings).

In such a scheme, the government first distributes a fixed number of tradable permits either freely, by auction, or both. The government then provides additional permits to anyone willing to pay a fixed ceiling or “trigger” price. The initial distribution of permits allows the government the flexibility to give away a portion of the right to emit CO₂, thereby satisfying concerns of businesses about government revenue increases. The sale of additional permits at a fixed price then gives the permit system the same compliance flexibility associated with a carbon tax.

With a combined price/quantity mechanism, it will be necessary to consider how both the trigger price and the quantity target should evolve over time. One possibility is to raise the trigger price over time in order to *guarantee* that the quantity target is eventually reached. A second possibility is to carefully choose future trigger prices as a measure of how much we are willing to pay to limit climate change. As we learn more about the costs of future emission reductions, however, this distinction between price and quantity controls will diminish. That is, once uncertainty about future compliance costs is reduced through experience, price and quantity controls can be used to obtain similar cost and emission outcomes.

Operationally, there are potential problems when this safety valve is used in conjunction with international emissions trading, as the Kyoto Protocol allows. In general, there would be a need for either harmonization of the trigger price across countries, or restrictions on the sale of permits from those countries with low trigger prices. Otherwise, there would be an incentive for countries with a low trigger price to simply print and export permits to countries with higher permit prices. This would not only effectively create low trigger prices everywhere, it would also create large international capital flows to the governments of countries with the low trigger prices.

Instead of harmonizing trigger prices, we could alternately set the trigger price low enough to avoid the need for international GHG trades. This may be a desirable end in light of concerns about the indirect economic consequences of large volumes of international GHG trade flows (this point has been made by Warwick McKibbin and Peter Wilcoxon; see Further Readings).

Finally, if we find it desirable to raise the trigger price rapidly, it will be necessary to limit the possibility that permits can be purchased now and held for long periods of time. Otherwise, there will be a strong incentive to buy large volumes of cheap permits now in order to sell them at high prices in the future. This problem is easily addressed by assigning an expiration date for permits as they are issued, for perhaps one or two years in the future.

VIII. Building domestic and international support for a price-based approach

While the safety valve approach is potentially appealing to businesses concerned about the uncertainty surrounding future permit prices, environmental groups will be wary of giving up the commitment to a fixed emission target. Such a commitment is already an integral part of the Kyoto Protocol. Ultimately, however, a strict target policy may lack political credibility and viability. Although a low trigger price would clearly rankle environmentalists as an undesirable loosening of the commitment to reduce emissions, a higher trigger price could allay those fears while still providing insurance against high costs.

Perhaps more controversial than the concept of a safety valve is the fact that a hybrid policy requires setting a trigger price. It extends the debate over targets and timetables to include, based on the trigger price, perceived benefits. Business interests will undoubtedly seek a low trigger price and environmental groups a high trigger price. I believe this is desirable. The debate will focus on the source of disagreement between different groups: namely, the value placed on reduced emissions. Rather than leaning on rhetoric that casts reduction commitments as either the source of the next global recession (according to businesses) or the costless ushering in of a new age of cheaper and more energy efficient living (according to environmentalists), it will be necessary to decide how much we are realistically willing to spend in order to deal with the problem.

While seemingly provocative in its challenge of the core concept of targets and timetables embedded in the Kyoto protocol, some form of the safety valve idea is already part of many countries' notion of their Kyoto commitments. European countries who are likely to implement carbon taxes must have some view as to how they will handle target violations if their tax proposals fail to sufficiently reduce emissions before the end of the first commitment period. Other countries who are considering either a quantity or command-and-control approach likewise must envision a way out if their actual costs begin to surpass their political will to reduce emissions.

Among the many “implicit safety valve” possibilities, one could imagine a more flexible interpretation of existing provisions, such as the Clean Development Mechanism or the use of carbon sinks. Alternatively, Article 27 specifies that parties can withdraw from the Protocol by giving notice one year in advance. A country that foresaw difficulty in meeting its target in the first commitment period could serve notice that it wishes to withdraw before the commitment period ends.

Implicitly, therefore, flexibility in meeting current commitments already exists. Countries can choose to massage their commitments using existing provisions, violate their targets and risk penalties (which have yet to be defined) or simply withdraw. In these cases, however, the outcome and consequence are unclear. The advantage of a price mechanism is that it makes the safety valve concept *explicit* and *transparent*. Establishing a price trigger for additional emissions allows countries, and in turn private economic decision-makers, to approach their reduction commitments with greater

certainty about the future. This not only improves the credibility of the Protocol but also its prospects for future success in reducing GHG emissions.

IX. Conclusions

The considerable uncertainty surrounding the cost of international GHG emission targets means that price- and quantity-based policy instruments cannot be viewed as alternative mechanisms for obtaining the same outcome. Price mechanisms will lead to uncertain emission consequences and quantity mechanisms will lead to uncertain cost consequences. Economic theory as well as numerical simulations indicate that the price approach is preferable for GHG control, generating five times the net expected benefit associated with even the most prudent quantity control. The essence of this result is that a rigid quantity target over the next decade is indefensible at high costs when the stock of GHGs is allowed to increase over the same horizon.

Importantly, a price mechanism need not take the form of carbon tax. The key feature of the price policy is its ability to relax the stringency of the target if control costs turn out to be higher than expected. Such a feature can be implemented in conjunction with a quantity-based mechanism as a "safety valve." A quantity target is still set but with the understanding that additional emissions (beyond the target) will be permitted only if the regulated entities are willing to pay an agreed upon trigger price.

This approach can improve the credibility of the Protocol and its prospects for successful GHG emission reductions. This last point is particularly relevant for ongoing climate negotiations. Should the emission incentives and consequences remain ambiguous and uncertain, or should they be made explicit and transparent? Specifying a price at which additional, above-target emissions rights can be purchased provides such a transparent incentive. The current approach does not. While ambiguity may prove to be the easier negotiating route, it may also be a disincentive for true action.

FURTHER READINGS

General

- Fischer, Carolyn, Suzi Kerr and Michael Toman. "Using Emissions Trading to Regulate U.S. Greenhouse Gas Emissions." RFF Climate Issues Brief #10-11, 1998.
- Harrington, Winston and Richard Morgenstern. "On the Accuracy of Regulatory Cost Estimates." RFF Discussion Paper 99-18, 1998.
- McKibbin, W. and P. Wilcoxon. "A Better Way to Slow Global Climate Change" Brookings Policy Brief No. 17 (available at <http://www.brookings.edu>), 1997.
- Newell, R.N. and W.A. Pizer. "Stock Externality Regulation Under Uncertainty," RFF Discussion Paper 99-10, 1998.
- Nordhaus, W.D. *Managing the Global Common*. Cambridge: MIT Press. 1994.
- Nordhaus, W.D. "Expert Opinion on Climate Change," *American Scientist* 82, 1994, p. 45-51.
- Parry, Ian. "Revenue Recycling and the Costs of Reducing Carbon Emissions." RFF Climate Issues Brief #2, June 1997.
- Pizer, W.A. "Prices versus Quantity Revisited: The Case of Climate Change." RFF Discussion Paper 98-02, 1998.

Technical

- Broecker. "Thermohaline Circulation, the Achilles Heel of Our Climate System: Will Man-Made CO₂ Upset the Current Balance," *Science* 278, 1997.
- Enting, I.G., T.M. Wigley and M. Heimann. "Future Emissions and Concentrations of Carbon Dioxide: Key Ocean/ Atmosphere/ Land Analyses." CSIRO Division of Atmospheric Research Technical Paper #31 (available at <http://cdiac.esd.ornl.gov>), 1994.
- Fischer, Carolyn, Suzi Kerr and Michael Toman). "Using Emissions Trading to Regulate U.S. Greenhouse Gas Emissions: An Overview of Policy Design and Implementation Issues," *National Tax Journal*, Vol. 51, No. 3, September 1998, p. 453-464.
- IPCC. *The Science of Climate Change Contribution of Working Group I to the Second Assessment of the Intergovernmental Panel on Climate Change*. Cambridge, UK, Cambridge University Press, 1996.
- Parry, Ian, Roberton Williams, III and Lawrence Goulder. "When Can Carbon Abatement Policies Increase Welfare? The Fundamental Role of Distorted Factor Markets," *Journal of Environmental Economics and Management*, Vol. 37, No. 1, January 1999.
- Roughgarden, T. and S. Schneider. "Climate Change Policy: Quantifying Uncertainties for Damages and Optimal Carbon Taxes," forthcoming in *Energy Policy*, 1999.