7

Sharing the Burden of Saving the Planet: Global Social Justice for Sustainable Development

*Lessons from the Theory of Public Finance**

JOSEPH E. STIGLITZ

Introduction

The risks of climate change resulting from the increase in atmospheric concentration of greenhouse gases have been well-documented. There are likely to be especially adverse effects on developing countries and particularly the poor within these countries. There is a global consensus that strong actions need to be taken to ensure that the world does not face *excessive* risk from an increase in the atmospheric concentration of greenhouse gases that would, say, lead to an increase in average global temperatures of more than 2 degrees Celsius. . This chapter is concerned with how the global community should respond to this global risk and in particular how the burden of preventing global warming—or more accurately, reducing the

*Earlier versions of this paper were presented at the World Congress of the International Economic Association, Istanbul, June 2009, and at a meeting on the Equitable Burden Sharing in Responding to Climate Change hosted by Brooks World Poverty Institute, Manchester, and jointly sponsored by the Initiative for Policy Dialogue on June 17–18, 2009. The author is indebted to the Ford, Mott, and Macarthur Foundations for financial support. The author served as a member of the 1995 assessment panel of the International Panel on Climate Change (IPCC).

risk—should be shared, especially between developed and less-developed countries. Almost surely, no matter what we do, there will be the necessity for adaptation, with significant costs borne especially by developing countries. How those costs should be shared is an important question that is beyond the scope of this chapter.

There are five other points of consensus that form the background for this chapter:

1. Global warming is a global problem, and it needs to be addressed globally. Unless all countries participate, there is a danger of leakage; reductions in emissions in one country may be more than offset by increases elsewhere (Stern 2007).
2. Global warming is a long-run problem. We are concerned not so much with the level of emissions in any particular year as with the long-run levels of atmospheric concentrations of greenhouse gases.
3. The costs of reducing the level of emissions (limiting the increases in atmospheric concentration of greenhouse gases) will be much lower if it is done efficiently. Efficiency implies *comprehensiveness*—we need to address all sources of emissions and explore all ways of reducing atmospheric carbon concentrations, including carbon storage and carbon sequestration.
4. There is considerable uncertainty about the level of "tolerable" increases in greenhouse gas concentrations and about the impact of particular policy interventions.
5. Global warming is a public good problem, so there is a risk of free riding. This means that there will have to be some system of credible enforcement.

There are five important corollaries of these points of consensus:

1. We need a global agreement, and a global agreement will require equitable burden sharing. Much of this chapter is concerned with exploring what this entails.
2. The shadow price of carbon should be approximately the same *in all uses, in all countries, and at all dates.* Current arrangements deviate in important ways from this principle. The (shadow) price of carbon in those countries that signed on to the Kyoto Protocol is higher than in other countries. The (shadow) price of carbon associated with deforestation is lower than in other uses. In many countries, the price of

carbon (reductions) associated with renewables, and especially ethanol, is higher than in other uses.

3. The costs of adjustment will be smaller if the adjustment is done efficiently, which means that the adjustment will be spread out over time. But that does not mean that the prices need to adjust slowly. An immediate adjustment in prices to reflect the true scarcity will result in a gradual adjustment in behaviors, reflecting an efficient response to the costs of adjustment. If there is an argument for gradual adjustment of prices, it is based on distributive concerns.
4. The fact that this is a long-run problem with considerable uncertainty means that whether we work through emission targets or prices, there will need to be adjustments over time. In an emission targets system, we will have to adjust the targets. In a carbon tax system, we will have to adjust the tax. Thus, the standard argument that, in the face of certain types of uncertainties, quantity targets are preferable to price interventions is of limited relevance.
5. We need to differentiate between systemic risk and risk faced by market participants. Uncertainties—and differences in beliefs about the nature of the risks—in fact provide an argument for mixed instruments, such as the safety valve, where, in the short run, there is a cap on the price faced by firms. Market participants are risk averse, and there is a cost to imposing risk on them. Intertemporal adjustments allow firm and individual risks to be spread out over time, and this greatly mitigates those risks. The fact that what matters is the long-run atmospheric concentrations means that the environmental costs of any limited temporary deviations from prespecified targets is likely to be small.

Insights from Public Finance

There are two more introductory remarks. The problem we are discussing has many of the features of those addressed within classical public finance. The preservation of our common atmosphere (preventing global warming) is a global public good.[1] It has to be financed. Standard theories of public finance provide clear formulations concerning equitable and efficient taxation.

Alternatively, we can think of carbon emissions as generating a global externality, and again, standard public finance theories discuss efficient and equitable ways of controlling an externality generating activity—including

the relative merits of corrective taxation and of regulatory interventions. Much of the literature has focused on the equivalence of the two systems of interventions, under certain conditions, and much of our analysis will make use of that equivalence. We will first analyze tax interventions because in doing so, the efficiency and equity implications become more transparent. We will then provide the interpretation for quantity interventions.

Theory of Second Best

Second, policy in this area—even more than in many other areas of economics—is a matter of the *economics of the second best*. Even governments that are committed to reducing emissions have limited control. Emissions are the by-product of every economic activity. Emissions are not just a matter of industrialization: The methane produced by animals is a major contributor to emissions. We have increasingly become concerned with deforestation, which contributes 20 percent of the world's emissions. But moving to building materials other than lumber may not help: 5 percent of the world's emissions come from the production of cement. Not only can't we control emissions perfectly we cannot even measure emissions perfectly.

There is a second important second-best consideration. There are, in fact, two important unpriced (or imperfectly priced) resources: (clean, fresh) water and (carbon in the) air. Many of the reform proposals involve, implicitly or explicitly, putting a price on carbon. But this may increase the importance of other distortions.

Biofuels provide an illustration of what is at issue. One of the responses in many parts of the world to the threat of global warming is to increase the production of biofuels, the production of which, in some parts of the world, makes extensive use of already very limited supplies of water.[2] At the very least, we need to be aware of this distortion.

Moreover, the increase in biofuels has contributed to the increase in the price of food. In this case, the *incidence* of the (hidden and implicit) tax on carbon is borne disproportionately by the poor in the world because they spend a larger fraction of their income on food, whereas the rich biofuel producers and corn producers in the United States are actually better off. This compounds the inequities already imposed by global warming: Those in the tropics—where a large fraction of the poorest live—are likely to be most adversely affected by global warming; but this response puts the burden of adjustment disproportionately on them.

One of the reasons that the economics of the second best is especially important in this context is that enforcing a global carbon regime will not be easy. Imagine the difficulties of enforcing a global income tax. Tax evasion would be rife. Whether we have a global carbon tax or a system of emission permits, carbon will have a price, and there will be incentives to avoid paying that price. Over the years, we have come to understand how better to enforce taxes; we will need to transfer some of these lessons to controlling carbon emissions.

Key Insights

This chapter has three parallel themes. The most important concerns burden sharing between developed and less-developed countries.[3] Any agreement has to be crafted in such a way that it does not adversely affect growth and poverty alleviation within the developing world. Beyond the normative perspective (in virtually any ethical framework, an agreement that put the burden of mitigation on developing countries to such an extent that it increased poverty—while the rich in the developed countries continued to consume in their current profligate style—would seem an anathema), it would be hard for any developing country to muster support for an agreement that was perceived to put their development agenda in jeopardy. This would be so even if it developed countries didn't have a historical responsibility for the increases in greenhouse gases that have occurred during the past two centuries, or even if the developed countries didn't bear especial culpability for their failures to live up to prior commitments.

All of this means that the costs of mitigation ought to be borne by the developed countries and that the developed countries ought to help the developing countries bear the costs of adaptation. Resources devoted to limiting emissions or to adapting to climate change are resources that could have been spent reducing poverty or promoting growth.

In a sense, it is unlikely that what emerges from any agreement is truly a "fair deal," given the "climate change" that has already been built into the system.

By now, there is a broad consensus that equitable distribution of the required global emission reductions—especially taking into account the imperatives for developing country growth and poverty alleviation—requires very large reductions by the developed countries, by 80 percent, 90 percent, or more below 1990 levels by 2050. Delay in large reductions

by the developed countries has large implications for equity: It means that, other things being equal, the developed countries will have "consumed" an even larger share of the "atmospheric pie," the total amount of cumulative emissions that are consistent with, say, two degrees Celsius global warming. And if these inequities are to be avoided, it means that the developed countries' cuts by 2050 will have to be all that much deeper. (And if the developed countries do not cut their emissions, there will be strong and inequitable adverse effects on developing countries.)

The second theme is methodological. One aspect, already noted, entails using insights and tools from public finance to provide insights into the merits of alternative approaches to addressing the problems posed by global warming. The second methodological insight is to view the problem of carbon management through the lens of the basic *carbon conservation equation*—which says that carbon molecules must either be in the atmosphere or in storage (below the ground, in the ocean, or in trees or other forms of terrestrial carbon). What is of concern (at least for global warming) is the carbon in the atmosphere. Individuals and firms have to be charged for putting carbon into the atmosphere, or rewarded for keeping carbon molecules in one of the other locations.

The third is substantive, the result in part of applying these methodologies. Four central insights (beyond those already described) are:

1. Distributive issues are central, within and between countries; issues of efficiency cannot be fully separated from those of equity.
2. Under standard normative approaches, fair burden sharing under a system of emissions rights would give more emission permits per capita to developing countries than to developed countries.
3. Because equal per capita emissions is thus the minimal acceptable allocation from the perspective of developing countries but is beyond the "maximal" acceptable allocation from the perspective of developed countries, achieving an agreement within the emissions rights framework is likely to prove elusive. A carbon tax may provide a better approach to achieving an equitable global agreement than the allocation of emission rights. We suggest one particular approach to the design of a carbon tax—carbon-added tax—that may have some advantages in implementation.
4. A concern for distributive consequences provides a rationale for at least partial reliance on regulatory measures. Although such measures may be less efficient than a uniform carbon price (achieved through a carbon tax or emission permits), the distributive impacts may be less severe.

Lessons from Public Finance: Tax Equivalence and Implementation

Public finance theory focuses on three key aspects of taxation: (1) efficiency—does the tax distort resource allocation (or how to design taxes to minimize distortions); (2) equity—how does the tax affect different groups (what is its *incidence)* and is the burden of the tax, in some sense, fair; and (3) administration—is the tax effectively enforced, at reasonably low transaction costs, and relatively immune from corruption.

There is, in general, a cost to raising taxes—as the old adage has it, there is no such thing as a free lunch. But taxes on polluting activities are an exception, for they increase economic efficiency at the same time that they raise revenues. Such taxes are referred to as corrective or Pigouvian taxes.[4] Such taxes are consistent with a basic principle of environmental economics, called the polluter pay principle: Making polluters pay for the pollution that they create is fair and efficient.

That is why economists have begun with the presumption that the best way to incentivize individuals and firms not to emit greenhouse gases is to impose a tax on their emissions, a carbon tax. Surprisingly, however, attention has shifted to another mechanism of reducing emissions, emission permits, which can be bought and sold. To emit, one must have an emission permit. This results in the imposition of a marginal cost associated with emissions, just as the case is with a carbon tax. As we will explain later in the chapter, the two systems can be made fully equivalent; in practice, though, they are not likely to be because of how emission permits are granted. This has large distributional consequences. These in turn have strong implications for the prospect of reaching an agreement. Before turning, however, to a comparison of these different instruments for reducing emissions, it is instructive to understand better the lessons of standard public finance.

In a world with perfect competition, it makes little difference whether we impose a tax on producers or consumers. The incidence of the tax (that is, who pays the costs) is the same, and the general equilibrium that emerges (that is, the output of each sector, the income of each individual) is the same. Public discussions, however, typically make a great deal of difference, partly because markets are not perfectly competitive and partly because transitions from one equilibrium to another are not instantaneous: How the tax is levied can make a great deal of difference in the transition.

In the case of carbon, the focus has been totally on production. China is being "credited" with exceeding the United States in emissions (though

its carbon emissions per capita are still markedly smaller). However, many of the goods that are produced in China and that account for considerable amounts of its emissions, are consumed in the United States. In terms of "consumption" accounting, America is still in the lead.

Whether one uses consumption or production based accounting makes a great deal of difference in the case of carbon, for two reasons. If one levies a tax (or imposes a system that is equivalent to a tax), the way that the tax is levied can have large distributional consequences. If a tax is levied on consumption, revenues are generated at the point of consumption; if levied on production, at the point of production. In a closed economy, it makes no difference. In an open economy, it can make a great deal of difference, because it affects who gets the tax revenues.

Second, if taxes are imposed on production, and some countries are more effective in enforcement—or impose a lower tax—then production, particularly of carbon-intensive goods, will gravitate to where it is, in effect, taxed less. This impedes efficiency: Production will occur not where it is most efficient but where taxes are the lowest. In the case of carbon, this is of particular concern because the objective of imposing the tax (restriction) is a *global* reduction. With such "evasion," total carbon emissions could actually increase as production shifts from locales with lower emissions but higher (effective) taxes to locales with higher emissions but lower effective taxes. (This is called the problem of "leakage.")

A Carbon-Added Tax

In the design of tax systems, problems of *enforcement* have taken on first order importance. The argument for the value-added tax (VAT) in the advanced industrial countries is that the system is self-enforcing, and thus there is greater compliance. Collection efforts can be focused on large firms that generate a large fraction of value added. Each firm in the production chain has an incentive to claim a deduction for goods purchased from others, which helps ensure that they reveal their purchases, forcing others to reveal their sales.

At the same time, the difficulty of enforcing the VAT *uniformly* in developing countries—it is virtually impossible to collect VAT revenues from the informal sector, including agriculture, which comprises a large share of GDP— has provided one of the strongest criticisms for its adoption there. Even though with full enforcement, such a tax is efficient, in practice, it

is highly distortionary—moving resources out of the "formal sector," the very sector that most developing countries wish to encourage (Emran and Stiglitz 2005).

A *carbon-added tax* (CAT) levied at each stage of production would have some of the same advantages that a value-added consumption tax has. Each producer would have to show receipts for the carbon tax paid on inputs into its production. (We frame the discussion in terms of a carbon tax; later, we will reframe the discussion in terms of a regime of emission permits.) The taxes levied at each stage of production would be passed on to consumers. It is *as if* the tax were imposed on consumers, but the problem with levying a tax directly on consumers is that there may be many ways of producing a good. We cannot look at a good and infer how much carbon was used in its production. A carbon value added tax will discourage production in more carbon-intensive ways and discourage the consumption of carbon-intensive goods.

If a firm could not produce receipts for carbon taxes on inputs, then a tax would be levied on the input, assuming it was made in the least carbon-efficient way. This would provide strong incentives for each firm to make sure that its suppliers complied with the carbon tax regime.

It would be easy to incorporate countries that failed to go along with the international regime. Producers in those countries would not be able to show carbon tax receipts. One could follow the procedure just described: A tax would be imposed on the input on the assumption that it was produced in the most carbon-intensive way possible. This by itself would provide a strong incentive for the country to impose a carbon tax, at least on exports. The cost to outside buyers would be the same, but the producing country's government would garner the revenue.[5]

Because most firms are unlikely to have two production lines—one for exports, one for domestic consumption—the tax would provide an incentive for reducing carbon emissions. But if exports are a small fraction of total production, the incentive is limited. This suggests a more aggressive approach, with a compensatory tax on the input designed to make up for the failure to impose the tax on output that is not exported.

Intergovernmental Distributional Implications

We noted earlier that the allocative effects of taxing consumers and producers are the same, but if production occurs in a different jurisdiction

from consumption, then the two taxes will differ in who receives the tax revenues. With a CAT, taxes are collected at each stage of production. This implies that the revenues go to the producing countries.

Because most carbon emissions are related to the burning of fossil fuels, it might seem that an efficient way of collecting the tax would be to have a global tax on fossil fuels, corresponding to the greenhouse gases that are emitted when they are used. There are many fewer points of production of oil, coal, and gas than there are points of usage (one of the arguments for the VAT).

It would thus seem efficient to have the taxes collected by the producers of fossil fuels; and obviously, these countries would prefer such a system because they would reap virtually all of the tax revenue—a tax in addition to the market price. As the tax increased, the market price (the "rents" they receive from their natural resources) would go down. But their total revenues would be essentially unaffected.[6] Interestingly, although the fossil fuel producers have been major opponents of doing something about global warming, under this regime, they would be fully compensated; but with their incentives for producing effectively unchanged, it is clear that such a system would do little for global warming.

Thus, even though it might seem *administratively* simpler to impose the tax at the point of production of coal, oil, or gas, or at the cutting of the forest, and so on, and such a tax would, in effect, capture most of the "carbon added" into the global system, any carbon tax system will have to focus on usage, that is, imposing the tax on the *use* of carbon (oil, coal, gas) at each stage of production.[7]

As we have noted, much of the policy debate has focused not on the question of the best way to implement a carbon tax but on whether a carbon tax (administered as a CAT or in some other way) is superior to a system of emission permits. Later in this chapter, we shall have a word to say about the relative merits of the two systems. For now, we note that the standard approach on emission permits is based on the "value added" approach—that is, emission permits are required at each stage of production. There has been less attention to the enforcement issue than in the context of taxation, but the same logic can easily be extended to emission permits, as both systems have the advantage of decentralized enforcement. Permits would be required at each stage of production. Each firm would be responsible for verifying that those from whom they bought inputs did so "legally," that is, holding the requisite carbon permits. If the supplier did not have valid permits, the firm would be "charged' for

using carbon on the assumption that the most carbon-intensive method of production had been used.

Equitable Burden Sharing

The key problem today in reaching an agreement is not the science: As we have noted, there is a growing consensus about the minimum that needs to be done—and that the minimum is much greater than what the world is doing today. The problem is how to share the burden of adjustment—and adjustment costs are likely to be large.

A scarce resource—carbon in the atmosphere—has been treated as if it were a free good. The market equilibrium that has emerged is, as a result, greatly distorted. Many of the key decisions that affect carbon emissions are long-run, such as power plants, housing, and transportation systems. Many of the decisions themselves are not totally market-driven; for example, land usage patterns are affected by zoning.

It is, of course, not just a matter of adjustment costs. Charging the social cost for something that has been treated as free will change relative prices. There will be winners and losers: The losers will want to be compensated; the winners will be reluctant to do so. In a sense, any change in the scarcity value of any factor of production has similar consequences; when these changes in relative prices are driven by market forces, we come to accept them—though those hurt are again more demanding of help than those who benefit are willing to share their newfound gains. But this seems somehow different because it is a political decision (though no less than the enclosure of common land or common knowledge[8] is a political decision).

If, say, through a high carbon price we succeed in ensuring that fossil fuels remain below the ground, then those who otherwise would have sold those fuels are clearly worse off.[9] With a credible program on global warming, the owners of oil and coal reserves will see the value of their assets diminish—regardless of the design of the program. The wealth of the oil exporters will also diminish. To be sure, there may be limited sympathy—they have done very well in the last few years, and unlike wealth that is the result of hard work, ingenuity, or savings, it appears to be largely the result of luck. We should expect that countries with large endowments of these resources will do everything they can to make sure that there is no agreement.

The same thing is true, of course, not just for countries but for companies—though companies have more of a choice. An important part

of their asset base is their skills and knowledge. BP, with its slogan *Beyond petroleum,* has suggested that a company can transform itself from an oil producer to an energy producer that is not dependent on fossil fuels. Still, responding to global warming will result in a decrease in the value of certain assets (just as *not* responding to global warming will result in the decrease in the value of other assets).

It is worth bearing in mind these *within*-country distributive effects because they play an important role in determining policies. America's response to global warming may be more determined by the impact on its oil companies and on its automobile industry, which has been geared toward high oil-consuming vehicles, than by a more balanced consideration of the country's national interest. As a major oil importer, America would benefit from the lower price of oil that a global agreement would bring about.

Still, for most of this chapter, as important as these within-country distributive effects are for political economy, I shall focus my attention on the cross-country distributive effects.

Externalities and Pareto Improvements

The fact that there are externalities means, of course, that there is scope for a "deal" that makes everyone better off. In principle, there is a "Coasian" agreement by which those injured by the polluters pay the polluters not to pollute, leaving the polluters and those suffering from the pollution better off. In this perspective, it should be easy to reach an agreement and because the developing countries, on average, are those harmed by global warming, it would entail the developing countries paying off the rich countries not to pollute, in contradiction to the principle of polluter pay. That the world is unlikely to reach agreement based on such a perspective seems obvious.

The Insufficiency of Improved Energy Efficiency

Much attention has been placed on the inefficiencies in energy usage in developing countries. Increasing energy efficiency will, it is widely believed, reduce emissions. This is presented as a win-win situation: The global environment benefits at the same time that the developing country saves on scarce resources. Such increases in energy efficiency are likely to go only a little of the way toward meeting the requisite reductions

in emissions. Indeed, whether increases in energy efficiency lead to an overall increase or decrease in emissions depends on whether achieving "economic" energy efficiency leads to an increase or decrease in energy usage, which in turn depends on whether the demand for energy has an elasticity that is greater or less than unity. More energy efficiency will lead to the price of energy falling; if the demand for energy is price elastic, then there will be a more than proportionate increase in energy usage so that emission levels will increase. Achieving energy efficiency is desirable, but it will not suffice.[10]

A Global Carbon Tax

The basic insight of public finance theory is that the global societal costs associated with reducing energy emissions can be minimized by the imposition of a global carbon tax. The current price of carbon is zero. Increasing the price of carbon from zero to the optimal price will, however, adversely affect some countries (carbon exporters). But the notion that it is inefficient to allow global warming means that the revenue generated by this corrective tax is more than sufficient to compensate them for the increase in the price of carbon. In appendix A, we provide an analytic framework within which one can calculate the amount of compensation each country must receive to make itself better off.

Indeed, there are many allocations of the tax revenues that can make every country better off, and much of the fight going on can be viewed as how to allocate the *typically implicit* tax revenues.

A system of carbon trading, with grants of emission permits based on, say, 1990 levels of emissions (as the Kyoto Protocol of 1997 effectively did), in effect gives emission tax revenues in proportion to 1990 levels of emissions. It does not, of course, directly give tax revenues; but it does grant emission permits, which have a market value, and the value of what each country receives is proportional to their emission allowances, which are roughly proportional to past emissions.[11] That means that the United States not only gets the single largest allocation but also gets the largest allocation on a per capita basis.

There is no ethical basis for such an allocation. Indeed, developing countries argue that because the North contributed disproportionately to the current build-up of greenhouse gases, their future allocations should be commensurately reduced.

The compensation on which attention has been focused is the direct extra costs associated with the higher price (inclusive of the tax) of fossil fuels. The discussion so far has not shifted to the broader issue of compensation for the implied changes in consumption prices, for example, for food.[12] It is thus worth noting that some developing countries could be worse off after the imposition of a global carbon tax in which they were full compensated for the increased price of their direct purchases of fossil fuel. Because the country might have to pay more for its imports of foods, it would have less money at its disposal for development.

The benefits of having a distorted global economy, where the price of carbon was zero, are distributed in complex ways, so much so that—not surprisingly—correcting this distortion will have ramifications for developed and developing countries. But allocating so much of the implicit tax revenues from the global carbon tax to the developed countries means that there is a much greater likelihood that more developing countries will see themselves worse off, once full account is taken of the indirect consequences for food as well as energy, which comprise such a large fraction of their market basket.[13] (Some of these, of course, will still be better off than they are in a regime with global warming; a full analysis also needs to take into account the benefits from reduced global warming.)

Financing a Global Public Good

The earth's atmosphere is a global public good. Thus, avoiding global warming—in other words, preserving the health of the atmosphere—is also a global public good, and given the large disparities between the rich and the poor countries, all (or at least most) of the costs of providing this public good *should* be borne by rich countries.[14] In this perspective, developing countries should be compensated for providing valuable environmental services—carbon storage—and for the additional costs of reducing carbon—of going beyond energy efficiency to carbon efficiency.

An Agreed-Upon Carbon Tax

One proposal that has received some attention is that the countries of the world agree upon a carbon tax level that would achieve the desired reduction in emissions. Each country would then keep the revenue for itself. In

effect, a carbon tax would substitute for taxes on work and savings; and under the principle that it is better to tax bad things than good things, such taxes would yield a double dividend.

Appendix B explains why, for most countries, we should expect this to suffice to provide adequate compensation—so that all countries are better off. In a sense, the distributional impacts are likely to be small. The "cost" of the carbon tax is the *difference* between the dead weight loss[15] of the carbon tax and the alternative tax (say, a wage tax). This number is likely to be small. But the *differential* incidence is the difference in this difference across countries—a number that is likely to be even smaller. In short, the advantage of the common carbon tax is that distributive consequences can be shunted aside.

Emission Permits

The alternative to a common global carbon tax is a set of agreed-upon emission limits. Efficiency can be obtained if these emission allowances (or permits) are tradable (that is, can be bought and sold). Later in this section, I will discuss some of the relative merits of the two systems. But one disadvantage of the emission permits is that it brings to the fore the distributive conflict.

We have already discussed the implicit—and unacceptable—allocation of emission permits under the Kyoto Protocol, in which those who had polluted the most get the most emission permits. The question of the allocation of emission permits is, of course, isomorphic to the question of the allocation of tax revenues.

Equal Emissions per Capita

One alternative, widely-discussed principle is equal emission permits per capita—that is, distributing the revenues equally among all the citizens of the world. This seems more philosophically acceptable than allocations based on past emissions. But there are two criticisms.

Most theories of social justice argue for a more progressive distribution of the revenues generated from the "sale" of a global natural resource—the right to emit carbon in the atmosphere—than equal per capita. Arguing that those who polluted more in the past have the right to pollute more in

the future is, to say the least, perverse; and because past levels of pollution are related to income, such a rule is clearly highly regressive.

The question can be viewed another way, from a property rights perspective: How should property rights in the atmosphere be allocated? Ronald H. Coase (1960) argued that it did not matter how one assigned property rights; all that mattered for economic efficiency was that there was a clear assignment. Though that proposition has come to be questioned, to achieve a global agreement among all the countries will require that the developing countries believe that the implicit assignment of property rights is, in some sense, fair, or at least acceptable.

Within democratic developing countries today, a property rights allocation that gives citizens any less than a proportionate claim is not likely to be acceptable. Indeed, there is an alternative approach that suggests that the developed countries get markedly smaller allocations on a per capita basis. The world has now agreed that carbon concentration should be limited to make it unlikely that the world will experience an increase in temperature of more than two degrees Celsius. This implies a maximum level of greenhouse gases in the atmosphere, that is, there is a limited amount of "carbon space." Developing countries are now arguing that each should be given a claim on this carbon space in proportion to their population. But, they add, the developed countries have already used up much of their carbon space, so that their remaining carbon space is much less (on a per capita basis); this implies that, going forward, they must have lower levels of emissions per capita.

Developed countries have two responses: First, any system based on population rewards countries for failing to limit their population; and second, why should they be held accountable for emissions that occurred before the risks of global warming were known?In American jurisprudence, firms are held liable for damages that occur before a particular risk is known because in doing so it provides strong incentives for firms to find out about the risks associated with their actions.

Both concerns may fairly be addressed by providing an allocation of the carbon space as of 1992, on the basis of population as of that date. Because the United States has used so much of its carbon space since then, it is still the case that under this normative principle the United States, going forward, would have fewer emission permits per capita than developing countries.

This approach has a positive incentive effect: It provides an incentive for the United States to reduce its emissions quickly. In the current

circumstances, it is as if the United States and some other developed countries are trying to "steal" as much of the global carbon space as they can before an agreement is reached.

The "problem" with even a rule based on the more modest allocation of equal per capita emission permits is the fear that it will lead to high levels of payments from developed to developing countries—at least for the foreseeable future. To be sure, as developing countries develop, differences in per capita emissions will be reduced, and so the scope for transfers will be reduced. A slow enough pacing-in of emission reductions might hold out the possibility that transfers could be kept to a moderate level. But, by the same token, a slow pace of emission reductions increases the inequities in the usage of the global carbon space.

But projections made on the basis of current rates of increases in emissions, say in China, may be misleading for at least two reasons: (1) Rapid paces of technological adaptation may lead to rapid increases in energy efficiency—the government is committed to making these changes; and (2) China has been (and will for some time continue to be) going through a resource-intensive phase of its development—focused on expanding housing and cars. It will eventually follow the pattern of other countries, shifting to the less resource-intensive service sector. Already, China is discouraging output in energy intensive sectors, particularly energy intensive exports (this, in turn, may in part be due to the system of attribution that "credits" China with emissions for products consumed elsewhere).

Corruption and Emission Permits

There is another problem with most systems of emission permits *within* countries: Any system in which the government allocates permits (which is equivalent to allocating money) is subject to corruption, either overt corruption, or the more subtle form, campaign contributions to induce the political process to adopt a "rule" that benefits particular parties.

There is an alternative—auctioning off emission permits. If the auction is held internationally, the system is identical to a system of global taxation in which the revenues are pooled together—and the international community must then decide on the allocation of revenues (see the earlier discussion). If the auction is held at the national level, it is equivalent to the system of an agreed-upon tax level, with revenues retained by each country.

Of course, the auction undoes one of the reasons given for the permit system: The possibility of receiving large amounts of money has provided political support for (or reduced opposition to) taking actions to reduce emissions.[16] But these political economy arguments *for* tradable emission permits are, at the same time, the main arguments against them; allocating a disproportionate number of permits to those currently engaged in polluting is the very reason that those who are not currently polluting—as much, on a per capita basis—will oppose it.

Risk, Permits, and a Carbon Tax

There is a second reason (besides its ability to "buy" support) that the emission permit approach has been in favor: Environmentalists like the seeming certainty that it provides. Given the agreed-upon level of emission permits, one knows the level of emissions. With a carbon tax, one can't be sure what the level of emissions will be.

However, what we are concerned about is climate change. Emissions translate into increased carbon concentrations, and carbon concentrations translate into climate change; there is uncertainty at each stage, so that there is a high likelihood that any agreed-upon emission levels will be revised as our scientific knowledge improves. So too, if we see that the agreed-upon tax is producing higher levels of emissions than desired and anticipated, we can increase the tax. In both cases, there will be sequential revisions.

Why Granting Emission Permits on the Basis of Past Emissions Overcompensates

There are other arguments for not granting emission permits on the basis of past levels of emissions, besides the obvious one that it rewards those with bad behavior, going precisely against the "polluter pay" presumption. In dynamic competitive markets, it *over-rewards* these past polluters; new firms entering the market will, for instance, not have these permits. It is their marginal costs—including the costs of buying the requisite pollution permits—that will determine market price. Prices will rise to reflect the marginal cost of pollution, so *efficient* firms are fully compensated in equilibrium. Thus, granting them pollution permits on the basis of past levels of

pollution overcompensates them. This may help explain the active support for these initiatives by these firms.

Distortionary Approaches to Mitigation

So far, we have considered two alternative, efficient ways of reducing emissions: a global carbon tax and a system of tradable emission permits. Both guarantee that there will be a single price of carbon, in all uses, in all countries. In fact, almost every country has deviated from this general principle by introducing, for instance, regulations on minimal usage of ethanol (United States), minimum fuel efficiency standards (United States), or providing subsidies to renewables (many developed countries).

How can these deviations be justified—particularly in the United States, even by administrations seemingly committed to free market principles? There are two bases for arguing for these distortionary interventions.

DISTRIBUTIVE CONCERNS

The first argument focuses on distributive concerns, a worry about the magnitude of price changes (say, induced by the carbon tax) required to elicit the requisite behavioral responses. When there are low demand or supply elasticities, large price changes may be required to elicit the desired changes in usage. A high enough price of carbon would lead to the requisite changes in carbon emissions; but the effect on the poor could be devastating. To be sure, one could offset these adverse effects using, for instance, revenues raised by the carbon tax or the auctioning of emission permits. But it is never possible to target perfectly, and many may be hurt in the process. And if the revenues have been committed to "buying off" politically powerful potential opponents of emission reductions (for instance, by providing emission permits on the basis of past levels of emissions), additional taxes will have to be levied to compensate those hurt indirectly; and there is a deadweight loss to these taxes.[17]

Regulatory approaches may be able to achieve large reductions in emissions with much smaller changes in equilibrium prices and, accordingly, with much smaller distributive impacts.

Part of the argument (for and against) these regulatory approaches may be that the impacts are less transparent. Requiring the use of renewables increases costs of production and leads to higher consumer prices; but it

may be harder to directly link the price increase with the regulation than in the case of a tax.

MARKET FAILURES

The second argument is that markets, by themselves, are not efficient, and government intervention is required to achieve efficiency. There may, for instance, be a coordination failure: Builders do not install energy efficient light bulbs as standard equipment; if they do, they know that consumers will be unhappy because they cannot easily replace them in local stores. And local stores do not stock these light bulbs because there isn't the demand. A government regulation requiring all new buildings to have energy efficient light bulbs solves the coordination problem. Stores will quickly perceive the demand and will stock them.

Innovation is based not only on prices today but on beliefs about future prices. Market expectations may not be rational. Each market participant may believe that there will be a technological breakthrough that will allow the economy to achieve its emission reductions with a low carbon tax. With a low carbon tax in the future, it does not pay most firms to invest heavily in carbon reducing innovation. (It is clear that American automobile manufacturers misjudged the probability distribution of future gasoline prices. Shareholders have borne some of the costs of this mistake—but so too does the rest of society when, as a result, there are excessive emissions. Of course, if they had to pay the full costs—through a carbon tax—society would have been compensated. But when a whole industry makes a correlated mistake, it may be too big to fail, and not only will there be a reluctance to impose the full carbon costs, there may even be a bail-out.)

Of course, innovation almost always entails externalities—there are learning spillovers, so that without government support or government mandates, there may be insufficient incentives to innovate.

For all of these reasons (and perhaps others) price signals alone often fail to induce sufficient shifts in investments—particularly research and development investments—so complementary policy measures such as regulations (agreed-upon standards across countries) can play an important role in responding to the challenges of global warming.

There is one arena where price signals explicitly play a more limited role: public investments. For instance, not only is there a need for more public transportation but cities also need to be redesigned and zoned to induce greater reliance on public transportation. This is an example where

market mechanisms by themselves will not suffice: There is a need for collective action, and prices may even give the *wrong* signal for what should be done. Changes in the design of cities can, themselves, lead to changes in preferences. There were changes in lifestyles (and almost surely preferences) in the United States in the 1950s following the construction of the superhighways; but more recently, there has been another change in lifestyle—an increased preference for urban living. Reducing emissions will require changes in the way we live and work—including where we live and work and the structures in which we live and work. And government policies may facilitate such changes.[18]

Global Agreements Around Standards

Perhaps a more hopeful approach (than a common tax or agreed-upon emission permits) to reaching a global agreement that will reduce emissions significantly is based on reaching an accord on standards, that is, for electricity generation, automobile emissions, cement, and so on. Such standards could embrace a large fraction of all emissions.

One of the reasons that an agreement can be reached is that the distributive impacts are less transparent and probably smaller, and compliance may be easier. An example is an agreement that no coal electricity generating plant will be constructed without offsetting carbon storage. Of course, there are still distributive consequences: Countries that rely more heavily on coal may face greater increases in energy prices. They will be disadvantaged relative to those who have hydroelectric resources.

If such agreements are not to have adverse effects on development, the incremental costs faced by developing countries of such carbon efficient technologies should be borne by the developed countries. The magnitude of the compensation would, however, be limited and be relatively easy to calculate. The magnitude of the compensation required would be reduced if the advanced developed countries fulfill their commitments to the transfer of technology.

Access to Technology

Efficient utilization of knowledge requires that it be made freely available. Knowledge is a quintessential global public good. But the patent system,

of course, tries to restrict the usage of knowledge as one way of compensating innovators.

The deficiencies in the patent system (especially as currently designed in the United States) are becoming increasingly recognized: Not only does it lead to an underutilization of knowledge, it may even have adverse effects on the pace of innovation.[19] Here, however, we are concerned with another aspect—the distributive impact and its effect on global warming and reaching an agreement. The refusal of the United States or other advanced industrial countries to transfer technology to developing countries may have large distributive consequences.

If developing countries sign on to a convention requiring them to reduce their emissions by a certain amount, by a certain date, they are thereby committing themselves to an increase in demand for emission reduction innovations. If certain countries have a comparative advantage in the production of these innovations, such a convention can induce large transfers from developing countries to developed countries—and it is understandable that the developing countries would object.[20]

With the developing countries feeling that they have repeatedly been shortchanged, not just by colonialism but also by international agreements (the poorest countries were actually made worse off by the Uruguay round), it is not surprising that they feel reluctant to sign on to an agreement that might result in large transfers from the developing countries to the developed.

Any equitable approach to global warming and to the financing of technological innovations that will succeed in reducing emissions requires that the financial burden rest on the developed countries.

In the 1992 Rio agreement, the developed countries made a commitment to the transfer of technology and to pay the incremental costs associated with carbon emission reductions. There was also a provision for compulsory licenses, so that the developing countries could not be "held up" in the manner described earlier. But there has been little (if any) use of the compulsory license provisions, and the developed countries have basically reneged on commitments to technology transfer and funding.

Developing countries also worry that any funding they do receive will, in effect, come out of existing aid budgets, that is, aid donors will in effect demand that the money that they had previously been providing to promote development be used to reduce emissions.[21] In effect, their growth will be sacrificed to provide a global public good.

Knowledge itself is a global public good, and knowledge to address a global public good is, in a sense, even more of a global public good. This

chapter has been concerned with the question of how to finance the costs of reducing global emissions. There is a similar question of how to share the burden of financing research to reduce the costs of addressing this public good and of what institutional arrangements will most facilitate the production and dissemination of knowledge. Incentivizing research through the patent system imposes a risk that the full benefits of any innovation will not be realized. Other mechanisms—including public support and a prize system—should be at the center of these research efforts (Henry and Stiglitz 2010).

National Security, Energy Independence, and Emission Reductions

The analysis so far has focused on conventional economic goods. Energy, however, is so important that many countries—including the United States—have expressed a concern about energy independence. A cutoff of supplies of energy would have a disastrous effect on the country. Countries can take actions to ensure that there is no cutoff of supplies within their boundaries, but there is little they can do to protect themselves against external shocks. These concerns are not just a matter of the imaginations of security experts, entrusted with thinking though worst-case scenarios. There have been oil boycotts in the past. Sea-lanes for shipping oil are vulnerable. Ukraine has interrupted the supply of gas traveling through the country from Russia on the way to the rest of Europe. Countries rightly worry about their vulnerability.

The problem is that different kinds of energy are not quickly substitutable. China and India have large coal stocks but must import oil and gas. Developing an economy that relies on imported oil and gas leaves the country vulnerable. Restrictions on emissions (or a global carbon tax) can impose a particularly large burden on such countries by forcing them to rely more on external energy sources.

It will be much easier to reach a global agreement on global warming if we can make progress in achieving greater international security.[22]

Terrestrial Carbon and Carbon Conservation

Terrestrial carbon—carbon that is embedded in trees, in agriculture, in meadows, and so on—provides a particularly difficult challenge, conceptually and

in terms of implementation.[23] Conceptually, it forces us to think through clearly stock/flow distinctions. Much of the discussion focuses on emissions—the flow of carbon into the atmosphere and how to limit it. Some scientists are hoping that the development of carbon storage technology will allow fossil fuels to be burned and the resulting carbon to be returned back below the ground. Carbon can, of course, enter into the atmosphere in a variety of ways—the melting of the tundra presents a major risk today.

What is of concern is the stock of carbon in the atmosphere. And that is affected not only by the flow into the atmosphere but also the flow out, for example, through the activities of plants.

We can thus look at the problem of the atmospheric concentration of carbon in two ways: at the dynamics—a flow into and out of the atmosphere or at the stocks—the amount of carbon "stored" below the ground (fossil fuels), in the ground (terrestrial carbon, carbon sequestered in trees), and in the atmosphere.

It is hard to monitor all of the flows into or out of the atmosphere. Consider deforestation. It occurs at millions of points on the globe. Moreover, only part of the wood from a tree that has been harvested will be used as fuel, and therefore contributes directly to amounts of carbon in the atmosphere. Carbon stored in wood used for furniture or construction enters the atmosphere only slowly, through decay. At the same time, cutting down a forest may lead to far more carbon entering the atmosphere than the carbon from the burning itself; carbon can be released from the soil (from the roots). Further, forests are major absorbers of carbon through photosynthesis; thus the elimination of trees also contributes indirectly to increases of atmospheric carbon. Those using wood as fuel should be charged for these indirect releases of carbon into the atmosphere; those using wood for long-lasting construction should be given some credit for the carbon storage.

It may be useful to think about how one might design a system if perfect monitoring were possible. When a tree is cut down, a charge would be made for the indirect emissions into the atmosphere. When the wood is burned, a charge would be made for the carbon entering the atmosphere. And when wood is used for construction, a charge would be imposed as the wood rots and the carbon enters the atmosphere.

In other words, given that our focus is on carbon *in the atmosphere,* a "toll" would be imposed on the individual responsible every time a carbon molecule enters the atmosphere—on those whose action "accounts" for the entry. (The charge would take into account the expected duration of the

carbon in the atmosphere—which is sufficiently long that it may be approximated by infinity.[24])

Alternatively, we could impose charges and provide payments based on "stocks." Those maintaining stocks of carbon below or on the ground would be rewarded. Thus, those maintaining forests (in which carbon is sequestered) would be paid for keeping their carbon molecules there, rather than in the atmosphere. The flow and stock systems could be made equivalent: An owner of fossil fuels would be rewarded for maintaining his carbon beneath the ground; the present discounted value of reduced payments, as a result of burning the fossil fuel, are equivalent to the "charge" for emissions. Of course, from a property rights perspective, the systems are quite different: One implicitly assigns the right to the owner of the forest to pollute and pays him not to pollute; the other gives him no rights to pollute and forces him to compensate should he pollute.

It should be clear that the "ideal" monitoring required for the implementation of either of these schemes is impossible. We will be looking for second-best approximations. One approximation that may do well—at least in the long run—focuses on the steady state, making use of the fact that forests are renewable. A forest takes out carbon from the atmosphere and stores it (not only in the tree itself but in the root system). In steady state, the tree and its products are decaying at the same rate that carbon is being taken out of the atmosphere. The problem with current biofuel policies is that, while recognizing that we are failing to take account of the cost of carbon in fossil fuels and the advantages of biofuels, we fail to take account of the carbon storage opportunity costs. *But if we give the forest "credit" for the carbon that it has stored (carbon that is not in the atmosphere)—including carbon that is stored in post-cutting uses (construction, furniture), we ameliorate this* to some degree. Over time, such a credit will incentivize switching land from its current production patterns (that pay no attention to carbon storage) to patterns that recognize the social value of carbon storage. Simultaneously, this entails an increase in the price of energy and in the tax on carbon. These price signals induce more and more land to be switched into uses that do better in carbon sequestration, and lead to less reliance being placed on fossil fuels for energy production. As we set lower limits on the levels of acceptable carbon concentration in the atmosphere, the switch from fossil fuels to renewables will occur more rapidly.

In appendix C, we contrast the optimal pattern of extraction of fossil fuels with the patterns actually observed, while in appendix D, I provide a

detailed analysis of the carbon conservation equation, including its implications for switching from fossil fuels to renewables.

Global Governance and Climate Change

The failure of the international community to reach a global agreement on climate change is, perhaps, the most vivid demonstration of the inadequacies in global governance. During the late nineteenth and twentieth centuries, the nation-state succeeded in protecting its citizens against some of the central risks that they faced. Global warming is a quintessential global risk. And delay in taking action may prove very costly. Effective action has to be global; but given the deficiencies in the current system of global governance, action adequate to what needs to be done has yet to be taken.

Perhaps worse still, the way negotiations and agreements—not just in the context of climate change but in other arenas of concern to developing countries—have proceeded so far may have actually undermined the kind of trust that is so necessary to reaching a global agreement. Distrust builds on itself, thereby compounding the difficulties of reaching global accords in areas of common concern.

For instance, in the trade arena, there was a widely perceived Grand Bargain leading to the Uruguay Round agreement in 1994. This entailed the developing countries giving the developed countries what they wanted—bringing into the ambit of trade agreements intellectual property agreement (the Agreement on Trade Related Aspects of Intellectual Property Rights or TRIPS) and financial services; and the developed countries in turn giving the developing countries what they wanted—elimination of agricultural subsidies and trade restrictions on textiles (the multi-fiber agreement). The developed countries got what they wanted; the developing countries had to wait a decade for the full elimination of the multi-fiber agreement and are still awaiting meaningful concessions in agriculture—especially on cotton. By the same token, at Doha, the developing countries reluctantly agreed to another round of trade negotiations, but in return, the developed countries agreed that it would address the imbalances of the past and called the round of negotiations that began a "development round." But within a few years, the development content was largely removed, so much so that it no longer deserved to be called a development round (Charlton and Stiglitz 2005). And the notion that it was supposed to redress the imbalances of the past was soon forgotten and the developed countries demanded

concessions from the developing countries commensurate with those that they were making. Moreover, the developing countries seemed to have to negotiate anew reductions in, say, cotton subsidies—even though the WTO had already ruled that the U.S. cotton subsidies were WTO-illegal.

So too, the failure of the developed countries to live up to the obligations they undertook earlier in climate change has undermined a sense of trust. In the 1992 Rio agreement, for instance, the developed countries agreed to make transfers of technology and to finance the incremental costs of emission reductions for developing countries. Yet in recent negotiations, the developing countries have had to renegotiate—as if these agreements had not previously been made. Rebuilding that trust is important, perhaps even necessary, for any meaningful negotiations going forward. One way in which that could be done would be for developed countries to undertake ambitious reductions domestically,[25] and to make and follow through on explicit, quantified, and binding commitments to technology transfer and to resources. So too, the institutional arrangements for the disbursement of funds must have the confidence of developing countries, which means, at a minimum, that they must have adequate voice within these arrangements.

Concluding Remarks

The world is engaged in a risky experiment, increasing to dangerous levels the atmospheric concentrations of greenhouse gases. Though we may not yet know the full consequences of this experiment, the risks are sufficiently great that there is a growing consensus that there must be marked reductions in the level of emissions. And given developing countries' aspirations of growth—and increasing evidence that many of these aspirations will be realized—the reductions within the developed countries will have to be all the greater. The total costs of meeting the requisite reductions will depend, to a large extent, on advances in technology. For the last two hundred years, much of the innovation in the West has been directed at saving labor; little has been directed at reducing emissions. This outcome is hardly surprising: With the atmosphere treated as if it were a free good, there were no incentives in place. This suggests that there may be ample opportunities for technological advances.

But the pace of innovation is uncertain, and it would be foolhardy to rely on such advances. It is imperative that wealthy countries change, as well, patterns of consumption—patterns that regrettably are all too often

emulated in the developing world. There is a need for a new economic model, one that centers less on the production of emission-intensive goods and more on other things that individuals and societies value. Changes in relative prices, reflecting the scarcity value of air and water, will help facilitate these changes but so too will other government policies.

This chapter has focused on how the burden of saving the planet should be shared between rich countries and poor and how reaching an agreement may be affected by the "framework," that is, whether the negotiations are over a common set of standards, a common tax, or a set of emission rights.

There is no question that there will have to be *global* reductions. That is not the question. The question is, upon whom should the incidence of the cost of adjustment be imposed? Avoiding global warming is a global public good. Standard public finance theory provides clear guidance about how to achieve such reductions in the most efficient way and how the burden should be shared. Clearly, the brunt of the burden (under virtually any welfare criterion) should lie with the advanced industrial countries. Indeed, these standard ideas suggest that even the approach often taken by developing countries—that there should be equal emissions permits per capita—puts an excessive burden on developing countries.

One way out of a political gridlock is to pass the costs on to someone not at the table—in this case, future generations. And that is what is at issue: whether we will continue to consume and produce as we have, preserving current living standards, at the expense of future generations. In the past, inequities in the global balance of power would have provided an easy alternative way out of the current global impasse: Pass the costs on to those too weak to defend themselves. The trade-offs are stark: between the living standards of the well-off today and those of the poor in the developing countries. The world has changed in the past decade, and this last strategy is not available.

One of the advantages of an agreed-upon common tax rate (with each country keeping its tax revenues) is that it reduces the scope for redistributive deadlock; most countries will, in fact, be better off moving from labor or savings taxes to a carbon tax—and the differences in the welfare costs are likely to be small.

A carbon tax may be the best way of avoiding the impasse confronting the world in addressing climate change, so evident in the failure to reach an effective agreement at the Copenhagen meeting. Unless some effective agreement is reached, the world will be facing enormous risk in the coming years.

Notes

1. A public good is something from which everyone benefits (consumption is "non-rivalrous.") It was first precisely defined by Paul Samuelson in 1954. The concept of global or international public goods—goods from which everyone in the world benefits—was first articulated in J. E. Stiglitz (1995).

2. More generally, many of the attempts to encourage renewables are far from well-designed. They are designed as much to increase the profits of the firms lobbying for them as they are to address the problems of global warming.

3. The views in the following paragraphs reflect the consensus reached at the Brooks World Poverty Institute/IPD Meeting in Manchester, June, 2009.

4. After the great Cambridge economist, A.C. Pigou, who first discussed them.

5. There is, of course, a problem of implementation, verification and credibility: How can one be sure that the level of emissions against which the tax was collected was accurately assessed?

6. In fact, in equilibrium with a fixed supply of oil with zero extraction costs, the price paid by consumers would remain unchanged—all of the oil will be extracted, with more of the value of the oil going to the government in the form of taxes and less in rents. In practice, with some private ownership and upward sloping supply curves, there will be some impact on total extraction.

7. Indeed, as we shall argue here, because we are concerned about long-run concentrations, we will almost surely want to keep large amounts of fossil fuels beneath the ground—with the optimal tax, rents will be zero for much of the existing supply. It is not surprising that owners of large amounts of fossil fuels are unhappy about this outcome. And most of the interventions discussed here do not focus on ensuring that they are fully compensated. There is, I suspect, widespread sentiment that it was luck that resulted in their wealth—the good luck of being born on land under which there was oil—and it is similarly the bad luck of the reality of global warming that is now taking that wealth away from them. Without this oil wealth, they may, of course, need assistance.

8. As a result of poorly designed intellectual property rights.

9. Although their economic losses may be partially offset by benefits arising from reduced global warming.

10. It is, of course, more likely that imposing a price on carbon will lead to reduced emissions of carbon. This will be the case so long as the "aggregate emissions curve" is downward sloping.

11. I say roughly because some countries have agreed to somewhat larger reductions from their base than others.

12. Nor have we accounted for the decreased rents received by owners of fossil fuels.

13. This is likely to be especially so in the future, as the increased importance of biofuels may result in food prices being more highly correlated with energy prices.

14. This is, of course, a normative statement, based on widely accepted normative assumptions, such as that the burden of taxes should be placed disproportionately on the rich.

15. A tax reduces the well-being of the taxed party at the same time as it increases the revenue of the government. Most taxes distort behavior, and thus impose a burden on the taxed party that goes beyond the value of the money transferred to the governmentf. This is the *dead weight loss*.

16. We can look at these payments in two different ways—as political bribes not to oppose the legislation to curb emissions, or as payments to ensure that the outcomes are Pareto superior. As we will explain, however, granting emission permits on the basis of past emissions provides overcompensation.

17. The problems may be exacerbated if monetary authorities subscribe to simplistic rules of inflation targeting; the large increases in energy prices then induce large increases in interest rates, which in turn lead to a slowing down of the economy and an underutilization of resources, with especially adverse effects on the poor.

18. There are other ways that government may affect the level of emissions. Standard welfare theory begins with the *assumption* of exogenous preferences. Yet we know that preferences themselves are endogenous, affected, for instance, by advertising and social processes. Government policies can help shape the evolution of preferences.

19. See, for instance, J. E. Stiglitz (2006b), *Making Globalization Work,* chapter 4. The adverse effects arise from several sources: (a) the patent system gives rise to monopoly power, which lowers levels of production, reducing incentives to innovate; (b) the patent system increases the cost of the most important input into innovation—knowledge; and (c) the patent system gives rise to a high risk of patent litigation, especially in the context of the patent thicket—where there is some probability that any innovation will trespass on others' intellectual property.

20. Assume, for instance, that with existing technologies, emission per unit of output is e_o. Assume the country signed an agreement to reduce emissions below the level of E*, that the international agreement has sufficient sanctions that the country will comply, and that in the absence of commitment, it would have produced an output of Q_o and emissions of $Q_o e_o$. To comply with its commitment, the country would have to restrict output to E^*/e_o. If the new technology lowers emissions per unit of output to e_1, sufficient that at Q_o the country can meet its obligations, then the owner of the new technology can extract a rent up to $[Q_o - E^*/e_o]$.

21. Confidence that the developed countries will live up to their commitments is undermined too by the fact that most have not lived up to previous commitments to provide 0.7 percent of their GDP for development assistance.

22. See appendix A for how we can incorporate these security costs into the analysis.

23. The equations and explanation of my conclusions for this section of the paper are included in appendix D.

24. That is, a carbon molecule can be thought of as renting space in the atmosphere. If the rent per unit time were *c*, and there were a decay rate of μ, and the interest rate is r, then the entry charge would be $c/(\mu + r)$. Of course, we don't care about how long any particular molecule stays in the atmosphere; we don't have to track each. We care about the average. If b= o zero (a molecule never leaves), then the entry charge is c/r.

25. Many of the developed countries did not live up to the commitments on reductions that they had made earlier; the United States has increased its emissions since 1992 almost unabated.