THE CHALLENGE OF CLIMATE CHANGE POLICY
A White Paper for Teaching[1]

William A. Sundstrom
Dept. of Economics
Santa Clara University
Santa Clara, CA 95053
wsundstrom@scu.edu

DRAFT of 5/21/09

The problem of global climate change poses a policy challenge of unprecedented scale and complexity. A number of features of the climate problem conspire to favor inertia—doing nothing, or doing too little too late. Of course, the option of "doing nothing" is illusory; the choice is really between doing something now to mitigate the greenhouse effect, and forcing future generations to do something to adapt to the consequences over the coming centuries. Contrary to claims one sometimes hears that efforts to limit greenhouse gases will be largely painless, paying for themselves in terms of improved efficiency and new "green" jobs, preventing the potentially catastrophic effects of dramatic global warming will in reality require significant sacrifices of wellbeing by the present generations. The goal of sensible climate policy must be to implement the necessary changes with minimal pain, paying particular attention to the impact on the poorer and more vulnerable elements of the world population. This section discusses some of the general considerations in analyzing the benefits and costs of climate change policy, and then explains and compares important policy alternatives, such as regulation or subsidization of energy sources, the carbon tax, and cap-and-trade programs.

1. Climate change as an environmental problem: commonality and uniqueness

At base, the climate change problem resembles many standard pollution problems. Greenhouse gases (GHG) are the undesirable byproduct of a variety of human activities—most notably the burning of fossil fuels—and in this regard are no different in principle from other pollutants: airborne particulates, sewage, toxic wastes, and the like. Modern economic society would be impossible without some minimal amount of environmental pollution, and thus a zero-pollution standard is unrealistic, but for well-understood reasons there is a strong tendency to

[1] This paper is part of a larger project to prepare teaching resources on climate change, joint with Stephen Smith, Dept. of OMIS, Santa Clara University. The generous support of the Center for Science, Technology, and Society at Santa Clara University is gratefully acknowledged.

produce too much pollution. Ignorance of the adverse consequences of environmental degradation is one reason, and in the case of GHG emissions this was a viable excuse for "business as usual" up until perhaps a decade ago, by which time the scientific consensus on anthropogenic climate forcing had formed. But perhaps more importantly, pollution in an unregulated environment constitutes what economists refer to as a *negative externality*: each individual source of pollutants suffers little of the cost of pollution damages caused by her or his actions. Well-meaning people may voluntarily undertake some of the efforts required to clean up: many of us will purchase compact fluorescent bulbs or buy a smaller car even if our direct personal incentive to do so is minimal. But inevitably, not everyone is well-meaning, and not every well-meaning person will go far enough. The potential for free riding on the sacrifices of others implies inadequate individual incentives to achieve the socially desirable reductions in GHGs. This fundamental incentive problem implies that an unregulated free market fails to achieve the optimality lauded by generations of economists, going back at least to Adam Smith. Indeed, externalities constitute one of the core justifications for government intervention in markets in modern mainstream economic theory.

Because of the basic similarity between the climate change problem and myriad other well-analyzed environmental policy problems, standard economic theory has a great deal to offer to discussions of policy design in the climate change area, as we shall argue. But there are obviously important aspects of climate change that, taken together, make it a unique and particularly daunting policy challenge.

Scale. Climate change is quite literally a global issue, and likely involves consequences for human life and ecosystems on an unprecedentedly large scale. Perhaps the only human-caused environmental threat of comparable scale would be global thermonuclear war. But unlike that threat, climate change is being exacerbated constantly by the everyday activities of every single human on the planet (which is not to deny that some contribute many times more than others). In this sense GHG emissions are the negative externality par excellence: every person contributes a tiny fraction to the climate change that will be experienced by every other person, but each person's individual selfish incentive to do something about it is vanishingly small, at least in an unregulated regime.

Time horizon and (near) irreversibility. The worst impacts of climate change will fall on humans several generations in the future. By the standards of geological time, the changes are taking place in the blink of an eye—indeed, the pace of change is one compelling piece of evidence that human activity is a significant contributing factor. But if "the past is a different country," so too is the future a century from now, at least by the standards of human decision-making. The elongated time horizon of climate change is related to the problem of persistence in the processes creating it. Were we to stop the flow of pollutants into a river immediately, in most cases the flow of water and natural processes of dilution and decomposition would result in immediate improvements in the health of the river, and often ecosystem recovery within a matter of decades if not years. In the case of climate change, even if human releases of GHG ceased entirely today, global temperatures would continue to rise for another half century; this means that climate policy is a lot like steering an aircraft carrier. Climate change is not truly irreversible, but any changes in behavior now have a long gestation period before their impacts are felt.

The long time horizon of the climate change issue poses significant ethical and economic challenges. In their everyday behavior, most humans exhibit a degree of impatience and myopia: present enjoyments are generally valued more highly than future. A consequence of this basic behavioral trait is that interest rates are positive: borrowing involves inducing someone to part with the use of her money today, and she will only be willing to do so if the *quid quo pro* is that she be paid back a larger amount in the future. Applied to social decisions, this logic implies that current expenses can be justified by future benefits only if those future benefits are sufficiently greater than the present costs to pay back "with interest." As we explain below, this *time-discounting* can have striking implications over very long time horizons. Applying a discount (interest) rate of the magnitude typically used in benefit-cost analysis for environmental and other policy applications makes it difficult to justify taking dramatic action to mitigate climate change today. But as we also argue, whether it is reasonable to apply such a large discount rate to decision-making across generations is highly debatable. In our view, an ethically defensible discount rate justifies immediate and drastic action to cut GHG emissions.

Uncertainty. Most human decisions, whether by individuals or large organizations such as governments, are made with imperfect information: the consequences of action are therefore fundamentally uncertain. The scale, complexity, and time horizon involved in climate change policy decisions serve to increase the degree of uncertainty. The sources of uncertainty include the sizeable variance in model predictions of global temperature change, uncertainty about the effects of any given global temperature change on physical and biological processes on earth (both globally and locally), uncertainty about the human, social, and economic costs of these changes, and uncertainty about the cost and effectiveness of potential technological and policy fixes for the GHG problem.

Making good policy decisions under imperfect information (uncertainty) involves two fundamental tasks. First, to the extent possible, the nature and size of the risks must be assessed. For example, what is a 95 percent confidence interval for the range of predicted temperature increases under business as usual? What is the worst-case scenario for coastal flooding as ocean levels rise, how many people would be displaced, and what is the likelihood of such an outcome? Second, having characterized the possible scenarios and the approximate probability of their occurrence, judgments must be made about how to take account of the risks in deciding what to do. Again drawing on the example of individual decision making, most people exhibit some degree of *risk aversion*: given the likelihood of occurrence of different outcomes, people give greater weight to bad outcomes than to good ones.[2] These motivations help explain the market for insurance, asset diversification, and many other social and economic phenomena. It seems reasonable that policy decisions should also reflect aversion to risk. But while an individual's risk aversion can be estimated by observing her behavior in risky situations and insurance markets, there is no agreed-upon scientific method for measuring "society's" appropriate degree of risk aversion. What is clear, however, is that policy makers must consider not only the most likely or expected future scenarios, but the full range of possibilities, and should in particular insure against the bad ones.[3] In this regard expenditures on climate change mitigation are analogous to defense expenditures, both representing insurance against large-scale societal risks.

Global politics. The global scale of the climate change problem poses a special set of challenges that go beyond the mere numbers of people involved and magnitude of costs mentioned above. These have to do with the fundamentally transnational character of the issue. Because there is no global government, there is no entity that can set and enforce policy responses appropriate to the scale of the problem. On both efficiency and equity grounds there is strong reason to seek internationally coordinated approaches to reducing GHG emissions. The very free-rider problems that plague negative externalities at the level of individual decisions apply at the level of nation-states as well. In a free-trade regime, for example, even if many countries undertake policies to reduce GHG emissions, countries that forego reductions could become "pollution havens," specializing in the production and export of carbon-intensive goods, with the result that net global emissions do not decrease. Fairness considerations are also important: given the huge disparities in average income and life chances across countries, is it fair to demand of the poor sacrifices that the rich managed to avoid by industrializing earlier?

The Kyoto Protocol was a first attempt to achieve coordination through treaty negotiations, but the goals were modest and binding only on the developed countries. Achieving binding agreements that include both developed industrial and poorer, developing nations has

[2] Extreme risk aversion would imply basing all decisions strictly on the worst-case scenario, no matter how unlikely.

[3] But for a cautionary note regarding this "insurance premium" argument, see Nordhaus, "Challenge of Global Warming" (2007), pp. 111-116.

proven to be a daunting task. For the foreseeable future, the most effective policies to mitigate climate change are likely to be implemented at the level of nations or even, as in the case of the United States, smaller political units such as state governments.

2. Comparing costs and benefits

Given the potential for catastrophic effects of climate change if we persist on a path of business as usual, it may seem a "no-brainer" that strong measures need to be taken to curb GHG emissions as soon as possible. But how strong? And given the dependence of modern economies on fossil fuel, is it not possible that the necessary measures will simply impose too much hardship on current generations to be morally defensible, let alone politically feasible? The questions have particular bite when it comes to imposing difficult choices on developing countries that have not yet enjoyed the material benefits afforded by a carbon economy.

In recent years economists have attempted to model the global economic consequences of alternative paths for climate policy. An important goal of such studies is to compare the costs of reducing greenhouse gas emissions, for example through conservation or the adoption of cleaner energy sources, with the costs of future climate damages if emissions are not reduced. The approach of benefit-cost analysis (BCA) is to argue that measures to mitigate GHG emissions should be undertaken so long as the additional benefits in terms of reduced damages exceed their economic costs. The estimates of costs and benefits are often derived from simulations that combine models of climate, energy technology, and markets and economic growth (such models are known as *integrated assessment models*, IAMs). Needless to say, these exercises are fraught with layers of uncertainty about the underlying physical and economic processes and require tricky judgments about the proper specifications and parameters. We summarize some of the findings here, and discuss some of the reasons that different analyses can come up with quite different prescriptions for the proper pace and magnitude of GHG reductions. Although thoughtful and highly qualified experts can and do disagree about these issues, one clear message has emerged from most recent research in the area: taking immediate action to begin reducing GHG emissions substantially, with a view toward eventually stabilizing global concentrations at a level that would prevent dire effects, passes a cost-benefit test. In other words, proponents of inaction cannot take shelter behind the economics.

Estimates of the dollar cost of climate change under business as usual scenarios can be expressed in a variety of ways. The total dollar figures involved are imponderably large, and for that reason are not easy to evaluate, even in an economic policy climate in which we have become somewhat accustomed to dealing with magnitudes in the trillions of dollars. Two useful ways of framing the costs of climate change and its mitigation are frequently used: costs expressed as a percentage of GDP or some similar measure of overall economic activity; or expressed in dollars per ton of carbon or carbon dioxide equivalent. The latter is particularly helpful, as we shall see, in thinking about policy measures that place an implicit or explicit price on carbon dioxide emissions.

The most widely cited and debated estimate of the costs of global warming and of the measures that might be taken to reduce it is the Stern Review, a report to the U.K. Treasury prepared by Nicholas Stern and his collaborators (Stern 2006). The methodology of the project is mainstream benefit-cost analysis: the predicted future time path of climate-change impacts is evaluated in monetary terms, and the net value calculated. For example, the costs associated with warming would include the predicted total value of crop losses due to drought, economic losses due to flooding and extreme weather events, health costs associated with increased incidence of tropical disease, estimated losses of value placed on species extinctions, etc. The value of some benefits to humans, such as more moderate winters at the higher latitudes, is also netted out of the costs.

In Stern's case, estimates are often presented as percentages of global gross domestic product (GDP)—the total annual value of the goods and services produced by the world economy. Expressing it this way should help assure non-economists that benefit-cost analysis is not based on the notion that money is all that matters: rather, it is an attempt to express real costs and benefits in terms of their implications for our material wellbeing. If the cost of an activity is 5 percent of GDP, it means that our ability to consume the goods and services we value is 5 percent less than it otherwise would be. Stern's bottom line is presented in the report's executive summary:

> Using the results from formal economic models, the Review estimates that if we don't act, the overall costs and risks of climate change will be equivalent to losing at least 5% of global GDP each year, now and forever. If a wider range of risks and impacts is taken into account, the estimates of damage could rise to 20% of GDP or more. In contrast, the costs of action – reducing greenhouse gas emissions to avoid the worst impacts of climate change – can be limited to around 1% of global GDP each year (Stern 2006, Executive Summary (short), p. vi).[4]

Stern's findings seemed to make the case that strong action against climate change could be justified by hard-headed economic logic, not just wooly-headed environmentalism. But the

---

[4] Stern's estimated costs of action are based on stabilizing global carbon dioxide-equivalent concentrations at 550 ppm.

results were immediately challenged by other highly respected economists working in the climate change field. Among the most important of these Stern skeptics is William Nordhaus, a Yale economist who has developed his own sophisticated climate-change model to calculate optimal GHG policy paths.[5] Nordhaus also favors taking action, but more modest and gradual action. His estimates of future costs of warming, and therefore the kind of action that can be justified to prevent it, are much smaller than Stern's. For example, Nordhaus shows that Stern's optimal carbon tax, which would discourage the use of fossil fuels by increasing their price, is about ten times larger than the one Nordhaus advocates.[6]

The wide range of cost estimates generated in the literature is shown in Tol (2005). Among results published in peer-reviewed journals, the mean estimate of the cost of damages resulting from the release of one more metric ton of CO2 (sometimes referred to as the *social cost of carbon*) was $43, with a standard deviation of $83.[7] The distribution of estimates is highly skewed to the right: that is, many peer-reviewed studies come up with estimates well below the mean of $43, but a few generate much higher estimates. The same pattern is also true of individual studies that provide a range of estimates. In other words, the estimates suggest that truly catastrophic costs are not likely, but possible. If policy makers (and world citizens) are risk averse and wish to avoid the most catastrophic outcomes, of course, greater weight should be placed on the estimates in the high estimates in the right tail.

Economists also disagree about the overall costs to society of reducing greenhouse gas emissions to a safe level. While Stern estimated that this could be accomplished at a cost of about one percent of global GDP, other estimates suggest somewhat greater costs. A recent study by researchers with the MIT Joint Program on the Science and Policy of Global Change, for example, predicts that the economic cost of an 80 percent reduction in CO2e emissions by 2050 is around 2 to 3% of global "welfare" (Paltsev et al 2009). Mitigation cost estimates vary depending on a variety of assumptions about the nature and pace of technological change and the ability of people and businesses to make adjustments in their behavior in response to changes in energy prices (see Keohane and Goldmark 2008).

[5] DICE, for Dynamic Integrated model of Climate and the Economy. See Nordhaus, "The Challenge of Global Warming" (2007).
[6] Nordhaus, "Critical Assumptions" (2007).
[7] Results appearing in peer-reviewed journals are vetted by anonymous experts in the field, and thus are widely considered more scientifically valid than those that are not peer reviewed. The average social cost of carbon estimate in non-peer-reviewed articles was higher than that in the peer-reviewed work.

Discounting in the economics of climate change

Surprisingly, the single factor most responsible for the large differences between cost estimates of future climate change damages has nothing to do with assumptions about temperature or weather changes, future technologies, or market responses. Rather, it is a single parameter: the discount rate used to weigh future against present monetary values. Nordhaus, following what is fairly standard practice in conventional benefit-cost analysis, uses a relatively high discount rate: approximately 5 percent per annum. Stern's discount rate is much lower: about 2 percent. Looking across a large number of studies, Tol (2005) shows that discounting accounts for much of the variation in estimates of the social cost of carbon, along with differences in how impacts are weighted across rich and poor countries.

Time discounting is used to take account of the fact that people typically value present consumption more highly than future. The discount rate, which captures this phenomenon quantitatively, works like an interest rate. The higher the discount rate, the more the present is valued over the future. And like interest, discounting compounds. Over the very long time horizons that characterize the climate change issue, compounding has huge effects. This can be seen by comparing the *present discounted value* (PDV) of a dollar of damages assumed to occur a century from now. This calculation answers the following question: To prevent $1 of damages a century from now, what is the most we should be willing to give up now, given the discount rate?

Three answers are provided in the table below. The "naïve egalitarian" wants to treat all generations equally, and therefore values a dollar a century from now the same as a dollar today. This implies zero discounting. (As we will argue below, it's not so clear that this is really what an egalitarian should do.) Under Stern's discount rate, that anticipated dollar of damages a century from now is worth only 14 cents today. But given the colossal magnitude of the predicted future damages, even 14 cents on the dollar implies that we should be willing to make huge sacrifices now. Perhaps the most striking result comes from using Nordhaus's 5 percent discount rate: the PDV of that future dollar nearly vanishes to less than a penny. This is the power of compounding over 100 years.[8] It is hardly surprising, then, that Nordhaus—along with the many economists who agree with his approach—advocates fairly modest measures to curb GHGs, compared with Stern and many environmentalists.

[8] The formula use for continuous compounding is PDV = exp(-rt), where r is the annual discount rate, and t is the number of years (100 in this case).

	Discount rate (%)	PDV of $1 of damages 100 years from now
“Naïve egalitarian”	0	$1.00
“Stern”	2	$0.14
“Nordhaus”	5	< $0.01

All of this naturally leads to the question of on what grounds a discount rate should be chosen. There are two broad approaches to thinking about the appropriate discount rate for social decisions like climate change policy. First, one might look to real-world market interest rates or returns on capital for guidance. This “descriptive” approach is one of Nordhaus’s justifications for using a higher rate: he notes that long-run historical real rates of return on corporate capital are on the order of 7 percent (“Critical Assumptions”, 2007, p. 202). Much hinges, however, on which asset returns are examined. The long-run real rate of return on “risk-free” government bonds, for example, tends to be considerably lower.

It might strike some as peculiar to base critical decisions about how we treat future generations—who after all have no say in the matter—on some market rate of return. Yet the practice can be defended. The standard justification for using a real-world rate of return runs something like this. Suppose we faced a stark choice between taking drastic action now, at considerable expense, to slow climate change, and taking no action, instead investing the money we would have spent in a safe, high-yielding investment fund, to be withdrawn in 100 years by our descendants, as they try to adapt to the changed climate. The logic of benefit-cost analysis rests on a hypothetical question we would pose those future citizens: would you rather have the lower temperatures, or take the cash? If Nordhaus has correctly estimated the future damages of climate change, what he is saying is that they would go for the cash.

Needless to say, the hypothetical “environmental bond” envisioned here would not actually exist. In order for the justification of the descriptive approach to work in the real world, then, we need instead to assume that all or most of the savings from doing nothing about climate change now would go into productive capital investments that yield the high rate of return to future generations, an assumption many have questioned.[9] Consequently, many economists, including Stern, have turned to the second, “prescriptive” approach. Here we do not appeal to real-world interest rates but instead try to derive the appropriate discount rate from ethical first principles and some assumptions about individual values.

[9] See Arrow 1995.

The prescriptive approach is usually explained with a bit of math, but the intuition is not especially difficult. Suppose that the wellbeing (or utility) of each generation depends on their consumption of goods and services, net of any costs imposed by environmental damages. Wellbeing increases with consumption per capita, but at a decreasing rate. This standard assumption (decreasing marginal utility) means that an additional dollar worth of consumption gives a bigger boost to the wellbeing of a poor person than it does for a rich person. Both introspection and a variety of evidence support this assumption, and it can also be defended on normative grounds as egalitarian.

Now suppose we take as our policy objective the maximization of total wellbeing across all the generations into the indefinite future, giving equal weight to the wellbeing of each person.[10] By treating all generations symmetrically, we are assuming that the *rate of pure time preference* is zero: that is, our comparisons of wellbeing give no preference to any generation simply because of its location in time. This idea is ethically appealing, because it embodies the ideal of impartiality. For example, in the influential justice theory of John Rawls (1971), distributional principles are to be chosen impartially from behind a "veil of ignorance"—as if we did not know which person we might become in real life, and had an equal chance of being anyone. Extending Rawls's logic to the intergenerational setting, if we could end up belonging to any generation, we would want each generation to be treated symmetrically, without any pure time preference.[11]

Even with zero time preference, however, the prescriptive approach generally implies that consumption streams should be discounted, so long as there is economic growth. This is a consequence of decreasing marginal utility. If each generation is expected to be somewhat richer (per capita) than the last, then a dollar is worth more to people now than it will be to people in the future, in terms of its impact on wellbeing. Thus the egalitarianism implicit in the assumption of decreasing marginal utility demands a positive discount rate in benefit-cost analysis: other things equal, the approach favors redistribution from rich to poor, and the people who come along earlier are the relatively poor ones, so long as economic growth continues. By the same logic, if climate change were expected to be so extreme as to lead to negative growth in the net value of consumption, reversing the established trend of the past two centuries, the appropriate social

[10] For simplicity I am implicitly assuming that the population is constant, which allows me to speak of the wellbeing of each generation as if it were a person. With population change, we'd have to decide whether we want to maximize the sum of individual wellbeing, giving more weight to more populous generations, or generational wellbeing. But the basic conclusions would be similar.

[11] Rawls was aware, however, of some paradoxical implications of assuming zero time preference. See Rawls 1971, Secs. 44-45.

discount rate would be negative. Even relatively pessimistic economic predictions, such as Stern's, do not generally yield such a prediction.

A more general version of the prescriptive approach allows for discounting future wellbeing directly—what the economists call pure time preference. It can then be shown that the appropriate social discount rate ρ for benefit-cost analysis is given by the sum of two terms:

$$\rho = \delta + \eta g$$

where δ is the rate of pure time preference, g is the annual growth rate of real consumption or income, and η is a measure of how quickly marginal utility is declining in income.[12] The second term captures the effect discussed above—that future consumption should be discounted when future generations are expected to be richer.

Stern assumes a very low (near-zero) value of δ, a value of η of about 1, and a 2 percent annual growth rate to arrive at his overall discount rate of about 2 percent.[13] Nordhaus ("A Review," 2007) argues that Stern's parameters are simply inconsistent with observed behavior, given a real-world rate of return of 5 percent and actual savings, and suggests that either δ, η, or both must be recalibrated to be made consistent with the real world. Nordhaus is convinced that a rate of pure time preference of δ = 3 percent is more consistent with how actual people make decisions, and he applies it in his preferred calculations. By changing δ from 0 to 3 percent in the Nordhaus model, Stern's call to drastic action morphs into Nordhaus's moderate gradualism.[14]

The basic issues at stake in the discounting debate should be fairly clear by now. A key question is whether decisions regarding climate change policy and the treatment of future generations should be based on the way current generations actually make similar decisions for themselves in their everyday lives—implying positive pure time preference—or rather based on the ethical ideal of impartiality—implying zero time preference. The divide between these views reflects something of a division between American and British public economists. The typical British view was expressed by the Cambridge economist and philosopher Frank Ramsey, who in his seminal 1928 paper on optimal savings wrote: "...it is assumed that we do not discount later enjoyments in comparison with earlier ones, a practice which is ethically indefensible and arises merely from the weakness of the imagination..." (1928, p. 543). A more recent statement, by the Princeton (but British-born) economist Angus Deaton in a short note on the debate over the Stern

[12] The elasticity of marginal utility with respect to consumption.
[13] Stern actually assumes a very small positive value of δ, justified not as time preference but to take account of the risk of human extinction.
[14] Stern's preferred value of 1 for the parameter η, which captures the extent to which marginal utility declines with wealth, is also relatively low. A lower value of η amounts to a greater tolerance of inequality, because it implies that an additional dollar is still worth a lot to a rich person.

review, argues that “Zero pure time preference, if it is a vice, is surely a minor one. Relying on markets to teach us ethics is very much worse” (2007, p. 4).[15]

Uncertainty and the possibility of catastrophic climate change

Another complicating factor in conducting benefit-cost analysis of climate change policy is the high degree of uncertainty. Uncertainty plagues estimates of both the costs of mitigation and the costs of damages due to warming. Climate change is far from unique in this regard, and there are standard approaches to incorporating risk into BCA and policy models. If policy makers, like most people, are risk averse, greater uncertainty about the potential future damages due to climate change should lead us to place somewhat more weight on the bad-case scenarios, and therefore undertake stronger measures to curb GHG emissions as a form of insurance. In a recent short essay on climate-change policy, for example, Kenneth Arrow (2007) suggests that, given the Stern Review estimates of the costs of climate change and the degree of uncertainty about them, action to stabilize carbon dioxide concentrations at around 550 ppm can be justified even assuming a rate of pure time preference as high as 8.5 percent.

An aspect of climate uncertainty that has entered the public consciousness but until very recently has received little attention in the policy analyses of economists is the possibility of truly devastating climate catastrophe. These are the quite unlikely but not impossible “doomsday” scenarios, in which business-as-usual GHG emissions lead to global temperature increasing on the order of 20°C or more. A temperature increase of that magnitude, over the tiny geological time of one or two centuries, would put extraordinary stress on global physical systems, organisms, ecosystems, and civilizations. Whether human society as we know it could survive in such a world is at least questionable.

The notion that societies have an obligation to avoid actions that could possibly cause catastrophic harm, if the science cannot rule it out, is at the core of the thinking underlying the so-called *precautionary principle*.[16] This principle has been met with skepticism by many policy analysts, especially in the United States, who fear the potentially paralyzing effects on technological progress of taking a view that seems to entail an unreasonably extreme degree of risk aversion. Instead, the standard approach of BCA noted above has been to apply expected utility (EU) analysis, taking account of risk aversion, but not allowing the worst-case scenario to completely dominate policy decisions. In most applications to climate change, the doomsday

[15] The British view goes back at least to the utilitarian philosopher Sidgwick (1907, p. 381).
[16] See http://en.wikipedia.org/wiki/Precautionary_principle .

scenarios are considered remote enough possibilities that they are largely ignored in the calculations.

Recent work by the economist Martin Weitzman (2007, 2008) has cast some doubt on the way conventional BCA has dealt with catastrophic climate change. Weitzman's argument is rather technical, but the intuition is fairly clear. Because the most extreme temperature changes due to anthropogenic warming would lie far outside the sample of observations science has used in climate modeling, and would represent very low-probability events, the underlying risk parameters are themselves highly uncertain estimates. In such a setting, the probability distribution of the potential climate-change damages, given the imperfect state of our knowledge, has "fat tails." In other words, catastrophe may be unlikely, but it is still too likely to ignore. And if the harm caused by catastrophic climate change is great enough, the risk of catastrophe may simply trump all other BCA considerations, including even a very high time discount rate and/or costs of mitigation that would impoverish current generations. Weitzman refers to his result as the "Dismal Theorem."

Beyond benefit-cost calculations?

We have noted a few times the ways in which the tradeoffs that benefit-cost analysis recognizes and seeks to quantify are similar in nature to small-scale tradeoffs and decisions that individuals must negotiate in daily life. Still, not surprisingly, many people are uncomfortable with basing weighty social decisions on comparisons of dollar valuations. There are several potential reasons for this queasiness, but most of them come down to one of two notions: first, that some values (ends) may be fundamentally incommensurable and thus cannot be reduced to a common dollar scale; and second, that there are competing ethical frameworks that may contradict or even trump the utilitarian calculus lying behind benefit-cost analysis.

Regarding the first issue, benefit-cost analysis offers well-known methodologies for estimating the willingness to pay for, say, the survival of an endangered species, or for a Chevy truck. This permits the comparison between a dollar of damages and a dollar of consumption. As Neumayer (2007) notes, however, it is at least plausible that were environmental degradation severe enough, no finite amount of compensation in terms of consumption would be sufficient to offset this "non-substitutable damage to and loss of natural capital." It is also debatable whether, even if a dollar value can be obtained for each good through some procedure, the resulting number helps us decide whether it is moral to allow the species to disappear if saving it would cost more trucks than it is "worth."[17] As for the second point, a rights-based framework might suggest that future generations have an inalienable right to inherit a minimally healthy global ecosystem from their ancestors (Neumayer 2007). These large issues in public philosophy lie well beyond the scope of this paper, let alone the expertise of its authors.

Political and global realities

The serious disagreements between economists about discounting and its implications for the optimal path of GHG reductions would be of practical concern if economists were in charge of climate policy, but of course they are not. GHG mitigation goals will be determined through real-world political processes, and for the time being there seems little threat that U.S. policy-makers will err in the direction of excessive reductions in emissions, even by Nordhaus standards. Further complicating the picture is the rapidly growing contribution to global emissions of developing, lower-income economies, led by China and India. China's rapid development of coal-fired electricity generation is indicative of the immense challenges that will be created by rapid economic development. As of 2007, an estimated 53 percent of greenhouse gases originated in developing nations. Given that climate change is a problem of the global commons, a coordinated international approach to solving it, involving both developed and developing countries, would make the most sense, and this indeed has been the objective of such efforts as Kyoto. But the developing nations have understandably placed a higher priority on economic growth and poverty reduction than on carbon emissions reduction, and they can call attention to the dramatically higher per-capita carbon emissions in the developed North and West, not to mention the rich countries' greater ability to pay, as justification of their own relative inaction.

Although industrialization is often associated with increased pollution, there is a view that economic development and growing affluence must eventually lead to improvements in environmental conditions, as people and their governments place increased value on a clean environment and the economy transitions to lower-polluting service industries. The graphical representation (see Figure 2.1 below) of this idea is the so-called *Environmental Kuznets Curve* (EKC). If the EKC hypothesis were generally correct, the sooner developing country incomes converged toward those in the West, the sooner would they bring their emissions under control. Unfortunately, empirical studies of the validity of the EKC hypothesis yield decidedly mixed results (Copeland and Taylor 2004). The relationship between pollution levels and per-capita income is not a simple one, and varies with the type of emission. Particularly troublesome for the climate change issue is that the kinds of pollution that appear to be most sensitive to rising

[17] On incommensurable goods, see Anderson 1993.

incomes are those with substantial local effects. In other words, if the Chinese behave like other folks, they are more likely to clean up their local sources of drinking water and urban smog before they start attending to carbon dioxide, the adverse effects of which are dispersed across the globe.

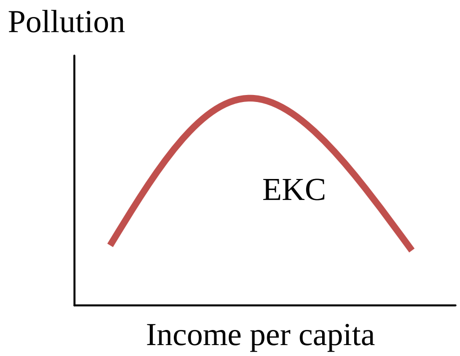

Figure 2.1: Environmental Kuznets Curve

That carbon dioxide emissions tend to grow with income up to a rather high level is confirmed by Figure 2.2, which plots per-capita CO2 emissions against per capita income (both variables logged) in 2004.[18] Each dot represents a country, with sizes proportional to population. For the countries of China (in red), Germany (orange), India (light blue), and the United States (yellow), the trails show the evolution of both variables between 1970 and 2004. The cross-section scatter of countries suggests a fairly tight and positive relationship between income and per-capita emissions, and the trails indicate that the economic growth of China, Germany, and India over the period was accompanied in each case by a substantial increase in per-capita emissions, in way roughly consistent with the cross-section snapshot. Interestingly, per-capita emissions in the United States are essentially flat over the same period, suggesting perhaps that the United States has reached the peak of the EKC for carbon dioxide. The evidence is also consistent with the possibility that some carbon-intensive manufacturing activities have been "outsourced" to industrializing nations.

[18] Source: Gapminder (http://www.gapminder.org/gapminder-world.html).

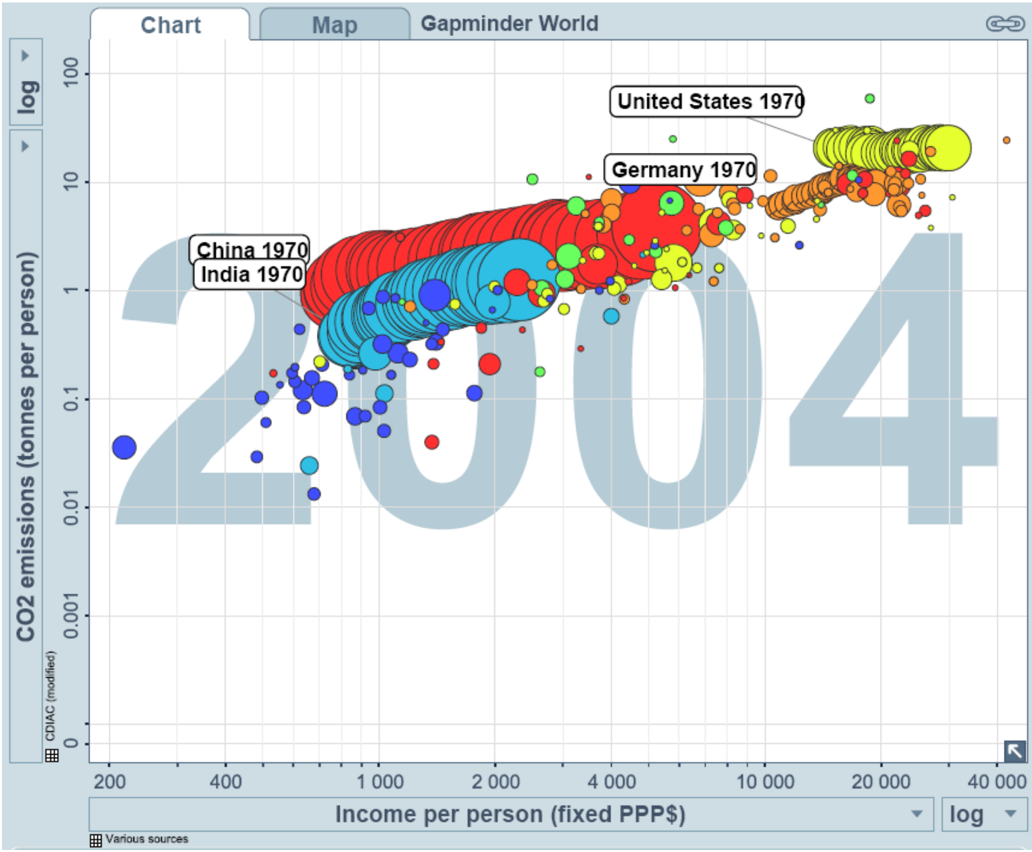

Figure 2.2: Per capita emissions and income (log scales)

Whatever the desirability of coordinated international action, and the Kyoto framework notwithstanding, further progress toward serious reductions in GHG emissions will likely originate in regional organizations (such as the EU), nation-states, or even sub-national units (such as states or cities). Targets and policies will be set by political processes, not by benefit-cost calculations or alternative ethical frameworks for public decision making. In the remainder of this paper, we turn to the guidance that economics can offer in the design of climate-change policies that can achieve the targets—however they are set—with minimal economic inefficiency, and proper attention to political and distributional concerns.

3. Designing climate-change policy

In many respects, the design of policies to reduce GHG emissions is a straightforward application of well-known pollution regulation principles. Alternative approaches include direct limitations on emissions, technology subsidies or standards, and market-based regulation, such as taxes or cap-and-trade. Because a large fraction of GHG emissions consists of CO2 generated by burning carbon-based fossil fuels, and there is a fairly direct correspondence between the amount of a specific fuel used and the amount of CO_2 released, regulation can actually target the sale and consumption of the fuels themselves rather than the resulting emissions.[19] Other aspects of the regulation of GHGs are more complicated. Methane gas released by farm animals or decomposition, for example, is less easily monitored. A complete policy should also encourage the sequestration of CO_2 from the atmosphere by plants or other mechanisms.

In this section we begin with an overview of the basic regulatory approaches to limiting negative externalities, such as GHG emissions. Throughout, we refer to the target of the regulation as quantities of CO_2 equivalents emitted into the atmosphere, understanding that in most cases the policy would actually be applied to the fuels that generate the emissions. We assume that the overall target for CO_2 emissions has been set, so policy design focuses on economic impacts and political feasibility.

Minimizing abatement costs

We have at our disposal a wide array of means of reducing our dependence on fossil fuels. These range from conservation (insulation, fuel-efficient cars) to low-carbon alternative energy technologies (solar, nuclear). To minimize the adverse economic impact of reductions in GHGs, climate change policy should aim to induce the mix of technologies and other measures that achieves the desired target at minimum total abatement (reduction) cost. Obviously, the least-cost methods should be used before resorting to costlier measures. Figure 3.1 depicts one set of estimates of the current cost of abatement for different methods, rank ordered from least to most costly. These estimates were prepared for the European energy company Vatenfall; cost is measured in Euros per metric ton of CO_2.[20] We use this figure not because we think it offers the best estimates, but because it helps illustrate important principles in efficient GHG regulation.

---

[19] The CO_2 produced by complete combustion of a unit of fuel is directly related to the carbon content of the fuel per unit. For example, complete combustion of a liter of regular gasoline would produce XXX kg of CO_2, compared with YYY kg for a liter of ethanol.

[20] Figure downloaded from http://www.vattenfall.com/www/ccc/ccc/577730downl/602590image/index.jsp 11/14/08.

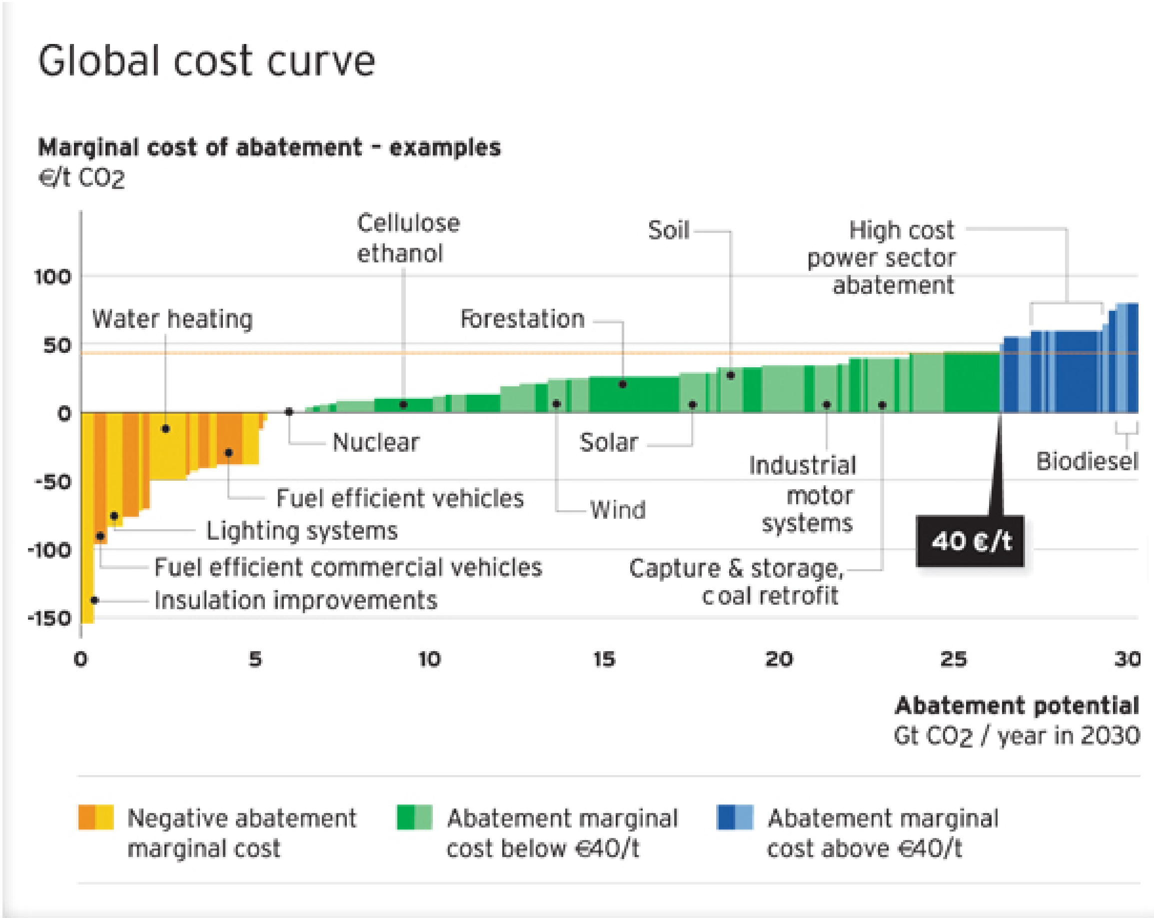

Figure 3.1: Abatement costs for alternative measures to reduce GHG emissions

The technologies with negative abatement costs (to the left of the diagram) would more than pay for themselves in reduced energy costs. Just to the right of those are technologies with low positive costs that can be considered “low-hanging fruit”: cheap abatement measures that should be adopted next. Moving to the right, as further reductions are undertaken, more and more expensive technologies would be required. The Vatenfall report assumes that atmospheric CO2 concentrations could be stabilized at a safe level if emissions were reduced by about 27 gigatons of CO2 per year by the year 2030, relative to business as usual. According to the cost estimates in the figure, doing this would require adopting technologies that would cost on the order of 40 €/t. To put this in perspective, 40 €/t works out to be about the equivalent of 50-60 cents per gallon of gasoline, given its carbon content and depending on exchange rates.

The variation in cost across different industries and technologies shows why it is important that policy lead to the adoption of the lower-cost measures. Compare, for example, the estimated abatement cost associated with cellulose ethanol (<20 €/t) versus biodiesel (>70). The difference of at least 50 €/t implies that a policy promoting biodiesel to reduce emissions by 1 gigaton would exceed the cost of the same reduction using cellulose ethanol by more than $75 billion per year.

If the government had perfect knowledge of the abatement cost schedule for all technologies and all energy users, it could simply mandate the adoption of the lowest-cost methods, up to the desired amount of reduction. Policies mandating specific levels or technologies of pollution control—referred to by economists as "command-and-control" regulations—are quite common, but in practice they seldom achieve least-cost methods. To understand why, first note that Figure 3.1 vastly oversimplifies the structure of abatement costs. For example, it shows solar energy as having uniformly moderate costs. In fact, the cost of solar energy production varies substantially by location: its cost per unit of energy is considerably lower in sunny places than cloudy ones. Similarly, energy conservation measures have widely varying costs across different users. Compare two hypothetical SUV owners, one of whom mostly drives by herself to and from work, while the other frequently transports seven members of the local soccer team. The true opportunity cost of replacing the gas-guzzler with a Prius is likely much greater for the latter than the former. The cost of converting between energy sources for manufacturing firms may depend on the nature of their business and the age of their facility. All of these complications make it extremely difficult to design policies that can identify the low-hanging fruit.

Regulations intended to reduce emissions can sometimes provide perverse incentives. Vehicle fuel-efficiency (CAFÉ) standards, for example, seek to reduce emissions by reducing fuel consumption per mile driven, but provide no incentive for decreasing driving. Indeed, by reducing the per-mile cost of driving, fuel efficiency standards can have the perverse effect of increasing total miles driven, thereby offsetting some of the desired pollution reduction. Similarly, low-carbon fuel standards (LCFS), currently being implemented by executive order in the state of California, provide incentives for the use of relatively low-carbon fuels, but by subsidizing carbon-based biofuels, they promote the use of fuels that still contribute to greenhouse gas emissions, even if their carbon content is lower than that of conventional fossil fuels.

Another practical difficulty is created by the ongoing evolution of technology. Technological changes are constantly reducing the costs of abatement, and at different and unpredictable rates. It would be a risky proposition for policy-makers to try to pick winners among these competing technologies. An efficient command-and-control approach would at the very least have to keep up with and adapt to changes in the relative costs of alternative abatement technologies methods.

For these reasons, command-and-control and technology mandates are likely to be inefficient approaches. Fortunately, there are well-known alternative policies that rely on the economic incentives and the market to pick the lowest-cost abatement methods. The best known of these are the carbon tax and cap-and-trade. These have become familiar terms in the public discussions of climate change. We now examine how these policies work, why they lead to cost-efficient reductions, and how they are similar and different.

Putting a price on CO2

The crux of both the carbon tax and the cap-and-trade policies is to put a price on CO2 emissions. If individuals and business firms have to pay for every unit of CO2 they release into the atmosphere, they have a direct financial incentive to reduce those emissions. The higher the price, the greater is the incentive. Furthermore, they have an incentive to find the methods of abatement that reduce their emissions at the least cost. These policies thus harness selfish motives to the search for low-cost abatement. Under the assumption that individuals and businesses will usually have a better idea of their own costs and opportunities than government would, decentralized decisions are more likely to lead to cost-effective solutions.

A carbon tax simply places a tax on the carbon content of fuels, and thus indirectly taxes the emissions that result from burning those fuels. Because fuel markets are already taxed and regulated, a basic carbon tax would be extremely easy to implement. Americans already pay various gasoline taxes, for example: a tax of $0.50 per gallon of gasoline would be the approximate equivalent of a $50 per ton tax on CO2 emissions.

The cap-and-trade system is perhaps less familiar than a tax. Under cap-and-trade, the government sets a certain limit (cap) on total carbon emissions per year. It then issues permits allowing the holder of a permit to emit a certain amount of carbon, where the total number of permits issued adds up to the cap. Holders of permits may use them to pollute, or they may sell them in a pollution permit market: this is th**e** "trade" component.[21] Like the carbon tax, cap-and-trade puts a price on CO2 emissions, but in this case the price is not set by government but determined by the forces of supply and demand in the permit market. A successful cap-and-trade policy was implemented for sulfur dioxide pollution in the United States during the 1990s to reduce the damage from acid rain (caused by coal-fired power plants). More recently, a cap-and-

[21] As in the case of the carbon tax, the permits would actually directly regulate fuel, rather than emissions.

trade policy for GHG emissions was implemented in Europe as part of the EU's effort to meet the targets of the Kyoto Protocol, and has been advocated by President Obama.

To understand how the market-based approaches work and in what ways they are similar, it helps to use some simple diagrams based on the abatement cost idea from Figure 3.1 above. The curve plotted out by the abatement costs in the figure is called the *marginal cost of abatement* (MCA) curve. In this context the term marginal means *incremental*: the cost of an additional reduction in CO2 emissions is the cost associated with the cheapest means available, given that all lower-cost methods have already been implemented. The MCA rises because once the low-hanging fruit have been plucked, further reductions require more drastic, costly measures.

It will ease the exposition to examine the quantity of CO2 emissions rather than CO2 reductions. This requires flipping the axis of the MCA figure from left to right. Figure 3.2 shows a stylized MCA curve after flipping the axis. "BAU" stands for the level of emissions under business as usual, with no policy to reduce GHG. Moving from right to left, emissions are reduced, which requires the use of increasingly costly abatement technologies. In the figure, I have removed the abatement opportunities that have negative cost (yellow areas in Figure 3.1), on the assumption that these should be adopted by individuals and businesses even without any further government regulation.

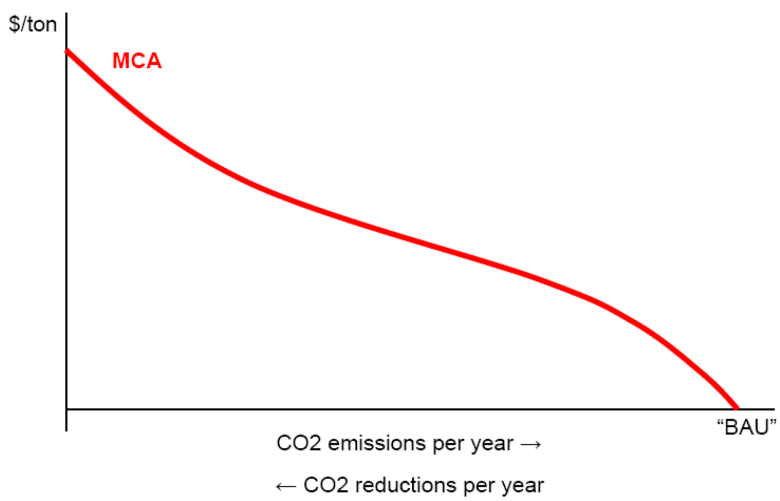

Figure 3.2: The MCA curve

Now suppose that people and businesses had to pay a certain price for the right to release a ton of CO2 into the atmosphere. They would then have to decide whether it was cheaper to pay the price and pollute, or avoid paying the price by paying the cost of reducing emissions.

Polluters with a marginal cost of abatement (MCA) greater than the price would find it cheaper to pollute; those with MCA below the price would be better off cutting their admissions and saving the price of the right to pollute. Consequently, the amount of CO2 emissions would end up being at a quantity where the MCA was exactly equal to the price charged for the right to pollute. This market equilibrium is depicted in Figure 3.3. Everyone to the left of E* would be releasing CO2, while those to the right would have eliminated their emissions using their low-cost technologies. In the jargon of economics, the MCA curve in Figure 3.3 is the *demand curve* for pollution rights.

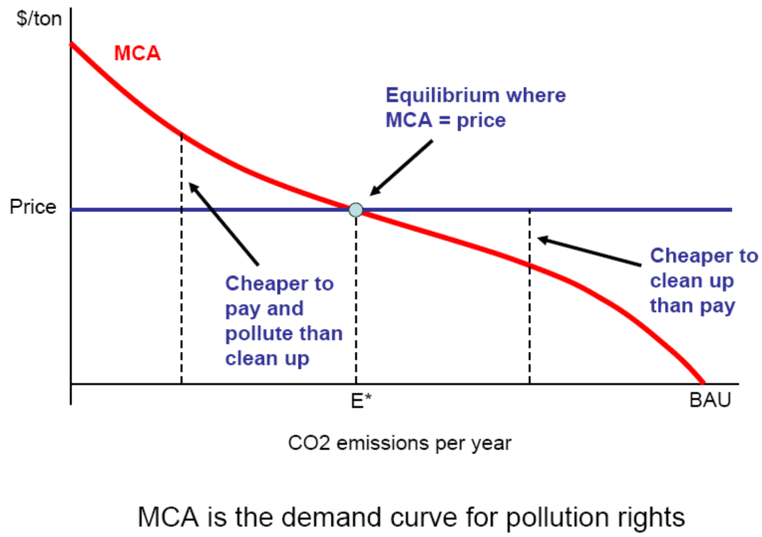

Figure 3.3: The demand curve for pollution rights

The carbon tax can now be understood as a government policy that simply sets a price for the right to pollute a ton of CO2. You are free to pollute if you are willing to pay the tax. Although the government sets the price, not the quantity, it can regulate the quantity indirectly if it knows roughly what the MCA curve looks like. By setting the tax low, it makes pollution more attractive; by raising the tax, it induces more polluters to reduce. This is depicted in Figure 3.4.

A critical feature of the carbon tax is that the GHG reductions are made by the individuals and businesses that have abatement costs below the tax, while those with costs above the tax simply pay the tax rather than reduce. The consequence is that everyone sorts themselves according to their marginal abatement costs; the incentives created by the price mechanism lead to private parties picking the low-hanging fruit without the government having to identify what they are.

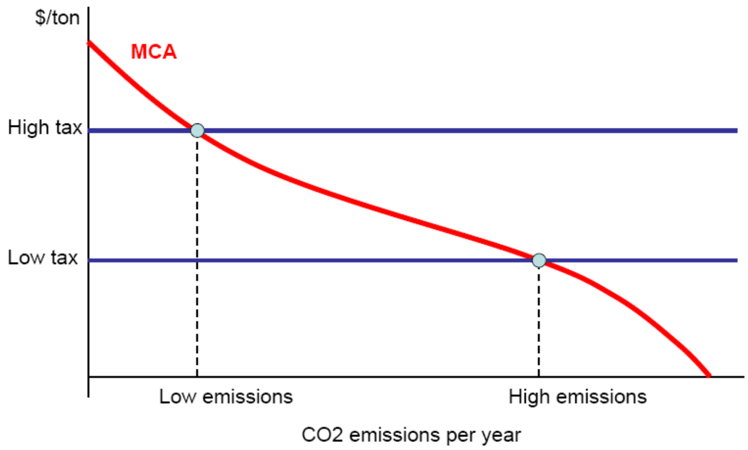

Figure 3.4: Response to changing carbon tax

Cap-and-trade can be shown on a similar diagram. Now the government sets a total amount of CO2 emissions that will be permitted. It then distributes that number of pollution permits. These could be auctioned off, given away, or a mix of both: in the EU carbon market, permits were given away to power plants based on their historical levels of pollution. No matter who gets the permits initially, however, the fact that they can be traded means that many of them will probably change hands. Once a price is established in the market, we know from the carbon tax example that people with MCA greater than the permit price will want to purchase permits, because it is cheaper for them to pay and pollute. Those with MCA lower than the permit price will be willing to sell permits; when they sell, they no longer have the right to pollute and will have to pay the cost of clean-up, but the income they receive from selling the permit more than makes up for the abatement cost. In sum, we have buyers and sellers of permits. The ultimate price is determined by the forces of supply and demand, and that price happens to be where the total number of permits crosses the MCA curve, as shown in Figure 3.5. And just as in the carbon tax case, the resulting carbon price creates an incentive for the low-hanging fruit to be picked first, leading to the most efficient mix of reduction efforts.

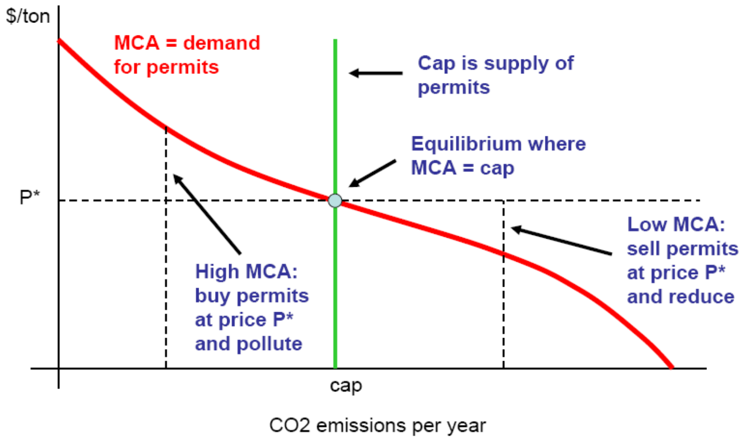

Figure 3.5: Cap-and-trade policy

Which is better: Carbon tax or cap-and-trade?

In the United States, market-based policies for reducing greenhouse gases seem to be tilting toward cap-and-trade rather than the carbon tax. For example, cap-and-trade is the policy of choice under the implementation plans for AB32, the state of California's major climate change law. Similarly, President Barack Obama has stressed the role of cap-and-trade in various statements he has made about climate change, and the major climate change legislation being considered by the U.S. Congress also embraces cap-and-trade. From a naïve political perspective, the greater popularity of cap-and-trade is not difficult to understand: a carbon tax would introduce a new and substantial tax, and taxes are always a hard political sell. But the contrast between the two policies is easily overstated. Indeed, the impact on final energy consumers would be—to a first approximation—identical under a carbon tax and a cap-and-trade program that set the same target for GHG reductions. The reason is easy to see from Fig. 3.5 above. The carbon tax that would achieve the same overall reduction in GHG emissions as the cap indicated by the vertical green line would have to be set at P*, resulting in the same price of carbon under either system. Because the price of carbon dictates the ultimate change in the price consumers pay for, say, electricity or gasoline, two policies that result in the same price for CO2 will lead to exactly the same price at the pump. Furthermore, both policies have the same efficiency properties, each providing the proper market incentive to pick the low-cost reduction measures.

The efficiency advantages of either policy are what really set them apart from other frequently considered policy approaches. Above we noted the mixed incentives for GHG reduction created by fuel-efficiency standards for automobiles. On the one hand, fuel-efficient cars release less CO2 per mile; but on the other, by reducing fuel costs, they encourage more driving. By contrast, a policy that directly prices carbon makes all fossil fuel consumption more costly: it thus leaves it up to consumers to decide how they will cut back on consumption—by buying a higher-mileage car, driving less, or both. What really matters is that their decisions reflect the true cost of their activity in terms of its contribution to global warming. The efficiency differences could be substantial. Recent estimates suggest that a policy that prices carbon could achieve a desired reduction in CO2 emissions at a half or less of the cost of less-efficient policies, such as CAFÉ or LCFS.[22]

In the real world, there are likely to be differences between the effects of a simple carbon tax and an equivalent cap-and-trade program that might make one in principle preferable to the other. But in practice, many of these differences could be reduced or eliminated by modifying the details of the policy. Here are a few of the important issues.

Uncertain marginal costs of abatement. In setting policy, the government cannot observe the true MCA, but must make an educated guess. As we have seen, the market outcomes under either a carbon tax or a cap depend on the location of the marginal cost of abatement (MCA) curve, which is also the demand for pollution rights. Suppose policy-makers believe that the MCA will be the solid red curve in Fig. 3.6 ("Expected MCA"). Then based on expectations, a cap or a tax could be used to achieve the same mitigation target. But suppose in practice that the MCA turns out to be higher—say, the dotted red curve ("Realized MCA"). Then if a tax has been fixed at the horizontal solid blue line, the resulting level of emissions will be where the tax crosses the realized MCA curve, Q*, as indicated by the vertical dotted green line: realized emissions are well above what was expected. On the other hand, if the policy had been a cap fixed at the vertical solid green line, the emissions would be fixed there, but the equilibrium price of permits would rise to the market equilibrium P* (the horizontal dotted blue line), much higher than anticipated.

[22] For recent examples, see Fischer and Newell (2008); Holland et al (2009).

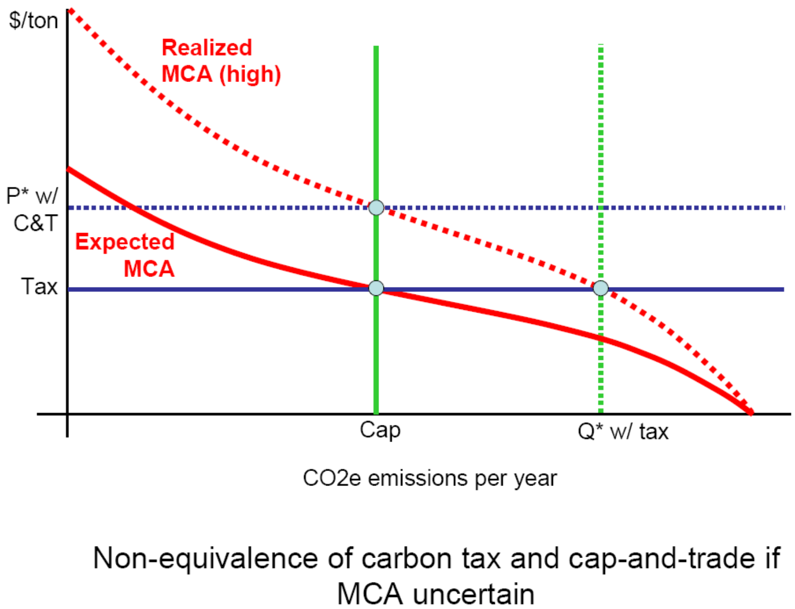

Figure 3.6: Environmental policy when abatement costs are uncertain

Which "mistake" is more costly: ending up with more pollution than expected at Q*, or ending up with a higher price of carbon than expected at P*? This depends on how quickly the costs of abatement rise relative to the costs of climate damages. If there is great concern about avoiding some threshold level of emissions, beyond which there is a serious threat of catastrophic warming, then the cap-and-trade policy offers an assurance that the threshold will not be crossed, no matter what MCA turns out to be. Of course, to the extent that either the cap or the tax could be adjusted over time in response to new information, this source of difference between the policies is less significant.

The double dividend. Advocates of a carbon tax frequently cite its potential for raising government revenue. This revenue could be used to finance increased government spending or to offset other taxes, either immediately or in the future (by reducing the government deficit). Whereas most taxes cause inefficiencies in the economy, a "green" tax such as the carbon tax is an exception, since it is expressly designed to reduce an inefficiency caused by an externality. Thus a carbon tax can potentially increase economic efficiency in two ways: first, by reducing an environmental problem, and second, by reducing inefficient taxes. This so-called "double dividend" reduces the net economic cost of attaining GHG reductions.[23] Revenue enhancement is

[23] Sometimes less appreciated in policy discussions is that by the same logic (turned around), the use of government subsidies to stimulate "green" behavior or technologies comes with a hidden extra cost, in the form of the inefficient taxes elsewhere in the economy that would have to be raised (or not lowered) to fund the subsidies. For a more sophisticated analysis of these issues, see Goulder (1999).

not, however, an advantage unique to the carbon tax: government revenue could also be generated under cap-and-trade by selling the pollution permits. Indeed, by selling all the permits at the equilibrium price P* in Fig. 3.5, the government could raise the same overall revenue that it would from collecting the equivalent carbon tax.

Distributional impact and environmental justice. Economists tend to focus on the efficiency properties of different climate change policies. Efficiency can be thought of, roughly, as relating to the size of the “economic pie”: the gap between total benefits and total costs, regardless of who receives them. By this standard, the better policy is the one that leaves the pie bigger. But as global citizens we are often at least as concerned about *distributional* effects: namely, how the pie is sliced, or the impact of climate policies on different identifiable groups or types of people. These groups can be defined along various lines. For example, one area of concern is the differential impact of GHG regulations on industries that are more or less intensive users or producers of fossil fuels. Beyond the issue of economic hardships that might be imposed on workers and investors in industries that experience large cost increases, there is the possibility that industries particularly hard-hit would mount a strong political opposition to climate policies.

If cap-and-trade becomes an important component of climate policy, the decision of how to allocate the emissions permits will have dramatic distributional implications. To see why, note that the United States currently emits somewhat more than 7 billion metric tons of CO2e per year. Suppose that a cap that cut emissions in half to 3.5 billion resulted in an equilibrium permit price of $50 per ton. Then if the government auctioned the 3.5 billion allowances (permits) and collected their market value of $50 per ton, it would collect some $175 billion per year. On the other hand, if the allowances were given away, perhaps in proportion to firms’ historical emissions levels, the $175 billion would be kept by firms in the private sector. These distributional differences notwithstanding, it should also be stressed that the impact on emissions is identical, however the permits are handed out, this being a function of the cap itself.

Perhaps the greatest concern about the distributional impact of both climate change and climate policy is the potential for adverse impacts on the poor, both within a given country and globally. These issues often come under the heading of *environmental justice*. It is well known that some of the most severe climate effects—especially drought—are expected to occur within the tropical zone, where they would affect some of the world’s poorest regions, such as sub-Saharan Africa. With respect to policy design, some critics of policies that would put a price on carbon argue that such policies are likely to be regressive, their impact falling disproportionately on low-income people. This would tend to be true, for example, to the extent that fuel- or energy-intensive goods and services are economic necessities and thus take a larger share of a poor person’s budget than a rich person’s. The net distributional impact of climate policies depends critically on whether governments undertake measures to offset adverse impacts on the poor, such as rebating some of the carbon tax or permit auction revenues as income tax credits.

Issues of implementation

Assuming that cap-and-trade in some form is likely to become the centerpiece of U.S. climate policy, various details of the implementation will have to be worked out. Of course, the biggest decisions involve the overall target for GHG emission reductions and the speed with which it is to be achieved. The leading climate change legislation under consideration by the U.S. Congress at the time of this writing is the Waxman-Markey bill, which would call for reductions of GHG emissions to 20% below 2005 levels in 2020, 42% below 2005 levels in 2030, and 83% below 2005 levels in 2050. Global reductions of this magnitude are often thought to be necessary to avoid dangerous warming. As we have seen, another key decision is how the allowances under cap-and-trade will be allocated to polluters—given away, auctioned, or some mix. These details have not been determined in the current draft of the Waxman-Markey bill.

Three additional aspects of cap-and-trade systems are likely to be quite important in practice. First there is the question of how to incorporate carbon sequestration under the trading regime. The simplest approach to implementing cap-and-trade in fuel markets is likely to involve regulating the consumption of fossil fuels based on their carbon content, rather than trying to monitor emissions directly. For example, carbon allowances would be allocated at the level of petroleum refiners, rather than individual drivers. But to the extent that carbon dioxide can be captured and sequestered before release into the atmosphere, or in fact removed ex post from the atmosphere (for example by reforestation), cap-and-trade should credit these reductions to the parties responsible. In essence, an additional CO2 allowance could be issued for any activity that caused a proven net reduction of one ton of CO2e in the atmosphere. While in principle cap-and-trade should treat emission and removal of GHG symmetrically, in practice monitoring and measuring sequestration present serious challenges.

Second, policymakers must decide the rules for banking or borrowing GHG allowances over time. Under cap-and-trade, allowances would generally be dated, permitting the release of a ton of CO2 during a specified time period—say, a particular calendar year or two. Banking would allow owners of allowances to hold some portion of them for use at a later time, while borrowing would allow polluters to borrow additional permits against future quotas. Because the global warming problem is largely caused by the total accumulation of GHGs over a long period of time, the main objective is reducing the total emissions, rather than the exact timing of reductions.

Banking and borrowing allow flexibility in the timing of reductions that should in principle permit cost reductions. To see why, suppose technological advances in energy conservation proceeded more rapidly than anticipated; with banking, polluters could adopt conservation measures early and bank their extra permits to be used later, when the cap becomes far more stringent. Without banking, the incentive for early adoption would be smaller. The key to successful implementation of banking and borrowing is confidence in the commitment of regulators to the path of planned emissions targets.

Third, and most challenging, is the problem of developing international agreements that will encourage GHG reductions in both developed and developing economies, the latter already accounting for more than half of global emissions. In principle, an efficient global climate policy would feature a single global price on GHG emissions. The logic is identical to the argument for using a single price within a country. If the price of carbon were different across countries, it would imply that some non-polluters in the country with the higher price would have a higher marginal cost of abatement than some polluters in the country with a lower price. In this case, the same level of global emissions could be achieved at lower total abatement cost if some emissions were shifted from the lower-price country to the higher-price country. This would occur in the market if the price of allowances were equalized across countries, through either a harmonized carbon tax or a global emissions trading scheme. The greatest obstacle to a global trading scheme is bound to be the difficulty of agreeing on the initial allocation of allowances across countries. An allocation based on historical emissions levels would naturally allocate a disproportionate share of permits to developed, energy-intensive economies, while a per-capita allocation would allocate the lion's share to populous less-developed economies. The distributional implications are enormous.

Subsidies for technology development

The economic argument for using a broad market-based scheme to regulate GHG emissions is compelling. Uniform pricing of emissions would place all abatement technologies and conservation measures on a level playing field, with adoption being driven by relative costs. In the simplest case, there would then be no need to subsidize particular technologies, and in fact subsidies would run the risk of encouraging adoption of unnecessarily high-cost solutions. Subsidies for corn-based ethanol production in the United States are a good example of this problem.

Under more realistic assumptions, however, there are at least a couple of justifications for subsidizing some new technology development. First, the private incentive to invest in research and development of new technologies in any field is often too low. The reason is that knowledge spillovers may prevent innovators from capturing the full return on their research investment. In this case a government subsidy for research, especially basic research that may have broad spillovers, is justified. Second, to the extent that international negotiations are unable to reach binding commitments to GHG reductions, advanced countries may find that the best way to promote cleaner energy technology is to push for rapid development of economical alternative energy sources and license them to developing countries at zero or low cost. This is an example of a second-best policy, when the first-best solution (imposing a stringent global cap) is politically infeasible.

4. Summary

Global climate change poses an unprecedented challenge to policymakers and citizens. The scale of the problem is daunting, and the need for action is urgent. Nevertheless, as we have shown, the benefits of preventing unchecked global warming are likely to vastly outweigh the necessary costs of the required measures. Furthermore, the basic policy choices are well understood. A rigorous cap-and-trade scheme, phased in over the next 30-40 years, coupled with technology development and assistance for developing countries, offers the best politically feasible path for meeting the challenge with minimal economic disruption.

References

Anderson, Elizabeth. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press, 1993.

Arrow, Kenneth J. (1995), "Intergenerational Equity and the Rate of Discount in Long-Term Social Investment," IEA World Congress, http://www-econ.stanford.edu/faculty/workp/swp97005.pdf (downloaded 2/16/07).

Arrow, Kenneth J. (2007), "Global Climate Change: A Challenge to Policy," *The Economists' Voice* 4:3, Article 2. Available at: http://www.bepress.com/ev/vol4/iss3/art2

Copeland, B.R. and Taylor, M.S., 2004. "Trade, growth and the environment." *Journal of Economic Literature* **42**, pp. 7–71.

Deaton, Angus, "Letter from America—On transatlantic vices, or Stern in America," *Royal Economic Society Newsletter*, n. 139 (October 2007), pp. 3-4

Fischer, Carolyn, and Richard G. Newell, "Environmental and Technology Policies for Climate Mitigation," Journal *of Environmental Economics and Management* 55 (2008), 142-62.

http://www.gapminder.org/gapminder-world.html

Goulder, Lawrence H., Ian Parry, Roberton Williams III, and Dallas Burtraw. 1999. The Cost-Effectiveness of Alternative Instruments for Environmental Protection in a Second-Best Setting. *Journal of Public Economics* 72 (1999):329-360.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel, "Greenhouse Gas Reductions under Low Carbon Fuel Standards?" *The American Economic Journal: Economic Policy*, 1 (2009), 106-146.

Keohane, Nathaniel, and Peter Goldmark, "What Will It Cost to Protect Ourselves from Global Warming?" Environmental Defense Fund (2008), available at http://www.edf.org/documents/7815_climate_economy.pdf

Neumayer, Eric, "A Missed Opportunity: The Stern Review On Climate Change Fails to Tackle the Issue of Non-Substitutable Loss of Natural Capital (LSE working paper, revised version April 2007).

Nordhaus, William, "Critical Assumptions in the Stern Review on Climate Change," *Science* **317** (5835), (13 July 2007), p. 201.

Nordhaus, William, "The Challenge of Global Warming: Economic Models and Environmental Policy," Yale University, July 24, 2007. p. 20 http://nordhaus.econ.yale.edu/dice_mss_072407_all.pdf

Nordhaus, William, "A Review of the *Stern Review on the Economics of Climate Change,*" *Journal of Economic Literature* 45:3 (September 2007), 686-702.

Paltsev, Sergey, John M. Reilly, Henry D. Jacoby, and Jennifer F. Morris, "The Cost of Climate Policy in the United States," MIT Joint Program on the Science and Policy of Global Change Report No. 173 (April 2009).

Ramsey, F.P., "A Mathematical Theory of Saving," *The Economic Journal* 38 (Dec. 1928), pp. 543-559.

Rawls, John. (1971), *A Theory of Justice* (Oxford: Oxford University Press).

Sidgwick, Henry. *The Methods of Ethics*, 7th edition, Chicago: University of Chicago Press, 1907 (1st edition 1874).

Stern, Nicholas. *Stern Review on the economics of climate change*. 2006 http://www.hm-treasury.gov.uk/independent_reviews/stern_review_economics_climate_change/stern_review_report.cfm

Tol, Richard S.J. "The Marginal Damage Cost of Carbon Dioxide Emissions: An Assessment of the Uncertainties." *Energy Policy* 33 (2005), 2064-74.

Weitzman, Martin L. "A Review of the *Stern Review on the Economics of Climate Change,*" *Journal of Economic Literature* 45:3 (Sept. 2007), 703-724.

Weitzman, Martin L. "On Modeling and Interpreting the Economics of Catastrophic Climate Change" (July 2008), http://www.economics.harvard.edu/faculty/weitzman/files/REStatFINAL.pdf , downloaded 8/11/08.