



CHAPTER I
Crisis and Opportunity

Worldwide, legislatures are beginning the long process of deciding whether to ratify and implement the December 1997 “Kyoto Protocol.” Widely hailed as a first serious step towards slowing greenhouse warming, the protocol requires each industrialized nation to cap its emissions at specific target levels. Those targets apply to the “budget period” of 2008–2012, and the protocol also envisions that nations will agree on caps for future budget periods. Although public pressure to do something about global warming is growing, legislators will weigh the cost of compliance before they ratify the Kyoto deal. One factor will loom large in the debate: whether governments will be able to buy and sell emission credits—a scheme known as emission trading.

Without emission trading, nations would be required to meet their Kyoto obligations entirely within their borders. In the United States, for example, compliance would require a Herculean effort. By the end of 1999, U.S. emissions had risen about 12% above 1990 levels and were on track to rise another 10% by 2008.¹ Yet the Kyoto Protocol requires a 7% cut *below* 1990

levels—in total, about a 30% cut. Turning the economy around to meet the Kyoto target could cost over \$1000 per household per year, which is similar to the annual spending on all federal clean air and water programs combined.² The cost would be high because most emissions of greenhouse gases come from burning fossil fuels for energy, and the economic lifetime of energy equipment like power plants, buildings, and automobiles is long (two decades or more). Compliance with a sharp 30% cut would force the premature disposal of some of the “capital stock” of energy equipment and retard significant parts of the U.S. economy. Electric power generation is especially vulnerable. About half of U.S. electric power is supplied by coal, which is the most greenhouse gas intensive of all fossil fuels. The time to implement easy changes has already passed. About four-fifths of the U.S. generating capacity that will electrify 2010 will already have been built by the end of the year 2000.

With trading, however, nations could lower the cost of capping emissions. For example, a trading system could allow U.S. firms to purchase emission credits overseas, which might be much cheaper than making all the needed emission controls at home. The Kyoto Protocol envisions three interrelated trading systems. One, known formally as emission trading, would allow an industrialized country to increase its emission cap by purchasing part of another industrialized nation’s Kyoto allocation. A second system, known as joint implementation (JI), would allow industrialized countries to earn credits when they jointly implement specific projects that reduce emissions. A third System, known as the Clean Development Mechanism (CDM), allows industrialized nations to earn credits for projects implemented within developing nations.³

Together, these three systems—emission trading, JI, and the CDM—could constitute a full-blown emission trading system that would allow firms to shop the world for the least costly ways to reduce emissions. Nations would have the ultimate responsibility for complying with treaties under international

law, and thus national governments would bear the ultimate responsibility for ensuring that the trading books balance. In practice, however, firms and individuals would probably do most of the trading. The economic appeal of trading is substantial. In the United States, full-blown trading could lower the annual costs by a factor of ten—to a more palatable \$100 per American household.⁴

As figure 1.1 illustrates, nearly all other advanced industrialized countries are in a similar situation. In Japan, emissions have risen more slowly than in the United States because Japanese economic performance for the last decade has been dismal. Nonetheless, Japan is not on track to comply with its Kyoto target (6% cut below 1990 levels) unless it can purchase credits overseas.

In Europe, emissions have actually declined slightly since 1990. Economic collapse and modernization in eastern Germany explain most of the reduction; in addition, energy market reforms in the United Kingdom have caused a shift from carbon-intensive coal towards carbon-light natural gas and zero-carbon nuclear power. Nonetheless, these events have not put the 15 countries of the European Union on track to comply with their Kyoto commitment to cut emissions 8% below the 1990 level during 2008–2012.⁵ That is why, after vigorously decrying emission trading as an American ruse to avoid any serious cuts in greenhouse gas emissions, the European Union (EU) is now developing a plan to implement a credit trading system that would enable it to meet some of its Kyoto obligation by purchasing emission credits overseas.⁶

After long opposing trading as a loophole that would let industrialized countries avoid their duty to slow global warming, a growing number of developing countries are also beginning to embrace trading cautiously. Exempt from Kyoto's regulatory obligations, developing nations have been wary of participating in international efforts to slow global warming because they fear that they, too, will be expected to implement costly controls on emissions of greenhouse gases in the future. Through the CDM,




FIGURE 1.1: Trends in emissions of carbon dioxide (CO₂) from combustion of fossil fuels for the United States, European Union (EU15), and Japan. Although several gases are responsible for global warming, CO₂ from fossil fuels is the most important and best documented. Shown are CO₂ emissions calculated from four semi-independent data sets for consumption of fossil fuels: (a) solid heavy lines, Oak Ridge National Laboratory (Marland et al., 1993; updated at www.cdiac.ornl.gov); (b) light dashed lines, EIA (1999); (c) solid light lines, BP Amoco (2000); and (d) dashed heavy lines, the IIASA/WEC data sets, which are based on International Energy Agency energy balance statistics and reported by Nakićenović et al. (1998). The estimates shown for CO₂ are computed using common heating values and carbon emission coefficients, and thus the variance in the emissions is a consequence of differences in the underlying data for combustion of fossil fuels. Also shown (◆) are data officially reported for the base year (1990) in the “national communications” by the European Union, Japan, and the United States to the Framework Convention on Climate Change. Bars during 2008–2012 illustrate the Kyoto targets, which are calculated from the reported base year statistics. Note that, although the figure shows only CO₂ from fossil fuels, the Kyoto Protocol includes six gases—these other gases account for much smaller shares of total global warming, and reliable time series for these data are not available. The figure also omits “sinks” of CO₂ due to growing forests and changes in land use, since governments have not agreed on accounting rules for including these sinks.

however, they could attract new investment into projects that limit emissions of greenhouse gases, stem other environmental problems such as urban smog associated with dirty fossil fuels, and modernize their energy systems.⁷

Thus a consensus is emerging worldwide that trading is the key to realizing the Kyoto Protocol. For the western industrialized nations, emission trading makes the protocol's targets and timetables appear cost effective and feasible. For developing countries, full-blown emission trading offers the best chance to benefit from a worldwide effort to slow global warming. Trading is the keystone of the architecture adopted in Kyoto.

This monograph explores the political, economic, and technical issues that policy makers must address prior to creating a complete emission trading system. It argues that, when viewed in totality, the hurdles to be cleared are so daunting that a sensible emission trading system is infeasible in the foreseeable future. It also argues that the diplomats who crafted the Kyoto Protocol have painted themselves into a corner. In Kyoto they achieved agreement by setting emission targets that would be politically impossible to implement without an emission trading system; yet they deferred discussion of all the details about how the system would operate. During their first high-level meeting after Kyoto—held in November 1998 in Buenos Aires—diplomats set a hopelessly optimistic timetable for resolving by late 2000 all 152 “elements” left outstanding in Kyoto.⁸ Individually, nearly every element—such as “compliance,” “reporting,” and “independent certification and verification”—is difficult to settle; together, the task is impossible. A longer timetable would make it easier to complete the job, but it would also shorten the time left between completion of the emission trading framework and the date when emission targets must be met.

With the clock ticking towards 2008, and the fate of the Kyoto Protocol hanging in the balance, what should be done? Should political leaders soldier on, ratify the protocol, and hope for the best? Should they retain the targets and trading architecture that

they created in Kyoto but stretch out the timetables to make it easier to comply? Or should they use Kyoto's troubles as an opportunity to construct a different framework for slowing global warming?

Most governments plan to soldier on, but that option has the least to recommend it because it forces countries to select among three dead ends. First, diplomats might make it easier to comply with the Kyoto caps on emissions by creating convenient accounting systems. Notably, the protocol includes language that allows countries to take credit for "sinks" that remove CO₂, the most important greenhouse gas, from the atmosphere. When plants grow they accumulate carbon in their trunks, stems, roots, and leaves—as well as in surrounding soils. Agricultural soils are important sinks. In the United States, for example, starting in about 1910—when tractors made it easier for farmers to plow deeper—intensive tilling has reduced the carbon content of soils. Since the 1950s, farmers have been shifting to "no till" techniques that have helped slow soil erosion while also fortuitously increasing the carbon content of soils. Forests are especially large sinks—forests are growing larger and denser in all the advanced industrialized countries, in part because intensive farming is reducing the need for cropland and some of the abandoned land reverts to forest.⁹

Luck and clever accounting could deliver large credits for these sinks. One data set suggests that the United States could offset about 14% of its current emissions if it were awarded full credit for "land-use change and forestry"—a significant down payment that could amount to nearly half of the required reduction during 2008–2012.¹⁰ The more credit awarded for CO₂ that plants and trees are already absorbing, the easier it is for nations to comply with the Kyoto Protocol targets without actually changing behavior.

But this strategy founders on the lack of widely accepted definitions, methods, and data for counting sinks.¹¹ Even if nations could agree on the necessary procedures, there would still

be enormous potential for cooking the books—only a monitoring program larger and more intrusive than anything ever attempted under international law could settle the inevitable disputes. Moreover, the carbon content of forests and soils varies naturally—decades of monitoring would be needed to be certain that a “sink” was not merely transient and deserved full credit.¹² Yet the commitment periods under international law are typically much shorter, such as the five-year “budget period” of the Kyoto Protocol.

A second dead end is for nations in deficit to earn credits overseas through the CDM. Diplomats still have not been able to agree on the rules that would govern the CDM system, and thus investors are still not sure whether and how they could earn credits through these mechanisms. Yet years of preparation, testing, and learning will be required to build a pipeline of sensible projects. Time has run out for firms and governments to earn large quantities of credits by investing in emission-reducing projects under the CDM.

Emission trading is a third way to ease compliance, but it also leads to a dead end. Governments must solve considerable technical problems that confound trading—which I discuss in the following chapters. But even if they succeed, this scenario poses a significant political problem. Russia and Ukraine are by far the cheapest sources of emission credits—not because the Russians and Ukrainians have had an epiphany about the risks of global warming but rather because their savvy negotiators got an emission target in Kyoto that far exceeds the likely level of emissions. Russia and Ukraine agreed in Kyoto to freeze emissions at 1990 levels, but the collapse of the post-Soviet economy in the early 1990s means that their emissions are already far below that target and unlikely to recover fully by 2008. Selling the windfall to nations in emissions deficit—notably the United States—could earn Russia and Ukraine perhaps \$100 billion.¹³ (About four-fifths of that windfall would flow to Russia.) Since the windfall is free—completely an artifact of the luck and skill of the diplo-

mats in Kyoto rather than the result of any effort to control emissions—these extra credits would squeeze out bona fide efforts to control emissions. That buys paper compliance but no reduction in global warming. No Western legislature will ratify a deal that merely enriches Russia and Ukraine while doing nothing to control emissions and slow global warming.

Nonetheless, the pressure to soldier on and preserve the Kyoto framework is strong. Bureaucratic inertia favors such muddling, since change is threatening to the international and national institutions that are now engaged full time in working on the Kyoto issue. Change would also endanger other cars that have attached themselves to the Kyoto train—for example, energy ministries in many countries have used the Kyoto framework as impetus for rekindling interest in energy policy. Intellectual inertia also favors keeping the Kyoto framework intact—since 1991, remarkably few analysts have examined any alternatives to the “cap and trade” architecture that was codified in the Kyoto Protocol.

These reasons explain why governments are now following the worst strategy—implementing all three of the Kyoto-saving devices simultaneously. Through clever “sink” accounting they hope to make the targets less onerous. Through JI and especially the CDM they hope to earn credits overseas. And through emission trading they aim to reduce the cost of full compliance. But, as we will see, creating an emission trading system requires creating emission permits worth hundreds of billions of dollars. Including difficult-to-measure sinks will undermine confidence in the value of those emission permits and give governments strong financial incentive to cook the books. Including the CDM offers a way to earn credits, but putting the CDM into practice will prove to be very slow and inefficient.

Governments cannot solve these problems unless they reopen the protocol—to tighten the targets for Russia and the other transition countries and to distinguish between fluxes of greenhouse gases that can be monitored accurately and those that are harder to count (notably, CO₂ sinks). Diplomats are loathe to do

that; they know that agreement in Kyoto was possible only because negotiators left in shadow the rules that would govern their imaginary emission trading system. Attempts to clarify and fix these rules will provoke deep disagreements and accelerate Kyoto's collapse.

Unable to keep the Kyoto deal as written, diplomats will thus try the next most attractive option: preserving the framework but stretching out the timetables to make compliance easier. That option merely delays the day of reckoning. The Kyoto framework is based on a fundamentally wrong assumption that it is best to slow global warming by setting strict targets and timetables for regulating the *quantity* of greenhouse gases emitted.

Regulating emission quantities is problematic because emissions are determined by factors such as technological change and economic growth that policy makers are unable to control or anticipate perfectly. If governments had control over all the factors that affect emissions then they could calibrate national behavior perfectly and comply with sensible targets, but in democratic market-based countries public administrators are neither omniscient nor omnipotent. The same logic obliges countries to adopt national trading systems that link with the international system.¹⁴ Because nations cannot be sure of their future emission levels, the only cost-effective way to balance the books is to allow international emission trading. Emission targets beget trading. Imposing strict limits on emission quantities requires a system for trading credits and debits.

The problem with trading is that it requires solving a nearly impossible problem before trading can begin: governments must allocate the emission permits. Because no nation knows its future level of emissions or the cost of controlling emissions, no nation will know how many permits it will need. Diplomats, properly trained to protect national interests, will seek allocations based on a worst-case perspective. They will imagine scenarios where their nation's future emissions and costs of control are much higher than expected. Each will demand a large share

of the total number of permits and feel harmed by the share awarded to other countries. The difficulty of allocating benefits and burdens is hardly new to international politics; allocation will confound any collective effort to slow global warming. But emission trading makes solving the allocation problem much harder—chapter 2 explores the three reasons why.

First, emission trading magnifies the stakes. Emission permits are semipermanent property rights. In any well-functioning market, property rights are much more valuable than the annual flow of payments based on those assets. For example, it costs more to buy a house than to rent it for a year. High stakes will make diplomats wary, which will cause them to be especially cautious in the first allocation because that will set the framework for subsequent adjustments and reallocations. That is why diplomats spent years in the Law of the Sea negotiations haggling over how to allocate deep seabed mining rights even though deep seabed mining was largely untested and mining rights had only hypothetical future value. By comparison, imagine how difficult it will be to allocate greenhouse gas emission permits of immediate value that are certain to be worth hundreds of billions of dollars.

Second, a key obstacle is gaining the consent of firms and governments that must pay the cost of acquiring permits. In earlier trading systems—such as the U.S. system for controlling acid rain, or New Zealand's system for tradable fishing quotas—policy makers eased this problem by “grandfathering” permits. They blunted political opposition by giving valuable property rights to the same entities that they most expected to oppose the scheme. But this strategy does not work for an international emission trading system. The trading system must be able to accommodate new entrants—developing countries—who will demand allocations that are quite different from the status quo. Emissions from developing countries are rising much more rapidly than those from the industrialized world, and governments in the developing world think that their future economic growth

demands much higher emissions still. For the developing countries, grandfathering is unacceptable.

Third, the economic efficiency of trading depends on the integrity of the emission permits, which are a form of property right. Yet it is extremely difficult to secure property rights under international law. To date, all significant experience with emission trading is within nations where the state is strong and able to impose the rule of law that is necessary to secure property rights. In contrast, international law has no central authority that can compel countries to remain part of a treaty. The high value of emission permits increases the likelihood that countries will attempt to defect, and the need for security of property rights increases the consequences of defection.

The typical pattern of international diplomacy magnifies the difficulty. Normally, diplomats craft international treaties and then must wait several years as their countries ratify the deal and bring it into legal force. Costly treaties that involve many countries typically require the longest time to ratify. That is problematic because the allocation is based on information about future emission levels and abatement costs that is imperfect. As time elapses, those factors also change; as the date for starting the system approaches, countries whose emissions are far above their allocation will seek reallocation or withdraw from the treaty. But new allocations will affect the costs for all others and unravel the agreement. Recent trends illustrate the problem. Projections of U.S. emissions made since the 1997 Kyoto conference have been markedly higher than the earlier projections that formed a basis for negotiations in Kyoto. The U.S. economy has been more robust than analysts expected on Kyoto's eve. In contrast, emissions from Russia and Ukraine remain lower than most experts anticipated, and thus their windfall allocation of emission permits is even larger.

I call these three the "cold start" problems. Trading requires first the creation and allocation of property rights. However, the constant threat of defection makes it extraordinarily difficult to

allocate and secure property rights. Yet secure property rights are the cornerstone of emission trading. Chapter 2 explores ways to solve the “cold start” problems, but I find no solution for the key problem: the weakness of international law.

Why did diplomats venture into this swamp by creating an architecture based on targets and timetables for emission quantities? At the time that diplomats were framing the Kyoto Protocol and its parent agreement—the 1992 Framework Convention on Climate Change—the working model for international environmental treaties was the 1987 Montreal Protocol on Substances that Deplete the Ozone Layer. The Montreal Protocol set targets and timetables for regulating consumption of ozone-depleting substances, and nearly all nations have since complied. The Montreal Protocol was widely seen as the most effective treaty in the history of international environmental diplomacy, and rightly so—it is why ozone-destroying substances such as chlorofluorocarbons are being phased out worldwide and why the thinning ozone layer is now poised to heal.

Barely a year after diplomats signed the Montreal Protocol they gathered again, in Toronto, for the first major international conference on political strategies for slowing global warming. The Toronto conference ended with the call for nations to cut CO₂ emissions 20% below 1990 levels by 2005. No major nation had a plan for how it would reach the Toronto target—and nearly all will fail to achieve that goal—but that did not slow subsequent efforts to set even more targets and timetables. In 1991, when negotiations on a formal global warming treaty began, many countries and pressure groups made the adoption of binding targets and timetables their central goal. They failed, and the 1992 Framework Convention on Climate Change included no clear targets and timetables. Advocates pushed the same agenda, with success, at the first meeting of the convention’s supreme decision-making body—the Conference of the Parties (COP-1), held in Berlin in 1995. The result was the “Berlin Mandate,” which gave the legal marching orders for the negotiating process

that led to the 1997 Kyoto Protocol. The Berlin Mandate specifically required that the Kyoto Protocol should set “quantified limitation and reduction objectives within specified time frames.” In plain language: the Kyoto Protocol should set targets and timetables for emission quantities.¹⁵

The architects of the Montreal Protocol never had to confront the central problems of trading. They established only an extremely limited emission trading system. The treaty expected that advanced industrialized nations would eliminate ozone-depleting substances on their own, rather than earning credits for overseas efforts. Initially, the goal was to cut consumption by half. With the ink on the Montreal agreement barely dry, incontrovertible evidence that these substances caused the ozone “hole” over Antarctica (and a lesser thinning of ozone worldwide) forced diplomats to tighten the goal to a complete phase-out of all major ozone-depleting substances. Compared with global warming, agreeing on the effort to control ozone-depleting substances was easier because the economic stakes were much lower. Producers and users of ozone-depleting substances soon found substitutes for nearly all applications of the most harmful compounds.

With the Montreal Protocol, diplomats matched the architecture of the treaty with the environmental ill they were trying to solve. Caps on emissions made sense because there was a widely agreed goal of avoiding a relatively clear and dangerous threshold. Scientists demonstrated that even tiny concentrations of chlorine and bromine in the stratosphere would trigger the ozone “hole.” Once the advanced industrialized nations agreed that the “hole” must be healed it was clear that essentially all uses of ozone-depleting substances must be eliminated. And once elimination was the collective goal there could be no debate over allocation—each nation individually had to achieve a phaseout. They wrangled over the timing and over how to phase out some minor ozone-depleting substances, but the central goal forced focus. Developing countries were more skeptical and op-

posed costly requirements to eliminate ozone-depleting substances. These countries are concentrated nearer the tropics where ozone depletion is less severe, and their governments were under much less political pressure to act—development, rather than costly environmental controls with distant benefits, was their aspiration. Once the advanced industrialized nations created a fund to compensate developing countries for the extra cost of complying with the Montreal Protocol phaseout and threatened trade sanctions against any country that did not participate, the developing nations shifted. Today, almost all are on track to eliminate nearly all their consumption of ozone-depleting substances.

Finally, the architects of the Montreal Protocol paid close attention to technical feasibility and economic costs. They created an “escape clause” that allowed countries to exempt important uses of ozone-depleting substances from regulation, which countries have often invoked in cases where it has been too costly to find substitutes. (Regular technical reviews kept countries from abusing the escape clause.) This provision in the protocol made it easier to allocate strict emission targets because it allowed countries to avoid extremely onerous commitments. Political support for protecting the ozone layer would have suffered badly if asthmatics had been forced to abandon medicines in metered-dose inhalers (MDIs). MDIs account for a tiny fraction of ozone-depleting substances, but finding substitutes for their chlorofluorocarbon propellant has proved much trickier than for most other uses of ozone-depleting substances.¹⁶

The Montreal Protocol loomed large when diplomats sought to build a regulatory regime to slow global warming. But they gave inadequate attention to whether the right lessons had been learned from the Montreal Protocol experience, and to whether the lessons were relevant for the global warming problem. Global warming diplomats should have taken more seriously the problem of allocating emission permits, which did not confound the Montreal process as much as it will when hundreds of bil-

lions of dollars of tradable assets are at stake. And they should have paid closer attention to the obscure but vitally important escape provisions of the Montreal Protocol, which made it easier to contain compliance costs. As we will see, provisions that make it possible to contain costs also make it easier to allocate emission permits.

Assuming that, somehow, diplomats might solve the allocation problem, chapter 3 examines other functions that are also necessary for an effective emission trading system—in particular, monitoring of compliance and enforcement.

Kyoto's architects gave little attention to the crucial role of monitoring. The protocol's targets apply to a basket of six greenhouse gases—carbon dioxide, methane, nitrous oxide, and other gases. The problem is that it is difficult to monitor emissions of most of these gases because the activities that cause the emissions are not well understood. The exception is carbon dioxide emitted during combustion of fossil fuels—that flux is extremely well measured and (luckily) also accounts for most of the increase in greenhouse warming.

All schemes to slow global warming must contend with monitoring problems, but they pose special challenges for emission trading. If the fluxes of some gases can't be measured accurately then permits can't be assigned reliably. The security of the underlying property rights erodes, and with this the efficiency of the trading system declines. A simple and effective solution to this problem would involve restricting an emission trading system to fossil fuel emissions of carbon dioxide—at least initially, until the monitoring problems for the other gases are fixed. Diplomats have resisted that because they erroneously think that the only way to address the entire global warming problem and to build an emission trading system is to lump all the gases together into a single system.

Enforcement is a perennial problem of international law, but emission trading potentially offers an elegant solution. If buyers were held responsible for the integrity of the permits they own

then the market would price permits according to their origin and risk of default. Since most permits would be used in advanced industrialized nations, where the rule of law is strong and legal institutions are efficient, buyers would be held accountable through their national legal systems.

It is odd, then, that a consensus is developing in favor of “seller liability,” which would hold Russia, Ukraine, and other major sellers liable for their own compliance—letting buyers off the hook the moment after the permit changes hands. This strange scheme would give sellers a strong incentive to flood the market with bogus permits, knowing that international institutions rarely muster the swift and painful enforcement that would be needed to avert the practice. Worse, if a penalty were imminent the seller could withdraw from the Protocol, pocketing the sale proceeds and leaving the bogus permits in circulation. A strong compliance mechanism could avert that outcome, but there is no precedent for such a mechanism in international environmental law. Moreover, the Kyoto Protocol prohibits the parties from adopting a compliance mechanism that imposes “binding consequences” unless governments formally amend the Protocol. Because amending the Protocol would reopen and unravel the Kyoto deal, diplomats have been doubly wary of crafting a compliance mechanism that is adequate to the task. Seller liability is like an autoimmune disorder; it creates incentives that tempt parties to undermine the trading system, and once over-selling begins the unraveling accelerates.

Should diplomats redouble their efforts to find solutions to trading troubles, or are better alternatives available? Chapter 4 explores these questions by comparing four major options for the architecture of a global warming treaty. It argues that the Kyoto approach of capping emissions at particular quantities makes sense only if the objective of international efforts to slow global warming is to avert a catastrophe that would be triggered by a certain accumulation of emissions in the atmosphere. Governments would identify the dangerous threshold, cap emissions

below the level, and allow trading so that firms could meet the cap at the lowest cost. Diplomats envisioned exactly this approach when they created the Framework Convention on Climate Change. In Article 2 they defined the central objective of international cooperation on climate change as to stabilize concentrations of greenhouse gases “at a level that would prevent dangerous anthropogenic interference with the climate system.” The approach appears to be elegant and sensible but is unworkable. It is not (yet) possible to identify particular thresholds that would trigger horrible climate changes. Worse, if governments set short-term emission caps too tightly they may force their economies to bear extremely high costs of cutting emissions more rapidly than can be achieved with the orderly turnover of the capital stock.

In theory, a better approach would focus on coordinating emission taxes. Governments would implement taxes that begin at a low level and rise over time. An international agreement would set the tax levels and a schedule for adjusting them. By controlling the price of emissions, this approach makes it easier for firms to anticipate the cost of emission controls and to plan long-term investments. It does not require that diplomats invent a hypothetical cap on the exact quantity of emissions. Because it makes it easier to contain costs, a tax system is economically more intelligent than a cap and trade system.¹⁷ Moreover, the tax approach may make it easier for governments to solve the thorny allocation problem because it does not require allocating and securing property rights in the form of emission permits. In practice, however, a tax system is extremely difficult to monitor and enforce. Governments would implement greenhouse gas taxes on top of existing distortions in their tax systems, making it hard to measure the practical effect of the new taxes. In principle, they could create an international regulatory body that would conduct inspections and run economic models to assess tax policy; in practice, such an institution would be much more intrusive and powerful than most governments are likely to tolerate.

Chapter 4 concludes that the best architecture for a global warming treaty is a hybrid of the trading and tax systems. Governments would set targets for emission quantities as well as targets for emission prices. Having limited quantities, governments would create an emission trading system. Unlike a textbook emission trading system, however, governments would also commit to selling additional permits at the target price. The cost of permits would therefore never rise above the target price.¹⁸

In contrast with a textbook emission trading system, the hybrid approach would make it much easier for governments to anticipate the cost of compliance. The hybrid system eliminates the possibility that compliance costs would be much higher than expected. By containing costs, the hybrid architecture greatly eases the “cold start” problem of emission trading. In a textbook emission trading system, governments will be risk averse because they fear the possibility that emissions will be higher, and abatement more costly, than expected. The hybrid system eliminates this worst-case scenario and makes it easier to agree on an initial allocation of emission permits.

Greater control over costs would also make it easier for governments to limit the financial flows that could occur when the system is switched on. If governments are more confident about the cost of abatement it will be easier for them to allocate permits according to marginal cost. Financial flows—such as the windfalls that would flow to Russia and Ukraine under the Kyoto scheme—arise when marginal costs differ markedly. When costs are contained it will be politically easier for governments to resist demands for extra headroom.

In a textbook emission trading system the number of permits remains fixed, and the demand for permits governs the price. In a textbook tax system the tax level governs the price, and the quantity of emissions varies. In a hybrid system both quantities and prices can vary. Demand for permits controls the price up to the target price; above that level, governments print new permits and the price is constant.

The target price would perform a function similar to the “escape clause” in the Montreal Protocol—if compliance proved too onerous the regulatory system would shift the goalposts to a tolerable position. In contrast, without an “escape clause” governments could be forced to bear unplanned burdens or, more likely, to tear down the goalposts by withdrawing from the treaty.

A hybrid approach would also be much easier to monitor and enforce. Unlike an emission tax, it would not require governments to implement extra taxes on top of existing distortions in their tax systems. Rather, the market would govern the price of emission permits, making it easy to determine whether governments are selling additional permits at the agreed price level. Requirements to mark the origin of every permit—as in government debt markets today—would make it easy to spot a government that floods the market with below-cost permits. As with a textbook emission trading system, buyer liability would be the best scheme for enforcement.¹⁹

Although this monograph is about the architecture for international agreements for controlling emissions, mainly emissions of carbon dioxide, I am mindful of the other dimensions to the global warming problem. In addition to whatever system is adopted for limiting emissions of carbon dioxide, four other types of policy are needed. First, governments must invest in knowledge. There is widespread agreement on the need to fund research and monitoring on the global warming problem itself. Also necessary are investments in long-term basic research in fields that are likely to make it easier to invent and apply new emission control technologies—physics, material sciences, nuclear engineering, and the like. Because basic knowledge is a public good—easily transmitted and difficult to appropriate—even if governments agree to control emissions, proper investments in basic research will not automatically follow. There is mounting evidence that advanced industrialized countries are already underinvesting in basic energy-related sciences.²⁰ Yet

over the last decade spending trends on such basic research in most advanced industrialized countries have been negative,²¹ despite articulate plans for how government can help reverse this ebbing tide.²² The challenges are large, especially as a great effort to increase spending should be part of an international knowledge strategy—because basic knowledge is an *international* public good—and not only set according to national priorities and institutions.²³

Second, governments must also invest in adaptation. Many effects of global warming—such as flooding from storm surges and higher sea levels—are unavoidable. Societies must anticipate and prepare to adjust to those effects. Adaptation policies make sense even without the fear of global warming. Most of the anticipated effects of global warming on humans are already within the realm of humanity's experience with nature—even without global warming, buildings flood, crops wither, and nature dominates the outdoors. The same policies that soften the blows of nature also ease adaptation to the effects of global warming.

Third, governments should make some investment in “geoengineering”—the ability to make large-scale interventions in the climate system to slow or reverse climate change.²⁴ Mirrors in space, for example, could reflect sunlight and cool the planet—unfurled at the same pace that greenhouse gases accumulate they could keep Earth's thermostat level. Critics have rightly worried that geoengineering gone awry could do more harm than good. Technological interventions often have unanticipated consequences, and vigilance is needed.²⁵ But equally sobering is that greenhouse warming could trigger nasty surprises in the climate system, and if we detect one of those surprises then geoengineering will be the only option for quick intervention. It is not palatable, but advance preparation through research can reduce the dangers.

Fourth, governments need to clarify the objective of their efforts to control emissions of greenhouse gases. Goals are needed to focus the effort—even draft goals that will require revision

but can focus debate during the interim. Yet the architects of the Kyoto system made two decisions that have deflated attention to proper goal setting. One decision derives from the conventional wisdom that treaties are the most effective instruments of international law. The consequence is that essentially all serious diplomatic discussions of goal setting have occurred within the context of negotiating two legally binding treaties—the Framework Convention on Climate Change and the Kyoto Protocol. Yet binding treaty negotiations are constantly shadowed by worries about compliance—a terrible atmosphere for debating uncertain and distant goals with uncertain economic consequences. More productive goal setting occurs in nonbinding frameworks where diplomats are less narrowly concerned with compliance and more likely to focus first on goals that make ecological and economic sense.

The other disservice of the Kyoto process was lumping all greenhouse gases into a single “basket” and treating them as freely interchangeable commodities. In reality, the gases have different lifetimes and merit distinct approaches. Methane is a strong greenhouse gas but lives only a relatively short time in the atmosphere (ten years). If our goal is to avoid climate changes in the next few decades then controlling methane is a quick way to get results,²⁶ but methane controls today are largely irrelevant for global warming problems that extend past 2020 or 2030. At the other extreme are sulfur hexafluoride (SF₆) and perfluorocarbons (PFCs). These greenhouse gases are thousands of times stronger in effect than CO₂ and linger in the atmosphere for thousands of years. Because of this long-term liability, and because firms can eliminate nearly all emissions of these gases at relatively low cost, governments should adopt policies to curtail these gases on a separate (and more stringent) timetable from the others.²⁷ Most climate policy is appropriately focused on CO₂, which causes most global warming that current and near future generations will experience, but it must not ignore the long-term liability of industrial society. The one-basket approach obscures these important distinctions, making it easier for most

CHAPTER ONE

policy makers to pretend that the time horizon of climatic effects does not matter.²⁸

Chapter 5 recapitulates the message. When the Kyoto Protocol fails, policy makers must ensure that they and the public learn the right lessons. Analysts are pinning Kyoto's imminent demise on the wrong factors—on fleeting political will, on the expectation that Kyoto's costs far outweigh its environmental benefits, and on the fear that Kyoto will create strong and intrusive international institutions that will harm national democracies and freedoms.²⁹ This monograph argues that, while these factors are important, the demise of the Kyoto Protocol is largely the consequence of its very architecture.

The danger is not that the Kyoto Protocol will collapse. Rather, it is that governments will not reckon with Kyoto's real problems—that they will try to muddle through by stretching out the timetables rather than rethinking objectives and strategy. The governments that crafted and signed the Kyoto Protocol, and the nongovernmental organizations (NGOs) that have encouraged them, feel—like mother and child—that to walk away would mean betrayal. But separation is the first step to real action.