MIT Joint Program on the Science and Policy of Global Change

Toward a Useful Architecture for Climate Change Negotiations

Henry D. Jacoby, Richard Schmalensee and Ian Sue Wing

Report No. 49
May 1999

The MIT Joint Program on the Science and Policy of Global Change is an organization for research, independent policy analysis, and public education in global environmental change. It seeks to provide leadership in understanding scientific, economic, and ecological aspects of this difficult issue, and combining them into policy assessments that serve the needs of ongoing national and international discussions. To this end, the Program brings together an interdisciplinary group from two established research centers at MIT: the Center for Global Change Science (CGCS) and the Center for Energy and Environmental Policy Research (CEEPR). These two centers bridge many key areas of the needed intellectual work, and additional essential areas are covered by other MIT departments, by collaboration with the Ecosystems Center of the Marine Biology Laboratory (MBL) at Woods Hole, and by short- and long-term visitors to the Program. The Program involves sponsorship and active participation by industry, government, and non-profit organizations.

To inform processes of policy development and implementation, climate change research needs to focus on improving the prediction of those variables that are most relevant to economic, social, and environmental effects. In turn, the greenhouse gas and atmospheric aerosol assumptions underlying climate analysis need to be related to the economic, technological, and political forces that drive emissions, and to the results of international agreements and mitigation. Further, assessments of possible societal and ecosystem impacts, and analysis of mitigation strategies, need to be based on realistic evaluation of the uncertainties of climate science.

This report is one of a series intended to communicate research results and improve public understanding of climate issues, thereby contributing to informed debate about the climate issue, the uncertainties, and the economic and social implications of policy alternatives. Titles in the Report Series to date are listed on the inside back cover.

Henry D. Jacoby and Ronald G. Prinn,
Program Co-Directors

For more information, contact the Program office:

MIT Joint Program on the Science and Policy of Global Change

Postal Address: 77 Massachusetts Avenue
MIT E40-271
Cambridge, MA 02139-4307 (USA)

Location: One Amherst Street, Cambridge
Building E40, Room 271
Massachusetts Institute of Technology

Access: Telephone: (617) 253-7492
Fax: (617) 253-9845
E-mail: *globalchange@mit.edu*
Web site: *http://web.mit.edu/globalchange/www/*

Printed on recycled paper

Toward a Useful Architecture for Climate Change Negotiations

Henry D. Jacoby, Richard Schmalensee and Ian Sue Wing[1]

Abstract

Years of hard bargaining have failed to produce a policy architecture to adequately address the complexities of climate change. Very likely, such a structure will have to be sought though improvement of the partial architecture developed to date within the Framework Convention on Climate Change. We identify key architectural features that have emerged in the Convention process, and then explore extensions that will be necessary if the current approach is to serve for the long term. An important task is to break the deadlock over accession of developing countries. To this end we propose further incorporation in the negotiations of concepts of burden sharing according to ability to pay that already seem to be embedded in the Convention. The implications of alternative versions of such an approach are illustrated with a set of simple model simulations.

Contents

1. Introduction

We find it hard to imagine a policy problem more daunting than climate change. Except in the unlikely event that greenhouse emissions are somehow proved harmless, the world faces an unending series of contentious negotiations over their control. To be successful, this process must survive over many decades in the face of changing economic conditions, evolution in scientific knowledge, and both the success and failure of prior negotiations and agreements. The prospects for achieving a sound, workable response over time will be increased if these negotiations are governed by a useful policy architecture: a unifying structure that restricts potential agreements in ways that both simplify negotiations and point them in desirable directions.

---

[1] Massachusetts Institute of Technology Joint Program on the Science and Policy of Global Change. An earlier version of this paper appeared in *Economic Modeling of Climate Change*, OECD Workshop Report, 17-18 September 1998, Paris.

Unfortunately, despite years of hard bargaining over emissions targets and nascent flexibility mechanisms, a useful policy architecture has not yet evolved from the climate negotiations. This failure is potentially more troublesome than the growing likelihood that Kyoto's emissions reduction targets will not be met,[2] because it may cast a longer shadow. There is no comparable alternative framework under discussion in policy circles, and so it seems probable that key elements of the Kyoto Protocol (United Nations, 1997a–for instance, its reliance on national emissions targets) will affect climate change negotiations long after its specific reductions have either been met or abandoned.

We of course cannot be certain that subsequent climate change negotiations will take the basic Kyoto framework as a point of departure, and we do not mean to argue that it dominates other approaches. There is, for example, a strong economic case for preferring a regime based on emissions taxes to one based on emissions limits, tradable or not.[3] But if the basic Kyoto framework survives, several important issues will need to be resolved before a useful architecture can emerge. It is accordingly useful to focus on improvements within the Kyoto framework, as we do here, as well as to seek superior architectures that may require abandoning the current approach and returning the negotiations to square one.

The accession of Non-Annex I parties has been kept off the negotiating agenda, for instance, and there is neither a commonly accepted model of what accession and subsequent obligation would look like, nor even a shared set of principles to guide discussions of the issue. If a Non-Annex I country were to join the existing regime of Kyoto emissions targets, what initial commitment would be appropriate and how would it relate to those of existing Annex I parties?[4] How should obligations evolve over time and in response to economic growth? Finally, and of critical long-run importance, how can one judge the consistency of particular choices about accession and subsequent commitment with the Framework Convention's fundamental goal of stabilizing atmospheric greenhouse gas concentrations?

Our exploration of these issues begins in Section 2, with an elucidation of the notion of a useful policy architecture, and then turns in Section 3 to a description of key elements of the Kyoto framework that we believe are compatible with this concept. Section 4 presents a class of potentially useful architectures that result from augmenting the Kyoto framework to address issues it currently ignores, these ideas are illustrated in Section 5 with a simple example. We close with some conclusions that follow from this exploration of the future of climate negotiations.

We preface our arguments with a disclaimer. It is not our contention that the example we present below, in which ability to pay is the sole determinant of incremental emissions control obligations, is optimal in any meaningful sense. Such a highly simplified, "cartoon" architecture is almost certainly too inflexible to serve as the focus of productive international negotiations. Our objective is more modest: to portray our conception of what a candidate useful architecture might look like, and how it might operate. In other words, this paper elucidates how such a structure might determine the decisions that must be made today, and how they in turn will

---

[2] For example, see the assessment by Portney (1999).

[3] See, for instance, Pizer (1997) and Schmalensee (1998b).

[4] The Kyoto Protocol defines an Annex B list of countries, which for purposes of this discussion can be treated as the same as Annex I in the original Convention.

determine the policy architecture that will shape future climate negotiations. On this basis we argue that the longer-run implications of decisions currently being taken deserve explicit attention, and serious thought should be given to the climate policy architecture that we will leave to our children.[5] If they need to reduce global emissions drastically, it will be tragic if we have not reached international agreement on an architecture for negotiations that eases their task. We hope this essay helps begin the serious intellectual work that must underlie any sound architectural agreement.

2. What Are Useful Policy Architectures?

The meaning we intend to convey by the term "architecture" follows Schmalensee (1998a), and it is close to its meaning in the software industry. Cusumano and Selby (1995, pp. 236-7) describe its meaning within Microsoft:

> A product architecture is the skeleton that is internal to a product; it defines the major structural components and how they fit together. A product architecture and the components that fit into it provide the backbone for implementing the product features (that is, doing the detailed design and coding). … The product architecture is also the skeleton that determines the long-run structural integrity of a product. Any evolutionary change in a product's functionality should not cause the underlying product architecture to unravel.

One can thus think of a software product's architects as making strategic decisions about system structure and interface specifications; programmers are left to work within these strictures. The decision to use the same code in Windows 98 to display both web pages and information about the computer's hard disk was architectural, for instance. As the passage above stresses, a good architecture permits considerable change in a product's functionality without compromising its structural integrity. On the other hand, if the desired change in functionality is fundamental enough, the product architecture must be modified or abandoned.

The term also is applied to building construction, of course. Definitions in that context that are close to what is intended here are (from Webster's Dictionary) "a unifying or coherent form or structure" and "a method or style of building." The conventions of gothic architecture, for instance, both limited the decisions to be made by church designers and steered that design process in fruitful directions. They remained in force for centuries, and proved flexible enough to permit the construction of a variety of distinctive buildings.

The creation of individual laws, treaties, protocols, or regulations is shaped by a *policy* architecture: a set of governmental structures and associated principles that constrain what can and must be decided. The fundamental policy architecture of the United States is set by the U. S. Constitution, as interpreted by the Supreme Court. Laws are drafted and debated with the understanding that, for instance, states cannot enter into international treaties, and *ex post facto* laws are forbidden. Such understandings are altered by constitutional amendments only on those infrequent occasions when it becomes widely accepted that they are pointing the society in undesirable directions.

[5] For a general discussion of the long-run aspects of the climate issue and their short-run implications for policy, see Jacoby, Prinn and Schmalensee (1998).

In many contexts important architectural elements also emerge from past policy debates, legislation, and regulatory decisions. Debates about the personal income tax, for instance, have long taken as given the desirability of such a tax, as well as of some degree of progressivity and the legitimacy of a distinction between deductible and other expenses. While elements of this sort, like constitutional provisions, may be disputed from time to time, they remain widely accepted and thus serve to simplify decision-making.

It is important to understand that a policy architecture doesn't have to be "best" in any sense—or even correct—in order to be useful. Often it may be impossible even to rank alternatives, let alone to devise one that is optimal. For example, it is debatable whether it is on the whole better to build tax policy around a personal income tax, as in the U.S., or using a value-added tax, as in the E.U. Furthermore, sound policy can be grounded on incorrect propositions, so long as they do not rule out desirable policy outcomes. For example, an element of the architecture guiding negotiation of the acid rain provision of the 1990 Clean Air Act Amendments was the implicit assumption that the environmental effects of a ton of sulfur dioxide were the same, no matter where in the continental U.S. it was emitted. This assumption is clearly false, but it helped simplify the negotiations, and it does not seem to have led the policy process seriously astray in substantive terms (Schmalensee *et al.*, 1998).

On the other hand, architectures can be plainly counter-productive. For instance, the Clean Air Act Amendments of 1970 required that national ambient air quality standards for seven criteria air pollutants be set at levels "requisite to protect human health" (U.S.C. 42 §7409(b)(1)). This architectural element, which has shaped almost 30 years of standard-setting, rests on the assumption that there are threshold levels of concentration below which each pollutant is safe and above which it is dangerous. In a world without such "bright line" thresholds, this architecture presumes information that doesn't exist, and in the process retards the search for understanding of the effects of alternative stringency levels. Getting the architecture wrong can thus have long-lived procedural and substantive costs.

With this definition in mind, we now turn to the Framework Convention on Climate Change or FCCC (United Nations, 1992)—the basic document that has provided a partial framework for international climate negotiations.

3. The Legacy of the Rio-Berlin-Kyoto-Buenos Aires Process

Like almost all international treaties, the FCCC, signed in Rio de Janeiro in 1992, is an agreement to which parties only adhere voluntarily. Placed as it is within the U.N. system, each step in the development of the Convention nominally involves almost all the world's nations, with the views of a dozen or so key states being crucial to any collective outcome. Much effort has already gone into this process, as it has evolved from the pre-Rio negotiations, through the first Conference of Parties to the Convention (COP-1) in Berlin, and to COP-3 in Kyoto and COP-4 in Buenos Aires. As a consequence, we believe it is highly likely, whatever the ultimate fate of the Kyoto Protocol, that subsequent stages will involve extension and/or modification of the structure created during these years of activity.

Taking that belief as the underpinning of this essay, we identify five key features of the structure that has emerged from the FCCC negotiations:

• Negotiation of near-term emissions limits,

- New commitments based on recent data,
- Provision for emissions trading,
- Atmospheric stabilization as a goal, and
- Allocation of burdens influenced by ability to pay.

In Section 3 we present additional elements that, if added to the mix, could make these features into a useful architecture.

3.1 Negotiation of Near-Term Emissions Limits

A process is emerging whereby nations negotiate emissions limits, and do so only for a specified future period with the procedure repeated over time. In contrast to other environmental agreements, such as the Montreal Protocol, it has not been possible to agree to a multi-decade path of commitments or the prohibition of particular activities. The first climate agreement to take this form was the statement in the original Convention that,

> [Annex I parties] shall adopt policies and take corresponding measures . . . with the aim of returning [by the end of the present decade] individually or jointly to their 1990 levels [of] these anthropogenic emissions of carbon dioxide and other greenhouse gases not controlled by the Montreal Protocol (United Nations, 1992, Article 2, paragraph 2 [a and b]).

The article made no reference to emissions *after* the year 2000.

Many of the developed nations agreed to this target. All but a few will fail to meet it. However, even before national performance under the year-2000 commitment was known, COP-1 produced the so-called "Berlin Mandate," which instructed the negotiators to seek a set of "quantified emissions limitation and reduction objectives" for the years 2005, 2010 and 2020 (United Nations, 1995). In eight meetings of the Ad-Hoc Group on the Berlin Mandate (formed to carry out the decisions made in COP-1) and in intense final negotiations at COP-3 in Kyoto, the system of emissions targets was refined to apply to a "commitment period" instead of a specific year, and a set of nation-by-nation reduction percentages was agreed for the first such period, 2008 to 2012.

If ratified and put into force, the Kyoto targets would have greater force, under international law and the domestic law of some countries, than the "aim" written into the original Convention. But even if some nations fail to agree to the Kyoto targets, or to live up to the commitments they do make, we think it likely that this sort of rolling negotiation will be repeated. Indeed, the Kyoto Protocol anticipates a second and subsequent budget periods, and specifies that consideration of commitments for the second period shall start at least seven years before the end of the first period, *i.e.*, by 2005. Moreover, the history of this agreement gives one no reason to think it likely that negotiations of near-term emissions limits will be rejected in the foreseeable future for a tax-based alternative, or other means of shifting the focus from outcomes to actions[6]—even if (as seems increasingly probable) the Kyoto Protocol itself is eventually recognized as a false start.

---

[6] For example, Schmalensee (1998a) has argued that it would be better to assess each nation's compliance with its obligations *ex ante*, by evaluating the likely emissions reductions from policies it has adopted, rather than *ex post*, by measuring actual emissions. The precedents established by other international agreements do not seem to rule out a shift in this direction, but neither do they suggest that such a shift is likely.

3.2 New Commitments Based on Recent Data

For both the Convention's original goal for 2000 and the Kyoto goal for 2008-2012, target reductions have been set with reference to emission levels in a near-past year. For most commitments 1990 has been made the benchmark, and procedures have been developed to seek consistent estimates for this year (United Nations, 1996). In a key exception, to be discussed further below, Annex I nations "under the process of transition to a market economy" (*i.e.*, states of Eastern Europe and the former Soviet Union) may choose a year other than 1990. This allows them to set their baselines at the higher emissions levels experienced before the economic decline of the late 1980s.[7]

Importantly, the benchmark is always the individual nation's own historical emission level, not some absolute standard (such as tons of carbon per capita or per dollar of GDP) or the performance of some other nation. It does not seem plausible to us that any such external standard will be agreed to by national governments. The negotiation of commitment levels in relation to a recent historical benchmark is consistent with the privileged role assigned to the status quo in most policy debates, and, if only for that reason, this pattern is likely to continue.

3.3 Provision for Emissions Trading

The Kyoto Protocol provides several mechanisms that allow a nation to seek credit for emissions reductions outside its borders. Joint Implementation applies to the nations of Annex I, and a Clean Development Mechanism has been introduced in the search for a way to foster transactions between Annex I and Non-Annex I parties. However, the most important potential feature, in terms of its future role in facilitating emissions reductions, is the provision for emissions trading. Much remains to be resolved regarding its definition, and these details must be worked out in the face of sharp opposition from some nations. Nonetheless, it seems to us that international trading of emissions, perhaps encumbered by some restrictions on volume, is emerging as a key architectural feature of the climate regime. Without trading (or a shift to a regime of harmonized taxes), we do not believe it will be possible to achieve atmospheric stabilization at tolerable cost (Yang and Jacoby, 1999).

3.4 Atmospheric Stabilization as a Goal

The ultimate objective of the Climate Convention, as stated in Article 2, is to achieve

> . . . stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system. Such a level should be achieved within a time-frame sufficient to allow ecosystems to adapt naturally to climate change, to ensure that food production is not threatened and to enable economic development to proceed in a sustainable manner. (United Nations, 1992)

Considered as an element of a durable policy architecture, this objective is problematic: it presumes the existence of a threshold level of greenhouse gas concentrations that at worst may not exist at all and at best may be impossible to infer from scientific observations for many

---

[7] The other exception to 1990 as the base year is the accounting of HFCs, PFCs and SF_6, for which parties may chose either 1990 or 1995.

decades. On the other hand, it is hard to escape the notion that if the climate change problem is serious, seeking to limit human intervention at some level is a logical part of the response,[8] and negotiations under the FCCC seem to be building toward a specific number for this target. At this point it appears that the level chosen, stated in CO_2 equivalents, will most likely be one of the five (350, 450, 550, 650 and 750 ppmv) used by IPCC Working Group I (IPCC, 1996) for construction of stabilization scenarios.

In a system of rolling negotiations of near-term emissions targets, a long-term concentration goal is at most a guidepost for judgments about whether emissions control efforts in aggregate are sufficiently stringent. For example, the FCCC called for COP-1 to review the "adequacy of commitments" (United Nations, 1992, Article 4.2(d)). The fact that the "aim" of emissions reduction agreed at Rio would make little progress toward the Article 2 goal, whatever the target level chosen, is likely to have influenced the conclusion by the COP that measures to date were not adequate, even though no numerical stabilization target had yet been selected. In future negotiations the goal of stabilization will likely play the same role, with its influence increased if and when a particular numerical target is agreed.

Indeed, because of its role in future negotiations, and because there is currently no scientific basis for establishing a "danger" threshold, this atmospheric target is the focus of high-level debate about the risks, benefits and costs of greenhouse-gas emissions control. So far as we know, no nations have argued for a target of 350 ppmv. The global concentration of CO_2 alone (ignoring other greenhouse gases) is already at around 360 ppmv, and the emissions reductions required to stabilize concentrations at this level over the next century or two is perceived as too costly relative to the benefits involved. By the same token, we have heard little or no support for a 750 ppmv target. With no emissions controls at all, most long-range forecasts show CO_2 concentrations rising to well above 750 ppmv, implying significant costs to stabilize even at this high level. However, among those who are willing to consider long-term targets at all, the implicit benefit-cost analysis seems to fix the desired goal below 750 ppmv.

Discussion thus focuses on the range of 450 to 650 ppmv. Groups with a high estimate of the risks of perturbing the atmosphere relative to the costs of foregoing carbon-producing activities (environmental NGOs among them) argue for 450 ppmv. Others, showing a stronger concern with the costs of restricting these activities, and perhaps less worry about the environmental risks, say that even 550 ppmv is unreasonably stringent, but perhaps 650 ppmv is plausible. The process that decision-makers go through to arrive at this result is neither transparent nor simple enough to model in an explicit, rational-actor framework. Nevertheless, given the national positions in the negotiations to date, it now seems most likely that the balance of forces will lead to the selection of 550 ppmv. This level has the attraction of being in the middle of what has become the standard range of numbers, making it appear a moderate compromise. Also, it is roughly the standard doubled-CO_2 case used in climate-modeling studies over the past two decades, giving it the advantage of familiarity. Indeed, the E.U. has already proposed it.[9]

---

[8] Here we focus on concentration targets in the long term. The rate of change also may be important. Its inclusion is a possible extension of the approach explored below.

[9] United Nations (1997b: 3), Paper no. 1 by the Netherlands (on behalf of the European Community and its Member States).

Of course, such a target level would not necessarily remain constant in the long run, wherever it might be set initially. Technological developments or lower economic growth could reduce the expected cost of emissions control, and scientific results (or a run of bad weather) could heighten fears of climate damage, creating demands to tighten the goal to a more "precautionary" level. The same forces could also operate to loosen the goal. Thus a policy architecture that uses such a goal as a guide to "adequacy" of near-term control commitments must be able to adapt to changes in the goal itself.

3.5 Allocation of Burdens Influenced by Ability to Pay

Here it is hard to see the principle within the mass of detail, but we cannot imagine a workable policy architecture for this or almost any other large cost-sharing problem that does not assign a central role to considerations of ability to bear the burdens. Certainly from the early stages of the climate negotiations nations have been classified, roughly, by income level when discussing who should make what sorts of commitments. For example, the Framework Convention divides parties into three categories, plus one aggregate:

- **Annex II**. This group includes the rich nations, and is roughly equivalent to the membership of the OECD in 1990.
- **Economies in Transition**. As noted earlier, this group encompasses the nations of Eastern Europe and most of the Former Soviet Union, which have much lower incomes than most OECD countries.
- **Annex I**. This is the aggregate, a combination of Annex II and the Economies in Transition.
- **Non-Annex I**. This is the developing world, some with per-capita incomes akin to the less successful of the Economies in Transition, but most being far poorer.

Both the original Convention text, and the Berlin Mandate instructions to the negotiators who produced the Kyoto Protocol, hold that any emissions control obligation applies to Annex I parties only. All discussion of how the lower-income Non-Annex I countries might assume similar commitments has thus far been ruled off the formal agendas of the Conference of Parties and its subsidiary bodies.

Also, a division is made within Annex I itself: from the outset the Economies in Transition were to be given "a certain degree of flexibility" within the process (United Nations, 1992, Article 2, paragraph 6), presumably to account for their lower economic status within the Annex I group. The provision noted above allowing a baseline other than 1990 is one aspect of this flexibility. Differences in income also underline those portions of the Convention providing for assistance for developing countries, including assistance with data development and analysis, aid to countries particularly vulnerable to climate change, and technology transfer.

Of course, criteria other than ability to pay necessarily enter any real negotiations. The so-called "polluter pays principle," for instance, has been generally accepted in the international arena for some decades. In the climate area this criterion becomes associated with the concept of "historical responsibility," which is based on a calculation of the national origin if the greenhouse gases now present in the atmosphere. This consideration becomes another input to debates over burden distribution, although its implications are not very different from those of current income because there is such a high correlation between the two. In addition, nations seek credit for any

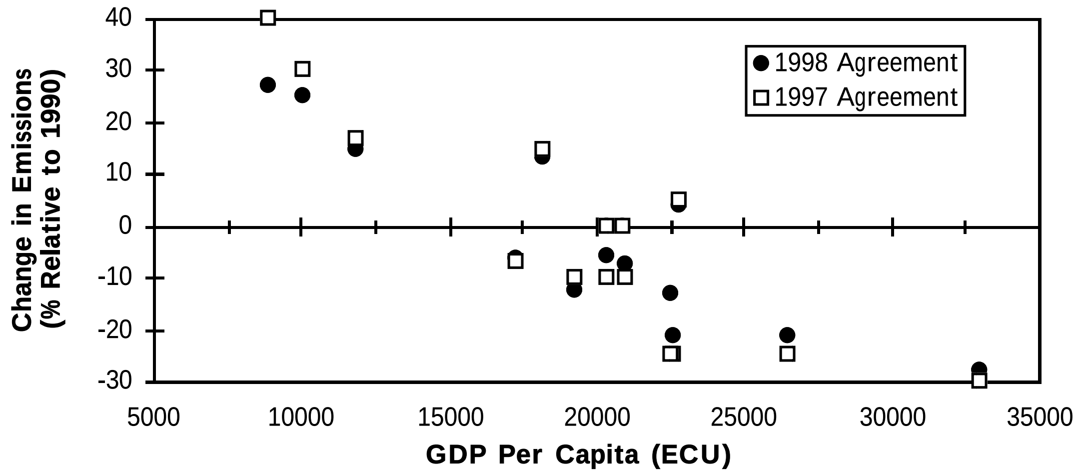

Figure 1. Implicit differentiation: E.U. burden sharing *vs*. ability to pay (2008-2012 period).

past controls that have helped contribute to an already-low emissions rate; some claim to be particularly disadvantaged by the choice of a base year; and some call attention to particularly high emissions control costs. Geopolitical factors matter as well, which helps explain why the Economies in Transition were included in Annex I as a bloc, even though some of them are poorer than several nations of the Non-Annex I group.

Nevertheless, when all nations participating in the negotiation are considered, it is hard to escape the notion that the most important single criterion thus far has been per capita income, treated as a measure of the economic ability to bear the burdens of emissions controls. We expect differentiation on this basis to become more explicit if serious limits on global emissions are adopted, simply because such agreements would require the participation of poor nations. With control obligations needed from notions with a wide range of incomes, significantly differentiated obligations would seem an obvious prerequisite to agreement.

A striking example of the apparent influence of income level can be seen in the results of the allocation among E.U. members of proposed Union level reductions. The E.U. went through this process twice: once before Kyoto, allocating an anticipated 10% Union-wide reduction below 1990 emission levels, and once after Kyoto to apportion the 8% reduction agreed to there. **Figure 1** shows the results of these intra-E.U. negotiations. Ability to pay clearly matters; equally clearly, it is not all that matters. Still, on average the richer countries are expected to make larger percentage reductions in emissions, and the poorest nations are even given room to increase emissions as they attempt to catch up economically.[10]

4. Developing a Useful Architecture

Although some elements of a policy architecture have been formulated in the Rio-Berlin-Kyoto process, important issues remain on which agreement has not been possible. As noted earlier, there is still great controversy over proposed flexibility mechanisms, and crucial details

[10] A plot of all nations, taking account of the Kyoto commitments and the EU internal allocation shows a similar but much more complicated pattern. It is a topic which merits further empirical analysis.

remain to be worked out on credit for carbon sinks. Also, the credibility of the entire policy regime requires advances in monitoring and enforcement. We have argued the importance of these topics elsewhere (Jacoby, Prinn and Schmalensee, 1998; Schmalensee 1998), and they remain essential elements of any useful architecture.

These shortcomings notwithstanding, we concentrate here on the no less essential task of extending the emerging policy architecture to handle issues of burden-sharing over time. We focus on three important issues that such an architecture must address and that the Kyoto framework does not:

- Accession of Developing Nations,
- A System of Rolling Baselines, and
- Imposing a Long-Term Stabilization Constraint.

We do not explicitly consider the need to induce the voluntary participation of an appropriately large coalition of nations, despite both its importance, and the good deal of game-theoretic attention it has received (Carraro and Siniscalco, 1993; Heal, 1994; and Barrett, 1994, 1998). Our reasons for this omission are several. Most of that literature has not dealt with the fundamental uncertainty that surrounds both the costs and, especially, the benefits of alternative climate regimes. For example, Barrett's (1994) model of a self-enforcing coalition relies on the reduction of *contemporaneous*, *global* benefits of abatement as punishment–inflicted upon defectors by remaining coalition members through cutbacks in their level of mitigation–to offset each country's individual incentive to free ride. In this framework it is unclear what equilibria might result when climate damage of uncertain magnitude and distribution occurs with a long time-lag after coalition formation. Also, one can either think of the need to build a coalition as imposing a set of constraints supplemental to the rules we describe below, or as necessitating a set of side payments. The latter we ignore in order to focus on the conditions under which countries might *accede* to a control regime, which are distinct from the incentives for countries to *comply* with the regime's provisions.[11] Finally, coalition-building on policy matters (on tax legislation, for instance) typically occurs either within the context of a given policy architecture or, less frequently, happens along with the construction of such an architecture. Thus, our point is that before developing countries even reach the position of being able to substantively disagree with the allocative provisions of an individual protocol, they first have to buy into the architecture that guides the evolution of the negotiations.

4.1 Accession by Developing Nations

The division of nations into two fixed camps on the emissions control question, Annex I *vs.* Non-Annex I, creates a negotiating deadlock that somehow must be resolved. It is not disputed that the stabilization goals discussed above cannot be met without emissions controls by major Non-Annex I parties (Jacoby, Schmalensee and Reiner, 1997). On the other hand, major Non-Annex I nations, particularly the poorest among them, have clear incentives to oppose any formal discussion of control obligations on their part. Climate may not rank high among their social (or even environmental) priorities, and even where it does rank high, nations may resist commitment

---

[11] For discussion of the design of incentives for countries to "opt in" to a greenhouse gas reduction regime, see Montero (1998).

on grounds of equity arguments familiar in North-South debate. Poor nations have good reason to doubt that developed countries would increase foreign aid enough to help much with the cost of obligations they are encouraged to assume.

Moreover, developing countries have good strategic reasons to postpone discussions of control obligations which include them as well as the rich nations. Climate change has already claimed a place on the wider foreign policy agendas not just of the United States and other developed countries but of large LDCs like China and India. These developing countries may well believe that their future bargaining power is strengthened (perhaps leading to larger aid flows, or more favorable conditions of ultimate entry into the control regime) if they face no agreed structure of universal obligation.

Smaller, wealthier developing countries like Chile or Singapore have a related but different set of incentives. They also are recipients of foreign policy pressures from the United States and other Annex I countries, but being small they may be more susceptible to those pressures. Also, some are close to membership or have already joined rich nation clubs like the OECD. Considerations of this sort surely go a long way to explain why nations like Romania and Estonia are in Annex I. As a result, some of these countries might be willing to accede voluntarily to Kyoto-style emissions targets, provided they are loose enough (providing sufficient “headroom” in climate jargon) not to threaten future economic growth. In Kyoto, however, a proposal to allow such voluntary accession was defeated on the opposition of China, India and others, for a complicated set of reasons very likely including the tactical ones suggested above.[12] Clearly, resolution of this deadlock is a key step in the evolution of an architecture for the long run.

We believe that the most promising approach to a solution is to focus negotiations on the development of rules that relate short-term emissions reductions obligations, if any, to national characteristics in ways that are perceived as fair. Countries cannot be forced to accept and discharge emissions control obligations. The only alternative, even if the perceived environmental threat is severe, is voluntary participation. We believe that a necessary condition for countries to be willing to accede to an international emission reduction treaty and to accept obligations to control emissions is that the (explicit or implicit) rules governing accession and control place a heavy weight on ability to pay. Grossman and Krueger (1995) find evidence that nations are more willing to protect their environments as they become richer (the so-called environmental Kuznets curve), although the nation-specific benefits from global emissions control are generally too speculative to be employed in negotiation. Prior reliance on ability to pay in climate negotiations, as discussed above, and widespread reliance on progressive tax regimes reinforce the plausibility of this basic approach.

4.2 A System of Rolling Baselines

As described above, discussions to date of emission targets have been based mainly on a 1990 baseline. In just the seven years from 1990 to COP-3 in Kyoto, however, economic conditions changed in unanticipated ways in many countries, complicating the negotiations.

---

[12] At COP-4 in Buenos Aires, both Argentina and Kazakhstan indicated their willingness in principle to assume emissions reductions obligations in the future. It is important to note, however, that these were basically press announcements, completely outside the COP agenda.

The continuing troubles of the former Communist states and the long and unanticipated U.S. expansion are obvious examples. By 2005, when discussions of the second commitment period are to begin, 1990 will be 15 years in the past. Some nations will be much richer than in 1990; other will not. Some may have made progress with emissions controls, others not. Discussion of relative burdens in terms of percentage reductions below 1990 will either be hopelessly artificial or will require a complex set of country-by-country corrections, in effect updating the baseline.

It seems clear that rolling negotiations on emissions ceilings, which seem to be clearly implied by climate negotiations so far, will require a rolling baseline. That is, at each round of negotiations nation-specific reductions will be defined using a baseline measured as close as practical to the date of new decision. The FCCC process, which has been taxed by the demands of establishing emissions baselines for 1990, will soon need to focus on establishment of a procedure and timeline for continual updating of these national inventories.

This observation about economic surprises also has implications for the policy analysis task. It is conventional in economic modeling studies to compute a long-term forecast of GDP in the absence of any policy and to measure burdens as departures from this counterfactual trajectory. We do this below. The procedure is reasonable in *ex ante* evaluation of a long-term policy. But burdens defined using some long-ago no-policy forecast can have little connection to perceived sacrifices and thus little relevance for negotiations. What will matter most is the cost of the next round of commitments and the abilities of different nations to bear them—even though those abilities may have been determined in part by the results of prior negotiations.

Finally, there is the question of what happens when a country fails to meet (or more than meets) its commitments for some particular budget period. Only time and experience will tell, but it seems likely that each new budget period will be free from the hangover of earlier periods. Countries that reduce emissions more than promised may gain some bargaining advantage in the next round, and those exceeding their quotas may lose it, but so many exogenous factors (such as recessions and booms) will have contributed to these results that it will be hard to require formal adjustments to future commitments based on previous emissions performance. It is worth noting, for example, that performance on the Rio commitment (to return to 1990 emission levels by 2000) seemed irrelevant to the negotiation of the Kyoto targets. Moreover, the COP quickly dropped the U.S. pre-Kyoto proposal that nations be automatically penalized in the subsequent budget period for any emissions above their commitments.

4.3 Imposing a Long-Term Stabilization Constraint

Much analysis of stabilization targets has been carried out using two sets of atmospheric CO_2 concentration profiles. The first was created by Working Group I (WG-I) of the IPCC for its Second Assessment Report (IPCC, 1996). The IPCC specified concentration trajectories which stabilize the atmosphere by some time in the mid-22nd century, and used a model of the carbon cycle (in an inverse calculation) to generate a corresponding emission path. To stimulate consideration of the economic effects of the choice of path, Wigley, Richels and Edmonds (1996) proposed an alternative set of paths (usually denoted WRE) that exhibited higher concentration levels in earlier periods but, assuming the same model of the carbon cycle, stabilized CO_2 concentrations at the same ultimate levels over the same time horizon. The WRE paths shifted the control burden to the future, assuming more drastic cuts in later periods.

The WG-I and WRE emissions paths underlie a growing body of research into the economic and environmental implications of the choice of an atmospheric concentration target. In model studies by Jacoby, Schmalensee and Reiner (1997), Yang and Jacoby (1997), and Manne and Richels (1997), and others, these trajectories are imposed as constraints on global emissions, for study of the welfare effects of alternative procedures for allocating emission reduction burdens among world regions. While these studies are instructive, their emphasis on *global* cost-effectiveness tends to underplay important distributional realities of climate change negotiations, and how they might be worked out over time. As we have discussed, we believe that, as at Kyoto, negotiations will be continue to be about emission levels in the near term, not about emission trajectories over several decades.

How can one connect negotiations about near-term emission reductions with a long-term stabilization objective? Put another way, how can one evaluate the “adequacy” of proposed emission control commitments in light of an agreed-upon stabilization objective in order to impose a useful discipline on negotiations? This basic question arises in any regime involving a sequence of negotiations about near-term actions, such as taxes or regulatory policies, or emission levels. Given the Rio-Berlin-Kyoto architectural elements described above, we believe there is a fairly natural answer in this context.

Suppose that the negotiating process can be constrained, perhaps by modification of the FCCC, to produce a rule relating commitments to, among other things, ability to pay. The natural analogy is to negotiations about U.S. income tax legislation, which produce a rule—albeit a very complex one—relating tax liability to, among other things, ability to pay. We would then propose that the adequacy of any proposed rule be judged by asking whether, if it were maintained over time, it would meet whatever atmospheric stabilization target had been selected in current or prior negotiations. Calculations of this sort would of course require an agreed-upon reference set of assumptions about future economic growth and technical progress, just as projections of income tax revenue do. There would, of course, be no presumption that future negotiations would simply adopt the same rule chosen in the past. A different rule leading to the same atmospheric concentration goal may have become more attractive, or a different concentration goal may seem more sensible. But at each round of negotiations the outcome would be required to be consistent with long-term atmospheric stabilization.

It is not totally straightforward to test whether a particular rule does in fact stabilize atmospheric concentrations of greenhouse gases. For simplicity we focus on CO_2. Let us define an emission path as a series of t positive carbon emission allocations $C(r, t)$ among regions r over a sufficiently long time horizon $[t_0,T]$:

$$C\,(r,t) > 0 \quad \forall r,t \in [t_0,T]$$

When these emissions propagate through the carbon cycle they result in a path of CO_2 concentration levels,

$$CO_2(t) > 0 \quad \forall t \in [t_0, T]$$

In the abstract, one might want to define stabilization of atmospheric CO_2 concentration at target level C^* in some future target date $t_s \leq T$ as follows:

$$CO_2(t_s) = C^*$$

and

$$\left.\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{CO}_2(t)\right|_{t=t_s} \leq 0$$

In practice, however, it is impossible to hit targets of this sort exactly. We thus replace these two conditions by the following:

$$\mathrm{CO}_2(t_s) = C^* + \varepsilon_l$$

and

$$\left.\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{CO}_2(t)\right|_{t=t_s} \leq \varepsilon_r$$

Determining whether the emissions path corresponding to a particular allocation rule results in a stabilizing concentration trajectory is thus analogous to a terminal condition problem. An allocation scheme will be consistent with stabilization at C^* over the horizon if at t_s the atmospheric CO_2 concentration exceeds the target level by not more than a tolerance ε_l, and exhibits a rate of increase of no more than ε_r ppmv per annum. To identify feasible emissions paths one must therefore specify ε_l and ε_r for desired values C^* and t_s.

In calculations below, $C^* = 550$ ppmv and $t_s = 2150$ (following the WG-I and WRE procedures for 550 ppmv), and an emissions path is said to stabilize concentrations at a particular target level if:

- The simulated concentration exceeds the target in 2150 by less than 10 ppmv, or has peaked at a level less than 10 ppmv above the target and has begun to decline ($0 < \varepsilon_l < 10$), and
- In 2150 the rate of change in concentration is less than 0.5 ppmv per year ($\varepsilon_r < 0.5$).

The tolerances that define this stabilization "window" can be made more or less strict depending on the needs of the analysis. For our purposes we need only to be able to define alternative paths that are roughly comparable in their climatic effects, so we can study the behavior of stylized burden-sharing schemes. Given the huge uncertainties in emissions forecasts over such a long horizon, and in the carbon cycle, efforts at great accuracy are not only misplaced but potentially misleading.

5. An Illustrative Model of Accession and Burden Sharing

To illustrate the basic architectural approach described above, we consider a situation in which *only* ability to pay matters. We seek neither to sell this approach as best, nor to trivialize the horse-trading on multiple criteria that will be part of any negotiation. Rather, we use it to demonstrate how a calculation relying on such an indicator could illuminate whether a particular set of commitments was consistent with a long-term stabilization objective. Suppose that each country (or region in our model analysis) becomes willing cut back its emissions once it exceeds a threshold level of per capita welfare (a welfare "trigger") stipulated for accession to the international regime. Otherwise, it follows its unconstrained (no-policy) emission trajectory. The rate at which each region agrees to reduce emissions (from a current-period baseline) in each period's negotiations is a function of the welfare threshold w^* and a per capita welfare data-point in the immediate past (w_{t-1}), as follows:

$$\eta_t = \left(\frac{\dot{C}}{C}\right)_t = \begin{cases} \gamma - \alpha\left(w_{t-1} - w^*\right)^{\beta} & w_{t-1}^{ref} \geq w^* \\ \eta_t^{ref} & \text{otherwise} \end{cases}$$

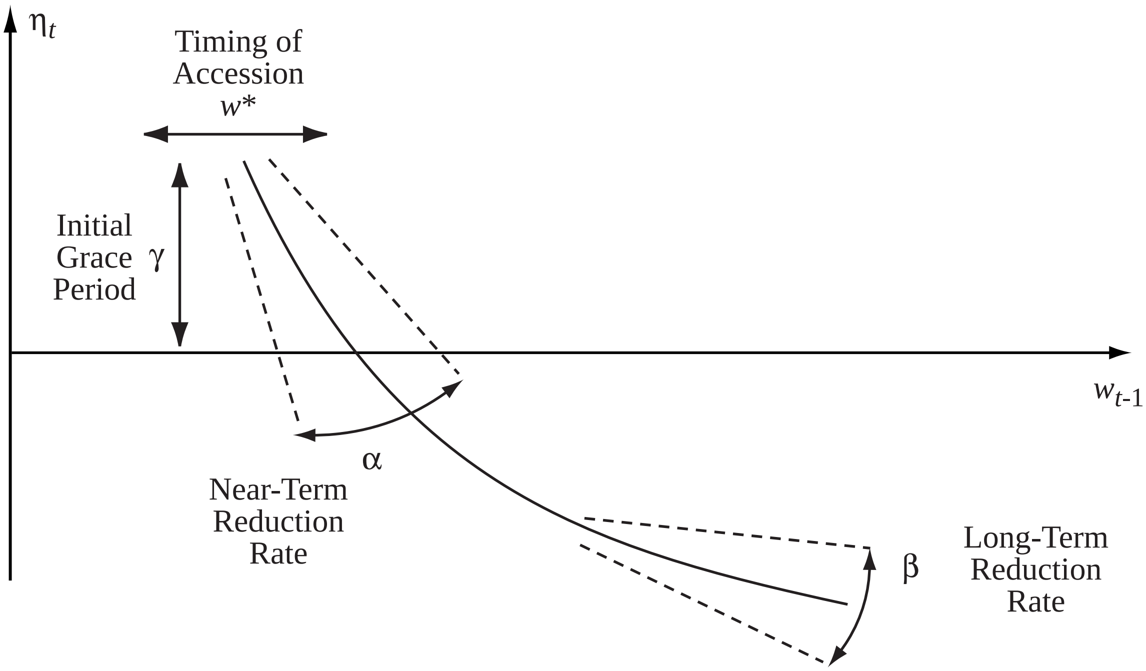

Figure 2. A sample burden-sharing scheme.

Figure 2 demonstrates how this formulation captures the stylized facts of accession. The coefficient α influences the overall rate of emission reduction but has its greatest effect on nations with incomes only slightly above w^*. The term β, on the other hand, strongly affects emission reduction rates when incomes are relatively far above the accession threshold. The term γ mitigates the effect of α and β; it determines the maximum rate at which regions should slow the rate of growth of their emissions prior to beginning absolute reductions. By varying w^*, we control the timing of the protocol's entry into force in the different regions. In order to ensure that the stabilization objective is met, of course, variations in w^* generally require offsetting variations in the other three parameters. Overall, then, the stringency, distribution, and timing of emissions reduction burdens can be determined by the choice of values for parameters α, β, γ, and w^*.

The functional form for the rate η_t is intentionally made as simple and transparent as possible, to focus attention on the larger issues of accession and subsequent obligation. We would not expect such a simple family of functions to govern real negotiations. We are in effect proposing an income tax system here, and we are analyzing it by examining functions relating tax obligations to income—even though real income tax systems base tax obligations on a host of additional factors as well. Here we focus on this four-parameter version for reasons of tractability, and we explore how it might serve to guide negotiators as they debate the many details that will inevitably engage their attention.

5.1 Calculation Procedure

To explore the implications of this illustrative structure for discussing burden allocation guidelines, we simulate a set of cases using the MIT Emissions Prediction and Policy Analysis (EPPA) model. EPPA is a recursive-dynamic CGE model based on OECD General Equilibrium

Environment (GREEN) model (Yang *et al.*, 1996). The world is represented by 12 nations and groups, as shown in **Table 1**. Each contains a utility-maximizing individual producer-consumer, each has eight production sectors and four consumption sectors, and all are linked by bilateral trade in both energy and non-energy goods. The endowments of each representative agent are updated at each step, according to assumed exogenous trends in rates of population growth, labor productivity increase, and technological change. The welfare index used here is one computed from national/regional consumption, and it is used to define w^*. In this discussion we speak interchangeably about this index and income, because of their tight correlation.

The effect of emissions on atmospheric concentrations is modeled by the carbon component of the MIT Integrated Global System Model (Prinn *et al.*, 1999). The atmospheric target is assumed to apply to CO_2 alone, thus avoiding the complexities of aggregation of the radiative effects of the trace gases. The essential issues of policy architecture are not disturbed by this common simplification. A logical extension of our approach would be to include other greenhouse gases, which are also forecast by the EPPA model (Reilly *et al.*, 1999).

In all previous applications, EPPA has been run to 2100, or for some shorter period. To explore stabilization calculations comparable to others in the literature, the time horizon of the EPPA model has been extended to 2150. There is an element of heroism (or folly) in pushing a model of this type (or indeed any economic model) so far into the future, and distant results should not be taken too seriously. What the model results do show is which of the burden-sharing schemes we formulate makes even approximate sense as a guide to rolling negotiations. As we have emphasized above, the idea is that parameter values would be agreed to only for the near term; their adequacy would be evaluated by considering whether maintaining them in place in the long run would result in stabilization.

Table 1. EPPA Regions

Annex I	
Annex II	
USA	United States
JPN	Japan
EEC	E.U.: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Portugal, Spain, Sweden, UK
OOE	Other OECD nations: Australia, Canada, New Zealand, the European Free-Trade Area (excluding Switzerland and Iceland), and Turkey
Economies in Transition	
FSU	Former Soviet Union
EET	Eastern European economies in transition: Bulgaria, Czechoslovakia, Hungary, Poland, Romania, Yugoslavia
Non-Annex I	
EEX	Energy-exporting developing countries: OPEC states & other nations exporting oil, gas, and coal
CHN	China
IND	India
BRA	Brazil
DAE	Dynamic Asian Economies: Hong Kong, Philippines, Singapore, South Korea, Taiwan, Thailand
ROW	Rest of World

5.2 Performance with No Emissions Trading

Experiments with this system reveal that as long as the welfare threshold, *w**, is not set too high, a number of parameter value combinations, if maintained over time, produce trajectories consistent with our definition of stability. Differences in the values of our four parameters, like differences in income tax rates, translate directly into the burdens imposed on nations at different income levels.

For the sake of brevity, we focus on the implications of two parameter value combinations, a High case with a relatively high value of *w** ($4500 in 1985 U.S. dollars) and a Low case with a relatively low value ($3,000) of this parameter. The values of the parameters are the following:

High Case: $w^* = \$4500$, $\beta = 0.25$, $\gamma = 0.015$, $\alpha = 0.03$
Low Case: $w^* = \$3000$, $\beta = 0.20$, $\gamma = 0.025$, $\alpha = 0.03$

Figure 3 shows the trajectories produced in a regime without inter-region emissions trading, assuming these parameter values are maintained over the entire period. The timing of accession under the various cases is shown in **Table 2**. In the Low case, Non-Annex I nations accept control commitments earlier, so these obligations can rise less rapidly with welfare and still achieve stability. In the High case, for instance, our model and reference case parameter values imply that China assumes obligations in 2030, while in the Low case it does not do so until 2050. In general, a decision to raise the threshold for accession (*w**) requires some combination of reduced "headroom" for the poorest countries with control obligations (lower γ) and a more rapid increase in obligations with income (higher α and/or β). These are the sorts of tradeoffs that are the stuff of politics, domestic or international; they can productively occupy negotiators for long periods.

Also shown in Figure 3 are the WG-I and WRE cases for 550 ppmv. We follow the same procedure as for our own cases, running WG-I and WRE emissions trajectories through the MIT carbon cycle model. The resulting concentration levels vary somewhat from those presented by

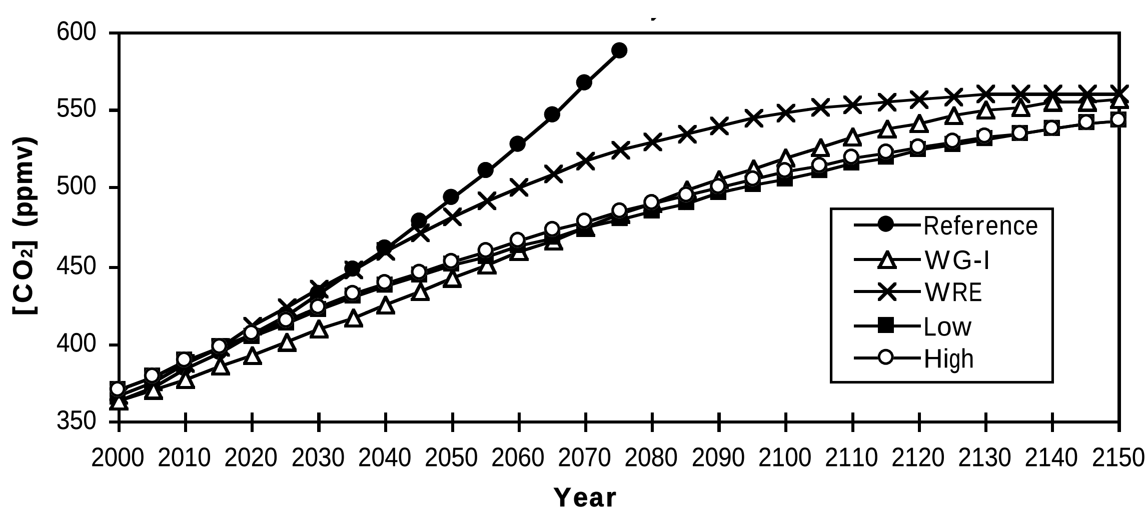

Figure 3. Stabilization of atmospheric CO_2 concentration at 550 ppmv: Low and high welfare thresholds (no trading).

Table 2. Timing of Accession as a Function of Welfare Trigger, *w**

Regions	No Permit Trade		With Permit Trade	
	High *w**	Low *w**	High *w**	Low *w**
EEX	2020	2015	2020	2015
CHN	2050	2030	2045	2030
IND	2090	2065	2085	2060
DAE	2015	2015	2015	2015
BRA	2035	2020	2035	2020
ROW	n/a	2130	n/a	2115

the IPCC (1996) and Wigley, Richels and Edmonds (1996), illustrating the uncertainty in models of the carbon cycle. Indeed, carbon cycle uncertainty is much larger than the differences between our estimates of WG-I and WRI concentration levels and those presented by their authors.

The paths resulting from our High and Low cases are similar to the WG-I result. They differ substantially from the WRE path, however, which imposes no controls in the early decades, and then strong reductions in the middle to latter part of the next century. Studies carried out in the Stanford Energy Modeling Forum and elsewhere (*e.g.*, Manne and Richels, 1997) tend to favor the WRE Path, but none of the control patterns then analyze is the result of a simple guideline for rolling negotiations of the type explored here.

Some discussions of long-term burden-sharing schemes take the position that fairness requires moving toward equal per-capita emissions across nations, and it is worth checking the performance of this structure of commitment under that criterion. The per-capita emissions implied by our High policy are shown in **Figure 4** for the United States (USA), Europe (EEC), Japan (JPN), China (CHN), and India (IND). The effect of the Kyoto commitments are clearly seen in the early years for members of Annex I. By the latter part of the next century, emissions per capita have roughly converged.

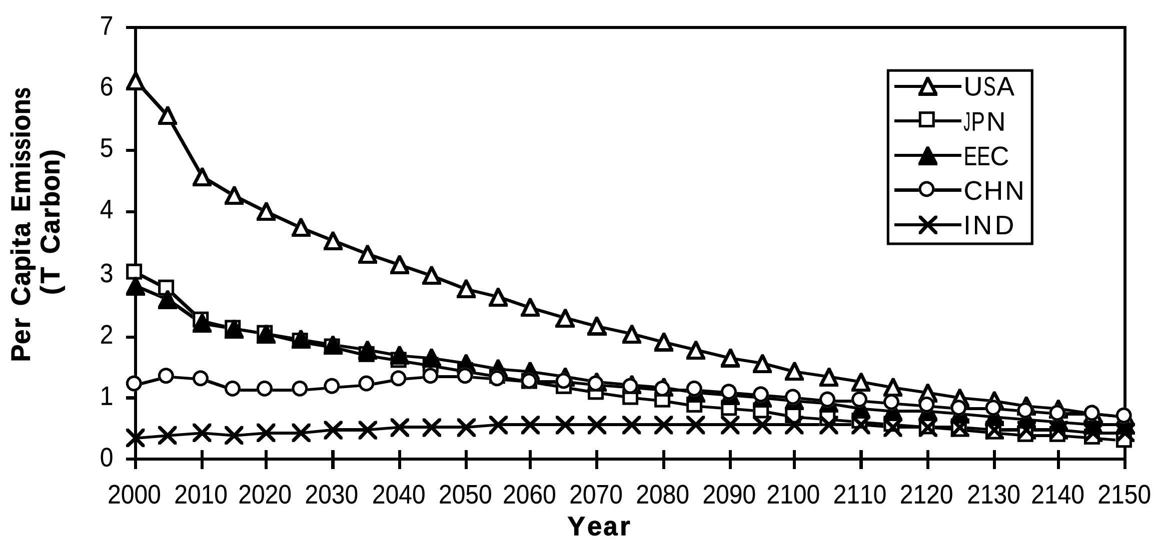

Figure 4. Convergence of regional per capita CO_2 emissions (*w** = $4500, no trading).

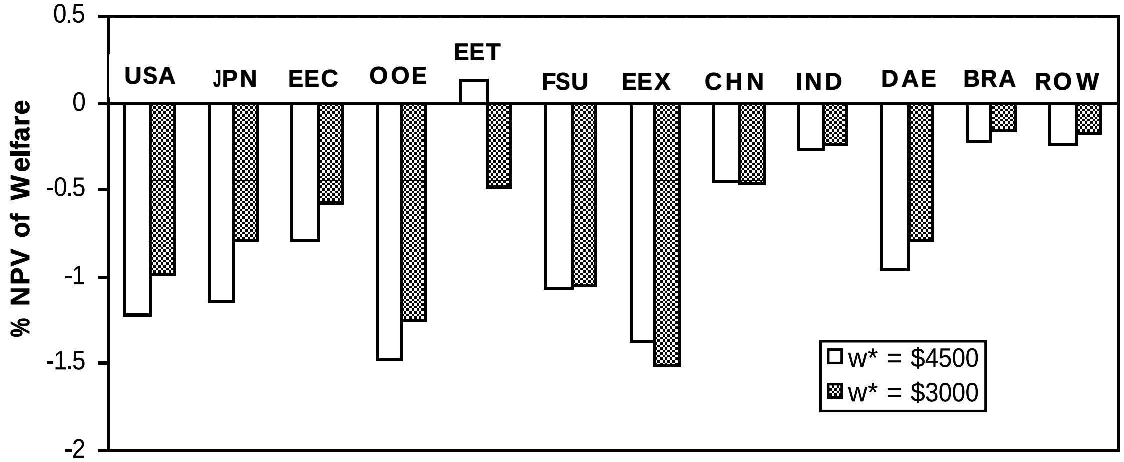

Figure 5. Effects of welfare threshold on distribution of policy costs (5% discount rate, no trading).

Figure 5 shows the discounted welfare costs relative to our reference, no-policy economic scenario for each of the EPPA regions (see Table 1) for our High and Low cases. For the parameter values chosen, a move from the High to the Low case (with Non-Annex I countries entering earlier and a lower long-term reduction rate, β) the wealthier nations (USA, JPN, EEC, OOE) all benefit. For the others the result is mixed, and this result is due to the effects of leakage. All the poorer countries are helped by the \$3000 as compared to the \$4500 entry threshold. But under the \$4500 trigger many were benefiting from carbon leakage (*i.e.*, a shift in comparative advantage and carbon emissions because of the greater stringency imposed on richer countries). With the shift to \$3000 trigger (putting less pressure on the USA, JPN, EEC and OOE), this economic benefit is lost. As a result, Eastern Europe (EET), energy exporters (EEX) and China (CHN) benefit overall from a shift to the Low case, but the others on balance do not. These global economic adjustments are a complex phenomenon, worthy of further investigation with EPPA and other models. It is important to note, however, that Figure 5 reflects the costs of trajectories along which the values of all four parameters are constant from the present through 2150. While this is a useful summarization device, we cannot imagine a century and a half of negotiations actually leaving burden-sharing arrangements unchanged.

5.3 The Effects of Emissions Trading

Inclusion of emissions trading in the architecture would have a substantial effect on burdens, even under the same structure of allocation. Only full global trading, without restriction, has been analyzed for this illustration, so it gives an upper bound on the difference from an architecture without trading. We focus on the High case, simply applying the High case parameter values to a calculation where trading is allowed. **Figure 6** shows that the addition of trading leads to a somewhat lower trajectory of emissions, although the 550 ppmv stabilization criterion is still met. (This also holds in the Low case.) Trading tends to raise incomes, activity levels, and thus emissions. On the other hand the higher incomes also cause some regions to enter the control system sooner, as shown in Table 2. For instance, more rapid income growth in China leads to its surpassing the accession threshold in 2045 instead of 2050.

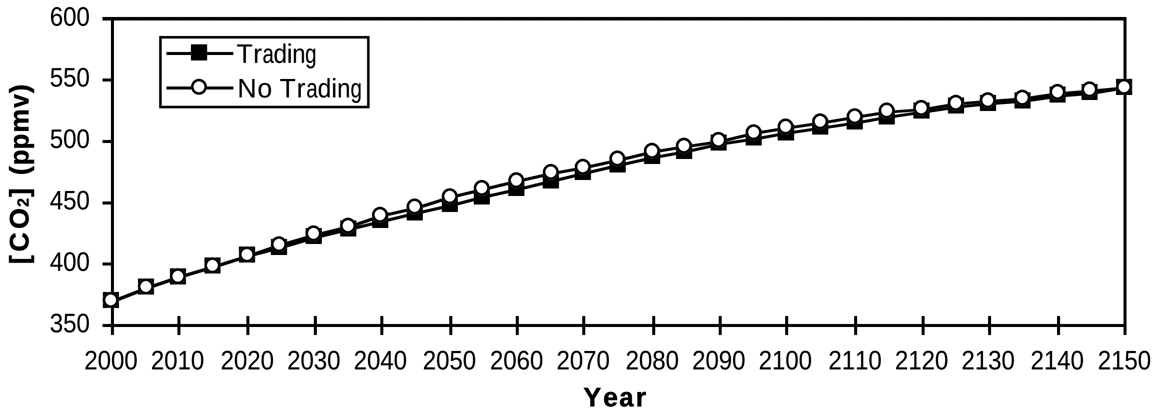

Figure 6. Stabilization of atmospheric CO_2 concentration at 550 ppmv: Effects of trading (w* = $4500).

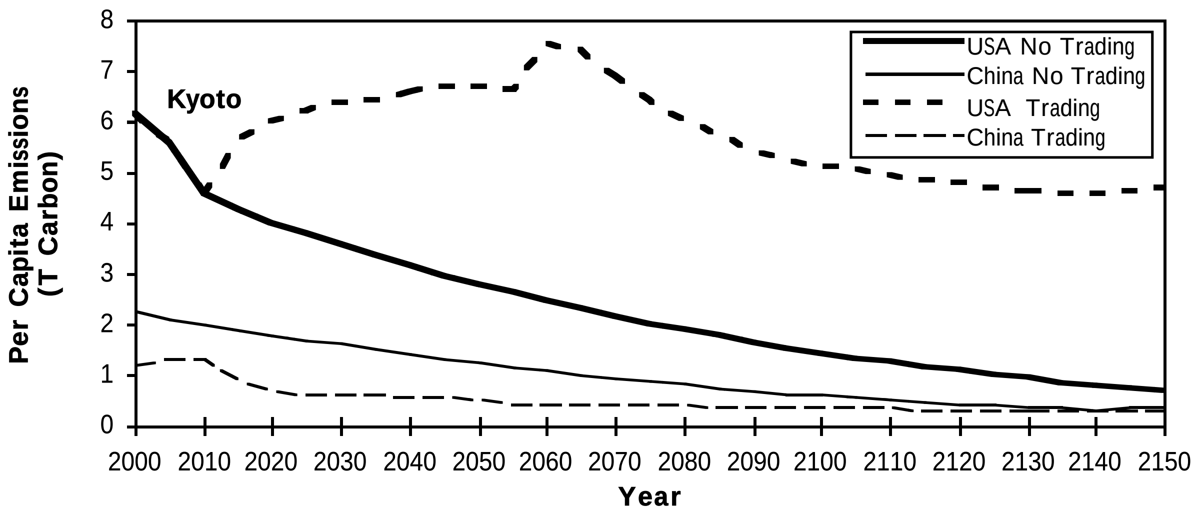

Figure 7. Representative per capita emissions (w* = $4500).

Figure 7 shows what happens to per-capita emissions in the U.S. and China as a consequence of trading in the High case. As shown in Figure 4, China's per-capita emissions entitlement surpasses that of the U.S. in about 2070. When those entitlements are made tradable, however, per-capita emissions in the U.S. are far above the U.S. entitlement, while China finds it beneficial to control emissions more than its entitlement implies. (The picture for the low case is similar.) In this scenario, per-capita emissions in the U.S. continue to rise until around the middle of the next century, after which they decline more rapidly than those in China.

Figure 8 shows the effect of unrestricted full global trading on the costs of the High case. Consistent with earlier work on this aspect of a control regime, the addition of trading lowers costs (or increases benefits) across all regions. This result holds as well for the Low case, not shown here. Clearly, this is a crucial component of climate architecture, as suggested above.

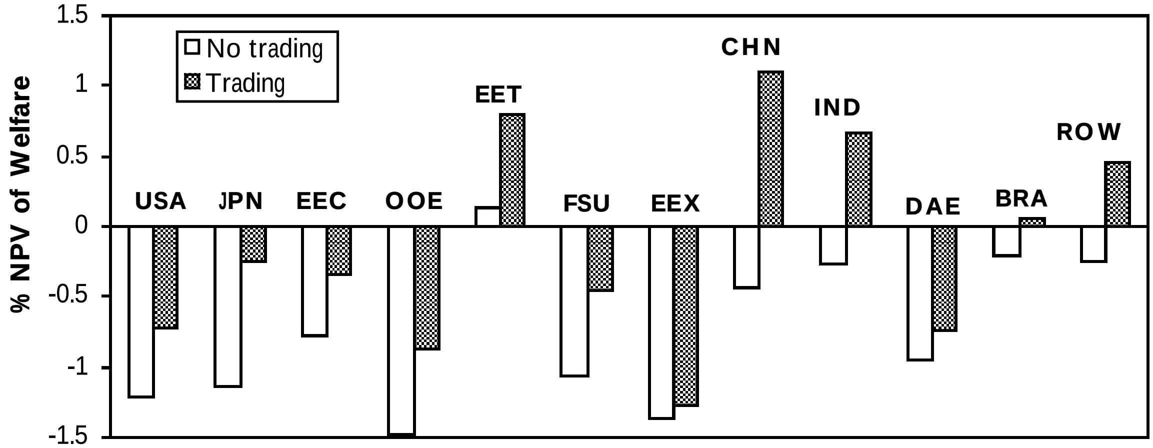

Figure 8. Effects of full global trading on distribution of policy costs
(5% discount rate, *w** = $4500).

6. Conclusions

We could easily devote more attention to the properties of the illustrative burden-sharing rule developed in Section 5. We do not do not do so, however, because we intend it to be only a simple (arguably simplistic) example of the sort of rule that might emerge from rolling emissions-reduction negotiations, each using a recent (and thus different) baseline. Such negotiations no doubt would produce more complex burden-sharing rules, just as the real income tax code is much more complex than a table of tax rates. The more complex the allocation produced by a particular negotiation, the harder it would likely be to verify that the underlying structure, if maintained over time, would lead to stabilization under reasonable assumptions. But we feel that some such test will prove necessary if the notion of a long-term stabilization target is to be made meaningful in a world in which it is impossible to commit to a long-term emissions trajectory.

We believe that in any fair regime, substantial differences in per-capita income will translate into substantial differences in emissions control obligations, though of course other factors will matter in the negotiations. This argues that climate negotiators should begin now to consider explicitly how differences in real GDP per capita (for which internationally accepted data generally exist) should affect accession and burden-sharing.

At the most fundamental level, though, the argument of this paper is that, like it or not, climate negotiators are both arguing about short-term national commitments and building a policy architecture that will focus and shape future negotiations. We have argued what elements the negotiations to date have contributed to a useful policy architecture, and we have explored one approach to completing such an architecture. But we do not argue that our approach is the best way to build on the Rio-Berlin-Kyoto framework; it is only the best we could think of. We would be surprised if hard work by smart people did not produce a better path to a useful policy architecture. This paper will have fully achieved its objective if it serves to stimulate such work.

7. REFERENCES

Barrett, S. (1994). Self-Enforcing International Environmental Agreements, *Oxford Economic Papers* **46**(5) Suppl., pp. 878-94.

Barrett, S. (1998). Political Economy of the Kyoto Protocol. *Oxford Review of Economic Policy* **14**(4), pp. 20-39.

Carraro, C. and D. Siniscalco (1993). Strategies for International Protection of the Environment. *Journal of Public Economics* **52**, pp. 309-328.

Cusumano, M.A., and R.W. Selby (1995). *Microsoft Secrets*. New York: The Free Press.

IPCC (1996). *Climate Change 1995—The Science of Climate Change*, Contribution of Working Group I to the Second Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, UK.

Grossman, G.M., and A.B. Krueger (1995). Economic Growth and the Environment, *Quarterly Journal of Economics* **110**(2), pp. 353-377.

Heal, G. (1994). Formation of International Environmental Agreements. In C. Carraro (*ed.*), *Trade, Innovation, Environment*. Dordrecht: Kluwer.

Jacoby, H.D., R. Prinn, and R. Schmalensee (1998). Kyoto's Unfinished Business. *Foreign Affairs* **77** (July/August), pp. 54-66.

Jacoby, H. D., R. Schmalensee, and D. M. Reiner (1997). What Does Stabilizing Greenhouse Gas Concentrations Mean? In B.P. Flannery, K.R. Kohlhase, and D.G.L. LeVine (*eds.*), Proceedings of the IPIECA Symposium *Critical Issues in the Economics of Climate Change*, pp. 225-244.

Manne, A., and R. Richels (1997). On Stabilizing CO_2 Concentrations: Cost Effective Emission Reduction Strategies, *Environmental Modeling and Assessment* **2**, pp. 251-265.

Montero, J.-P. (1998). Optimal Opt-in "Climate" Contracts, *Journal of Applied Economics* **1**(2), pp. 363-384.

Pizer, W. A. (1997). Prices v. Quantities Revisited: The Case of Climate Change. Resources for the Future, Discussion Paper 98-02.

Portney, P.R. (1999). The Joy of Flexibility: U.S. Climate Policy in The Next Decade. Keynote Address to National Energy Modeling System/Annual Outlook Conference, U.S. Energy Information Administration, March 12.

Prinn, R., H.D. Jacoby, A. Sokolov, C. Wang, X. Xiao, Z.L. Yang, R.S. Eckaus, P. Stone, A.D. Ellerman, J. Melillo, J. Fitzmaurice, D. Kicklighter, G. Holian, and Y. Liu (1998). Integrated Global System Model for Climate Policy Assessment: Feedbacks and Sensitivity Studies, *Climatic Change*, **41**(3-4), pp. 469-546.

Reilly, J., R.G. Prinn, J. Harnisch, J. Fitzmaurice, H.D. Jacoby, D. Kicklighter, P.H. Stone, A.P. Sokolov and C. Wang (1999). Multi-Gas Assessment of the Kyoto Protocol. MIT Joint Program on the Science and Policy of Global Change, Report No. 45.

Reiner, D.M., and H.D. Jacoby (1997). Annex I Differentiation Proposals: Implications for Welfare, Equity and Policy. MIT Joint Program on the Science and Policy of Global Change, Report No. 27.

Schmalensee, R. (1998a). Greenhouse Architectures and Institutions. In W. Nordhaus (*ed.*), *Policy Analysis for Decisionmaking about Climate Change.* Washington: Resources for the Future, pp. 137-158.

Schmalensee, R. (1998b). Tradable Emissions Rights and Joint Implementation for Greenhouse Gas Abatement: A Look Under the Hood. In C.E. Walker, M.A. Bloomfield, and M. Thorning (*eds*.), *The Impact of Climate Change Policy on Consumers: Can Tradable Permits Reduce the Cost?* Washington: American Council for Capital Formation, pp. 39-55.

Schmalensee, R., P.L. Joskow, A.D. Ellerman, J.-P. Montero, and E.M. Bailey (1998). An Interim Evaluation of Sulfur Dioxide Emissions Trading. *Journal of Economic Perspectives*, **12**(4), pp. 53-68.

United Nations (1992). Framework Convention on Climate Change. New York: United Nations.

United Nations (1995). Conclusion of Outstanding Issues and Adaption of Decisions. Framework Convention on Climate Change, Conference of the Parties, Berlin, 7 April.

United Nations (1996). Revised 1996 IPCC Guidelines for National Greenhouse Gas Inventories, FCCC/SUBSTA/1996/18/Add. 1, 16-18 Dec.

United Nations (1997a). Kyoto Protocol to the United Nations Framework Convention on Climate Change, Conference of the Parties, 3rd Session, FCCC/CP/1997/L.7/Add.1, 10 Dec.

United Nations (1997b). United Nations Framework Convention on Climate Change Ad Hoc Group on the Berlin Mandate, 7th Session, FCCC/AGBM/1997/MISC.2/Add.1, 27 June.

Webster, M., and D.M. Reiner (1997). Beyond Emissions Paths: Rethinking the Climate Impacts of Emissions Protocols in an Uncertain World, MIT Joint Program on the Science and Policy of Global Change, Report No. 31.

Wigley, T.M.L., R. Richels, and J.A. Edmonds (1996). Economic and Environmental Choices in the Stabilization of Atmospheric CO_2 Concentrations, *Nature* **379**, pp. 240-243.

Yang, Z.L., R.S. Eckaus, A.D. Ellerman, and H.D. Jacoby (1996). The MIT Emissions Prediction and Policy Analysis Model, MIT Joint Program on the Science and Policy of Global Change, Report No. 6.

Yang, Z.L., and H.D. Jacoby (1997). Necessary Conditions for Stabilization Agreements, in C. Carraro (ed.), *International Environmental Agreements on Climate Change*, Kluwer Academic Publishers, Lancaster, UK, forthcoming.