GLOBAL ENERGY POLICY CENTER

Research Paper No. 10-07

International Climate Games: From Caps to Cooperation

Peter Cramton and Steven Stoft

20 August 2010

Abstract

Greenhouse gas abatement is a public good, so climate policy is a public-goods game and suffers from the free-rider incentives that make the outcome of such games notoriously uncooperative. Adopting an international agreement can change the nature of the game, reducing or exacerbating the uncooperative tendencies of the players. We analyze alternative international agreements as variations of the public-goods game, and examine the incentives for cooperation under each alternative.

The addition of cap-and-trade rules to the basic public-goods game is found to polarize the free-rider incentives of that game, encouraging those who would abate the most to target even higher abatement levels and those who would abate the least to target lower, and even negative, abatement levels. Such polarization between developed and developing countries is familiar from both the Kyoto and Copenhagen climate summits.

Since cap-and-trade rules decrease cooperation by developing countries, developed countries are led to reject the game's outcome and in the process prevent agreement on a set of quantity targets. To break this deadlock and shift the equilibrium toward cooperation, a modification of the public-goods game based on price rather than quantities is needed. This involves a global price target and equity transfers via a Green Fund that rewards adoption of and compliance with such a target. The Nash equilibrium of one such game is analyzed for a group of three countries similar to the United States, China and India.

Contents

International Climate Games: From Caps to Cooperation

Peter Cramton and Steven Stoft*

20 August 2010

1. Introduction

Ignorance of the warming effect of greenhouse gasses (GHGs) has contributed to climate problems, but no one expects that scientific knowledge alone will lead to control of GHG accumulation in the atmosphere. There is one root cause of this pessimism. GHG emissions cause an externality—an effect that is not priced by private markets—and consequently GHG abatement is a global public good.[1] Because each country is only a small part of the global economy, the resulting free-rider incentives dominate national strategies. Little abatement will be provided without an agreement on international cooperation.

Although this view is widely accepted, the problem of how to design an agreement to overcome free-rider problems has been largely ignored by negotiators, environmentalists and climate scientists. This is surprising in light of the cooperative failures observed starting with the Kyoto summit.[2] Despite clear warnings from a number of prominent economists (for example, Schelling 2002; Barrett 2003; Stiglitz 2006; Nordhaus 2008, Barrett and Toman 2010), there appears to be little recognition that there is a science of cooperation and that changing the rules of a game can make the play of the game more—or less—cooperative. Without such understanding, the negotiating strategy of the major players has been to induce cooperation by invoking scientifically determined numbers with vehement and often moralistic insistence.

However, the behavior of Copenhagen negotiators was highly strategic in nature. This is expected by game theory but assiduously ignored by even the best standard analysis of the negotiation process (Blair 2009). Ignoring strategic interactions, negotiators have pursued cap-and-trade rules that degrade the limited tendencies toward cooperation found in the public-goods game. This is a remarkable accomplishment, because the public-goods game is known for producing even less cooperative results than the prisoners'-dilemma game.[3]

The current framework for a global agreement calls for binding national targets and a global market for emission reductions.[4] This produces what we will call the cap-and-trade game, which is a modification of the classic public-goods game. Although the current failure to cooperate is bad news, the fact that the failure was the result of a grossly flawed negotiating framework is good news. Cooperation may still be well within reach if the game is changed to foster cooperation.

---

* Peter Cramton is professor of economics at the University of Maryland. Steven Stoft is director of the Global Energy Policy Center. We thank Roger Sparks for helpful discussions of the games presented here and seminar participants at the World Bank for helpful comments. We are grateful to the National Science Foundation for funding. Documentation and updates are available at www.global-energy.org.

[1] "It is generally understood that the climate system is a global public good par excellence" (Bradford 2008).

[2] "Most countries do not have country-wide emission limitations under the Kyoto Protocol" (Bosi and Elis 2005).

[3] "Consistent with previous findings, we observe that the average contribution [in a four-person public-goods game] declines in the control experiment" (Rand et al. 2009).

[4] "In the context of the Kyoto Protocol, most industrialised countries have adopted emissions obligations" (Bosi and Elis 2005).

This paper makes three fundamental points. The design of the climate policy game plays a key role in determining whether countries cooperate. The cap-and-trade game induces the types of uncooperative strategies that have become ubiquitous. Some other modifications of the public-goods game exhibit far more cooperative equilibrium behavior and hold much greater promise for an effective climate agreement.

Four Climate-Policy Games

To make these three fundament points, four global climate-policy games are investigated, starting with the standard public-goods game. In each case, countries choose the value of a single strategic variable, and a unique Nash equilibrium is discovered.[5] The four games are:

Table 1. The Four Games

Name of Game	Strategic Variable		Description
public-goods	Abatement	A_j	Independent action without a treaty
cap-and-trade	Quantity Target	T_j	Abatement + purchased credits = T_j
price-target	Proposed Price Target	P_j^T	Lowest choice becomes global target, P^T
Green-Fund	Proposed Price Target	P_j^T	Green-Fund payment proportional to P^T

In the first three games countries are identical except for size; in the Green-Fund game, countries also differ in their monetary valuation of climate benefits.

In the public-goods game analyzed here, countries abate, without any international agreement, to the point where their marginal abatement cost equals their marginal climate benefit. Since larger countries capture a higher proportion of the benefits of abatement (simply because they have more people), they abate to a higher level of marginal cost, also known as their domestic carbon price, P_j^G. Since this price does not take into account the benefits to other countries of their abatement, this price falls below the optimal price. Since the benefit functions are all assumed to be linear (except when noted), marginal benefit does not change with abatement and the optimal price, which is the same everywhere, is $P^O = \sum P_j^G$, where the sum is taken over all N countries.

The global cap-and-trade game was first analyzed by Helm (2003). Holtsmark and Sommervoll (2008) extended and improved the analysis. Carbone et al. (2009) analyzed the game by embedding it in a computable general-equilibrium model, which included trade in fossil fuel and energy-intensive goods between six regions. Godal and Holtsmark (2010) argue that the six-region model makes a number of assumptions that cause global cap-and-trade to perform better within the model than should be expected in practice.

The global cap-and-trade game is strategically dissimilar to the national cap-and-trade game (with which this paper is not concerned) because the global game lacks a government capable of allocating carbon allowances or permits. The global game focuses on national strategies for choosing targets. In the global cap-and-trade game there is a global market in abatement, which will be called the carbon market. But the resulting carbon price, P^*, is suboptimal in the resulting cap-and-trade equilibrium, and equals P^O/N, where N is the number of countries. With identical countries, this is the same as the domestic carbon prices in the public-goods game. But the outcome can easily be worse than the public-goods outcome as can be seen by considering a world with two countries one of which is 99 times larger than the other. The larger country, because it internalizes almost all the climate benefits will set its carbon price near P^O, and this will cover almost all emissions, while the tiny country will set its price near zero. In the cap-and-trade game, both will face a world carbon market with a price of $P^O/2$. This is much worse in 99 percent of the world and better in only one percent.

But this is not the main problem with cap and trade. The problem, which surfaced at the Kyoto summit and has been worsening every since, is the antagonistic split between rich and poor countries—or Annex I and non-Annex I countries as defined by the Kyoto protocol.

This polarization begins with such natural differences as those between large and small countries and between rich and poor countries. But the cap-and-trade framework exacerbates and monetizes polarization. Large and rich countries place a higher value on abatement, which means a higher price on carbon. Poor countries place

[5] A Nash equilibrium is a set of strategies such that no country can benefit by unilaterally changing its own strategy.

a lower value on abatement. But cap and trade creates a global carbon market that presents all countries with the same carbon price. Rich countries see it as low and target more abatement under cap-and-trade. But poor countries see it as high and target less abatement. In fact, as will be shown, the global market makes it rational and profitable for many of them to set an abatement target of less than zero and sell all "abatement" above that target level to the rich countries.[6]

To rich countries, this equilibrium seems unfair. They reject the perfectly rational response of poor countries to the proposed international framework, and so the negotiations deadlock.

The price-target game takes place in the same idealized world as the cap-and-trade game but depolarizes the negotiations and results in an optimal global carbon price instead of one N times too low. But in a more realistic world where rich and poor countries have different preferences, the price-target game also fails. The Green-Fund game corrects the problem that arises when poor countries prefer a low carbon price for income-related reasons. In this case, linking Green-Fund payments to the global price target encourages poor countries to favor a higher price. The higher the global price target, the greater their Green-Fund payments. This linkage also contributes to depolarization.

The Value of Strategic Analysis

That the cap-and-trade game fails to enhance cooperation should not surprise. It was not designed to do so. In fact it was designed for domestic use, a situation in which cooperation is enforced by a national government. But in the international setting, where there is no such authority, countries are free to act strategically, and they do (Finus 2002). Consequently, we need to replace the cap-and-trade framework with one that accounts for strategic incentives.

Some may protest that the cap-and-trade game is poorly modeled by non-cooperative game theory, because that assumes countries follow their self-interest. Instead, proponents of caps assume that setting targets will be done cooperatively, taking into account the scientifically determined needs of the planet. But the history since Kyoto suggests that the assumption of cooperation is unwarranted. Cooperation is much more apt to appear and survive if it is consistent with self-interested (non-cooperative) play.

The present analysis is not merely prescriptive; it is also predictive. It predicts that, when international cap and trade is the negotiating framework, non-Annex I countries will reject meaningful targets, that they will favor the U.N.'s Clean Development Mechanism, and that national governments will not be helpful to, and sometimes will obstruct, the U.N. as it tries to enforce its "additionality" constraint. It also predicts that Annex I countries will be willing to set higher targets if they can buy offsets or emission reduction credits in the world market. It further explains why non-Annex I countries want Annex I countries to set higher targets.

None of these predictions is particularly surprising. What is surprising is that the approach that correctly makes these simple predictions is not given more credence than the approach that gets them wrong—the approach based on a view that environmental concern and climate science will guide national policies. The advantage of the present approach is not so much in showing that today's uncooperative behaviors are easily explained by a game theoretic analysis of cap and trade, but that this analysis opens the door to thinking systematically about cooperation and how to attain it.

2. The Global Public-Goods Game

Greenhouse gas abatement is a global public good. It is non-rivalrous—one person's climate benefit does not detract from another's. And, no country can be excluded from global climate benefits. Although the following discussion could apply to any global public good, we call the public good "abatement," meaning a reduction in the emissions of greenhouse gases below some well-defined business-as-usual level. In the standard public-goods game, each country decides its own abatement level, and the benefits that accrue from abatement depend only on the global total of all national abatements.

[6] Although a negative abatement target may seem outlandish to those unfamiliar with the economics of cap-and-trade, it is essential to some highly respected proposals (Frankel 2010; Bradford 2008).

This process can be, and often is, considered a one-shot game in which each country chooses an abatement level, and nature determines the payoff function. In fact, the current practice is to attempt to make the climate game into a one-shot game by locking in targets as far into the future as possible. In the one-shot game, each country chooses a single abatement level, A_j. The marginal costs of national abatements are increasing and the marginal benefits are (weakly) decreasing in total abatement. In this case the game has a unique Nash equilibrium. This can be found by maximizing the national net benefits of the public-goods game:

$$NB_j = B_j(A) - C_j(A_j), \text{where } A = \sum A_j. \tag{1}$$

Throughout the paper, all sums are taken over all countries, $1 \leq j \leq N$. The functions $B_j(A)$ and $C_j(A_j)$ are the national climate benefit and abatement cost functions.[7] Each country maximizes its net benefit by setting $B_j'(A) = C_j'(A_j)$, or marginal benefit equal to marginal cost. To obtain simple and illuminating results, we will assume throughout the paper that each country has a constant marginal benefit, B_j', but these marginal benefits will differ among countries. It will be useful to think of the marginal cost of abatement as a carbon price, P_j, even though countries may not price carbon explicitly. Hence in the public-goods equilibrium, $P_j^G = B_j'$ (superscript G will denote values in the public goods game that will be compared with similar values in other games).

The socially optimal outcome is found by maximizing, for each national abatement, A_j, the global net benefit

$$NB = \sum [B_j(A) - C_j(A_j)].. \tag{2}$$

This optimization requires that each country take into account the total global benefit provided by its abatement instead of taking account of only its own benefit. Naturally, the socially optimal abatements are much greater than the abatements that constitute the Nash equilibrium. For future reference, optimizing equation (2) over all A_j results in a uniform marginal cost of abatement, which can be thought of as the global price of carbon. This optimal price is given by

$$P^O = \sum B_j' = \sum P_j^G. \tag{3}$$

If the marginal benefit of abatement is constant, then comparing equations (1) and (2) shows that with identical countries the public-goods game produces a uniform carbon price of $P^G = P^O/N$.[8] The optimal price must take account of the benefit in all N countries, but the public-goods price accounts for only one.

As a point of reference, suppose countries are identical except for size, and the cost functions are quadratic and the benefit functions linear. Then, in the Nash equilibrium, a country that is responsible for 20 percent of the emissions under business as usual will abate by one fifth the optimal amount, and smaller countries will abate proportionally less. If Europe is counted as one country, a somewhat optimistic assumption, total abatement under these standard but idealized assumptions, for the world's Nash equilibrium, would be about 12 percent of optimal.

A Nash equilibrium defines the outcome of a game that is played non-cooperatively, meaning that each player considers only its own benefits and participates in no coalitions. All of the games considered will be non-cooperative. However we will refer to *outcomes* as being more or less cooperative depending on how far the abatements are from the socially optimal levels. Abatement levels of 12 percent of the optimal level constitute an uncooperative outcome. Fortunately, some games have highly cooperative outcomes, even though they are played in the non-cooperative manner just described.

In the climate-change context, a global public-goods game may be more appropriately modeled as a repeated game, because countries can change their abatement strategy in response to the strategies of other countries. For example, Europe could abate by more than its Nash level for a while, and then if other countries did not follow suit, it could abate by less than its Nash level. Another country could adopt the classic tit-for-tat strategy, in which case it would mimic Europe, but with a delay. If many countries did this, Europe would succeed in leading the world to a cooperative outcome.

---

[7] The abatement-cost approach is equivalent to the emission-cost approach, but aligns better with policy discussions that talk in terms of 80% abatement by 2050, and take business as usual to be the origin.

[8] This relationship holds if marginal benefits under cap-and-trade equal marginal benefits under an optimal policy. Marginal benefits of this year's abatement is a long-term benefit stream. The early part of the stream will be nearly identical under either policy because the stock of CO_2 changes slowly. If current policies are viewed as temporary, they will be seen as changing far-future CO_2 stocks very little. In this case, marginal benefit could be seen as nearly identical under the two policies.

However, in the 20 years since the world began trying to cooperate on abatement, cooperation has not been sustained. In fact there is no clear evidence that we have done even as well as the public-goods Nash equilibrium. This is not surprising. Experiments with public-goods games typically show that, even with repeated play, the outcomes are far from the social optimum. Unlike with two-person prisoners'-dilemma games, cooperation does not tend to build up with repeated play, but instead tends to die out as the players learn more about the free-rider possibilities (Rand et al. 2009).

Because of the historical and experimental evidence and because there is no theoretical case for believing the one-shot Nash equilibrium will not be predictive, we believe that it does provide a useful guide to the likely outcome. Moreover, we suggest that it is reasonable to be skeptical that games with even less cooperative Nash equilibria will produce better outcomes if implemented. Similarly, it is reasonable to believe that games with much more cooperative Nash equilibria will produce outcomes that are more cooperative than those observed to date.

3. The Global Cap-and-Trade Game

The outcome of the public-goods game is inefficient for two reasons. First, much too little abatement takes place, and second, the abatement that is undertaken is not undertaken efficiently. Abatement proceeds in each country until marginal abatement costs equal marginal climate benefits. Since larger countries have higher marginal benefits they sustain higher marginal costs, which means they could achieve the same abatement more cheaply if they could pay small countries to abate.

Global cap and trade implements a global market for abatement that allows high-marginal-cost countries to purchase abatement from low-marginal-cost countries, thereby making efficient the abatement that is undertaken. This solves what is most likely the smaller of the two efficiency problems. The other problem—that too little is abated—is the cooperation problem.

Cap and trade has long been promoted for its efficiency, but as the Copenhagen climate summit made all too apparent, cooperation is presently the larger problem.[9] This paper is concerned first with the implications of cap-and-trade for cooperation and second with finding what other modifications of the basic public-goods games might improve cooperation more.

Because cap and trade allows countries to, in-effect, abate by paying others to abate, a country's move in a one-shot global-cap-and-trade game consists of choosing a target, T_j, instead of a level of domestic abatement, A_j, which is instead determined by the carbon market. A target can be met with either domestic abatement or with abatement purchased in the global cap-and-trade market at the global market-clearing price of abatement (price of carbon), P. It will be assumed for now, that countries fulfill their targets and do not cheat. Consequently the, net-benefit of cap and trade for country j is

$$NB_j = B_j(A) - C_j(A_j) - P\cdot(T_j - A_j). \tag{4}$$

The new term is the cost of purchasing abatement or the revenue from selling abatement. This equation holds even if the private-sector abaters in a given country both buy and sell abatement from and to other countries.

The game proceeds in two stages. First all countries choose a target, T_j, just as they choose an abatement level, A_j, in the public-goods game. Second, the global market distributes the abatement efficiently so that the marginal cost of abatement in all countries equals the global price of abatement, P.

Cap-and-Trade: The Optimal Price Divided by N

Optimizing net benefit with respect to the choices of the target, T_j, with other targets held constant, results in the following simple relationship, derived in the appendix, which will be used shortly to explain polarization:

$$B_j' - P - P'\cdot(T_j - A_j) = 0, \tag{5}$$

[9] "On the notion of cooperation leading to IEAs [international environmental agreements], we can distinguish between two views. One is economic theoretic, and the other is game theoretic" (Parkash and Tulkens 2008).

where B_j' and P' are, equivalently, derivatives with respect to either total abatement or T_j. From an individual country's perspective, equation (5) says that increasing its target will increase its climate benefit but impose two costs. First, there is the direct marginal cost, P, of paying for a higher target, and second, the increase in global abatement will increase the global carbon price and hence the cost of all purchased abatement.

Summing equation (5) over all countries, dividing by N, and using equation (3) gives

$$P^* = (\sum B_j')/N = (\sum P_j^G)/N = P^0/N, \tag{6}$$

because the sum of ($T_j - A_j$) is zero. This result depends on the marginal benefit of global abatement being the same in the cap-and-trade game and in the optimal outcome. While this may be realistic if the current policy is viewed has having little effect on the long-run accumulation of greenhouse gases, a more realistic assumption holds that B_j' decreases as global abatement increases. As will be seen shortly global abatement under cap-and-trade can be more or less than under a public-goods regime. To resolve this ambiguity we consider the case when cap-and-trade is helpful, in other words, the case in which it increases abatement relative to the public-goods outcome. The meaning of equation (6) is summarized in Proposition 1.

Proposition 1. In the Classic global cap-and-trade game, the global price, P^*, equals $1/N$ times the optimal price if the marginal value of abatement is constant, and is less than the average public-goods price if marginal abatement decreases with abatement and cap-and-trade serves to increase global abatement.

The qualifier "Classic" anticipates a discussion at the end of this section and refers to a cap-and-trade game that encompasses all abatement activity. Proposition 1 places the cap-and-trade outcome in a restrictive box. If it manages to improve abatement, it must do so within the constraint that it lowers the unweighted country-average price of carbon. Since there are more small countries than large countries and since small countries, ceteris paribus, have a lower public-goods price, it is difficult for cap and trade to attain a lower carbon price than the average public-goods price. This indicates that cap-and-trade is not likely to increase abatement, a conclusion reached by Holtsmark and Sommervoll (2008) with similar reasoning. In a world with identical countries, cap-and-trade opportunities provide no perverse incentives and the outcome is identical to the public-goods game. However, a problem arises when countries differ in size.

Global Cap-and-Trade Polarization

The source of the cap-and-trade problem can be seen by solving equation (5) for the target values:

$$T_j = A_j^* + (P_j^G - P^*)/P'. \tag{7}$$

Note that B_j' has been replaced by P_j^G because countries using carbon pricing in the public-goods game would set price equal to B_j' (recall that B_j' is assumed to be constant). Since polarization compares equilibrium values under cap-and-trade with those in the public-goods game, ambiguous variables are distinguished by "*" when they describe the cap-and-trade equilibrium and by superscript G for public-goods variables. All derivatives are evaluated at the cap-and-trade equilibrium. Since P' is the derivative with respect to T_j it is also the derivative with respect to total abatement A. Inverting this derivative gives

$$T_j = A_j^* + (P_j^G - P^*)\, \mathrm{d}A/\mathrm{d}P. \tag{8}$$

Now consider a country with $P_j^G > P^*$. The second term on the right in equation (8) would approximate the increase in A_j as the country's carbon price increases from P^* to P_j^G if $\mathrm{d}A/\mathrm{d}P$ were $\mathrm{d}A_j/\mathrm{d}P$ instead. In other words the second term would be approximately $A_j^G - A_j^*$. So, in this case, T_j would equal the abatement in the public goods game, A_j^G. But since the derivative is not country-specific, but is instead the derivative of global abatement with respect to the global price, it is greater than $\mathrm{d}A_j/\mathrm{d}P$, and T_j is therefore greater than A_j^G. There are other secondary effects to consider, as well as the case where $P_j^G < P^*$, and these are considered in the appendix, which proves Proposition 2.

Proposition 2 (polarization). If the marginal benefit of abatement is constant and the marginal cost of abatement is increasing and quadratic, then a country with a domestic carbon price that is higher (lower) in the public-goods game than the cap-and-trade price will set a cap-and-trade target that is higher (lower) than its public-goods abatement.[10]

In part, this is good news. High-abatement countries pay for even more abatement under cap-and-trade because they can buy abatement more cheaply. In fact this effect is well known, and frequently taken advantage of. The European Emissions Trading System allows the purchase of UN-Certified Emission Reductions from developing countries at a price that is intended to be near P^*. Allowing this has long been recognized to facilitate the acceptance of higher abatement targets for the EU.[11] Similarly, allowing the purchase of billions of tons of foreign offset under the Waxman-Markey bill was a key factor in gaining acceptance of a higher target in the bill passed by the U.S. House of Representatives. In other words, developed countries are setting higher targets because they can take advantage of cheap abatement in developing countries, just as polarization in the cap-and-trade Nash equilibrium predicts.

The only puzzle concerning this result is why it is not much stronger. Why don't countries with a high marginal benefit of abatement keep raising their target without limit, since abatement can be purchased at a price of P^* which is below their marginal benefit? Of course, if they did, the global price would eventually rise to their marginal benefit, but equation (6) tells us that in the cap-and-trade Nash equilibrium, high-abatement countries set targets that leave the global price much lower than their marginal benefit. Why do they set such a low target?

The answer involves the P' term in the denominator of equation (7). That equation tells us that the derivative of the global carbon price limits the targets of high-abatement countries. The cost of increasing an abatement target is not just P^*; there is another component. Raising an abatement target raises the demand for abatement and hence the price of abatement. And this raises the cost of all the abatement that a country is purchasing. When the sum of these two cost components reaches P_j^G, which is the marginal benefit to country *j* of more abatement, the country decides its target is high enough.

Low-abatement countries

Now consider a small country with a low marginal benefit. It will find that it must pay more in the world market than it is willing to pay in the public-goods game. Consequently, the small country's naïve analysis leads it first to conclude that it should target no abatement at all. But it will still find itself abating because other countries will pay it to abate. At that point it realizes that merely by setting a negative target it appears that it will be paid more for exactly the same amount of actual abatement, because it is paid $P \cdot (A_j - T_j)$ and its abatement, A_j, is determined by the world price and not by its choice of T_j. So why doesn't this lead to a complete collapse of the cap-and-trade market, with every country that has a marginal benefit even slightly below P^* setting its target to some arbitrarily negative value?

The answer is the reverse of the reason that high-benefit countries do not set their targets arbitrarily high. When a country sets its target lower, it has more "abatement" (relative to its target) to sell. This increases its sales of abatement, but also lowers the global price it is paid for abating. At some point, it cannot make up in volume what it loses in price reductions, so it limits its supply of abatement—or more accurately its supply of abatement credits.[12]

It may seem that the idea of changing a target to account for its effect on the global carbon price—either up or down—is simply too subtle for national governments to comprehend. This may be. But if they ignore this effect, then there is no limit on the amount of polarization that appears reasonable. Of course, there would be political limits in both directions. But if the political limits are binding then all countries are pushing the polarization effect as far as they dare politically. This may, in fact, be a more accurate description of the anti-cooperative behavior

[10] This result is similar in appearance to Helm's Proposition 1, but differs crucially in that it compares cap-and-trade abatement with public-goods abatement.

[11] "[T]he CDM makes it easier for Annex I Parties to commit to deeper emission cuts" (Schneider 2008).

[12] Although the idea of raising the global price of carbon by limiting the supply of carbon credits from low-emission countries may seem farfetched, it has been discussed by Schneider and Cames (2009) as a possible U.N. policy in the context of a Sectoral Crediting Mechanism (SCM). "On the supply side, introducing a cap on the issuance of credits from a SCM could … ensure reasonable [high] credit prices."

that is taking place. But still the Nash equilibrium explains why we may be bumping into political constraints on polarization.

In any case the Nash targets of low-abaters are likely to be well below doing nothing at all (meaning the emission cap will be above business as usual). Such negative targets constitute a national-policy that illustrates the famous additionality problem faced by the UN with its Clean Development Mechanism. As was expected from the outset, developing countries are quite happy to sell "abatement" that includes measures that are not "additional," but which are measured from a target level of emissions that exceeds business as usual.[13] Again, simple game theory has predicted a major problem found in practice.

Two Global Cap-and-Trade Examples

The simplest example of polarization has two countries, with one twice as large as the other. The two net-benefit equations are:

$$NB_1 = 2A - (A_1)^2 - P\cdot(T_1 - A_1) \tag{9a}$$

$$NB_2 = A - 2\,(A_2)^2 - P\cdot(T_2 - A_2). \tag{10a}$$

The second country is smaller, with half the benefit and twice the cost per unit abatement. From equation (6) we know that the price will be the average marginal benefit or $P^* = 1.5$. Setting marginal cost equal to price gives $A_1 = 0.75$, and $A_2 = 0.375$. Using equation (7) the target values can be calculated as shown in Table 2.[14]

Table 2a. A Two Country Example (Identical Countries Except for Size)

	Optimal Abatement			Public-Goods Game			Cap-and-Trade Game			
j	A_j	MC_j	NB_j	A_j	MC_j	NB_j	T_j	A_j	$P^* = MC_j$	NB_j
1	1.5	3	2.250	1	2	1.5	1.125	0.75	1.5	1.125
2	0.75	3	1.125	0.25	1	1.125	0	0.375	1.5	1.406
Total	2.25		3.375	1.25		2.625	1.125	1.125		2.531

A second example of polarization is more extreme, with the low-benefit country setting a negative abatement target. This time the countries differ qualitatively and not just by size. Again the cap-and-trade price is the average marginal benefit and thus half of the optimal price. The two net-benefit equations are:

$$NB_1 = 2A - 2\,(A_1)^2 - P\cdot(T_1 - A_1) \tag{9b}$$

$$NB_2 = 0A - (A_2)^2 - P\cdot(T_2 - A_2). \tag{10b}$$

Table 2b. A Two Country Example (One Country Does Not Value Abatement)

	Optimal Abatement			Public-Goods Game			Cap-and-Trade Game			
j	A_j	MC_j	NB_j	A_j	MC_j	NB_j	T_j	A_j	$P^* = MC_j$	NB_j
1	0.5	2	2.5	0.5	2	0.5	1.00	0.25	1.0	0.625
2	1	2	−1.0	0	0	0	−0.25	0.50	1.0	0.500
Total	1.5		1.5	0.5		0.5	0.75	0.75		1.125

These examples illustrate several points. First, changing the game from a simple public-goods game to a cap-and-trade game can reduce or increase the total abatement of the Nash equilibrium (Helm 2003). (If countries differ only by size, then Holtsmark and Sommervoll's (2008) Proposition 2 implies that abatement will decrease.) Second, cap and trade can reduce or increase total net benefit. Third, it can increase polarization—that is, it can

[13] See Wara and Victor (2008). Schneider and Cames (2009) describe this concern most explicitly, "An overestimation of the crediting baseline could result in crediting considerable amounts of business-as-usual emissions, undermining the environmental integrity of the mechanism and potentially resulting in a collapse of the global carbon market."
[14] See spreadsheet available at www.global-energy.org/lib/2010/10-07.

cause high-abatement countries to target an even greater abatement level and low-abatement countries to target even less abatement. Fourth, the change of rules can be good for low-abatement countries and bad for high abatement countries.

The example in Table 2b bears some similarity to the analysis of Carbone et al. (2009). After dividing the world into six regions (USA, Japan, Western Europe, China, Former Soviet Union (FSU), and "Rest of World"), they find that the Europe-China-FSU coalition maximizes abatement compared to all other such coalitions that could implement cap and trade. However, the FSU adds very little, and the Europe-China coalition provides about 99% as much abatement. In example 2b, country 1 is similar to Europe with a strong desire for abatement, but high costs, while country 2 is similar to China in the six-region model with no desire for abatement and low abatement costs.

As a consequence, cap and trade allows Europe to purchase abatement cheaply from China, so it sets its cap-and-trade abatement target above its public-goods abatement level. China realizes it can profit from selling more carbon permits if it sets a negative abatement target, and so it does. In spite of this, both countries increase their net benefit, though China gains four times more than Europe. In the Carbone-et-al. analysis, which includes trade in fossil fuel, China has an additional motive for cooperation. By abating more, it uses less oil which helps lower the world price of oil, which reduces the cost of its oil imports.[15] We will return to this six-region analysis shortly, but first we must discuss a broader model of cap and trade.

Cap and Trade with Out-of-Market Strategies

The cap-and-trade game just presented is what might be termed the "Classic" game. It assumes implicitly that all abatement will be purchased in the global market, so that cap-and-trade is all encompassing. But this has not been the case, nor have there been proposals to make this a rule. As a consequence, countries have strategies available that are not considered in the above analysis. Countries can fulfill any part of their target abatement commitment outside of the global carbon market (Tol 2008; Godal and Meland 2010; Godal and Holtsmark 2010).[16]

To see why out-of-market strategies make sense, note that the high abatement countries limit their targets because of the indirect cost of raising the global price of carbon. But, starting at the Classic equilibrium, if they abate outside the market, their cost of abatement is only their domestic marginal cost of abatement, which is at first the global carbon price and includes no indirect cost of raising the global price. That makes out-of-market abatement less expensive than a purchase of abatement in the carbon market. Thus, they should abate more than under Classic cap and trade. In fact, a high-abatement country—one with $P_j^G > P^*$—should be willing to abate up to the point where the marginal cost of abatement is P_j^G, exactly as it does in the public-goods game, provided it can accomplish this abatement outside of the global market.

Less obviously, low-abatement countries should also want to adjust their domestic abatement to the level they would choose in the public goods game. To see why this would increase their net benefit relative to the Classic cap-and-trade equilibrium, denote their target and domestic abatement level by T_j and A_j respectively. Under Classic cap and trade, low-abatement countries sell abatement and consequently set $A_j > A_j^G$, say by an amount ΔA. And although they could reduce only T_j, this would reduce the global demand for abatement and hence the price of abatement, damaging abatement profits. However, with out-of-market strategies, the country could lower T_j and also prevent abatement exports from increasing. This would mean subtracting ΔA from both T_j and A_j, so that carbon-permit exports, $A_j - T_j$, remained unchanged. With exports unchanged, the price of abatement in the world market would not be affected, but domestic abatement and its cost would decrease.[17] This decrease in cost would be greater than the decrease in climate benefits until abatement fell to the public-goods level, which the country finds optimal.

---

[15] This is a key effect which aligns energy and climate policies for the U.S. and China, and which has been overlooked by negotiators. It can be used to benefit most climate policies not just cap and trade.

[16] Tol considers out-of-market policies with effects quite similar to the ones considered here. "The country with the lowest safety valve would effectively set international climate policy. It would flood the market with exported permits. The revenues may induce a veritable race to the bottom." Godal and Meland study the exercise of market power on both sides of the permit market. Godal and Holtsmark consider this exact out-of-market effect in detail.

[17] A country can adjust A_j down by failing to implement business-as-usual abatement policies such as requiring scrubbers on coal plants, ending fossil-fuel subsidies, or building transmission for wind power. More simply, it might limit the extent of its domestic cap-and-trade market as the EU has done.

In conclusion, rational countries would chose the same levels of domestic abatement as in the public-goods game, but there will still be trade in the carbon market. This means that the marginal cost of abatement will not be equalized across countries, and abatement will not be conducted efficiently. It also means that the market price cannot equal the global marginal cost of abatement but will instead be determined by the exercise of market power. This market-power game appears difficult to analyze, but since abatement is determined outside the market-power game. The market will serve only to transfer wealth from countries that set their targets above their abatement level to countries that set them below. This conclusion, which is substantiated in more detail by Godal and Holtsmark (2010), is summarized in the following result.

Result. If out-of-market control of abatement is allowed and unrestricted in the cap-and-trade game, the Nash equilibrium will entail the same national abatement levels as the public-goods game, so abatement will no longer be efficient. The global price of carbon will be determined by the exercise of market power on both sides of the market, and will have no effect other than to transfer wealth.

In other words, high-abatement (large or rich) countries will withhold abatement demand from the carbon market and supply themselves at an above-market price in order to avoid raising or, as the developing countries would see it, to suppress the global carbon price.[18] (This price suppression has been analyzed by Schneider (2008) in the context of the CDM mechanism.[19]) Similarly, low-abatement countries would restrict the supply of abatement in order to raise the global price of abatement. Of course, only large countries will have enough market power to make much difference. The net result, if the game were played without political pressure, would likely be rich countries paying poor countries to "abate" from their low targets up to what they would have done anyway.

Real-World Cap-and-Trade?

As discussed above, Carbone et al. (2009) have analyzed a six-region model of the world using a computable general equilibrium model. They find that non-cooperative cap-and-trade can, under the right circumstances, induce the world to abate at a level half way between the no-policy public-goods equilibrium and the optimal abatement level. While their model represents a significant leap forward in the realistic modeling of cap-and-trade and appears to reach an optimistic conclusion, it should be interpreted in the context of its assumptions.

First, the CGE model tacitly assumes Classic cap and trade, as just discussed. So far there is no indication that any country has contemplated conforming to this restriction which would require both Europe and China to give up all abatement incentives other than cap and trade. Of course this would be inefficient, since quite a few abatement measures—advanced research, building standards, and so on—are best achieved by other means. But if countries act strategically as assumed by Carbone et al., and if they are not required to implement Classic cap and trade, then Table 2b would show no increase in abatement.

A second cooperative problem is explicitly assumed away in the CGE analysis. Although many possible coalitions could form, and although these coalitions provide different payoffs to the regions in the resulting coalitions, coalition formation is not analyzed. Instead it is simply assumed that "nature," a short-hand term for international negotiation, will select the most globally beneficial, stable, coalition. In other words, "stage 0" of the game is assumed to be cooperative. Since the problem from the start has been the difficulty of the international negotiating process, this seems to sidestep the central question.

The problem of assuming cooperation during stage 0 seems particularly acute give the striking asymmetry of the proposed outcome. It seems likely that many European citizens would protest that they should not bear nearly the full global burden of climate policy, while the U.S., Japan, and the rest of the world stand on the sidelines. Since, by assumption, during stage 0 they realize this will be the outcome, agreement seems unlikely.

The basic dynamic discovered by the CGE model does seem to point in the right direction. Europe has in fact set the highest targets and has paid China a considerable amount of Certified Emissions Reductions under the UN's Clean Development Mechanism. But experiments with public goods games and experience with global

[18] "[S]ome developed countries have expressed reluctance to purchase vast amounts of international offsets for the purpose of compliance, and in some cases provided quantified limits to such purchases (see below). The general policy trend seems to be towards a restricted, not unlimited, demand for credits post-2012" (Baron et al. 2009).

[19] "The limitation on the use of CDM results in two different prices … The limitation on the use of the CDM mainly reduces the rents of the CER suppliers."

negotiations indicate that the Europe-China coalition is not likely to be a path to solving half of the climate problem.

4. Better Games

Modifying the public-goods game by adding cap-and-trade rules improves the game in two ways. It makes the abatement that is undertaken efficient and hence less expensive, and this encourages high-abatement countries to abate more. Unfortunately, the new rules also encourage low-abatement countries to commit to less abatement, and this effect is so strong that there will likely be little net benefit in the Nash equilibrium.

But the actual effect of changing the game is even worse than what would occur if it were played with Nash strategies. The outcome in practice is that the game is rejected. The players simply choose not to play such a game.[20] The problem stems from the polarization it causes.

In theory, the high-abatement countries should reject the cap-and-trade rules because accepting them will likely reduce their net benefit. Of course, if they did accept the cap-and-trade game, then the developing countries would accept it and set their targets to zero or lower, and they would profit from all of the abatement they undertook. This increases their net benefit relative to the public-goods game in which all players make some level of costly contribution to the public good. In theory, then, we should expect the low-abatement countries to maximize the social pressure they put on developed countries to accept the cap-and-trade rules, set high targets, and allow the developing countries to pursue their Nash strategies. Such pressure has been prominently on display in recent years to such an extent that the Indian edition of Reuters chose the headline "U.N. climate chief faces widening rich-poor split," to warn Christiana Figueres of her central problem just after her appointment as the U.N.'s chief climate negotiator.[21]

The theory presented so far follows the traditional free-rider analysis in which the size of countries is the only determinant of relative abatement levels. This means that there is absolutely no justification for polarization on the part of low-abatement countries, because these countries are no different in characteristics such as income, wealth, or cost per unit abated. In this model, cap-and-trade simply causes countries that free-ride the most in the standard public-goods game to become even more anti-social under cap and trade. Instead of simply shirking, they actively exploit the larger countries, choosing to exacerbate the public goods problem so that they can sell permits to the more cooperative countries. In short, cap and trade transforms deadbeats into racketeers. (Again, this has nothing to do with poor countries, but concerns a model where countries differ only by size.) Cap and trade also encourages more cooperative behavior from the most cooperative, so it is not all bad.[22] But it is hard to imagine why such a system would be considered desirable were its strategic implications understood.

Rich versus Poor

But what does game theory have to say about the rich-poor polarization that is making headlines? The logic of equation (7), which says that high-abatement countries will increase their targets, still holds. High abatement countries see a relatively high benefit from abatement, and trade allows them to purchase abatement at a lower cost. It does not matter why they view abatement as more beneficial. It may be because they are rich and have the luxury of being more concerned with the distant future, or because their country is large. In either case, when abatement becomes cheaper, they will choose to abate more.

Similarly, low-abatement countries may value climate benefits less or view the cost of abatement as greater because their utility of income is greater than that of richer countries. In either case, facing a global carbon price above their domestic public-goods valuation of abatement causes them to commit to less abatement.

---

[20] Pizer (2008) analyzes this same potential dynamic between Europe and the United States (playing the role of a low-abatement country) and concludes that Europe would not want to play. He asks, "Does the European Union really want to legitimize what it may believe is an inappropriately weak target in the United States? Not only legitimize, but allow the United states to actually make money off their weak choice?"

[21] "Rich-poor rifts stall progress at U.N. climate talks," Gerard Wynn, Bonn, Germany (Reuters), 31 May 2010.

[22] "Nevertheless, because it lowers mitigation costs, sectoral crediting might still indirectly help achieve more ambitious targets by encouraging Annex I countries to adopt more stringent objectives" (Berniaux et al. 2009).

So the same polarization result found for countries that differ by size also holds for countries that have varying abatement propensities caused by income or geographical differences. Cap and trade exacerbates polarization no matter what the initial cause of the difference. It makes cooperators more cooperative and it pays free riders to free-ride more aggressively. But, if the free-riding is caused by poverty and resentment at having to pay to clean up a problem caused by rich countries, the outcome may appear more equitable. However, as we will see shortly, equity concerns can be addressed in a much more productive fashion.

A Global Price Target

Cap and trade, as an international treaty, encompasses three fundamental flaws. It targets quantities instead of price. It requires individual targets instead of a single collective target. It ignores the necessary equity transfers—sometimes referred to as Green-Fund payments. Cap and trade also has benefits, but these have been cataloged frequently and will not be revisited here.

The use of quantity targets is problematic primarily because they make the second flaw—the use of individual targets—nearly inevitable. However, quantity targets also have the problem of increasing the abatement-cost risk. This is justified as necessary in order to reduce risk to the environment. But this argument fails on two counts. First the environmental risk imposed by price targets is short-run and hence of negligible importance for an environmental problem with a characteristic time scale of fifty to one hundred years. There is plenty of time for corrections, which will be needed under either price or quantity targeting. Second, the risks that quantity targets impose on potential treaty participants jeopardize cooperation to such an extent—both before and after agreement—that they increase abatement uncertainty. Moreover the risk premium associated with quantity targets raises the cost of abatement.

Technically, the use of a single collective target could be achieved as easily with a quantity target as with a price target, and the second flaw could be avoided. For example the agreement could specify that a single global cap would be selected and each country would be given a target equal to that cap times its share of the global population. This sort of approach has long been advocated by India, and has some theoretical justification. So far, even Europe has not seemed willing to discuss such a proposal and the chance of adoption by the United States would seem to be infinitesimal.

Other single-target quantity approaches might have broader appeal, but so far there is not even a candidate worthy of discussion.[23] With a single price target, the initial candidate is obvious to all. In fact its virtues are extolled even by the advocates of cap and trade. Each country should adopt the same price of abatement. This is too simplistic to work without an auxiliary policy. But with such a policy, equal-pricing becomes a true focal point for cooperation. The contrast with the quantity approach could not be more stark.

The third cap-and-trade flaw, concerning equity transfers, is generally denied, because cap-and-trade may provide financial transfers from large or rich countries to small or poor countries. But these transfers lack any systematic structure—"financial burdens among countries would be determined according to whatever determines such things in international negotiations" (Bradford 2008).[24] Also, they encourage free riding, or more accurately, profit seeking. In other words, cap-and-trade equity transfers discourage cooperation by developing countries.[25] This incentive needs to be reversed. But, after more than a decade of offering to pay for bad behavior, the transition to constructive incentives will not be easy.[26] It is, however, necessary.

---

[23] "A thorough review by Lecocq and Crassous (2003) shows country preferences on allocational rules to be very unstable and conditional on baselines assumptions and time horizons" (Hourcade et al. 2008).

[24] Here Bradford is describing his own cap-and-trade proposal, but except for Frankel's (2010) proposal, this is generally an accurate characterization of most cap-and-trade proposals. Barret and Toman (2010) make a similar point, "there is nothing approaching consensus today on an approach to this issue [income transfers] and no clear candidate solutions in sight."

[25] "[T]he actual experience under the CDM has had perverse effects in developing countries—rather than draw them into substantial limits on emissions it has, by contrast, rewarded them for avoiding exactly those commitments" (Victor and Wara 2008). "There is a concern that the CDM could create a perverse incentive not to introduce a sector-wide GHG-friendly policy that could otherwise make sense for the country, but would prohibit future CDM credits to be earned in that sector" (Bosi and Ellis 2005). "Crediting mechanisms such as the CDM may also reduce the incentives for non-Annex I countries to take on binding targets in the future" (Burniaux et al. 2009).

[26] "Another incentive problem is that the large financial inflows developing countries may benefit from under a future CDM may undermine their willingness to take on binding emission commitments at a later stage" (Berniaux et al. 2009).

We next present a simple price-target game. This works perfectly in an idealized world of countries that differ only by size. Although success in such a world may not be surprising, its sharp contrast with the cap-and-trade outcome provides encouragement. We then extend the price-target game to include a Green Fund coupled to the price target. This is analyzed for a world that includes low-emission countries that place a higher subjective value on both abatement costs and funds gained by selling abatement.

5. International Commitments and National Policies

Before developing the price-target approach, a brief diversion is required. The price target about to be discussed is an international commitment, not a national policy. In particular it is not a carbon tax, nor is it a requirement for a carbon tax. In fact, the original proposal, which is modeled here as a game, was designed specifically to allow compliance *without* the use of a carbon tax but with only cap-and-trade.

So far, we have considered two types of national commitments, to an abatement level and to a target. And, although we have focused on a system of national cap-and-trade policies, these commitments could, as well, be accomplished through carbon taxes, carbon caps, standards or subsidies. This flexibility helps make a strong national commitment more acceptable. A commitment to a price target is almost as flexible—it can be met with cap-and-trade, carbon taxes, or feebates, but it does not count subsidies (except for fossil subsidies, which are counted against carbon revenues).[27] We see this as an advantage.[28]

The discussion of international commitment is often accompanied by a tacit assumption that the form of the commitment will be reflected in national policy, for example, that a commitment to an international quantity cap will be fulfilled with a national cap-and-trade policy. In fact, although there is some advantage to such matching, actual practice indicates countries make extensive use of the flexibility provided by international commitments. The United States may abate more with fuel efficiency standards than with a cap-and-trade arrangement. Europe makes extensive use of "feed-in tariffs" and other subsidies, not to mention its $400-per-ton carbon tax on liquid fossil fuel.

But the key point regarding international commitments is that their presumed effects on national policy have been almost the sole focus of analysis, while their effects on the international agreement have been ignored. Yet, the problems caused by the choice of international commitment stems almost entirely from its impact on international negotiating strategies. To bring some balance to the discussion, this paper reverses the standard bias and focuses entirely on the neglected international aspects of international commitments.

6. Price-Target and Green-Fund Games

In the basic price-target game, there is one global price target, P^T, and all countries are required to price (not tax) carbon at that level. Such a system has been proposed by Cooper (2008) and by Cramton and Stoft (2009). In the latter proposal, countries that under-price can buy credits from countries that over-price in any given year, just as countries that under-achieve a quantity target can buy credits from those that over-achieve. This possibility will be discussed in the section on enforcement.

Having noted the possibility for such flexibility, we will simply assume, for the time being, as we did with cap and trade, that countries fully comply with their commitments one way or another. Because there is only one global target instead of many national targets, that target must be decided through a collective process, and that process is the price-target game. For simplicity we assume that each country names its preferred global price target, $P_j{}^T$, and that the lowest of all named targets is selected as the global price target. Selecting the lowest price means that no country is asked to sign a stronger treaty—with a higher price target—than it named. In practice, some countries, such as oil exporting countries, may name an unacceptably low price. This may be addressed by

[27] "A carbon price should be applied as widely as possible across the major emitting countries and sectors, preferably starting with the removal of fossil fuel subsidies" (Berniaux et al. 2009).

[28] In the United States the most costly "abatement" program has been ethanol subsidies, and they have likely resulted in negative abatement at the global level. For a global program, counting subsidies would be disastrous. See Hutchinson, et al. (2010).

setting the global price target at the named price level of, say, the 75th-percentile voter. Votes would be arranged in descending order and weighted by country emissions, so that at least 75 percent of emissions will be covered by countries willing to set a carbon price at least as high as the global price target.

The Price-Target Game

The first result concerns optimality of the price-target game in a world with countries that are identical except for size. This game is defined by the strategic variables, $P_j{}^T$, and the following equations:

$$P^T = \min\{P_j{}^T\} \tag{11}$$

$$C_j{}'(A_j) = P^T \tag{12}$$

$$NB_j = B_j(A) - C_j(A_j)\,,\ \text{where } A = \sum A_j \tag{13}$$

Note that P^T controls every A_j, so if a country's choice of $P_j{}^T$ matters to the outcome (if it is the lowest choice), then it controls every A_j. This is similar to the problem of finding the optimal global price P^O, which also controls A_j in every country. In fact, with countries that are identical except for size, finding the optimal global price is simply a scaled up version of the problem that every country faces in the price-target game. Hence the Nash equilibrium of this game is the optimal price, P^O, as will now be demonstrated.

When countries are identical except for size, each is related to the global economy by a scale factor, s_j. This means for each j, $B(A) = s_j\, B_j(A)$, $A = s_j\, A_j$, and $C(A) = s_j\, C_j(A_j)$. In the Nash equilibrium of the price-target game, countries maximize NB_j, which is equivalent to maximizing $s_j\, NB_j$, which equals $s_j\, B_j(A) - s_j\, C_j(A_j)$. But $s_j\, B_j(A) = B(A)$, which is just $\sum B_j(A)$. And $s_j\, C_j(A_j) = C(A)$, which is just $\sum C_j(A_j)$. So the price-target problem for each country, j, is the same as maximizing

$$NB = \sum\,[B_j(A) - C_j(A_j)], \tag{14}$$

which is just the condition for the optimal global price, as given by equation (2). Hence, each country chooses the optimal global price for $P_j{}^T$, and $P^T = P^O$, which proves Proposition 3.

Proposition 3. In the price-target game with counties that differ only in size, every country will favor the socially optimal value of the global price target, and the Nash equilibrium price will be optimal.

The problem is that countries differ in ways other than size. This is also a problem for any similar scheme based on a single global quantity target. And it is also a problem for schemes based on national targets. But the advantage of a global price target is that the simplest possible formula for national targets already corrects for country size and emission levels. A country with twice the emissions must collect twice the carbon revenue. It also handles growth appropriately. If India were to unexpectedly become as rich as the United States, any effective cap would create impossible problems, but a global price target would create none.

The Green-Fund Game

Still, the price-target game just analyzed would fail because poor countries would vote for a low global target. One solution to this problem is to link Green-Fund payments to the level of the global price target.[29] In this way poor countries can be given an incentive to choose a reasonably high target. The rule that will now be analyzed sets

$$\text{Green-Fund Payment} = G\cdot(e - e_j)\; n_j\, P^T \tag{15}$$

where e_j is the country's emissions per capita, e is the global average emissions per capita, n_j is country j's population, P^T is the global price target, and G is the Green-Fund parameter. In this version of the game, G will be set as part of the game's rules, but in practice it would be chosen in a political process. This will be discussed shortly.

Adding equation (15) to the net benefit formula defines the Green-Fund game, but its analysis requires a model of a world in which low-emission countries are low-price countries—meaning that in the public goods game

[29] "The first measure we consider to enhance the success of global climate treaties is transfers aimed at balancing strong asymmetries among the actors involved in climate change" (Eyckmans and Finus 2008).

they would choose a level of abatement corresponding to a below-average domestic carbon price. To this end we assume that low-emission-per-capita countries are poor and that poor countries see abatement costs and Green-Fund payments as magnified by a factor of e/e_j. In the context of a linear benefit function and a quadratic cost function, this leads to the following net benefit function:

$$NB_j = b_j A - (e/e_j)[\, c_j A_j^2 - G\cdot(e_j - \mathrm{e})\, n_j P^T\,]. \tag{16}$$

Although this model is too specialized to be viewed as realistic, it provides an informative contrast with cap-and-trade results which suffer from severe free-rider problems and polarization of the different effort levels that arise from both country size and country income.

To add a bit more realism to the exercise, we next assume that only low-emission countries perceive costs as magnified by e/e_j, and that countries with $e_j > e$ perceive costs normally. In other words the Green-Fund game is defined by equations (11) and (12), and by equation (16) for low-emission countries and by equation (16b) for high-emission countries:

$$NB_j = b_j A - c_j A_j^2 - G\cdot(e_j - \mathrm{e})\, n_j P^T. \tag{16b}$$

This makes it impossible for such a simple Green-Fund rule to produce an optimal outcome. Nonetheless, the outcome can be highly cooperative as shown in this example with three countries with emissions and per-capita-emissions roughly equal to those of the United States, China, and India. Table 3 shows the Nash equilibrium for a Green-Fund game with $G = 3.6\%$ (explained below) and the b_j and c_j values set so that, in the absence of cost magnification, the optimal carbon price is \$30/ton and the optimal abatement is 20 percent. Table 3 also shows the populations and per-capita emissions.[30]

Table 3. The outcome for a Three-Country World with a Simple Green Fund

	Stage 1			Stage 2					
	n_j	e_j	Voted P_j	P^T	A_j	Benefit	Cost	G. Fund	NB
Country	Billions	t/cap./yr	\$/ton	\$/ton	B ton	\$B/year			
U.S.	0.3	18	\$26.04		0.94	\$28	−\$12	−\$4	\$12.21
China	1.2	5	\$30.00	\$26.04	1.04	\$31	−\$14	\$0	\$17.69
India	1.0	1.1	\$26.04		0.19	\$6	−\$2	\$4	\$11.30
Total/Avg.	2.5	5			2.17	\$65	−\$28	\$0	\$41.19

For simplicity, this example ignores an advantage of the Green-Fund, which is that it provides an additional incentive for abatement. Because abatement reduces emissions per capita, it reduces the amount that high-*e* countries must pay, and increases the amount that low-*e* countries receive from the Green Fund. (In spite of this, the Fund remains revenue-neutral.) Although the simple rule for allocating Green-Fund payments cannot produce an optimal outcome, with the right choice of Green-Fund parameter (3.6% in this example) it does produce a nearly efficient outcome.

Besides abating emissions, the policy also results in a transfer of funds from high-emission to low-emission countries. Since low-emission countries have been assumed to benefit more from cost reductions and from income from the sale of abatement, this transfer increases global net benefit. Moreover, in spite of paying India double the cost of its abatement policy, the Green Fund charge is only \$0.95 per ton of emissions above the global per-capita average. That comes to 3.4 cents per person per day in the United States.

Setting the Green-Fund Parameter

Green-Fund payments are proportional to both the Green-Fund parameter, G, and P^T. All countries will likely agree that the Green-Fund parameter should not be too low. Poor countries want Green-Fund payments and rich

---

[30] Also the b_j and c_j values are set so that abatement benefits are proportional to total emissions and abatement costs are inversely proportional to total emission. The exact values used can be found in the solution spreadsheet at www.global-energy.org/lib/2010/10-07.

countries will want poor countries to vote for a higher price target. All countries will also agree that the Green-Fund parameter should not be too high. If it is, rich countries will refuse to vote for a high price target in order to avoid excessive Green Fund payments to poor countries. Nonetheless poor countries will favor a higher parameter value than rich countries, so there is room for conflict.

In the present example, the three countries would favor three different values of the Green Fund. Because China has average emissions per capita, it will neither pay nor receive Green-Fund payments, so it wishes to induce a price-target that is as close to socially optimal as possible. Hence it votes for the ideal green fund parameter, which results in P^* = \$26.04 as shown above. The United States favors a lower Green-Fund parameter, which would result in India voting for P^* = \$22.26 and setting the global target a bit further from the optimal value of \$30. India would, in turn, vote for a higher Green-Fund parameter, which would triple its net benefit, even though the United States would then hold P^* down to \$14.22.

This suggests that countries that are nearer to the mean value of emissions per capita should be given more weight in deciding the Green-Fund parameter. This is not difficult to accomplish.[31] But no matter what value is chosen, countries can still protect themselves by their choice of the global price target. Once the global price target is set in the non-cooperative voting game, we assume that the treaty is enforced.

7. Enforcement / Inducement

The public goods game requires no treaty and no enforcement to support its Nash equilibrium. However, the Nash equilibria of the other games all require treaties. And, in each case it is possible for countries to profit from violating the treaty, provided this does not cause the treaty to unravel.[32] In the present games, if prompt unraveling could be guaranteed, this in itself would provide enforcement, because the country that reneged would find itself worse off. This is a standard result for repeated prisoners'-dilemma and public-goods games.[33]

But prompt unraveling cannot be guaranteed, and in reality, some countries would prefer unraveling, hence some other form of treaty enforcement is necessary. This has been emphasized by economists (Barrett 2007; Kosfeld et al. 2009; Stiglitz 2006), but enforcement is often shunned by non-economists because the concept is associated with punishment of those who renege.[34] However, recent experimental evidence (Rand et al. 2009) suggests that rewards are at least as effective at inducing cooperation in a public-goods games, so inducement may be a more appropriate term than enforcement. Fortunately the Green-Fund game presents a no-cost opportunity to use rewards as inducements to cooperate, especially for low-emission countries.

Why Enforcement Matters

Signing a treaty obligates a country to fulfill a commitment that is generally not in its self interest. Otherwise, the treaty would not be needed. The treaty is signed only because others commit to doing something not in their strict self-interest but in the interest of the signer. Hence, a treaty is an agreement to trade commitments. As with any trade, each party is unwilling to step forward unless it is reasonably sure the other party will as well. But to think that enforcement simply prevents reneging is to miss two important benefits of enforcement. Parties that have every intention of complying will not want to commit to a strong treaty that is not enforced, because they fear

[31] A simple way to do this and preserve incentive compatibility is as follows. Sort participating countries by emissions per capita. Determine the percentile of the country closest to the mean emissions per capita. Let each country report its preferred Green-Fund parameter and sort these reports. Set the Green-Fund parameter equal to the report at the percentile equal to 100 minus the percentile determined from the mean emissions. If the distribution of emissions is not skewed, then this rule selects the median report. Selecting the report based on the specified percentile, rather than the median, accounts for any skewness in the distribution of emissions per capita.

[32] It may appear that the cap-and-trade game does not require enforcement, because countries will not find it in their self interest to reduce their self-selected quantity targets given a functioning global carbon market. However, they will find it in their self-interest to renege by allowing their domestic part of the global carbon market to malfunction, for example by failing to enforce the appropriate retirement of carbon permits.

[33] For a theory of self-enforcing treaties in dynamic games see Dutta and Radner (2004).

[34] Finus (2002) reports on a number of studies of non-compliance with international environmental agreements and concludes that the problem is nearly ubiquitous and that there is "a need to sanction non-compliance."

being taken advantage of. Hence the second reason for enforcement is to encourage commitment by cooperative parties.

The third benefit of enforcement is to discourage commitments by uncooperative parties. Without enforcement, such parties may sign on even though they know they may renege. Then they will later renege and this will tend to unravel the treaty. This can cause more damage than if they had never signed. In short, a treaty is simply a contract, and unenforced contracts are worth little.

What to Enforce

Proposals for a harmonized global carbon tax typically propose that countries commit to the global tax rate. But, just as with commitment to a quantity target, perfect compliance with the target cannot be expected even for countries with the best of intentions. Since punishment for a small miscalculation is counterproductive, a tolerance band is often proposed in such circumstances. But this leads players to target the lower limit of that band, and the problem of punishment for a small miscalculation is reproduced at the lower limit.

Cap-and-trade solves this problem by allowing countries to purchase carbon abatement credits at their fair market value. In effect there is no punishment for non-compliance with the national quantity commitment; there is simply a choice of whether to comply physically or financially. This eliminates the problem of disproportionate punishments for miscalculation. The option of financial compliance eliminates the need to differentiate between a mistake and deliberate reneging.

This same technique can and should be used with carbon pricing. In fact it is just as essential since countries may well price carbon with national cap-and-trade policies which are quite uncertain with respect to their revenue collection in any given year. Such a system of financial compliance has been proposed (Cramton and Stoft 2010, 2009; Stoft 2009). The global price target is multiplied by a country's emissions to find its carbon revenue target. If a country falls short of this revenue target, it must purchase carbon revenue credits from a central exchange. The cost of these credits would be adjusted so that the exchange remains revenue neutral and the countries are provided with a strong enough incentive that global carbon revenues remain on target. In other words, just as with carbon abatement credits, the price of carbon revenue credits would clear the market and enforce the global target—but not the individual national targets.

This is also the first case of using rewards to enforce the agreement. As with carbon-abatement credits, countries that collect more than their target carbon revenues will be able to sell revenue credits to the exchange and earn a monetary reward. Some countries may find carbon revenues a useful source of funds or they may find limiting oil dependence useful as many European countries have. In either case the opportunity to earn foreign exchange may induce compliance purely out of self-interest without regard for the benefits of maintaining the treaty.

But as noted, the requirement to purchase carbon revenue credits when a country falls short of the global price target does not enforce national targets. In this case revenue credits simply provide an alternative, and more efficient, form of compliance. But this flexible rule for compliance must itself be enforced.[35]

Enforcing Flexible Compliance

The question is then how to enforce national compliance with the flexible rule. Fortunately, compliance of the lowest-emission countries can be nearly guaranteed simply by linking Green-Fund payments to compliance. These countries will profit by complying. Again we have enforcement with rewards.

Compliance by average- and high-emission countries will require the threat of some unavoidable penalty. But before considering that, it should be remembered that what will be enforced is a treaty that has been signed voluntarily because it is in the best interest of each signer to have the treaty enforced rather than have no treaty at all. This is the result of designing a treaty that includes enforcement and then having countries state how strong they prefer the treaty to be and then sign a treaty that is no stronger. Although countries will still find it profitable to renege if they can do this without harming the treaty, they will realize that reneging might, in fact, cause the

[35] As explained by Cramton and Stoft (2009), revenue credits will likely cost far less than the value of revenue they represent. This is because paying another nation is far more costly than collecting and keeping domestic revenues. Hence even collecting 10 percent of the revenue target would likely generate enough funds for financial compliance.

treaty to unravel. In this case they would find reneging unprofitable. As a consequence the treaty should be reasonably easy to enforce.

In spite of this optimistic assessment, it is still important that enforcement be needed only rarely, for punishments always strain an organization. It is also important that enforcement be successful since failure sends a signal that the treaty is not enforceable. Both points argue for a particularly strong mechanism. If the penalty is strong enough, it will succeed. And if players know it will succeed, the penalty will never be assessed, but nonetheless will support the cooperative equilibrium.

A penalty must be strong for another reason as well. There is always some chance that the penalty will not be applied. If the penalty for stealing $100 is a $100 fine, and there is some chance of not being caught, then stealing becomes a rational policy. Since there will always be measurement ambiguities and procedural difficulties with any global carbon policy, there will always be some chance of avoiding enforcement. For this reason the penalty should be significantly larger than the gain from reneging.

Fortunately, as Stiglitz (2006) has pointed out, trade sanctions for the enforcement of environmental policies have been approved by the WTO. In a case against the U.S., the WTO court made a point of emphasizing the right of the United States to enforcing its own environmental policy with a ban on imported shrimp from countries that did not utilize U.S.-style turtle-excluder devices. An international climate treaty would have a significantly greater claim to a right to apply trade sanctions for enforcement.

8. Conclusion

Climate negotiators recognize that, at root, the climate problem is a free-rider problem. Each country would prefer that the others do more. But the free-rider problem is not immutable, and that has been overlooked. The design of agreements can exacerbate it or alleviate it. Instead, the implicit assumption has been that it must be overcome by moral suasion and scientific argument, both of which are weak instruments when dealing with national self-interest. Because of this oversight, both Kyoto and Copenhagen negotiations followed a global cap-and-trade framework, which has proven ineffective. We find that adopting cap-and-trade rules polarizes the free-rider incentive and discourages cooperation.

Attempting to convince developing countries to adopt binding emission targets is made more difficult by specifying that the agreement will entail global cap and trade. Without an agreement, countries (and even sub-national regions) are willing to abate to an extent that varies greatly. Creating a global market and allowing trade reduces the marginal cost of abatement for those that would otherwise abate most, but raises it for those that would abate least. The result is that high-abaters raise target abatement levels and low-abaters, lower them, often to negative values.

This polarization prevents agreement on the national targets that must accompany an operational cap-and-trade agreement. Low-abaters resent being asked to do more than their Nash strategies, while high-abaters reject the rational choices of low-abaters to target even less than they would abate without cap and trade.

This diagnosis provides a ray of hope. The Copenhagen summit failed, not because cooperation is beyond reach, but because the proposed agreement was reaching in the wrong direction. The agreement's structure was adopted without any consideration for its impact on cooperation. Instead efficiency and certainty were the goals.

To enhance cooperation, the agreement should use a single global target instead of reproducing the free-rider problem with national targets. The target should be a global carbon price, because there is an agreed best relationship between national prices and the global target—they should be equal. There is no hint of such an agreed relationship for a global quantity target. The agreement should include equity transfers to compensate low-emitting poor countries for meeting the global target. These transfers should reinforce cooperation with the price target rather than make free-riding more profitable.

Such a design, in contrast to the mandatory cap-and-trade design, mitigates free-rider incentives and creates a game with an equilibrium that is highly cooperative. And, it allows countries to choose cap-and-trade if they wish. International cooperation is the only way to solve a global public goods problem, and the only way to achieve cooperation is to design for cooperation.

Appendix

The Global Carbon Price under Cap and Trade

Proposition 1. In the Classic global cap-and-trade game, the global price, P^*, equals $1/N$ times the optimal price if the marginal value of abatement is constant, and is less than the average public-goods price if marginal abatement decreases with abatement and cap-and-trade serves to increase global abatement.

Proof: Begin with the definition of net benefit in the cap-and-trade game:

$$NB_j = B_j(A) - C_j(A_j) - P\cdot(T_j - A_j) \tag{A1}$$

Global abatement, the sum of A_j, is denoted by A. The first-order Nash condition requires that the derivative of NB_j with respect to T_j be zero, when other T_k, for $k \neq j$, are held constant. But since other T_k are constant, $\mathrm{d}A/\mathrm{d}T_j = 1$, and we have

$$\mathrm{d}B_j/\mathrm{d}A + \frac{d}{dA_j}[P^*A_j - C_j(A_j)]\frac{dAj}{dT_j} - \mathrm{d}P/\mathrm{d}A\cdot(T_j - A_j) - P^* = 0, \tag{A2}$$

where P^* is the Nash-equilibrium price. Because the carbon market sets price equal to marginal cost, the derivative of the term in square brackets is zero. Defining $B_j' = \mathrm{d}B_j/\mathrm{d}A$, and $P' = \mathrm{d}P/\mathrm{d}A$, gives

$$B_j'(A) - P^*(A) - P'(A)\cdot(T_j - A_j) = 0. \tag{A3}$$

Now, sum over all countries and divide by N. Because targets are met, T_j and A_j both sum to A, and the third term in (A3) sums to zero. So,

$$P^* = (\textstyle\sum B_j'(A^*))/N. \tag{A4}$$

Note that $B_j'(A^*)$ the cap-and-trade equilibrium abatement level, A^*, will differ from the public-goods and optimal levels of global abatement. However, if $B_j'(A)$ is constant this does not affect comparisons of P^* with the average public goods and optimal prices. However, a more realistic assumption holds that the marginal benefit of abatement, $B_j'(A)$, decreases as abatement increases. These observations lead to two conclusions.

By equation (3), $P^O = \sum B_j' = \sum P_j^G$, where P^O is the optimal global price. So

If the marginal benefit of abatement is constant:

$$P^* = P^O/N = (\textstyle\sum P_j^G)/N. \tag{A5}$$

However if cap-and-trade increases global abatement, and marginal benefit of abatement decreases with abatement,

$$P^O/N < P^* < (\textstyle\sum P_j^G)/N. \tag{A5'}$$

Polarization

Proposition 2. If the marginal benefit of abatement is constant and the marginal cost of abatement is increasing (quadratic), then a country with a domestic carbon price that is higher (lower) in the public-goods game than the cap-and-trade price will set a cap-and-trade target that is higher (lower) than its public-goods abatement.

Proof:

Lemma 1. If the marginal benefit of abatement is constant and the marginal cost of abatement is increasing, then a country with $P_j^G > P^*$ will set a cap-and-trade target $T_j > A_j^G$, its abatement in the public-goods game.

First, note that all derivatives are evaluated at the cap-and-trade Nash equilibrium. Since $\mathrm{d}A/\mathrm{d}T_j = 1$, $\mathrm{d}P/\mathrm{d}T_j = \mathrm{d}P/\mathrm{d}A$. Also, $\mathrm{d}B_j/\mathrm{d}T_j$ can be interpreted as the carbon price in country *j* in the public-goods game and denoted by P_j^G. The Nash condition, (A3), for country *j* can then be re-written as

$$P_j{}^G - P^* = (\mathrm{d}P/\mathrm{d}A)(T_j - A_j). \quad \text{(A6)}$$

This can be re-written as

$$T_j = A_j(P^*) + (P_j{}^G - P^*)\,(\mathrm{d}A/\mathrm{d}P). \quad \text{(A7)}$$

Next, assume that the slope of $A_j(P)$ decreases as price increases, which is the same as assuming that the marginal cost of abatement increases with abatement. This tells us that the average slope of $A_j(P)$ over a range of prices is less than the slope at the lowest price in the range. In other words, for a high-price country, with $P^G > P^*$,

$$\frac{A_j(P^G) - A_j(P^*)}{P^G - P^*} < \left.\frac{dA_j}{dP}\right|_{P^*}. \quad \text{(A8)}$$

But $\mathrm{d}A_j/\mathrm{d}P$ is certainly less than $\mathrm{d}A/\mathrm{d}P$, since other countries besides *j* will respond to changes in the global price of carbon. Hence we have

$$\frac{A_j(P^G) - A_j(P^*)}{P^G - P^*} < \left.\frac{dA}{dP}\right|_{P^*}. \quad \text{(A9)}$$

which rearranges to

$$A_j{}^G < A_j(P^*) + (P^G - P^*)\,(\mathrm{d}A/\mathrm{d}P). \quad \text{(A10)}$$

And comparing (A7) to (A10) shows that, for high-price countries,

$$T_j > A_j{}^G. \quad \text{(A11)}$$

Lemma 2. If the marginal benefit of abatement is constant and the marginal cost of abatement is quadratic, then a country with $P^G < P^*$ will set a cap-and-trade target $T_j < A_j{}^G$, its abatement in the public-goods game.

A different approach is needed for the case where $P_j{}^G < P^*$, and we will use a proof by contradiction. Hence we begin by assuming that $T_j > A_j{}^G$. All variables that pertain only to the public-goods equilibrium will be denoted with a superscript *G*. Net benefit in the cap-and-trade game is given by

$$NB_j = b_j \mathrm{A} - c_j A_j^2 - P\cdot(T_j - A_j), \quad \text{(A12)}$$

and (letting prime denote $\mathrm{d}/\mathrm{d}T_j$) its derivative with respect to T_j is

$$NB_j' = b_j - 2c_j A_j A_j' - P'\cdot(T_j - A_j) - P\cdot(1 - A_j'), \quad \text{(A13)}$$

We wish to evaluate this at the cap-and-trade Nash equilibrium, and show that it is not zero, which is a contradiction of the Nash condition. At the Nash equilibrium $P = P^* = 2c_j A_j$ and the r.h.s. simplifies to

$$NB_j' = b_j - P'\cdot(T_j - A_j) - P^*. \quad \text{(A14)}$$

In the public-goods Nash equilibrium, $b_j - P_j{}^G = 0$, and we subtract this from the r.h.s. to find:

$$NB_j' = P'\cdot(A_j - T_j) - P^* + P_j{}^G. \quad \text{(A15)}$$

Also, in the public-goods Nash equilibrium $P_j{}^G = 2c_j A_j{}^G$, so

$$NB_j' = P'\cdot(A_j - T_j) - 2c_j(A_j - A_j{}^G). \quad \text{(A16)}$$

Now define $\Delta A_j = A_j - A_j{}^G$. Since $P^* > P^G$ and abatement increases with price, $\Delta A_j > 0$, which will be used later. Also, by assumption $T_j > A_j{}^G$, so $A_j - T_j < A_j - A_j{}^G$, which shows that

$$NB_j' < P'\cdot(A_j - A_j{}^G) - 2c_j\Delta A_j, \text{ or}$$

$$NB_j' < P'\cdot\Delta A_j - 2c_j\Delta A_j. \quad \text{(A17)}$$

Since $P' = 2c_j A_j'$,

$$NB_j' < 2c_j(A_j' - 1)\Delta A_j. \quad \text{(A18)}$$

But $A_j' = \mathrm{d}A_j/\mathrm{d}T_j < \mathrm{d}A/\mathrm{d}T_j = 1$, and from above , $\Delta A_j > 0$. So

$$NB_j' < 2c_j(A_j' - 1)\Delta A_j < 0. \tag{A19}$$

This is the contradiction we were seeking since we have evaluated NB_j' at the Nash equilibrium and know it must be zero at that point. Hence our assumption that $T_j > A_j^G$ is false. This proves Lemma 2, and combining this with Lemma 1 proves Proposition 2.

References

Baron, Richard, Barbara Buchner, and Jane Ellis (2009) "Sectoral Approaches and the Carbon Market," OECD/IEA information paper.

Barrett, Scott (2003) *Environment and Statecraft: The Strategy of Environmental Treaty-Making*, Oxford Univ. Press, London.

Barrett, Scott (2007) "Proposal for a New Climate Change Treaty System," *The Economists' Voice*, 4:3, Article 6. Available at: http://www.bepress.com/ev/vol4/iss3/art6.

Barrett, Scott and Michael Toman (2010) "Contrasting Future Paths for an Evolving Global Climate Regime," *Global Policy*, 1:1 , 64-74.

Burniaux, Jean-Marc, Jean Chateau, Rob Dellink, Romain Duval, and Stephanie Jamet (2009) "The Economics of Climate Change Mitigation: How to Build the Necessary Global Action in a Cost-Effective Manner," OECD Economics Department Working Papers, No. 701, OECD.

Blair, Tony (2009) "Doing the Deal: Key Elements for a Copenhagen Climate Agreement," Climate Group, Office of Tony Blair. Available at http://www.theclimategroup.org/_assets/files/BTCD---Doing-the-Deal-13-10-09-Final_1.pdf.

Bosi, Martina and Jane Ellis (2005) "Exploring Options for 'Sectoral Crediting Mechanisms'," OECD/IEA information paper.

Bradford, David F. (2008) "Improving on Kyoto: Greenhouse Gas Control as the Purchase of a Global Public Good," Chapter 2, in Roger Guesnerie and Henry Tulkens, eds., *The Design of Climate Policy*, MIT Press.

Carbone, Jared C., Carsten Helm, and Thomas F. Rutherford (2009) "The Case for International Emission Trade in the Absence of Cooperative Climate Policy," *Journal of Environmental Economics and Management*, 58, 266–280.

Chander, Parkash and Henry Tulkens (2008) "Cooperation, Stability, and Self-enforcement in International Environmental Agreements: A Conceptual Discussion," Chapter 8, in Roger Guesnerie and Henry Tulkens, eds., *The Design of Climate Policy*, MIT Press.

Cooper, Richard N. (2008) "The Case for Charges on Greenhouse Gas Emissions," Harvard Project on International Climate Agreements, Discussion Paper 08-10.

Cramton, Peter and Steven Stoft (2010) "Price Is a Better Climate Commitment," *The Economists' Voice*, 7:1, Article 3. Available at: http://www.bepress.com/ev/vol7/iss1/art3.

Cramton, Peter and Steven Stoft (2009) "Global Carbon Pricing: A Better Climate Commitment," Global Energy Policy Center Research Paper, 09-06. Available at: http://www.global-energy.org/lib/.

Dutta, Prajit K. and Roy Radner (2004) "Self-Enforcing Climate-Change Treaties," *Proceedings of the National Academy of Sciences*, USA, 101-14, 5174-5179. Available at: http://www.pnas.org/content/101/14/5174.

Eyckmans, Johan and Michael Finus (2008) "Transfer Schemes and Institutional Changes," Chapter 6, in Roger Guesnerie and Henry Tulkens, eds., *The Design of Climate Policy*, MIT Press.

Finus, Michael (2002) "Game Theory and International Cooperation: Any Practical Application?," in Christoph Bohringer, Michael Finus, and Carsten Vogt, eds., *Controlling Global Warming: Perspectives from Economics, Game Theory and Public Choice*, Chapter 2, Edward Elgar Publishing, Inc. Available at: http://books.google.com/books?id=4aLcob3DSDoC.

Frankel, Jeffrey (2010) "An Elaborated Proposal for a Global Climate Policy Architecture: Specific Formulas and Emission Targets for All Countries in all Decades," in Joseph E. Aldy and Robert N. Stavins, eds., *Post-Kyoto International Climate Policy: Implementing Architectures for Agreement*, Cambridge University Press.

Godal, Odd and Bjart J. Holtsmark (2010) "International Emissions Trading and Endogenous Taxes," Discussion Papers 626, Research Department of Statistics Norway.

Godal, Odd and Frode Meland (2010) "Permit Markets, Seller Cartels and the Impact of Strategic Buyers," *B.E. Journal of Economic Analysis & Policy*, 10: 1, Article 29. Available at: http://www.bepress.com/bejeap/vol10/iss1/art29.

Helm, Carsten (2003) "International Emissions Trading with Endogenous Allowance Choices," *Journal of Public Economics*, 87, 2737–2747.

Holtsmark, Bjart J. and Dag Einar Sommervoll (2008) "International Emissions Trading in a Non-cooperative Equilibrium," Discussion Papers 542, Research Department of Statistics Norway.

Hourcade, Jean-Charles, P. R. Shukla and Sandrine Mathy (2008) "Untying the Climate-Development Gordian Knot," Chapter 5, in Roger Guesnerie and Henry Tulkens, eds., *The Design of Climate Policy*, MIT Press.

Hutchinson, Emma, Peter W. Kennedy, and Cristina Martinez (2010) "Subsidies for the Production of Cleaner Energy: When Do They Cause Emissions to Rise?" *B.E. Journal of Economic Analysis & Policy*, 10: 1, Article 28. Available at: http://www.bepress.com/bejeap/vol10/iss1/art28.

Kosfeld, Michael, Akira Okada, and Arno Riedl (2009) "Institution Formation in Public Goods Games," *American Economic Review*, 99:4, 1335–1355.

Nordhaus, William D. (2008) *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press.

Pizer, William A. (2008) "Economics versus Climate Change," Chapter 10, in Roger Guesnerie and Henry Tulkens, eds., *The Design of Climate Policy*, MIT Press.

Rand, David G., Anna Dreber, Tore Ellingsen, Drew Fudenberg, Martin A. Nowak (2009) "Positive Interactions Promote Public Cooperation," *Science*, 4 September, Vol. 325.

Schelling, Thomas C. (2002) "What Makes Greenhouse Sense?" *Foreign Affairs*, 81:3, May/June.

Schneider, Lambert (2008) "A Clean Development Mechanism (CDM) with Atmospheric Benefits for a Post-2012 Climate Regime," Öko-Institut Discussion Paper, 25 September.

Schneider, Lambert and Martin Cames (2009) "A Framework for a Sectoral Crediting Mechanism in a Post-2012 Climate Regime," Report for the Global Wind Energy Council. Öko-Institut, May.

Stiglitz, Joseph E. (2006) "A New Agenda for Global Warming," *The Economists' Voice* 3:7, Article 3. Available at: http://www.bepress.com/ev/vol3/iss7/art3.

Stoft, Steven E. (2009) "Flexible Global Carbon Pricing: A Backward-Compatible Upgrade for the Kyoto Protocol," European Union Institute Working Papers, RSCAS 2009/35. Available at: http://papers.ssrn.com/abstract=1438182.

Tol, Richard S. J. (2008) "Economics versus Climate Change: A Comment," Chapter 11, in Guesnerie, Roger and Henry Tulkens, eds., *The Design of Climate Policy*, MIT Press.

Wara, Michael and David Victor (2008) "A Realistic Policy on International Carbon Offsets," Program on Energy and Sustainable Development, Stanford University, Working paper 74, April.