

4 Global Carbon Pricing

Peter Cramton, Axel Ockenfels, and Steven Stoff*

Why Is Cooperation So Important?¹

Narrow Self-Interest

Do we need cooperation? The 2°C goal, or even more ambitious goals, could be met purely due to national self-interests if technological progress is fast enough—if the price of nonfossil energy falls quickly enough relative to the price of fossil energy—and if countries would rationally address the domestic damage being done by greenhouse gas emissions. It has been suggested that this will be the case, and so the only real problem is getting countries, such as the United States, to recognize and act in their self-interest (for a discussion of this view, see Cooper et al., chapter 1, this volume; for an example, see Green, 2015).

Of course, it would be convenient if technology is about to allow humankind to achieve its goals through pure self-interest. Such a technology miracle would imply that no transfers from rich to poor will be needed (although they would still be laudable). Also, no additional round of negotiations would be necessary to ratchet up agreements. The reason is that there is no need for agreements to get countries to do what is in their narrow self-interest. International education might be necessary, but we do not make international agreements to provide city parks or clean drinking water, public goods that present no international free-rider problem.

However, no one in Paris seemed to believe either of these conclusions, which would follow from the assumption of a timely technology rescue. Indeed, no one has offered any proof of such a miracle. Actual

* Ockenfels thanks the German Research Foundation (DFG) for financial support through the Research Unit “Design & Behavior” (FOR 1371).

developments, plans, and behaviors suggest that there is no relief in sight, but that we should rather expect the opposite: a continuing global “coal renaissance” together with increasing greenhouse gas emissions (Covert et al., 2016; Edenhofer and Ockenfels, chapter 9, this volume; Edenhofer et al., 2016).

There are more reasons that the world’s nations will need to cooperatively take individual costly actions to achieve the greater collective benefit that will result from meeting (or at least coming closer to) the Paris objectives. It is true that CO₂ abatement policies to address local pollution can help mitigate global warming. But the damages from local pollution and global warming due to CO₂ emissions are additive. That is, if a country suffers a local negative externality equivalent to \$60/ton of CO₂ and the negative climate externality is \$50/ton, a carbon price of \$60/ton is not optimal. Instead, only a carbon price of \$110/ton would efficiently address both externalities. So, *even if* all countries abate to fully address domestic pollution, they would not sufficiently address climate change. This is independent of technological progress—unless and until self-interest takes us to a corner solution where the negative climate externality vanishes.

In other words, narrow self-interest in local pollution issues will always provide too little mitigation incentive by the exact amount of the negative climate externality at every point in time. Also note that CO₂ causes no local externalities whatsoever, so local incentives are helpful only to the extent that CO₂ continues to correlate with other pollutants such as particulates and SO₂. This means that technology may gravitate toward solutions, such as scrubbers and natural-gas substitution, which target the local pollutants and have a limited or even at times a negative impact on CO₂ mitigation. Furthermore, history indicates that local pollution has been significantly undermitigated, especially in the early stages of economic development. So even in the case where narrow self-interest should be sufficient, realism would suggest a cautious approach.

Besides climate, an additional international externality calls for cooperation on carbon abatement. In a global economy, unilateral abatement reduces a country’s competitiveness (and, by related mechanisms, often comes with carbon leakage effects), so that it is, in fact, not in a country’s self-interest to fully address local pollution. Regardless of how one looks at this: narrow self-interest is unlikely to solve the climate dilemma, although it can certainly provide climate benefits and help with kick-starting global

carbon pricing. In other words, self-interest can be part of the solution as well as the central problem.

Ambition and Aligned Self-Interest

Christiana Figueres called 2014 the Year of Climate Ambition. Ten thousand UNFCCC web pages tell us that ambition is essential for a strong agreement. The UNFCCC newsroom informs us that developed countries are expected to lead the global drive to raise ambition. Ambition is what we want. But how do we get it if narrow self-interest is not enough?

Elinor Ostrom, a political scientist, won the 2009 Nobel Prize in economics for a lifetime studying “common-pool resource dilemmas” (such as global warming). She worked in the field, analyzed a thousand field studies by others, did game-theory experiments, and developed her own theories. She never mentions ambition. Instead, in her report on climate policy to the World Bank (Ostrom, 2009), she says her research on collective action identifies a “necessary central core of trust and reciprocity.”

Indeed, reciprocity is what changes self-interest. I will do X for you if you do Y for me. It is not in your self-interest to hand \$20 to your local cleaners. and it is not in their self-interest to clean your coat. But if you want your coat cleaned and they tell you that would cost \$20, then you may well decide to change their self-interest and make them want to clean your coat. Or perhaps that’s too much money. So you may negotiate to see whether you can change their self-interest at a lower cost. You are good at changing the self-interest of others.

It’s trickier for a group to change its self-interests. That requires a special form of reciprocity, a *common* commitment. I will do X if all of you also do X. (Of course, X can be a rule that takes circumstances into account in the same way for all players.) It’s trickier still when there is no government to organize or enforce the common commitment. But we know it’s possible. Ostrom’s (2009) central point is that people can self-organize what she calls “self-governance” when there is no government to do it for them.²

All disciplines dealing with human cooperation find that the reciprocity of a common commitment—I will if you will—is the key principle underlying collective human cooperation. Ostrom goes on to note that, “Trust and reciprocity are mutually reinforcing. A decrease in either can generate a downward cascade leading to little or no cooperation” (Poteete et al., 2010). In other words, insufficient reciprocity will not lead to an “upward spiral of

ambition,”³ as is hoped for with “pledge-and-review” (the approach exercised in Paris). This is the crucial lesson for international climate policy after Paris: ambitious aspirations mean little, and trust cannot be legislated, but reciprocity can be designed into a treaty. If that design is effective, then trust will follow, and then ambition.

This chapter, and this book, is about how to design a climate treaty that builds on reciprocity.

The Paris approach omitted built-in reciprocity. So it leaves out Ostrom’s (2009) “necessary central core.” The consequence is that the “Intended Nationally Determined Contributions” (INDCs) submitted in Paris are far from being ambitious enough to solve the climate dilemma, and there is no hope that, without a reciprocal, common commitment, pledge-and-review will succeed (Cooper et al., chapter 1, this volume). This is one conclusion of every chapter in this book. Instead of ambition, the chapters collected here focus on designs that, when put into practice, produce reciprocity.

What Ostrom (2009) and many others find in every successful cooperative system are rules governing everyone’s contributions to, and use of, the common resource, as well as penalties for breaking those rules. These rules and penalties are the reciprocity mechanisms, and they build trust. Exhortation to be ambitious is helpful but not enough. Common rules and sanctions for breaking them are required in real-world settings.⁴

Free-Riding and Cooperation

The atmosphere is a common-pool resource, a type of public good, and so it is susceptible to overuse. It’s a *global* public good. But imagine for a moment that it was not. Imagine that the United States could only damage its own climate and the same for other countries. What would change? We would still need climate science. But there would be no reason for international negotiations. There would be no reason for any other country to be upset with US or Chinese emissions because they would not affect anyone else. Domestic pollution and domestic climate change could be addressed fully by narrowly self-interested countries.

So, the reason for international negotiations is the public-goods nature of the atmosphere. The essence of a common-pool resource is that everyone has access to the common pool, and hence all can overuse it to their own

benefit and to the detriment of others. This is the definition of free-riding. So the only reason that international negotiations over ambition levels are needed, or make sense, is because of free-riding.

Deniers of Free-Riding Sometimes we have noticed climate advocates denying the importance of free-riding. One reason given is that climate deniers make use of the free-riding argument. But your opponent's use of a fact does not make the fact wrong.

Another reason for dismissing the importance of free-riders is a simple lack of understanding. One highly placed climate policy expert has argued that free-riding means doing little and expecting to "benefit sufficiently from other countries' mitigation." The expert then pointed out that there is currently not much action from which to benefit. In other words, if there are few actions to free-ride on, then free riding can't be a big problem.

This is backward. The main characteristic of the most severe free-rider problem is that when we all try to free-ride, there is no one left to give rides. In the most severe version of the free-rider problem, there is no free-riding!

But this confusion runs even deeper. The destruction of trust is the most pernicious aspect of free-riding. It causes those with no inclination to free-ride to act just like free-riders. This is the insidious mechanism that causes the unraveling of cooperation. Here's one way that could happen.

One Bad Apple Spoils the Bunch Consider 10 equal countries, nine of which do not want to free-ride but also don't want to be taken advantage of. However, one nice but poor country (a "classic" free-rider) will not cooperate even if others do their part. Also suppose that:

- One country will tolerate no defecting (free-riding) countries
- One will tolerate one defector
- One will tolerate two defectors
- And so on down to the most tolerant country that will tolerate eight defectors.

What happens? Well, the classic free-rider country defects, so the zero-tolerance country defects, so two have now defected, and the country that

will only tolerate one defector defects, and so on down to complete unravelling. Although only 1 in 10 was a free-rider in the classic sense of wanting to do nothing, disaster ensued. Also notice that, in the end, no one had anyone to free-ride on, even though free-riding was what caused the whole problem.

This example has only one equilibrium—no cooperation. A common commitment by itself will not fix this version of the problem because the poor country would not be willing to sign any commitment that involves ambitious contributions by everyone. But a common commitment that included a green climate-fund payment to the poor country could well work and achieve total cooperation. So reciprocity could overcome free-riding.

In other examples, everyone will cooperate if enough others cooperate. But the world can still get stuck with no cooperation if there's no trust to start with. But then all it takes is a common commitment to serve as a coordinating mechanism to shift everyone from no cooperation to full cooperation.

Not Being a Sucker Ostrom (2009) described another aspect of the problem in her climate report to the World Bank. “When participants fear they are being ‘suckers’ for taking costly actions while others free-ride, more substantial effort is devoted to finding deceptive ways of appearing to reduce emissions while not doing so.” This is what pledge-and-review will lead to. The free-rider problem is so essential that at the start of her classic book, *Governing the Commons*, she defines her “central question” as finding out “how a group ... can organize and govern themselves to obtain continuing joint benefits when all face temptations to free-ride.”

These are not theoretical judgments. For example, as Ostrom (2009) reported, in Törbel, Switzerland, the common-commitment rule is “no citizen can send more cows to the alp than he could feed during the winter,” and this is still enforced by “substantial fines for any attempt by villagers to appropriate a larger share of grazing rights.” Those two reciprocity mechanisms prevent free-riding and generate trust. There are numerous other field studies like this.

What the US Senate Understood During the Kyoto negotiations, developing countries said “you go first” to the others. The US Senate voted 95 to 0 to say, in effect, “No. You must come along with us. You can’t free-ride.” Even if this is not what the senators were thinking, there is a profoundly true message in that unprecedented vote. The message is, “Make a reciprocal deal we can trust.”

Some claim the senators’ expressions of concern were a cover for baser motives. But suppose that was true of all 95 senators. No one is more calculating and no one listens to the electorate better than professional politicians. When they calculate the same answer 95 to 0, you’d be a fool not to listen. This is what they understood:

Americans have a powerful and abiding fear that they will be taken advantage of in the international arena. They don’t *trust* “foreigners.” So do not depend on their altruistic cooperation. They want a reciprocal deal they can trust.

That is the message of the Senate’s vote on Kyoto. That is not just what the senators said, that is what they were betting their careers on.

In 2015, the main argument in the United States against the Iran nuclear treaty is, “You can’t trust Iran” or any of our partners in the negotiations—China, France, Russia, the United Kingdom, and Germany. Nothing has changed. But the United States is hardly alone in this. During the 2009 Copenhagen climate negotiations, China became convinced and angry that the United States, Europe, and other developed countries were actively trying to turn the developing world against it.

Similar issues came up in Paris. Japan’s Paris pledge has been attacked by the Climate Action Network, a network of more than 950 nongovernmental organizations (<http://www.climateactionnetwork.org/fossil-of-the-day>). They claim that Japan is “using smoke and mirrors (shifting baselines) to fake ambition.” That’s a claim of free-riding. China’s pledge is for slightly less than what they found they needed to do to curb domestic pollution (<http://climateparis.org/china-emissions-pledge>). When they announced it, they made no claim that it was motivated by concern for the climate. This may be free-riding or fear of it. But we are not blaming China; it’s just what one should expect from rational players who have no good reason to trust other players.

Those advocating national climate policies face climate-change deniers pointing out that other countries could free-ride on our efforts. In defense,

they tend to deny the free-rider problem and make up baseless theories, such as Al Gore's notion that, "If we lead, China will follow." In the long run, it will be far more constructive to solve the free-rider problem—the most central problem of cooperation—rather than deny there is a problem. Solving the free-rider problem is the heart of the solution proposed in this volume.

The Problem and the Solution

For 20 years, almost all climate negotiators have agreed on the need for strong climate-change mitigation. Even before Paris, there was a strong consensus that 2°C should be the goal. But this aspiration has not been translated into commitments and actions.

The Problem

The real problem is not the climate or the lack of climate-science knowledge, and it's not the lack of a common aspiration or goal—Paris achieved that. It's not even the lack of blueprints for global action. The trouble is that negotiations end in acrimony or hollow victory statements. Paris was not different in that respect. So the problem is to find and fix the cause of these negotiation failures.

A better approach to negotiation is needed, and so we have made "how to negotiate" the focus of this book. This focus requires a distinction often overlooked. Two things matter most to the success of a negotiation: what outcome you aim for and how you go about getting there. Everyone knows this, but it is easier to focus on what you want than on how to structure the negotiations. So the "how" part is usually ignored and almost never analyzed systematically. In fact, the "how" part is so important that Weitzman (2014, chapter 8, this volume), Cramton and Stoft (2012a, 2012b), and Cramton et al. (chapter 12, this volume) argue it is decisive. But "what" and "how" are interrelated, and that just adds to the tendency for the "how" part to be forgotten.

Consider free-riding. As discussed in the previous section, an agreement that makes free-riding attractive will break down. As Ostrom (2007, p. 201) explains, it will "generate a downward cascade leading to little or no cooperation." But she is not concerned with *how* to negotiate, so this is purely a consideration of *what* works. But if negotiators see that an agreement allows

free-riding, it is hard to negotiate a strong commitment within that framework. So the *potential* to free-ride—the “what” part—affects the negotiation process—the “how” part.

This is not a general principle. A treaty that will work poorly may be easy to negotiate because it demands little, whereas one that would solve the problem brilliantly may be impossible to negotiate because of coordination problems—parties can’t agree on who should play which role.

Also, an agreement must cover three separate areas—abatement, burden sharing, and enforcement. It must get the “what” and “how” parts right in each area. But to avoid being too ambitious, we will only briefly speak to enforcement and leave that question mostly to Nordhaus (chapter 7, this volume). It is important to note, however, that enforcement is not a substitute for a common commitment but rather a complement (see MacKay et al., 2015). This leaves the two closely related areas: abatement and burden sharing. The challenge is to find a treaty design that solves the free-rider problem for abatement and the fair burden-sharing problem, as well as to find a way to negotiate the two solutions that will lead to a strong treaty and not to a deadlock or weak commitments.

The Pledge-and-Review Non-Solution

Pledge-and-review was first proposed by Japan in a memo to the UNFCCC in 1991 and was much discussed starting in 1992. It was the model for the Paris Agreement. Intended Nationally Determined Contributions (INDCs) are the pledges, and these will be reviewed and, it is hoped, upgraded occasionally. These are individual commitments, not common commitments, and so they do not address free-riding. But the situation is actually much worse than this statement implies.

The Kyoto Protocol was based on individual commitments, too, but all countries committed to some percentage (not a common percentage) of emission reduction below their 1990 emissions level. These commitments provided some hope of comparability. But the Paris pledge-and-review commitments are quite free-form even for developed countries and entirely free-form for developing countries. Hence, they are essentially impossible to compare.

Being both individual and noncomparable opens the door wide to free-riding. Many countries will do their best to either lock in a free-ride or prevent others from free-riding on them. Either strategy results in weak

pledges. Hence, this approach will fail. For further discussion of actual outcomes in Paris, see Cooper (chapter 5, this volume) and the references cited therein. For a further general discussion of why individual commitments will fail to solve the climate dilemma, see Cramton et al. (chapter 12, this volume), Weitzman (chapter 8, this volume), and Gollier and Tirole (chapter 10, this volume).

Problems with Cap-and-Trade

Global cap-and-trade is likely the oldest of the three major approaches (global carbon pricing, pledge-and-review, global cap-and-trade), although the early (standard) versions were not global and worked quite differently. There are variations of global cap-and-trade, but for simplicity we will mostly adhere to the most up-to-date one, presented by Gollier and Tirole (chapter 10, this volume).

This approach first negotiates a global cap, Q , which is a quantity limit on emissions. However, no country is responsible for it, and at this stage, nothing has been decided about how to share responsibility for it. So what looks like simplicity has likely left us in a worse negotiating predicament than the one faced by the Kyoto negotiations. In Kyoto, countries just signed up for whatever abatement reduction they wanted. So naturally they reached agreement, just as they did in Paris with INDCs. However, that would not likely work under this global cap-and-trade approach because voluntary pledges probably would add up to more than the agreed Q . It is easy to be ambitious for the group when choosing Q , and again it is easy to allocate to your own country a lot of permits and explain that others should be taking fewer.

To solve this problem, it is necessary to agree on a formula for allocating any Q that the players select. Weitzman (chapter 8, this volume) explains why this is nearly impossible, and we examine the 20-year search for such formulas and find that little progress has been made. Moreover, because the formula is negotiated after Q is chosen, every country will evaluate the formula by computing its share of Q under the formula and judging the formula on that basis. This will make a successful negotiation even more difficult than choosing a formula first (as was tried unsuccessfully in Kyoto).

There are many reasons that agreeing on a formula is difficult. There is no agreeable focal point for a “fair” allocation of Q to all countries. Moreover,

with this mechanism, assuming that the global Q is actually credible, there is no built-in reciprocity that might help to discipline countries: the “I will do more if you do more” principle is not applicable because in the constant-sum-game that is being played when allocating a fixed Q to many countries, “I will actually do less if you do more.” The best strategy in this kind of game, even when players would in principle be willing to reciprocate others’ cooperation, is to be competitive (Bolton and Ockenfels, 2000). This makes it impossible to build trust. Of course, this in turn implies that no ambitious global Q can be made credible.

In addition, we will find that a strong global cap-and-trade would cause enormous trading risks. Concern over such risks will make negotiating a strong treaty even more difficult as poor countries seek large permit allocations for protection. In short, the chances of negotiating a global cap-and-trade agreement appear to have been receding for 20 years, and if a strong agreement were ever implemented, it seems likely to unravel due to unfair realized outcomes of trading risks.

The Solution: A Global Carbon Price Commitment

A global carbon price commitment evolved from the idea of a global carbon tax and (to our knowledge) was first published by Cooper (2008) and then Stoft (2008). One key feature in these publications was the idea that countries could commit to a minimum price and then meet that commitment with either standard cap-and-trade or fossil-fuel taxes. Compared with a uniform global tax on carbon, this allows tremendous flexibility, which is clearly necessary for political reasons.

When fleshed out, the proposal can be seen to be quite similar to global cap-and-trade but with quantities replaced by prices, as is shown in table 4.1.

Table 4.1

Comparing Global Commitments: Cap-and-Trade versus Carbon Pricing

Global Commitment	Target	Missed Target? Buy or Sell:	Allowed National Policies
Quantity cap, Q	$Q_i = ?$	Emission credits	Everything
Price, P	$P_i = P$	Pricing credits	Cap, tax, or similar

Notice that no rule for allocating global target Q to country i is provided. This missing part of the specification is one of the main criticisms of global cap-and-trade. Also notice that global cap-and-trade allows any national (or regional or local) abatement policy, and hence it does not require the pricing of carbon emissions, whereas global pricing provides a great deal, but not complete, flexibility and does require carbon pricing.

Besides flexibility, there is a need for burden sharing. This would appear to be crucial for any effective global climate agreement. Hence, the defining features of global carbon pricing can be summarized as follows:

1. Negotiate green-climate-fund payments, $G_i =$ some formula, for each country i .
2. Negotiate a global price-floor, $P(t)$, to be flexibly met by each country.

The price path, $P(t)$, is the “common responsibility” of all countries, whereas climate-fund payments, G_i , are their “differentiated responsibilities” under the UNFCCC. It is essential that P be a common commitment to solve the free-rider problem with trust and reciprocity as described by Ostrom. However, P could be a flexible commitment, provided all countries are granted the same flexibility. For example, if a country's price is too low in one year, then it could be allowed to buy pricing credits⁵ from another country that has overperformed. Also note that $P(t)$ should be adjusted every five years or so as ambitions increase.

Price is defined flexibly within a country as total carbon revenues divided by total emissions. Carbon revenues include the market price of freely allocated cap-and-trade permits because these credits price carbon just as effectively as a tax. Because price is defined in terms of carbon revenue, tax rates on fossil fuels can vary from fuel to fuel and between customer classes.

It is important to negotiate the green climate fund first because the results determine what common price countries will accept. If the climate fund is meager, then poor countries will demand a low carbon price, whereas if it is too generous, rich countries will attempt to stymie the price negotiations. Consequently, we propose that the generosity of the green fund be determined with only one goal in mind—to maximize the global carbon price. To do this, it must satisfy both rich and poor reasonably well so that the second phase, price negotiations, will succeed. To arrive at such a fair compromise, we suggest that the generosity of the climate fund be decided by a group of countries that are midway between donors and recipients. These

countries will care most about the success of the climate policy and worry least about their climate fund payments, either into or from the fund.

Summary This completes the sketch of the solution—the path to a strong and negotiable climate treaty. First, a common climate-fund formula is negotiated, guided by the goal of achieving the strongest possible climate commitment. This means it must be seen as reasonably fair by rich and poor alike. Then a common price commitment is negotiated, which prevents free-riding.⁶

A Simple Treaty to Change Self-Interest

When it is not kept in check, free-riding produces the tragedy of the commons. The Kyoto negotiations hoped to keep it in check with a uniform requirement (X%) for emission reduction. The search for a commonly acceptable X failed, but the idea of a common X was based on good intuition. So failure led to a nearly permanent loss of the intuition that a common commitment is needed. Global carbon pricing searches along a different path and finds the common commitment that eluded the Kyoto negotiators. This section explains the basic idea of that common commitment and how it defeats free-riding.

Few doubt that self-interest is a powerful force, and few believe we can redirect it to solve the tragedy of the commons that now controls climate change. This pessimism is a bit surprising considering that Elinor Ostrom's work (and the work of many others) focuses on explaining how communities, many of them poor, have been doing this for centuries. In this section, we sketch a simple climate treaty that demonstrates the basic principles of promoting cooperation. It begins with an independent-commitment game among 10 purely self-interested countries and argues that they would be trapped as expected by the tragedy of the commons. We then change one rule in the negotiation game and nothing else. The result is that, although the 10 selfish countries remain as selfish as ever, they cooperate perfectly. The new rule is a common price commitment.

A Simple Global Climate Model⁷

First, we need a simple model of the climate-policy world. Imagine that the world has 10 identical countries. If you weight countries by their size (so

the little ones don't count much), the average-size country has about one-tenth of the world's population. This turns out to be a good first model to use to find out how countries might cooperate or free-ride on each other's efforts.

Suppose that one ton of emissions does \$50 of climate damage to the whole world but only \$5 of damage to each country. So the world is better off if it stops emissions worth less (in nonclimate benefits) than \$50/ton, and an individual country is better off stopping emissions worth less than \$5/ton.

A carbon price of \$50/ton will stop emissions that cause a net loss to the world, but a single selfish country will only want to price its carbon at \$5/ton in order to allow local projects worth \$6/ton. These bring it a net benefit of \$1 but do \$45 of damage to the rest of the world. Remember that the tax itself does not cost the country anything, it just shifts money around. (From here on we will often drop the per-ton units.)

Two Climate Treaties

Enforcement. Can these 10 countries negotiate a strong treaty? The answer is "yes, if there is enforcement." So assume there will be enforcement but only of voluntary agreements. We won't force anyone to comply with an agreement they don't like. Imagine that if one country reneges, this enforcement will be carried out by the other nine countries (perhaps with trade sanctions). So if you voluntarily sign a deal to cut emissions by 20%, it will be enforced. But if you don't like the treaty, you can just not sign it, and then nothing will be enforced. That's pretty mild enforcement—saving face might even be a strong enough motive to accomplish this. For simplicity, we assume that all countries are selfish.

An Individual-Commitment Treaty In the spirit of the 2015 Paris conference, all countries could agree to the following Individualistic treaty:

Treaty: Each country will pledge to implement a carbon price of its choosing.

Of course, this will be enforced as discussed earlier. There would be long negotiations and discussions first, but nothing would be known for sure until the written pledges are opened simultaneously on the deadline.

Selfish countries would set a carbon price no *lower* than \$5 because that would allow emissions that directly caused them more climate damage

than the non-climate benefit they would gain. But self-interest will hold them down to that level.

So individual pledges made simultaneously will lead to a complete lack of cooperation. It takes reciprocity to prevent this outcome—and that's missing with the individualistic treaty.

Outcome: All countries would pledge \$5; that is, no country would cooperate to address climate change—the same outcome as if there were no negotiations.

A Common-Commitment Treaty. Suppose instead that the 10 countries ask their treaty-design team to invent a treaty and a way to negotiate it. The team reports back: Every country should pledge their “conditional price” with the understanding that they will only implement that price if all others pledge that much. Otherwise they will implement the lowest conditional price pledged by any country. With this treaty design, once the pledges are in, the lowest pledge becomes the common global price commitment. This will be enforced for all countries because all have voluntarily agreed to that particular price—the lowest “conditional price” that was pledged. So, as before, any country can pledge \$5 without any penalty, so it can fully defect.

Treaty: Every country agrees to price carbon at the lowest price pledged by any country.

With this treaty, all countries will pledge to price at \$50, so that the global carbon price is at its optimum. First, we check that any other outcome would make all countries worse off. Suppose the lowest price is lower, say \$40. Then companies in all countries would emit carbon that only had a \$41 benefit to them (\$1 more than the tax they pay). But with identical countries, all would do this, and again with 10 countries each emitting 1 ton, the damage is \$50 per country, but they only get \$41 in local benefit from burning the fossil fuel. All the countries would lose. Similarly, if the lowest price is higher than \$50, say \$60, then this would stop them from all using a ton of fossil fuel that had a \$59 local benefit. But such a project is good because, with 10 of these there is only \$50 of climate damage per country, and they all gain a net benefit by \$9 (a \$59 local benefit and \$50 climate cost). So a \$50 carbon price is as high as anyone would want.

To see why everyone pledges \$50, observe that one's pledge can only be relevant for the outcome if one has submitted the lowest price among all players. So if your pledge was lowest and it was less than \$50, then it would benefit everyone (including you) to raise it. If your pledge was lowest and above \$50, then it would benefit everyone to lower it. Just in case your pledge is the controlling minimum pledge, you will want to pledge \$50, which is optimal for all.

Outcome: All countries would pledge \$50, and that would be the global price of carbon—all countries would fully cooperate to optimally address climate change.

Conclusion

In a simple world with identical, completely selfish countries, and with enforcement of voluntary commitments, we can change the negotiation game and thereby change selfish *behavior*—even though the intention remains selfish. This can be done by introducing a common commitment into the negotiations. This changes the outcome from no cooperation to full cooperation. Note that the results in this section are robust: They hold when players are completely selfish as well as under weaker assumptions about players' motivations (e.g., even if most players are conditional cooperators and would be willing to match others' average contributions, there would be a complete lack of cooperation with an individualistic treaty as long as at least one player is at least "to some extent" selfish) (see Bolton and Ockenfels, 2000, for a more formal treatment). Note, however, that both a common commitment and some enforcement are required. After discussing some practicalities and fairness, we will return and discuss how to make this simple story more realistic.

Does Global Cap-and-Trade Price Carbon?

For economists, the central goal of global cap-and-trade, as well as standard, local cap-and-trade, has been economic efficiency—its ability to get the job done much more cheaply than traditional command-and-control approaches. It does this by (supposedly) putting a "uniform price" on carbon emissions. Standard cap-and-trade actually does do this—by requiring emitters to have emission permits.

But global cap-and-trade, which works like the Kyoto Protocol, will fail to achieve the central objective of cap-and-trade for the same reason that Kyoto did—emitters are *not* required to have permits for their emissions or to price carbon at all. Here we explain this essential flaw in global cap-and-trade and how a global carbon price commitment would avoid this.

Why Price Carbon Emissions?

Until recently, emitting carbon (by which we mean emitting CO₂ from fossil-fuel or certain other easy-to-track greenhouse gases) has generally been free. Disposing of carbon into the atmosphere, unlike taking your garbage to the dump, had a price of zero. But carbon emissions turned out to have a cost, which is increasing as the concentration of atmospheric carbon increases.

As with any good, when the price is too low, it is overused. However, burning carbon produces valuable services, and the damage from disposal may be only \$30 or \$40 per ton at present. Certainly, no one would suggest we immediately stop driving, heating, and using electricity. We must admit there is a tradeoff. Economics shows that if carbon is correctly priced and we are rational, we will make that tradeoff perfectly. The net benefit (value minus damage) will be maximized. The result won't actually be perfect, but to a reasonable approximation, it will maximize net benefit—carbon benefits minus climate damages.

Price Matters In 1972 and 1973, US CO₂ emissions *rose* by 4.6% and 4.2%, respectively. In October 1973, the Arabs declared an oil embargo, and oil prices rose sharply. In 1974 and 1975, CO₂ emissions *declined* by 3.5% and 4.2%, respectively. As figure 4.1 shows, at the end of 12 years—which is when the Saudis stopped propping up oil prices and took back their market share, emissions were down not up. Doubters claim this was caused by the introduction of nuclear power, but if all those plants had emitted as much CO₂ as coal plants, then emissions would have been only 9% higher and still would have been down not up from 1973 levels. In fact, even in this fictitious (no nuclear) scenario, they would have been down about 37% relative to a trend line based on GDP. Note also that the decline was not at all caused by a decline in GDP.

Prices work in an uncountable number of ways. For example, US refrigerators were made more efficient. But this was the direct result of work by the

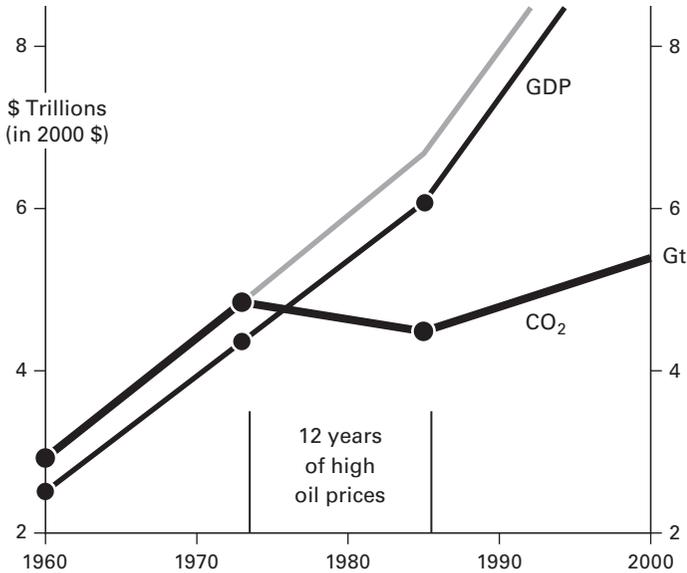


Figure 4.1
Oil price impact on US CO₂ emissions

brilliant particle physicist Art Rosenfeld, who explained, “I was prompted by the 1973 Organization of Petroleum Exporting Countries (OPEC) oil embargo to switch to improving energy end-use efficiency, particularly in buildings.” That was after he realized “why we in the United States used so much energy; oil and gas were as cheap as dirt or water, and so they were treated like dirt or water.” Art Rosenfeld’s brilliant work on energy efficiency was a direct result of higher carbon prices.

Of course prices also change individual consumer decisions, but it is important to remember their impact on politics, how cars are advertised, regulatory hearings, and the environmental movement. The impact is enormous, and most of it is long run so it is not immediately apparent.

Why Carbon Charges Are Cheap It is cheap to price carbon. For simplicity, think of a \$100 billion per year carbon tax. How much does that cost the country? The tax itself costs nothing. The money collected is not lost to the country and can either be returned directly to its citizens or used in place of some other tax—as a “tax swap.” As long as revenues are not returned in proportion to the tax collected, the carbon charge will still do its job. For

example, if a gasoline tax collects \$500 per person on average, and everyone is given a refund of \$500 regardless of their gasoline usage, every dollar of tax saved by buying less gasoline will be kept, and the incentive to use less is unaffected.

A direct and complete refund is the best way to prove the carbon charge costs nothing. It is also fair because it is equivalent to saying each person has an equal right to the atmosphere, and those who use more should have to buy extra rights from those who use less and not just take the rights for free (usually from the poor). Nonetheless, a tax swap will likely be politically more popular, and a swap also demonstrates that the tax itself is free (but only in the short run).

There is, however, an indirect and nearly invisible cost to pricing carbon. No matter how it is done or what is done with the revenues, pricing carbon will cause reduced emissions—abatement, and abatement is costly. It will cost somewhere between nothing and the price of carbon because there is no use in paying more—it's better to just pay the charge. So the standard formula for that cost is $P \times A/2$, where P is the carbon price and A is the amount abated. If a \$30 price reduces emissions from 1 billion tons to 0.8 billion, then $0.8 \text{ billion} \times \$30 = \$24 \text{ billion}$ will be collected in revenue. However, because abatement is only 0.2 billion, the cost of abatement will be only $\$30 \times 0.2/2 = \3 billion , eight times less than the carbon revenues collected (and recycled).

So to summarize, the carbon charge itself costs nothing. It just causes money to change hands. In contrast, the induced abatement does have a cost. However, people will be quite creative in figuring out how to minimize this cost and will consider far more possibilities than regulators possibly can. This is what makes carbon pricing much cheaper than regulatory subsidies (see Edenhofer and Ockenfels, chapter 9, this volume, for a similar argument regarding the German “Energiewende”).

A uniform global carbon price is certainly not a full solution to the climate problem, but it is by far the broadest, simplest, and most efficient (cheapest) partial solution. So it should be high, perhaps even highest, on the list of important policies to implement. This is the well-known, traditional economic argument for pricing carbon.

A New Reason to Price Carbon The point of this section is that global cap-and-trade fails at its mission—pricing carbon emissions. But we don't want

to give the impression that that is the only mission of global carbon pricing. The broken climate negotiations suggest an arguably more important reason to price carbon. As we just saw, free-riding and the fear of it have prevented the world from taking meaningful action to limit climate damage. To overcome this problem, we need a common commitment. As it turns out, global carbon pricing makes an ideal common commitment, and nothing else seems to work. So this is the new and probably most important reason to price carbon. We will return to this topic later.⁸

How Can We Price Carbon Emissions?

There are two well-known ways that governments can price carbon emissions: cap-and-trade and taxing fossil fuel. Although both of these could be operated as a global *policy*, it would require complex international institutions that presently seem quite improbable. Cap-and-trade would require that all emitters in all countries trade permits in one unified market. The European Union (EU) emissions trading scheme (ETS) is such a policy, but it only covers half of emissions and only exists because the EU has a government, which the world does not. A global tax would require that all fossil fuels be taxed at the same rate. This seems nearly as impossible as global cap-and-trade.

Consequently, it is far more realistic to consider simple global commitments instead of complex global policies. Global cap-and-trade only means committing to a set of country-specific permit allocations and to restricting emission to permits. A global price commitment only means national commitments to a global price. In both cases, countries would choose from such policies as domestic cap-and-trade, the EU ETS, various fossil fuel taxes, and other pricing policies. None of these policies would be governed from the top.

The Kyoto Protocol is a model for global cap-and-trade. It covers a broader territory than the EU, and so the Protocol has no corresponding government. A similar protocol could potentially support a global cap-and-trade treaty. But the Kyoto Protocol is a form of cap-and-trade that *does not price carbon emissions*, and neither would global cap-and-trade. We explain this next.

Global Cap-and-Trade The Kyoto Protocol implements a small version of global cap-and-trade, but it does not implement anything like the EU ETS,

California's cap-and-trade, or Regional Greenhouse Gas Initiative (RGGI) in the eastern United States. All these markets require emitters to own carbon permits. They all price carbon emissions. The Kyoto Protocol does not.

Under the Kyoto Protocol, governments, not emitters, must own all the carbon permits for their whole country even if the government were to emit nothing. This creates a disconnection. Kyoto's international carbon permits are called Assigned Amount Units (AAUs). If the AAU market worked (which it does not) and priced AAUs at, say, \$30/ton, then that would not mean that any emitter would be charged \$30 for one ton of carbon emissions. The government must curb emissions, but it can do that however it pleases. It is not required to price carbon. Of course it is allowed to price carbon, and it might do so. But if the EU ETS is any guide, it will not price carbon emissions at the price of international carbon permits—the AAU price.

Suppose a government has 1.2 billion AAUs and its country is emitting only 1 billion tons of carbon. It can sell its 0.2 billion AAUs at the global market price of, say, \$30/ton, and it doesn't need to cut back on its emissions at all. Therefore, it has absolutely no need to price carbon. So why is there an almost universal pretense that global cap-and-trade would price all carbon emissions at a uniform AAU price?

In effect, the following is the economic story behind this pretense. The government will freely choose to tax carbon usage at \$30 per ton (the AAU price). It's a clever tax because it will cause the country to emit less than 1 billion tons, and this will free up more permits, say 0.1 billion more, to sell to other countries. Now the government can sell a total of 0.3 billion permits on the international market for \$9 billion instead of selling 0.2 billion permits for \$6 billion.

It's a nice theory. But for the government, there is a cost. It must impose a potentially unpopular carbon tax (or permit requirement) that collects \$30 times 0.9 billion tons, or \$27 billion domestically. Some will ask, why? Especially when the country has more than enough permits to start with.

It didn't work like the economists' theory predicted under the Kyoto Protocol. Russia and other East-European countries didn't do that. The same theory says that all of Europe would have done this, but no country did. The EU did implement the ETS, and some countries did impose a carbon tax, but not for this reason. Mainly, they chose to subsidize solar and efficiency measures, causing "implicit carbon prices" to range from 0 to 1,000 €

or even more, instead of being uniform (Gollier and Tirole, chapter 10, this volume). In other words, in a major real-world test of global cap-and-trade, it did not price carbon emissions—it did not come even remotely close to accomplishing its central purpose.

A Global Carbon-Price Commitment A global carbon price commitment is a commitment made by countries to price their domestic carbon emissions, on average, at least as high as the agreed-on global carbon price. Like global cap-and-trade, a global carbon price commitment does not specify national policies. It would allow the EU ETS or fossil taxes, or a combination (as actually exists), and various other policies. But unlike global cap-and-trade, a global carbon price commitment will not count pure command-and-control policies. This does not mean countries cannot continue such policies. A global carbon price commitment simply ignores them.

In fact, command-and-control policies could even continue to be part of a separate international pledge-and-review system. Hopefully, the most wasteful of these policies would die out, and the beneficial, well-designed ones that plug holes in the carbon-pricing approach would be retained. (Later we will see that the climate fund used with a global carbon price commitment can provide incentives for good, nonprice policies.) A global carbon price commitment would also prominently raise the question of how much it is really costing a country to abate carbon. This will expose the wastefulness of some policies, and we believe this will strongly encourage greater efficiency through the use of price-based approaches.

A strict version of a global carbon price commitment would allow only true carbon pricing, such as cap-and-trade, fossil taxes, and bonus-malus (AKA feebate) pricing schemes. But at the start, it might be necessary, for political reasons, to count renewable energy subsidies as well. These would be credited only for carbon abated and not for the money spent on subsidies. An estimate would be made of the carbon price the country would have needed to abate as much without the renewable subsidies. It would then be credited with that carbon price.

So a global carbon price commitment, although not interfering with command-and-control measures, would actively encourage countries to engage in the efficient carbon pricing to meet their commitments. At first it would likely allow borderline pricing policies (like the renewable subsidies just mentioned), but eventually a global carbon price commitment would

become far more efficient than global cap-and-trade is likely to be. But even then, a carbon price commitment will be tremendously flexible. Different emissions from different fuels and different polluters can be priced differently. Automobile efficiency can be subject to highly effective bonus-malus pricing to address consumer myopia. Of course countries can use cap-and-trade or even join regional or global cap-and-trade networks. Also price credits could be traded and banked from one year to the next.

Although this flexibility will not lead to a completely uniform price by any means, it should lead to a much more uniform cost of abatement (implicit price) than global cap-and-trade, and it would make sense to put some restriction on the nonuniformity of carbon prices used to meet the global price commitment.

Figure 4.2 illustrates both the similarities and differences between the two approaches. Both allow all possible climate policies, but global carbon pricing will not count pure command-and-control policies toward the global commitment, whereas global cap-and-trade will. Instead, carbon pricing

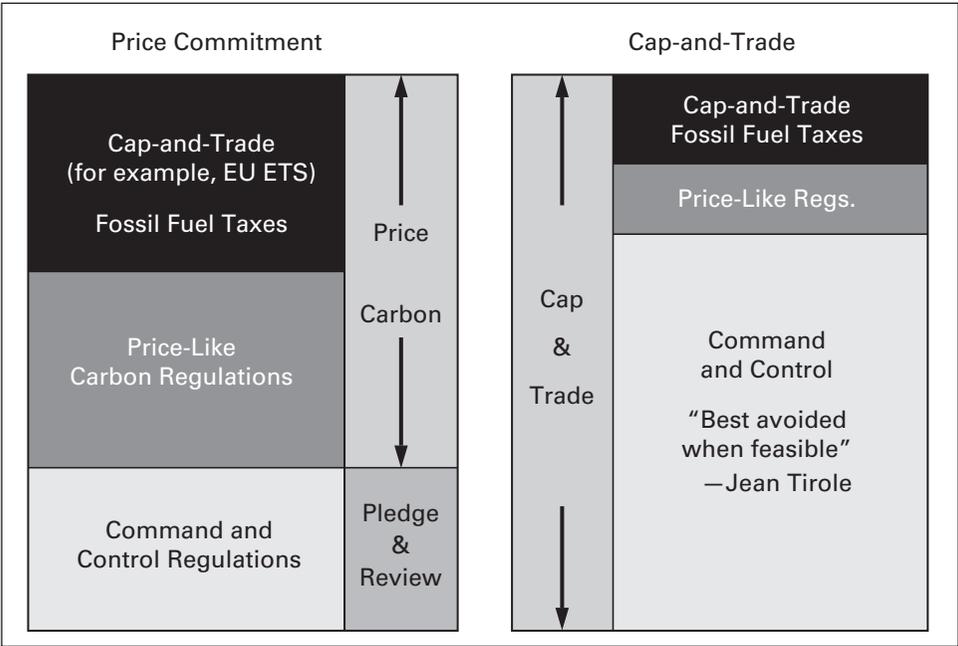


Figure 4.2
Pricing of carbon *emissions*

will accommodate parallel pledge-and-review commitments, whereas global cap-and-trade will not. At least for a while at the beginning, policies such as subsidies for renewables would have their abatements counted toward pricing. In these ways, global carbon pricing will shift the policy mix away from command-and-control and toward pricing carbon emissions (although it will provide some incentive through climate-fund rules for good nonprice policies).

The end result will be that global cap-and-trade will do little, if anything, to price carbon emissions, whereas global carbon pricing will do much to shift national policies in this direction by only giving credit for true carbon pricing or for a price-equivalent value of measured abatement.

Fairness

Perhaps the toughest problem for international negotiations is burden sharing. Who will bear the cost of stopping further atmospheric damage? The essence of the problem would seem to be fairness, a subject notoriously susceptible to bias. But the only focal point for fairness is to divide rights to the atmosphere equally among the entire human population.

This proposal might be as fair as possible for a simple rule, but there are two decisive arguments against making use of it. First, it will continue to be blocked by a large number of claims for alternative “fair” ways to share the burden. Second, it would surely be blocked by many richer countries.

Because this approach is decisively blocked, we suggest considering a focal principle that is not based directly on fairness but that would nonetheless contain checks and incentives that would pull it substantially toward a fair solution. This principle can be stated as follows.

Burden-sharing principle: The costs of improving the climate should be assigned in such a way that the climate is most improved.

This is not such a new idea. It has been long employed by Frankel (1998) in his quest for an effective common-commitment formula for the allocation of free carbon-emission permits. It has also been used by Cramton and Stoft (2010, 2012a,b) in their green climate-fund design for a number of years.

This principle has advantages, some important. First, its only built-in bias is toward countries that will be harmed most by climate change—because its goal is to maximize climate benefits. In principle, this could be

problematic. But in reality, there are stronger opposing biases, such as one caused by free-riding on future generations.

Second, the principle provides concrete guidance on how to negotiate equity transfers. In the next section, we use this principle to help allocate payments into and from the climate fund.

A More Realistic Treaty Design

We have now demonstrated how a common commitment to a carbon price could lead to cooperation in a highly simplified world. That involves a commitment to the lowest pledge submitted, which works perfectly, although it sounds like a weak approach. From here on we will refer to it as consensus voting because the lowest pledge is also the highest pledge that could elicit a consensus (100%) in its favor. This section introduces the two toughest problems facing climate negotiators:

1. Low motivation
2. Fairness questions

For various reasons, some countries may want to do much less than others; hence, they might vote for a carbon price as low as zero. As a result, consensus voting is ruled out as an effective negotiation tool. Second, some countries will need help with the costs of abatement, so a climate fund will need to be included in negotiations.

Preventing Too Low a Price

The first problem, low motivation, arises for several reasons. First, poor countries may heavily discount the future because they are so concerned with the present. Second, some northern countries may be somewhat ambivalent about being warmer. Also, countries that export oil may want the carbon price to be zero because a higher price suppresses the demand for oil and reduces their profits.

Because some countries may want too low a price, the minimum-price rule of the previous section would produce an unacceptably low carbon-price commitment if applied to all countries. To fix this, we must limit voting rights by excluding, in some way, the lowest votes from the determination of the global price. This can be done by forming a “coalition of the willing,” which is essentially the same as the idea of a Climate Club

suggested by Nordhaus (chapter 7, this volume, and 2015). Such a coalition could be formed through informal negotiation or a formal voting procedure.

Although it is easy to think of a procedure, such as forming the coalition from countries that vote for the highest global price, there are subtle incentive problems with many of these, so an informal procedure may be best. It is clear, however, that if the coalition has some power to reward those who cooperate with its pricing policy or punish those who do not, it will be easier to form a large coalition that agrees on a high price. In fact, enforcement is recommended by Nordhaus (chapter 7, this volume), Stiglitz (chapter 6, this volume), and Gollier and Tirole (chapter 10, this volume).

Fortunately, enforcement may be less necessary than is predicted by standard economics. Ostrom (1990) finds that “the fines assessed in these [common-pool governance] settings are surprisingly low. Rarely are they more than a small fraction of the monetary value that could be obtained by breaking the rules.”

In any case, we will postpone the discussion of enforcement until the next section and simply assume here that we can form a coalition that excludes uncooperative countries. Within this coalition, we can use consensus voting to select the highest global price that is acceptable to all.

A Green Climate Fund

The problem of burden sharing has been the most divisive and was addressed by the UN’s famous phrase in its Framework Convention on Climate Change (UNFCCC, 1992), stating that countries have “common but differentiated responsibilities.” The interpretation of this phrase has been most contentious. A global carbon-price commitment resolves this tension by making carbon pricing the common responsibility of all countries. Differentiated responsibilities are then handled by differing contributions to and receipts from a green climate fund.

Carbon prices should not be differentiated for two reasons.⁹ First, this wastes money. More important, as demonstrated in the prior example agreement, if coalition members commit to an enforceable common price, this eliminates free-riding at least within the coalition. This simplifies the negotiation and greatly strengthens the outcome. Because fixing the 20-year-old negotiation impasse is our primary objective, it’s best not to undo the progress we’ve made so far.

So differentiation of responsibilities should be accomplished with a climate fund and not by differentiating prices. This poses the obvious problem of how much each country should contribute or receive, which could be as complex a problem (although smaller according to Weitzman, chapter 8, this volume) as deciding carbon permit allocations under global cap-and-trade. But just as that problem is solved by replacing individual commitments with a common commitment, so can the climate-fund problem be solved with a common-commitment formula.

The most obvious climate-fund formula was invented independently by both cap-and-trade and carbon-pricing proponents and can be found in Stoft (2008, 2010), Cramton et al. (2010, 2012a,b) and Gollier and Tirole (chapter 10, this volume). Admittedly, it is too simple, but it is quite instructive and demonstrates most of the good properties that could be approximated with a more realistic formula. The formula is¹⁰:

Payment into the climate fund by country i , $G_i = g \times X_i \times P$, where g is the generosity parameter, X_i is the excess emissions of country i , and P is the global price.

Excess emissions are defined as emissions above what would occur if the country had the global-average per capita emissions rate. Negative values of G_i (resulting from below average per capita emissions) indicate a payment from the climate fund. Because there is a high correlation between emissions per capita and income per capita, this formula transfers funds from rich to poor countries. However, a realistic formula would need to be a bit more complicated to compensate for various anomalies.

The climate-fund formula has one primary effect and three beneficial side effects:

Primary effect: because of g , the formula makes successful negotiations possible.

1. It provides an incentive for poor countries to vote for a higher level of P .
2. It provides every country with an extra incentive to reduce emissions.
3. It can be used as a friendly enforcement mechanism for compliance with P .

The primary effect of the formula is to simplify the n -dimensional problem of negotiating payments for n countries into a one-dimensional

problem of negotiating g , the overall generosity of payments from rich to poor. Although this is essential, the side effects are also surprisingly beneficial and important.

The first side effect, an incentive to vote for a higher P , is most essential. As already noted, poor countries will tend to favor a low value for P . The climate-fund formula overcomes this problem because poor countries see that if the price is doubled, their climate-fund payments will double. This gives them an incentive to pledge and lobby for a higher carbon price, P .

The second beneficial side effect, an additional incentive to abate, happens automatically because any additional abatement reduces a country's excess emissions. So a country with high emissions would pay less and one with low emissions would receive more. The third benefit is activated by making a rule that the climate fund is paid only to countries that are in compliance with the global carbon price. This also makes rich countries feel like they are "getting something for their money," which makes these payments more palatable. However, requiring contributions from developed countries still makes them less likely to participate, but techniques described next, for maximizing the coalition's chosen price, should still help produce the strongest possible price agreement.

How the Carbon Price and the Climate Fund Interact

We have now specified a two-phase negotiation process that works as follows:

1. Negotiate climate-fund generosity g (payment = $g \times X_i \times P$).
2. Negotiate a "Climate Club" price P for a "coalition of the willing."

Breaking the negotiation in half, as this does, is enormously beneficial. Otherwise, when countries attempt to reduce their own burden, they unintentionally destroy the climate ambition of the treaty. This happens under cap-and-trade negotiations, where individual-country permit allocations add up to the total cap. As an additional benefit, these two negotiation phases also interact beneficially. These are the two complementary interactions:

1. The climate-fund negotiation over g is used to improve P .
2. Subsequent negotiation over P makes the g -outcome acceptable.

First, note that negotiating g does not require unanimous agreement because countries can protect themselves in phase 2. To understand

interaction 1, note that if too high a level of g is selected, then rich countries will pledge a low P to hold down their payments into the climate fund (see the equation for the payment into the climate fund above). Similarly, if too low a level of g is selected, poor countries will pledge a low P to avoid the abatement costs that come with higher values of P . So with either extremely high or low values of generosity g , one group or the other will opt for a low global price.

If either rich or poor pledge too low a price, then the global carbon price will be too low. So, with a coalition of rich and poor, both extremes of g must be avoided so that neither group will pledge too low. So by setting g at the right intermediate value, the highest possible price will be agreed on. This is in keeping with the burden-sharing principle suggested earlier.

Consequently, it is best if g is determined by countries that do not have a conflict of interest regarding climate-fund payments. These will be countries that have near-zero excess emissions and hence participate little in the climate fund. Such countries will be inclined to focus on getting a successful climate treaty with a high carbon price.

The second interaction guarantees that countries in the coalition of the willing will find both their climate-fund obligation and the global price acceptable. If they did not, then they would have pledged a lower value of P and that value would have become the coalition's agreed-on value. In the extreme, this could lead to a price of zero and no climate-fund payments, but for two reasons this should not happen. First, the coalition will be selected to contain cooperative countries. Second, offering the protection of the second interaction will make those selecting g especially careful to select a reasonable value. The result should be that few countries feel they need protection from the chosen g . So they will be willing to vote for a high global price in phase 2 of the negotiations.

Conclusion

We have considered two factors that tend to weaken a climate treaty: low motivation and fairness questions. Our strategy has been to avoid disrupting the common price commitment that serves to solve the free-rider problem. To maintain this common price, we have separated the "willing" from the unwilling and handled "differentiated responsibilities" with a climate fund.

To simplify climate-fund negotiations, we suggest using a formula. To make this easier to negotiate, we allow countries to opt out of a climate-fund

formula they perceive as unfair by not joining the coalition, although there will eventually be some penalties for opting out. This also motivates construction of a fair formula so that few will opt out.

Why Price Negotiations Work Better Than Quantity Negotiations Compare this process with negotiating a Kyoto-style treaty. Such a treaty allocates quantities, Q_i , of free carbon permits to participating countries. These quantities serve double duty. The total of all Q_i determines world carbon emissions, and individual Q_i s determine wealth transfers to each country. If your Q_i is higher by one ton, then you will be richer by the price of one ton of carbon. So every country will be paid handsomely to negotiate as high a Q_i as possible, which means every country is paid to do all they can to increase global carbon emissions.

Price negotiations eliminate this incentive to obtain a free ride from the negotiations. Some will see this flexibility and accommodation as a weakness and will want to enforce the “scientifically correct” commitment. This view is backward. A heavy-handed approach will only produce conflict or, at best, a treaty that quickly unravels. The source of this weakness is the lack of a global government. Given this weakness, procedures that reduce risk and eliminate adverse free-riding incentives will produce the strongest possible sustainable treaty.

Climate Clubs, Enforcement, and Reciprocity

Some say enforcement is the key to cooperation. This is half true, but we’ve already seen that enforcing independent commitments does not produce cooperation. Think of a road speed limit. If we independently set our own speed limits, then there would not be much use enforcing them. But if the limit applies to all, we have good reason to choose a moderate limit, and then enforcement is necessary and effective.

Fortunately, there are gentle types of enforcement, such as “internal enforcement” discussed next, social pressure, and rewards for poor countries that comply with the common commitment. These will be particularly useful at the beginning. Nordhaus (chapter 7, this volume) discusses Climate Clubs (coalitions) and a strong style of enforcement but makes little mention of a common commitment. However, as he explains, he still considers a common price commitment essential.

No other chapters in this volume focus on clubs, but several of them agree (and none disagree) that some enforcement, probably trade sanctions, will eventually be necessary. This section shows why Climate Clubs and carbon price commitments are just two different views of the same carbon-pricing-with-enforcement policy, although Nordhaus (chapter 7, this volume) advocates stronger enforcement than some of the other authors.

Internal Enforcement

Scott Barrett (1994) discusses self-enforcing international environmental agreements, and the first type he discusses uses what Nordhaus (chapter 7, this volume) refers to as “internal enforcement.” This enforcement is particularly gentle. To understand it, suppose the world consists of countries that are identical except for their size. This means that if one country is half the size of the other, then the larger country will experience twice the climate damage; if it abates the same amount per capita, then it will incur twice the abatement cost. Now suppose that the world has the following cost and benefit functions:

$$C = c A^2 \text{ and } B = b A,$$

where A is global abatement, C is the cost of global abatement, B is global benefit from reduced climate damage, and lowercase c and b are fixed parameters. These assumptions are typical and are the ones used by Nordhaus (chapter 7, this volume).

In this world, a single country acting on its own will realize that its own abatement will improve the climate and bring it some climate benefit. However, most of the benefit of its efforts will accrue to others. As it turns out, if a country is one-tenth the size of the world, then it will abate only 1/10 of what would be optimal, and it will do this by setting a carbon price 1/10 as high as it should. Of course the analogous rule holds for any other size country.

A Coalition of Two Equal Countries In a world with these cost and benefit formulas, would two identical countries be better off if they formed a coalition and made a common (enforceable) commitment? They will realize that a higher price would make more sense than when they acted alone because if they raise their joint price, then they will benefit from their own extra abatement and the same extra abatement from the other country.

As mathematics confirms, they will be better off together, even though they spend more on abatement. Neither will want to leave and have their small coalition fall apart. We will call this “internal enforcement” because it only depends on the externalities driving the agreement and not on some “external” threat, such as trade sanctions.

The Coalition Size Limit for Internal Enforcement Sticking with our simple model, if a large country and a tiny country try to form a coalition, then the large country will already have a fairly high price, but the tiny country would have a low price. In an equal-price coalition, that means the little country would have to work much harder than it would selfishly choose to on its own. Of course this extra effort would benefit the large country much more than the small one. So tiny countries will not want to form or stay in coalitions with large countries. They will prefer to free-ride on the large country.

As it turns out, if the small country is bigger than half the size of the large country, then it will benefit from joining in a coalition, but if it is smaller, it will lose. Similarly, if three identical countries form a coalition, each will view itself as just half as big as the other two put together and will be indifferent about being in the coalition. Hence, three identical countries make a coalition that is just on the brink of falling apart.

If the three countries differ in size at all, then the smallest will want to leave and free-ride on a coalition of the other two. This is what Nordhaus (chapter 7, this volume) terms the “small-coalition paradox.” It shows that internal enforcement is not strong enough to realistically hold more than two countries together, which is an argument for common commitments with external enforcement.

A Real-World Coalition

Fortunately, the real world may be more cooperative than the world of standard economic models. As noted previously, Ostrom (1990) found that penalties holding together successful public-goods arrangements were generally much weaker than what could be gained by cheating. There is now much evidence that weak and strong reciprocal interactions stabilize more cooperation than is generally predicted by standard economics that assumes static interaction. Let us consider just one possible outcome of this type that could be quite useful for getting started.

If the EU can continue to act as a single country, then the world will have three large countries: China, the United States and the EU, accounting for about half of all emissions. Together these might be a good nucleus for a Climate Club. China, being the largest, would have a positive incentive to stay in the coalition, and the EU would (according to the “small coalition paradox”) prefer to leave and free-ride on China and the United States. However, reciprocity, together with face saving and public pressure for good behavior, may push Europe into such a coalition. So internal enforcement, although not as strong as we need it to be, just might be strong enough to stabilize a coalition covering half of all emissions.

External Enforcement

Although Nordhaus (chapter 7, this volume) discusses internal enforcement, his Climate Clubs all rely on external enforcement. In particular, he favors trade sanctions that are simple yet powerful. These would be employed by Club members against those outside the club. Although some World Trade Organization (WTO) policies would need to be changed, he advocates placing a tariff of something like 5% on all goods sold by non-members to those in the Climate Club.

This approach is certainly logical in that failing to price carbon is a much larger problem than is indicated by the relatively small amount of carbon embodied in a country’s exports. However, we would not like to depend on this legally complex approach to get started, and we do not believe this is necessary. Later, when climate measures need to be stricter and climate problems are more obvious, this approach may be needed and may be possible.

In the meantime, as was mentioned previously, a substantial climate fund can be made use of as external enforcement, and it would likely be far more acceptable to developing countries. As noted, it would pay climate funds only to countries that meet the common price commitment. Also “internal enforcement,” discussed earlier, will help stabilize the initial coalition.

Reciprocity and Enforcement

Trust and reciprocity are essential to (and what Ostrom (2009) calls the “necessary central core”) of successful collective action. Broadly speaking, positive reciprocity means responding kindly to kind actions, whereas

negative reciprocity means responding unkindly to unkind actions. Both responses can act as enforcement.

Economics distinguishes two fundamental types of reciprocity: weak and strong. (Both can be positive or negative.) Strong reciprocity refers to actions that are similar to altruism, in that they do not serve narrow self-interest and often serve the common good. Weak reciprocity (more common) is motivated by narrow self-interest to gain better treatment by others. This is, of course, the point of any deliberate system of enforcement. Having explained this, we will discuss them all together and call any combination of them simply “reciprocity.”

We have suggested several ways of including reciprocity in the design:

1. Coalition members only commit to a price as high as others.
2. Poor countries that join the coalition will be rewarded.
3. If a county does not join the coalition, then it could be subject to trade sanctions.

This approach to treaty design should not be surprising because all disciplines dealing with human cooperation find that reciprocity is the key principle underlying cooperation. It is the most robust pattern of cooperation seen in laboratory, field, and theoretical studies of free-rider situations, and it is consistently found to stabilize higher cooperation levels. This has been thoroughly explained by scholars across all disciplines dealing with human cooperation (Bowles and Gintis, 2013; Fudenberg and Tirole, 1991; Kraft-Todd et al., 2015; and Ostrom, 1990, among many others). Without reciprocity, a public goods dilemma such as climate change will result in the tragedy of the commons.

Group Reciprocity Requires a Common Commitment

With only two parties, it is quite common for reciprocity to be asymmetrical—I will fix your sink if you pay me \$100. With three parties, it becomes difficult. You may suggest I will do x if Alice does y and Charlie does z. But Charlie may think he should do less and Alice more, which would be OK with Alice if you did more too. So the negotiation quickly becomes more complex as the number of parties increases.

Under pledge-and-review, every pledge will be of a different type, and all will need to compare the others' pledges to their own. But in reality, they will not find that worthwhile and will just focus on their own contribution. So there will be little, if any, reciprocity.

The complexity of many individual commitments makes effective reciprocity impossible without a common commitment. This could, in theory, take many forms, but none based on emissions quantities has been found to garner even modestly broad support. However, supporters of cap-and-trade and carbon pricing both agree that carbon prices around the world should be equal. That is the entire point of the “trade” in cap-and-trade, and that is all that is needed for a common price commitment—every country should commit to the same price.

Conclusion

Reciprocity is the key to cooperation, and enforcement is a form of reciprocity. To utilize reciprocity in a group, a common commitment is required. As previously suggested, this will need to take the form of a global carbon price. Simply agreeing to this common commitment is a form of reciprocity: “I will implement the global price if you will, and I won’t if you won’t.” Trade sanctions and climate-fund payments are negative and positive forms of reciprocity that can stabilize and enlarge a Climate Club or a coalition of the willing.

One reason that negotiating a global carbon price strongly facilitates a common agreement is that a *uniform* price is efficient and fair and thus a salient focal point for the negotiation. A focal point greatly reduces the complexity of multiparty and multi-issue negotiations and thus enables successful coordination and cooperation (Schelling, 1960). There is no salient focal point when negotiating global cap-and-trade.

However, many other actions, such as tech transfer or support in various international forums, could be used informally to help stabilize and strengthen a climate treaty. The real message of this section is that all countries should stop thinking in terms of altruistic climate aid and think instead about reciprocal actions of many kinds to encourage and stabilize a strong climate commitment. But none of these will work well until we have a global treaty based on common commitments. This is the most fundamental message shared by all experts contributing to this book.

The Enormous Risk of Trading

We turn now to one of the most serious but rarely mentioned problems of global cap-and-trade. So far, all workable forms of global cap-and-trade

require long-term allocations of permits to countries, so these must be based on long-term estimates of future business-as-usual emissions. When these estimates prove wrong, countries can find themselves with surprising windfall gains and losses, which have nothing to do with good climate policy.

Global cap-and-trade needs international trading of carbon permits for two reasons. First, some countries need to be given extra permits to sell as a way of reducing their burden. Second, some countries can abate more cheaply, so countries where abatement is costlier can (and should) take advantage of this efficiency gain by buying permits from them. This is equivalent to one country paying another to abate on its behalf.

Sometimes we may want to place no burden on a country by giving them all the permits we think they will need. However, by accident, we may give them too few permits, which could force them to spend a lot of money buying permits from foreign countries (or abating excessively, which would cost them even more). We call such trading “prediction-error” trading.

To understand the following example of prediction-error trading, it is useful to first understand two concepts: business-as-usual emissions and business-as-usual targets. Business-as-usual emissions are simply the emissions that would occur without a climate policy. Suppose this is one gigaton per year of emissions. In that case, a business-as-usual *target* gives the country one gigaton of free carbon permits per year. This means that it can ignore this climate policy and continue to emit one gigaton per year with no consequences because it has enough permits. But if it’s smart, it will realize that it can find some cheap ways to abate more carbon; by taking these opportunities, it will find it has leftover permits, which it can sell to other countries at a profit. In this way, giving a country a business-as-usual target keeps it safe—it can do nothing and have no cost, and it also gives it an opportunity to abate and make some profit selling permits.

Frankel on the Safety of Business-as-Usual Targets

Jeffrey Frankel served on the US President’s Council of Economic advisors during the Kyoto negotiations. In July 1998, he wrote, “Let us consider a plan under which developing countries commit to their ‘business-as-usual’ emission paths in 2008–2012.” He considered that a bit more generous than was likely because, at that time, cap-and-trade advocates were favoring stricter targets than business-as-usual.

Then he wrote about countries such as China: “The first thing you should notice is that this system is not going to hurt you.” He explained exactly what we explained earlier about a business-as-usual target. However, his explanation and ours both apply to targets that actually do match business-as-usual emissions, and this is not what his claim of “not going to hurt you” applies to. He was talking about setting a target in 1998 for the Kyoto compliance period of 2008–2012, which is what the Kyoto treaty did. So there is no reason to believe there actually would be a perfect match between the so-called business-as-usual target (the allocations of free permits) and the actual future business-as-usual emissions of various countries.

China Comes Up 29 Billion Permits Short

Because Frankel mentions China in this discussion, let us look at how China might have fared. The US Department of Energy’s (DOE) *1999 International Energy Outlook* predicted that China’s CO₂ emissions in the target years would total 7.5 billion tons. In reality, they turned out to be 36.6 billion tons. So China would have been short 29.1 billion permits. At a permit price of \$30/ton, China would have had to spend \$874 billion buying permits, mostly from developed countries, had it not cut emissions.

Of course they would have found some emission that could be cut more cheaply than buying permits, so that might have brought the bill down to, say, \$600 billion, but then again trying to buy an unexpected 20 billion permits from the market might well have sent the price up above \$30/ton. In any case, the Chinese might have taken issue with Frankel’s assurance that “this system is not going to hurt you,” especially when they realized their rich trading partners would be selling them permits at the marginal cost of abatement, which is always higher than the average cost. Hence, rich countries would have profited from China’s bad luck, quite possibly by more than \$100 billion.

Frankel Proposed a Fix and Then Dropped It

To be fair, a few pages after estimating that, “If China were to join, it would capture almost \$4 billion a year” in gains from trade,¹¹ Frankel does issue a warning: “One problem is the uncertainty of the business-as-usual path. It is difficult to forecast ten years ahead what a country’s emissions would be in the absence of policy change.”

He then suggests, “I have a possible response to this problem. It is a suggestion to index the emissions target, to such variables as GDP in the year 2007.” This would have helped, but the GDP prediction for 2007 from back in 1999 was only 26% short of the mark, whereas the CO₂ prediction was 80% too low. This correction technique, although helpful, would only have eliminated roughly one-third of the problem.

Frankel also suggested, in 1998, that the business-as-usual path could be determined by “objective experts,” which would seem to correspond to the DOE forecast used previously. In 2014, when describing his most recent formulas for determining future free permit allocations for a global cap-and-trade system, he suggests, business-as-usual “is defined as the path ... countries would experience in the absence of an international agreement, preferably as *determined by experts’ projections*” (italics added). So 16 years later, he has settled on the method (experts’ projections) that produced the 29-billion-ton underestimate of the business-as-usual target, as still the best estimation method he can come up with. The point is not to fault Frankel but rather to indicate that the problem of setting even roughly accurate business-as-usual targets appears unsolvable.

Comparing Global Price versus Global Cap Commitments

Suppose that China had agreed to a global carbon price commitment instead of global cap-and-trade in 1999. Let us add some detail that, although speculative, is in no way implausible. Rather it consists of exactly the sort of assumptions the Chinese should have made, and perhaps did make, when determining whether to accept the type of “binding commitment” they were being asked for. Suppose, to make comparison simple, that the *expected* carbon price under either global cap-and-trade or a comparable global carbon price commitment would have been \$30/ton. Further assume that, given the surprising increase in China’s business-as-usual emissions, the global carbon price under cap-and-trade would have risen to \$45/ton and a \$30 carbon price would cause 20% abatement and a \$45 price would cause a 30% abatement.

Under either system—a cap or a price—there would be abatement costs, which are reasonably and traditionally calculated with the following cost-of-abatement formula: $C = P \times A/2$, where P is the carbon price and A is the abatement.

With this formula, we calculate the cost of abatement as \$247 billion under cap-and-trade and \$110 billion under carbon price commitment. The cost is less under a price commitment because the price cannot rise, whereas under a cap the surprise emission shock pushes it up to \$45 billion. But there would still be 25.6 gigatons of unabated emissions under cap-and-trade, only 7.5 of which China would have permits for. So it would have had to buy permits for 18.1 billion tons of emissions at \$45/ton, for a cost of \$817 billion, and a total cost under global cap-and-trade of \$1.1 trillion over the five-year life of the Kyoto treaty.

But remember, some abatement cost (\$22 billion) would have been expected under either system if the 1999 emission prediction had been accurate. Under cap-and-trade, the permits for the abated tons could have been sold at a profit of \$44 billion. The net gain (trading cost minus abatement cost under cap-and-trade) would have been \$22 billion if everything had turned out as expected.

The net *unexpected* cost under a carbon price commitment would have been $\$110 - \$22 = \$88$ billion (unexpected minus expected abatement costs). The final result is that the unexpected cost to China would have been more than 12 times greater under global cap-and-trade than under a matching global carbon price commitment, and it would have been more than \$1 trillion greater.

It should also be noted that, although the unexpected cost of \$88 billion (over five years) is still fairly large under a global carbon price commitment, this cost would have gone mainly toward cleaning up their coal industry and solving a major internal pollution problem. The \$817 billion spent on purchasing carbon permits from, say, the United States and EU would have caused unimaginable domestic political recriminations had it been carried out. These costs are illustrated in figure 4.3.

In figure 4.3, areas represent costs, and the sloped line represents the demand-curve for carbon emissions. The higher the price of carbon, the lower the emission level. The rectangle shows the cost to China of purchasing permits after doing extra abatement due to the unexpected \$45/ton cost of permits. China's unexpected abatement cost is the combined area of the two trapezoids. The smaller trapezoid (on the right) is the cost of unexpected abatement under a fixed global carbon price of \$30. The area of the triangle represents the cost of abatement that was expected under either system due to a \$30 carbon price and the DOE-predicted level of

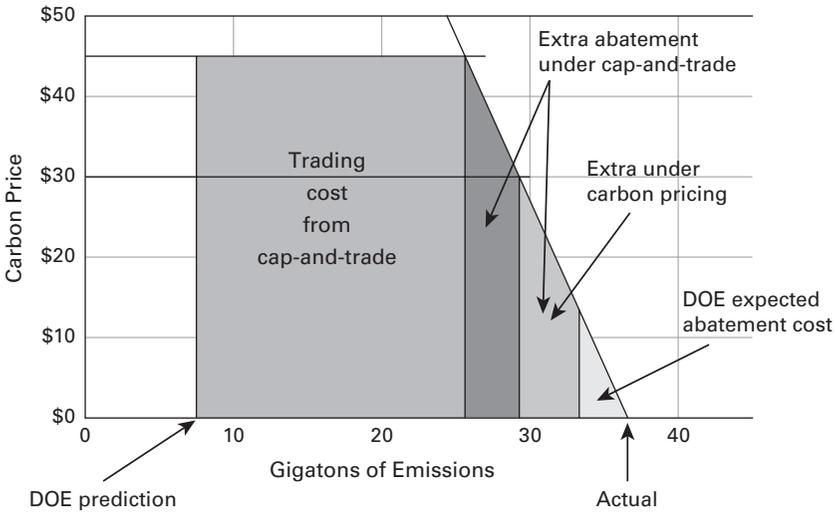


Figure 4.3
Prediction-error trading costs

emission. (The triangle has the right area but has been moved and reshaped from where its area was calculated—at the DOE prediction.) From China’s perspective, abatement costs in the two trapezoidal areas have considerable benefit and, hence, low net cost. The trading costs under cap-and-trade are far larger, have no benefit, and carry a considerable political liability.

Conclusion

The previous example was not cherry picked. It was chosen by a leading advocate of global cap-and-trade, who simply was uncertain of what the future would bring. It should be noted that this is not the only surprising change in business-as-usual emissions we have witnessed. There was also the collapse of the Soviet Bloc and the Fukushima disaster, among others. Global cap-and-trade is designed so that it turns such unexpected shocks into huge windfall gains or losses, which will inevitably destroy any treaty with an effective carbon price based on this approach.

Almost all comparisons of global capping and pricing made by capping advocates have used what is called a “certainty equivalent” model. This ignores all prediction errors as if being right on average was the same as being right all the time. Yet two of the most serious problems with global cap-and-trade are due to price and business-as-usual uncertainty. Both of these are rendered invisible by certainty-equivalent models.

In the case of business-as-usual uncertainty, it seems likely that, as Cooper (chapter 5, this volume) explains, developing countries will want caps that assume a business-as-usual CO₂ growth rate something like China's because that might happen, and they do not want to be seen as projecting anything less than stellar growth. They also do not want to risk having to buy billions of dollars' worth of carbon permits from rich or rival countries. Accepting permit allocations that accommodate such hopes and fears will result in a global cap that is far too high to have any significant effect on the climate.

Does Cap-and-Trade Have a Record of Success?

Two systems claim to achieve a fairly uniform carbon price: global cap-and-trade and global carbon pricing. We have already made several comparisons, but one misconception still needs to be addressed. Has global cap-and-trade already been widely implemented and found to work quite well? If so, what is the point of analyzing its shortcomings?

The Kyoto Protocol is global cap-and-trade. It allocates international emissions permits (AAUs) and sets up a system for trading them. The argument for trying this was largely that standard cap-and-trade had been tried in the United States, and it worked well. In fact, it did work well for curbing sulfur emissions, but that argument is without merit. Global cap-and-trade and standard cap-and-trade are fundamentally different.

1. Standard cap-and-trade is run by a government, whereas global cap-and-trade is not.
2. Standard cap-and-trade can subsidize participants with somewhat hidden transfers, whereas global cap-and-trade transfers are far more transparent.
3. Standard cap-and-trade prices carbon emissions, whereas global cap-and-trade does not.

The only track record for global cap-and-trade is the dismal record of the Kyoto Protocol.

There's No Government

Under standard cap-and-trade, the government sets and enforces the cap. Under the Kyoto Protocol, because there is no global government, no one even talked about what the cap would be, and no one enforced it. The cap

was determined indirectly, not by the Protocol, but by the sum of the AAUs eventually allocated to those who ratified the treaty. This is like having the coal-fired power plants discuss among themselves but decide individually how many sulfur emission permits each would get.

Profits from Permits, but Not for Countries

It is often claimed that global cap-and-trade will be good for hiding compensating transfers to polluters (Gollier and Tirole, chapter 10, this volume). This is true but could be misleading because these systems hide some of their transfers in a way that will not work under global cap-and-trade.

Standard cap-and-trade causes companies to raise their prices (due to the “opportunity cost” of not selling the permits, which we will not explain here). The result is that standard cap-and-trade can actually increase the profits of polluters without any money passing from the government to the polluters. This will not work for international financial transfers, however. Poor countries cannot profit by raising prices on their own citizens. As Gollier and Tirole (chapter 10, this volume) point out,

To be certain, the transfers made under national cap-and-trade programs are different in their economic and political nature from international payments for international permits.

... transfers associated with an allocation of free permits are not that hard to compute and one would imagine that politicians (privately or publicly) opposed to an ambitious climate change agreement would quickly publicize the numbers (if unfavorable to the country) so as to turn their domestic public opinion against the agreement.

In fact, under the Kyoto Protocol, AAU trading became so controversial that Japan had to publicly deny purchasing AAUs from countries previously in the Soviet Bloc (Cramton et al., chapter 12, this volume). The argument that standard (domestic) cap-and-trade demonstrates that global cap-and-trade can hide international transfers from the rich countries and their citizens while making them transparent to the poor countries and their citizens is questionable.

Global Cap-and-Trade Will Not Price Carbon Emissions

Previously, we showed that global cap-and-trade does not *require* countries to price emissions. In reality, under Kyoto, the AAU market was so

illiquid and secretive that there was no “market price,” and the price of few transactions was known. This did not result in any carbon pricing policies at all. The main Kyoto Protocol compliance policies were subsidies and requirements for wind, solar, and energy efficiency. These do not put a price on emissions even though the implicit cost of saving carbon ranged as high as 800 euros per ton (Gollier and Tirole, chapter 10, this volume). *Global* cap-and-trade is unlikely to cause much pricing of carbon emissions, unlike *standard* cap-and-trade, which *requires* the pricing of carbon emissions.

Why Global Cap-and-Trade Negotiations Cannot Succeed

The most decisive flaw in global cap-and-trade is that a strong treaty could never be negotiated, and if it could be, it would unravel. The three main parts to this argument are:

1. Trading risk would unravel a strong global cap-and-trade treaty
2. Free-style permit negotiations would likely end in deadlock
3. No common-commitment formula can be found to replace freestyle negotiations

Note that we cannot rule out a weak global cap-and-trade agreement—one that has too little impact on the climate to warrant attention. However, we ignore this possibility because it is essentially useless, and instead we focus exclusively on the problems of strong global cap-and-trade agreements.

The Risks of Prediction-Error Trading

Earlier we discussed prediction-error trading risk in detail and concluded that unexpected shocks to business-as-usual emissions would lead to deflections. This point is not necessary for the present argument because we will argue next that a strong treaty could not even be negotiated. However, this risk is reason enough not to embark on such an adventure. The cost in time (decades) and effort to put such a system into place should not be squandered on one that would have disintegrated in just 10 years had it been built 15 years ago.

We also argued that the knowledge of such individual country risks would drive the developing countries (and likely others) to demand

larger allocations of carbon permits than they will likely need just to protect against risk (see Cooper, chapter 5, this volume). This factor, in addition to the ones we are coming to, will weaken any global cap-and-trade treaty.

Why Freestyle Negotiations End in Deadlock

As explained by Gollier and Tirole (chapter 10, this volume):

Free-style negotiations among n countries are exceedingly complex and are very likely to lead to a deadlock ... [concerning] the allocation of free permits among countries under cap-and-trade.¹²

The extreme complexity they mention is only half the reason that deadlock is inevitable, but it is still decisive. Such complexity is obvious from the dozen or so different variables that Kyoto negotiators attempted to account for when they tried to invent formulas for allocating permits (Depledge, 2000). Many factors were ignored—for example, access to renewable resources.

But Weitzman (2014, chapter 8, this volume) and Cramton and Stoft (2010, 2012a,b) emphasize a different problem with freestyle negotiations—free-riding. With freestyle negotiation of permit allocations, every country's self-interest is to gain more free permits. This dramatically weakens freestyle commitments. As noted, the risks of prediction-error trading will make countries even more aggressive in their demands for free permits.

On top of the free-riding problem of freestyle negotiations, we have the extreme complexity noted by Gollier and Tirole (chapter 10, this volume), which includes the mixing of climate policy with burden sharing. This makes it that much easier to find excuses to hide behind when free-riding.

But none of these effects leads directly to a deadlock. Instead, they only seem to lead to weak national commitments, Q_i . Because the sum of all such commitments is the global cap Q , there will be a high (weak) global cap.

But Gollier and Tirole (chapter 10, this volume) make the first step in their negotiation process the selection of Q , the global cap, and assert that this should be consistent with a 2°C limit. Because the outcome of this first step in the negotiation does not commit any country to do anything

in particular, all will want to show their “ambition” by agreeing to a tight cap, probably consistent with the 2°C limit, as has been the case with some previous aspirational agreements, including COP21’s Paris agreement.

But with a 2°C cap locked in place, the weak individual-country commitments, which sum to a weak global cap, can now be seen to lead to deadlock. Deciding the global cap by two completely different processes, one that leads to unchecked optimistic aspirations and one that leads to nearly unchecked self-serving caution, will never produce consistent results. Hence, deadlock is inevitable.

Why a Common-Commitment Formula Fails for Quantities

As just seen, a freestyle negotiation leads to deadlock, but we have argued that a common-commitment simplifies negotiations and solves the free-rider problem. So why doesn’t this work for a quantity-based treaty? We first noted that a 20-year search for a common quantity commitment has turned up no satisfactory proposal. This history of failure is no accident.

The root of the problem is the nature of the quantity approach. Every allocation of free permits plays two contradictory rolls. Permits are money with which to solve the burden-sharing problem, and, collectively, permits must curb emissions. In “theory,” they could do both at once, but that requires the allocation of permits by a fair world government with perfect foresight.

Kyoto’s Formulas The Kyoto negotiations first tried a simple rule—equal percentage reductions from 1990 emissions levels. When that failed, they went on to try nine more-complex rules (Cramton et al., chapter 12, this volume; Depledge, 2000). But all of these failed as well, and countries were left to choose their own commitments—a freestyle negotiation indeed.

Frankel’s Formula After Kyoto, it was obvious that no acceptable allocation rule was in sight, and including the developing countries would make finding an acceptable rule far more difficult. It was also obvious from the US Senate’s 95 to 0 vote that, without the developing countries, the United States would not join. Understanding this, Jeffrey Frankel (1998) took up

the challenge and worked on politically acceptable allocation formulas. These evolved over the next 16 years and are quite sophisticated (Bosetti and Frankel, 2014). They specify free permit allocations in terms of several parameters, including business as usual emissions, emissions in 1990, and, for the initial-year formula, emissions in the year the country signs the cap-and-trade agreement. So far, there does not seem to be much interest in these formulas, which may not be as transparent as required for acceptability.

Stiglitz (2006a, 2006b) argued that it would be impossible to find a formula for free permits that the world could come close to agreeing on, and Weitzman (chapter 8, this volume) has taken a similar position. History seems to be confirming these judgments.

Conclusion

The Kyoto negotiators knew they needed a common-commitment formula and invented 10 of them. They could not agree on any. After Kyoto, it became clear the problem would become far more difficult because developing countries would need to be included. Realizing this, Frankel began proposing formulas in 1998 that covered all countries. There has not been much interest in these perhaps because of their complexity, and after 16 years there seems to be less interest than ever.

The possibility that an easily agreed-on common-commitment formula for global cap-and-trade will someday be discovered cannot be ruled out with certainty. However, it seems that after 20 years of failure and a general loss of interest, it is time to take global cap-and-trade off the table.

Problems with Global Carbon Pricing

Although a global carbon price commitment is a more direct and simple approach to carbon pricing than is global cap-and-trade, a strong enough version of a carbon price commitment will still be difficult to implement. But difficult is better than impossible. Here we examine the points that may need the most attention from negotiators and researchers. The key problem areas are enforcement and climate-fund transfers.

Enforcement

There are always two parts to enforcement: monitoring and incentives. To enforce, you must find out whether the party is in compliance. That's monitoring. To get them to comply, there must be an incentive. The incentive can be a carrot for compliance or a stick for lack of compliance.

Incentives The incentive problem is much the same for any climate commitment. There are social-pressure incentives and there are financial incentives. It is not clear whether the former will be strong enough, and it is not clear that the latter can be implemented. This is equally true of global cap-and-trade or a global carbon price commitment. But one thing is certain: in either case, the problem is much worse without a common commitment. As we argued earlier, such a commitment is almost certainly impossible for global cap-and-trade.

In fact, without a common commitment, strong enforcement is counterproductive. If you think you don't want to drive faster than 70 mph, then you might commit to that individually out of a spirit of cooperation with weak enforcement, but with strong enforcement, say a \$10,000 fine, you will certainly not commit to anything under 90 mph, "just in case." Although enforcement may be hard to arrange, at least with a global carbon price commitment, it is of some use and not counterproductive.

Monitoring The primary challenge for monitoring price is the possibility that a government will cook its books with regard to revenues collected from carbon charges. For most countries, this should not be a problem because they will either provide reasonably reliable public data (most of the Annex I countries under the Kyoto Protocol) or they will be poor countries receiving some climate-fund assistance that can be withdrawn if they do not fully open their books.

For the problematic countries, and there may be a couple of large ones, there are three recourses. First, if they do not open their books to careful auditing, they could be deemed noncompliant regardless of claims concerning carbon revenues. Second, four international organizations—the World Bank, the International Monetary Fund, the International Energy Agency, and the WTO—already conduct similar audits. In fact, such audits would

be needed to monitor global cap-and-trade with regard to carbon pricing of exports—one of the most difficult segments of society to monitor. Of course, whichever organization performs this function will need additional funding, but that will be a small burden relative to other costs.

Finally, the price of most fossil fuels has easily visible public indicators. The price of gasoline is no secret, and that accounts for roughly one-third of fossil-fuel use. The price of electricity to residential and commercial users can also be discovered easily, as can the price of heating fuels to these groups. Monitoring will not be perfect, but with a little effort, it may well be as good as or better than the monitoring of emissions.

Green Climate Fund Transfers

We have discussed how to allocate responsibility for and benefits from a climate fund. A higher level question is perhaps just as difficult. How can significant funds be transferred from rich to poor countries without triggering too much political opposition in wealthy countries, especially those on the hook for larger transfers, due to their wealth and high emission levels?

Hiding Transfers One approach is to hide the transfers. This is often cited as a benefit of cap-and-trade, but as explained earlier, this is largely based on a fallacious analogy between global and national cap-and-trade programs.

However, Frankel suggests that poor countries could be given free permits, and they could give the permits to private companies, who would then sell them to private companies in rich countries. In this way, the financial transfers would be kept private and less visible than the financial transfers between governments that are generally envisioned for the Green Climate Fund.

This method would not be as surreptitious as it might seem because companies in the rich country will be required to return the permits to the UN in order to make use of them, and the UN will be required to keep a full accounting. This transfer will be made public, at which point the press will write stories about how much money went where. It may take a few years before this information is fully utilized by the forces that wish to gain political advantage from stopping the transfers, but that outcome seems inevitable.

If this ruse is thought to be effective, then a similar process could be arranged under a global carbon price commitment. If the United States had been allocated a responsibility for \$10 billion of climate fund contributions and the global carbon price was set at \$20/ton, then half a billion carbon-price credits could be issued and marked as redeemable in the United States only. These could then be distributed to poor countries, which would give them to their businesses, which would sell them to US businesses, which would then not be charged for that many tons of carbon emissions.

One advantage of price-based climate-fund transfers (as opposed to permit transfers under a cap) is that their value would be far more predictable. For example, with global cap-and-trade, the carbon price would drop precipitously in the case of a global slowdown. In this case, permit transfers to poor countries would suddenly become far less valuable and perhaps nearly worthless. However, the global carbon price might spike while a developed country is in recession, and it would find itself making double or triple climate-equity payments at just the wrong time. Such risk would not be present under a global carbon price commitment.

Making Transfers More Appealing Jonathan Gruber, an economist who consulted on the design of President Obama's Affordable Care Act, is now famous for explaining that the "Lack of transparency is a huge political advantage" for "getting the bill passed." In the long run—and no policy is longer run than climate policy—attempting fairly transparent deceptions involving tens of billions of dollars may prove counterproductive. Better approaches are available.

The first principle for making equity transfers more palatable is to make sure they are reciprocal. Traditionally, this would mean requiring the money be used for some approved "green" project, hopefully related to climate. Unfortunately, history has shown that this leads to corruption—witness the Clean Development Mechanism, the Joint Implementation Mechanism, and even the enormous subsidies for corn ethanol in the United States and elsewhere.

The basic formula for reciprocity should be that equity transfers are conditional on compliance with either a global cap-and-trade or a global carbon price commitment. This will provide funders with far more assurance

that they are getting something worthwhile for their money while providing a useful incentive mechanism for enforcing compliance.

A number of other standard techniques are available for making transfers more palatable. One is to require funds to be spent in the donor country. This would not be possible with global cap-and-trade. Another way is to earmark tax receipts from the most unpopular domestic polluters to be used for equity transfers.

Summary and Conclusion

If steady progress was being made with global cap-and-trade, then even a promising new approach would seem questionable. But after 20 years of real-world testing and academic theorizing, no obvious progress can be seen. Our discussion leaves global cap-and-trade with four decisive failures, all of which are addressed by global carbon pricing.

Carbon Pricing Eliminates Huge Trading Risks

Global cap-and-trade needs to lock in targets for a decade or two. During this time, business-as-usual emissions change unpredictably. As shown previously, this can be extremely risky for participating countries. This leads to demands for more generous targets or even refusal to participate. If a strong treaty were ever implemented, then it would lead to defections and unraveling. Global carbon pricing nearly eliminates this source of risk.

Carbon Pricing Actually Does Price Carbon Emissions

Neither the Kyoto Protocol nor global cap-and-trade, as specified in this volume, requires that emitters acquire emission permits. Instead, governments must own permits similar to Kyoto's AAUs. In idealized economic theories, the price of AAUs would be transmitted, with the help of government regulations, to actual carbon emissions. There has been no sign of this under the Kyoto Protocol, and there is no reason to believe things would be different under a newly proposed global cap-and-trade policy. In contrast, global carbon pricing would require countries to price carbon emissions to meet the global carbon price commitment. So global carbon pricing would strongly promote efficiency, and global cap-and-trade would do little to promote efficiency.

Pricing Rewards Environmental Ambition

A global cap, if it works as intended, will control the total emissions of the participating countries. If one country emits less, that will free up permits so other countries can emit more. If one emits more, others will be constrained to emit less.

The consequences are obvious. If any country, province, social group, or individual voluntarily does more than is in their narrow self-interest, *it will not benefit the climate at all*. All such altruistic efforts will be negated by the market. Ambitious action by some will simply make it cheaper for those who are not ambitious to do less, and they will do less. If they do not do enough less to negate all environmental ambition, then the market will depress the price of carbon even more and make sure selfish people do even less. The cap will be met.

Global carbon pricing does not discourage ambition at all. Extra abatement does not change the price faced by nonambitious groups and individuals, so the ambition of others does not encourage them to do any less. The result is that every ton of ambitious abatement reduces global emissions by a full ton.¹³

Pricing Stops Free-Riding in the Negotiations

Climate change is a problem of managing the collective commons, and the essence of that problem is that countries can free-ride on the use of the atmosphere. Requiring them to pledge some action, even if the action is to join a global cap-and-trade agreement and choose a “target”—an allocation of emission permits—does not prevent free-riding. Instead of free-riding by just emitting, countries can now free-ride by taking a high target and either emitting more or profiting from selling extra permits.

To stop free-riding, we need to replace individual commitments with a common commitment. For 20 years, Kyoto negotiators and academic economists have tried to find ever more complex formulas to create a common quantity commitment, with no signs of progress. Global carbon pricing provides an obvious solution. All countries should commit to price at the same global price. There is still a problem of negotiating climate-fund transfers, but decoupling these two problems greatly simplifies them and largely insulates climate policy from disputes over monetary transfers.

Conclusion

Global carbon pricing was designed to facilitate negotiation and cooperation. To many this will seem backward—it should have been designed “for the climate.” But the real problem is not the climate; the real problem is people—and their lack of cooperation. After 20 years of pretending to do what is right for the climate, and actually doing almost nothing, it is time for a change. We should design the negotiations and our policy goal to maximize cooperation and accept that we cannot do better than the best we can do. Unfortunately, COP21 in Paris was a step back from this perspective. Paris focused on nonbinding, nonenforceable, incomparable “intended nationally determined contributions,” which is the opposite of a reciprocal, common commitment. As a result, contributions do not add up to what is required, and carbon pricing was hardly mentioned in the final agreement. This is a recipe for inaction, and thus disaster.

Carbon pricing is a simple idea. But the change of focus from supposedly scientific round numbers, 1 trillion tons, 2.0 degrees, 450 ppm (or some say 350), to a focus on how people cooperate makes all the difference. Elinor Ostrom spent her life studying how people actually solve common-pool resource problems. She found the answer was always “trust and reciprocity,” not numerology. Virtually all cooperation research agrees. Global carbon pricing is designed to build trust with reciprocity.

Frequently Asked Questions About “Global Carbon Pricing”

Q1. Does it mean a global tax? No. It does not require that any carbon taxes or fossil fuel taxes be implemented. See Q3.

Q2. What is it? An agreement between countries responsible for most of the world’s greenhouse gas emissions to price their own fossil-fuel emissions at least as high (on average) as the agreed-on global price, P .

Q3. What does “to price their own emissions” mean? Quite simply, a country’s average carbon price—carbon revenue per unit emissions—must be at least as high as the global carbon price. The simplest way would be with a carbon tax, which could be used to replace other taxes. Cap-and-trade could also be used, as well as other methods. Renewables could be given credit based on carbon saved and the global price.

Q4. What does “on average” mean? Countries could price gasoline at one level, diesel at another, and coal at another. All that matters is (total carbon

revenue)/(total carbon emissions) $\geq P$. There could even be averaging from one year to the next.

Q5. Who would set the global price? It would be negotiated by a “coalition of the willing,” AKA a Carbon Club. This coalition will be a group of countries that encompass most emission and are willing to set a reasonably high price.

Q6. Why does a price agreement help? It forms a common commitment, so each country in the coalition is saying, “We will price carbon at P if all of you will too.” Read the preface to see how this works.

Q7. Is it fair to poor countries? A green climate fund is needed because without one there would be no international payments. This negotiation must be separate but related. The UNFCCC requires “common but differentiated responsibilities.” The global price is the common part and the climate fund is the differentiated part.

Q8. Why not stick with global cap-and-trade? There’s a reason it has been getting less popular for 20 years. It was accidentally designed to be hard to negotiate. The idea was to make it safe for the climate but risky for countries. Global pricing was scientifically designed for cooperation, and it can be adjusted to hit climate targets just as well, probably better, than global cap-and-trade.

Q9. Who’s in favor of it? Everyone on the list of contributors to this book is in favor of global carbon pricing. The authors have alternative views on how best to implement it.

Q10. With your green climate fund, how big would the transfers be from rich to poor? At the start, a high-end estimate might be €36 billion per year, and a low-end estimate might be €5 or €10 billion. But this is, of course, speculative. It will be determined by negotiation, not science, so it can’t really be calculated. Negotiators will balance rich-country reluctance against poor-country needs and demands.

Note that the high end is about one-third of what US Secretary of State Clinton promised at Copenhagen. To put this in perspective, this is about one-tenth of 1% of the rich country’s GDP. This is for a €30/ton carbon price. Eventually, it would need to go much higher, but by then the world will likely have seen enough to be willing to spend more.

Consider the high estimate first. World CO₂ emissions from fossil fuel are a bit less than 36 billion tons. China has said it doesn’t need climate-fund

subsidies, so that leaves about one-third of the emissions (12 billion) coming from poor countries that need climate funds. A tremendously strong start would be a €30/ton carbon price, and that might reduce emissions by as much as 20%, or by 2.4 billion tons in poor countries. Some abatement will be cheap to free, and some would cost as much as the €30 carbon prices, so on average the cost would be about €15/ton, for a total cost of $15 \times 2.4 = €36$ billion/year. So the high-end number assumes that rich countries pay 100% of the costs and somewhat more because when the poor countries stop subsidizing fossil fuel, that actually saves them money (it prevents waste).

But €30/ton is a high starting price, and 100% is a high subsidy rate, and not all of these countries will join and need subsidies (e.g., some of the OPEC countries). In fact, it may be necessary to begin quite slowly. But after 20 years without any global cooperative agreement, a slow but solid beginning would be enormous progress. Also remember that without any transfer from rich to poor, little is likely to get done.

Q11. What carbon price do you think the EU countries, for example, would vote for? This brings up the central advantage of global carbon pricing. But first, note that the United Kingdom is already paying more than £100¹⁴ and more per ton of carbon saved, and the OECD¹⁵ finds that feed-in-tariffs cost an average of €169 per ton saved, and there are other subsidies on top of those.

Second, notice that a £100 carbon tax, if implemented as a tax shift, would be close to free. The tax that was shifted away would return as much revenue as the carbon tax collected, and the distortions and inefficiencies of the old tax would be eliminated. These would roughly match the cost of carbon abatements, and those abatements would have the added benefit of reducing damage from domestic pollution. So not even counting the climate benefit, this policy might produce a net benefit.¹⁶

Now turn to our best feature. Global carbon pricing is not an individualistic approach. The EU would not be doing this without major partners, at least the United States and China, and probably more, even at the start. The agreement would be that all countries price as high as the global price. Now we have no illusion that the EU will suddenly impose a £100 carbon charge. More likely, it will be inclined toward some timid level, such as €25/ton. But with our proposal, the EU would then realize that if it advocates €35, then getting that accepted would bring China, the United States,

and others along with it. So why not advocate €35? You only have to do that if you gain the satisfaction of finally bringing the United States and China along with you, and bringing them along would at least quadruple the impact of that €10 increase.

So we can't predict the EU's price proposal, and we certainly cannot predict what the coalition of the willing will agree to, but we can tell you that even a €30 price on carbon could save a lot of money while doing far more good for the climate than current policies.

Q12: What if some countries have large-scale, relatively cheap, and measurable carbon-capture potential (e.g., afforestation)? How could that potential be harnessed with a global carbon price? This will require an add-on mechanism, but a fairly simple one once the negative emissions become measurable. The measurement process would supply the negative-carbon facility with a one-ton carbon credit for each ton captured. The add-on mechanism would require that all private carbon emitters can use a negative-carbon credit (a negaton) in place of buying a carbon credit in their cap-and-trade market or in place of paying their fossil-fuel tax. Every country in the climate coalition would be required to allow this.

The negative-carbon credits would be purchased by those subject to the highest carbon prices anywhere within the coalition. Competition would then set the price of negatons of carbon at the highest carbon charge imposed, and because the global price, P , is the average of all such charges, the price of a negaton would always be higher than P .

Q13: Do you think global carbon pricing is all that is needed? No. Although global carbon pricing facilitates cooperation and is an essential climate policy, it is of course not the only policy needed to effectively address climate change. Investments in green research are needed, too, and there is a role for some command-and-control style regulation, such as building standards. But the lack of a common commitment on carbon pricing is the primary source of the problem, and so correcting this is what this chapter and book is about.

Notes

1. This first section does not cover global carbon pricing directly. Rather, it orients the reader to a different way of thinking about international climate negotiations. For a more direct approach, start with "The Solution."

2. This can often be better than having it done by a higher government. However, Ostrom (2012) does not say this is always best: “People want to make me argue that community systems of governance are always the best: I will not walk into that trap.” At the global level, there is no government. The chapters in this book discuss how the countries of the world can self-organize a global system of climate *governance* that builds in reciprocity and allows trust to develop. For a global public good, a global governance (not government) is required.

3. Ad Hoc Working Group on the Durban Platform. Non-Paper, ADP.2014.6. July 7, 2014.

4. We often refer to Elinor Ostrom (although she did not directly speak about cooperation among governments) because she has pioneered some of the most important cooperation research and was particularly concerned with solving common resource problems without top-down help from a government, which is the problem faced at the international level. Also, her work on the role of reciprocity for cooperation has been robustly supported by the general science of cooperation, and it is relevant for, and fully consistent with others’ view of, international climate negotiations (see, e.g., MacKay et al., chapter 2, this volume).

5. These would likely be carbon-revenue credits—in other words, credit for collecting, say, \$1 million of revenue from carbon charges (see Stoft, 2009).

6. It would likely be best to start the negotiation process in a small group of big emitters so that certain basics are agreed on before involving the UN (see chapter 12, this volume).

7. See MacKay et al. (2015) for a similar treatment.

8. We note that what is proposed here is fully multilevel, in other words, “polycentric,” as Ostrom uses that term. There would be an international agreement on price, but each country would decide *all* of the pricing details and what to do with revenues and monitoring and enforcement details. Countries would be free to delegate responsibility to provinces, and provinces could delegate to cities. Most actual change would happen at the local and even individual level.

9. We agree with Parry (chapter 3, this volume) that local carbon-related externalities are a reason for allowing countries to use different carbon prices, but we think this is best handled by a uniform price commitment and some form of tradable pricing credits.

10. Gollier and Tirole (chapter 10, this volume) give the formula in terms of permits; our equation converts it to dollar transfers by using the price of permits.

11. It would cleverly engage in some cheap abatement and sell the resulting excess permits.

12. As they note, this also applies to negotiations over the Climate Fund contributions.
13. Of course, other aspects of the problem, particularly land use, technological progress, and the dissemination of accurate scientific information, are all consistent with carbon pricing.
14. Take the CFD strike prices for new onshore wind or new nuclear, roughly 90 per MWh, or recent RO bands (roughly 45 pounds per MWh subsidy). If this is substituting for gas-generated electricity at 450kg/MWh, then we have a subsidy of 45 pounds per 0.45 t CO₂ abated, or 100 pounds per ton.
15. "Climate and Carbon: Aligning Prices and Policies," OECD Environment Policy Paper 1, OECD Publishing, Paris, 2013.
16. An economy-wide tax shift involving a carbon price of 100 pounds per ton and reductions of other taxes would have little effect on the budgets of a typical family, although there might need to be distributional corrections, such as already exist (e.g., winter fuel payments targeted at all old people and other social mechanisms to protect those in "fuel poverty").

References

- Barrett, S. 1994. Self-enforcing international environmental agreements. *Oxford Economic Papers* 46:878–894.
- Bolton, G. E., and A. Ockenfels. 2000. ERC: A theory of equity, reciprocity and competition. *American Economic Review* 90 (1): 166–193.
- Bosetti, V., and J. Frankel. 2014, August. Sustainable cooperation in global climate policy: Specific formulas and emission targets. *Climate Change Economics* 5 (3).
- Bowles, S., and H. Gintis. 2013. *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton, NH: Princeton University Press.
- Cooper, R. N. 2008. The Case for Charges on Greenhouse Gas Emissions. Discussion Paper 08–10, Harvard Project on International Climate Agreements, Belfer Center for Science and International Affairs, Harvard Kennedy School.
- Covert, T., M. Greenstone, and C. R. Knittel. 2016. Will we ever stop using fossil fuels? *Journal of Economic Perspectives* 30 (1): 117–138.
- Cramton, P., A. Ockenfels, and J. Tirole. 2017. Policy Brief—Translating the collective climate goal into a common commitment. *Review of Environmental Economics and Policy*. In press.
- Cramton, P., and S. Stoft. 2010. International climate games: From caps to cooperation. Research Paper No. 10-07, Global Energy Policy Center.

Cramton, P., and S. Stoft. 2012a. Global climate games: How pricing and a green fund foster cooperation. *Economics of Energy & Environmental Policy* 1 (2).

Cramton, P., and S. Stoft. 2012b. How to fix the inefficiency of global cap and trade. *Economists' Voice*.

Depledge, J. 2000. The origins of the Kyoto Protocol. Prepared under contract to the UNFCCC, FCCC/TP/2000/2.

Edenhofer, O., J. C. Steckel, M. Jakob, and C. Bertram. Reading the writing on the wall: Coal and the Paris Agreement. Mimeo.

Frankel, J. 1998, July. The Kyoto Agreement on mitigating global climate change: Gains to all from participation by developing countries. Keynote speech, Executive Program on Climate Change and Development, Harvard Institute for International Development.

Fudenberg, D., and J. Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.

Green, F. 2015. Nationally self-interested climate change mitigation: a unified conceptual framework. Working Paper No. 224, Centre for Climate Change Economics and Policy.

Kraft-Todd, G. T., E. Yoeli, S. P. Bhanot, and D. G. Rand. 2015. Promoting cooperation in the field. *Current Opinion in Behavioral Sciences* 3:96–101.

MacKay, D., P. Cramton, A. Ockenfels, and S. Stoft. (October 15, 2015). Price carbon: I will if you will. *Nature* 526:315–316.

Nordhaus, W. 2015. Climate clubs: Overcoming free-riding in international climate policy. *American Economic Review* 105:4.

Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, UK: Cambridge University Press.

Ostrom, E. 2007. Collective Action Theory. In *The Oxford Handbook of Comparative Politics*, ed. Carles Boix and Susan C. Stokes, 186–208. Oxford: Oxford University Press.

Ostrom, E. 2009. A polycentric approach for coping with climate change. World Bank Report, WPS5095.

Ostrom, E. 2012. *The Future of the Commons*. London: Institute of Economic Affairs.

Poteete, A., M. Janssen, and E. Ostrom. 2010. *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton, NJ: Princeton University Press.

Schelling, T. C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Stiglitz, J. E. 2006a. A new agenda for global warming. *Economists' Voice* 3:7. doi:10.2202/1553-3832.1210.

Stiglitz, J. E. 2006b. Saving the Planet. In *Making Globalization Work*. New York: W. W. Norton & Company.

Stoft, S. 2008. *Carbonomics: How to Fix the Climate and Charge It to OPEC*. Nantucket, MA: Diamond Press.

Stoft, S. 2009. "Flexible Global carbon pricing: A backward-compatible upgrade for the Kyoto Protocol.," Working Paper No. RSCAS 2009/35, European University Institute Working Paper No. RSCAS 2009/35.

Stoft, S. 2010. *Dépasse Copenhague: Apprendre à coopérer*. Translated by Mathieu Thomas. Paris, L'Harmattan.

United Nations Framework Convention on Climate Change (UNFCCC). 1992. Accessible at: <https://unfccc.int/resource/docs/convkp/conveng.pdf>

Weitzman, M. L. 2014. Can negotiating a uniform carbon price help to internalize the global warming externality? *Journal of the Association of Environmental and Resource Economists* 1 (1/2): 29–49.

