Environmental & Resource Economics (2005) 31: 21–34

Piercing the Veil of Uncertainty in Transboundary Pollution Agreements*

CHARLES D. KOLSTAD
Department of Economics and Bren School of Environmental Science and Management, University of California, Santa Barbara, CA 93106-5131, USA (e-mail: kolstad@bren.ucsb.edu)

Accepted 18 November 2004

Abstract. This paper addresses the question of how uncertainty in costs and benefits affects the difficulty of reaching a voluntary agreement among sovereign states. A measure of "difficulty" is constructed related to side-payments necessary to make an agreement a Pareto-improving move. Using a simple model, it is shown that uncertainty actually makes agreement easier.

Key words: international environmental agreements, pollution, side-payments, transboundary uncertainty, treaties

JEL classifications: Q5, H4, D7, D8

1. Introduction

Environmental problems are plagued by uncertainty – uncertainty regarding physical and biological processes as well as uncertainty regarding societal costs and benefits. Certainly the public policy debate over climate change has focused on how certain we are about the problem. Some have used uncertainty to justify taking action now, before it is too late, and others have used uncertainty to justify delaying action. The economic literature suggests that uncertainty in and of itself should have little effect on action, provided decision-makers are risk neutral.[1]

At the international level, the problem is compounded by the fact that a group of decision-makers (countries) must voluntarily agree to solve the environmental problem. Although uncertainty (and learning) may influence the decision of individual countries, there is an issue as to how uncertainty affects the agreement process, if at all. For instance, if in fact the costs and benefits do not fall equally on all countries, it may be easier to forge an

*An earlier version of this paper was presented at the Conference on *Risk and Uncertainty in Environmental and Resource Economics*, Wageningen, The Netherlands, June 2002.

agreement before uncertainty is resolved, revealing who the winners and losers may be.

More specifically, consider a set of countries contributing different quantities of pollution which in the aggregate generates a global externality. This externality affects all of the countries, though to different degrees and possibly even some positively and some negatively. Think of greenhouse gas emissions leading to global warming. The "solution" for the environmental problem is that each country has to reduce emissions and the cost of doing so has to be shared in some equitable fashion. How much each country reduces emissions depends on their costs of emission control; how much each country contributes towards the cost of emission control depends on each country's benefits from controlling the problem as well as other issues of equity. An agreement to solve the problem would contain both an actions component, prescribing the emissions control undertaken by each country, as well as a sharing component, stipulated how costs are to be shared.

In this paper, we address the question of the effect of learning (increased information) on the *difficulty* of reaching a cooperative agreement on an environmental externality. The hypothesis, which is borne out in theory, is that learning increases the difficulty. It is easier to reach agreement behind the "veil of uncertainty"[2] about winning and losing than it is after the veil is lifted. This conclusion must be qualified since the concept of difficulty of reaching an agreement is not well defined. After all, with side-payments, any agreement with positive net benefits can be supported by an individually rational agreement among parties.[3] Without the possibility of side-payments, reaching agreement may be much more difficult.

Recognizing that side-payments can be politically difficult to implement, we measure the difficulty of reaching an environmental agreement by the gross size of side-payments necessary to make all parties voluntarily agree to a movement to the Pareto frontier. We find that uncertainty decreases, in expectation, this measure of difficulty. Put differently: difficulty, so measured, increases as uncertainty is reduced. This has a number of important implications. For one thing, it suggests that learning should induce earlier rather than later agreement. This stands in contrast to few results for the single decision-maker, that learning can justify postponing action. On the other hand, this result does not suggest that more action should be taken than is justified on cost–benefit grounds, only that delay can be very costly, particularly if delay results in a second-best agreement.[4]

2. Background

The literature on uncertainty in the context of international agreements spans economics and political science. The political science literature tends to be somewhat general, though deep, whereas the economics literature tends to build on idiosyncratic models of the bargaining process.

Young (1994) adopts the concept of the "veil of uncertainty" from Brennan and Buchanan (1985), who developed it for analyzing the emergence of constitutional rules in a society. Young (1994) suggests that uncertainty can be "good", serving to facilitate agreement on the core of international environmental agreements. To quote Brennan and Buchanan (1985, p. 30) "The uncertainty introduced in any choice among rules or institutions serves the salutary function of making potential agreement more rather than less likely...[An individual] will tend to agree on arrangements that might be called 'fair' in the sense that patterns of outcomes generated under such arrangements will be broadly acceptable, regardless of where the participant might be located in such outcomes." Though the authors are persuasive, neither offers an analytic version of their arguments.

Cooper (1989: p. 181) analyzes a century of public health agreements and comes up with the opposite conclusion: "So long as costs are positive and benefits uncertain, countries are unlikely to cooperate systematically."

Iida (1993) takes a game theoretic approach to international agreements, providing a nice review of how asymmetric information has entered into this literature, though primarily in the context of macroeconomic agreements. He distinguishes between *strategic* uncertainty, which is uncertainty about the types of opponents, and *analytic* uncertainty, which amounts to uncertainty about your own payoffs (as well as the payoffs of others). Depending how one defines "type", this may or may not involve uncertainty about one's own type. Although this type of uncertainty (analytic) is the focus of his paper, Iida points out that strategic uncertainty has dominated the literature on international agreements. Iida's interpretation of analytic uncertainty (sometimes termed *model* uncertainty) is that there are underlying characteristics of the international economic system which are unknown to all agents – these characteristics will be revealed *ex post* and will determine payoffs. Iida argues, through the use of a simple example, that analytic uncertainty will tend to retard international cooperation. Other literature is mixed (e.g., Frankel and Rockett 1988) as to whether analytic uncertainty tends to facilitate or retard international macroeconomic agreement.

Helm (1998) comes closest to analyzing the problem of this paper. He considers the case of an international agreement on acid rain, though much of the paper is independent of the application. He repeats many of the arguments above regarding the veil of uncertainty, and goes on to construct a simple two-country model of cooperation and non-cooperation In his example, he confirms Young's (1994) hypothesis that uncertainty is favorable for cooperation. He then goes on to investigate a repeated game and shows that generally a trigger strategy can support a cooperative equilibrium. What is not clear from the analysis is how the veil of uncertainty *facilitates* cooperation.

Na and Shin (1998) compare cooperation from both an *ex ante* (before uncertainty is resolved) and *ex post* (after uncertainty resolved) perspective. Their model is quite specific, though they do conclude that countries are unequivocally better off with *ex ante* negotiations. This is not quite the same as saying *ex ante* negotiations are easier. Further, the result depends on their very specific assumptions about cooperation. In their model, countries have costs of abatement and benefits from collective abatement. Cooperation is defined as non-cooperative bargaining among stable coalitions. They show that *ex ante*, when all countries view themselves as identical in expectation, the grand coalition (in a three country example) is stable and supports a true joint benefits maximum equilibrium. *Ex post*, after uncertainty has been resolved, the grand coalition is no longer stable since one or more countries may have an incentive to defect. Since bargaining is non-cooperative among coalitions, the joint payoff is bound to be lower. And thus their result.

Ulph and Maddison (1997), following up on earlier work by Ulph and Ulph (1996), focus on the role of learning on aggregate utility. This is an extension of work done by a variety of authors on the effect of learning on current period emission control when there is a single decision-maker (e.g., Kolstad 1996). Naturally, the cooperative equilibrium is analogous to the single decision maker and consequently they find information is always valuable. In examining non-cooperative Nash equilibria, they find more ambiguity – information may have negative value. These results are interesting and important but somewhat tangential to the problem being investigated here.

3. A Model of Agreements

Our interest here is in the negotiation of an environmental agreement – reaching international agreement on a common environmental problem. For simplicity, we consider two countries and a global externality. This is the simplest model possible that captures the essence of the problem and is generalizable to the n-country case. However, one shortcoming of focusing on only two countries is that we are unable to investigate the issue of coalition formation, fundamental to the analysis of Na and Shin (1998).[5]

The structure of the problem is two countries trying to reach a voluntary agreement regarding a common externality. There are two points of time, an *ex ante* period with uncertainty about payoffs, and an *ex post* period where uncertainty is resolved. The question we ask is simple: is it easier for the countries to reach an agreement prior to uncertainty being resolved or after uncertainty is resolved? To answer this question, we add structure to country payoffs and the allowable class of agreements. Further, we provide a definition of "difficulty to agree". The game theoretic literature is relatively silent on how difficult it is to reach an equilibrium; the literature tends to focus more on what those equilibria may be.

Although the problem here is couched as a two period problem, we are not really looking at a dynamic decision problem with information being acquired over time. It is more accurate to look at this as a comparative statics analysis, comparing two cases, one where an agreement is forged in a condition of uncertainty and the other where the agreement is made after the state of the world is revealed and that state is uncertain *ex ante*.

Focusing on two countries is somewhat restrictive since much of the richness of the literature on international agreements concerns three or more countries, where there are opportunities for free-riding (e.g., Chander and Tulkens 1997) and defection from coalitions. Nevertheless, much can be learned from a two-country model; extensions can follow.

3.1. QUANTIFYING THE DIFFICULTY TO AGREE

In examining a possible agreement between the two countries, we are limiting consideration to a fairly narrow set of agreements. Specifically, we are excluding contingent contracts, whereby obligations depend on what state-of-the-world is revealed. Further, we exclude side-payments. Including either or both of these would largely eliminate bargaining problems. Political realities make it unlikely that a contingent agreement will be pursued. Similarly, direct side-payments between countries are generally very difficult, politically.

Because payoffs are monetized, we can view the first-best outcome as an agreement which maximizes the joint payoff. This implicitly assumes risk neutrality on the part of decision-makers. The threat point is the non-cooperative Nash equilibrium. It may or may not be the case that the joint payoff maximum represents a Pareto improving move from the threat point. We represent the difficulty of reaching agreement as the value of the side-payment necessary to make the joint payoff maximum a Pareto improvement, which is then easily reached through cooperative agreement.

This might seem like an inconsistency. On one hand we are precluding side-payments from feasible agreements yet we are using the magnitude of side-payments as a measure of agreement difficulty. But if we look at the process of negotiation as having to overcome a hurdle represented by the size of the potential side-payment, the measure makes more sense. A reluctant country must be cajoled or offered advantages in other treaties or in some other way persuaded to overlook the side-payment otherwise necessary to reach an agreement. Thus the size of the side-payment is a measure of the difficulty that will be encountered in reaching an agreement. If no side-payment is necessary, reaching an agreement should not be a problem; agreement is individually rational without any transfers. If the side-payment is small, then it should be relatively easy to convince the potential recipient to reach an agreement. If the side-payment is very large, then it may be very difficult indeed to overcome this barrier.[6]

This is represented graphically in Figure 1 where there are two countries, W and L, located at point B, the *status quo* or threat point. Think of country L as actually benefiting from global warming; thus any emission control by either country will only make L worse off (a loser). On the other hand, country W can benefit significantly from modest rollback in emissions (a winner). Country W can benefit but only at the expense of country L. Thus the joint payoff maximum involves some loss in payoff for L and a large gain for W. If country W pays country L an amount AB, then country L will be indifferent between staying at B *versus* moving to C. Thus the payoff necessary to facilitate a move to a joint payoff maximum is AB. We term this minimal side-payment necessary to facilitate a move to the joint payoff maximum, the "agreement-inducing side-payment".

This concept of "difficulty to agree" is critical to the results of this paper. It is not the perfect measure. Payoffs to agreements are often difficult to quantify; further there may be asymmetries in how different countries view the same payoffs. There may be other issues that are important to a country, such as maintaining sovereignty, preserving a way of life that depends on environmental quality, or distributional issues.

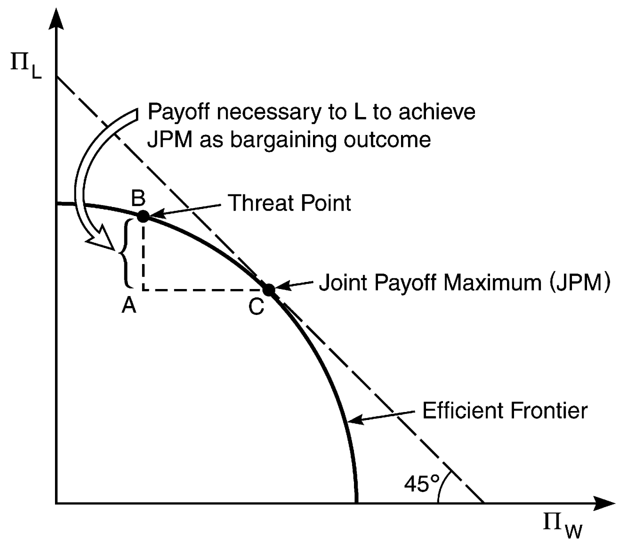

Figure 1. Illustration of payoff necessary to achieve agreement.

Furthermore, the assumption here that side-payments are not possible is somewhat of an oversimplification. Side-payments, though difficult, are possible. The Montreal Protocol for ozone protection contains provisions for payments from the developed world to the developing world to help defray the costs of complying with the protocol. The proposed Kyoto Protocol to prevent climate change contains several such provisions. For instance, the Clean Development Mechanism (CDM) is effectively a side-payment to the developing world – the developed world pays for some activities in the developing world designed to reduce global carbon emissions. However, the CDM is a disguised payment, much easier than a direct payment. Furthermore, one of the primary objections of the US Government to the Kyoto Protocol is that it provides such payments to Russia. Even disguised, side-payments to Russia are not very palatable politically. Another vehicle for disguised side-payments is "issue linkage" – connecting a concession on one agreement to a gain on another, unrelated agreement.[7] For instance, Krutilla (1967) suggests that in negotiations between the US and Canada over development of the Columbia River basin, Canada received an attractive deal on gains from development in exchange for providing better defense support to the US.

All of these nuances of negotiations aside, the fact that side-payments are difficult, suggests that the magnitude of the side-payment necessary to make an agreement in the best interest of all parties is a measure, albeit imperfect, of the magnitude of the obstacles that stand in the way of reaching agreement. Thus, in our simple model, we use this proxy for "difficulty of reaching agreement". In comparing two potential agreements, the one with the higher agreement-inducing side-payment will be considered to be the potential agreement that is more difficult to finalize.

3.2. A SIMPLE TWO-COUNTRY MODEL

Let there be two countries, A and B, and two time periods, between which nature moves and reveals the state of the world. The state of the world affects the damages from the externality (equivalently, the benefits from abatement). Each country may undertake abatement activities, q_i. Aggregate abatement, $Q = q_A + q_B$, determines benefits to each country i, $\theta_i Q$, where θ_i is a random variable, drawn independently (for each i) from a uniform distribution over [0,1], which implies a mean of 0.5 and a variance of 1/12. Other distributions would work and countries need not share the same distribution, though the essence of the problem would remain unchanged. Basically, this is our representation of the uncertainty associated with damages from the environmental externality. Neither country is aware of what in fact the benefits of abatement may be – a country's payoff from abatement is uncertain.

Costs of providing abatement are known (this could of course be made uncertain) and are given by $c_i q_i^2/2$, where c_i is a cost shifter for country *i* and without loss of generality, we assume $c_A \leq c_B$. One country may be able to abate cheaply; another more expensively. This is a very specific type of cost function. The only difference between the two countries is in the slope of the marginal cost of emission control. Obviously other simple linear representations of costs would also be possible.

This terminology results in the following payoff function for each country (which is similar to the formulation of Na and Shin 1998):

$$\Pi_i = \theta_i Q - c_i\, q_i^2/2 = \theta_i(q_A + q_B) - c_i\, q_i^2/2, \qquad i = A, B. \tag{1}$$

Ex ante, before uncertainty is resolved, expected gross benefits from abatement are the same for both countries, $Q/2$. *Ex post*, after uncertainty is resolved, the two countries will in all likelihood have different benefits. Thus the payoff function used for bargaining will be different, depending on whether the bargaining is done before or after uncertainty is resolved.

4. Analysis

We are interested in comparing two situations: uncertainty and full information. In the context of our model, uncertainty is manifest in the random variable θ_i. Although *ex ante*, before uncertainty is resolved, both agents will have the same expectations on θ_i, the countries will behave differently because of differences in costs. We measure the agreement-inducing side-payment assuming uncertainty over θ_i. We then consider the *ex post* case, where nature has moved and revealed the value of θ_i to each country. In this case, the agreement-inducing side-payment is a function of realized θ_A and θ_B. We then consider the *ex ante* situation again and compute the expectation of the agreement-inducing side-payment over the random variables θ_A and θ_B. We then compare the two measures: the agreement-inducing side-payment in an atmosphere of uncertainty (behind the veil of ignorance) and the expected value of the agreement-inducing side-payment that will apply after uncertainty is resolved.

It is straightforward to show that the Nash equilibrium, given the payoff structure in (1), is

$$q_i^N = \theta_i/ci, \qquad i = A, B \tag{2a}$$

$$\Pi_i^N = \theta_i^2/(2c_i) + \theta_i\theta_{<i>}/c_{<i>}, \qquad i = A, B \tag{2b}$$

where $<i>$ means the index other than *i*. If θ_i is not known (as is the case before uncertainty is resolved), it is appropriate to take expectations over θ_i in Equation (2).

The joint payoff maximum can be obtained in an equally straightforward, though slightly more computationally tedious way, by summing the payoffs in Equation (1) over the two countries, and solving the first order conditions:

$$q_i^{\mathrm{C}} = (\theta_{\mathrm{A}} + \theta_{\mathrm{B}})/c_i, \qquad i = \mathrm{A}, \mathrm{B} \tag{3a}$$

$$\Pi_i^{\mathrm{C}} = \Pi_i^{\mathrm{N}} + \theta_i^2/c_{<i>} - \theta_{<i>}^2/(2c_i), \qquad i = \mathrm{A}, \mathrm{B}. \tag{3b}$$

Although the joint payoff ($\Pi_{\mathrm{A}}^{\mathrm{C}} + \Pi_{\mathrm{B}}^{\mathrm{C}}$) must be at least as high as the total payoff under non-cooperation ($\Pi_{\mathrm{A}}^{\mathrm{N}} + \Pi_{\mathrm{B}}^{\mathrm{N}}$), it may easily be the case that for one of the countries, non-cooperation yields a higher payoff than cooperation, at least without side-payments. Let T_{i} be the minimum non-negative payment to country i necessary to make cooperation more attractive than non-cooperation, $\Pi_i^{\mathrm{C}} + \mathrm{T}_i \geq \Pi_i^{\mathrm{N}}$:

$$T_i = \max\{0, \Pi_1^{\mathrm{N}} - \Pi_i^{\mathrm{C}}\}, \qquad i = \mathrm{A}, \mathrm{B}. \tag{4}$$

Clearly, either T_i is zero or $T_{<i>}$ is zero; both may be zero but both cannot be positive (otherwise, the joint payoff maximum would be less than the non-cooperative joint payoff). By construction if T_{i} is positive, the gains to $<i>$ must be sufficient to make the payment T_i, while still keeping the payoff from cooperation higher than the payoff from non-cooperation:

$$\Pi_i^{\mathrm{C}} - T_{<i>} + T_i \geq \Pi_i^{\mathrm{N}}, \qquad \text{for } i = \mathrm{A}, \mathrm{B}. \tag{5}$$

We now turn to the question of the agreement-inducing side-payments in an atmosphere of uncertainty. We will consider two cases: one with no resolution of uncertainty, where any agreement must be reached under conditions of uncertainty; and a second case where the agreement is made after resolution of uncertainty. We are interested in comparing the aggregate expected side-payments in these two cases. As we have argued here, whichever case has lower expected side-payments will be presumed to be the easier agreement to forge.

Inspection of Equation (3b) indicates that country B always does better under cooperation, no matter what the realization of θ_i (due to the fact that $c_A \leq c_B$) is. Thus the side-payment (if any) goes to country A. The expected value of that (conditional on the side-payment being positive) is easily computed:

$$E(\Pi_i^{\mathrm{C}} + \Pi_i^{\mathrm{N}}) = E(\theta_{\mathrm{A}}^2)/c_{\mathrm{B}} - E(\theta_{\mathrm{B}}^2)/c_{\mathrm{A}} = (2c_{\mathrm{A}} - c_{\mathrm{B}})/(6c_{\mathrm{A}}c_{\mathrm{B}}) \tag{6}$$

Thus the side-payment necessary to achieve agreement under persistent uncertainty without learning (the expected agreement-inducing side-payment with uncertainty – EAISPU) is

$$\mathrm{EAISPU} = \begin{cases} (c_{\mathrm{B}} - 2c_{\mathrm{A}})/(6c_{\mathrm{A}}c_{\mathrm{B}}) & \text{if } c_{\mathrm{A}} < c_{\mathrm{B}}/2 \quad (7\mathrm{a}) \\ 0 & \text{otherwise} \quad (7\mathrm{b}) \end{cases}$$

Depending on the resolution of uncertainty, *ex post* side-payments may be necessary to reach agreement. We are interested in the expected value of the agreement-inducing side-payments, where the expectation is taken over the random variables θ_A and θ_B. By inspection of Equations (3b) and (4), we can see that the agreement-inducing side-payment for country i will be positive if and only if

$$\theta_i^2/c_{<i>} < \theta_{<i>}^2/(2c_i) \tag{8a}$$

or, equivalently

$$\theta_{<i>} > \theta_i\sqrt{2c_i/c_{<i>}} \tag{8b}$$

Consider the expected size of side-payments to country A:

$$\begin{aligned} E(T_A) &= \int_0^1 \int_{\theta_A\sqrt{2c_A/c_B}}^1 T_A d\theta_B d\theta_A \\ &= \int_0^1 \int_{\theta_A\sqrt{2c_A/c_B}}^1 \left[\theta_B^2/2c_A - \theta_A^2/c_B\right] d\theta_B d\theta_A \\ &= \left[\frac{1}{6c_A} - \frac{1}{3c_B} + \frac{\sqrt{2c_A/c_B}}{6c_B}\right] \end{aligned} \tag{9}$$

Using a symmetry argument, the expected side-payment to B is

$$E(T_B) = \frac{1}{6c_B} - \frac{1}{3c_A} + \frac{\sqrt{2c_B/c_A}}{6c_A} \tag{10}$$

This implies that in expectation, if agreement is deferred until after uncertainty is resolved, the expected agreement-inducing side-payment with learning (EASIPL) is given by

$$\text{EAISPL} = E(T_A) + E(T_B) = \frac{c_A\left[\sqrt{2c_A/c_B} - 1\right] + c_B\left[\sqrt{c_B/2c_A} - 1\right]}{6c_Ac_B} \tag{11}$$

This brings us to the main result of the paper:

Proposition: In a two-country bargaining arrangement as described above, the expected value of side-payments necessary to induce agreement after uncertainty is resolved, is strictly greater than the expected value of side-payments necessary to induce agreement when uncertainty persists and is not resolved:

$$\text{EASIPL} > \text{EASIPU} \tag{12}$$

Proof. There are two cases: (1) $c_A < c_B/2$ and (2) $c_A < c_B/2$. For case (2), EASIPU is 0 from Equation (7b). Furthermore, the terms within each radical in Equation (11) are greater than or equal to 1 (and the argument of the second radical is strictly greater than 1). Thus EASIPL > 0.

For case 2, we have from Equation (7a);

$$\text{EAISPL} - \text{EASIPU} = \frac{c_A\left[\sqrt{2c_A/c_B} + 1\right] + c_B\left[\sqrt{c_B/2c_A} - 2\right]}{6c_A c_B} \tag{13}$$

Because $c_A < c_B/2$, $\sqrt{2c_B/c_A} > 2$, which implies the RHS of Equation (13) is positive. ■

In interpreting Equation (11) or (13), consider the simplest version of it: the two countries have identical cost coefficients. With $c_A = c_B = c$, Equation (11) reduces to

$$\text{EAISPL} = \frac{\sqrt{2} - 1}{c} \tag{14}$$

which is, of course, positive. Furthermore, the larger the value of c, the smaller are the expected side-payments. The interpretation of this result is that when c is larger, the significance of the uncertainty in benefits is reduced. Referring back to Equation (4a), we see that larger costs induce smaller levels of abatement, as one would expect. With smaller levels of abatement, the significance of uncertainty in gross benefits is reduced; thus the reduction in the expected agreement-inducing side-payments.

Another question concerns the heterogeneity of costs. We would expect agreement to be more difficult, the more divergent are costs for the two countries. Let $\Delta = c_B - c_A$. In order to not confound the problem by introducing an overall increase in costs, consider

$$\frac{\partial(\text{EAISPL} - \text{EASIPU})}{\partial\Delta} = \frac{\partial(\text{EAISPL} - \text{EASIPU})}{\partial c_B} - \frac{\partial(\text{EAISPL} - \text{EASIPU})}{\partial c_A} \tag{15}$$

Differentiating Equation (13), we obtain (after some manipulation),

$$\frac{\partial(\text{EAISPL} - \text{EAISPU})}{\partial\Delta} = \frac{1}{6\sqrt{2}} c_A^{-2.5} c_B^{-2.5}\{c_A c_B(c_B - c_A) + 3(c_B^3 - c_A^3)\} + 6c_A^{-2}c_B^{-2}(2c_B^2 - c_A^2) \tag{16}$$

which is clearly positive when $c_B \geq c_A$.

The interpretation is consistent with the results that the more heterogeneous the countries, the more difficult it may be to reach agreement.

5. Conclusions

With a simple two-country model, we have investigated the question of whether uncertainty retards or enhances reaching voluntary agreement on transboundary externalities. This has been suggested in the literature but never explored in a formal analytic model. One problem in answering this question relates to how one measures the difficulty of reaching agreement. In this paper, we offer the size of "agreement-inducing side-payments" as a measure of the difficulty of reaching agreement.

Our results confirm prior suggestions in the literature. Uncertainty makes agreement easier. Future work would include investigating other representations of payoffs and including a more explicit treatment of the dynamic nature of commitment in an international environmental agreement.

Acknowledgements

I am grateful to Scott Barrett and Alistair Ulph for enlightening conversations; naturally, they bear no responsibility for the quality of this paper. The paper also benefited from valuable suggestions from two anonymous referees. Able research assistance from Sofia Franco is gratefully acknowledged.

Notes

1. The story is different if learning is taking place over time and action can have irreversible consequences. In this case, the environmental protection decision by individual decision makers can be biased upwards or downwards, depending on the nature of the learning process and the irreversibilities. See Kolstad (1996), among others.
2. Brennan and Buchanan (1985), among others, use the term "veil of uncertainty" in this context, a variant of the Rawlsian "veil of ignorance".
3. If the group as a whole is better off, then there must be a way to share those gains so that each country individually is better off.
4. The "timing" of regulatory action in the presence of learning has been addressed by a number of authors. A key variable is the extent to which there are irreversibilities, either environmental or in terms of pollution control capital investments. These irreversibilities operate on the timing question in opposite directions with respect to the base case of timing control based purely on costs and benefits without any learning: environmental irreversibilities call for the acceleration of control whereas abatement capital irreversibilities imply the optimality of the delay of control.

5. The bulk of the literature on international environmental agreements focuses on voluntary participation in an agreement – how many countries are in and how many are out (e.g., Ulph 2004; Barrett 1994). That issue is not treated here, since we assume two countries only.
6. We are unaware of other work that uses side-payments as a measure of difficulty-to-agree. There are, however, other papers that explicitly use side-payments as a determinant in equilibrium. See, for example, Kowalczyk and Sjöström (1994) and Petrakis and Xepapadeas (1996).
7. Issue linkage has been addressed in the political science literature – see Haas (1980) and McGinnis (1986) for example. Folmer and von Mouche (1994) have addressed the problem from a game theoretic perspective – interconnected games.

References

Barrett, S. (1994), 'Self-Enforcing International Environmental Agreements', *Oxford Economic Papers*, **46**, 878–894.

Brennan, G. and J. M. Buchanan (1985), *The Reason of Rules: Constitutional Political Economy*. Cambridge University Press, Cambridge, UK.

Chander, P. and H. Tulkens (1997), 'The Core of an Economy with Multilateral Environmental Externalities', *International Journal of Game Theory* **26**, 379–401.

Cooper, R. N. (1989), 'International Cooperation in Public Health as a Prologue to Macroeconomic Cooperation', in R. N. Cooper et al., eds., *Can Nations Agree?* Brookings Institution, Washington, DC, pp. 78–254.

Folmer, H. and P. von Mouche (1994), 'Interconnected Games and International Environmental Problems', *Annals of Operations Research* **54**, 97–117.

Frankel, J. A. and K. Rockett (1988), 'International Macroeconomic Policy Coordination: When Policy Makers Do Not Agree on the True Model', *American Economic Review* **78**, 318–340.

Haas, E.B. (1980), 'Why Collaborate – Issue Linkage and International Regimes', *World Politics* **32**, 357–405.

Helm, C. (1998), 'International Cooperation behind the Veil of Uncertainty', *Environmental and Resource Economics* **12**, 185–201.

Iida, K. (1993), 'Analytic Uncertainty and International Cooperation: Theory and Application to International Economic Policy Coordination', *International Studies Quarterly* **37**, 431–457.

Kolstad, C. (1996), 'Fundamental Irreversibilities in Stock Externalities', *Journal of Public Economics* **60**, 221–233.

Kowalczyk, C. and T. Sjöström (1994), 'Bringing GATT into the Core', *Economica* **61**, 301–317.

Krutilla, J. V. (1967), *The Columbia River Treaty: The Economics of International River Basin Development*. Johns Hopkins University Press for Resources for the Future, Baltimore, MD.

McGinnis, M. D. (1986), 'Issue Linkage and the Evolution of International Cooperation', *Journal of Conflict Resolution* **30**, 141–170.

Na, S., and H. S. Shin (1998), 'International Environmental Agreements under Uncertainty', *Oxford Economic Papers* **50**, 173–185.

Petrakis, E. and A. Xepapadeas (1996), 'Environmental Consciousness and Moral Hazard in International Agreements to Protect the Environment', *Journal of Public Economics* **60**, 95–110

Ulph, A. (2004), 'Stable International Environmental Agreements with a Stock Pollutant, Uncertainty and Learning', *Journal of Risk Uncert*. **29**, 53–73.

Ulph, A., and D. Maddison (1997), 'Uncertainty, Learning and International Environmental Policy Coordination', *Environmental and Resource Economics* **9**, 451–466.

Ulph, A. and D. Ulph (1996), 'Who Gains from Learning about Global Warming?', in Ekko C. van Ierland and Kazimierz Gorka, eds., *Economics of Atmospheric Pollution*, Ch. 3. Springer-Verlag, Berlin, pp. 31–67.

Young, O. (1994), *International Governance: Protecting the Environment in a Stateless Society*. Cornell University Press, Ithaca, NY.