Uncertainty and Learning in a Strategic Environment: Global Climate Change*

Erin Baker†
University of Massachussetss, Amherst

May 21, 2003

Abstract

Global climate change is rife with uncertainties. Yet, we can expect to resolve much of this uncertainty in the next 100 years or so. Therefore, current actions should reflect the value of flexibility. Nevertheless, most models of climate change, particularly game-theoretic models, abstract from uncertainty. Additionally, most analysis of climate change abstracts completely from the idea that the uncertainties may not be perfectly correlated across nations. A model of the impacts of uncertainty and learning in a non-cooperative game shows that the level of correlation of damages across countries is crucial for determining optimal policy.

JEL classification: C72; D81; Q28

Keywords: Climate change; Noncooperative games; Information and uncertainty; Environmental policy

*Running Head: Uncertainty and Learning in Climate Change

†Correspondence Address:Erin Baker, 220 ELab, University of Massachusetts, Amherst,MA 01002

1 Introduction

How should the uncertainty inherent to global climate change impact optimal policy? Policy makers and advocates seem to fall into two groups regarding the implications of uncertainty. One group argues that the most prudent course is to wait for more knowledge and then act. The other group invokes the precautionary principal, saying the world should act to reduce emissions now, before we are engulfed in catastrophe. One way to find a middle ground between these two groups is to consider the impacts of what – and how much – we expect to learn.

Climate change research is proceeding rapidly. It can reasonably be expected that a great deal – if not all – of the scientific and economic uncertainty will be resolved in the next 100 years or so. As more is learned about climate change, policies can be tailored accordingly. Therefore, current actions should reflect the value of flexibility. The question is, what preserves more flexibility – reducing emissions now or waiting for information?

A selection of papers have considered the effect of learning about the climate on near-term emissions decisions. They have found that the possibility of learning combined with the ability to alter future behavior implies that a hedging strategy is optimal: reduce emissions a small amount now, then wait and see what happens. All of these papers, however, assume a single decision maker.

This paper considers the impact of learning on near-term emissions decisions in a non-cooperative game. Given that the United States has determined not to sign the Kyoto Protocol, and that most developing nations are exempt from the agreement, a non-cooperative game is an accurate description of the current situation. Current policy decisions influencing climate change are made independently by each nation. Understanding the non-cooperative equilibrium is an important part of crafting a cooperative agreement, since the non-cooperative equilibrium is often the "threat-point" that holds a cooperative agreement together.

In the model presented in this paper two non-cooperative players attempt to maximize utility by choosing a level of emissions, taking the other's emissions as given. The damages from climate change increase in the stock of emissions. The exact relation between emissions and damages is assumed to be initially uncertain with information revealed over time. The players know that they will be learning and, therefore, use a framework of sequential decision making under uncertainty. Each player makes decisions in the short run under uncertainty, knowing that the uncertainty will be (partially) resolved in the future and that future decisions will be made under greater certainty (See figure 1).

The model addresses the following questions. How does an increase in the amount the decision maker expects to learn affect equilibrium emissions? How does an increase in uncertainty affect equilibrium emissions? The answer to both of these questions depends on how the damages are correlated across countries. Most analysis of climate change abstracts from the idea that the uncertainties may not be perfectly correlated across nations. The single most common measure of climate change is global mean temperature, which clearly is identical for all nations, yet individual nations can have temperature variations very different from the global mean. Regional analyses indicate that not only will damages differ from region to region, but that it is also difficult to predict to what degree regional damages will vary. Thus, this paper considers three cases. If damages are highly positively correlated across countries then first period equilibrium emissions increase in both uncertainty and learning. If damages are partially independent or negatively correlated across countries, the results are reversed. Hence, a third question is addressed: how does a change in the perceived level of correlation of damages across countries affect emissions?

The rest of the paper is organized as follows. In the next section this work is put in the context of previous literature. In section 3 the model is described. In section 4 three propositions that answer the questions posed above are presented. Section 5 concludes with some policy interpretations.

2 Human-induced Climate Change

2.1 Uncertainty and Learning in Climate Change

The most comprehensive climate change models are the large and complex integrated assessment models such as DICE [15], RICE [17], and MERGE [11]. These models have been influential in the policy debate over responses to climate change, as they have provided important insights into the costs and benefits of mitigation. It has been difficult, however, to include uncertainty in such complex models. Nevertheless, two topics related to uncertainty, learning, and climate change have been covered in the literature: the value of information and the optimal timing of emissions reductions. A number of papers (see [16] and [12]) have concluded that there is considerable value in gaining information about climate and damages. Hammitt *et al.*[6] consider the optimal sequential decision strategy in a simple decision analytic framework, and Scott *et al.* [22] emphasize that an adaptive policy of "act, then learn, then act" is optimal. Kolstad [8] looks explicitly at the effect of learning about the climate in the face of irreversibilities. He finds that the possibility of learning causes a downward bias in control, i.e. an increase in emissions over the no-learning case. He also finds little evidence that the stock effect affects current emissions. Similar results are found in Ulph and Ulph [25] and Manne [10], although Gollier *et al.*[4] point out that these results can be reversed by considering utility functions with high levels of prudence.

2.2 Game Theory and Uncertainty

Since climate change exhibits both uncertainty and multiple non-cooperative players, a game-theoretic model with uncertainty in the state of nature is most appropriate. Nevertheless, very few such models exist. In the literature on fisheries, Sandler and Sterbenz [21] find that harvest uncertainty will reduce exploitation in a tragedy of the commons game if players are risk averse and if all actions are *ex-ante*. Looking at climate change, Na and Shin [14] compare the expected utility of nations when coalitions are formed *ex ante* versus *ex post*. They find that since countries are more likely to be facing similar conditions *ex ante* the possibility of coalition formation is enhanced the sooner negotiations take place. Their one-period model does not consider learning or commitment issues when the negotiations take place *ex ante.* Using a numerical model, Hammitt and Adams [5] find that the expected benefits of a non-cooperative solution are very close to the expected benefits of the cooperative solution in a game with perfect learning and perfect correlation across players.

In the real options literature, there have been some recent attempts to model the effect of strategic interactions on the value of waiting to make irreversible decisions (See [26],[7], and [23]). A key assumption in all these papers is that there is a single uncertain variable that effects all players. Reingannum [18], on the other hand, assumes that the random variables affecting each player are independent in her work combining game theory and uncertainty to analyze R&D races. In climate change, however, the level of correlation between players turns out to be an important issue.

The importance of the level of correlation of damages was first pointed out by Ulph and Ulph [24]. Using a simulation they show that emissions under perfect learning can be lower than under no learning when damages are negatively correlated. This paper uses analytical methods to extend their results to closed-loop games, and demonstrate the more general effects of an *increase* in learning, as well as the effects of an increase in uncertainty.

3 The Model

Consider a two-period closed-loop dynamic game with uncertainty in the state of nature. The uncertainty is assumed to be resolved before the 2nd period, allowing players to adjust their behavior. Thus, it is a model of sequential decision making under uncertainty.

The game is restricted to two periods to focus on the effect of closed-loop strategy and uncertainty. The two periods can be thought of as *now* when there is uncertainty about the nature of damages caused by global climate change, and *later*, when that uncertainty will be resolved. There are two players, X and Y, representing individual nation states. Each player creates emissions[1] through energy use and other economic activity. Emissions are released into the atmosphere and dissipate very slowly at rate γ. Player X has emissions x_i, for periods $i = 1, 2$ and Player Y's emissions are represented as y_i. The stock of emissions in the second period is $s = \gamma (x_1 + y_1) + x_2 + y_2$. Damages are assumed to be zero in the first period. The time discount factor from one period to the next δ, is common to both players.

The players are trying to balance the private benefits of emissions against the uncertain damages caused by climate change, as follows

$$b(x_1) + \delta \, \mathrm{E} \left[b\left(x_2\right) - \varepsilon D(s) \right] \tag{1}$$

The benefits from emissions are represented by the function $b(x)$, the net benefits of the energy use that creates an emission level of x. $b(\cdot)$ is assumed to be strictly concave – implying that the marginal cost of reducing emissions is increasing – and to have a unique maximum point, commonly referred to as the business as usual level.

$\varepsilon D(s)$ represents the damages from climate change as a function of the stock s of carbon in the atmosphere. $D(\cdot)$ is increasing, strictly convex, and deterministic – implying that the *shape* of the damage function is known. The uncertainty is represented by ε, a stochastic shift parameter that multiplies the deterministic portion of the damage function, $D(\cdot)$. ε can be indexed by x or y – the uncertain damages are not necessarily the same for both players. Damages increase in ε for any given level of emissions: If $\varepsilon = 0$ then there are no damages; if ε is large then damages will be high when the stock of emissions is high.

The players use a hedging strategy in a non-cooperative framework. The solution concept is feedback Nash equilibrium. Emissions in the second period depend on the realized value of the random variable and on first period emissions. In the first period, emissions are chosen recognizing

[1] Emissions are defined as a scalar measure of carbon-equivalents.

that the second period emissions of both countries depend on first period emissions. Thus, the model is "closed loop". The game is played with complete information; specifically the distributions of the random variables ε_x and ε_y are common knowledge.

2nd period equilibrium emissions are characterized by the first order conditions:

$$b'(x_2) = \varepsilon_x D'(s) \quad (2)$$

$$b'(y_2) = \varepsilon_y D'(s) \quad (3)$$

The stock $s = \gamma(x_1 + y_1) + x_2 + y_2$ can be written as a function of the first-period emissions and the two stochastic shift parameters, since second period emissions are completely determined by these variables In the first period each player maximizes the benefit from first-period emissions *plus* the discounted expected benefit of the second period emissions *minus* the discounted expected damages from the stock. Again, each player takes the other player's first-period emissions as given, while the other player's second-period emissions are a commonly known function. $\mathrm{E}\{\cdot\}$ is the expectation operator. Player X's optimization problem is as follows.

$$\max_{x_1} b(x_1) + \delta \,\mathrm{E}\{b(x_2(x_1 + y_1, \varepsilon_x, \varepsilon_y)) - \varepsilon_x D(s(x_1 + y_1, \varepsilon_x, \varepsilon_y))\} \quad (4)$$

Player Y solves a similar problem simultaneously. The first order condition derived from (4) is

$$b'(x_1) = \delta \,\mathrm{E}\left\{\varepsilon_x D'(s(x_1 + y_1, \varepsilon_x, \varepsilon_y))\left[\gamma + \frac{\partial y_2}{\partial x_1}\right]\right\} \quad (5)$$

Marginal benefits equal expected marginal damages. X's first period emissions are based on X's damage alone, rather than on the total damage $\varepsilon_x + \varepsilon_y$. This is the tragedy of the commons. Note that the marginal effect of X's first period emissions on X's second period emissions drops out of the equation by the envelope theorem[2]. This is not true, however, of the effect of X's first period emissions on Y's second period emissions. This is the effect of strategic behavior seen in a "closed loop" model. Each player realizes that an increase in first period emissions will cause the other player to reduce second period emissions (i.e. $\frac{\partial y_2}{\partial x_1} \leq 0$; see appendix for proof). Therefore internalized expected marginal damages (the right-hand side of (5)) are *additionally* lower than in a typical one-period or open-loop Tragedy of the Commons. When players are sophisticated there is not only the tactical tragedy of the commons, there is also a strategic increase in emissions reflecting the knowledge that higher emissions now will force their opponents to emit less in the future.

4 Results

We determine whether initial emissions, x_1 will increase or decrease with uncertainty. Since the payoff functions are linear in the random variables ε_x and ε_y, this is equivalent to determining whether emissions will increase or decrease with a more informative signal (see Blackwell [2] for definition), i.e. as more is expected to be learned in the sense described in Epstein [3] (see Baker [1] for proof).

[2] $(b'(x_2) - \varepsilon_x D'(s))\frac{\partial x_2}{\partial x_1} = 0$ by the first order conditons in equation 4.

Define an increase in risk to be a mean-preserving spread (see Rothschild-Stiglitz [19]). The expected value of a convex function increases in risk, while the expected value of a concave function decreases. Thus *expected* marginal damages will increase (decrease) with uncertainty if marginal damages (the quantity inside the brackets in (5) above) are convex (concave) in ε. Since b is strictly concave, first period emissions will increase (decrease) as expected marginal damages decrease (increase)[3]. The results are presented in the following propositions. $b(\cdot)$ and $D(\cdot)$ are assumed to be quadratic in each case[4]. See the appendix for the proofs.

Proposition 1 *Perfect Correlation: Let $\varepsilon_x = \varepsilon_y = \varepsilon$ where ε is a non-negative random variable. Then symmetric first period equilibrium emissions (x_1^*, y_1^*) are **increasing** in uncertainty and informativeness.*

Exploitation of the commons may increase in uncertainty when there is learning.

Proposition 2 *Independence: Let ε_x and ε_y be independent non-negative random variables. Then first period equilibrium emissions are **increasing** in own-uncertainty (or in increased learning about ε_x) when $\varepsilon_y < -\frac{2}{3}\frac{b''}{D''}$ almost surely, and **decreasing** in opponent's uncertainty (or in increased learning about ε_y) unconditionally.*

Proposition 3 *Perfect Negative Correlation: Let $(\varepsilon_x, \varepsilon_y) = (\varepsilon, 2-\varepsilon)$ where ε is a symmetric, non-negative random variable with mean 1. Then symmetric first period equilibrium emissions (x_1^*, y_1^*) are **decreasing** in uncertainty and informativeness.*

What do these propositions tell us? Proposition 1 simply echoes the results in the single-decision maker literature: if the irreversibility of emissions doesn't bite, then the possibility of learning causes emissions to increase. When damages are perfectly correlated, this result is not substantially changed by including a strategic framework.

Propositions 2 and 3 reverse the findings in the single decision maker literature and expand the results from Ulph and Ulph [24]. In a strategic situation, the possibility of an opponent's learning can be damaging enough to outweigh the benefits of own-learning. In a game, the level of correlation across players fundamentally affects behavior in the face of uncertainty. The implication for climate change policy is that if the coefficient of correlation across nations is low enough, the possibility of learning may indicate an upward bias in control (i.e. a decrease in emissions) rather than the downward bias found in the literature. This result stands even without strong irreversibility and without assuming risk averse players. Insight into the different results are provided in the next sections.

4.1 Single Decision Maker

In order to understand the results from the non-cooperative game, it is useful to first consider the single decision maker case. Assume for a moment that 2nd period emissions are not constrained to

[3] See Rothschild-Stiglitz [20], Laffont [9], or Baker [1] for a proof.

[4] This assumption is sufficient but not necessary. If the third derivatives do not vanish, however, a statement of the proposition becomes very murky.

be non-negative – that emissions could be reduced below zero, perhaps by using carbon sinks such as forests. The methodology in this paper confirms that in this case optimal first period emissions increase in uncertainty and in informativeness.

To understand this, think of an extreme example, where damages from climate change will either be zero or “high”. In the second period, if the single decision maker learns that damages are zero, then second period emissions will be set at a business-as-usual level, independent of first period emissions. The benefits forgone by reducing emissions in the first period are gone forever. On the other hand, if the decision maker learns that damages are in fact “high”, then second period emissions will be reduced according to the damages, but also taking into account the total stock of emissions and therefore first period emissions. Thus, to the degree that first period emissions were too high (in hind sight) this can be partially made-up for in the second period. Therefore, the decision-maker wants to “err on the side of too high”. Because emissions can be reduced in the 2nd period, the bad outcomes are ameliorated. Therefore a *mean preserving spread* of the distribution of ε decreases the expected value of the damages, inducing an increase in 1st period emissions.

Figure 2 illustrates this in another way. If the decision-maker was not able to respond by decreasing emissions in the second period, then the total stock of emissions would be constant in the random shift parameter ε, and the marginal damages would be linear. If, however, the decision maker can reduce emission in response to the realization of ε then the marginal damages will be concave in ε as shown. The expected value of a concave function decreases in risk, so the expected marginal damages decrease in risk.

Now, extend the intuition to an increase in learning. It is only in response to learning that the stock of emissions are changed. In some sense, the more that is learned, the more concave the curve will be. Expected marginal damages are lower with more expected information, because the decision maker will reduce emissions as he learns that damages are increasing.

What if there is a constraint on 2nd period emissions? In that case, the decision maker reduces emissions as damages get worse, until the constraint is reached. Figure 3 shows marginal damages as a function of the random shift parameter when emissions are constrained to be non-negative. The curve is concave as long as emissions can be reduced, but when the constraint bites and emissions can no longer be reduced, the stock of emissions becomes constant and the marginal damages follow a linear path. Thus marginal damages are neither convex nor concave in the random shift parameter. This means that expected marginal damages will increase with some increases in risk (or informativeness) and decrease with other increases in informativeness.

If we expect to gain knowledge about catastrophic damages before it is too late, then we would like to have the flexibility to react to what we learn. This means that 1st period emissions would have to be lower to give us room to maneuver. On the other hand, if we only expect to get a vague signal that tells us which half of the distribution we will be in, then the constraint won’t bite in the 2nd period. In that case we would want to increase 1st period emissions, following the argument given above.

Note that if the upper bound on the support of the random shift parameter were low enough, then the marginal damages would be concave in the relevant region, implying that emissions increase in risk (and learning). This is important because the large numerical models of climate change such as those used in [8],[17], and [10], use a high damage/low probability event as a proxy for the right-hand tail of the distribution. What figure 3 shows is that if the “high damage” scenario is not high

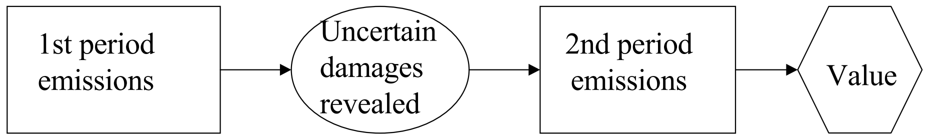

Figure 1: Sequential decision making under uncertainty

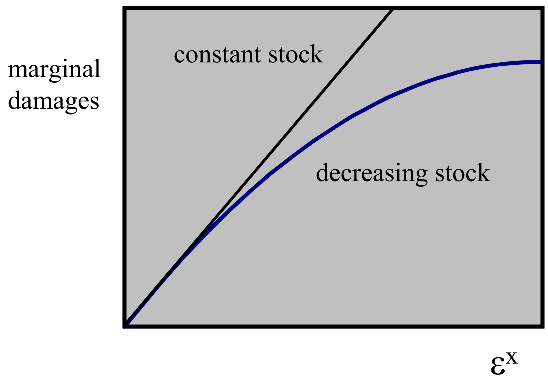

Figure 2: Marginal damages as a function of the random shift parameter ε.

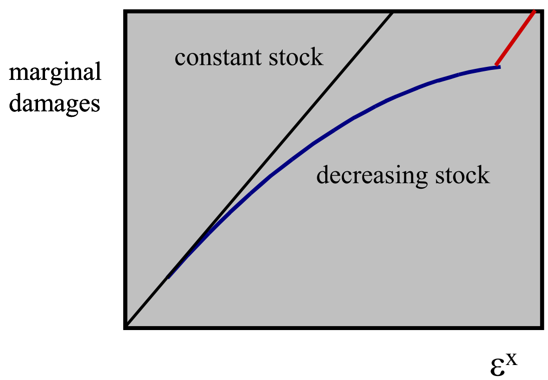

Figure 3: Emissions are constrained to be non-negative.

enough, then the results from the numerical models will be over-general. If we expect that we may learn that the magnitude of damages is very high, then we want to leave ourselves the flexibility to reduce emissions down to an optimal level in the 2nd period. To do this we may need to reduce emissions in the first period.

Figure 3 also illustrates the drawbacks to considering only the effects of perfect information. The fact that optimal emissions with perfect learning are higher than optimal emissions with no learning does not imply that emissions are monotone increasing in learning. Ulph and Ulph [25] present a sufficient condition for emissions in the perfect learning case to be higher than in the no learning case. But figure 3 clearly shows that the *only* sufficient condition for emissions to increase monotonically in learning is for ε_x to be bounded. It is important not to generalize results based on perfect learning to get insights about incremental increases in learning.

Emissions are not constrained to be non-negative for the remainder of the paper. Thus the results for perfect correlation (which echo the results in the single decision maker literature) are weaker, but the results on independence and perfect negative correlation are stronger.

4.2 Perfect Correlation

The results for the perfect correlation case follow the intuition for the single decision maker results. Since damages are perfectly correlated, both players will reduce emissions as damages get worse, thus marginal damages will be concave in the random variable, as in figure 2.

4.3 Independence

Emissions increase in own-risk (and own-learning) following the logic given above: player X will decrease his emissions as the magnitude of the damages gets higher. Player Y will increase emissions in response, but not enough to dampen the effect i.e. $\frac{\partial y_2}{\partial x_2} > -1$.

On the other hand, a change in one's opponent's shift parameter affects one's own damages only through the stock of emissions. When an opponent's realized shift parameter increases, he will reduce his second period emissions. But, he will reduce emissions less and less for each increase, since the marginal costs of reducing emissions are increasing, i.e. since $b(\cdot)$ is assumed to be concave. Thus the stock of emissions will be decreasing and convex in ε_y. As long as $D''' \leq 0$, this translates into convex marginal damages for X (See figure 4). Thus, expected marginal damages are increasing in risk. Again, the intuition extends to learning. The more player Y expects to learn, the more he will be able to adjust emissions in the second period, the more convex the marginal damages will be.

4.4 Perfect Negative Correlation

When damages are perfectly negatively correlated, then the stock of emissions is constant in the random shift parameter. Every increase by Player X in response to worse damages is countered by an increase by player Y in response to less severe damages. This implies that, in an open loop game, marginal damages are linear in the random variable (see the top line in figure 2), and thus information and uncertainty have no impact on the first period equilibrium. In the closed loop game, however, marginal damages in equation (5) are multiplied by a strategic term $\left[\gamma + \frac{\partial y_2}{\partial x_1}\right]$. This term accounts for the effect of X's first period emissions on Y's second period emissions, and is increasing – i.e. the impact of x_1 on Y's second period emissions is getting smaller as X's damages get worse. Thus X's damages are being compounded, therefore X reduces first period emissions in response to greater risk and/or informativeness.

5 Conclusion

Most policy related work on climate change has focussed on one of two issues: calculating the optimal policy in the event of a global agreement or attempting to determine ways to ensure a global agreement. This paper looks at optimal policy when a global environmental agreement is lacking, as is currently the case. The optimal policy in the face of uncertainty and learning for an individual decision maker is to "go slow" – reduce emissions modestly while pursuing better and faster information about climate change. Yet, this may not be the optimal policy for individual nations today. The results in this paper show that the optimal policy in a non-cooperative environment depends on the correlation of damages across countries. In fact, if damages are negatively correlated the optimal policy is reversed from the individual decision maker case: the more we expect to learn, the lower emissions should be.

Given that the climate change literature tends to focus on global variables such as global mean temperature, it is quite likely that policy makers are unconsciously assuming that damages are perfectly correlated across countries. To the degree that this assumption is wrong current emissions should be more aggressively reduced, particularly in the absence of an international agreement.

The results in this paper are relevant to questions about forming international agreements. Each of the players behaves as if risk averse to the other players uncertainty. This implies that each player would choose to buy insurance (against fluctuations in the opponent's damages) from some independent broker. This suggests a promising avenue for research on international environmental agreements.

Moreover, simply recognizing that the final distribution of damages is uncertain is very important in the design of a successful international agreement. Just as the papers looking at optimal policy have emphasized preserving flexibility, it is similarly important that any agreement be flexible enough to survive surprisingly uneven damages in the future. The Kyoto protocol, for example, is very flexible in one sense – it is a very short term agreement. This allows the emissions targets to be adjusted as new information about average damages is acquired. On the other hand, it is perhaps too flexible to be useful if the distribution of damages is revealed. In the absence of a pre-existing agreement low-damage nations will have very little incentive to continue participation.

There seems to be little understanding of the correlation (or lack of it) of climate change damages across countries. Most analyses give the impression that damages are perfectly correlated across countries. Yet regional analyses and common sense imply otherwise. For example, it is not clear whether the average temperature in Europe will increase or decrease with a moderate increase of global mean temperature. A better knowledge of how damages are correlated could aid greatly in making efficient decisions about emissions today as well as helping countries come to an agreement in the future.

A Proofs

A.1 Proof of Proposition 4

Proposition 4 $\frac{\partial y_2}{\partial x_1} \leq 0$

Proof. Totally differentiate the first order conditions for x_1 and x_2 to get

$$\begin{aligned} \left(b'' - \varepsilon_x D''\right) dx_2 - \varepsilon_x D'' dy_2 - d\varepsilon_x - \varepsilon_x D'' \gamma dx_1 - \varepsilon_x D'' \gamma dy_1 &= 0 \\ -\varepsilon_x D'' dx_2 + \left(b'' - \varepsilon_y D''\right) dy_2 - d\varepsilon_y - \varepsilon_y D'' \gamma dx_1 - \varepsilon_y D'' \gamma dy_1 &= 0 \end{aligned} \tag{6}$$

Applying Cramer's rule

$$\frac{dy_2}{dx_1} = \frac{\varepsilon_y D'' b'' \gamma}{\left[\left(b'' - \varepsilon_x D''\right)\left(b'' - \varepsilon_y D''\right) - \varepsilon_x \varepsilon_y D''\right]} \leq 0 \tag{7}$$

■

A.2 Proof of Proposition 1

Proposition 1:*Perfect Correlation: Let $\varepsilon_x = \varepsilon_y = \varepsilon$ where ε is a non-negative random variable. Then the symmetric first period equilibrium emissions (x_1^*, y_1^*) are increasing in uncertainty and informativeness.*

Proof. If

$$f(\varepsilon) = \varepsilon D'(s) \left[\gamma + \frac{\partial y_2}{\partial x_1}\right] \tag{8}$$

is concave, it implies that x_1 is increasing in risk, holding y_1 constant. Consider

$$\begin{aligned} f''(\varepsilon) &= 2D'' \left[\gamma + \frac{\partial y_2}{\partial x_1}\right] \frac{\partial s}{\partial \varepsilon} + 2D'(s) \frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon} + 2\varepsilon D'' \frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon} \frac{\partial s}{\partial \varepsilon} + \\ & \varepsilon D'' \left[\gamma + \frac{\partial y_2}{\partial x_1}\right] \frac{\partial^2 s}{\partial \varepsilon^2} + \varepsilon D'(s) \frac{\partial^3 y_2}{\partial x_1 \partial \varepsilon^2} \end{aligned} \tag{9}$$

The derivatives are calculated by totally differentiating the first order conditions for x_2 and y_2 (see the proof for proposition 4) and then applying Cramer's rule, as follows:

$$\begin{aligned} \frac{\partial s}{\partial \varepsilon} &= \frac{\partial x_2}{\partial \varepsilon} + \frac{\partial y_2}{\partial \varepsilon} = 2 \frac{D'(s) b''}{\left[\left(b'' - \varepsilon D''\right)^2 - \varepsilon^2 D''^2\right]} \\ \frac{\partial^2 s}{\partial \varepsilon^2} &= 4 \frac{D'(s) D'' b''^2}{\left[\left(b'' - \varepsilon D''\right)^2 - \varepsilon^2 D''^2\right]^2} \end{aligned} \tag{10}$$

$$\frac{\partial y_2}{\partial x_1} = \frac{\varepsilon D'' b'' \gamma}{\left[(b'' - \varepsilon D'')^2 - \varepsilon^2 D''^2\right]} \tag{11}$$

$$\frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon_x} = \frac{\gamma b''^3 D''}{\left[(b'' - \varepsilon D'')^2 - \varepsilon^2 D''^2\right]^2}$$

$$\frac{\partial^3 y_2}{\partial x_1 \partial \varepsilon_x^2} = \frac{2\gamma b''^4 D''^2}{\left[(b'' - \varepsilon D'')^2 - \varepsilon^2 D''^2\right]^3}$$

Plugging the partial derivatives into (9) and simplifying gives

$$f''(\varepsilon) = \frac{2D'' \gamma b''^3 D'}{\left[(b'' - \varepsilon D'')^2 - \varepsilon^2 D''^2\right]^3} \left[-4\varepsilon b'' D'' + 2\varepsilon^2 D''^2 + 3b''^2\right] < 0 \tag{12}$$

Since the players are symmetric the same argument holds for y_1 holding x_1 constant. The equilibrium can be shown to be stable (see author for details) and thus the comparative statics results hold for the equilibrium. ■

A.3 Proof of Proposition 2

Proposition 2: *Independence: Let ε_x and ε_y be independent non-negative random variables. Then first period equilibrium emissions are increasing in own-uncertainty (or in increased learning about ε_x) when $\varepsilon_y < -\frac{2}{3}\frac{b''}{D''}$ and decreasing in opponent's uncertainty (or in increased learning about ε_y) unconditionally.*

Proof. First look at ε_x. The derivatives are calculated as in the proof of Proposition 1 above.

$$\frac{\partial s}{\partial \varepsilon_x} = \frac{\partial x_2}{\partial \varepsilon_x} + \frac{\partial y_2}{\partial \varepsilon_x} = \frac{\varepsilon_y D' D'' + D'(b'' - \varepsilon_y D'')}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]} = \frac{D' b''}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]} \tag{13}$$

$$\frac{\partial^2 s}{\partial \varepsilon_x^2} = \frac{2D' D'' b''^2}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]^2}$$

$$\frac{\partial y_2}{\partial x_1} = \frac{\varepsilon_y D'' b'' \gamma}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]}$$

$$\frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon_x} = \frac{\gamma \varepsilon_y b''^2 D''^2}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]^2}$$

$$\frac{\partial^3 y_2}{\partial x_1 \partial \varepsilon_x^2} = \frac{\gamma \varepsilon_y b''^3 D''^3}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]^3}$$

Substituting the expressions for the derivatives into 9 shows that

$$f''(\varepsilon_x) = \frac{D'(s) D'' \gamma b''^3}{\left[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D''\right]^3} \left[2b''^2 - 2\varepsilon_x D'' b'' - 2\varepsilon_y^2 D''^2 - 3\varepsilon_x \varepsilon_y D''^2\right] < 0 \tag{14}$$

$$\varepsilon_y < -\frac{2}{3}\frac{b''}{D''}$$

Now look at ε_y. We want to show that (8) is convex in ε_y. Consider

$$f''(\varepsilon_y) = 2\varepsilon_x D'' \frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon_y} \frac{\partial s}{\partial \varepsilon_y} + \varepsilon_x D'' \left[\gamma + \frac{\partial y_2}{\partial x_1}\right] \frac{\partial^2 s}{\partial \varepsilon_y^2} + \varepsilon_x D'(s) \frac{\partial^3 y_2}{\partial x_1 \partial \varepsilon_y^2} \tag{15}$$

Note that $\left[\gamma + \frac{\partial y_2}{\partial x_1}\right] = \gamma b'' \frac{b'' - \varepsilon_x D''}{[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D'']} > 0$ and that $\frac{\partial s}{\partial \varepsilon_y} = \frac{\partial s}{\partial \varepsilon_x} < 0$ and $\frac{\partial^2 s}{\partial \varepsilon_y^2} = \frac{\partial^2 s}{\partial \varepsilon_x^2} > 0$. The remaining derivatives are as follows:

$$\begin{aligned} \frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon_y} &= \frac{\gamma b''^2 D'' (b'' - \varepsilon_x D'')}{[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D'']^2} < 0 \\ \frac{\partial^3 y_2}{\partial x_1 \partial \varepsilon_y^2} &= \frac{\gamma b''^3 D''^2 (b'' - \varepsilon_x D'')}{[(b'' - \varepsilon_x D'')(b'' - \varepsilon_y D'') - \varepsilon_x \varepsilon_y D'']^3} > 0 \end{aligned} \tag{16}$$

Therefore each term in (15) is positive and the right hand side is convex in ε_y.

The reaction functions are increasing in the risk of own-damages and decreasing in the risk of other's-damages. It is a straightforward application of Theorem 3 from Milgrom and Roberts [13] to show that first period equilibrium emissions increase for the player whose risk is increasing and decrease for the opponent (whose risk is held constant). ■

A.4 Proof of proposition 3

Proposition 3: *Perfect Negative Correlation: Let $(\varepsilon_x, \varepsilon_y) = (\varepsilon, 2 - \varepsilon)$ where ε is a symmetric, non-negative random variable with mean 1. Then symmetric first period equilibrium emissions (x_1^*, y_1^*) are decreasing in uncertainty and informativeness.*

Proof. The partial derivatives are as follows:

$$\frac{\partial s}{\partial \varepsilon} = \frac{\partial x_2}{\partial \varepsilon} + \frac{\partial y_2}{\partial \varepsilon} = \frac{D'(s)}{b''} - \frac{D'(s)}{b''} = 0 \tag{17}$$

Thus the stock of emissions is constant in ε.

$$\frac{\partial y_2}{\partial x_1} = \frac{(2 - \varepsilon) D'' \gamma}{b'' - 2D''} \leq 0 \tag{18}$$

$$\frac{\partial^2 y_2}{\partial x_1 \partial \varepsilon} = \frac{-D'' \gamma}{b'' - 2D''} > 0 \tag{19}$$

$$\frac{\partial^3 y_2}{\partial x_1 \partial \varepsilon^2} = 0 \tag{20}$$

Thus, 9 is positive, and x_1 is decreasing in uncertainty. The players are symmetric, so the rest of the proof follows the logic of Proposition 1. ■

References

[1] Erin Baker. *Increasing Informativenss and Increasing Risk: Equivalence Theorems.* PhD thesis, Stanford University, Palo Alto, 2002.

[2] David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.

[3] Larry G. Epstein. Decision making and the temporal resolution of uncertainty. *International Economic Review*, 21(2):269–283, June 1980.

[4] Christian Gollier, Bruno Jullien, and Nicolas Treich. Scientific progress and irreversibility: An economic interpretation of the 'precautionary pricipal'. *Journal of Public Economics*, 75:229–253, 2000.

[5] James K. Hammitt and John L. Adams. The value of international cooperation for abating global climate change. *Resource and Energy Economics*, 18:219–241, 1996.

[6] J.K. Hammitt, R.J.Lempert, and M.E. Schlesinger. A sequential-decision strategy for abating climate change. *Nature*, 357:315–318, 1992.

[7] N. Kilatilaka and E. Perotti. Strategic growth options. *Management Science*, 44:1021–1031, 1998.

[8] Charles D. Kolstad. Learning and stock effects in environmental regulation: The case of greenhouse gas emissions. *Journal of Environmental Economics and Management*, 31:1–18, 1996.

[9] Jean-Jacques Laffont. *The Economics of Uncertainty and Information.* MIT Press, Cambridge, Massachusetts, 1989.

[10] Alan S. Manne. Hedging strategies for global carbon dioxide abatement: A summary of poll results - EMF 14 subgroup analysis for decisions under uncertainty, 1996. Energy Modeling Forum, Stanford University.

[11] Alan S Manne, Robert Mendelsohn, and Richard G. Richels. MERGE: A model for evaluating regional and global effects of GHG reduction policies. Technical report, Electric Power Research Institute, 1993.

[12] Alan S. Manne and Richard G. Richels. *Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emission Limits.* MIT Press, Cambridge, MA, 1992.

[13] Paul Milgrom and John Roberts. Comparing equilibria. *The American Economic Review*, 84:441–459, 1994.

[14] Seong-Lin Na and Hyun Song Shin. International environmental agreements under uncertainty. *Oxford Economic Papers*, 50:173–185, 1998.

[15] William D. Nordhaus. *Managing the Global Commons.* MIT Press, Boston, 1994.

[16] William D. Nordhaus and David Popp. What is the value of scientific knowledge? an application to global warming using the PRICE model. *The Energy Journal*, 18:1–45, 1997.

[17] William D. Nordhaus and Z. Yang. A regional dynamic general equilibrium model of alternative climate-change strategies. *American Economic Review*, 86:714–165, 1996.

[18] Jennifer F. Reinganum. The timing of innovation: Research, development, and diffusion. In Richard Schmalensee and Robert D. Willig, editors, *Handbook of Industrial Organization*, pages 850–908. North-Holland, 1989.

[19] Michael Rothschild and Joseph Stiglitz. Increasing risk: I. a definition. *Journal of Economic Theory*, 2:225–243, 1970.

[20] Michael Rothschild and Joseph Stiglitz. Increasing risk II: Its economic consequences. *Journal of Economic Theory*, 3:66–84, 1971.

[21] Todd Sandler and Frederic P. Sterbenz. Harvest uncertainty and the tragedy of the commons. *Journal of Environmental Economics and Management*, 18:155–167, 1990.

[22] Michael Scott, Ronald Sands, Jae Edmonds, Albert Liebetrau, and David Engel. Uncertainty in integrated assessment models: Modeling with MiniCAM 1.0. *Energy Policy*, 27:855–879, 1999.

[23] L. Trigeorgis. *Real Options.* MIT Press, Cambridge, MA, 1996.

[24] Alistair Ulph and David Ulph. Who gains from learning about global warming? In E. Van Ierland and K. Gorka, editors, *The Economics of Atmospheirc Pollution*, pages 31–67. Springr-Verlag, Heidelberg, 1996.

[25] Alistair Ulph and David Ulph. Global warming, irreversibility and learning. *The Economic Journal*, 107:636–650, 1997.

[26] Kevin Zhu. *Strategic Investment in Information Technologies: A Real-Options and Game-Theoretic Approach.* PhD thesis, Stanford University, Palo Alto, 1999.

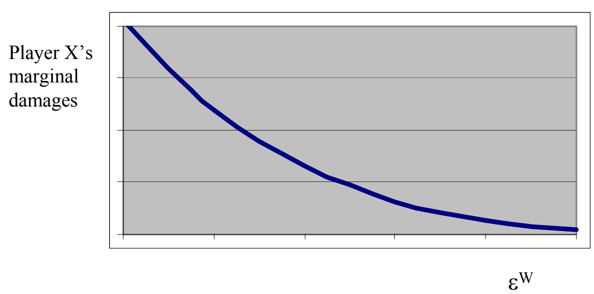

Figure 4: Marginal damages as a function of opponent's shift parameter