

VU Research Portal

Adaptive Play in a Pollution Bargaining Game

van der Goes, V.A.N.M.

published in

Proceedings of the 10th International Conference on Artificial Evolution (EA-2011)
2011

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

van der Goes, V. A. N. M. (2011). Adaptive Play in a Pollution Bargaining Game. In *Proceedings of the 10th International Conference on Artificial Evolution (EA-2011)* Springer Verlag.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Adaptive Play in a Pollution Bargaining Game

Vincent van der Goes

Vrije Universiteit Amsterdam, The Netherlands,
vgoes@feweb.vu.nl

Abstract. We apply adaptive play to a simplified pollution game with two players. We find that agents with longer memory paradoxically perform worse in the long run. We interpret this result as an indication that adaptive play may be too restrictive as a model of agent behaviour in this context, although it can serve as a starting point for further research on bounded rationality in pollution games.

Keywords: evolutionary economics, game theory, adaptive play

1 Introduction

Global environmental problems, such as pollution, provide real-world examples of a *social dilemma*. Countries involved in an environmental issue could often benefit from mutual cooperation. However, if they choose to do so, any country in the cooperation might benefit from *free riding*. That is, each country could possibly improve net welfare by defecting and profiting from the efforts at reducing pollution, without reducing its own pollution.

So it is in the mutual interest of countries involved in such an environmental problem to find a way to facilitate cooperation. To this end, countries have signed treaties known as *International Environmental Agreements* (IEA's).

In particular, the case of greenhouse gas emissions is receiving a lot of attention today. Greenhouse gas emission is a form of *transboundary pollution*, i.e. the impacts of emissions are global and do not depend on the location where they originate. As reduction of greenhouse gas emissions is very costly, it provides a good example of a social dilemma. Agreeing on an appropriate IEA has proven difficult in practice.

In the literature this has sparked a discussion about the question why negotiations on reduction of emissions are so difficult and what kind of policy instruments could help on reaching an agreement. For a recent overview of the literature in this area we refer the reader to [2] and [3].

A common way of modelling emission reductions is as a repeated game. The dynamics of a realistic model are highly complex, as there are several complicating factors at play. The game changes over time, as greenhouse gas concentrations build up and new technologies become available.

Also, there is a high degree of uncertainty, especially considering the long time scale at which the climate responds [6]. First of all *systematic uncertainty*: damage as a function of emissions is hard to predict. Secondly *strategic uncertainty*:

agents can not predict the behaviour of other agents very well.

In this paper we explore the possibility of modelling bargaining for reduction of emissions of a transboundary pollutant, while assuming only bounded rationality and imperfect information. We focus on the question whether agents will be able to coordinate on an efficient outcome under these conditions. In particular, we look at the effect of asymmetry in information.

The agents in our model will bargain on pollution abatement levels, using the *adaptive play* mechanism [9] to optimize their results. In adaptive play, agents are *myopic*: payoff in future rounds is not taken into account. Instead, the agents remember the earlier choices of other agents, assume they will make similar choices at present and search for a best response. The agents will also make occasional errors. This is to account for the fact that agents in the real world will occasionally act irrational, or there might be an unexpected external factor playing a role.

The bargaining game in our model is defined in two steps. The first step in the definition of our model is a toy model of transboundary pollution, the *pollution game*. The model is attractive for study and has been widely used for analysis with perfect information and rational agents, e.g. [1] [8]. We make a further simplification to the model, as we limit it to two agents.

Within the context of this pollution game, we assume that agents can make binding agreements on abatement for one round. At the beginning of each round, the agents will play a bargaining game over their abatement levels during that round. In the bargaining game both agents simultaneously make a proposal for the abatement levels. If the two proposals are mutually compatible, an agreement is made for that round. If the proposals are not compatible with each other, the agents have failed to make an agreement in that round and they will both choose their abatement levels according to the Nash equilibrium of the pollution game, instead.

We investigate the behavior of this model by running simulations. We find that agents are able to reach Pareto efficient outcomes, even though this is not true for all parameter settings. However, when one agent has a longer memory of past actions than the other, our behaviour model predicts that he will be worse off in the long run. This is a rather paradoxal finding. There is very little difference between our model and the bargaining model in [10], except that the underlying pollution game has a different payoff structure. Yet, in that study agents with longer memory were consistently better off.

This leads to two conclusions. First, it shows that the result in [10] does not generalize very well to other classes of bargaining games. Secondly, it is an indication that adaptive play may compromise the rationality of agents too strongly to be suitable as a model for this particular bargaining game.

The remainder of this paper is organised as follows. Section 2 describes the bargaining game. The pollution game is defined and analyzed in section 2.1 and in section 2.2 it is extended with a bargaining mechanism. Section 2.3 defines the behaviour of the agents. Section 3 contains the experiments. The simulations

are motivated and detailed in section 3.1. The experimental results are given in section 3.2 and section 4 concludes and suggests paths for further research.

2 Model

2.1 The Pollution Game

We model negotiations between two agents on emissions of a pollutant. The pollutant is transboundary and damage functions depend only on total emissions. Agents have the options to *abate*, lowering their emission for one period at a cost. The emphasis of the study is on learning dynamics. For this purpose we employ a strongly simplified pollution model, known as the standard model [1] [8]. We assume that costs of abatement are independent between agents and do not change over time. Further, it is assumed that the pollutant does not accumulate in the environment, even though greenhouse gases are in reality a stock pollutant. Also, all agents have the same damage function.

Agents are labelled $i = 1, 2$. They have an fixed unabated emission level $a/2$, where a is the aggregate emission level. They can spend on abatement, resulting in benefits for every agent. Hence, the benefits to agent i depend on the total abatement level Q . We assume linear marginal abatement benefits $B_i(Q)$. This leads to a quadratical benefit function, as specified in equation 1.

$$B_i(Q) = b(aQ - Q^2/2)/2 \quad (1)$$

The marginal benefit of the first unit of abatement is $ba/2$, with b a positive parameter. When total abatement equals total emission, the marginal benefit of further abatement is zero.

The costs of abatement $C_i(q_i)$ for agent i depends only on the abatement q_i of that agent, where $Q = q_1 + q_2$. The marginal costs of abatement are assumed to rise linearly, leading to the quadratic cost function in equation 2.

$$C_i(q_i) = cq_i^2/2 \quad (2)$$

where c is a positive parameter. The resulting net payoff vector or *utility* vector is shown in equation 3.

$$u_i(q_1, q_2) = B_i(q_1 + q_2) - C_i(q_i) \quad (3)$$

The Nash Equilibrium In the *Nash equilibrium*, both agents optimize their own payoff, without considering the possibility of mutual benefits from cooperation. This is known as *non-cooperative* behaviour.

In the pollution game of section 2.1, the Nash equilibrium can be calculated by taking the partial derivatives of the utility functions of both agents with respect to their own abatement level. Solving for the first order conditions yields a single symmetric Nash equilibrium, given in equation 4.

$$q_i = q_D = \frac{a}{2(1 + c/b)} \quad (4)$$

Here the Nash equilibrium is denoted as q_D , because it will serve as the *disagreement point* in the bargaining game in section 2.2. The Nash equilibrium yields equal utility $u_D = u_1(q_D, q_D)$ to both players.

The Pareto Frontier The Nash equilibrium of this game is not Pareto optimal. Both agents could increase their payoff by cooperation. Therefore, this game is an example of a *social dilemma*, similar to the *prisoner's dilemma*. If the agents cooperate, they will ideally choose Pareto optimal abatement levels. Additionally, both agents should receive greater payoff than in the Nash equilibrium. The symmetric solution on the Pareto frontier, the *Nash bargaining solution* (q^*, q^*) (see equation 5), satisfies both conditions. In [7], it is shown that under certain assumptions, the Nash bargaining solution will be the result of negotiations between two agents equal bargaining power.

$$q^* = \frac{a}{2 + c/b} \quad (5)$$

However, it should be noted that in our simulations the agents will not have perfect information, nor will they necessarily have identical bargaining power. The complete Pareto frontier can be characterized by maximizing the weighted average of u_1 and u_2 in equation 6, with weight α .

$$\alpha u_1(q_1, q_2) + (1 - \alpha)u_2(q_1, q_2) \quad (6)$$

The function in equation 6 has a single maximum at $(q_1^*(\alpha), q_2^*(\alpha))$ in 7.

$$\begin{aligned} q_1^*(\alpha) &= \frac{(1 - \alpha)a}{1 + 2\alpha(1 - \alpha)c/b} \\ q_2^*(\alpha) &= \frac{\alpha a}{1 + 2\alpha(1 - \alpha)c/b} \end{aligned} \quad (7)$$

In the special case that $\alpha = 1/2$, the solution in equation 7 reduces to the Nash bargaining solution (q^*, q^*) . The extreme points of the Pareto front are the cases $\alpha = 0$ and $\alpha = 1$. In these cases, one agent does not abate and the other agent chooses an abatement level of a .

2.2 The Bargaining Game

The pollution game described in section 2.1 is assumed to be repeated for an unlimited number of rounds. Under such conditions, cooperation is possible, according to the folk theorem [5]. However, the incentive for agents to cooperate,

rather than to play the Nash equilibrium, is that cooperation in this round may be rewarded by the other agent in later rounds.

In this work we will look at agents with no forward looking abilities. With such limited rationality, agents will not be able to cooperate unless some bargaining mechanism is added to the game.

Therefore, we add a simple bargaining system. Let agent $-i$ denote the opposite agent from agent i . Instead of choosing abatement levels directly, both agents simultaneously make a proposal $p_i^t = (\hat{q}_{i,1}^t, \hat{q}_{i,2}^t)$ for the abatement levels in round t . Here $\hat{q}_{i,i}^t$ is the level of abatement that agent i is offering himself. However, in exchange, agent i demands an abatement level of at least $\hat{q}_{i,-i}$ from agent $-i$. Only if the offers and demands of both agents are compatible, the negotiations result in cooperation for round t . Otherwise, the agents fail to reach an agreement in round t and they will play the Nash equilibrium of the pollution game instead. Denote the boolean value of this requirement as $\Delta \in [\mathbf{true}, \mathbf{false}]$:

$$\Delta(p_i^t, p_{-i}^t) = \forall i : \hat{q}_{i,i}^t \geq \hat{q}_{-i,i}^t \quad (8)$$

If condition 8 is met, both agents are willing to abate at least as much as the other agent demands. But it is not yet clear what the final agreement will be. Will q_i^t be set to the level $\hat{q}_{i,i}^t$ that i offered himself, or to the level $\hat{q}_{-i,i}^t$ that the other agent demanded? In general, we could imagine the final result of negotiations to be any value in between the two. We assume that the agents will share the difference in a manner agreed upon prior to negotiations. The final abatement level of agent i will be a weighted average of the abatement level that agent i offered and the abatement level that was required, as in equation 9. The effect of varying λ will be analysed in sections 3 and 4.

$$q_i^t = q_i(p_i^t, p_{-i}^t) = \begin{cases} \lambda \hat{q}_{i,i}^t + (1 - \lambda) \hat{q}_{-i,i}^t & \text{if } \Delta(p_i^t, p_{-i}^t) \\ q_D & \text{otherwise} \end{cases} \quad (9)$$

Note that if requirement 8 fails, the agents abort negotiations and instead revert to the disagreement point, which is the Nash equilibrium of the pollution game.

The final utility level of agent i in round t , $u_i^t(q_i^t, q_{-i}^t)$, then follows by inserting these values into equation 3:

$$u_i^t = u_i(q_i^t, q_{-i}^t) = B_i(q_i^t, q_{-i}^t) - C_i(q_i^t) \quad (10)$$

Nash Equilibria of the Bargaining Game The negotiation process thus defined constitutes a new game, where the actions agents have to choose from are the proposals they make. In order to distinguish this game from the pollution game itself, we will refer to it as the *bargaining game*. Where the pollution game has only a single Nash equilibrium, the bargaining game has many.

Consider the solution where, for $i = 1, 2$, $\hat{q}_{i,i}^t = 0$ and $\hat{q}_{i,-i}^t = a$. This is a degenerate Nash equilibrium. Since requirement 8 is not met, the result is that both

players play the Nash equilibrium $q_i^t = q_D$ in the pollution game. Neither agent can improve his own utility by deviating unilaterally.

The game also has many non-degenerate equilibria, where both players end up with higher utility than in the Nash equilibrium of the pollution game. For example, consider $\hat{q}_{i,i}^t = \hat{q}_{i,-i}^t = q^*$. This leads to the Nash Bargaining Solution of the pollution game. Both agents have improved their utility over u_D and neither agent could improve his utility further by a unilateral deviation from his negotiation strategy.

Any non-degenerate equilibrium of the bargaining game must be Pareto dominant over the Nash equilibrium of the pollution game. If it does not, then at least one agent is worse off than u_D and therefore has an incentive to let the negotiations fail instead, by lowering his abatement offer.

If $\lambda < 1$, any non-degenerate Nash equilibrium, with at least one player ending up with a higher utility than u_D , must have $\hat{q}_{i,i}^t = \hat{q}_{-i,i}^t$. If $\hat{q}_{i,i}^t < \hat{q}_{-i,i}^t$ then the negotiations fail. But if $\hat{q}_{i,i}^t > \hat{q}_{-i,i}^t$, then agent $-i$ has an incentive to increase $\hat{q}_{-i,i}^t$, since by equation 9 this would increase the abatement of i , but leave his own abatement unchanged. As an increase in the abatement of i increases the benefits for $-i$ but not his costs, the utility of agent $-i$ would increase.

Provided that $0 < \lambda < 1$, any solution of the bargaining game that satisfies both of these conditions is a Nash equilibrium of the bargaining game. Neither player has an incentive to change his own proposed abatement $\hat{q}_{i,i}^t$, because it would either let the negotiations fail or it would cause him to increase his abatement level further above q_D , lowering his utility. Similarly, neither player has an incentive to change the abatement level he demands from the opponent, $\hat{q}_{i,-i}^t$.

Now that we identified the non-degenerate Nash equilibria (at least for $0 < \lambda < 1$), we can calculate the corresponding range of values for q_i^t .

The extreme points of the region that Pareto dominates q_D are on the Pareto frontier of the pollution game, as defined by $q_i^*(\alpha)$ in equation 7. We have the constraint that $q_i^t \leq q_D$. Solving for α in equation 7, we find that the region of non-degenerate Nash equilibria is bounded by the values q_{min} and q_{max} in equation 11.

We will assume that agents only consider values for their proposals in the range $[q_{min}, q_{max}]$. It could be argued that the agents have insufficient information to know q_{min} and q_{max} ahead of time. However, if our agent model is applied to the Pollution Game itself, rather than to the Bargaining Game, it will converge to the unique Nash equilibrium (q_D, q_D) , so it is not altogether unreasonable to postulate that agents are aware of $q_{min} = q_D$, at least.

$$\begin{aligned}
 \alpha_{min} &= \frac{\sqrt{b^2 + 2bc} - b}{2c} \\
 \alpha_{max} &= 1 - \alpha_{min} \\
 q_{min} &= q_D \\
 q_{max} &= q_D * \frac{\alpha_{max}}{\alpha_{min}}
 \end{aligned}
 \tag{11}$$

2.3 Adaptive Play

In this section we describe how agents decide what proposals to make in the negotiations of section 2.2. Agents are assumed to have bounded rationality. They do not have any forward looking abilities and do not know each others payoff functions. Instead, they form a model of the behaviour of the other agents, in order to find the best response to their actions.

We apply the *adaptive play* mechanism [9]. In adaptive play, agents assume that the other agents draw their proposals from a fixed probability distribution. In other words, they assume that the other agent plays a mixed strategy that doesn't change over time.

Agents have a limited memory m_i . Only actions chosen by the other agent in the latest m_i rounds are remembered. The agents then assume that the other agent will play any of the actions that he played during the last m_i rounds, each with equal probability.

Based on this internal model of the opponent, the agents search for an optimal response. They will maximize the expected final utility of their proposal p_i^t . Let (x, y) be a candidate proposal for p_i^t . It is evaluated by the expected value of the utility that it could yield:

$$f_i^t(x, y) = \frac{1}{m_i} \sum_{t'=t-1}^{t-m_i} u_i(q_i((x, y), p_{-i}^{t'}), q_{-i}((x, y), p_{-i}^{t'})) \quad (12)$$

Where q_i is calculated as in equation 9 and u_i as in equation 3. The function $f_i^t(x, y)$ serves as a fitness function.

The agent will then make the proposal that maximizes this fitness function:

$$p_i^t = \underset{(x, y) \in [q_{min}, q_{max}]^2}{\operatorname{argmax}} f_i^t(x, y) \quad (13)$$

In case there are multiple values of (x, y) that maximize $f(x, y)$, one value is drawn from the set with a uniform random distribution.

However, in adaptive play, agents are allowed to make occasional errors. These are introduced to take into account that agents in the real world occasionally behave unpredictably, because of factors external to the game. With a small probability ε , equation 13 is ignored and p_i^t is instead drawn at random from $[q_{min}, q_{max}]^2$ with uniform distribution. This is known as a mutation.

Of course, equation 12 is undefined during the first few rounds of play. For $t = 1, \dots, \max(m_1, m_2)$, p_i^t is again drawn at random from $[q_{min}, q_{max}]^2$ with uniform distribution.

2.4 Stochastic Stability

As shown in section 2.2, the bargaining game has many Nash equilibria. It is easily verified that all of the non-degenerate Nash equilibria of the game are also evolutionary stable strategies. Hence, the question is how to define or identify

the most likely outcome of the game.

We are interested in finding the *stochastically stable strategies* of the game. Stochastically stable strategies are a refinement of the concept of evolutionary stable strategies. Any stochastically stable strategy is also an evolutionary stable strategy, but the reverse is not true.

In an evolutionary system such as our multi-agent model, an evolutionary stable strategy is stable in the short run, as neither agent has an incentive to deviate unilaterally. However, over increasingly long time scales there is an increasing probability that mutations disturb the equilibrium and cause the system to enter a different equilibrium. On sufficiently long timescales, the system behaves ergodic, switching back and forth between different evolutionary stable strategies. The timescale at which this kind of behaviour is typical is called the *long run*, or sometimes the *ultra-long run* [10].

Stochastically stable strategies are the evolutionary stable strategies that occur the most frequently in the long run, when the mutation rate ε approaches zero. For a full definition we refer the reader to [4].

3 Simulations

We will use simulations to gain insight into the stochastically stable strategies under different parameter settings of the model. The model will be run for a large but finite number of rounds T_{MAX} , under different parameter settings.

We will then analyze the results, using autocorrelation in the series under different time lags to find empirically what timescale constitutes the long run, i.e., at what timescale the system behaves ergodic.

3.1 Experimental Setup

The primary question that we wish to address is the role of information: what happens when one agent has an information advantage over the other? What if one agent has a longer memory? We have set up a series of five experiments (E_3 and $E_6 \dots E_9$) to compare the behaviour of the system under varying differences in memory.

The second question we wish to address is: how sensitive are the outcomes to a change in λ ? For this purpose, we have set up a second series of experiments, $E_1 \dots E_5$, where λ was varied. Experiment E_3 serves as the benchmark, as it has equal memory for both players and a value of λ of 0.5, splitting the remainder in equation 9 equally.

Table 1 provides an overview of all the parameter settings that have been used in the experiments. The ratio of b and c has been chosen such that it maximizes the potential gain from cooperation [1]. The mutation rate ε has been chosen as small as feasible to approximate the stochastically stable equilibria.

Table 1. Overview of Parameter Values

parameter	experiment								
	E_1	E_2	E_3	E_4	E_5	E_6	E_7	E_8	E_9
a	1								
b	1.64								
c	1								
ε	0.01								
T_{MAX}	10^7								
λ	0	0.25	0.5	0.75	1	0.5	0.5	0.5	0.5
m_1	8								
m_2	8	8	8	8	8	6	4	2	1

3.2 Results

In this section we will analyze the output, series of resulting utility values u_i^t . They are not time independent, but strongly autocorrelated. That is also to be expected, since a stochastic equilibrium is defined as the equilibrium that the is dominant in the long run. In order to find out what the long run is in this game, we need to find out at what time lag the autocorrelation between subsequent proposals approaches zero.

Autocorrelation can be expressed as a function of timelag, using formula 14:

$$R_i(\tau) = \frac{E[(u_i^t - \bar{u}_i)(u_i^{t+\tau} - \bar{u}_i)]}{\sigma^2} \quad (14)$$

In figure 1 all the autocorrelation plots of the experiments are shown. As expected, autocorrelation vanishes at sufficiently long timescale. From the figure we estimate a time lag of $\tau = 50.000$ as the typical long run.

Now we can sample measurements from each of the series at intervals of 50.000 rounds and treat them as approximately independent. Figure 2 shows a histogram of the resulting sample for benchmark experiment E_3 .

As can be seen from figure 2, the measurements follow a roughly bell shaped curve, with a group of outliers at the disagreement utility U_D . The outliers result from rounds where the negotiations failed. Since the interval $[q_{min}, q_{max}]$ does not contain degenerate equilibria, failed negotiations are not part of a stochastically stable equilibrium. Therefore, we will ignore those measurements in the remainder of the analysis.

We assume by the bell shaped form of the rest of the distribution that there is a single stochastically stable equilibrium and estimate it by taking the mean value. The mean values for each experiment, after dismissing outliers, can be found in table 2.

In the benchmark experiment E_3 , with equal memory and $\lambda = 0.5$, agents have on average been following the Nash Bargaining Solution ($u^* = 0.3141$) nearly exactly. Indeed, a t-test shows no statistically significant difference from




Fig. 1. Autocorrelation as a function of timelag, for both agents in each of the experiments.




Fig. 2. Histogram of utility values sampled from E_3 .

a normal distribution with mean u^* .

When λ differs from 0.5, as in experiments E_1 , E_2 , E_4 and E_5 , it clearly negatively affects the results of negotiations. For each of those experiments, average utility remained significantly below u^* , all at p -values below 10^{-5} .

Table 2. Average observed utility values (outliers removed)

	E_1	E_2	E_3	E_4	E_5	E_6	E_7	E_8	E_9
λ	0	0.25	0.5	0.75	1	0.5	0.5	0.5	0.5
m_2	8	8	8	8	8	6	4	2	1
u_1	0.3098	0.3129	0.3140	0.3131	0.3101	0.3125	0.3125	0.3131	0.3129
u_2	0.3099	0.3128	0.3142	0.3133	0.3102	0.3149	0.3148	0.3141	0.3141

Agents with shorter memory length have a clear advantage in the long run. Experiments E_6 to E_9 all seem to favour the less informed agent 2. However, the results from those experiments show a wide spread, with long tails. This might explain why a t -test shows no statistical significant difference between both agents in E_8 and E_9 . In E_6 and E_7 , however, the difference is significant at the 1% confidence level.

The average utilities observed in the experiments with asymmetric information are all below the Pareto frontier. For example, in E_6 agent 1 had an average utility of 0.3125. It follows from 7 that a pareto optimal solution with this u_1 should have $u_2 = 0.3157$. The actual average u_2 was lower, even though a t -test shows no statistical significance. The same holds for E_7 to E_9 .

4 Conclusions and Outlook

Even though our model is strongly simplified, especially as it is a two agent model, it nonetheless offers important insights into myopic bargaining on pollution games.

As often in adaptive play, agents are able to coordinate on a Pareto efficient outcome, without knowing each others payoff functions. However, this is not guaranteed and depends on parameter values of the model, moreso than in many other applications of adaptive play.

Surprisingly, it turns out that having a shorter memory is an advantage in the long run. This must be due to the structure of the pollution game, as adaptive play has been applied in this fashion to a similar bargaining game [10], where adaptive play favours the agents remembering more of the past under any parameter settings.

This result indicates that adaptive play may be too limiting for studying pollution games. While we believe that evolutionary economics is generally a promis-

ing venue for studying bounded rationality, the challenge for future research is to understand the behaviour of agents that are imperfect, yet less short sighted than in adaptive play.

5 Acknowledgements

The author would like to thank Cees Withagen for reviewing the manuscript. Also, thanks to Guszti Eiben and Daan van Soest for their helpful suggestions.

References

1. Barrett, S.: Self-Enforcing International Environmental Agreements. *Oxford Economic Papers* 46, 804–878 (1994)
2. Barrett, S.: *Environment & Statecraft: the Strategy of Environmental Treaty-Making*. Oxford University Press (2005)
3. Finus, M.: Game Theoretic Research on the Design of International Environmental Agreements: Insights, Critical Remarks, and Future Challenges. *International Review of Environmental and Resource Economics* 2, 2967 (2008)
4. Foster, D., Young, P.: Stochastic evolutionary game dynamics. *Theoretical Population Biology* 38:2, 219-232 (1990)
5. Fudenberg, D., Maskin, E.: The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54, 533-554 (1986)
6. Kolstad, C.D.: Systematic Uncertainty in Self-enforcing International Environmental Agreements. *Journal of Environmental Economics and Management* 53, 68-79 (2007)
7. Nash, J.: The Bargaining Problem. *Econometrica* 18, 155-162 (1950)
8. Rubio, S., Ulph, A.: Leadership and Self-Enforcing International Environmental Agreements with Non-Negative Emissions. unpublished manuscript, university of Valencia (2003)
9. Young, P.: The evolution of conventions. *Econometrica* 61, 57-84 (1993)
10. Young, P.: *Individual strategy and social structure*. Princeton University Press, 113-129 (1998)