Bargaining Dynamics in an Uncertain World

(MSc Thesis - Extended Summary)

Mateus Valadas Preto Bravo Cordeiro *
*Departamento de Engenharia Informatica and INESC-ID
IST-Taguspark, Universidade Tecnica de Lisboa,
2744-016 Porto Salvo, Portugal
mateus.cordeiro@tecnico.ulisboa.pt

***Abstract*—Recent conventions aiming to reduce greenhouse gas emissions failed to secure robust agreements in this regard. Changes by a single country do not have enough significance to mitigate this problem. Saving the planet requires cooperation. Understanding how cooperation occurs may be one of the keys to solve this problem. Game Theory allows the study of this type of conflicts between individual and collective interests. This work focuses on the study of cooperation in bargaining situations, developing models with new mechanisms and studying their effect in cooperation. The models are centered on the dynamics of the negotiations in the presence of uncertainty in collective goals and the countries' decision processes. Studying this mechanisms is possible using multi-agent systems simulations. We conclude that several factors and their combination can greatly affect the outcome and feasibility of negotiations between countries. We show that using a player's full set of moves to predict the next one produces better results than using a subset. Irrationality in the player's choices influences negotiations in a negative way, although the players can adapt to erratic moves. We verified that players have trouble adapting to variations in the collective goal, be those changes random or depending on the player's performance. Analysis of different risk taking strategies also proves that risk averse populations thrive better than risk prone ones. Risk prone societies, however, obtain better results when the collective goal is uncertain, surpassing even risk averse societies, so being risk prone may be beneficial in the presence of uncertainty in the collective goals.**

I. Introduction

People negotiate to protect what is theirs, or to satisfy their interests. Negotiations can take multiple forms and serve multiple purposes, and all of them seek individual or collective gain. One of the major challenges in science is to identify the mechanisms underlying cooperation in climate change agreements. Weather agreements are feasible or not, is dependent on the participant's will to cooperate, that is, to endow others some benefit to help satisfy theirs needs. Understanding how successful agreements are obtained benefits from the vast literature in cooperation and altruism.

Cooperation occurs when a party benefits another, being subjected himself to a cost. A self-interested individual weights his winnings and losses to make an informed decision. For this individual, cooperation is not an option, because there is not benefit in loss. A rational decision is waiting to get benefits from other's cooperation, while benefiting no one. This, however, is not supported by the empirical evidences, as cooperation prevails in both the nature and in our society. Animals will group themselves in herds, allowing others to eat from their source. People take time and energy from their lives to recycle. How does this cooperative behavior prevail, in a world in which the rational choice for our actions seems to be selfishness? The answers to these questions can be unveiled by framing cooperation in simple mathematical models, as often carried out in Game Theory.

Game Theory studies mathematical models of conflict of interests between individuals. These models are analogies to situations in which agents interact. The agents encapsulate a behavior, and the game is played. In Game Theory, cooperation problems often involve a tension between the individual interest and social desirability, and are described as Dilemmas. These dilemmas can be simple situations where two individuals compete to their best interests, or they can be more complex, and emulate situations comprised of multiple individuals. The last are called N-Person Games. Moreover, often these type of dilemmas occurs in large populations, raising complex ecologies of choices. This is often the case within problems with ecology and theoretical biology. The population counterpart of game theory is called Evolutionary Game Theory [1].

Evolutionary Game Theory attempts to map the mathematics of Game Theory to biological systems, introducing the complexity inherent to population dynamics to classic Game Theory. It associates an agent's success to its fitness, meaning fit agents are prone to thrive, be imitated and reproduce their ideas, much like in natural selection. A player's strategy becomes now less relevant, as our main focus is the model's evolution and player's adaptation to the environment.

International summits regarding climate changes are a major phenomenon studied by Game Theory [2]–[6], as they can be encapsulated in N-Person Games, in which players can represent countries or coalitions, and targets can be set for their individual or collective reductions.

This work focused on bargaining situations which are likely to occur in large-scale international climate agreements, using an existing model as main framework [7]. In this model, players can propose possible reductions on their greenhouse gas emissions and, as a whole, they must reduce enough emissions to stop or reduce their effect in climate changes.

On each round, the players propose their reductions based on their expectations of the other player's proposals. If the players meet the set target, the negotiations are successful. Otherwise, they fail. A player's payoff depends of the success of this negotiation. The model will be further explained in Section III.

Section II will present an overview and analysis on the state of the art studies relevant to this work. Section III, as stated before, describes the model under study, and Section IV presents the changes introduced to this model and an analysis of their impact on negotiations. At last, Section V concludes this work, providing an overview on the work completed and some additional notes and prospects for future work on this topic.

II. Related Work

Cooperation is one of the main areas of study by Game Theory. Cooperation and climate agreements can be framed in the broader field of cooperation studies, which, in the last decades, have successfully identified a large range of mechanisms capable of promoting cooperation in large populations, from cells to humans. This section presents some relevant mechanisms that promote cooperation.

A. Mechanisms for Cooperation

Most of the following mechanisms for cooperation are studied in the context of the Prisoner's Dilemma. Here, players can incur in a cost c to provide a benefit b. Table I is the payoff matrix for this dilemma. This Prisoner's Dilemma explains why cooperation may fail in the presence of two individuals. The best outcome for a player is to Defect, and wait for the other player to Cooperate. This way, the cost of cooperating is avoided, while having a chance of receiving from the other player's cooperation.

Table I: Prisoner's Dilemma payoff

	Cooperate	Defect
Cooperate	benefit - cost	- cost
Defect	benefit	0

Nowak summarizes, in a recent review [8], the main mechanisms that enable cooperation between individuals. The first rule is named Kin Selection. Kin Selection is defined as cooperation between related individuals that share genes. A specie's goal is thriving and surviving, and an individual's main purpose is passing its genes to the next generations, ensuring the survivability of its species. While simple and easy to understand, this rule does not explain why cooperation occurs between unrelated individuals. To make up for this, Direct Reciprocity is introduced.

Direct Reciprocity assumes that individuals will have multiple encounters and remember each other. In those encounters, an agent will choose to cooperate instead of defecting, because the other may return the favor. This mechanism proved to be great as a strategy for a Prisoner's Dilemma tournament, proved by Axelrod with his strategy Tit-For-Tat [9], winning the tournament twice. This Tit-For-Tat strategy is as simple as simple can be: start by cooperating, and on all subsequent rounds, imitate your opponent's move. While this strategy may seem great, it has a major failure when dealing with unexpected moves [10], [11] that can lead to an unending set of defections. This problem was faced creating the Win-Stay, Lose-Shift strategy [12]. This strategy keeps the simple behavior of the prior one, stating that when you have a good round, you maintain your move, and change it otherwise. Comparison between both strategies proved that the latter is superior to maintain cooperation [13]. Direct Reciprocity explains why cooperation occurs between individuals that are not related that face repeated encounters, but does not explain why it exists in situations where repeated encounters are not assumed, like donating. Thus, the Indirect Reciprocity rule was created.

Indirect Reciprocity states that, in those situations, one that cooperates does it expecting, from this behavior, good reputation and the prospect of being rewarded for it. Unlike Direct Reciprocity, the choice between cooperation and defection is not straightforward, as individuals must monitor every other individual in the network and think ahead when making a decision. Due to this complexity, Indirect Reciprocity is nearly exclusive to humans, mostly because the rest of the living beings have a limited rationale and do not gossip.

Cooperation can lead to success not only individually, and Group Selection explains why this happens. Nowak and Traulsen [14] developed a simple model, in which a population is divided into groups, and an individual's chance of reproducing is proportional to its fitness. Whenever a group reaches a certain size threshold, it splits in two, eliminating another existing group, for population control. That study showed that pure cooperative groups grow faster than pure defector groups, while in mixed groups, defectors reproduce faster.

B. Mechanisms for Cooperation in Bargaining Situations

The previous mechanisms concern cooperation dilemmas in a rather general setting, yet mainly concerning the famous prisoner's dilemma game. Differently, here we will address a cooperation dilemma which is often framed in the context of bargaining situations.

Punishment is the most studied mechanism in the context of bargaining situations. It is a kind of sanctioning applied to a defector. Punishment in the real world is in the form of fines, refusal to trade, and others. In Game Theory, it is usually carried out as a fine or some similar mechanism. Fehr and Gächter concluded [15] that even if there are no palpable benefits for a punisher and that punishment is expensive, it will usually be carried out, due to negative feeling towards the defector. That study also concluded that the presence of punishment can reduce the degree of defection of some players. Punishment is also closely related to Reputation. It can even be said that Reputation drives Punishment. An individual's reputation can be modeled and studied as a number or image that increases whenever help is provided and decreases otherwise. The effect of Reputation was also studied by Pacheco, Santos and Chalub [16], using the concept of stern-judging, a policy that awards cooperation to individuals with good reputation and discourages it to ones with bad reputation.

One of the major mechanisms that affects cooperation is the Awareness of Risk. Santos and Pacheco [17] showed that perception of a collective risk enhances the probability of successful negotiations. They studied a model in which individuals have wealth and can either donate some or keep it and if, as a whole, a population of N individuals fails to reach M cooperators, there is a probability *r* (risk) that everyone loses all of their possessions. It comes to no surprise that higher values of *r* cause negotiations to be more successful. The concept of risk leads us to the conclusion that success in international summits for climate changes is highly dependent on the risk perception by all the parties.

Globalization lead to a blend of races, cultures, ideologies, wealth and size in countries. This disparity is often responsible for difficulty in agreement between parties. Studies in this area [18]–[20] showed that under the right circumstances of Risk Perception mechanism and positive homophily, i.e, influence by successful peers, cooperation can overcome defection. In bargaining situations, the existence of social diversity changes the negotiation's dynamic, because wealthier/larger parties have to make larger concessions to make up for the smaller parties.

This section shows that the work done in cooperation and the mechanisms that enabled it is quite extensive, even in bargaining situations. Another trend is the common reference of climate changing issues and how societies or groups can understand its implications and how to succeed in negotiations. For that reason, my work will be focused on expanding existing models, since a large set of the aforementioned mechanisms are already implemented. The next section details how this is achieved in practice.

III. METHODS

This section provides an overview on the model and some of its details and intricacies. It is an extension to Nash's bargaining game [21] with N players and, for that reason, is an N-Person Bargaining Game. In bargaining games, agents make offers that can either be accepted or rejected. When accepted, a successful negotiation is achieved and the players are rewarded. A player's goal is maximizing its payoff. However, because it is dealing with a multi-agent environment, the implementation of this goal is not obvious. A possible way uses the information about opponents' past actions to derive a best response. I expand upon this model and alter its behavior and mechanisms to study how cooperation is affected.

A. A Bargaining Game of International Climate Negotiations

Departing from the seminal work by Smead et al. [7], each player has a strategy set that ranges from [0, 1]. 1 represents the Business As Usual (BAU), meaning it will emit 100% of its current emissions, and 0 means no emissions at all. d_i represents player d's demand for a proportion of its BAU or, equivalently, its proposed reduction.

As a group, the N players must meet a global emission reduction target T $\in$ [0, 1]. Higher values of T imply a smaller decrease in emissions, making negotiations easier. Lower values mean the opposite. Success in negotiations is achieved when, as a whole, the player's total demands are equal or inferior to the target reduction T. For a set of demands, we can define $t = \sum_{i=1}^{N} d_i$ as the player's total demands. The payoff function for each player is then defined as:

$$\pi_i(d_i, t) = \begin{cases} d_i & \text{if t} \leq \text{T.N} \\ \delta d_i & \text{if t} > \text{T.N} \end{cases}$$

If the target value T is not met, the player's payoff is devalued by a factor $\delta \in [0, 1]$. This value represents both the importance of reaching an agreement and the payoff devalue for a player's demand. Lower δ values devalue payoffs more than higher values. This function means that players will favor demands that are closer to their BAU, unless that demand threatens to collapse an agreement.

Adaptive agents are represented through a player's learning process. The historic of a player's demands is visible to every other player. After each round of negotiations, the agents can change their demands, based on the new information they have. This learning process is modeled using expectations on the other players' next demands. Expectations on each player are made by averaging their previous demands, although other alternatives could be devised. After forming expectations on all the other players, an agent presents its demand with the goal of maximizing its payoff. The first round demands are randomly generated for each player from the interval [δ, 1].

The game will continue for 100 rounds, and the final demands are analyzed at this point, deeming the game either successful or not.

1) Size in players: Players can have varying sizes, as do countries. The model incorporates this differences attributing a size to each player. A player's size is an integer $s_i \in$ [1, ∞]. The payoff is now updated to the following:

$$t = \sum_{i=1}^{N} d_i s_i$$

$$\pi_i(d_i, t) = \begin{cases} d_i & \text{if t} \leq \text{T} \sum_{i=1}^{N} s_i \\ \delta d_i & \text{if t} > \text{T} \sum_{i=1}^{N} s_i \end{cases}$$

2) Restrictions: Initial demands are assigned on the first round from the interval [δ, 1]. If the negotiations' initial demands are too high, the negotiations usually break down, because the further away they are from the objective, the more collective action is needed to achieve that goal. Thus, mechanisms that instigate initial demands closer to the target value will increase the negotiation's success rates. The model allows for these restrictions on the initial demand to be upper bound, lower bound, or both. Each kind of restriction affects the range of values from which the first demand is chosen from. For upper bounds, the interval is [δ, 1 - *r*]. For lower bound, its [δ + *r*, 1]. When restricting both, its [δ + *r*, 1 - *r*]. These restrictions are only applied to the first round.

3) Side Deals: In real negotiations, countries can form coalitions or have some prior agreements on what values they can or can not propose. The model allows for this, and the player's demands can be restricted with a maximum demand, and it lasts for the whole game. This mechanism reduces the player's initial demands and average demands during the game. When active, a player's strategy range is [$\delta\ Max_i$, Max_i]. It can also be applied to a subset of players, instead of the whole field.

IV. Results

This section presents the changes in the model's mechanisms and behavior. The results from running computer simulations on the model are also presented here, as well as their analysis.

A. Simulation Parameters

The model is implemented in Java. The model's parameters are defined in an input file, and the results written to another output file.

Each datapoint is simulated 10^4 times and, unless stated otherwise, the plot's parameters are the following: 8 players (size 1 for all), agreement importance δ = 0.1, target value T = 0.5, and any other mechanism is disabled.

B. The role of group size, agreements' importance and heterogeneity in collective action

This section presents plots that were present in the original model and were reproduced with the new implementation, as a way to prove that it is similar to the original version.

1) Player Count and Agreement Importance: The authors of the original model tested how different values for the agreement importance and number of players affect the outcome of the negotiations. Figure 1 presents the results to this simulation. It is possible to conclude that both features have a large influence in the success of the negotiations. A larger number of players makes negotiations harder to achieve, as well as higher values for the agreement importance, because higher values imply less devalue to the player's payoff.

2) Player size heterogeneity: Heterogeneity in the player's size also has an influence in the success of negotiations. Table II shows that. Heterogeneity in the player's size helps reaching positive results because it introduces additional constraints to the player's move set, resembling games with fewer players, thus increasing the success rates. Further analysis also indicates that successful negotiations are more likely to occur in situations where larger players give up a larger part of their emissions, which is natural to observe, given that those players make up a larger portion of the total emissions. This can be observed in Table III.

Table II: Frequency of success in an eight-player negotiation where players 1 and 2 (s_1 and s_2) are varied in size. s_1 = s_2

Sum of s1 and s2	25% of BAU	40% of BAU	50% of BAU	66% of BAU
Freq. of success	0.1850	0.2400	0.3070	0.6460

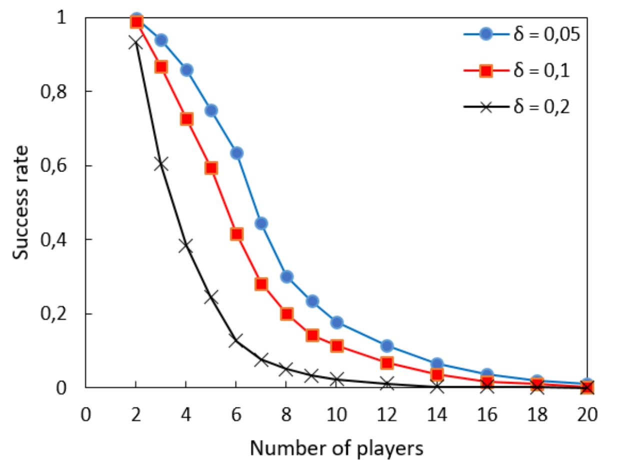

Figure 1: Frequency of reaching a feasible solution as a function of the number of players involved, and the effect of the agreement importance in reaching successful negotiations.

Table III: Average demand in successful agreements as a function of an agent's size

Agent size	1.0	2.0	3.0	4.0	5.0	6.0
Average final demand in successful negotiations	0.688	0.607	0.556	0.506	0.476	0.421

3) Initial Demands Restrictions: The initial demands are very important to the success of negotiations. If too many players start with high demands, the collective effort by the players will have to be higher, in order to succeed. Mechanisms that incentive initial demands to be closer to the target value improve the success rate of the negotiations. Figure 2 shows exactly that. Restrictions of both minimum and maximum demands force initial demands to be closer to the target value, thus increasing the probability of success in a negotiation. Restrictions on the minimum demand simply increase the average initial demand, making negotiations harder. On the other hand, restrictions on the maximum demand reduce the average initial demand, helping in negotiations.

4) Side Deals: Side deals work similarly to restrictions, except they last for whole game. The original authors tested how this influences the outcome of negotiations. Table IV presents the results of these tests, using my implementation. Constraining a subset of large players is not as effective as constraining an equivalent subset of smaller players. This suggests that prior agreements between smaller players is a better technique than prior agreements between larger ones.

C. Learning Process

The original model [7] has a player average the previous demands of another player in order to create expectations on him. Here, two new options of creating expectations are introduced. The first is the simplest expectation possible: the player expects the others will propose the same demand as in the previous round. The second option has the player average

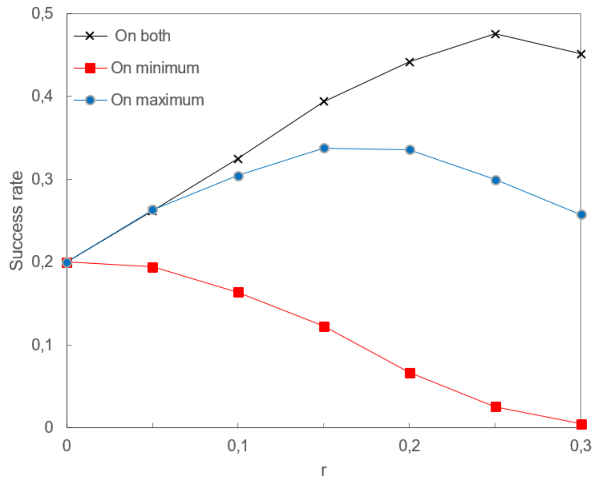

Figure 2: The effect of restricting initial demands.

Table IV: Success rates under side deals in an eight-agent simulation, where to of which comprise 25% of the total player size.

Scenario	Success Rate
Baseline (no prior reductions)	0.2980
10% prior reduction from two largest	0.3980
10% prior reduction from six smallest	0.4400
5% prior reduction from all	0.3920
20% prior reduction from two largest	0.5460
20% prior reduction from six smallest	0.7400
10% prior reduction from all	0.5960

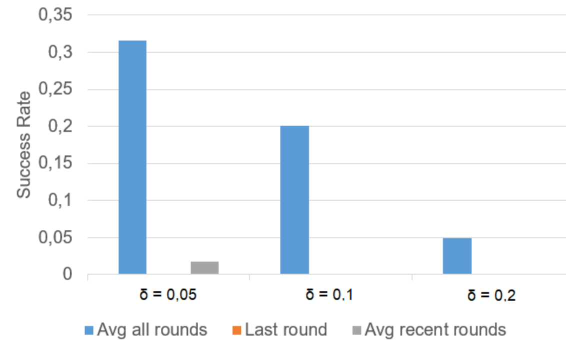

Figure 3: Success rate of negotiations using different learning processes. The orange label *Last round* exists to show that there is no chance of success when using this strategy.

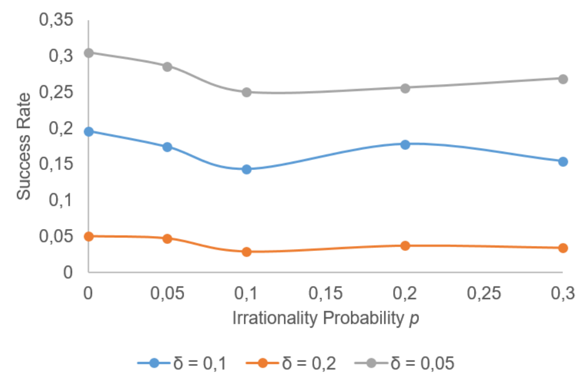

Figure 4: Frequency of success with different probabilities of irrationality.

the last demands from a player, more specifically, the last 8. The main advantages of both options are the simplicity in the expectation gathering, because now the players do not need to analyze the full set of previous demands to make a decision. The main disadvantage is the loss of information on previous demands and the susceptibility to oscillations on the rest of the field's demands. The results of simulating this new learning methods are presented in Figure 3. We can observe that none of the new strategies does, in fact, improve the chances of successful negotiations. When using the last round only to create our expectations, the outcome will surely be failure. When averaging the most recent rounds, success is, once again, rarely achieved.

D. Irrationality

Game Theory seldom studies irrational players, preferring to use strict payoff maximization strategies. Empirical knowledge, however, proves that theoretical studies can sometimes pale to conform to reality due to the adoption of purely rational agents. For that reason, the model incorporates irrationality. Every player has a probability p of being irrational. Whenever the player chooses his demand, with probability p, he will select a random demand from his strategy set $[\delta, 1]$. This mechanism is not expected to increase the success probability of the negotiations, but rather study how badly erratic players can affect the outcome. This mechanism's influence was simulated, and the results can be observed in Figure 4. Surprisingly, irrationality does not have a large influence in the outcome of the negotiations. The largest oscillation in the success rates is around 5%. This proves that the model and the players can adapt to irrational players and mitigate the effects of their unexpected moves.

E. Time-dependence and the decrease in target values.

Environmental issues are a pressing problem and, if nothing is done on each year, achieving satisfactory results becomes harder for the next years, and if more than needed is done, the following year will be slightly easier. That is exactly what is encapsulated in this mechanism. After each round of negotiations, their success is analyzed, and if they are below the target value T (meaning a temporary achievement has been achieved), this value is increased by a factor λ, making the negotiations easier. If the negotiations fail, the target value T is reduced by that factor λ. The target value has a ceiling, which is the original target value, i.e. it can not go above it. Do note that the game only terminates after round 100, and temporary successful negotiations are not

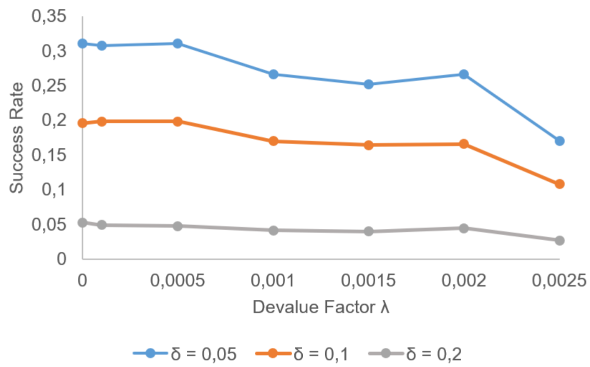

Figure 5: Different oscillation values for the common goal affect the success rate of the negotiations.

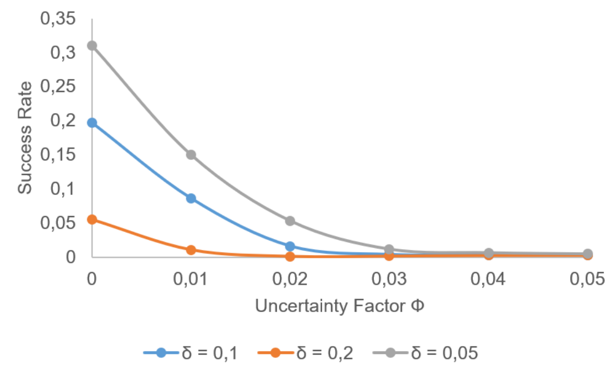

Figure 6: Feasibility of success using different values for ϕ.

counted as successes. Figure 5 presents the results of the simulations using different values of λ. For lower values of λ, the success rates of the negotiations is constant. When this value is increase, however, the negotiations break down. This can be justified with the existence of an upper limit in the target value. If 20 consecutive rounds are failed, it is hard to recover from the harder goals and players will likely opt for their BAU. On the opposite, if 20 consecutive rounds are successful, no change will occur to the target value due to its upper limit.

F. Uncertainty in Collective Goals

Uncertainty in negotiations is one of the major themes of this work. Following the previous mechanism's trend, and attempting to model reality more accurately, this mechanism brings another level of uncertainty to the collective goal. This one, however, is not dependent on the player's ability to reach agreements and is independent on the game's flow. A factor of target value uncertainty ϕ is declared when the game starts and, for every round, the target value T will oscillate between [T - ϕ, T + ϕ]. This increased uncertainty to the target value brings another level of dynamic to the game, rewarding different game strategies. Figure 6 is a plot with the results of simulating the model with different values of ϕ. The results clearly show that this added uncertainty does not bring success in any way. Increasing uncertainty by itself is not so interesting, because the higher the values of ϕ, the lower the chances of success. The main interest of this mechanism is its combination with other ones, as it will be presented ahead.

G. Risk Aversion

This mechanism introduces some realism to the players. Not all people are alike in terms of their willingness to take risks. Some like to gamble and take risks, while others believe that avoiding unnecessary risks is better. The model simulates this property in the strategies people take, in regards to risk. Players can be either risk prone or risk averse. The first will be optimistic and believe that the other players will concede more than expected, while the second group believes that they will concede less than what they expect. Risk aversion/proneness is

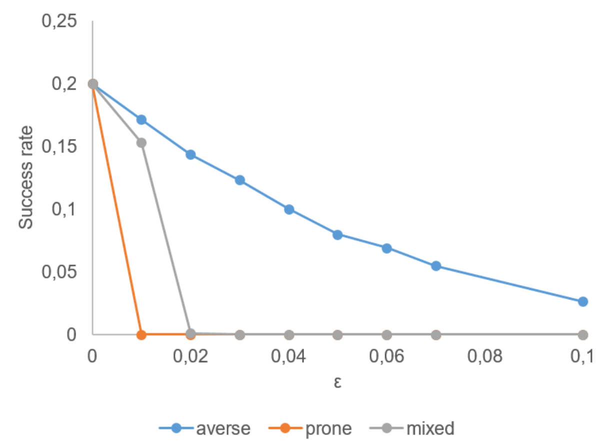

Figure 7: Different strategies and values of ϵ affect the outcome of negotiations. Each line represents a different risk aversion strategy used by the players.

represented by a factor ϵ, and is either added or removed after gathering expectations on the other players. Let ad_i represent player i's average demand. Player *j*'s expectation e_j using this mechanism is the following:

$$e_j = \frac{\sum_{i=1,i\neq j}^{N}(ad_i)s_i + \epsilon}{N-1} \tag{1}$$

The simulations were ran using 3 distinct populations: one containing only risk averse people, one containing only risk prone people, and other with an even distribution between risk averse and risk prone players. Figure 7 represents the success rate of the negotiations when using any of the strategies. A society composed of risk prone players never reaches success. A society that contains both risk prone and risk averse players is slightly more successful, but it still is barely possible to attain success. Analyzing risk averse players, the success of negotiations is better kept when ϵ grows, but it still provides worse results than using no risk aversion strategy at all. This is due to the fact that players can be so pessimistic about the results expected that they feel that keeping their BAU is a better option.

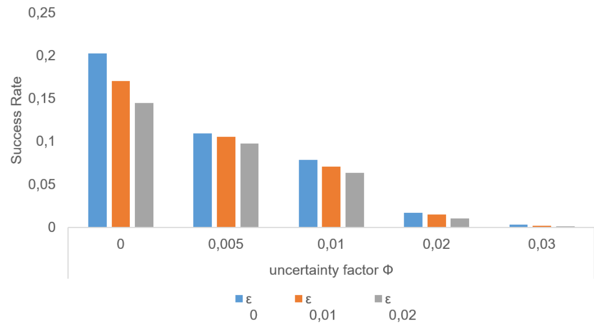

Figure 8: Different strategies for risk aversion combined with uncertainty in collective goals affect the negotiations, when dealing with a risk averse population.

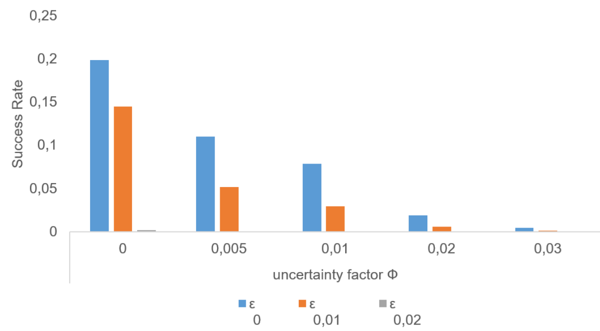

Figure 9: Different strategies for risk aversion combined with uncertainty in collective goals affect the negotiations, when dealing with an even distribution between risk averse and risk prone players.

H. Risk Aversion with Uncertainty in Collective Goals

While simply introducing uncertainty in the target value has some interesting properties that can be studied, pairing it with risk aversion leads us to some more interesting results. In this simulation of the model, the outcome is more unpredictable. Risk prone players, for example, can now get away with lower concessions if they are lucky enough to see a raise in the target value. The opposite can also happen, however, causing the negotiations to crumble.

Figures 8, 9 and 10 present the results for these simulations, using different kinds of populations. Figure 8 does not present surprising results, and follows the same trend of Figure 6. Figure 9 follows the same trend, except with even lower value of success. The results are truly interesting, however, when analyzing a risk prone population. When there is no uncertainty, any population composed of risk prone individuals will fail to ever achieve an agreement. On the other hand, when in the presence of uncertainty on the target value, populations of risk prone agents are able to achieve agreements and even outdo populations that are not prone to risk. This result is important, showing that being prone to risk may be beneficial, if in the presence of uncertainty in the collective goal.

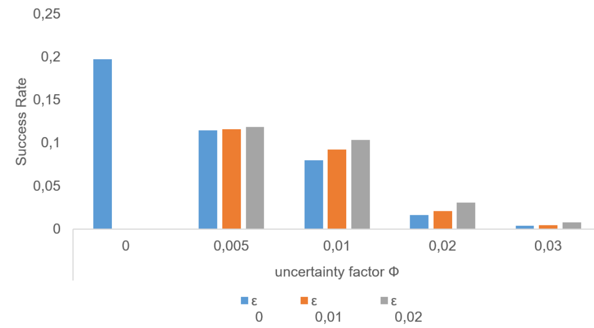

Figure 10: Different strategies for risk aversion combined with uncertainty in collective goals affect the negotiations, when dealing with a risk prone population.

V. Conclusion

The main goal of this thesis was studying different mechanisms that can influence cooperation, either positively or negatively. Cooperation is studied in a negotiation environment in which countries discuss their concessions relative to greenhouse gas emissions. An existent model was studied, due to its characteristics and already implemented features. After that, the model was expanded with newer features and mechanisms that allowed further study on cooperation. The main expansions on the model focused on the effects of uncertainty in the objective, irrationality on the players and different learning mechanisms by the players. This chapter summarizes the results and contributions obtained in this work.

A. Contributions Overview

1) Effect of Different Learning Methods: We studied how different methods of gathering expectations influence the outcome of negotiations and concluded that averaging the historic of demands provides better chances of success than the other methods studied.

2) Effect of Irrationality in the Player's Choices: Analysis of the effect of different likelihoods of irrationality by the players in each round concluded that the model is able to adapt to such erratic moves and mitigate its negative effects, up to a certain point.

3) Effect of Uncertainty in Collective Goals: We studied how uncertainty in the collective goal influences the player's moves and outcome of the negotiations and concluded that it decreases their success rate. However, combining this mechanism with others can improve the changes of successful negotiations.

4) Effect of Time-Dependence and Decrease in Target Values: Oscillations in the collective goal based on the success or lack of it in previous rounds showed that the feasibility of negotiations is mostly maintained, up to certain values of λ.

5) Effect of Risk Aversion: We studied how different strategies regarding the willingness to take risks affects the game's dynamic and its outcome. We concluded that risk averse societies tend to do better than their counterpart. Risk prone societies never achieve success in this model. This result is different, however, when uncertainty in the target value is introduced for these societies. In the presence of this mechanism, negotiations are possible, and their likelihood of success surpasses that of risk averse players.

B. Future Work

1) New Learning Methods: The expectation gathering methods tested in this work were rather simplistic. It can be interesting to try new methods and see how they affect cooperation in the model.

2) Introduction of Punishment and/or Reputation: Some work has been done in this area, as described in Section II. The introduction of more complex mechanisms such as punishment and reputation has potential to enhance our knowledge base.

3) Research Cooperation Mechanisms that allow Cooperation with Large Numbers of Players: We concluded that the number of players in the game has a major influence on the outcome of negotiations. Higher numbers of players mean that, with great probability, the negotiations will fail. International negotiations often deal with an amount of countries in the order of dozens, and even hundreds. The Paris Agreement, for example, was signed by 196 countries. Researching mechanisms that can explain how cooperation is encouraged in such situations may be one of the keys for the saving of the planet.

4) Combination of mechanisms: More conclusions as to how mechanisms such as irrationality, collective goal uncertainty, risk aversion and time-dependence are related can be extracted combining them. These mechanisms are mostly based on uncertainty, so it could certainly shed more light as to how it affects cooperation in bargaining situations.

REFERENCES

[1] J. M. Smith, "Evolution and the theory of games," in *Did Darwin Get It Right?* Springer, 1988, pp. 202–215.

[2] M. Finus, "Game theory and international environmental cooperation: any practical application?" *Controlling global warming: Perspectives from economics, game theory and public choice*, pp. 9–104, 2002.

[3] J. Heitzig, K. Lessmann, and Y. Zou, "Self-enforcing strategies to deter free-riding in the climate change mitigation game and other repeated public good games," *Proceedings of the National Academy of Sciences*, vol. 108, no. 38, pp. 15 739–15 744, 2011.

[4] T. Dietz and J. Zhao, "Paths to climate cooperation," *Proceedings of the National Academy of Sciences*, vol. 108, no. 38, pp. 15 671–15 672, 2011.

[5] J.-C. Altamirano-Cabrera and M. Finus, "Permit trading and stability of international climate agreements," *Journal of Applied Economics*, vol. 9, no. 1, p. 19, 2006.

[6] R. Dellink, M. Finus, and N. Olieman, "The stability likelihood of an international climate agreement," *Environmental and Resource Economics*, vol. 39, no. 4, pp. 357–377, 2008.

[7] R. Smead, R. L. Sandler, P. Forber, and J. Basl, "A bargaining game analysis of international climate negotiations," *Nature Climate Change*, vol. 4, no. 6, p. 442, 2014.

[8] M. A. Nowak, "Five rules for the evolution of cooperation," *science*, vol. 314, no. 5805, pp. 1560–1563, 2006.

[9] R. M. Axelrod, *The evolution of cooperation: revised edition*. Basic books, 2006.

[10] D. Fundenberg and E. Maskin, "Evolution and cooperation in noisy repeated games," *The American Economic Review*, vol. 80, no. 2, pp. 274–279, 1990.

[11] R. Selten and P. Hammerstein, "Gaps in harley's argument on evolutionarily stable learning rules and in the logic of "tit for tat"," *Behavioral and Brain Sciences*, vol. 7, no. 1, pp. 115–116, 1984.

[12] M. Nowak and K. Sigmund, "A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game," *Nature*, vol. 364, no. 6432, pp. 56–58, 1993.

[13] L. A. Imhof, D. Fudenberg, and M. A. Nowak, "Tit-for-tat or win-stay, lose-shift?" *Journal of theoretical biology*, vol. 247, no. 3, pp. 574–580, 2007.

[14] A. Traulsen and M. A. Nowak, "Evolution of cooperation by multilevel selection," *Proceedings of the National Academy of Sciences*, vol. 103, no. 29, pp. 10 952–10 955, 2006.

[15] E. Fehr and S. Gächter, "Cooperation and punishment in public goods experiments," *The American economic review*, vol. 90, no. 4, pp. 980–994, 2000.

[16] J. M. Pacheco, F. C. Santos, and F. A. C. Chalub, "Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity," *PLoS computational biology*, vol. 2, no. 12, p. e178, 2006.

[17] F. C. Santos and J. M. Pacheco, "Risk of collective failure provides an escape from the tragedy of the commons," *Proceedings of the National Academy of Sciences*, vol. 108, no. 26, pp. 10 421–10 425, 2011.

[18] F. C. Santos, M. D. Santos, and J. M. Pacheco, "Social diversity promotes the emergence of cooperation in public goods games," *Nature*, vol. 454, no. 7201, p. 213, 2008.

[19] V. V. Vasconcelos, F. C. Santos, J. M. Pacheco, and S. A. Levin, "Climate policies under wealth inequality," *Proceedings of the National Academy of Sciences*, vol. 111, no. 6, pp. 2212–2216, 2014.

[20] G. Gosnell and A. Tavoni, "A bargaining experiment on heterogeneity and side deals in climate negotiations," *Climatic Change*, vol. 142, no. 3-4, pp. 575–586, 2017.

[21] J. F. Nash Jr, "The bargaining problem," *Econometrica: Journal of the Econometric Society*, pp. 155–162, 1950.