TÉCNICO LISBOA

Bargaining Dynamics in an Uncertain World

Mateus Valadas Preto Bravo Cordeiro

Thesis to obtain the Master of Science Degree in

Information Systems and Computer Engineering

Supervisors: Prof. Francisco João Duarte Cordeiro Correia dos Santos
Prof. Fernando Pedro Pascoal dos Santos

Examination Committee

Chairperson: Prof. Nuno João Neves Mamede
Supervisor: Prof. Francisco João Duarte Cordeiro Correia dos Santos
Member of the Committee: Prof. António Manuel Ferreira Rito da Silva

February 2018

Acknowledgments

I want to thank you, the one reading right now. After all, if it weren't for you, this work would serve no purpose. I also want to acknowledge some people that were, directly or not, involved in the building of this product I present to you today. If you find yourself in the set {friends from school, university, hometown, Erasmus and friend's friends, family, thesis advisors}, allow me to thank you from the bottom of my heart, for either your support or presence in my life.

Those who belong to the group {parents, girlfriend, grandparents, special buddies} deserve a special mention, because you were my pillars, not only in this last months, but during all my life. I am your making, and you are mine.

For my grandmother. Your forth boy is almost a man.

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.

Keywords

Cooperation; Game Theory; Uncertainty; Risk Aversion; Irrationality; Bargaining;

Resumo

As mais recentes cimeiras que visam reduzir as emissões de gases de efeito de estufa não estabeleceram acordos suficientemente robustos. Mudanças por um país apenas não tem impacto suficiente para mitigar este problema. A salvação do planeta requer cooperação. Perceber como occore a cooperação é fundamental à resolução deste problema. A Teoria de Jogos permite o estudo deste tipo de conflitos entre interesses sociais e individuais. Este trabalho foca-se no estudo da cooperação em situações de negociação, desenvolvendo modelos com novos mecanismos e estudando o seu efeito na cooperação. Os modelos são centrados na dinâmica das negociações na presença de incerteza nos objetivos coletivos e nos processos de decisão dos países. O estudo destes mecanimos é possível usando simulações em sistemas multi-agente.

Concluímos que vários fatores e a sua combinação podem ter um grande efeito no resultado e na viabilidade das negociações entre países. Mostramos que a utilização de todas as jogadas de um jogador para prever a sua próxima produz melhores resultados que utilizando apenas um subconjunto. Irrationalidade nas escolhas dos jogadores influencia negativamente as negociações, apesar dos jogadores conseguirem adaptar-se a jogadas inesperadas. Verificámos que os jogadores têm problemas na adaptação a variações no objetivo coletivo, sejam estas aleatórias ou dependentes do desempenho destes. Análise de diferentes estratégias de disposição ao risco provou que populações cautelosas prosperam melhor que populações que arriscam. As últimas, no entanto, obtêm melhores resultados quando o objetivo global é incerto, ultrapassando até sociedades cautelosas. Ser dado ao risco pode então ser benéfico na presença de incerteza nos objetivos.

Palavras Chave

Cooperação; Teoria de Jogos; Incerteza; Aversão ao Risco; Irracionalidade; Negociações;

Contents

1 Introduction **1**

1.1 Thesis outline 5

2 Related Work **7**

2.1 Mechanisms for Cooperation 9

2.1.1 Kin Selection 9

2.1.2 Direct Reciprocity 9

2.1.3 Indirect Reciprocity 10

2.1.4 Group Selection 10

2.2 Mechanisms for Cooperation in Bargaining Situations 10

2.2.1 Social Diversity 10

2.2.2 Risk Perception 11

2.2.3 Punishment 11

2.2.4 Reputation 11

2.3 Analysis and Relevance 12

3 Model **13**

3.1 A Bargaining Game of International Climate Negotiations 15

3.1.1 Accounting for size 16

3.1.2 Restrictions on initial demands 16

3.1.3 Side Deals 16

3.2 Implementation 17

3.2.1 Overview 17

3.2.2 Inside a round 18

3.3 Preliminary Results 18

3.3.1 Simulation Design 19

3.3.2 Results 19

3.3.2.A Player Count and Agreement Importance 19

3.3.2.B Player Size Heterogeneity 20

3.3.2.C Initial Demands Restrictions 21
3.3.2.D Side Deals 21
3.3.3 Additional details 22
3.3.3.A Ending Round 22
3.4 Preliminary Results Analysis 23

4 Results and Contributions **25**
4.1 Learning Process 27
4.2 Irrationality 27
4.3 Time-dependence and the Decrease of Target Values 28
4.4 Uncertainty in Collective Goals 30
4.5 Risk Aversion 30
4.5.1 Risk Aversion coupled with Collective Goals Uncertainty 32

5 Conclusion **35**
5.1 Future Work 39

List of Figures

3.1 UML representation of the code's classes . . . 17
3.2 Input file example . . . 18
3.3 High-level representation of the code's flow . . . 19
3.4 Variation in Player Number and Agreement Importance. Original . . . 20
3.5 Variation in Player Number and Agreement Importance. New version . . . 20
3.6 Varying *r* on initial demands. Original . . . 22
3.7 Varying *r* on initial demands. New version . . . 22

4.1 Expectations with different learning processes . . . 28
4.2 Results of varying *p* . . . 29
4.3 Results of varying λ . . . 29
4.4 Results of varying ϕ . . . 30
4.5 Results of varying ϵ . . . 31
4.6 Effect of combining ϵ and ϕ in a risk averse population . . . 32
4.7 Effect of combining ϵ and ϕ in a risk mixed population . . . 32
4.8 Effect of combining ϵ and ϕ in a risk prone population . . . 33

List of Tables

1.1 Prisoner's Dilemma payoff 3

3.1 Variation of proportion of player size. Original 20
3.2 Variation of proportion of player size. New version 20
3.3 Average Demand as function of agent size. Original 21
3.4 Average Demand as function of agent size. New Version 21
3.5 Effect of side deals. Original 22
3.6 Effect of side deals. New version 23

1

Introduction

Contents

1.1 Thesis outline 5

Bargaining is part of life. Negotiations involve presenting proposals and making concessions to satisfy interests. The feasibility of reaching an agreement is often built on the will of the participants to cooperate, that is, to endow others some benefit to help satisfy their needs. Understanding how socially desirable bargaining solutions are promoted can benefit from the extensive literature on cooperation and altruism.

Cooperation occurs when a party benefits another, being subjected to a cost. A self-interested individual will always weight his winnings and losses when he makes an informed decision. Cooperation seems to be out of the realm of possibilities for this individual, because there is no benefit in loss. The rational decision is waiting to get benefits from other's cooperation, while benefiting no one. This is by no means ground breaking information, as Darwin has previously stated in his Natural Selection theory [1] that the fittest individuals will survive and thrive if, and only if, they develop traits that allow them to reproduce in a competitive environment. This theory supports the claim that rational individuals should be selfish instead of cooperative. This is, however, not the case, as we can observe a multitude of situations where cooperation occurs in the nature. From cells forming tissues to animals living in organized herds, cooperation is everyone in the nature. Even in our own society, the simple act of recycling can be used as an example. How does cooperation prevail then, in a world in which the rational decision 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 is the study of mathematical models of conflicts of interest between individuals. These models are analogies to situations in which agents interact. The agents encapsulate some sort of behavior for those situations based on what is either known or believed to be correct. In Game Theory, cooperation problems often involve a tension between the individual interest and social desirability, being usually described as Dilemmas.

A better way to explain what Game Theory actually studies is with an example. There is a Donation Game that captures the essence of cooperating, where players can incur in a cost c to provide a benefit b. When this game is played by two individuals, we have a Prisoner's Dilemma [2], in which individual interests tend to hinder the best social outcome. This Prisoner's Dilemma explains why cooperation may fail in the presence of two individuals. In this situation, two individuals act simultaneously and have the option of either Cooperating or Defecting. The set of possible payoffs is described in Table 1.1.

Table 1.1: Prisoner's Dilemma payoff

	Cooperate	Defect
Cooperate	benefit - cost	- cost
Defect	benefit	0

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. In this case, any change on his decision will grant him a negative payoff compared to its current situation. When both players face this decision, we are in the presence of a Nash Equilibrium [3]. A Nash Equilibrium is a situation in which a change of strategy will benefit neither of the players. However, in this case, the best scenario is the cooperation of both players. Thus, we are in the presence of a social dilemma in which cooperation is not the rational choice, however constituting the socially desirable action.

Game Theory has focused much of its studies in models of interaction between two parties, instead of multiple party models, due to the technical difficulties in the analysis of the later, as W.D Hamilton best describes *"The theory of many person games may seem to stand to that of two-person games in the relation of sea-sickness to a headache"* [4]. The Prisoner's Dilemma can be extended to an N-Person Game, and is usually referred to as Public Goods Game. In this game, players possess wealth and can offer some to a public fund. After each round, the amount of money in said fund is multiplied by a factor larger than 1, and split equally between all players. Once again, as in the Prisoner's Dilemma, the best strategy for a rational player is Defection as a player's wealth will increase as much as a cooperator, and without the cost of Cooperation. The best collective strategy is, then again, collective cooperation. How can populations evolve in a way that enables cooperation? Evolutionary Game Theory may help us with this question.

Evolutionary Game Theory, first proposed by John M. Smith [5], is an attempt of mapping 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.

Bargains are also frequently studied by Game Theory. These negotiations can be comprised of two or more agents, with single or multiple rounds of negotiations, but the constant is that each player tries to maximize its payoff. The Ultimatum Game [6] is an example of a bargaining game. In it, a Proposer is given something, which he will split, as he prefers, with a Receiver. The Receiver can either accept the offer, with both agents keeping their share, or refuse the offer, in which case they both losing everything. This dilemma can also be extended to an N-Person Game , and it has been studied extensively before [7–9]. Indeed, bargaining situations, more often that not, occur in a group context, like international summits.

In international summits regarding climate changes, countries often negotiate reductions in greenhouse gas emissions. In these summits, targets are set for the reduction of the participating countries' current emissions. These targets are based on negotiations between countries. Using Game Theory, this situation has been studied before [10–15], and my thesis will continue to do so, expanding upon an existing N-Person Bargaining Game [16].

The model I will be expanding is a version of Nash's Bargaining Game [17] for N players. On each round, each player decides how much he is willing to reduce on his emissions, based on his expectations of the other players' decisions. After each round, if the average reduction of all players meets a set target, the negotiation is considered successful. Otherwise, a new round will take place. The player's payoff depends on the success of this negotiation. I will detail the intricacies of the model in Section 3.

The model will be expanded introducing new variables and complexity to the game. One of the major topics of study of this work is uncertainty, thus the majority of the variables introduced have the goal of emulating it. In the new version of the model, players now have an irrationality probability p for their moves. They can also create expectations on other players using new learning processes/methods. Players can be either risk averse or risk prone when gathering such expectations on the field, and have a cautious or venturesome strategy introduced by the risk aversion factor ϵ. The new variables also affect the game, and now, the collective goal has two oscillating factors that can either depend on the player's moves and ability to reach feasible agreements (λ) or be entirely independent on the game's flow, randomly oscillating the target value (ϕ).

1.1 Thesis outline

The current chapter introduces the topics covered in the thesis, and points the reader in the direction followed by this work. Chapter 2 presents the state of the art in regards to cooperation and bargaining, mainly. In Chapter 3 the model under study is introduced, studied and proved to be in line with the model the thesis is based on. Afterwards, in Chapter 4 the model is expanded and tested. The results of this tests are discussed there. At last, Chapter 5 presents an overview of the work conducted, as well as some final notes and prospects on future work.

2

Related Work

Contents

2.1 Mechanisms for Cooperation 9

2.2 Mechanisms for Cooperation in Bargaining Situations 10

2.3 Analysis and Relevance 12

Chapter 1 made it clear that cooperation does occur, but no reason other than empirical knowledge is given to explain how it evolves. 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 topic is thoroughly studied in Evolutionary Game Theory [18, 19]. The situations are modeled defining the agents' behavior and choosing the strategies that can be taken and then, through repeated iterations, the negotiation outcome is collected.

This section will present the work that has been done in this area that is relevant to my thesis, focusing mainly on the emergence of cooperation and the mechanisms that enable it.

2.1 Mechanisms for Cooperation

2.1.1 Kin Selection

Kin selection is defined as cooperation between related individuals, more specifically, individuals that share genes. Hamilton's rule [20] states that the probability of sharing a gene must exceed the cost-to-benefit ratio of an altruistic act. While interesting and clearly observable in a society, it pales to explain how cooperation exists amongst societies in which the agents are not related.

2.1.2 Direct Reciprocity

Direct reciprocity assumes that individuals will have multiple encounters and will remember each other. In those encounters, an agent will choose to cooperate instead of defecting because in later rounds the other may cooperate as well. This mechanism is a great strategy for the Prisoner's Dilemma, as Axelrod proved with his strategy Tit-For-Tat [21]. Tit-For-Tat is the simplest strategy possible in the game: you start by cooperating, and on all the subsequent rounds, you imitate your opponent's move. This strategy is so good, that Axelrod actually won two iterations of this tournament. While this theory looked to be the undisputed champion, posterior work concluded that it has flaws, especially when you consider abnormal or unexpected moves [22,23]. Tit-for-tat's issue is the impossibility of error correction when the opponent unexpectedly defects, leading to an unending set of defections. To face this problem, the strategy Win-Stay, Lose-Shift was created [24]. As Tit-For-Tat, this strategy is simplistic and states that whenever you have a good round, you maintain your current move, shifting it otherwise. Comparison between both strategies proved that the latter is superior to maintain cooperation [25].

2.1.3 Indirect Reciprocity

While direct reciprocity explains how cooperation evolves in repeated encounters, it does not explain why it happens in scenarios where retribution is not available, like donating. The benefit one expects from this behavior is good reputation and the prospect of being rewarded for it. Creating a model that encapsulates indirect reputation is slightly more complex than direct reputation, because it requires not only random encounters between agents who do not know each other, but also viewers who can attest for one's cooperation [26–29]. In even more complex models, this reputation can be spread to the rest of the agents in the model via gossip [30]. Unlike direct reciprocity, the choice of cooperation or defection is not straightforward. Individuals must monitor every other individual in the network and think ahead when making a decision. This complexity makes indirect reciprocity much more common in humans than in other species, mostly because animals do not gossip and have a limited rationale.

2.1.4 Group Selection

All the previous mechanisms apply to individuals, but do not explain how groups can be successful. Group Selection has the answer. Nowak and Traulsen [31] 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 threshold, it can split in two, eliminating another existing group, for population control. This study has shown that pure cooperative groups grow faster than pure defector groups, while in mixed populations, defectors reproduce faster.

2.2 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.

2.2.1 Social Diversity

Not all individuals are the same, and there are great disparities among nations in wealth and size. That disparity can make agreements hard to achieve. Studies in this area [9, 32–34], however, showed that under the right circumstances, such as risk perception 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.

2.2.2 Risk Perception

Santos and Pacheco [35, 36] showed that the perception of a collective risk enhances the probability of successful negotiations. In their model, each individual has some wealth and can either donate some (cooperate) or keep it (defect). If, as a whole, a population of N individuals fails to reach M cooperators, with probability r (risk), everyone will lose their wealth. Their result in an unsurprising increase in cooperation when the perceived risk of loss is raises.

Another study in the same topic [37] of risk awareness shows that under high risk of loss, players tend to cooperate, as in the previous model. In this model, players are given money and, during 10 successive rounds, are asked to donate either 0, 5% or 10% of it to a global fund. If at the of the 10 rounds the groups manages to reach a threshold, each player keeps the remainder of their money. If they fail to achieve that value, they will lose their money with probability p. The results showed that when p is 90%, half of the groups are able to achieve negotiations. With p at 50%, only one tenth of the groups managed to succeed and for p at 10%, no group was able to reach the threshold.

Both studies reveal that success in negotiations targeting climate changes is highly dependent on risk perception. Thus, to improve cooperation, people need to be convinced that failure to negotiate will lead to complete loss.

2.2.3 Punishment

In an attempt to study how free-riding, i.e, defection, affects one's feelings towards a defector, Fehr and Gächter developed a model [38] that lead to the conclusion that even if there are no palpable benefits for a punisher and punishment is expensive, it will usually be carried out, due to negative feelings towards the defector. They also concluded that the presence of punishment opportunities can reduce the amount of free-riders and/or their degree of defection.

2.2.4 Reputation

The previous section's Indirect Reciprocity mechanism can also be applied to bargaining situations, as it has been studied to some extend before [39, 40]. An individual's reputation can be modeled as a number that goes up when help is given to another individual with good reputation or when help is refused to one with bad reputation, and goes down when the opposite happens. While similar to punishment, reputation acts in a different way, as there is no direct form of punishment, and the possible cooperator actually saves money when refusing to help an individual with bad reputation.

2.3 Analysis and Relevance

This section shows that the work performed in cooperation and the mechanisms that enable it is quite extensive, even in bargaining situations. Another trend we observed was the common reference of the climate changing issue and how we can model societies/groups to understand its implications and how to succeed in its negotiations. For that reason, we chose to expand upon Smead's model [16], since it already implements a large set of the aforementioned mechanisms. The next section will detail how this is achieved in practice.

3

Model

Contents

3.1 A Bargaining Game of International Climate Negotiations 15
3.2 Implementation 17
3.3 Preliminary Results 18
3.4 Preliminary Results Analysis 23

This section will provide an overview on the model [16] and details in its implementation and intricacies. It is an extension to Nash's bargaining game [17] with N players, so we can define it as an N-Person Bargaining Game. Typically, in bargaining games, agents make offers that can either be accepted or rejected. When those offers are accepted, a successful negotiation is achieved and the agents are rewarded. An agent's goal is maximizing its payoff. However, given the fact that we are dealing with a multi-agent environment, the implementation of this goal is not obvious at all. A possible way, underlying the process of Fictitious Play implemented in this work, uses the information about opponents' past actions in order to derive a best response.

This bargaining game also aims to replicate, to a certain point, the dynamics involved in negotiations between countries for climate changes.

While my model is heavily based on Rory Smead et al.'s model [16], I expand upon it and alter its behavior and mechanisms. First I will describe the original model and prove that my implementation is correct, and will then proceed to detail my changes and results.

3.1 A Bargaining Game of International Climate Negotiations

The model aims to replicate negotiations between countries, so these countries are modeled in the players. Each player has a continuous strategy set that ranges from [0, 1], and it has the following meaning: 1 represents the Business As Usual (BAU), meaning 100% of its current emissions and 0 means no emissions at all. Let d_i represent a player's demand for a proportion of their BAU emissions 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, meaning negotiations will have a larger success rate. Lower values will, on the other hand, make negotiations harder, thus reducing their success rate.

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 given 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, for a given demand set, is the defined as:

$$\pi_i(d_i, t) = \begin{cases} d_i & \text{if t} \leq \text{T.N} \\ \delta d_i & \text{if t} > \text{T.N} \end{cases} \tag{3.1}$$

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. 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.

To represent learning agents during negotiations, each player undertakes a learning process. All demands done by a player are visible to every other agent. After each round of negotiations, players may learn and change their demands. The learning process is modeled using expectations on other players. The expectations are formed accessing their former demands and averaging all the values. After forming expectations on all other players, a player's demand is presented in a way to maximize its payoff, taking those expectations into account.

On the first round, as there are no previous demands to learn from, a random value is generated for each player, and that value is in the range [δ, 1]. All the demands are gathered simultaneously. During the course of the game, a player's strategy set is also [δ, 1].

The next subsections present mechanisms implemented by the original authors, to allow for some further study on the model.

3.1.1 Accounting for size

Countries differ in size, wealth, volume of emissions and other parameters. A 5% emission reduction in a country like China is much more important than a 5% reduction from Portugal. To account for these differences, the model incorporates the concept of player size. A player's size is represented by $s_i \in$ [1, ∞[, and this new feature affects the payoff function. The new updated function is as follows:

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

$$\pi_i(d_i, t) = \begin{cases} d_i & \text{if t} \leq \text{T} \sum_{i=1}^{N} s_i \\ \delta d_i & \text{if t} > \text{T} \sum_{i=1}^{N} s_i \end{cases} \tag{3.2}$$

3.1.2 Restrictions on initial demands

At the start of the game, each player is assigned a random first demand, ranging from [δ, 1]. If too many players start with high demands, the negotiations usually break down, because the further away the initial demands are from the objective, the more collective action is needed to achieve that goal. With that said, any mechanism that instigates initial demands to be closer to the target will increase the success rate of the negotiations. Thus, restricting the initial demands will likely be beneficial. They can either be upper bound, lower bound, or both. Each of them will affect the interval from which the initial demand is chosen. For an upper bound, this interval is [δ, 1 - *r*]. For lower bounds, the interval is [δ + *r*, 1]. When restricting both, the initial demand will be chosen from the interval [δ + *r*, 1 - *r*].

These restrictions apply only to the first round of negotiations, and for the rest of the game, a player's strategy range is the usual [δ, 1].

3.1.3 Side Deals

In real negotiations for climate changes, countries can have some prior negotiations or coalitions with other countries. To account for that, it has the possibility of emulating side deals. These restrict the maximum demand that a country can present not only on the first round, but in the whole game. In the end, it works in the same way as the restriction mechanism, but it lasts for the entirety of the game. This new feature will reduce, on average, the player's initial demands and restrict their demands during the game.

When this mechanism is active, the player's strategy range is [$\delta\ Max_i$, Max_i].

3.2 Implementation

Before the introduction of any new mechanism, I implemented the model and verified its correctness in comparison to the original, previously described model.

The code is written in Java, using the Object Oriented paradigm. Figure 3.1 is a UML representation of its classes.

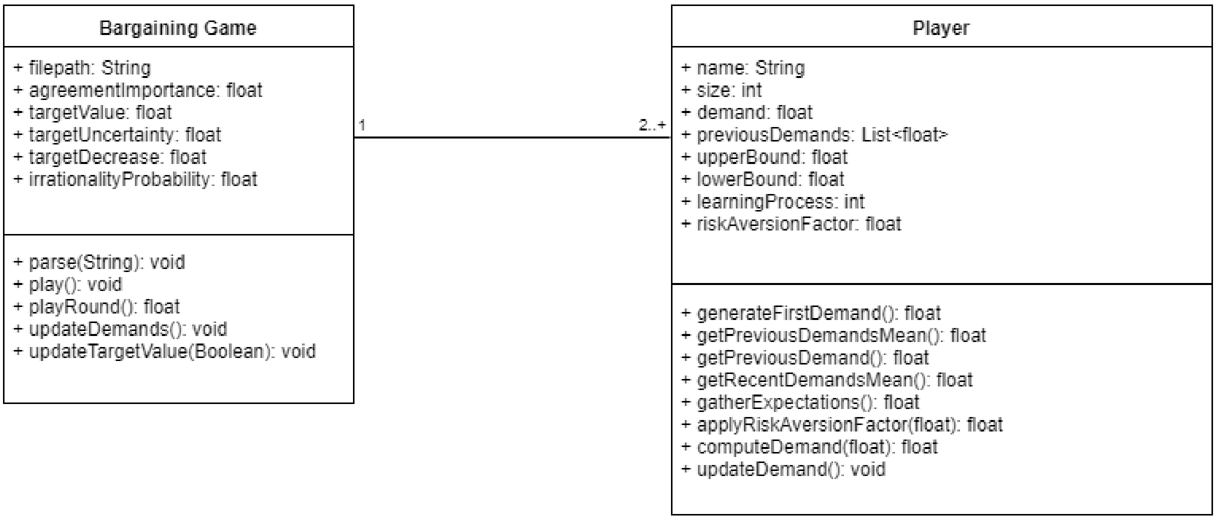

Figure 3.1: UML representation of the code's classes

The main class receives, as input, a *filepath* to the game's configuration file. Figure 3.2 is an example of an input file with a Target Value of 0.5, Agreement Importance of 0.1 and 8 players of equal size, with no other relevant mechanisms active. The code then analyses this file and starts the game. The program's flow is described in Figure 3.3.

3.2.1 Overview

The game starts by creating all the *Player* objects and generates, for each of them, an initial demand. After that, the negotiations ensue. On every single round, each player computes its demand and then, at the same time, the players reveal them.

```
target_value 0.5
agreement_importance 0.1
initial_restriction f
decreasing_target 0.0
oscilating_target 0.0
irrationality 0.0
learning_mode 1
ending_round 100
risk_aversion_type n
risk_aversion_factor 0.0
A 1 0.0 0.0
B 1 0.0 0.0
C 1 0.0 0.0
D 1 0.0 0.0
E 1 0.0 0.0
F 1 0.0 0.0
G 1 0.0 0.0
H 1 0.0 0.0
```

Figure 3.2: Input file example

After 100 rounds of negotiations, the set of demands is gathered and compared to the target value. If that set of demands is lower than the target value + 0.01, the negotiation is considered successful. Otherwise, it is considered unsuccessful.

3.2.2 Inside a round

Upon a player's creation, its first demand is generated. Every other demand computation takes into account a player's expectation of the rest of the field's next demand. Expectations on a specific player are formed using the mean of its previous demands. After estimating the rest of the field's moves, a player must decide its next demand. That calculation is made by subtracting the expectations to the target value. This determines the remainder of the available emissions given the expectations of the rest of the field. The player's response is this remainder, given that it falls within his strategy set. If it does not, the player will choose 1 for his demand, or Max_i if its moves are restricted due to side deals. Having the goal of maximizing its payoff, and knowing the payoff function, the player will choose the demand. To simulate simultaneous proposals, the player's demand will only be updated after everyone decides their next move.

3.3 Preliminary Results

To prove that my implementation of the model is correct, I emulated all of the original paper's plots, ran them using my code and compared the both. This section presents the plot's comparisons.

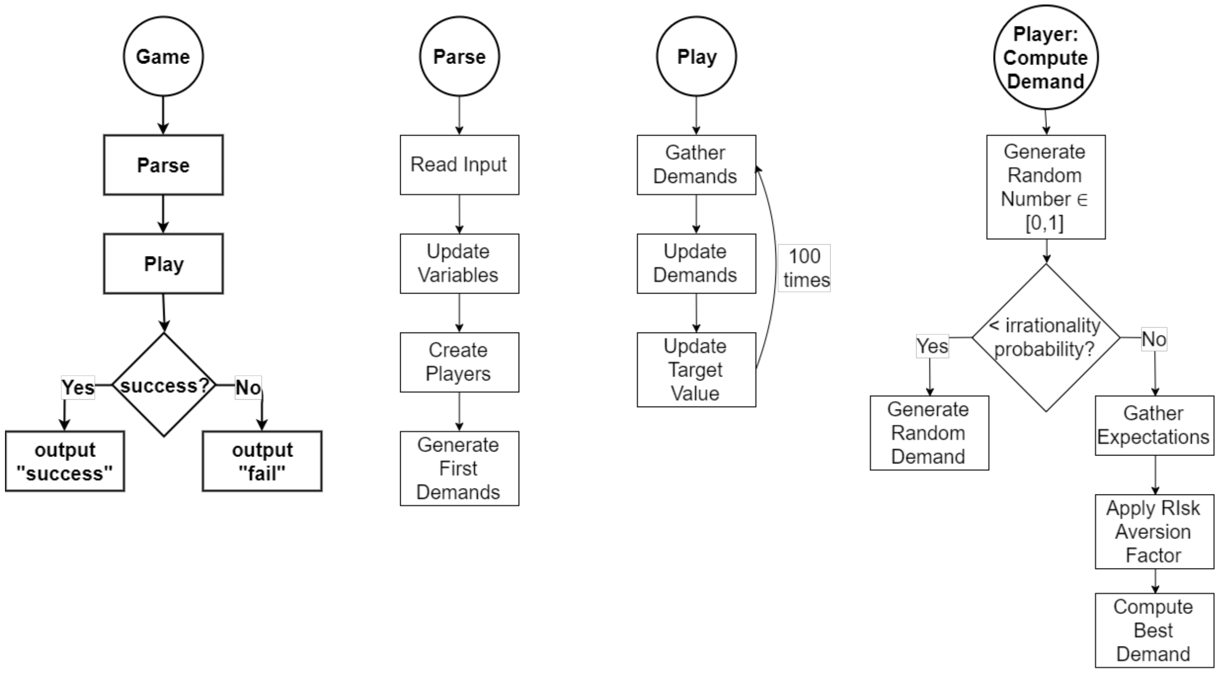

Figure 3.3: High-level representation of the code's flow

3.3.1 Simulation Design

Each datapoint for each plot was simulated 10^4 times. Unless stated, the following values are constant for the plots:

- Agreement Importance δ = 0.1
- Target Value T = 0.5
- Rounds in a Negotiation = 100
- Number of Players = 8
- Player Size = 1

The results of each negotiation are written in a file, stating whether is was successful or not and, in some cases, the average player demand.

3.3.2 Results

3.3.2.A Player Count and Agreement Importance

The success of a negotiation depends on a multitude of factors. An increase in the number of players will decrease its chance of success. The agreement importance δ also impacts the success rate. Larger values of δ will raise the minimum initial demands and have a smaller impact in the payoff function, reducing the success probability.

Figs. 3.4 and 3.5 compare my plot to the original one. Both plots follow the same trend, and have similar values.

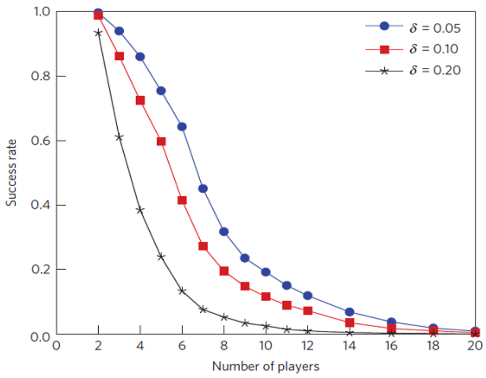

Figure 3.4: Success rate as a function of the number of players involved and agreement importance. Original paper's version.

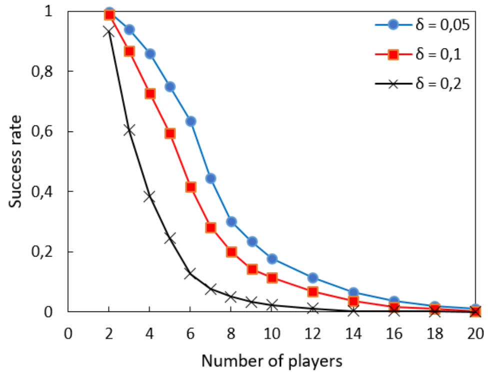

Figure 3.5: Success rate as a function of the number of players involved and agreement importance. My implementation's version.

3.3.2.B Player Size Heterogeneity

Heterogeneity in the player's size also affects the success of negotiations. Increasing the size increases the chance of successful negotiations, as observed Table 3.1. Heterogeneity helps reaching positive results because it introduces additional constraints to the players' move set, resembling games with fewer players, thus increasing the success rate. Table 3.2 presents my model's implementation results on the same simulations, which resemble the original one.

Table 3.1: Success rate in an eight-player negotiation with players 1 and 2 making up a proportion of the global BAU. Original paper's version.

Sum of s1 and s2	25% of BAU	40% of BAU	50% of BAU	66% of BAU
Freq. of success	0.1982	0.2173	0.2885	0.6143

Table 3.2: Success rate in an eight-player negotiation with players 1 and 2 making up a proportion of the global BAU. My implementation's version.

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

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 3.3. Table 3.4 is my version of this table which, once again, is similar to the original.

Table 3.3: Average demand in successful negotiations as function of the agent's size. Original paper's version.

Agent size	1.0	2.0	3.0	4.0	5.0	6.0
Average final demand in successful negotiations	0.687	0.606	0.545	0.510	0.473	0.426

Table 3.4: Average demand in successful negotiations as function of the agent's size. My implementation's version.

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

3.3.2.C Initial Demands Restrictions

The initial demands are very important to the success of the negotiations. If too many players have high demands, the collective effort by the players in order to succeed will be higher, meaning the negotiations will have a smaller chance of success. It is then important to introduce mechanisms that inspire initial demands closer to the target value, at least in the upper bounds.

The introduction of restrictions on initial demands proved to be beneficial, as long as it is, at least, on the maximum value. On Figure 3.6 we can observe, however, that restrictions on the minimum value only will produce worse results, as that is only raising the average initial demand. Restrictions on both minimum and maximum demands will produce the best results, because it narrows the range of possibilities to values closer to the Target Value. Figure 3.7 is my version of this simulations, and it is very similar to the original version.

3.3.2.D Side Deals

Side deals are modeled as restrictions that last for the whole game, and can be applied to subsets of players, instead of the whole field.

Analysis of this mechanism produced interesting results, that can be seen in Table 3.5. Constraining, for example, 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. Table 3.6 contains the results obtained by my implementation, and they are in line with the original ones.

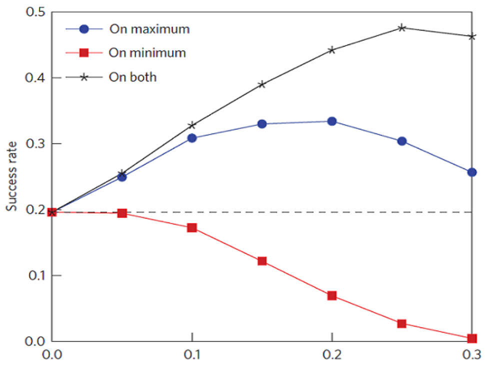

Figure 3.6: Success rate of introducing restrictions on the initial demands. Original paper's version.

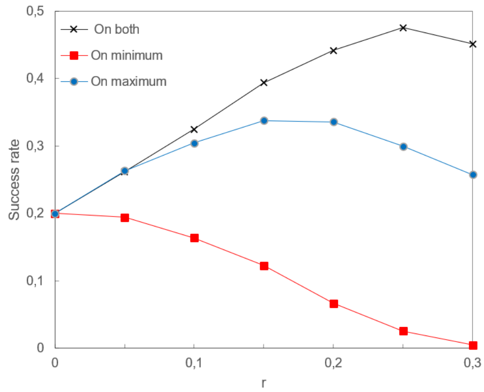

Figure 3.7: Success rate of introducing restrictions on the initial demands. My implementation's version.

Table 3.5: Success rates under different side deals. Negotiations with 8 players, 2 of which constitute 25% of the global BAU. Original paper's version.

Scenario	Success Rate
Baseline (no prior reductions)	0.2918
10% prior reduction from two largest	0.3921
10% prior reduction from six smallest	0.4440
5% prior reduction from all	0.4063
20% prior reduction from two largest	0.5480
20% prior reduction from six smallest	0.7128
10% prior reduction from all	0.5967

3.3.3 Additional details

3.3.3.A Ending Round

The negotiations carry on for 100 rounds, when the final results are gathered and the success of the negotiations is evaluated. This number may seem random and different values for the ending round can affect its final result. With this in mind, the model was tested for ending round values of 75, 100, 120 and 150, and they all produce similar values, with a sample amplitude of under 0.5% in success rates. For this reason, all the simulations are ran with the ending round value set for 100.

Table 3.6: Success rates under different side deals. Negotiations with 8 players, 2 of which constitute 25% of the global BAU. My implementation's version.

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

3.4 Preliminary Results Analysis

Comparison between my implementation's results and the original implementation show that both versions have nearly the same results for every plot. This observation leaves me comfortable for further enhancements in the model. The next chapter details the mechanisms and changes I introduced to the model and the results obtained in this study.

4

Results and Contributions

Contents

4.1 Learning Process 27
4.2 Irrationality 27
4.3 Time-dependence and the Decrease of Target Values 28
4.4 Uncertainty in Collective Goals 30
4.5 Risk Aversion 30

The previous chapter described the model and proved that my implementation of the model is in line with the author's original results. This section will focus on presenting the new features introduced in the model and studying their effect on the negotiation's success. The main focus of these features is introducing uncertainty, be it on the player's moves or the behavior of the model.

4.1 Learning Process

In the original model, a player creates expectations on the rest of the field using their previous demands to calculate its most optimal demand. For each player, he will average all of its demands, since the first round. Here, I present two new methods of creating expectations and study their impact in the success rates.

The first process is making the decision based on the previous demand only. The main advantage of this learning process is its simplicity, meaning that agents will make decisions much faster, while losing some information. The main disadvantage with this learning process is that a prediction is always subject to oscillations in the rest of the field's demands and their own predictions as well.

The second process used for the creation of expectations combines the original version with the one described just above: it averages the most recent rounds. This learning process has the advantage of looking only at the most recent demands, meaning that on round 90, a player will not take the first or tenth demands of each player into account, which are probably outdated at that point. The main disadvantage, much like in the previous learning process, is its tendency to be affected by recent oscillations in the player's demands. In the model, this was defined as the 8 most recent demands and, when there are fewer than 8 demands, the expectation is the average of the existing ones.

Figure 4.1 presents the results obtained. The results are, in a way, predictable. While ignoring the older demands may seem like a valid and interesting method of creating expectations, it misses some important information, and is prone to unpredicted oscillations in the player's recent demands. That disadvantage vastly outweighs any advantage in both learning processes.

4.2 Irrationality

This new mechanism could also be related to the Learning Process, but given its relevance, it was given a section of its own.

Game Theory seldom studies irrational agents, as it prefers to use strict payoff maximization strategies. Empiric knowledge, however, proves that theoretical studies can sometimes pale to conform to reality due to the adoption of purely rational agents. For that reason, irrationality was introduced to the model. Each player has a probability of being irrational. Whenever choosing a demand, before gathering

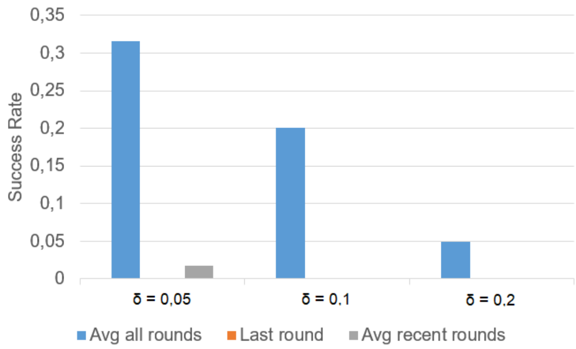

Figure 4.1: Success rate of negotiations with players using different learning processes. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

the expectations, a random number $\in$ [0, 1] is selected. If that number falls below the player's irrationality probability *p*, no expectations will be gathered and a random demand will be chosen from its strategy set ([δ, 1] unless restrictions are applied).

This experiment is not expected to increase the cooperation between the players, given that they will be dealing with erratic and unpredictable players, whose plays may or may not jeopardize the collective goal. More than that, irrational moves will change their demand average, further confusing the rest of the players. The goal of this study is, therefore, studying how badly do unpredictable players affect cooperation.

Figure 4.2 presents the results to this experiment and, surprisingly, irrationality does not have a large impact in the success rate of negotiations, as the largest oscillation is in the 5% range. This shows that, in most cases, the model can adapt to irrational players and greatly mitigate the effect of those in negotiations.

Upon a first analysis of this mechanism, one may think that higher values of *p* would lead, in the limit ($p = 1$), to a convergence on some value, independently of δ, given that the players would stop looking at factors like the target goal and agreement importance. This, however, is not the case because the demands must be chosen between δ and 1, therefore the model is dependent of δ, even when dealing with complete irrationality from the players.

4.3 Time-dependence and the Decrease of Target Values

Environmental issues are a growing concern and if nothing is done, it will be impossible to reverse the problems of greenhouse gas emissions. On the other side, if the countries emit less greenhouse causing gases, the issue becomes less pressing. The mechanism introduced here attempts to introduce some further dynamic to the model. After each round of negotiations, if the demands are not enough to satisfy the target value, the target value T is lowered by a factor λ. If, on the contrary, the demands

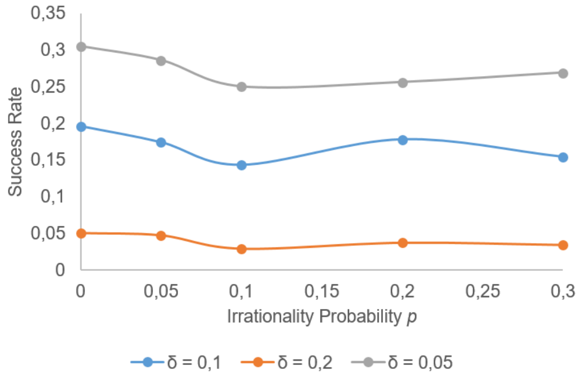

Figure 4.2: Success rate of negotiations under different values of probability *p* of demand irrationality. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

satisfy the target value, it will be raised by the same factor, while having an upper limit equal to the original target value. The expectations for this mechanism are not of increased success rates, given that the upper limit of the target value is its original value.

This mechanism aims to explain how using a dynamic target, defined by the player's demands, affects the success rate and the player's moves. Figure 4.3 presents the results of this experience. We observe that, for lowers values of λ, the success rate of the negotiations remains the same. When the devalue is increased, however, the negotiations breakdown with much more ease. This can be explained due to the existence of an upper limit to the target value. 20 consecutive successful negotiations in the beginning of the game do not make the negotiations easier, whereas a single fail will increase the needed effort to succeed. When the devalue factor increases, this problem only escalates, and the players may be faced with situations in which their BAU will provide a higher payoff than the minimum offer possible.

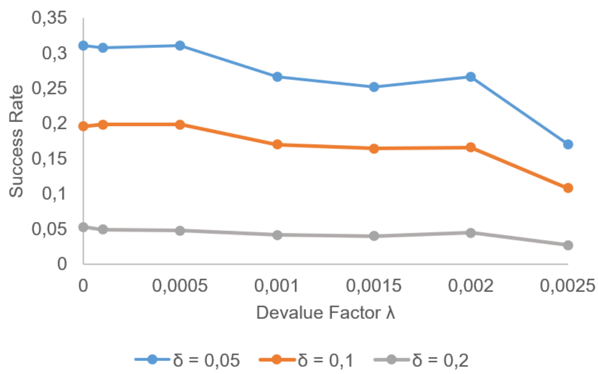

Figure 4.3: Success rate of negotiations under different values of λ. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

4.4 Uncertainty in Collective Goals

Uncertainty in this thesis is also an attempt of modeling reality in a more accurate way. In the same trend as the previous mechanism, and attempting to introduce both realism and uncertainty to the game's dynamic, a mechanism was introduced to oscillate the target value. However, unlike the previous one, in which the oscillation depends on the players' demands and their success, this one is completely random and independent on the flow of the game.

Upon game launch, the target value uncertainty ϕ is set and, for every round, the target value will T will be in the interval [T-ϕ, T+ϕ].

Figure 4.4 presents the the success rate with different values of ϕ.

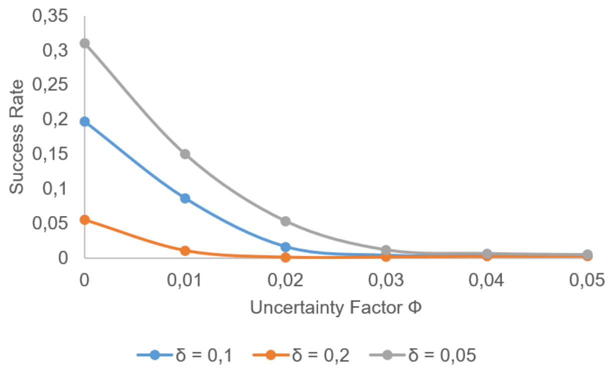

Figure 4.4: Success rate of negotiations under different values of ϕ for target value uncertainty. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

This data can be analyzed on its own, and it comes as no surprise that this uncertainty does not bring any improvements on the success of negotiations. However, this uncertainty feature has more interest when coupled with other mechanisms. The next sections will contain some examples of cross-mechanism simulations that include the target value uncertainty.

4.5 Risk Aversion

Humans are not all alike. However similar we may look or act, there are always differences between people. Some people do not like going to the casino, while others do. And from these people who go to the casino, some prefer to pay a fixed amount and keep all of their earnings, however small they may be. Others prefer to keep playing until they either have a satisfactory profit or no money at all. The group one fits in depends on his personality and his circumstances. In a lottery in which you have to bet all of your money for a chance of winning one million dollars, the poor man who has two dollars will easily take the bet, while the multimillionaire will not see this lottery as a profitable investment.

The model introduces a mechanism that simulates the existence of different people in regards to their willingness to take risks. This section introduces risk averse and risk prone people. A player will be risk averse if he considers that the expectations gathered on the other players are too optimistic and that they will demand values closer to their BAU. This situation would cause, in the eyes of this player, a global collapse and, therefore, a smaller payoff for him. For that reason, the player may concede a larger part of his BAU, hoping to achieve a successful negotiation. On the opposite, the risk prone player is an optimist one and considers that the rest of the field will actually lower their demands, creating room for him to increase his demand, giving him a larger payoff.

Risk aversion is represented by a factor ϵ, that is either added or removed after gathering the expectations. Let ad_i represent player i's average demand. Player *j*'s expectations e_j on the rest of the field will, in this case, be the following:

$$e_j = \frac{\sum_{i=1, i \neq j}^{N} (ad_i) s_i + \epsilon}{N-1} \tag{4.1}$$

The simulations were ran using 3 distinct populations. One containing only risk averse players, other with risk prone players only and other with an even distribution between the two types of players regarding risk strategy.

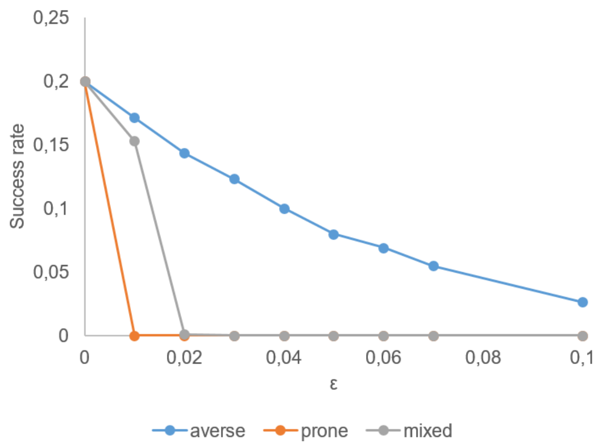

Figure 4.5: Success rate of negotiations using different strategies of risk aversion and different values for ϵ. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

A society of risk prone individuals, as can be observed in Figure 4.5 never achieves success, being the most unsuccessful group of players. Mixed populations, much like in the risk prone case, rarely achieve success, except for low values of ϵ, and even then it is not the best strategy. When dealing with risk averse populations, however, the success rate does not decay as fast with the increase of ϵ. This happens because the demands will mostly be conservative. Risk averse populations fail to succeed more frequently because their expectations can be so pessimistic that keeping their BAU may feel like a better strategy. All of the three different risk aversion strategies produce worse results than using no strategy at all.

4.5.1 Risk Aversion coupled with Collective Goals Uncertainty

While simply introducing uncertainty in the target value has some interesting properties that can be object of study, pairing it with risk aversion leads us to much more interesting conclusions. In this version of the model, the outcome is much more unpredictable. Risk prone players, for example, can now get away with higher demands if they are lucky enough to see a raise in the target value, due to its uncertainty. The opposite, however, can also happen and cause negotiations to crumble.

Figures 4.6, 4.7 and 4.8 present the results to this simulation. The first one, for a risk averse field, does not present any results beyond expectation, as it follows the same trend as Figure 4.4. When dealing with a mixed population, the results are, yet again, similar, but with lower success rates for simulations in which ϵ is different than 0.

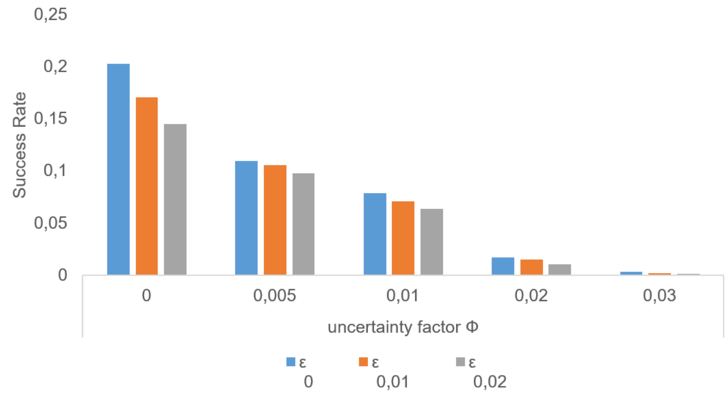

Figure 4.6: Success rates for negotiations using different values of ϵ and ϕ for risk aversion and target value uncertainty. Represents risk averse players. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

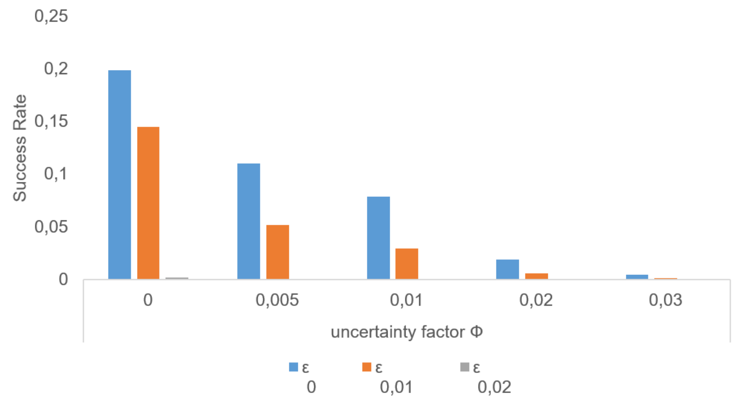

Figure 4.7: Success rates for negotiations using different values of ϵ and ϕ for risk aversion and target value uncertainty. Represents risk averse and risk prone players in the same negotiations, with an even distribution.

The results, however, get more interesting when studying a risk prone population. When there is no uncertainty, a 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 (Figure 4.8) and even outdo populations that are not prone to risk. This result is important, showing that being prone to risk may be beneficial in the presence of uncertainty in the target value.

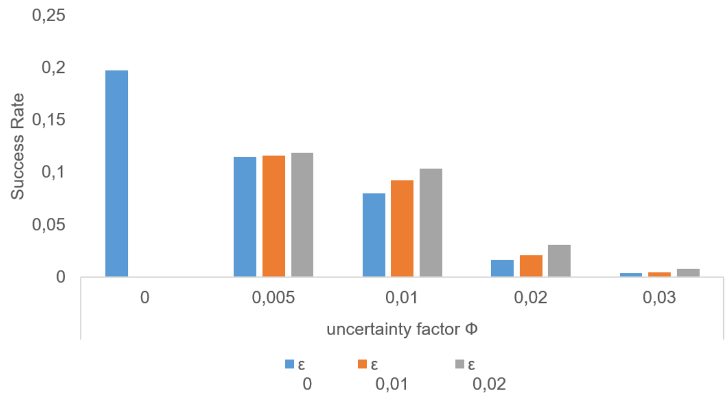

Figure 4.8: Success rates for negotiations using different values of ϵ and ϕ for risk aversion and target value uncertainty. Represents risk prone players. Simulation parameters: T=0.5, δ=0.1, #players=8 (size 1 each), all other mechanisms deactivated.

5

Conclusion

Contents

5.1 Future Work 39

The main goal of this thesis was studying different mechanisms that can influence cooperation, either positively or negatively. Besides that, cooperation is also 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. This last chapter presents an overview on the work performed, the main results and contributions obtained, as well as some final commentaries and prospects for further work.

The dissertation focused on studying how different mechanisms affect the outcome of negotiations. These results are presented in the topics below.

Influence of Number of Players

The number of players present in a negotiation is one of the most prevalent factors that affect the feasibility of negotiations. We concluded that the higher the number of players, the harder it is to succeed. This result confirms the conclusions obtained in [16].

Effect of Agreement Importance

The agreement importance is one of the cornerstones in the model. Changing it has a significant influence in the outcome of the negotiations. The more important a negotiation is considered, the lower the value value for agreement importance, because it is seen as a payoff devalue in case of failure in achieving negotiations. We concluded that lower values, i.e. more important negotiations cause an increase in the success rates. This result confirms the conclusions obtained in [16].

Effect of Heterogeneity

The model has the possibility of having players of varying sizes. This feature changes the dynamics of a negotiation, and we concluded that having players of different sizes is a positive feature, as it allows for greater chances of success. We also concluded that successful negotiations are more likely to occur in situations in which larger players give up a larger part of their emissions. These results confirms the conclusions obtained in [16].

Effect of Restricting Initial Demands

The initial demands are selected randomly, and have a significant role in the success of the negotiations. Restricting these proved to increase the success rates, as long as it is applied to the maximum value, or to both the maximum and minimum values. If only the minimum is restricted, negotiations tend to crumble. We also concluded that restricting the maximum demand or both maximum and minimum demands has a threshold after which this mechanism is counter-productive. These results confirms the conclusions obtained in [16].

Effect of Side Deals

Evaluation of the effect of side deals in negotiations lead to the conclusion that prior agreements between players promote success in negotiations. We also concluded that constraining subsets of large players is not as effective as constraining equivalent groups of smaller players. These results confirms the conclusions obtained in [16].

Effect of Different Learning Methods

The model allowed for different learning methods by the players. Learning method refers to the process each player undertakes when gathering expectations on the rest of the players. We concluded that, in comparison to analyzing the average of the most recent rounds and the most recent round, averaging the whole historic of demands provides a better chance of success in negotiations.

Effect of Irrationality in the Player's Choices

The existence of irrational players was studied, using different values of p for the probability of making an irrational move. We concluded that such moves influence the game's outcome and dynamic, but this influence is, for the most part, nearly negligible and proved that the players, as a whole, can adapt, up to a point, to erratic moves and still achieve similar rates of success.

Effect of Uncertainty in Collective Goals

A major point of study in this thesis was uncertainty. We studied how uncertainty in the target value affects the game, the players' moves, and the negotiation' outcome. The uncertainty factor ϕ was applied to the target value, and made it oscillate randomly between a set interval [T-ϕ, T+ϕ]. We concluded that this uncertainty is prejudicial to the negotiations and the success rates drop drastically with the introduction of this mechanism. We also concluded that this mechanism may allow for better success rates when combined with other mechanisms.

Effect of Time-Dependence and Decrease in Target Values

Following the trend of the previous mechanism, we also studied how oscillations affected the outcome of the negotiations. The target value's oscillation is dependent on the player's success in the previous round, resembling a sort of meritocracy. We concluded that, up to certain values of λ, the feasibility of achieving negotiations does not change. For higher values of λ, however, negotiations tend to break down.

Effect of Risk Aversion

We studied how different strategies regarding willingness to take risks affect the game's dynamic and outcome. This was achieved using risk prone and averse players. We came to the conclusion that such strategies bring bad results in regards to the negotiations, with risk averse societies producing better results than others. We also studied the effect of combining risk aversion with uncertainty in collective goals, and concluded that this compound mechanism benefits risk prone societies. Under this setting, pure risk prone populations are able to succeed in negotiations in the presence of uncertainty in the target value, and their strategy is rewarded. This phenomenon, however, does not apply for risk averse or mixed populations.

5.1 Future Work

The summary of results and contributions with this work was presented in this chapter. However large or small they may have been, there is always room for further enhancements. Here, we present work that can be built on top of this dissertation and possible shed some light in the mysteries that surround cooperation.

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.

Introduction of Punishment and/or Reputation

Some work has been done in this area, as described in Chapter 2. The introduction of more complex mechanisms such as punishment and reputation has potential to enhance our knowledge base.

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.

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.

Bibliography

[1] C. Darwin and W. F. Bynum, *The origin of species by means of natural selection: or, the preservation of favored races in the struggle for life*. AL Burt, 2009.

[2] K. Sigmund, *The calculus of selfishness*. Princeton University Press, 2010.

[3] J. F. Nash *et al.*, "Equilibrium points in n-person games," *Proceedings of the national academy of sciences*, vol. 36, no. 1, pp. 48–49, 1950.

[4] W. D. Hamilton, "Biosocial anthropology," in *Biosocial anthropology*. Malaby Press, 1975, pp. 133–155.

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

[6] W. Güth, R. Schmittberger, and B. Schwarze, "An experimental analysis of ultimatum bargaining," *Journal of economic behavior & organization*, vol. 3, no. 4, pp. 367–388, 1982.

[7] F. P. Santos, "Evolution of fairness under n-person ultimatum games," Master's thesis, Instituto Superior Técnico, 2013.

[8] F. P. Santos, F. C. Santos, A. Paiva, and J. M. Pacheco, "Evolutionary dynamics of group fairness," *Journal of theoretical biology*, vol. 378, pp. 96–102, 2015.

[9] F. P. Santos, J. M. Pacheco, A. Paiva, and F. C. Santos, "Structural power and the evolution of collective fairness in social networks," *PloS one*, vol. 12, no. 4, p. e0175687, 2017.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] S. Barrett, *Environment and statecraft: The strategy of environmental treaty-making: The strategy of environmental treaty-making*. OUP Oxford, 2003.

[16] 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.

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

[18] D. Challet and Y.-C. Zhang, "Emergence of cooperation and organization in an evolutionary game," *Physica A: Statistical Mechanics and its Applications*, vol. 246, no. 3-4, pp. 407–418, 1997.

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

[20] W. D. Hamilton, "The genetical evolution of social behaviour. ii," *Journal of theoretical biology*, vol. 7, no. 1, pp. 17–52, 1964.

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

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

[23] 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.

[24] 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.

[25] 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.

[26] M. A. Nowak and K. Sigmund, "Evolution of indirect reciprocity," 2005.

[27] O. Leimar and P. Hammerstein, "Evolution of cooperation through indirect reciprocity," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 268, no. 1468, pp. 745–753, 2001.

[28] H. Ohtsuki and Y. Iwasa, “The leading eight: social norms that can maintain cooperation by indirect reciprocity,” *Journal of Theoretical Biology*, vol. 239, no. 4, pp. 435–444, 2006.

[29] F. P. Santos, F. C. Santos, and J. M. Pacheco, “Social norms of cooperation in small-scale societies,” *PLoS computational biology*, vol. 12, no. 1, p. e1004709, 2016.

[30] P. G. Lind, L. R. da Silva, J. S. Andrade Jr, and H. J. Herrmann, “Spreading gossip in social networks,” *Physical Review E*, vol. 76, no. 3, p. 036117, 2007.

[31] 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.

[32] 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.

[33] 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.

[34] 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.

[35] 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.

[36] F. C. Santos, V. V. Vasconcelos, M. D. Santos, P. Neves, and J. M. Pacheco, “Evolutionary dynamics of climate change under collective-risk dilemmas,” *Mathematical Models and Methods in Applied Sciences*, vol. 22, no. supp01, p. 1140004, 2012.

[37] M. Milinski, R. D. Sommerfeld, H.-J. Krambeck, F. A. Reed, and J. Marotzke, “The collective-risk social dilemma and the prevention of simulated dangerous climate change,” *Proceedings of the National Academy of Sciences*, vol. 105, no. 7, pp. 2291–2294, 2008.

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

[39] M. Milinski, D. Semmann, and H.-J. Krambeck, “Reputation helps solve the ‘tragedy of the commons’,” *Nature*, vol. 415, no. 6870, pp. 424–426, 2002.

[40] 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.