



Pre-coalitions in international politics

Using software agents to simulate strategic a priori coalitions
in diplomatic negotiations and evaluate their sufficiency

Jesse Kaptein

Pre-coalitions in international politics

Using software agents to simulate strategic a priori coalitions in diplomatic negotiations and evaluate their sufficiency

Master thesis submitted to Delft University of Technology
in partial fulfillment of the requirements for the degree of

Master of Science

in Complex Systems Engineering and Management (CoSEM)
Faculty of Technology, Policy and Management by

Jesse Kaptein

To be defended in public on

May 25th, 2018

Student number:	4153758	
Proposed duration:	October 1, 2017 – March 1, 2018	
Thesis committee:	Prof. dr. ir. A. (Alexander) Verbraeck	TU Delft, chairman
	Dr. S. W. (Scott) Cunningham	TU Delft, 1st supervisor
	Dr. M.E. (Martijn) Warnier	TU Delft, 2nd supervisor
External coach:	Dr. G. (Georgios) Pierris	Deloitte

Preface

The thesis in front of you concludes my Master of Science in Complex Systems Engineering and Management with the specialization track Information Architecture at the University of Technology in Delft. Although it had its highs and lows, the majority of time I really enjoyed completing this final graduation project. Hereby, I would like to thank my graduation committee – Alexander Verbraeck, Scott Cunningham and Martijn Warnier – for the time and effort they have put in guiding me through this final project. In particular, I thank Scott Cunningham for introducing me to this subject, the many interesting discussions we had and the freedom he gave me to create my own research project.

In addition, I want to thank all the wonderful colleagues within the Analytics and Information Management service line in Deloitte for their input and the amazing experience I had. More specifically, I would like to express my gratitude towards Stefan van Duin, giving me this great opportunity. Moreover, I thank Georgios Pierris for his contributions, wisdom and all the fantastic, slightly strange and off-topic conversations we had. Of course, finally, I would like to express my gratitude towards my parents, supporting me for the past six years.

*Jesse Kaptein
Delft, May 2018*

Executive summary

In everyday life, negotiations take place anywhere, anytime. They are also predominantly present in international politics: negotiations among governments in various intergovernmental organizations. Crucial in these negotiations is the formation of coalitions, both before and during the actual negotiations. While negotiating, these coalitions are often the result of ongoing discussions, whereas *ex ante*, they can be seen as strategic choices of negotiators. Political scientists provided the *balance-of-power* theory, rationalizing this choice in terms of "balancing" – allying against more powerful states to avert its hegemony – and "bandwagoning" – aligning with more powerful countries to harvest some of its gains [157, 175, 177]. Unfortunately, until now, they have only been applied in situations of immediate threat. Additionally, multilateral bargaining models have been developed predicting winning coalitions under strict assumptions, suppressing many relevant variables and neglecting strategic choices in the pre-negotiation phase. Automated negotiation might bring both fields together. However, currently, they lack automated environments supporting both multilateral multi-issue negotiation using alternating offers and the strategic choice of forming coalitions *ex ante*. Therefore, the objective of this study is to develop an automated negotiation environment supporting complex political contexts, able to generate pre-coalition arrangements and evaluate their sufficiency. By doing so, the following research question is answered:

Research question

When are pre-coalitions sufficient in international political negotiations given a state's own set of objectives and specific power position?

In contrary to most bargaining models, our automated negotiation environment is build upon three literature reviews in which we focus on key elements of actual diplomatic negotiations; define strategic actions actors have within such negotiations and; determine the behavior underlying these actions and decisions. We implemented the final negotiation environment in a Python package named "negopol", which is essentially a two-stage bargaining model in which agents start to form pre-coalitions by 'bargaining' over ideological values and power positions, after which multiple rounds of actual negotiations are held.

In this thesis, we have defined sufficient pre-coalitions as coalitions established prior to actual negotiation phase and substantially improving the quality of outcome of international political negotiations. In turn, the quality of outcome depends on how well states meet their own set of objectives. In this thesis, self-interested negotiators are assumed to primarily aim for maximizing (absolute) personal gains or utility. However, there are also one or more secondary objectives, including to maximize social welfare; minimize time to agreement; maximize political support; or maximize the likelihood of agreement. Without political support, for instance, ratification and implementation of outcomes is often impossible as "(...) *agreement may be reached, but the loser might end up wanting to sabotage its implementation.*" [108, p. 27-28].

Additionally, international political negotiations themselves can be decomposed in many components, ranging from the regimes structuring the negotiation process, voting procedures they impose and the actors they include. Following automated negotiation linguistics, this can be abstracted in both the negotiation domain (i.e. what is discussed?) and the negotiation setting (i.e. with whom/how to negotiate?). As international political negotiations vary widely in domains and settings, we derived a hypothetical context from the balance of power theory assuming states have three types of power positions – i.e. weak, secondary and strong – in which weak states are significantly less powerful than strong states, and, as introduced by Waltz, secondary states are: "(...) *major powers, rather than great powers*" [177, p. 44].

In this context, simulations have shown weak states benefit most from forming pre-coalitions as on average all their objectives increase. In addition, forming pre-coalitions is also sufficient for secondary states, although their personal gains improve less. In contrary, strong states mainly benefit by increased political support and time objectives, whereas their median personal gains actually decrease. Though, the variety of simulation results indicate in some negotiation environments – i.e. domains and settings – pre-coalitions

may become (more) sufficient. Especially strong, and to a lesser extent secondary states should be extremely aware at what time they make concessions during the negotiation process [108, p. 59-63]. So, if they benefit from the actual negotiation phase, they should perhaps not make concessions by forming pre-coalitions beforehand. And when they do, they have the resources to be more picky [177, p. 12-13] and follow the best pre-coalition strategy. Obviously, this brings us to answering the main research question.

Answering the main research question

To answer this question, we distinguish between the negotiation environment in which pre-coalitions may become (more) sufficient and what strategies states can follow to increase its sufficiency. In this thesis, we assume these strategies to be either balancing, bandwagoning, a combination of both or staying alone.

First, for the **environment**, simulations have shown many factors affect the sufficiency. For all types of states, pre-coalitions become more sufficient if the magnitude of negotiations increases and, hence, more subjects are discussed. Overall sufficiency increases as well when the number of states at the negotiation table increases. Both are in line with most political science theory [65, 175, 177]. More interestingly, simulations also show negotiating behavior of other states is increasingly important for the overall sufficiency of pre-coalitions for both secondary and strong states. Especially in environments in which other states tend to negotiate irrational and risky, pre-coalitions may increase the quality of outcome significantly by increasing secondary and strong states' political support objectives. Particularly for strong states, pre-coalitions are sufficient if there are (many) other strong states to negotiate with. As such settings are often more competitive [42, p. 568-575], pre-coalitions can substantially increase the personal gains of a single strong state. Additionally, if other states have good outside options – and, therefore, are less willing to negotiate – pre-coalitions become more sufficient as well. In such settings, a priori coalitions might break the deadlock by quickly becoming minimal winning.

In many other negotiation environments also weak states, and, to a lesser extent, secondary states, may take the 'gamble' by forming pre-coalitions. This, as the majority of simulation experiments have shown there is a substantial chance that pre-coalitions become sufficient and, thus, increase the overall quality of outcome. On the contrary, strong states should be cautious forming pre-coalitions, as there is a fair chance they benefit more from the actual negotiations. Unless, they opt for specific **strategies** rationalizing the pre-coalition choice in a certain negotiation environment. For strong states, these include:

- Whenever other states have on average poor outside options, strong states should *never balance* with other strong states. Instead, they must *bandwagon* with weak or secondary states for pre-coalitions to become sufficient, especially in terms of personal gains;
- Whenever there are many secondary states at the negotiation table, strong states should *only partly bandwagon* with weak states to create pre-coalitions sufficiently increasing personal gains.

The latter leaves room for *balancing* with other strong states, hence a mixed strategy. Thereby, it partly supports the scarce research on 'great-power balancing' (e.g. [93] and [166]). In addition also **weak agents** may prefer distinct strategies:

- With many states at the negotiation table, overall it becomes beneficial for weak states to almost *fully balance* with other weak or *mainly balance* with secondary states;
- With many secondary states at the negotiation table, pre-coalitions become sufficient if weak states *partly bandwagon* with stronger states.

Finally, we have also found simulation evidence for two strategies **secondary agents** should follow:

- With many subjects to discuss at the negotiation table, pre-coalitions become sufficient for secondary states when *mainly balancing* with other secondary states and *partly balancing* with weak states;
- When states appear to be more risk taking and secondary states aim for large political support, they should *mainly balance* with weak or secondary states for pre-coalitions to become sufficient.


Table of Contents

Preface	iii
Executive summary	v
1 Introduction & Background	1
1.1 Problem exploration	1
1.1.1 Bargaining and coalition formation in literature	2
1.1.2 Knowledge gaps and problem statement: what's the problem?	3
1.2 Problem specification	4
1.2.1 Objective: generate pre-coalition arrangements and evaluate sufficiency	4
1.2.2 Scope: coalitions in multilateral multi-issue non-cooperative negotiations	5
1.2.3 Research questions: what to research?	5
1.2.4 Related literature and studies on ex ante coalition partners	6
1.3 Thesis outline	7
1.3.1 Scope, strategic decisions and behavior	7
1.3.2 The two-stage model	8
1.3.3 Experiments and simulation results	8
1.3.4 Putting the pieces together	8
2 Characteristics of International Political Negotiations	9
2.1 Methodology: systematic literature review	9
2.2 Architecture of political international negotiations	9
2.2.1 Agreements: negotiation in political realm	10
2.2.2 Actors: who conduct the negotiations?	10
2.2.3 Issues: what are we talking about?	11
2.2.4 Regimes: structure through intergovernmental organizations	12
2.2.5 Process: is everything a bargaining game?	12
2.2.6 Power: inequality in international negotiations	13
2.2.7 Strategies: defining negotiating behavior	14
2.2.8 Time: macro and micro impact	14
2.3 Discussion: is this it?	15
2.4 Conclusion & summary	15
3 Coalition Formation in International Politics	17
3.1 Pre-negotiation coalitions: why are states working together?	17
3.1.1 To face threatening situations	17
3.1.2 To become more decisive	18
3.1.3 To find like-minded states	19
3.2 Coalitions while negotiating: why do states remain together?	19
3.2.1 Rationality and the core: fundamental game-theoretic principles	20
3.2.2 Riker's minimal winning coalitions	21

3.3	Conclusion & summary	22
4	Fundamentals of Automated Negotiation	23
4.1	Protocol: rules governing the negotiation process	23
4.1.1	Negotiation among multiple parties	23
4.1.2	Negotiation over multiple issues	24
4.1.3	Time and deadlines	25
4.1.4	Selection function	25
4.2	Scenario: outcome space, preferences and utility	25
4.3	Strategy: bid to negotiate	26
4.3.1	Tactics: different actions to achieve your goals	27
4.3.2	Negotiate over issues with discrete domains	28
4.3.3	Acceptance strategy	29
4.4	Conclusion & summary	29
5	A Two-Stage Automated Coalition Formation Model	31
5.1	Methodology: why use a simulation model?	31
5.2	Requirements: do we need to model everything?	32
5.3	Core model logic: how does it work?	33
5.3.1	From ideological proximity models to pre-coalitions	34
5.3.2	Selecting agents and coalitions	34
5.3.3	Making offers in coalitions with divergent strategies and tactics	36
5.3.4	From utility functions to accepting offers and leaving coalitions	38
5.4	Formal model definition and initialization	40
5.5	Implementing the formal model	41
5.5.1	Going the Pythonic way	42
5.5.2	Tool structure: what have we build?	42
5.5.3	Look-and-feel of the tool: an example.	43
5.6	Code quality and verification	45
5.7	Conclusion & summary	46
6	Negotiation Objectives & Experimental Setup	47
6.1	Defining sufficiency: when are pre-coalitions good enough?	47
6.1.1	Maximum social welfare: utilitarian, egalitarian and Nash product	47
6.1.2	Maximum self-interestedness: only personal gains?	48
6.1.3	Conclusion: sufficient pre-coalitions and quality of outcome	49
6.2	Experimental design	49
6.2.1	Settings of negotiation environment	49
6.2.2	Domains in negotiation environment	50
6.2.3	Decisions: composition of pre-coalitions	51
6.2.4	Relations between environment, decisions and quality of outcome	51
6.3	Discussion: is this the right experimental design?	51
6.4	Conclusion & summary	52
7	Simulation Results	55
7.1	Raw output: before quality measures are calculated.	55
7.2	Forming pre-coalitions or staying alone?	56

7.3	Sufficiency of pre-coalitions	58
7.3.1	In different negotiation environments	60
7.3.2	By changing the composition	66
7.4	Conclusion & summary	70
8	Discussion of Results	73
8.1	Fitted Random Forest models	73
8.2	Overall sufficiency of pre-coalitions	73
8.3	Sufficiency in different negotiation environments	74
8.4	Sufficiency of balancing or bandwagoning strategies	76
8.5	Model validation: is the right model constructed?	79
8.5.1	Two-stage model.	79
8.5.2	Experimental design	80
8.5.3	Taking the next steps in validity.	80
8.6	Conclusion & summary	80
9	Conclusion & Discussion	83
9.1	Short recap	83
9.2	Answering the main research question.	84
9.2.1	The balance-of-power theory: a reflection	85
9.3	Scientific contributions	85
9.4	Social relevance: why should we bother?	86
9.4.1	Who should use the two-stage simulation model?	86
9.4.2	Link to master program	87
9.5	Discussion: nobody's perfect	87
9.5.1	Limitations of this thesis.	87
9.5.2	Future research directions	88
	Bibliography	91
A	Search Terms in Systematic Literature Review	101
A.1	Characteristics of International Political Negotiations	101
A.2	Coalition Formation in International Politics	102
B	Issue Classification	105
C	Overview of Assumptions	109
D	Model Requirements	117
E	Class Diagram of the Two-Stage Model	123
F	Code Testing and Verification	125
F.1	Testing tool quality: is it any good?	125
F.1.1	Functionality: assessment of requirements.	126
F.1.2	Examining non-functional requirements	129
F.1.3	Assessment of non-functional requirements	134
F.1.4	Conclusion of tool quality	136
F.2	Verification: have we computed the correct model?	136
F.2.1	Examining inputs and outputs.	137

F.2.2	Theoretical prediction: simple three-agent model147
F.2.3	Pushing the tool to the limit: simple three-agent model148
F.2.4	Theoretical prediction: multi-agent model.148
F.2.5	Pushing the tool to the limit: multi-agent model151
F.2.6	Testing tool variability154
G	Experimental design	157
G.1	Factors in the experimental design157
G.2	Validity of the experimental design158
G.2.1	Negotiation domain158
G.2.2	Negotiation setting159
H	Simulation results	163
H.1	Perspective of weak agents163
H.1.1	Correlation between factors163
H.1.2	Fitness of Random Forest models166
H.1.3	Partial dependency plots.167
H.1.4	Partial interactions plots.171
H.2	Perspective of secondary agents176
H.2.1	Correlation between factors176
H.2.2	Fitness of Random Forest models178
H.2.3	Partial dependency plots.179
H.2.4	Partial interactions plots.184
H.3	Perspective of strong agents.189
H.3.1	Correlation between factors189
H.3.2	Fitness of Random Forest models191
H.3.3	Partial dependency plots.192
H.3.4	Partial interactions plots.196



Introduction & Background

I'm gonna make him an offer he can't refuse.
Vito Corleone, in *The Godfather* (1972)

In everyday life, negotiations take place anywhere, anytime. Whether couples negotiating over who will do the domestic chores or labor unions demanding for more wages, almost always situations arise in which parties with opposing interests need to make a joint decision. Negotiations are also predominantly present in international politics: negotiations among governments in various intergovernmental organizations, such as the EU. Political scientists describe many strategies governments adopt while negotiating (e.g. [12, 42, 119, 184]). Among which allying with others – and thus forming coalitions – is said to be crucial [41, 74, 117, 136].

A simple example covering coalition formation include three friends going on holiday and bargaining over the *destination* and means of *transport*. Alex and Bob want to go by plane to Greece and Turkey respectively, while Carl wants to go to Greece by car. Because Carl finds the means of transport less important than the destination, he decides to support Alex by forming a coalition to go to Greece. Thus, he makes a concession by going by plane in favor of his more preferred holiday destination. Having a worse bargaining position, Bob feels the "pressure" to join Alex and Carl's coalition and, therefore, this bargaining game most likely results in going to Greece by plane. Although this example is way too simple for representing political negotiations, it raises some important questions. For example, if Bob had approached Carl first, would the outcome of the negotiation be different? Did Carl's decision to join Alex result in the most sufficient coalition for Carl? Or was there another way to get what he preferred most? This thesis provides possible answers by simulating international political negotiations using intelligent software agents able to form coalitions both before and during actual negotiations.

Before doing so, this chapter first explores the problem in order to delineate knowledge gaps which, in return, are used to create a problem statement. Accordingly, this statement is translated to a research objective, main research questions and scope. We conclude by presenting related work and providing a detailed thesis outline.

1.1. Problem exploration

So, what is actually meant with coalitions in international political negotiations? To start with, the term *negotiation* has a variety of definitions in academic literature. In the field of automated negotiation, it is often defined as the way "(...) to reach an agreement, and in particular, agreement in the presence of conflicting goals and preferences" [49, p. 1]. In political sciences, a predominant definition is defined by Odell. He states that negotiations refer to a "sequence of actions in which (...) parties address demands and proposals to each other, whatever particular steps it may include in a given case" [118, p. 10]. An important definition stemming from bargaining theory is that negotiation is "any process through which the players on their own try to reach an agreement" [113, p. 2]. Because in this thesis all aforementioned perspectives are used, the following combined definition is used:

Negotiation is a sequence of actions in which two or more parties, each with conflicting goals and preferences, explore their options and try to find mutually acceptable outcomes.

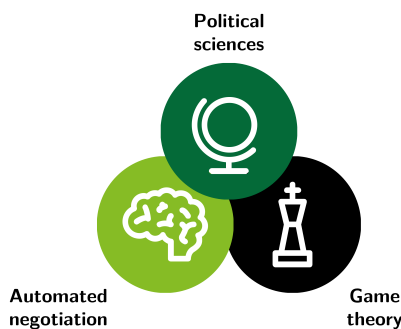


Figure 1.1: Three different research perspectives this study takes.

Note that because negotiation and bargaining often refers to each other, these terms are used interchangeably. The term *political negotiations* then refers to negotiations among individuals who act in institutions on behalf of others [101, p. 1] and, thus, represent a larger group of individuals with a (more or less) homogeneous set of goals and preferences. These institutions do not necessarily have to be political parties, but can also be companies, interest groups, labor unions or lobbyists. In our context, *international negotiations* then relates to negotiations among governments [74, p. 266] and, therefore we define:

International political negotiations as negotiations among governments in which governments are represented by individuals who act in institutions on behalf of their national government.

Note that we sometimes also refer to them as 'diplomatic negotiations' (see also [108]). These types of negotiations often take place in intergovernmental organizations addressing many different issues. Examples of such organizations include the United Nations and European Union. States participating in such organizations are often addressed as *member states*. Therefore, we define a:

Coalition as a subset of member states that work together in a partnership to achieve mutual beneficial goals [182], often prior to actual negotiations themselves.

It is important to note that coalitions are not subsets of member states having identical goals and preferences. So, most likely before a partnership is established various concessions have been made already. We have observed this phenomenon in the introduction of this chapter, in which Carl makes a concession by going by plane, prior to the negotiation with Bob.

1.1.1. Bargaining and coalition formation in literature

We now introduce three different fields of research that have been analyzing coalition formation, namely i) *game theory*; ii) *political sciences*^[1] and iii) *automated negotiations* (see Figure 1.1). In general, **game theory** on coalition formation proceed along two broad lines of reasoning [142, p. 14-15]. First, in cooperative game theory the unit of analysis is the group, rather than individuals. Although cooperative game-theoretic models still include individual preferences, they describe the joint capabilities of a group, rather than individual actions of its members. Moreover, classical cooperative literature generally takes the formation of large or largest possible (grand) coalition for granted [67, p. 67]. Especially in political negotiations, one can imagine this is not always possible. As Aumann and Dreze (1974, [9, p. 233]) stated: "(...) the very act of 'acting together' may be difficult, costly, or illegal, or the players may, for various 'personal' reasons, not wish to do so".. Therefore, there has also been a non-cooperative approach to coalition formation, in which [86, p. 9]:

"(...) the unit of analysis is the individual participant in the game who is concerned with doing as well for himself as possible (...). If individuals happen to undertake behavior that in common parlance would be labeled as 'cooperation' (...) then this is done because such cooperative behavior is in the best interests of each individual singly."

This leads to a bargaining approach of coalition formation, described by a process in which proposers alternately try to form coalitions by making offers. In return, receivers need to accept or reject the proposal,

^[1]Political sciences is used as an umbrella term for numerous (sub)fields of study such as, among others, international relations, political theory, public administration and public policy.

which eventually may lead a coalition. Many, mainly bilateral bargaining models have been proposed using this concept of "alternating offers". Arguably, the most influential model was introduced by Rubinstein [148] and extended many times thereafter (e.g. [126, 141, 169]). All these models provide clear outcomes of which coalitions rationally should form and how to divide the winnings.

However, at a cost, as most of these models suppress many factors "(...) *in order to produce a tractable problem that can be handled mathematically.*" [26, p. 67]. We argue that these simplifications have become to severe and, among others, negatively influences practical use. More specifically, these models often assume actors act rationally, have complete information, large cognitive capabilities, and follow strict rules [84, p. 153]. Moreover, game theory is mainly focused on equilibrium outcomes in infinite time horizons, neglecting the process of negotiation itself; strategic decisions made before or during this process; and, occasionally, even (irrational) mistakes. This, in sharp contrast to real negotiations, which are characterized as unruly processes with "capricious", "unpredictable" and "chaotic" behavior [35, 36, 79].

More descriptive theories of coalition formation are provided by **political sciences**. Negotiations are analyzed through case studies and coalitions are explained as strategic choices [173, p. 6-7], among others caused by different power positions governments have. For example, Dür and Mateo describe the formation of defensive coalitions aimed at blocking compromising proposals of more powerful opponents [42, p. 562-570]. In another case study, Karayanidi explains that it might be beneficial to form alliances between countries with weaker power positions in order to not become "policy takers"^[2] [74, p. 268-269]. However, she also notes that Drahos has shown the opposite of this statement is also true, and thus that these alliances can also lead to worsened bargaining power [41, p. 81-87]. This is precisely why these descriptive theories might explain specific use cases effectively, but applying them in other negotiation settings become difficult [84, p. 153]. Finally, stemming from international relations scholars, more generalizing theories describing the importance of "balancing" – allying *against* more powerful countries to avert its hegemony – and "bandwagoning" – allying *with* more powerful countries to harvest some of its gains – have been discussed extensively in threatening situations (e.g. [157, 175, 177]). Unfortunately, it is not known to what extent these two concepts explain coalition formation in (less threatening) international political negotiations.

Finally, emerged from artificial intelligence, **automated negotiation** is dedicated to design software agents (ultimately) negotiate on our behalf [49, p. 1-2]. Being more process-oriented, these software agents are build upon many game-theoretic concepts, but negotiate by making alternating offers and, therewith, (re)act upon and learn from each other. Automated negotiation is only recently dealing with the formation of coalitions. Majority of work has been focusing on social optimal (welfare) solutions [80, 150, 158, 159]. Moreover, while bilateral negotiations have been studied within automated negotiation for years, *multilateral* environments are relatively new. It is more challenging to maintain the same useful results, while adjusting multilateral negotiations to real-world situations [84, p. 166]. Thus, the rules in multilateral negotiation that guide the interaction amongst software agents – also known as the *negotiation protocol* – become extremely domain dependent.

1.1.2. Knowledge gaps and problem statement: what's the problem?

Different problems arise when analyzing the strategic decisions of forming coalitions in international negotiations. Many political scientists have discussed the choice of allying with other states by trying to find rational explanations of these choices in various case studies. As these theories are extremely domain dependent, they can hardly be adopted in different negotiation settings [84, p. 153]. The more generalizing balance-of-power theory rationalizes the strategic choice for coalition partners, but it is normally applied in threatening – near war – situations. Plainly, international political negotiations are not conducted in conflict situations, at least not the vast majority of the times. But also in a more negotiation setting, choosing coalition partners is an essential strategy for governments to reach their goals.

As international negotiations often cover long periods of time, an important question raises [82, p. 547]: are political coalitions in international negotiations sufficient over time? The notion of sufficiency (or "stability") has received major attention in most bargaining models by finding equilibrium solutions. However, analyzing these solutions suppresses many relevant factors and, therefore, not captures the "capricious", "unpredictable" and "chaotic" behavior of negotiations. As finding equilibrium solutions has become common practice, Bolton and Corson (2012, [23, p. 85-86]) state there is a sincere gap between theoretical literature and empiric and experimental research. Either data needs to be gathered and/or simulated with the theory

^[2]In contrary to a "policy shapers", a policy taker is a party that is (almost) forced to except a proposed policy [77, p. 65].

in mind. Finally, as these game-theoretic models mainly focus on predicting results of negotiations – both in terms of formed coalitions and negotiation outcomes – they neglect strategic choices states make *ex ante*.

Intelligent agents guided by game-theoretic concepts and negotiating in dedicated environments could be used to represent international political negotiations better. However, only little progress has been made in modeling multilateral multi-issue negotiations in the field of automated negotiation [88, p. 27-28]. Moreover, academic literature has only come up with few protocols supporting the formation of coalitions [84, p. 166] which are not applicable to international political negotiations. Neither do those environments consider the strategic action of choosing coalition partners before the actual negotiations take place.

As a result, international political negotiations cannot be adequately represented yet, which hinders deeper systematic understanding of coalition formation in this context. Therefore, we need to combine the knowledge of game theory, political sciences and automated negotiation together. This enables us to create negotiation environments in which autonomous agents can strategically choose pre-coalition partners, negotiate together, make concessions and trying to form the winning coalition while negotiating. Hence, simulating key aspects of international political negotiations. In this environment, *ex ante* coalitional arrangements can be generated computationally and evaluated respectively. Though, current automated negotiation environments do not support coalition formation in the more complex international political context. They do, for example, not incorporate intrinsic differences^[3] among negotiating countries, only partly include contrasting power positions and not at all represent the messy process of political decision making [49, p. 97-120]. Together, this leads to the following problem statement:

Problem

Game-theoretic models predict coalitions under strict assumptions, suppressing many relevant variables and neglecting strategic decisions in the negotiation process. Theories of allying with other states are described in political sciences extensively, but adopting them in other settings is cumbersome. Automated negotiation might bring both fields together by simulating pre-coalitional arrangements and evaluating their sufficiency. However, we lack automated negotiation environments supporting (more) complex political contexts and allowing to strategically choose coalition partners *ex ante*.

In the above statement, the "complex political context" refers to a multilateral multi-issue negotiation setting, in which coalition formation is both part of the process as well as a strategic choice in advance.

1.2. Problem specification

The following sections describe the thesis objective, scope and introduces the main research questions. Moreover, we briefly discuss some related literature and studies.

1.2.1. Objective: generate pre-coalition arrangements and evaluate sufficiency

The results of exploring the problem and identifying knowledge gaps as outlined above have motivated the following concrete research objective:

Research objective

The objective of this thesis is to *develop* an automated negotiation environment supporting complex political contexts, able to *generate* pre-coalition arrangements and *evaluate* their sufficiency.

Thus, the objective consists of two distinct aspects, namely *creating* an automated negotiation tool and *using* it to generate and evaluate pre-coalition arrangements. The first takes a more design point of view, whereas the second an analysis point of view. Developing the tool is not an objective on its own but merely a means to perform the analyses. Therefore, this tool i) must allow us to answer the research question(s) in this thesis; ii) it should allow others to review our work and iii) it should allow others to extend our work.

^[3]Such as economical, sociodemographic and cultural differences.

1.2.2. Scope: coalitions in multilateral multi-issue non-cooperative negotiations

As international negotiations involve multiple parties and issues to negotiate over, this research solely focuses on coalition formation in multilateral multi-issue negotiation settings. Moreover, as automated negotiation assumes agents – and thus the governments they represent – are self-interest by nature, we mainly use non-cooperative game theory to describe agents' bidding, accepting and coalition strategies. However, this thesis also recognizes that as negotiations are conducted at the highest diplomatic levels, agents are well aware of their surroundings. Therefore, strategic demands and maximum personal gains are also (partly) viewed from a cooperative perspective. Moreover, other game-theoretic concepts such as inside- or outside options^[4] and exogenous risk are left out of scope. Moreover, the only strategic action agents have prior to actual negotiations is choosing their pre-coalition partners. Thus, strategies such as changing bringing more subjects to the negotiation table, change the order of the negotiation agenda or introducing new players are out of scope. Thus, in short:

Scope

This thesis focuses on coalition formation in *multilateral multi-issue* negotiations using *self-interested* agents with *incomplete information* and the *alternating offers protocol*. The most relevant strategic choice these agents have is selecting *ex ante coalition partners*.

This does not imply, however, agents do not bluff, act aggressively, or always tell the truth. Of course, similar to actual negotiations, they have limited information and have to learn opponent's preferences and strategies. In addition, we focus on developing a tool that can easily be used and extended by others. However, we will not focus on user interfaces, user interactions and test the tool with potential users as it would require the development of an extensive interaction design study which is out of scope of this research.

1.2.3. Research questions: what to research?

Based upon the previous sections the following central research question is formulated:

Research question

When are pre-coalitions sufficient in international political negotiations given a state's own set of objectives and specific power position?

To answer this main research question the following supporting questions are formulated:













Supporting questions

- 1) What characterizes international political negotiations?
- 2) Which theories have been developed to explain coalition formation in (international) political negotiations?
- 3) What are the main components of automated negotiation environments?
- 4) How can the formation of pre-coalitions in international political negotiations be modeled using automated negotiation?
- 5) Which objectives make pre-coalitions sufficient?
- 6) How do pre-coalitions influence state' set of objectives in different negotiation environments?

We have already introduced the perspectives this study uses. An oversight for each supporting question is given in Table 1.1, accompanied by the way questions are answered and means of presentation.

^[4]Except that the coalition formation process itself lead to dynamic or changing outside options over time, which are in scope of this research. See [113, p. 99-135] for more information on outside options.

Table 1.1: Answer specification of supporting questions

Question	Answered by	Means of presentation	Perspectives
1)	Literature review, decomposition of characteristics	Text, visuals	
2)	Literature review, comparison of theories	Text, visuals	 
3)	Literature review, breakdown of components	Text, visuals	
4)	Requirements specification, formal model definition, pseudo algorithms, example run	Text, visuals, model	  
5)	Decomposition of objectives, experimental design	Text, visuals	 
6)	Simulation results	Text, plots, statistics	  

1.2.4. Related literature and studies on ex ante coalition partners

Related literature on strategically choosing coalition partners ex ante has thoroughly been described in political sciences. Though, in this thesis, we mainly focus on the "balance-of-power theory" and its "bandwagoning" extension as they are more general theories of alliance formation. Both concepts – introduced by Kenneth Waltz (1979, [177]) and extended by Walt (1985, [174]) – motivate the strategic choice of forming coalitions in conflict situations. They often use different power positions of states as their main line of argument, but, as we will see in Section 3.1.1, this view is slightly flawed. However, as it is still the theory's major component, this thesis considers three distinct power positions, namely those of weak, secondary and strong states^[5]. As introduced by Waltz, secondary states are: "(...) *major powers, rather than great powers*" [177, p. 44]. So, they are certainly not as powerful as strong powers but definitely hold more power over weak states.

Of course, alternatives to the power-of-balance theory have also been considered in political science literature. The theory of "soft-balancing" introduced by Pape (2005, [130]) describes the decade after the collapse of the Soviet Union in which the US became the sole strong power. In this period it is observed that – contrary to what the balance-of-power theory prescribes – secondary states almost never balanced against the US. Instead, secondary states use international institutions to oppose the US's power [130, p. 9-10]^[6]. The limitation of this theory is that its viewed from a unipolar strong power, which becomes cumbersome analyzing international political negotiations in, for example, the European Union. Another line of reasoning is provided by Leffler (1992, [92]) in his book "A Preponderance of Power: National Security, the Truman Administration, and the Cold War". He describes US's greatest concern in the immediate postwar period that industrial resources and military power of Western Europe might fall the Soviet Union [92, p. 10-15]. Therefore, they developed the "strategy of preponderance" which implied creating a strong dependency on the US forcing Western-Europe to bandwagon with them. Though, one can argue whether the strategy of preponderance can be seen as an alternative to the balance-of-power theory rationalizing alliance formation, or an economic and/or military strategy opposing the Soviet Union.

Compared to the well-developed literature in political sciences, there is little work on ex ante coalition partners in game theory. There are some models analyzing a slightly different setting, as they have been studying pre-electoral coalitions. For example, in [111] the author studies the formation of new heterogeneous parties from multiple groups of homogeneous parties. This, to examine new parties' influence on voting behavior and, thus, the outcome of post-electoral negotiations. Persson, Roland and Tabellini (2007, [133]) study a similar model, but also empirically show government spending decreases when pre-coalitions do establish. Moreover, Bandyopadhyay, Chatterjee and Sjöström (2011, [13]) have created a negotiation model showing the influence of pre-coalitions on post-electoral negotiations. They also name some important parameters affecting the pre-electoral strategic choice, including the trade-off between more powerful pre-coalitions and accepting some ideological distant parties.

Finally, to our best knowledge, there are no automated negotiation environments developed supporting the strategic decision of choosing coalition partners ex ante. In fact, they not even support coalition formation in general. There are environments supporting various negotiation domains and protocols. For example, the

^[5]Which are at the same time the three perspectives through which the main research question is answered.

^[6]The author describes, among others, how these features have helped the international opposition to the US war against Iraq.

state-of-the-art GENIUS^[7] negotiation environment is an open architecture allowing researchers to design and evaluate agent strategies [97, p. 51-53]. As it is originally developed for bilateral negotiation strategies, the vast majority of implemented protocols is only suitable for negotiations with two agents. There have been other (less established) environments proposed supporting multilateral and/or multi-issue negotiations, including CloudNeg [128] and RECON [7]. Yet, these environments both strongly dependent on the e-commerce domain they are applied in. Moreover, they do not support multi-issue negotiations [7, p. 163].

1.3. Thesis outline

A detailed overview of this thesis is given in the subsections below, including the way the chapters relate to each other. For those readers in a hurry, we recommend to take the following 6 essential steps:

Quick read guide

1. Skip Chapter 2 entirely but read the conclusion and summary in Section 2.4;
2. Read Section 3.1.1 and Section 3.3 of Chapter 3 and the conclusion of Chapter 4 in Section 4.4;
3. Focus on reading Chapter 5 entirely;
4. Read the introduction of Section 6.1, Section 6.1.3 and Section 6.2 completely;
5. Skip Chapter 7 but read discussion of results and conclusions in Chapter 8;
6. Read the answer to the main research question in Section 9.2 followed by our two main scientific contributions in the second paragraph of Section 9.4.

1.3.1. Scope, strategic decisions and behavior

Current multilateral bargaining models face a sincere gap between theoretical literature and empiric and experimental research, as stated by Bolton and Corson (2012, [23, p. 85-86]). This is due to the fact that they are evolved from bilateral models, rather than real-life negotiations. Therefore, in Chapter 2 contours of the multilateral bargaining model are sketched by performing a systematic literature review on most essential elements – or "architecture" – of international political negotiations. Subsequently, assumptions are derived *scoping* the model;

In Chapter 3 we explore existing political and game-theoretic theories on coalition formation. By doing so, we provide strategic options and (rational) motivations for forming pre-coalitions in the first place – i.e. why are states working together? Furthermore, we link and adjust a fundamental cooperative game theoretic principle (the *core*) to different actions negotiators have when being or becoming part of a coalition – i.e. why do states remain together? By doing so, important assumptions are derived regarding the *strategic decisions and actions* negotiators have.

Finally, in Chapter 4 we lay the foundations of automated negotiation. This chapter explains how negotiators interact, who is allowed to make offers, how to evaluate these offers and which strategies can be used to negotiate. We also introduce some limitations agents have while negotiating, primarily caused by their cognitive limitations and social and/or cultural factors. Thus, this chapter mainly derives assumptions regarding how negotiators *act* during the negotiations.

Combined together, these three chapters – and the assumptions within each chapter^[8] – narrow the scope of multilateral bargaining model; provide strategic decisions and actions agents have within the model; determine the behavior underlying these actions and decisions. Note that the aforementioned requirements are translated to concrete (software) requirements, such that a dedicated multilateral multi-issue negotiation environment could be implemented with Python. This "translation" is documented in Appendix D.

[7] "General Environment for Negotiation with Intelligent multi-purpose Usage Simulation", see [97, p. 48].

[8] See Appendix C for an overview.

1.3.2. The two-stage model

In Chapter 5, we discuss the conceptual two-stage model developed to simulate (pre-)coalition formation in international political negotiation. In the first stage agents can form coalitions *ex ante* based on a simple bargaining game. In the second stage, they are able to conduct the actual negotiations by settling issues. Before we introduce the formal model in Section 5.4, we discuss the core model logic in Section 5.3, briefly present the requirements from which the two-stage model is derived in Section 5.2 and motivate our choice for a simulation model in Section 5.1.

Furthermore, in Section 5.5 we discuss the implementation of the conceptual model, motivate our choice for Python and discuss the structure of our code. Moreover, we perform an example run to give the readers an impression as well as check code quality and perform several verification tests in Section 5.6. Note that this section is supported by Appendix F.

1.3.3. Experiments and simulation results

In the first section of Chapter 6, from the (formal) two-stage model, we derive output parameters of interest for negotiators in real international political negotiations. This output finally determines whether pre-coalitions are sufficient or not. In the subsequent section, we create an experimental design used to investigate which factors – e.g. number of negotiators and how they behave – influence the sufficiency of pre-coalitions.

Then, in Chapter 7, we present the results of the simulation experiments. Relations within the experimental design are analyzed by fitting several Random Forest regression models, keeping three questions in mind: should pre-coalitions be considered in the first place? If so, in what negotiation environments are pre-coalitions sufficient? And, how should pre-coalitions be composed given their environment? Each subsequent section answers each question.

1.3.4. Putting the pieces together

Analyses results will simply remain results as long as they are not translated to the relevant context. Therefore, in Chapter 8, we link results to international political negotiations as mainly described in Chapter 2 and Chapter 3. This leads to partial conclusions which, in return, are aggregated to answer the main research question in Section 9.2. Finally, in Section 9.4 we address societal relevance and scientific contributions, whereas in Section 9.5 we discuss the limitations of our work and provide future research directions.

2.

Characteristics of International Political Negotiations

If you don't like how the table is set, turn over the table.
Francis (Frank) Underwood, in House of Cards (2014)

By means of a systematic literature review, this chapter will answer the first supporting question: what characterizes international political negotiations? The purpose of this chapter is to derive assumptions scoping the eventual automated negotiation environment, without suppressing major factors.

2.1. Methodology: systematic literature review

Decades of political science research has resulted in an astounding body of knowledge. So vast that randomly searching studies may lead to severe problems including, but not limited to, no clear overall picture, biased selection of studies and conflicting results [81, p. 3-4]. Systematic literature reviews try to address these by identifying, evaluating and integrating highly-cited studies in a structured way. Therefore, the selection process is objective, transparent and repeatable to some extent [81, p. vi]. Table 2.1 presents this selection process. Please consider Appendix A for a description of the different steps in this table.

2.2. Architecture of political international negotiations

Academic research on international negotiations is often hard, as they happen behind closed doors, especially in multilateral inter-state bargaining [108, p. 23]. As a consequence, many theories have been described based upon evidence heard through the grapevine or gathered via memoirs and other written accounts. Taken this side note into account, the following subsections describe the architecture of international negotiations.

Table 2.1: Search and selection process of the papers and books.

Selected by	Databases				
	IEEE Xplore	Web of Science	Scopus	ABI/INFORM Complete	Miscellaneous
Meta data	15	1549	27	343	-
Title	3	92	2	24	-
Abstract	2	12	2	13	-
Full article	1	8	2	9	8
Total (duplicates)			28 (2)		
Final papers/books			26		

2.2.1. Agreements: negotiation in political realm

In Section 1.1, we carefully delineated coalitions in international political negotiations. However, we did not explicitly state to what extent political negotiations differ from Alex, Bob and Carl's holiday negotiation. So what does the extension "political" add? Before answering, let us first consider the concept of *deliberative negotiation*, which implies [101, p. 92-93]:

"(...) negotiation based on processes of mutual justification, respect, and reciprocal fairness. Such negotiation includes elements of arguments on the merits made by advancing considerations that the other parties can accept; searching for zones of agreement and disagreement; and arguing about the terms of fair processes as well as outcomes."

Thus, pure deliberations are processes of mutual justification, whereas pure negotiations are situations in which individuals give something in exchange for something else, while each time trying to give the least and get the most. Deliberation is aimed at establishing agreement and clarifying conflicts in situations of common interests [101, p. 94]. Of course, these types of situations rarely occur in political negotiations. These are characterized by more distributive means of negotiation, which implies making compromises. This happens in negotiations in which parties have conflicting goals. Mansbridge & Martin argue that both on national as well as international level, political negotiations have shifted from fair to power-based compromises [101, p. 100-103]. Traditionally, fair compromises include asking opponents to give up something, while offering concessions yourself. Parties do so in relatively open situations with mutual respect, without taking advantage of their opponents. In contrast, in this thesis, we argue that most international political negotiations are nowadays characterized by pure distributive negotiations, in which involved parties opt for maximum goals, resulting in power-based compromises. Current polarization of Europe, rise of populism and Trump's "America First" policy are developments supporting this argument. In these circumstances, negotiating parties *"(...) are merely trying to exercise power, exploit institutional advantages, and gain as much as possible at the expense of the other* [101, p. 98]. Thus, we refer to:

Political negotiations as pure negotiations in which issues are conflictive by nature (i.e. if some party gains another party loses) and valued differently. Moreover, parties use their distinct power positions to make strategic demands in which they aim for maximum goals, therefore resulting in power-based compromised agreements (if any).

Most political negotiations contain one or more elements of the aforementioned negotiation types (see [101, p. 100]), but this thesis solely focuses on pure negotiations. As a result, we explicitly assume that:

Assumption 1: Self-interested actors

Actors^a involved in international political negotiations are self-interested and want to maximize their own (absolute) gains. If they engage in behavior that can be labeled as 'cooperation', then this will be in their personal best interest.

^aTo be defined in the following subsection.

This assumption plays a central role in theories of selfish cooperation among states in political literature [110, p. 470] as self-interested behavior has become a self-fulfilling prophecy. If actors assume their opponents act self-interested, they will do the same as it is in their best interest [109].

2.2.2. Actors: who conduct the negotiations?

International negotiations involve numerous (non-)state actors bargaining over multiple issues in many institutions over various years [40, p. 17]. Many researchers have described which actors are involved in international negotiations and how they behave (e.g. [40, 60, 90, 138, 146]). In this thesis, we adopt the *two-level intergovernmental game*, introduced by Putnam (1988, [138]). He assumes the state is not a unitary actor [138]. Instead, domestic state and non-state actors^[1] constrain the behavior of negotiators on an international level by issuing specific goals (and sometimes even strategies; see Figure 2.1). They have the power to do so because any tentative agreement on an international level needs domestic ratification [138, p. 437]. However, international negotiators monopolize the external representation of a state because

^[1]Examples of such actors are the national parliament, companies, political parties, interest groups and many more.

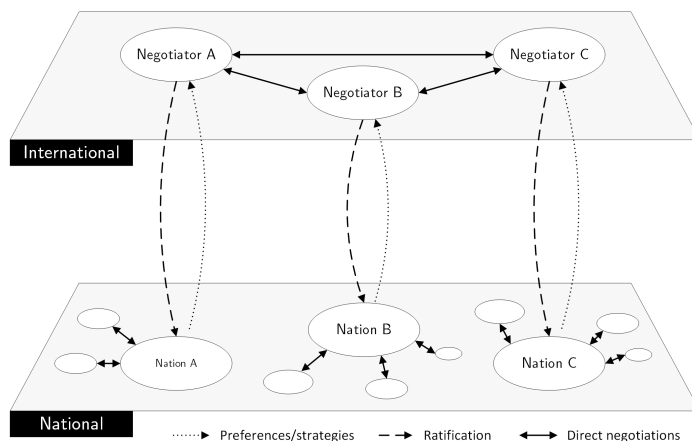


Figure 2.1: Two-level intergovernmental perspective (with three nations)

they have the ability to aggregate and shape the preferences of domestic actors [40, p. 21]. Therefore, they are often represented by politicians (e.g. prime ministers and cabinet ministers) or diplomats [108, p. 24]. In the remainder of this study, we adopt the following assumption:

Assumption 2: Omniscient negotiators

Negotiators (e.g. politicians or diplomats) on an international level represent the state in such a way that their preferences are already the results of ongoing domestic negotiations.

2.2.3. Issues: what are we talking about?

International negotiations discuss many subjects. Recent examples include conferences on climate change – resulting in among others the Paris Agreement – and EU’s refugee crisis meetings. These subjects are often divided in many (sub)issues. For example, during the refugee crisis, the European council has discussed to enhance border control [120], allocation of refugees [32] and fund partnerships with third countries to house refugees locally [121]^[2]. Table 2.2 presents an oversight of the different types of issues, supported by their outcome domain and examples in literature. They are described by drawing an analogy between them and the simple holiday negotiations discussed in the introduction of the first chapter.

Table 2.2: Types of issues under consideration.

Type	Domain	Examples and literature
Binary	Discrete (binary)	Location of IGO headquarters (UN, EMA, EU) [44], war [4, 186], independence [186]
Indivisible resources	Discrete	Intellectual property rights [74], quality standards [183], legislation [15], means of infiltration [4]
Unrestricted resources	Continuous	Greenhouse gas emissions [40, 134, 178], import tariffs [89]
Limited (shared) resources	Continuous	Budget allocation [5, 42], rare materials [125], refugees [32], EU fishing quotas [152], dividing territory [108, 186]

A *binary* issue is best compared with the "winner-takes-it-all". Imagine that during their holidays, Alex, Bob and Carl have collectively won a price in a singing-contest. Leaving Greece, the three friends need to negotiate who will keep the price. The only possible outcome results in one of the friends keeping the price, while others remain empty handed^[3]. In contrary, *indivisible* issues have fixed number of options. Both the

^[2]These examples are derived from a (real-time) time line provided by the European Council (see <http://www.consilium.europa.eu/en/policies/migratory-pressures/history-migratory-pressures/>).

^[3]Of course, assuming the price cannot be divided in parts or shared for alternating periods of time.

holiday destination and the means of transport are already good illustrations. Moreover, *unrestricted* issues can theoretically take any value between a certain range. Pretend that Alex, Bob and Carl have just finished their study and have not applied for a job yet. Then, bargaining over the holiday duration is unrestricted. Of course, assuming that they will not go on holiday for the rest of their lives. Finally, a *limited* issue cannot take any value, because it is either naturally or politically restricted. Imagine the three friends have a fixed and shared holiday budget of €1000. As they spend €450 on a hotel, only €550 is left to negotiate upon other issues restricted by the budget. The following is assumed for the rest of this study:

Assumption 3: Issue types

All to be negotiated issues can be represented by any of the identified issue types – i.e. *binary*, *indivisible*, *unrestricted* and *limited* – or be adjusted in such a way that they can.

2.2.4. Regimes: structure through intergovernmental organizations

International negotiations differ in the sense that there is no overarching mechanism that can intervene when negotiations fail to reach outcomes [108, p. 49]. Because globalization and international conflicts have forced governments to work together, many intergovernmental organizations have emerged. Literature addresses such overarching organizations as *regimes* [38, 78, 83], setting out "(...) principles, norms, rules, and decision-making procedures around which actor expectations converge in a given issue-area" [83, p. 1]. The negotiations conducted in these regimes use various voting rules. Posner and Sykes (2014, [167]) provide a survey on the major rules used by organizations such as the EU, WTO and UN (see Table 2.3).

Table 2.3: Major voting rules in intergovernmental organization (adopted from [167])

Decision rule	Explanation	Deployed
Unanimity (or consensus)	All members need to agree	WTO, International Seabed Authority and climate conventions
Simple majority	More than half of members need to agree ("50% + 1")	United Nations General Assembly, International Court of Justice, International Criminal Court
Weighted majority	The majority of votes where each member can have multiple votes	World Trade Bank, International Monetary Fund, European Council
Super majority	Way more than half of members need to agree (e.g. 60%, 66% or 75%)	Often in European Commission, Council Parliament
Plurality	The alternative with the highest number of votes wins	Choosing location for IGO headquarters by European Council

As outcomes of international negotiations are imposed on all involved parties – no matter what you voted – intergovernmental rules are often more complex. Of course, many alternatives to the above voting rules exist. The EU, for example, uses besides the simple majority and unanimity also *qualified majority*. This implies that more than 55% of member states, representing at least 65% of the EU population, have to vote in favor of new legislation [122]. Historically, veto rights are also common in intergovernmental organizations.

2.2.5. Process: is everything a bargaining game?

In his highly cited book, Barrett distinguishes five phases in international negotiations [16]: i) pre-negotiations, (ii) negotiations, (iii) ratification, (iv) implementation and (v) re-negotiations of which only the first two phases are in scope of this study (see Figure 2.2). We adopt his view by defining a single negotiation process as cycling through these phases once.

Of these phases, pre-negotiations are highly relevant, as they define the rules, actors involved, issues and agenda, thereby shaping the overall negotiation process [178, p. 389-390]. Regimes structure this rather capricious pre-negotiation phase, which shows their importance in determining negotiation outcomes before actual negotiations take place [108, p.391]. Moreover, coalition formation both occurs in this stage, as well as in the actual negotiation phase [149, p. 4-5]. Therefore, we assume the following concerning the pre-negotiation phase:

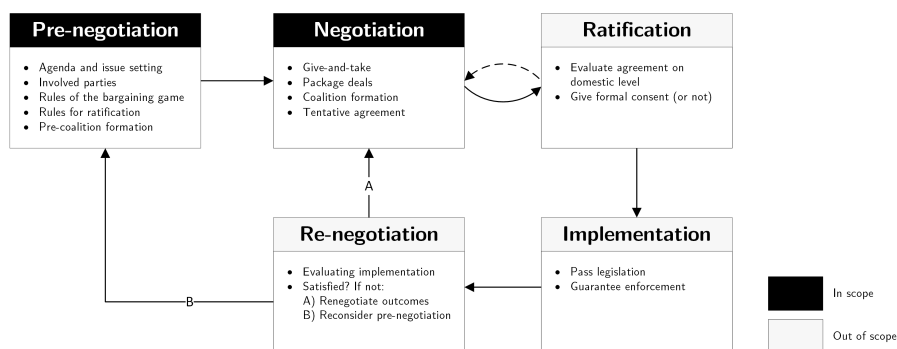


Figure 2.2: Five phases of international negotiations.

Assumption 4: Strategically choose pre-coalitions

Regimes have already defined the bargaining table^a, voting and negotiation rules, issues under consideration and bargaining agenda. Pre-coalition formation is the only strategic action states can take *before* the actual negotiation takes place.

^aActors involved in the negotiation process.

In the second phase the actual negotiations take place. This is a process of "giving-and-taking", mainly characterized by *package deals* [108, p. 80]: creation of common interests by combining different issues in a single offer. These deals are so important that arguably the EU could never have survived without them [108, p. 62]. Because of modern technologies, the actual proposals can be submitted orally, by email or text messages, which has reduced physical burdens of negotiations [30, p. 19-38]. However, as issues become more global – e.g. climate change – the (physical) negotiation rooms are still overcrowded [30, p. 23]^[4]. In this phase, coalition formation is more policy driven, which means that negotiators get to know each other's preferences regarding certain issues and form coalitions accordingly [149, p. 4-5]. While debating and making proposals, countries try to convince others to join their coalition in order to decide upon an issue in their favor. Thus, the actual voting rule is often achieved in this phase, while in the prior phase minority^[5] pre-coalitions are established.

The previous phase results in a tentative agreement, which means that it is an accepted outcome of international negotiations, but need ratification on domestic level [108, p. 36-43]. Thereafter, countries implement and enforce the outcomes by passing legislation. This process may create dissatisfaction nationally, which increases the chance for re-negotiations of either the pre-negotiation or negotiation phase. Both the ratification and implementation and re-negotiations are not in scope of this research.

2.2.6. Power: inequality in international negotiations

Factors influencing the above processes are so vast, that a comprehensive understanding is nearly impossible [108, p. 27]. However, many researchers argue *power asymmetry* of the countries is one of the most crucial factors (e.g. [42, 59, 108, 184]). In negotiations, power is defined as a party's ability to change the probability of negotiation outcomes [101, p. 91]. This can be [59]:

- *Behavioral*, which implies that the power is obtained because of the experience, network and resources of the negotiator representing a state;
- *Structural*, originating from the resources of the country being considered. Examples include economic (e.g. GDP) or natural resources (e.g. oil), but also the number of votes or veto rights in an intergovernmental organization;
- *Issue-specific*, which means that given the to be negotiated issues, a country can have very strong outside options or "best alternative to a negotiated agreement" (BATNA).

^[4]For example, during Paris' climate conference in 2015, one particular room hosted 193 representatives from UN member states, accompanied by delegates from domestic climate agencies and intergovernmental organizations

^[5]Minority in the sense that they do not adhere to the voting rule in place.

Most researchers argue differences in bargaining power have significant impact on strategies used by negotiators [42, p. 558-559], both in the pre-negotiation and actual negotiation phases. Moreover, as we will see in Chapter 3, power asymmetry is also used in political sciences to rationalize coalition formation, as well as to explain a party's ability to make proposals and/or influence the bargaining agenda [172, p. 345-354]. Because power comes from such a variety of sources, we assume:

Assumption 5: Defining power

Negotiators are all equally experienced and, therefore, have the same *behavioral power*. Moreover, *structural power* is derived from the voting rule in place (e.g. seats in the European Parliament). Finally, *issue-specific power* is defined implicitly from actors preferences and opponents' offers.

2.2.7. Strategies: defining negotiating behavior

Negotiation strategies and tactics have been studied thoroughly for decades. These studies have classified strategies in various ways, such as distributive versus integrative bargaining [176], value claiming versus creating [118] and, more recently, soft bargaining versus hard bargaining [42]. Hard bargaining describes tactics resulting in *conflict*, whereas soft bargaining describes tactics used to *cooperate*. Both play a significant role in coalition formation. Table 2.4, provides examples of both.

Table 2.4: Typology of negotiating tactics (adopted from [42, p. 563])

Hard bargaining (conflictive)	Soft bargaining (cooperative)
Making a commitment of not giving in	Signaling flexibility
Criticizing the other side	Making a conciliatory statement, praising the other side
Instigating of joining a defensive coalition	Seeking partners for compromise
Making a threat	Making a proposal for compromise

Not all of these tactics apply to the actual negotiation phase of an international negotiation process (see Section 2.2.5), neither are they exhaustive. Therefore, the following is assumed:

Assumption 6: Negotiation strategy

Negotiators deploy both cooperative and conflictive strategies aimed at maximizing their goals, but adjusted for the amount of risk they want to take for ending in no agreement.

Thus, negotiators are willing to cooperate if it helps them to maximize their overall goals. There is also a close relation between bargaining power and strategies negotiators use. Dür and Mateo show more powerful states most likely opt for hard bargaining strategies. Finally, repeated interaction^[6] also affects a negotiator's strategy as might need each other more than once [42, p. 568-575].

2.2.8. Time: macro and micro impact

In general, time influences negotiation in two distinct ways [108, p. 79]. From a *macro* perspective, past events affect current negotiations. Actors may use previous outcomes as leverage for future negotiations or adjust their expectations accordingly. Time also influences the emerge and evolution of regimes: WOI resulted in the League of Nations, whereas the UN followed WWII [112, p. 438]. Subsequently, EU's expansion has continuously changed the voting procedures [75, p. 700-701]. Both lead to the following assumptions:

Assumption 7: Stability of regimes

Regimes are fixed and stable during a single negotiation process, but may change thereafter.

^[6]Which is defined as looping through the negotiation process – as defined in Section 2.2.5 – multiple times.

Assumption 8: Negotiation memory

Past events *within* a single negotiation (as defined in Section 2.2.5) process influences a negotiator's future negotiation strategy. Past events occurred in *previous* negotiation processes are already represented by a negotiator's preferences and strategies.

From a *micro* perspective, time is both a constraint and driver of bargaining process. Time limits the negotiation process because of deadlines set in the pre-negotiation phase. In addition, time also creates pressure to reach agreements as countries may have their own deadline as result of domestic issues.

2.3. Discussion: is this it?

No, of course not. In the first chapter, we already mentioned negotiations are unruly processes with "capricious", "unpredictable" and "chaotic" behavior [35, 36, 79], which in itself is hard to capture in a handful of architectural components. Meerts (2015, [108]) has extensively surveyed both qualitative and quantitative approaches trying to capture the essence of diplomatic negotiations. He concludes that according qualitative approaches negotiations are often too complex and too contextual to put in a single model (or framework). Citing Berridge et al. (2001, [21, p. 77]): "(...) *different [negotiation] circumstances require different approaches*". Meerts argues that this view is too limited as a combination of both approaches helps to understand negotiations te best "(...) *in the sense that the quantitative method is an addendum to the qualitative approach*" [108, p. 47]. So, capturing most relevant components of negotiations in a single architectural framework is beneficial, as long as you do not forget the more complex setting it describes.

2.4. Conclusion & summary

This chapter answers the first supporting question, namely: what characterizes international political negotiations? We have seen that the architecture of international political negotiations consists of eight components influencing one another. From literature, architectural elements are deduced to describe international political negotiations from a high level perspective. These include:

- The type of *agreements* these negotiations result in. This thesis assumes political negotiations come to power-based compromises. However, in practice, negotiations occasionally may also partly result in fair compromised, (partially) integrative or informed consent agreements;
- The *actors* involved, which are state or non-state and acting on (inter)national or lower levels;
- The type of *issues* discussed. Four types are distinguished based upon possible outcomes they can result in. These are either discrete (binary/indivisible) or continuous (unrestricted/limited resources);
- The *regimes* through which the negotiations are structured and governed, such as the European Union. These institutions also often impose certain decision rules, of which the most well-known are unanimity, plurality or simple, weighed or super majority rule;
- The negotiation *process* shaped by the aforementioned regimes. In general, this process is a cycle consisting of five phases, namely i) pre-negotiation; ii) actual negotiations; iii) ratification; iv) implementation and v) re-negotiation. Only the first two are in scope of this study;
- The *power* positions of actors involved. This power is, for example, derived from negotiator's experience (behavioral), from the decision rule (structural) or from subjects under negotiation (issue-specific);
- The *strategies* used, which vary extremely. This study assumes they are hard or soft, which implies conflicting or cooperating strategies respectively;
- The influence of *time*, which can either have a micro or macro impact.

Note that political negotiations are marked as unruly processes with "capricious", "unpredictable" and "chaotic" behavior [35, 36, 79]. Thus, the above eight elements only capture its structural essence. Which is relevant, as long as you do not forget the more complex setting it describes.

3.

Coalition Formation in International Politics

Just when I thought I was out, they pull me back in.
Michael Corleone, in *The Godfather Part III* (1990)

This chapter answers the second supporting question: which theories have been developed to explain coalition formation in (international) political negotiations? Again, a systematic literature review is performed similar to the previous chapter. By doing so, assumptions are derived defining the strategic decisions and actions negotiators have. Table 3.1 presents the selection procedure of articles (see also Appendix A.2).

3.1. Pre-negotiation coalitions: why are states working together?

In the pre-negotiation phase, we follow Van Deemen (1997) in his view that coalitions are ex ante strategic choices, rather than the result of negotiations [173, p. 6-7]. This implies pre-coalition formation comes down to selecting a 'good' coalition from a set of possible coalitions. So why should we choose to be part of a coalition in the first place? "*Public policy is a public good (or bad)*" [17, p.189], implying no matter if you are in the final winning coalition, you have to accept the outcomes as they are applicable to everyone. Those who are not part of the winning coalition might have the "luck" outcomes are beneficial. However, being part of the winning coalition makes someone decisive, increasing the chance of being "lucky" [65, p. 22-26].

3.1.1. To face threatening situations

The balance-of-power theory is used to rationalize the strategic choice of forming alliances in threatening situations [175, p. 17]. It describes two options states face when creating coalitions, by assuming three different power positions: weak, secondary and strong. As introduced by Waltz, secondary states are: "(...)

Table 3.1: Search and selection process of the papers and books.

Selected by	Databases				
	IEEE Xplore	Web of Science	Scopus	ABI/INFORM Complete	Miscellaneous
Meta data	17	123	212	39	-
Title	8	34	45	25	-
Abstract	3	15	20	6	-
Full article	1	9	13	5	1
Total (duplicates)			29 (0)		
Final papers/books			29		

major powers, rather than great powers" [177, p. 44]. So, they are certainly not as powerful as strong powers but definitely more powerful than weak states.

First, states can **balance** with equal or slightly less powerful countries to form alliances against threatening situations. They do so to protect themselves from states with superior power resources. Two reasons underpin balancing behavior. Joining less powerful countries avoid being dominated by allies, while aggregated power resources help blocking aggressive opponents.

Secondly, states may **bandwagon** to form alliances with stronger opponents. For doing so, weaker states need to make concessions in order to harvest some of the powerful state's gains [107, p. 161-162]. Rationally, a weaker country should align itself with more powerful states "(...) *because the latter can take what it wants by force anyway*" [107, p. 163]. So, if you deem the goals of powerful states acceptable, than at least you are maximizing your own minimal demands, as it will probably not become any worse. This is comparable with a *maxmin* strategy in game theory (see [96, p. 15-16]).

Historic evidence shows that balancing is more common than bandwagoning [175, p. 20]. This should come as no surprise as bandwagoning assumes we are certain about the more powerful state's preferences, which makes it more risky. Especially considering the additional resources brought by bandwagoning might result in uncertain behavior [175, p. 19-21]. However, the weaker the state, the higher the chance of bandwagoning behavior. In addition, states are also more likely to bandwagon when there are limited like-minded countries^[1]. In addition, "*although power is an important part of the equation, it is not the only one*" [175, p. 19]. Similar ideology, geographic proximity, aggregated power and aggressive intentions are also influencing this strategic decision [26, 175]. Finally, "(...) *status quo states should avoid provoking countervailing coalitions by eschewing threatening foreign and defense policies*" [174, p. 4]. So, the willingness to change the current situation is another important factor influencing ex ante coalition formation.

Since its existence, the balance-of-power theory has been subjected to major criticism [157, p. 204]. First of all, both theories stand in sharp contrast to each other. They provide mutual exclusive predictions which coalition partners states will choose [174, p. 8]. Moreover, balancing and bandwagoning tend to be naturally framed solely in terms of power. This is a serious flawed view which often leads to contrasting predictions and actual historic events [174, p. 16]. Therefore, in this thesis, balancing and bandwagoning are framed in terms of power *but* applied in different negotiation settings with varying ideological distances, aggressive behavior, offensive capabilities and more (see Chapter 6). More specifically, we say the theory defines two distinct options to form pre-coalitions, but does not prescribe the negotiation settings in which apply.

3.1.2. To become more decisive

"*Power is the currency of politics*" as parties exchange power resources to reach mutual beneficial goals [57, p. 374]. Coalitions are means to exchange these resources. So, choosing (powerful) coalition partners is an important strategy to increase one's decisiveness and, hence, its chances of being lucky.

Cooperative game theory has been widely studying a party's decisiveness in bargaining games mainly in terms of *voting power* – i.e. the ability of parties to control the negotiation outcome "(...) *through the resources made available to the voter by the decision rule*"^[2] [50]. Voting power is generally divided in two categories. First, *I-power* describes the likelihood parties are decisive – and thus are part of the winning coalition – given the voting rule and assuming parties accept proposals independently. The Banzhaf index is an example of I-power (see [14]). Secondly, *P-power* defines a party's expected share of the pie as the main source of voting power. The Shapley-Shubik index is an example of P-power (see [156]). More similar indexes are studied in literature^[3]. As each of these indexes give an indication of a party's "true" power, they are used as a measure to analyze coalition outcomes from a cooperative perspective [163, p. 177-183].

However, in Chapter 1 we already described that by doing so, many factors in the actual negotiation process are neglected. Therefore, in this thesis it is not assumed that there is a direct relation between a state's true power in the bargaining game and its final winnings. Instead, we assume an implicit relationship. If states form ex ante coalition, their true power might increase disproportionately. Therefore, not only can they make greater strategic demands, also the likelihood increases that their respective offers are considered by others. Therefore, the pre-coalition may strengthen its position in the actual negotiation process. So, for the remainder of this study we assume that:

^[1]More hypotheses concerning (conditions favoring) balancing or bandwagoning are enumerated in Walt (1987) [175, p. 30-31].

^[2]This is closely related to the structural power described in Section 2.2.6.

^[3]We refer to Felsenthal and Machover (1998), who surveyed various power indexes extensively [50].

Assumption 9: True power position

The a priori true power position of a state or coalition of states^a influences both its strategic demands and likelihood their respective demands are considered in the bargaining process.

^aWhich can be defined by both Shapely-Shubik and Banzhaf indexes depending on the bargaining situation.

Note that whenever a pre-coalition is formed, there might still be power asymmetry between its members influencing the demands a coalition makes. Van Deemen (1997) suggests the use of *power excess* (*pow*) to measure the control of a party in a coalition, which is defined as [173, p. 171-174]:

$$pow(i, C) = p_i - p_{C-\{i\}}$$

with $i \in A$ is and $C \subseteq A$. In the above definition, i denotes a country from all involved countries A , C denotes a coalition as subset of A and p_i is a power measure. Aligning with Van Deemen (1997), we assume that:

Assumption 10: Power excess

The ability of a party to make proposals within a coalition is defined by his/her *power excess*, in which the power measure p_i is a priori defined according to some power index.

3.1.3. To find like-minded states

A limitation of power indexes is that they do not incorporate the parties' policy preferences, so every coalition forms with equal probability. To encounter this deficiency, connectivity models have been developed [164, p. 1140-1144]. Dating back to the French Revolution, "game-theoretic like" methods^[4] were already applied to explain why individuals would (not) work together based upon a one-dimensional ideological scale from left (liberal) to right (conservative) [163, p. 197]. Figure 3.1 shows a fictitious example. For this one-dimensional scale, Axelrod introduced the concept of *connectivity*: the only coalition that can form are coalitions in which parties are bound together on the ideological scale [10].




Figure 3.1: Fictitious one-dimensional scale (e.g. economical) for ordering the political parties a_1 , a_2 , a_3 , a_4 and a_5 .

The one-dimensional scale and concept of connectivity is rather limited, as beliefs are hard to map. This has accelerated the use of ideological proximity models for analyzing coalitions (e.g. [127, 155]). They use multiple dimensions – i.e. cultural, economical and religious – to describe the distances between parties (see Figure 3.2 for two and three dimensions respectively). In contrary to one dimension, with multiple dimensions proximity becomes a distance. Quite often the Euclidean distance is used to denote the probability p a coalition will form^[5] [22, p. 689-690]. Thus, probabilistic models emerge in which the likelihood of coalition formation increases as the ideological distance decreases [27, p. 1150-1151]. It is not hard to imagine that if states want to form coalitions, they will look for ideological close members as these coalitions are presumably more stable.

3.2. Coalitions while negotiating: why do states remain together?

In contrary to pre-negotiation stage, the formation of coalitions in the actual bargaining stage is less of a strategic choice and more a consequence of the negotiations, convincing other parties to join (or leave) pre-coalitions. This raises the next question: why do pre-coalitions remain together during the actual negotiations? Therefore, we deduce some fundamental game-theoretic principles for coalitional bargaining.

^[4]Arguably, modern game theory began with the introduction of Von Neumann and Morgenstern's book the "Theory of Games and Economic Behavior" in 1944, well after the French Revolution.

^[5]On a one-dimensional scale, a coalition has either $p = 1$ if it is connected or $p = 0$ if it is not. In contrary, in ideological proximity models a distance metric is used such that $p \in [0, 1]$.

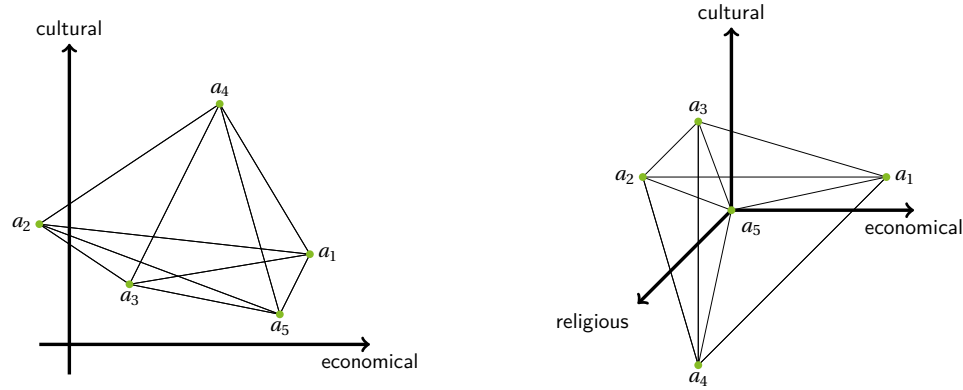


Figure 3.2: Fictions two-dimensional (left) and three-dimensional (right) ideological proximity model describing the connections between political parties a_1 , a_2 , a_3 , a_4 and a_5 .

3.2.1. Rationality and the core: fundamental game-theoretic principles

In this study, rationality is referred to as sensible and justifiable explanations of decisions. Cooperative game theory defines two fundamental properties for rational parties that are also important for bargaining situations. Let us define x_i as the share of the pie offered to party i , $v(i)$ as the lowest guaranteed outcome for i and A all the parties involved [163, p. 150-151]. Then:

Property 1

Individual rationality: $x_i \geq v(i)$ for all i . As $v(i)$ is a party's lowest guaranteed outcome, it would be irrational to expect any outcome below this value.

This implies, for example, that if France currently receives €1 billion from EU subsidies for agriculture, it would not expect anything less in future negotiations, as the *status quo* makes them better off. Assuming, of course, that future negotiations solely involve agriculture subsidies and do not consist of package deals.

Property 2

Collective rationality: $\sum_{i \in A} x_i = v(A)$. As $v(A)$ is equal to the total pie to be divided, we cannot have $\sum_{i \in A} x_i > v(A)$. In addition, it would be irrational to have $\sum_{i \in A} x_i < v(A)$, as all parties can increase their outcomes by forming the grand coalition^a.

^aThe coalition containing all involved parties.

So, if pie π is not completely divided yet, there is an outcome making at least one party better off without harming the opponent. For example, if the EU agricultural budget consists of €10 billion and there is an agreement that only divides €9 billion, this agreement is considered to be irrational as the remaining €1 billion can be divided in such a way that no one would be worse off. Both these properties define the *core* as:

The core

The *core* of a bargaining game is the set of all payoffs (x_1, \dots, x_m) such that for all $C \subseteq A$, $\sum_{i \in C} x_i \geq v(C)$, while adhering to both individual as well as collective rationality.

So, if C is the coalition that originates during the actual negotiation process, i is in C only if he gets a higher payoff from being part of the coalition. If this is true, i either **joins** or **stays** in coalition C . The opposite also holds, which means that whilst bargaining if $i \in C$ and if the offer of coalition C to another party $i+1$ is lower than v_i , then i 's rational decision is to **leave** the coalition. Therefore, we consider the core as a justification why parties join/remain together during the coalitional bargaining process.

However, there are some practical objections of using the core in a complex, political setting. First, states are assumed to be self-interested parties who are only working together for their own gains, which contradicts the underlying cooperative game theoretic assumptions of the core [86, p. 9]. In addition, diplomats negotiating on behalf of states can be *irrational* for a single issue in package deals. Therefore, he accepts lower offers than his guaranteed offer if the package deal as a whole is more beneficial. Consequently, we loosen the *collective rationality* constraint by $\sum_{i \in A} x_i \leq v(A)$, as long as the solution is of benefit for the party under consideration. In addition, we loosen the *individual rationality* constraint as well by saying that in package deals with m issues, for i 's payoff vector (x_1^i, \dots, x_m^i) the following holds: $\sum_m x_j^i \geq \sum_m v_j(i)$. Note that we can only do this by assuming that:

Assumption 11: Transferable utility

Utility is transferable between both *issues* and negotiating *parties*.

This implies a common currency valued equally by all parties, such as money [163, p. 130]. Together with the previous assumption and property 1 and 2 this leads to the following assumption:

Assumption 12: Joining/leaving rationality

Self-interested parties will only join or stay in a coalition if his gains are at least as much within the coalition as it could benefit by himself [159, p. 223], or in any other case they will leave.

Finally, rationality assumes that diplomats have complete information of the negotiation setting. Of course, many economists have rejected this assumption and stated that agents (i.e. diplomats) are bounded rational, which assumes that when making decisions, they are limited by the knowledge of the bargaining problem, information they have, cognitive limitations, time available to make a decisions or social and cultural factors [8, 115, 153]. Their decisions are noisy: given similar situations, people may make different offers. Thus, for the remaining of this study we assume that:

Assumption 13: Bounded rationality

Agents are bounded rational, which implies they make decisions given the (incomplete) information, cognitive limitations and available decision time.

3.2.2. Riker's minimal winning coalitions

In 1962, William Riker introduced a fundamental principle – known as the "size principle". It states "*(...) [parties] create coalitions just as large as they believe will ensure winning and no larger*" [144, p. 33, 47]. This seems rational, as the addition of redundant coalition members results in a decrease of the original member's maximum payoff [163, p. 196-197]. However, some empirical studies not support this principle. Coalitions may be over-sized such that they are more stable – i.e. to reduce the within coalition disagreement [26, p. 73-74]. Moreover, to be certain decisions survive ratification negotiators may prefer redundant coalition partners as long as it does not change their winnings (much). Therefore, we assume that:

Assumption 14: Stopping criteria

Negotiations end if either the public deadline is passed or members of the winning coalition agree its size is adequately "large enough"^a.

^aWe refer to Chapter 5 to see how this is modeled in our tool.

So, coalitions might **accept** new members even though they are already minimal winning. But why do they accept new members in the first place? Rationally, a coalition accepts a new member only when the "value" of the newly formed coalition is at least greater or equal to the original coalition [85, p. 62]. In our case, aforementioned value refers to the combined utility derived from all payoffs (x_1, \dots, x_m) that the

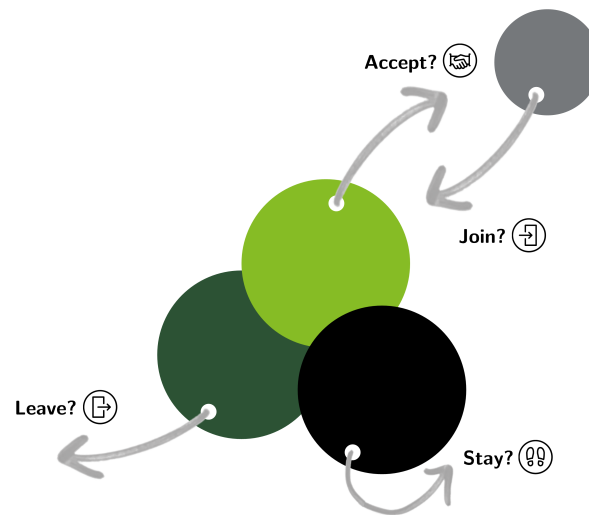


Figure 3.3: Different actions for individual parties and members of a coalition while negotiating.

coalition would have received if the negotiation ends at that moment in time. With this final statement, we have defined all actions coalition members have (see also Figure 3.3).

3.3. Conclusion & summary

This chapter answers the second supporting question: which theories have been developed to explain coalition formation in (international) political negotiations? The answer to this question is twofold. First, there are theories and models explaining the choice of coalitions *ex ante*. Political science scholars provide the balance-of-power theory *to face threatening situations*. Basically, when facing a threat, weak, secondary or strong states can either balance – i.e. allying with the weaker side – or bandwagon – i.e. join the strong side. Furthermore, states may form pre-coalitions *to become more decisive*, which implies that power resources are aggregated to increase states' true power position in the negotiations. Cooperative game theory studies states true power positions extensively in terms of power indexes (i.e. Shapley-Shubik). Lastly, game theory provides ideological proximity models *to find like-minded states*. The closer they are together in these models, the more likely they form coalition *ex ante*.

Secondly, cooperative game theory provide rational explanations why states remain together during the actual negotiation phase. Basically, they face four options while negotiating: states may either *accept* (new) coalition members or *join* a coalition, or they might decides to *stay* or *leave*. The core describes that whenever a state gets a higher payoff from being part of the coalition he either chooses to stay or join. In contrary, if the offer he gets is lower than his guaranteed payoff (i.e. outside option), it is individual rational to leave the coalition. Finally, new members are accepted whenever the coalition is not adequately large enough. So, either till minimal winning (i.e. Riker's size principle) or slightly larger than minimal winning.

By defining these four strategic actions we follow more recent literature on coalitional bargaining. Traditionally only the emerge – both the 'join' and 'stay' actions – of coalitions is modeled using alternating offers (e.g. [31, 67, 123]). Later, this was extended by Ray and Vohra's [143] Equilibrium Binding Agreement (EBA), also modeling the collapse – i.e. the 'leave' action – of emerged coalitions. Comparable to our model, other authors further refined this model by including the re-emerge of collapsed coalitions, similar to the 'accept' decision in our model (e.g. [39, 55]). On top of that, we describe the *ex ante* strategic choice of creating coalitions, potentially influencing all four actions during the actual negotiation process. This is distinctly different from other recent coalition bargaining models.

4.

Fundamentals of Automated Negotiation

I'm not bad. I'm just drawn that way.
Jessica Rabbit, in *Who Framed Roger Rabbit* (1988)

This chapter answers the third supporting question: what are the main components of automated negotiation environments? This is done by performing a concise literature review. Concise, because there has already been some work done taking helicopter view of this field. Though, to lesser extent in multilateral multi-issue negotiations. Conceptually, a *negotiation setting* in the field of automated negotiation consists of three major components: the *protocol*, *scenario* and negotiation *strategies* that are used by negotiating parties [49, p. 4-7]. We will describe them in more detail in the following sections.

4.1. Protocol: rules governing the negotiation process

The protocol determines the set of rules governing the negotiation process and, therefore, it specifies the actions each agent has [11, p. 22]. In a political domain, it should deal with two dimensions, namely *multiple issues* and *multiple parties*. Moreover, it defines negotiation *deadlines* and *selects* of initial proposers.

4.1.1. Negotiation among multiple parties

Arguably, the most influential protocol is the *alternating offers* procedure introduced in Rubinstein's classical bilateral negotiation model [148]. This protocol describes negotiation as a sequence of offers between two parties. Party a_1 begins by making an offer to a_2 . This is 'generated' according to some function Ω_{a_1} depending on his strategy (see Section 4.3.1). According to party a_2 's acceptance function α_{a_2} , he will either accept or reject a_1 's proposal. If opponent a_2 rejects he will generate a counteroffer according to Ω_{a_2} after which a_1 has the decision to accept according to α_{a_1} . If a_1 rejects, the negotiation moves to the next stage until either of the parties accept. Multilateral extension of this protocol is not without difficulties. It leads to multiple optimal solutions and, hence, no unique outcome to the bargaining problem [126, p. 63-65]. By altering the rules of the game, two *multilateral* protocols have been proposed regaining uniqueness.

In their work, Chae and Yang studied a protocol in which multiple parties engage in series of bilateral negotiations [29, p. 88-90]. Their procedure starts by choosing an initial proposer a_i . Then, a_i selects a second party to make an offer. If this second party a_{i+1} accepts the proposal, he has to wait until the end of the bargaining round, while the initial proposer continues making proposals to all other parties. If a_{i+1} had rejected, however, he will become the initial proposer. This game ends when a single offer is accepted by (enough) remaining parties according to a decision rule, similar to those presented in Section 2.2.4.

Another slightly distinct procedure is presented by Krishna and Serrano [87, p. 64-66]. They describe a protocol in which an agent makes an offer to all other parties simultaneously. Agents receiving the offer are allowed to exit the bargaining game by partial (dis)agreements. This means that when a particular voting rule is achieved and, thus, enough parties agree with their share, an agreement is reached. When the voting rule is not achieved, a new party is chosen to be the proposer, after which he will make an offer to all parties. In contrary to Chae and Yang, this protocol is therefore 'truly' multilateral.

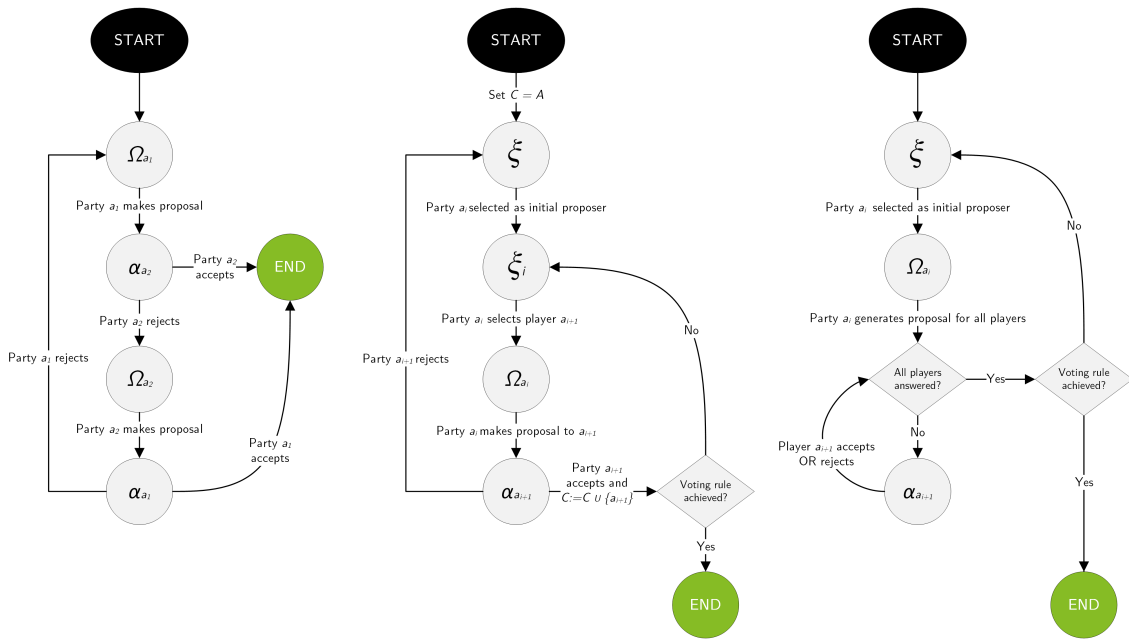


Figure 4.1: The bilateral alternating offers protocol introduced by Rubinstein (left) and two multilateral offers protocols proposed by Chae and Yang (middle) and Krishna and Serrano (right).

Figure 4.1 shows the flow diagrams of Rubinstein’s protocol and the two multilateral extensions described above. Notice the increased complexity of both multilateral bargaining procedures, as they now not only rely upon functions to generate and accept proposals, but also on voting rules and agent selection function(s).

4.1.2. Negotiation over multiple issues

There are four high-level procedures for negotiating over multiple issues addressed in literature [51, 135], which are considered to be:

- 1) *Package-deal procedure.* The parties tackle all the issues directly, at once. They alternate offers existing of a complete package of issues that specifies how each one of the issues should be settled. If the opponents’ do not agree, one of them suggests a new package and so forth;
- 2) *Simultaneous procedure.* Negotiations over individual issues are handled separately and independently of each other. In a political context, every party sends a negotiator N_j , solely dealing with negotiating

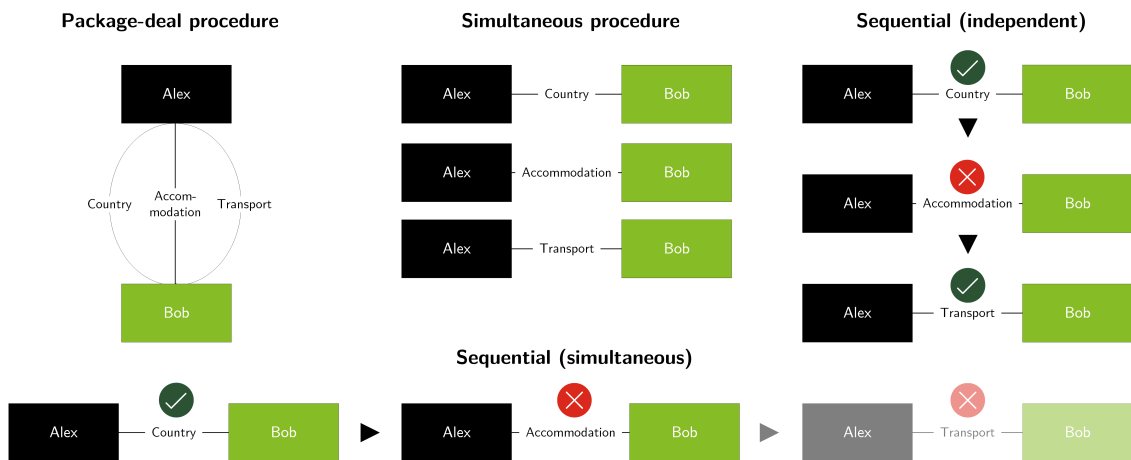


Figure 4.2: Four key procedures for negotiating over multiple issues (figure is partly adopted from [66, p. 11]).

on issue j and negotiator N_{j+1} to negotiate on $j+1$ and so on. None of the outcomes of a single issue can affect the other outcomes, as at the time of negotiation the other outcomes are unknown;

- 3) *Sequential procedure with independent implementation.* Negotiators consider one issue at the time in this procedure. Therefore, it strongly depends upon a *negotiation agenda*. Note that because of this, outcome of issue j most likely effects the outcome of $j+1$ making issues interdependent;
- 4) *Sequential procedure with simultaneous implementation.* This procedure is similar to the previous one, except that an agreement is only struck if agreement is reached upon *all* subsequent issues.

The above procedures are schematically presented in Figure 4.2, again by drawing an analogy between Alex and Bob's holiday negotiations introduced in the first chapter of this thesis.

4.1.3. Time and deadlines

Another key characteristic of protocols is defining negotiation deadlines, which refers to a certain time t before an outcome should be established. If the deadline passes negotiating agents fall back on their "best alternative to no agreement" or BATNA [140]. If the deadline is part of the protocol, it is assumed public knowledge. However, each agent may also have their own, private deadlines, which is not shared [11, p. 23]. In both cases, the deadline can be defined as the maximum number of rounds per issue [162], or be set globally as an overall constraint [11, 47]. For the remainder of this study, we assume that:

Assumption 15: Public deadlines

Deadlines are hard, political targets set prior to the negotiation and, thus, public information.

4.1.4. Selection function

The multilateral protocols introduced before use *randomly* selecting agents to regain uniqueness of the bargaining solution. More specifically, an initial proposer is chosen randomly, but is only selected for the second time all other players have been chosen already. In political negotiations, this random selection is not realistically and, therefore, other suggestions have been made in literature.

Simon and Rausser introduce a function selecting initial proposers randomly, but according to some kind of *access vector* [141, p. 1]. This vector describes a political concept representing actors' power in terms of resources and relationships [141, p. 3]. Many others (e.g. [3, 154, 160, 169]) have subscribed the use of access vectors and defined it more specifically as: "(...) *players' ability to influence the process*" [154, p. 4].

Both Kalandrakis and Herings et al. use a Markov process to describe the probability a player is chosen the next round given he was a proposer in previous periods ([72, p. 407-408] and [64, p. 308]). The transition probabilities for this process are then given by matrix $W(A)$, so that $W_{a_1 a_2}(A)$ gives the probability the next proposer is a_2 if current proposer a_1 , with both $a_1, a_2 \in A$.

Assumption 16: Selecting initial proposer

Selecting an agent can either be done *randomly*, according to their *power position* or a combination of both in which the selection function is a *Markov chain*. These functions are used to select i) initial proposers and ii) opponents. The selection function for i) and ii) might differ for single negotiations.

4.2. Scenario: outcome space, preferences and utility

The negotiation scenario consists of two major components, namely the negotiation domain and preference profiles [11, p. 14]. The negotiation *domain* – or outcome space – describes all possible outcomes. So, the domain consists of all to be negotiated issues, in which each issue itself has certain possible outcomes. To reach consensus, agents must agree upon a specific value x_j for each issue j . Thus, the final agreement on m different issues is a vector $x = (x_1, \dots, x_m)$.

Every outcome x_j is valued differently for individual agents. Preferences for these outcomes is described in their *preference profiles*. A profile is nothing but an ordering of (possible) outcomes. This can be a ranking (e.g. Alex prefers Greece before Turkey), but often a *utility function* is used [11, p. 24-25]. This function –

further denoted with u – simply assigns a value to every outcome x_j and orders them accordingly. So, in the holiday example $u_{Alex}(Greece) > u_{Alex}(Turkey)$. Note that preference profiles belong to an agent's private information. Literature distinguishes two types of utility functions dealing with multiple issues:

- *Linear additive.* The most common used method is simply to compute the combined utility U by weighting all m issues as follows such that $U(x) = \sum_{j=1}^m w_j \cdot u_j(x_j)$. In which w_j denotes the weight for issue j and u_j is the utility function for issue j given outcome x_j .
- *Non-linear.* Many studies (e.g. [71, 98, 145]) show that assuming independent issues is too simple for real world problems. Therefore, non-linear utility functions with interdependent issues are proposed.

The addition of interdependencies increases the outcome space tremendously resulting in multiple optima and harder to find solutions to bargaining problems. Moreover, utility functions u_j may differ for every agent, as well as for every issue, which increases complexity for the combined utility function U even more. Therefore, for computational ease, we assume:

Assumption 17: Independent issues

Issues are independent and the combined utility function U can be defined by a (weighted) linear combination of the utilities u_j of all independent issues j .

For determining utilities, preference profiles consist of both target and reservation values [11, p. 24-26]. A *target value* is something agents aim for, which can both be low or high – i.e. when Bob buys flight tickets, he aims at getting a low price while airlines want to sell the tickets for as much as possible. A *reservation value* is the value still reasonable to an agent for reaching agreement. Bob might have a maximum budget of €100 for each ticket, while the airline would not go lower than €95. Hence, both values together result in a range of potential outcomes. For Bob, this might be [€60, €100] and for the airliner [€95, €140]^[1]. Literature commonly evaluates each of these outcomes differently for various moments in time. Roughly two methods are proposed in (bilateral) bargaining models [113, p. 42-55]:

- *Fixed bargaining costs:* utility for issue j is determined by $u_j(x_j) - c \cdot t$, where c is the *fixed cost* of bargaining and t the time that have passed;
- *Fixed discounting factor:* utility for issue j is corrected by $u_j(x_j) \cdot \mu^t$, where $\mu^t \in [0, 1]$ is the fixed discounting factor.

Recently, also functions are introduced consisting of both a deterministic and random component, capturing the "(...) idiosyncratic and unpredictable factors of human behavior [e.g. preferences]" [19, 391-392, 396-401]. This results in the utility function $\tilde{u}_j = u_j + \epsilon$ in which ϵ is a random agent-specific component representing capricious behavior. Therefore, we assume that:

Assumption 18: Discounted utility with uncertainty

The utility function for agent i depends on the outcome x at time t – denoted by $U_i(x, t)$ – is described by $U_i(x, t) = \mu_i^t \cdot u_i(x) + \epsilon$, in which u_i is the agents private utility function and $d_i \in (0, 1]$.

4.3. Strategy: bid to negotiate

Before introducing strategies, we first need to describe *tactics*. According to [45, p. 6-19], offers are generated by a combination of functions, named tactics. Each tactic τ generates an offer, based upon a specific component of the negotiated object such as time, the left-over pie and negotiation behavior. A weighted combination of tactics result in the eventual offer. Hence, if y tactics are used and the final proposal x at time t for issue j from party a_1 to party a_2 is denoted with $x_{a_1 \rightarrow a_2}^t$ then:

$$x_{a_1 \rightarrow a_2}^t = \omega_{j1} \tau_{j1}(t)[j] + \omega_{j2} \tau_{j2}(t)[j] + \dots + \omega_{jy} \tau_{jy}(t)[j] \quad \text{with } \forall j \in I: \sum_{i \in [1, y]} \omega_{ji} = 1$$

^[1]The **bold** values indicate the target value. Note that the way offers are generated is explained in the following section.

Given a set of tactics, an agent adjusts his combination of tactics to suit a particular issue j . Hence, for all z issues, a party has its own weighted combination of tactics at time t , given by the matrix Ω :

$$\Omega_{a_1 \rightarrow a_2}^t = \begin{bmatrix} \omega_{11} & \dots & \omega_{1z} \\ \vdots & \ddots & \vdots \\ \omega_{y1} & \dots & \omega_{yz} \end{bmatrix}$$

A *strategy* is then defined as any function f such that:

$$\Omega_{a_1 \rightarrow a_2}^{t+1} = f(\Omega_{a_1 \rightarrow a_2}^t)$$

Thus, a *strategy* expresses the way a party changes his weights ω of the different tactics τ over time. So, to formulate strategies, we have now shifted the problem to the design of tactics.

4.3.1. Tactics: different actions to achieve your goals

This section describes three families of tactics introduced by Faratin et al. [45]. First, **time-dependent** tactics are designed such that offers are mainly dependent upon time t and deadline θ_j for issue j . Following the definition of [45, p. 14], an offer is generated as follows:

$$x_{a_1 \rightarrow a_2}^t[j] = \begin{cases} \min_j^{a_1} + \alpha_j^{a_1}(t)(\max_j^{a_1} - \min_j^{a_1}) & \text{If } a_1\text{'s utility decreases with a higher share of } \pi \\ \min_j^{a_1} + (1 - \alpha_j^{a_1}(t))(\max_j^{a_1} - \min_j^{a_1}) & \text{If } a_1\text{'s utility increases with a higher share of } \pi \end{cases}$$

In which $\min_j^{a_1}$ and $\max_j^{a_1}$ denote party a_1 's reservation and target values respectively. Thus, a family of time-dependent strategies is then defined by adjusting the function $\gamma_j(t)$:

- Polynomial: $\gamma_j(t) = \kappa_j^{a_1} + (1 - \kappa_j^{a_1})(\min(t, \theta_j)/\theta_j)^{1/\beta}$
- Exponential: $\gamma_j(t) = \exp(1 - \min(t, \theta_j)/\theta_j)^{\beta \ln \kappa_j^{a_1}}$

This leads to an infinite number of possible tactics. Two distinct sets for different values of β have been described by [45, p. 14-15]. First, *Boulware tactics* are either polynomial or exponential functions with $\beta < 1$. In this tactic, an agent stays relatively close to his initial offer the majority of time, taking the risk of no agreement. In contrary, *Conceder tactics* have $\beta \gg 1$. Here, an agent takes less risk and admits to an opponent's offer easily. See the Figure 4.3 to see this distinct behavior for different values of β [2]. Next, **resource-dependent** tactics are very similar to time-dependent ones if time is the only "resource" that is scarce. Hence, they are modeled in the same way apart from $\gamma_j(t)$ now depending on a resource such as money or parties involved. Formally, this can be modeled as:

$$\gamma_j(t) = \kappa_j^{a_1} + (1 - \kappa_j^{a_1})e^{-resource(t)}$$

In which $resource(t)$ is the resource dependent function. The latter is relevant for the way in which we want to model political negotiations as the main resource – pie π (e.g. budget)– diminishes over time and parties might adjust their tactics accordingly. Finally, **behavior-dependent** tactics is the last family proposed by Faratin et al. [45, p. 17-18] and depends on the opponent's (counter)offering behavior. Within this family, three different *subfamilies* have been proposed, each imitating the opponent's tactic in some distinct manner:

- *Relative Tit-for-Tat*: a party copies his opponent's offering behavior in relative terms, based upon $w \geq 1$ steps ago. Thus, more formally and in terms of $\min_j^{a_1}$ and $\max_j^{a_1}$, this tactic can be formulated as:

$$x_{a_1 \rightarrow a_2}^{t_{n+1}}[j] = \min(\max(\frac{x_{a_2 \rightarrow a_1}^{t_{n-2w}}[j]}{x_{a_2 \rightarrow a_1}^{t_{n-2w+2}}[j]} x_{a_1 \rightarrow a_2}^{t_{n-1}}[j], \min_j^{a_1}), \max_j^{a_1}) \quad (4.1)$$

[2]Note that in this figure, t_{max} denotes the deadline θ_j for issue j .

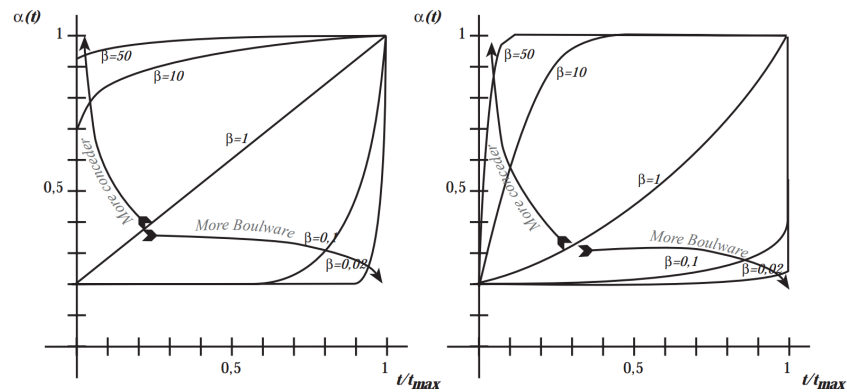


Figure 4.3: Polynomial (left) and exponential (right) time-dependent tactics for different values of β (adopted from Faratin et al. [45, p. 15]).

- **Random absolute Tit-for-Tat:** this tactic copies (in absolute terms) opponent's behavior, only with an added random component. Let M denote the maximum amount by which an agent can adjust his absolute imitative behavior. Then, we can formulate this tactic as:

$$x_{a_1 \rightarrow a_2}^{t_{n+1}}[j] = \min(\max(x_{a_1 \rightarrow a_2}^{t_{n-1}}[j] + (x_{a_2 \rightarrow a_1}^{t_{n-2w}}[j] - x_{a_2 \rightarrow a_1}^{t_{n-2w+2}}[j]) + (-1)^s R(M), \min_j^{a_1}), \max_j^{a_1}) \quad (4.2)$$

where $R(M)$ generates a random number on the interval $[0, M]$ and

$$s = \begin{cases} 0 & \text{If } a_1 \text{'s utility decreases with a higher share of } \pi \\ 1 & \text{If } a_1 \text{'s utility increases with a higher share of } \pi \end{cases}$$

- **Averaged Tit-for-Tat:** similar to above, only now an agent computes his average relative offer based upon a predefined time-window of size $v \geq 1$. More formally, this tactic can be defined as:

$$x_{a_1 \rightarrow a_2}^{t_{n+1}}[j] = \min(\max(\frac{x_{a_2 \rightarrow a_1}^{t_{n-2v}}[j]}{x_{a_2 \rightarrow a_1}^{t_n}[j]} x_{a_1 \rightarrow a_2}^{t_{n-1}}[j], \min_j^{a_1}), \max_j^{a_1}) \quad (4.3)$$

4.3.2. Negotiate over issues with discrete domains

So far, we have only been discussing tactics in continuous issue domains. However, in Section 2.2.3, we have also identified issues with discrete domains. In [49, p. 92-94], Fatima et al. proof for binary issues that two agents a_1 and a_2 with deadline θ and discount factor $\mu \in (0, 1]$ that the initial offerer a_1 always has the disadvantage as the binary issue will (rationally) be allocated to a_2 . Thus, depending on μ_{a_1} and β_{a_1} , the a_1 's bidding function (see Figure 4.3) for binary issues transforms to a step function with one step at $0 < t \leq \theta$. Before discussing indivisible issues, we first assume:

Assumption 19: Ordering of alternatives

Agents have a distinct ordering of preferences for indivisible issues with discrete values, either on a *ordinal* (ranking)^a or *interval* scale.

^aFor convenience, we use an inverted scale in which the highest number denotes the most preferred alternative.

So, going back to the holiday negotiation, suppose for accommodation three options were selected: *tent*, *villa* and *hotel*. Alex' ordering is as follows:

$$\text{hotel} \preceq_{Alex} \text{villa} \preceq_{Alex} \text{tent} .$$

As he does not prefer some options way more than others, his correspondent rankings are $r_{hotel}^{Alex} = 3$, $r_{villa}^{Alex} = 2$ and $r_{tent}^{Alex} = 1$. Moreover, Bob has the same ordering of options, however, he really detests the tent. So,

his "ranks" can for example correspond to $r_{hotel}^{Bob} = 10$, $r_{villa}^{Bob} = 8$ and $r_{tent}^{Bob} = 1$. For this example, Alex' and Bob's bidding functions are still stepwise, though – again depending on μ_i and β_i – they differ slightly. For $\beta = 1$, Alex's bidding function will have equal steps in both height and width, while Bob's first and second step are significantly larger than his last step (see [49, p. 92-94 & 114-119]).

4.3.3. Acceptance strategy

Until now, we only discussed bidding strategies. However, all agents also deal with the question whether to accept or reject offers. In (bilateral) models, the acceptance strategy is often modeled in such a way that offers are accepted if they are more advantageous than those agents would make themselves if they were initial proposer^[3]. More formally, let $x_{a_1 \rightarrow a_2}^j$ be the offer a_1 makes to a_2 for issue j and $x_{a_2 \rightarrow a_1}^j$ be the offer a_2 would have generated if he was the initial proposer. Then a_2 accepts a_1 's offer if and only if [131, p. 10-13]: $u_j^{a_2}(x_{a_1 \rightarrow a_2}^j) \geq u_j^{a_2}(x_{a_2 \rightarrow a_1}^j)$, in which $u_j^{a_2}$ denotes the utility a_2 derives from a_1 's offer. The acceptance strategy together with the tactics in Section 4.3.1 introduce the following assumption for the remainder of this study:

Assumption 20: Strategy profile

An agent's strategy profile is represented by i) bidding strategy which is a weighted combination of *time-dependent*, *resource-dependent* and *behavior-dependent* tactics and ii) acceptance strategy defining when to accept offers.

As tactics rely upon agents' initial offers there is a need for making another assumption. In Section 2.2.1, we have described political negotiations are *pure* negotiations in which parties make strategic demands aiming for maximum goals. In their empirical research, Lewicki & Robinson have shown that these strategic demands express themselves in among others "(...) *misrepresentation of position to an opponent*" [95, p. 666-667]. This implies making high^[4] opening offers to increase the likelihood of maximizing goals. Therefore, we assume:

Assumption 21: Opening strategy

An agent's initial offer is always centered around its *target values* and, in our tool, drawn from a skewed normal distribution.

It is actually very common for negotiators to begin making high initial offers, even in excess of the target value [137, p. 27]. A reasonable explanation for this is that they are motivated by position loss – implying that we are reluctant in giving up a good bargaining position – or image loss [56]. The latter because lower offers might indicate weakness.

4.4. Conclusion & summary

This chapter answers the third supporting question: what are the main components of automated negotiation environments? A general automated negotiation environment consists of three components, namely a protocol, scenario and strategies. The *protocol* describes the set of rules governing the negotiation process. This includes how agents interactions, issues are agreed upon or deadlines are passed. In a political domain, the protocol should support agent interactions in two dimensions – both multiple issues and agents – making the number of possible interactions immense. Finally, the protocol also provides a decision rule and (public) deadlines to determine if issues end in agreement (or not).

Furthermore, the *scenario* is a specific instance of the negotiation environment that describes the negotiation domain – set of all possible outcomes based upon the subject and issues under consideration – and an agent's preferences towards those outcomes. These preferences consist of target values (i.e. what is aimed for), reservation values (i.e. what is acceptable) and how to value everything in between. In theory,

^[3]Note that because of the cognitive limitations and incomplete information, agents might make mistakes and reject offers, even if they are more advantageous. See the last assumption of Section 3.2.1.

^[4]Or low, if an agent's utility decreases with a higher share of π (see Section 4.3.1 on page 27).

this valuation can be a certain ordering, but often a utility function is used. In this thesis, we assume linear additivity, implying issues are valued independently. But, of course, also non-linear functions have been suggested in literature. Moreover, discount factors to account for time in the valuation, whereas a random component is used to represent capricious behavior.

Finally, *strategies* define agent's actions for achieving its goals. Thus, it describes the proposals an agent makes over time and how they react upon opponent's offers. A bidding strategy consists of a weighted (and changing) combination of tactics, in which we distinguish three families: time, resource and behavior dependent tactics. These are weighted to match certain issues and may vary because of an opponent's actions or because time passes. An agent's strategy profile is completed by his acceptance strategy. This is modeled in such a way that offers are accepted if they are more advantageous than offers an agent would (theoretically) make himself if he were the initial proposer. Though, because of the cognitive limitations and unpredictable behavior, agents might make mistakes and rejecting offers, even if they are more advantageous.

The non-cooperative nature of automated negotiation environments and behavioral assumptions of Chapter 3 impose *"rigid procedural rules on the timing of moves (...) as well as strong behavioral restrictions regarding the strategies players employ"* [171, p. 36-37]. For example, in [66, p. 15-27], many other tactics of real negotiations are characterized, among which lying, threatening and insisting. Additionally, protocols introduced in this chapter do not support argumentation: offers are solely evaluated based upon the derived value. Finally, although utility functions are widely applied in economic literature, there is also much criticism. In a political context, for example, governments not always maximize their utility but rather minimize their regret, as loss of face or historic events may be present in international negotiations (see [124]).

So, although quite advanced, the automated negotiation components described in this chapter still only capture the basic elements of negotiation procedures and behavior. In our opinion, this is adequate for analyzing (pre-)coalition arrangements in multi-issue negotiation settings as the formation of coalitions is most likely not affected by one specific tactic at a distinct moment in time. Rather, they are the result of various offers over time. However, for examining distinct tactics at specific moments in time techniques in this chapter may not be adequate.

5.

A Two-Stage Automated Coalition Formation Model

My mom always said: life was like a box of chocolates. You never know what you're gonna get.
Forrest Gump, in Forrest Gump (1994)

This chapter discusses the designed simulation tool and thereby answers the fourth supporting question: how can the formation of pre-coalitions in international political negotiations be modeled using automated negotiation? In return, this tool is used to generate coalitional arrangements and evaluate their sufficiency. The next section motivates our choice for using simulation models based upon automated negotiation environments. Then, most essential model requirements are described. Finally, we define the formal model, describe its core logic, code structure and conclude with an example and verification results. For consistency, in this chapter we refer to *tool* if we indicate the entire application – as displayed in Figure 5.1. With *model* we indicate the automated negotiation environment complemented with the pre-negotiation phase. Lastly, with a *run* we imply a single simulation of the model as displayed in Figure 5.2.

5.1. Methodology: why use a simulation model?

Answering the main research question entails analyzing vast amounts of possible strategic choices – i.e. pre-negotiation coalition partners – in a highly complex, capricious and unpredictable environment. In general, there are three approaches to get enough data to evaluate these choices sufficiently, namely by means of i) literature review surveying existing researches; ii) empirical research and iii) simulation.

In Section 2.2, we stated that performing academic research in international negotiations is hard, as they often happen behind closed doors [108, p. 23]. In addition, international political negotiations take decades, making thorough empirical research relatively scarce. Although some theories are identified trying to prescribe the strategic choice of choosing (pre-)coalition partners, they lack applicability in different, less conflictual domains. Moreover, to improve generality, game theory has provided many models predicting which coalitions will form and how they should divide their winnings among its members. Unfortunately, these methods rely upon strict assumptions. It is, therefore, unlikely that data derived from existing political or game theoretic studies are sufficient to answer our research question.

Performing *empirical research* would perhaps have provided this data, but is infeasible considering limited time for completing this thesis and difficulty of getting access to international political negotiations. Moreover, there are infinitely possible negotiation setups, strategic choices and uncertain variables. It is unrealistic – considering the time and limited resources – that we could observe them all. Of course, the scope of our research could be altered in such a way that it might become empirical achievable. However, this will have negative consequences for our research questions.

Because of limitations of both (current) academic and empirical research, *simulation* becomes an obvious choice. Simulation is extremely popular and widely used for studying complex environments. The advantages for our research are as follows (based upon [91, p. 114-115]):

- Political negotiations are so complex they cannot be represented by deterministic mathematical models and evaluated analytically. Thus, simulation is often the only way of gathering data;
- Simulation models easily adapt and compare alternative negotiation domains or different negotiation protocols within a single domain;
- Simulation studies allow us to have a better control over the experimental setup, something that is not possible when performing empirical research;
- As international political negotiations often take decades, simulation allows us to study this environment in compressed time over and over repeatedly.

Unfortunately, simulation is not without drawbacks [91, p. 115]:

- Each run of stochastic simulations only generates estimates of model's true value based upon certain input parameters. Therefore, simulation is less useful for optimization – as it most likely not yields in "true" optimal outcomes – if deterministic models of international political negotiations exist;
- Simulation models are time-consuming to develop^[1];
- The vast amount of data gathered by simulation studies is often interpreted with greater confidence than is justified by the results. Especially if one considers the (strong) assumptions made beforehand.

Deduced from the literature reviews in the previous chapters, we conclude that sufficient (pre-)coalitions in international political negotiations cannot be found analytically, justifying the use of a simulation model. To mitigate the disadvantages, the model is developed under concrete requirements (see Section 5.2). Finally, we carefully deducted and documented all our assumptions throughout this thesis^[2].

Of course, in order to perform simulations, we need a *model*. To create this, we have mainly taken an *automated negotiation (AN)* perspective. This field has its roots in MAS and AI, and is ultimately interested in creating computer programs that negotiate on our behalf [49, p. 2]. For doing so, intelligent agents are created negotiating in a dedicated environment and able to learn from, and interact with each other.

Although this field has considerable overlap with agent-based modeling (ABM), the dedicated negotiation environment is one of the reasons to prefer AN. In contrary to ABM, AN already provides techniques, tools and strategies solely devoted to negotiation, potentially saving development time. Moreover, ABM is created to gain insight into collective behavior of agents, who do not necessarily have to be intelligent. In ABM, agents have to obey clear rules, rather than finding practical solutions [116, p. 480-482]. In contrary, AN agents continuously adapt to their environments. The latter resembles political negotiations best, as negotiators continuously adjust their strategies to adapt to new situations. Moreover, AN always takes a self-interested view, in which the agents may cooperate, but only when it is beneficial to them [84, 85, 150, 159, 171]. Again, we claim that this resembles international political negotiations better (see Section 2.2.2).

5.2. Requirements: do we need to model everything?

As we develop a model to analyze (pre-)coalition formation, there is a need to structure the design process in more detail. In Appendix D we have defined concrete *software requirements*, derived from our research scope, objective and various assumptions. Requirements are defined as "(...) a property that must be exhibited by something in order to solve [analyze] some problem in the real world" [24, p. 1-1-1-2] and thus constraints the final model so that it is able to sufficiently examine the main research question, given limited development time. As dozens of requirements are formulated, in this section we only give a brief overview of requirements that *must* be achieved. The model must be able to:

- Facilitate multiple negotiation domains, representing (real) international political negotiations;
- Create pre-coalitions based upon agent's ideological distances and power positions;
- Provide a negotiation environment in which multiple self-interested agents – whether or not member of a coalition – can negotiate with each other over multiple subjects and issues;
- Incorporate agents with divergent strategies and characters, closely resembling political negotiators;

^[1]As we have experienced.

^[2]For a complete overview of assumptions, see Appendix C.

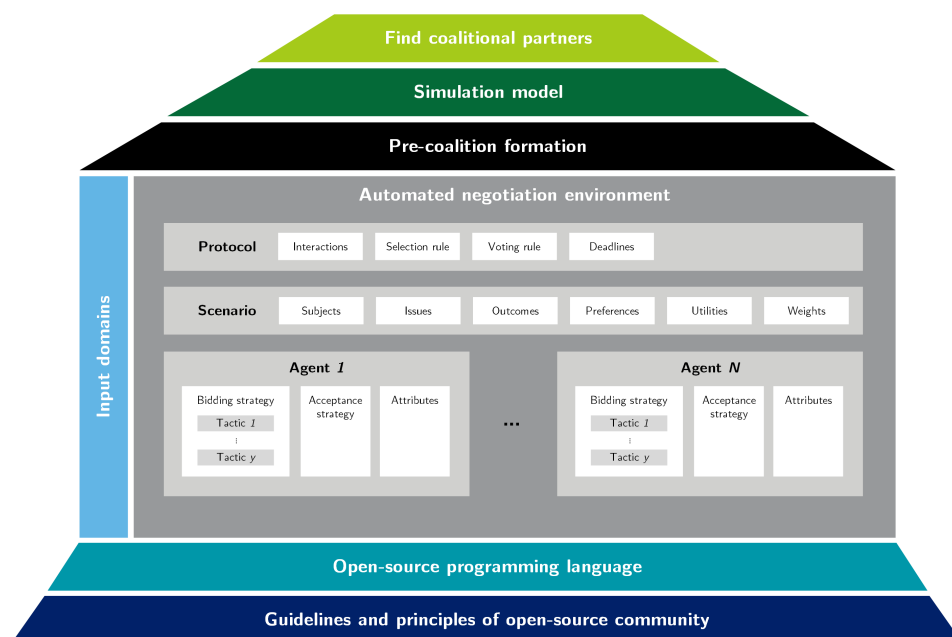


Figure 5.1: Conceptual framework of the simulation tool.

- Generate coalition arrangements computationally by means of simulation and evaluate their sufficiency in order to find sufficient ex ante coalitions;
- Investigate influence of the pre-coalition, negotiation domain and setting on the negotiation outcomes.

Moreover, allowing other researchers to use it, we must also adhere to two non-functional requirements:

- The model must be build using a widely used open-source programming language (in our case Python);
- Adhering to principles and guidelines set out by that particular community (in our case PEP 8)^[3].

The requirements above lead to the conceptual framework presented in Figure 5.1.

5.3. Core model logic: how does it work?

The main model consists of two stages: pre-coalition formation stage and actual negotiation stage. As we have limited the pre-negotiations stage to the formation of pre-coalitions exclusively – and therewith assuming agenda and arena setting are already resolved (see Section 2.2.5) – the negotiation stage consists the vast majority of model logic as shown in Figure 5.2. The model starts with **forming pre-coalitions** based on agents' ideological values and power positions, modeled by a short bargaining game in which agents offer partnerships to other agents, while the latter accept if the "ideological distances" are not too large.

When pre-coalitions are established (or not) the **actual negotiations** begin by entering the first bargaining round. A *round* is defined as a negotiation over a single subject. Accordingly, one subject might contain multiple issues. So, if we take the subject 'climate change' for example, two potential issues are 'reduction of CO₂ emissions' and 'increase in sustainable energy production'. So, the pre-coalitions enter the automated negotiation environment by first selecting a subject from a clearly defined agenda. Subsequently, an initial proposer coalition c_1 is chosen^[4]. This coalition proposes a single offer covering all issues (see Section 5.3.3). Then, coalition c_1 selects an opponent coalition, c_2 , to whom the offer is made. In return, c_2 can accept or reject the offer. Only when he accepts, both coalitions merge together to form a tentative grand coalition for this particular subject. After merging, the proposing coalition continues making offers to every other coalition of his choice. If the final tentative coalition that forms – if any – is sufficiently large enough to surpass the voting rule, the first round of negotiations end. If not, an initial proposing coalition is chosen

^[3]For the style guide, see <https://www.python.org/dev/peps/pep-0008/#id9>.

^[4]Note that "model-wise" a coalition might contain one agent, which contradicts to our natural interpretation of a coalition.

again and the above steps are repeated, in which every step^[5] resembles another opportunity in which negotiators actually meet each other – i.e. conferences and summits. As soon as two coalitions are merged, it is assumed they work together in the next step until one of them – or agents within the coalition – decide cooperating is not fruitful any more (see Section 5.3.4).

Negotiations on a single subject continue until its public deadline is overdue. If this is the case, the entire negotiation ends in no agreement, even if other subjects have not been discussed. However, if a coalition is sufficiently large enough, agreement is made and agents enter the next round, discussing new subjects. This is repeated until agreements are made upon all subsequent issues or if one of the deadlines is passed. Thus, this procedure resembles the sequential simultaneous implementation of subjects, whereas issues within each subject are dealt with by "package-deals" (as described in Section 4.1.2). Finally, negotiations also have an overall deadline. Similar to, for example, the hard Brexit deadline, in which UK is due to leave the EU on 29th of March 2019. So, negotiations also end in no agreement if the overall deadline is not met.

Some of the functions and decisions in the above high-over model logic need a more detailed description, as they are less obvious than they appear at first glance. These functions are highlighted in green in Figure 5.2 and explained more thoroughly in the following subsections.

5.3.1. From ideological proximity models to pre-coalitions

In Section 3.1.3 we have discussed how ideological proximity models are used to predict which coalitions might form given the ideological differences of agents. From them, probabilistic models are derived in which the likelihood of a coalition increases as ideological distances decreases [27, p. 1150-1151]. The flowchart in Figure 5.4 presents the algorithm implemented in our model to form pre-coalitions. Before explaining the algorithm, we first want to highlight the transformation from proximity models to cumulative distribution functions (CDF). Consider a setting with agents a_1 to a_5 differing in three-dimensional ideological values: cultural (c), economical (e) and religious (r) as displayed in the top-left corner of Figure 5.3. By taking the Euclidean distance^[6] of every data point in this model, the distance matrix is obtained in the top-right corner of Figure 5.3. Subsequently, as a measure of proximity, we transform this matrix by applying $f(x_i) = 1/x_i$ to every point in the distance matrix. Then, by standardizing this matrix such that every row sums to 1, the probability density function (PDF) is obtained. The final step of transforming the PDF to CDF is evident.

In Figure 5.5, the resulting CDF is used in the following way: first agent i is chosen according some selection procedure ζ . In turn, i chooses a potential coalition partner, agent $i+1$, and offers him a tentative agreement. The latter accepts this agreement if $CDF_{i+1}(g)$ returns i – with $g \sim Rand(0,1)$. If not, agent i is allowed to repeat the above steps ψ times, which more or less resembles the number of bi- or multilateral meetings negotiators have prior to summits. Of course, every agent in the model is allowed, but not obliged, to form pre-coalitions, so after agent i reaches ψ the next agent is chosen. Finally, the target coalition c_T is the pre-coalition of interest. So, no matter what happens during this phase, we assume that the target coalition always emerges as we are interested in it's sufficiency.

5.3.2. Selecting agents and coalitions

Both the previous subsection and high-level description of the core model logic in Section 5.3 subscribe the need of a (realistic) selection function ζ . In some settings, the first proposal is not necessarily made by the most powerful agent or coalition, but simply by the one *randomly* speaking first. However, in most political negotiations this is not the case as more powerful agents are most likely to be considered first. We have assumed *power* is derived from how decisive someone is given the decision rule in place. Decisiveness is then calculated by either the Shapely-Shubik or Banzhaf index, depending on the bargaining setting (see Section 3.1.2). Although assuming power-based selection procedures seem legitimate in international political negotiations, it does not tell the full story. When more powerful parties dominate negotiations one after another, weak actors might feel distressed and try to claim the podium as well (e.g. by seeking media attention). To model this, we have followed Herings and Predtetchinski by creating a *Markov Process* in which the initial transition matrix is still based on power, but updates by the number of times agents are selected as initial proposer [64]. Our tool allows users to select either of the three selection functions. This, to examine its effect and to allow choosing the most suitable procedure for specific negotiation settings.

^[5] Denoted with t (time) in the model.

^[6] We have chosen for the Euclidean distance, as this is by far the most used distance measure. However, theoretically every distance measure known to Python can be used.

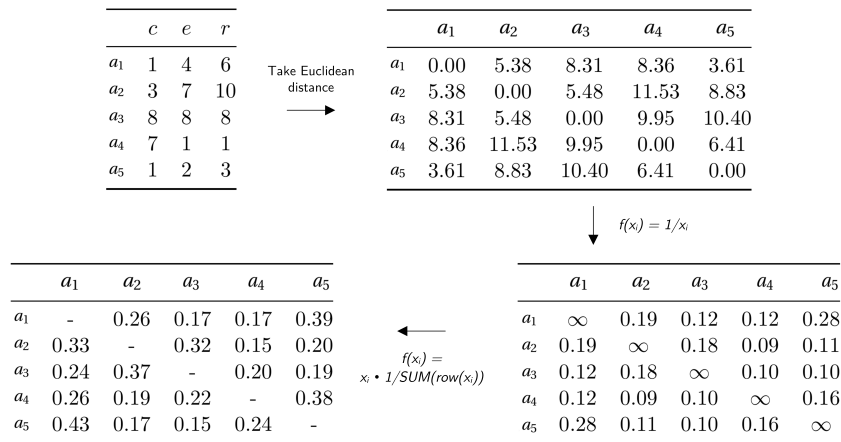


Figure 5.3: Transforming proximity model to probability density function (PDF).

5.3.3. Making offers in coalitions with divergent strategies and tactics

When working together, making new offers is a difficult process, both in practice and theory. Not every actor within a coalition can make offers to its will, not even if they are more powerful compared to others. Although in reality many factors play a role – e.g. parties who initiated the coalition might be more powerful than actors who recently joined – we limited ourselves to how much a party contributes to the overall decisiveness of the coalition. In Section 3.1.2, we have already seen that other researchers have tried to model this power asymmetry *within* coalitions in terms of *power excess* [173, p. 171-174]. After defining an agent’s power excess, the overall procedure of generating coalitional offers is straightforward. Put simply, for a single subject, all coalition members generate offers according to their strategy profiles, after which a combined offer is created by weighting all individual offers^[7]. This procedure is displayed in Figure 5.6.

Let us now explain how each coalition member *individually* generates offers. An integral part of negotiation strategy is making concessions. In fact, without them, negotiations would not exist [94]. Thus, individual bidding strategies vary in different levels of concession behavior. In Section 2.2.7, we have shown classic strategies are distinguished in two directions: cooperative and competitive. As negotiations are highly dynamic, these two directions are rather limited as renewed positions of negotiators might change another party’s behavior [94]. Therefore, Baarslag argues negotiators deploy a mixture of the two directions [11, p. 150-152] (see Figure 5.7). For example, an agent may be extremely competitive if he faces (strong) competitive opponents, while being cooperative with opponents open for cooperation. Thus, he *matches* his negotiation strategy. Subsequently, agents may cooperate with both sincere competitors as well as co-operators and, therefore, they *concede*. Following the same rationale, we can also have agents who always *compete* or cooperate with competitors while exploiting cooperative opponents (*inverters*).

In our model, we adhere to this extended classification. However, we limit the mixtures between the two directions. First, every agent has its own strategy profiles for the four classifications presented in Figure 5.7. This profile resembles a weighted combination of tactics as described in Section 4.3. Moreover, agents have initial preferences to either of the four profiles, which are defined by probability distributions. Finally, they have preferences how to change from one profile to another based upon opponent behavior, also described by probability distributions. For simplicity, we assume uniform probabilities. To align this with the theory in Section 4.3, every agent *i* has four $m \times n$ strategy profiles $\Omega_{i \rightarrow i+1}$ – each for every classification in Figure 5.7 – defining his bidding strategy to *i* + 1. Going from one profile to another is then described by:

$$\Omega_{i \rightarrow i+1}^{t+1} = f(\Omega_{i \rightarrow i+1}^t)$$

In which *f* is given by a uniform distribution. Finally, we need to describe how bidding strategies – consisting of multiple tactics – result in one bid for a single issue. This procedure is displayed in Figure 5.8. Given an opponent coalition *k* + 1, every agent in coalition *k* generates a bid for every agent in coalition *k* + 1 using tactic τ_n . As *k* + 1 is considered one entity, every agent in *k* should return one bid – *bid*_{*n*} – for a single issue. If agent *i* in coalition *k* prefers high values (so the *target* value is high), he takes the maximum

^[7]This procedure is among others advocated in [159, p. 31-32], [43, p. 174-176] and [132, p. 397-407]

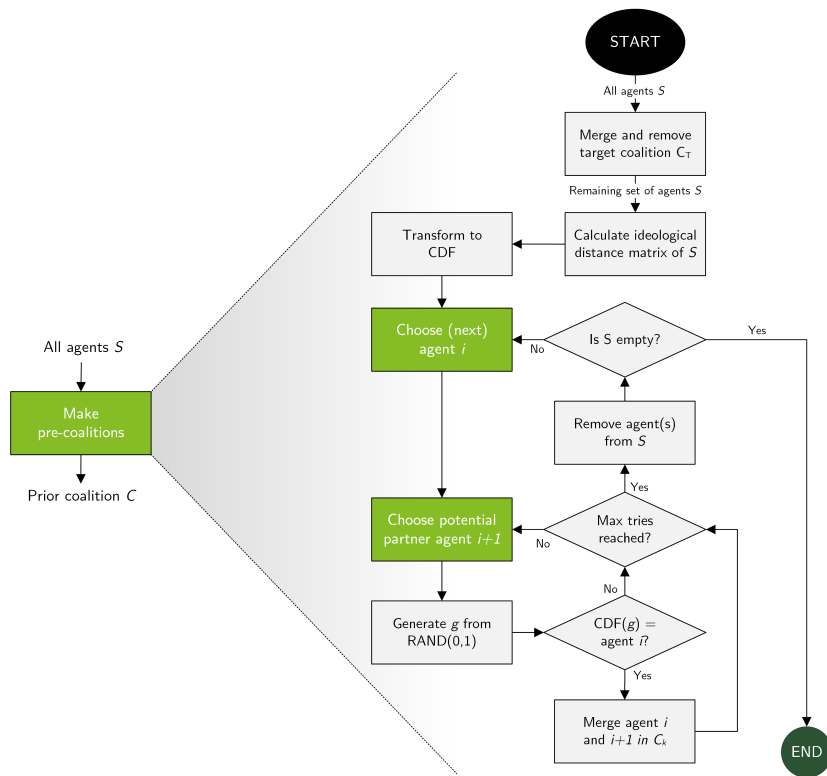


Figure 5.4: Low-level flowchart of the function forming pre-coalitions.

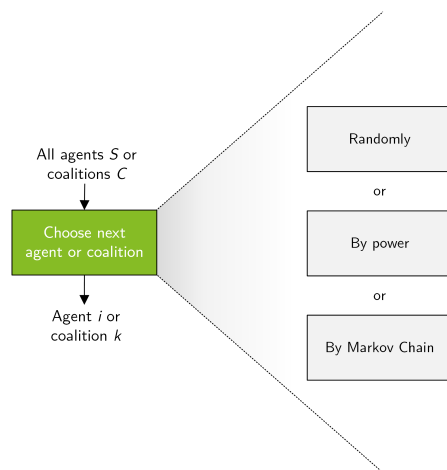


Figure 5.5: Options to select (next) agent or coalition.

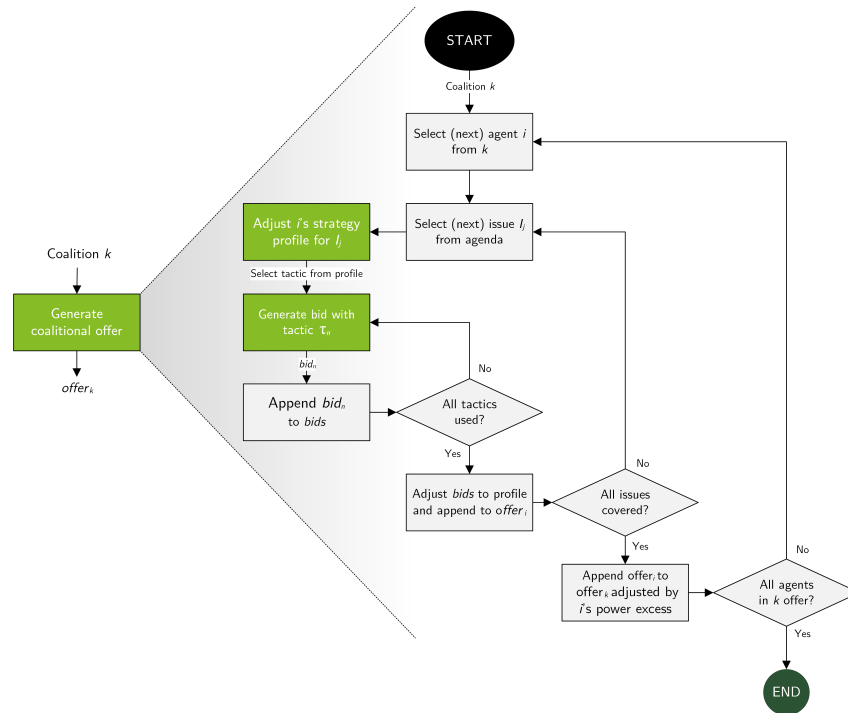


Figure 5.6: High-level flowchart of the function to generate offers in a coalition.

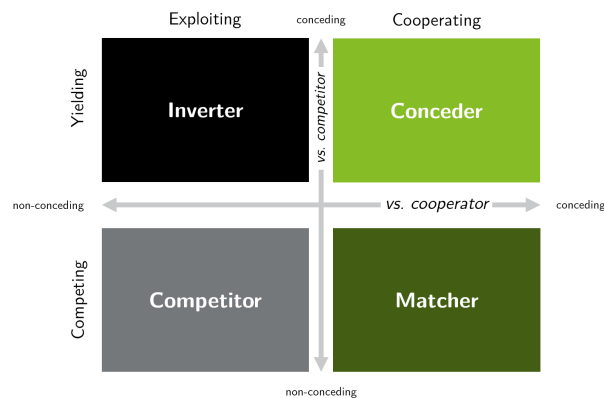


Figure 5.7: Conceding behavior against purely cooperative or competitive opponents (adopted from [11, p. 151]).

of all individual bids and returns bid_n . Of course, the opposite holds when aiming for low values. In return, bid_n is then aggregated to a single *coalitional* offer, as also coalition k is considered as one entity. Thus, only one offer is proposed from k to $k + 1$.

5.3.4. From utility functions to accepting offers and leaving coalitions

In Section 3.2 we defined coalitions (and its members) have four options when facing an offer. The coalition might *accept* the offer and *join* their coalitions together. After doing so, agents within the coalition can *stay* or decide to *leave*. Either of those actions depend on how agents evaluate opponent's offers. In line with most economic theory, we assume that this evaluation is based upon an agent's utility function. So before describing the procedure that leads to one of the previous actions, we first describe the utility functions in our model.

In line with our assumption in Section 4.2, our model presumes the utility function $U(x, t)$ consists of a deterministic $u(x)$ and stochastic component ϵ . As utility is adjusted for the time in which the offer is made,

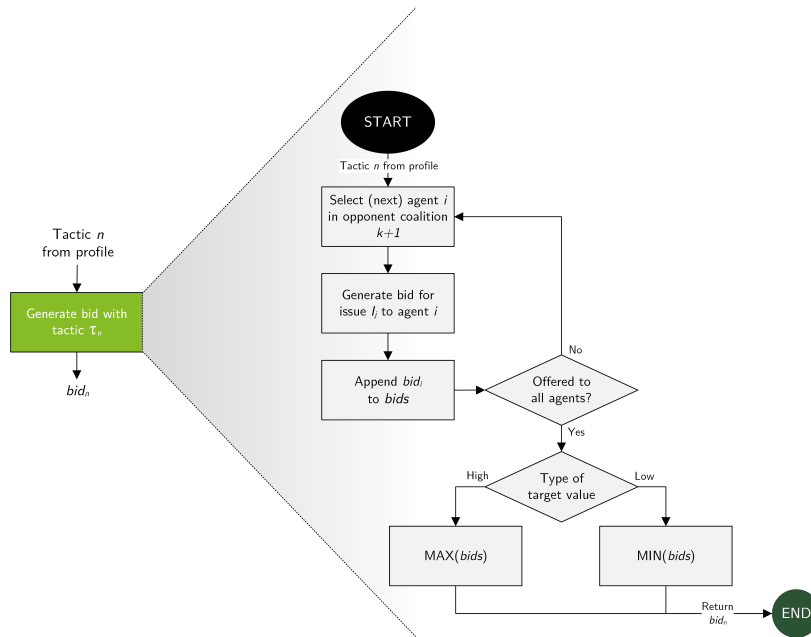


Figure 5.8: Low-level flowchart of the function to generate bids from an agent's tactic.

the final function for agent i results in $U_i(x, t) = \delta_i^t \cdot u_i(x) + \epsilon$. Although the user of our tool is able to choose utility functions, for the remaining of this study we assume that $u_i(x)$ is described by the Euclidean distance and $\epsilon \sim N(0, \sigma)$. The origin of using distance utility functions dates back to Wold [180], but evolved ever since. In short, it selects a reference point (or vector) – denoted with x_R – and measures the (Euclidean) distance to the upper bound of each element in the choice set [6, p. 65]. We justify the use of distance utility functions as the choice set represents the proposed offer x_k , whereas the reference point x_R can be interpreted as an agent's goals to which he wants to stay as close as possible. Therefore, in a negotiation settings, this method is intuitively appealing in contrary to, for example, money-metric utilities [6, p. 65-66].

Unfortunately, using Euclidean distances also comes at a cost, especially if x is high dimensional (thus one subject contains many issues)^[8]. A well-known problem occurs when distances for two nearly identical high-dimensional offers are calculated. To marginalize its effects, while initializing the model, one can apply two countermeasures. First, the user can split one high-dimensional subject in a series of lower dimensional subjects. Though, this decreases the ability to make package-deals and does not (always) reflect reality (see Section 2.2.5). Secondly, the issues within a subject should have comparable outcome spaces so that changes in one issue would equally impact the overall offer.

Now we have introduced which utility functions are implemented, we explain the "acceptance" procedure^[9] more closely. When coalition $k+1$ faces an offer from k , it will first produce a hypothetical offer itself. This offer serves as a baseline and can be interpreted as follows: given the elapsed time, historic offers and deadlines, what offer will coalition $k+1$ deem acceptable? After doing so, both utility from k 's offer as well as from the hypothetical offer is determined. This is done by defining *collective utility*, simply summing all individual utilities weighted for their power excess. If the coalitional utility derived from opponent's offer is higher than from $k+1$'s own offer – thus, $U_{k+1}(x_k) > U_{k+1}(x_{k+1})$ – coalition $k+1$ will join k .

After merging, every member agent should decide for himself whether the newly formed coalition is acceptable or if it is better to leave. So, they first evaluate their individual utility for both offers. If the utility for his coalitional offer is higher than for the final accepted offer, it is *individual rational* to leave the new coalition (see Section 3.2.1). However, as discussed, negotiators are incapable of always acting

^[8]Famously referred to as the "curse of dimensionality", introduced by Richard Bellman in 1957 [18].

^[9]Although we refer to this function as the "acceptance" procedure, it describes all actions agents might take after facing an offer, so *accepting*, *leaving*, *staying* or *joining*.

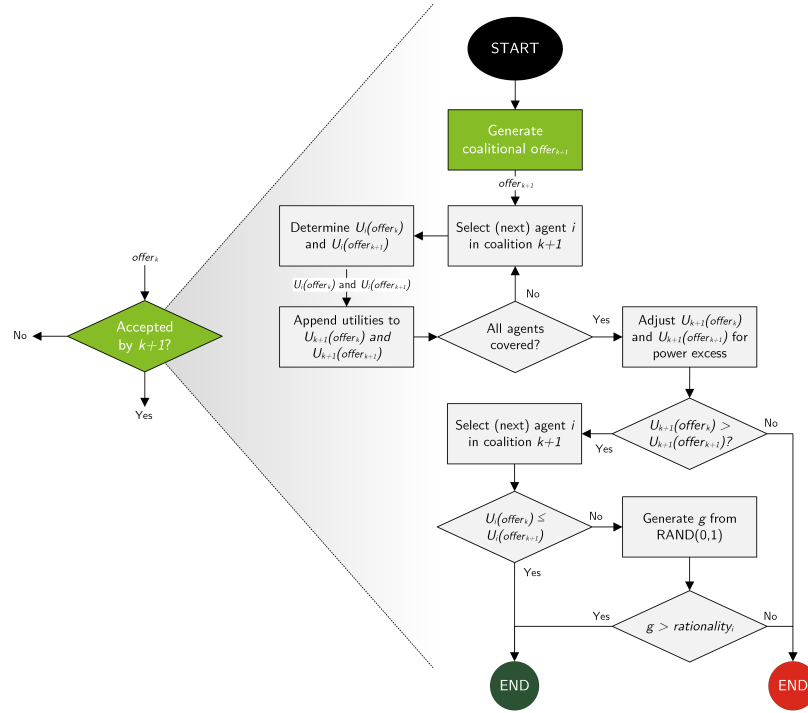


Figure 5.9: Low-level flowchart of accepting offers from an opponent coalition.

rational and, therefore, make mistakes. In our model, we have implemented a rationality parameter which provides an artificial threshold representing how rational agents are. Only if this threshold is surpassed – i.e. $g \sim \text{Rand}(0, 1)$ is greater than $\rho \in [0, 1]$ – then agents actually leave the new coalition.

5.4. Formal model definition and initialization

In the previous sections we have described core model logic and thereby introduced many input parameters. This section defines the formal model consisting of both the pre-coalition stage and automated negotiation environment, which is partly based on existing multilateral bargaining models with the alternating offers protocol (i.e. [49, 141, 154, 161, 169]). Our negotiation model can be described by a tuple $(A, P, B, R, \Delta, D, M, I, S, O, W, T, R, u, U, \theta, \Theta, \xi, \chi, \tau, \alpha, \Omega, \Lambda, C, \psi)$ in which:^[10]

Agent (attributes)

- $A = \{a_1, \dots, a_n\}$ denotes the set of n negotiating agents, indexed by i ;
- $P = \{p_1, \dots, p_n\}$ in which $p_i \in \mathbb{R}_{>0}$ is a value describing the different power positions of the agents;
- $B = \{\beta_1, \dots, \beta_n\}$ in which $\beta_i \in \mathbb{R}_{>0}$ describes agent i 's risk behavior while negotiating. If $\beta \rightarrow 0$ an agent will be extremely risk seeking, while for $\beta \rightarrow \infty$ he is exceptionally risk averse;
- $R = \{\rho_1, \dots, \rho_n\}$ in which $\rho_i \in [0, 1]$ describes agent i 's rationality. If $\rho \rightarrow 0$, agent i shows capricious behavior, while for $\rho \rightarrow 1$ an agent will act highly rational;
- $\Delta_i = \{\delta_1, \dots, \delta_z\}$ in which $\delta_k \in \mathbb{N}$ denotes agent i 's private deadline for subject s_k ;
- $D = \{d_1, \dots, d_n\}$ in which $d_i \in [0, 1]$ describes agent i 's overall disagreement utility^[11];
- $M = \{\mu_1, \dots, \mu_n\}$ in which $\mu_i \in [0, 1]$ describes agent i 's discount factor;

^[10]A **bold** parameter indicates that it is either a vector or a matrix.

^[11]The disagreement utility can be derived from agent i 's overall satisfaction for the *status quo* (see [113, p. 10-12]).

Scenario^[12]

- $I = \{I_1, \dots, I_m\}$ is a vector containing all m issues I , indexed by j ;
- $S = \{s_1, \dots, s_z\}$ describes a set of z subjects, indexed by k . Note that each subject s_k is a vector containing an unique subset of issues $s_k \subseteq I$. Note too that in our model, the order of subjects in s also accounts for the negotiation agenda;
- $O = \{o_1, \dots, o_z\}$ is a set describing possible outcomes of z subjects, indexed by k . Each outcome is a matrix which columns depend on the number of issues in each subject^[13] and the rows depend on the outcome space of each issue j . Note that an offer within the outcome space is denoted with \mathbf{x} ;
- $W_i = \{w_1^i, \dots, w_m^i\}$ in which $w_j^i \in \mathbb{R}_{>0}$ is a value describing agent i 's weight for issue j ;
- $T_i = \{t_1^i, \dots, t_z^i\}$ describes agent i 's target value or goal for each outcome o_k . Thus, every target value is a vector of which it's length equals $|s_k|$;
- $R_i = \{r_1^i, \dots, r_z^i\}$ describes agent i 's reservation value for any of the outcomes. Again, each reservation value is a vector of which it's length equals $|s_k|$;
- $u_k^i: o_k \times r_k^i \times t_k^i \mapsto \mathbb{R}$ describes agent i 's utility for subject s_k as a function of the outcome o_k , target value t_k^i and reservation value r_k^i ;
- $U_i: u_k^i \times W_i \mapsto [0, 1]$ denotes a party's cumulative utility for all subjects;

Protocol

- $\theta = \{\theta_1, \dots, \theta_z\}$ in which $\theta_k \in \mathbb{N}$ describes the deadline for each subject k ;
- $\Theta \in \mathbb{N}$ is the overall deadline for all subjects;
- $\xi: A \times P \mapsto a_i$ denotes the selection function for choosing agent i to be the (initial) proposer;
- $\chi: A \times P \mapsto [\text{True}, \text{False}]$ describes the voting rule for the negotiation model;

Strategies

- $\tau = \{\tau_1, \dots, \tau_y\}$ is a set of available negotiation tactics;
- $\alpha_i: U_i \times D \mapsto [\text{True}, \text{False}]$ is agent i 's acceptance function;
- $\Omega_i = \begin{bmatrix} \omega_{11} & \dots & \omega_{1z} \\ \vdots & \ddots & \vdots \\ \omega_{y1} & \dots & \omega_{yz} \end{bmatrix}$ is a matrix describing i 's strategy profile for subject k based upon tactics τ ;

Pre-coalition formation

- $\Lambda_i = \{\lambda_1^i, \dots, \lambda_j^i\}$ describes agent i 's ideological or intrinsic characteristics;
- c_T is the target pre-coalition of agents as a subset of A . Hence, $c_T \subset A$;
- ψ denotes the amount of tries an agent is allowed to do to form a pre-coalition.

Only the final three parameters are used in the pre-coalition stage, whereas the remainder is used to define the automated negotiation environment. By adding two more parameters – N and σ – a simulation tool is described with the number of simulations and uncertainty factor respectively. Naturally, the simulation tool consists of two or more instances of the two-stage model defined above. While initializing the different instances, we assume that all unknown input parameters X ^[14] are normally distributed with uncertainty σ . Thus, every instance of the two-stage model is a realization of $\tilde{X} \sim N(X, \sigma X)$ with $\sigma \in \mathbb{R}_{>0}$. In addition, the initial coalitional state of each instance consists of n , single agent, coalitions – ranging from $\{c_1, \dots, c_n\}$.

5.5. Implementing the formal model

This section discusses the implemented two-stage coalition formation model by motivating our choice for using Python, briefly discussing code structure and verifying a simple negotiation setting.

^[12]Of which the tuple (I, S, O) describes the negotiation domain and (W, T, R, u, U) the preferences of every agent i within that domain (see Section 4.1.)

^[13]Such that $\sum_{k=1}^z |s_k| = m$.

^[14]In the model, we consider all agents' preferences, characteristics and strategy profiles to be uncertain.

5.5.1. Going the Pythonic way

There is many tooling available able to implement the two-stage model, ranging from generic programming languages (i.e. Python and Java) to dedicated negotiation environments and agent-based modeling software in between (e.g. NetLogo). The dedicated GENIUS environment is an open architecture allowing researchers to design and evaluate agent strategies [97, p. 51-53]. It is developed for bilateral negotiations and, therefore, vast majority of protocols is not suitable for political negotiations. Moreover, it does not support the strategic choice of forming (pre-)coalitions. The less dedicated NetLogo is a programmable environment used for modeling complex systems with hundreds of agents over time [179]. NetLogo's major advantage is that it is build "...to have a low threshold for beginners" [170, p. 19] and is therefore widely used in social sciences. As a consequence, NetLogo demands its users to follow the concepts upon which the environment is build. Therefore, it lacks adaptability. In addition, it is not created for large numbers of simulations with lots of interactions [54, p. 355-356]. Creating our two-stage model in NetLogo can therefore become cumbersome.

So, for our research, we propose Python: a high-level language for general-purpose programming. Python is increasingly popular among most data scientists, both within academics and businesses. More specifically, in recent years it has seen a tremendous increase in popularity – partly due to the rise of artificial intelligence – and is almost certain going to follow the same upwards trajectory [28]. Moreover, it's community has produced thousands packages with functionalities ranging from visualization to machine learning. Thus, Python resembles our objective to create a reusable negotiation tool build upon many well-known packages.

5.5.2. Tool structure: what have we build?

To implement the formal model above, we have created our own Python package, called "negopol". This package consists of the following modules, again dividend in smaller sub-packages (or *directories*):

dir: *negotiation*

Contains the **automated negotiation environment** in which:

- *__init__.py* initializes the negotiation environment by importing relevant packages and modules;
- *domain.py* contains classes who together define the main components of negotiation domains;
- *protocol.py* defines procedures for interaction between entities and rules deciding if negotiations succeed;
- *round.py* executes a single negotiation round, thus covering one subject;
- *strategies.py* defines tactics entities are allowed to use;

dir: *entities*

Contains the **negotiation entities** in which:

- *__init__.py* initializes entities by importing relevant packages and other dependencies;
- *agents.py* contains the class defining agents, its attributes and methods;
- *coalition.py* contains the class defining coalitions, its attributes and methods;

dir: *simulation*

Defines the **simulation model** in which:

- *__init__.py* initializes simulations by setting up the automated negotiation environment and it's entities;
- *model.py* creates two-stage model by forming pre-coalitions and initiating instances of rounds;
- *reporting.py* aggregates simulation results and makes them available for reporting and visualization;
- *scenario.py*^a generating different negotiation settings.

^aThe latter is explained in more detail in Chapter 6 and Appendix G.

For a detailed oversight of relationships between the automated negotiation environment and its entities, see Appendix E. The two-stage model can be initiated from a *main.py* file in which input domains and other relevant parameters are specified before creating instances of the *negopol.Model* class. This instance has dedicated class methods – i.e. *Model.simulate(*args)* and *Model.scenario(*args)* – to perform simulation experiments. Their results are only presented when calling methods such as *Model.outcome* and *Model.utility*.

5.5.3. Look-and-feel of the tool: an example

This section gives an impression of how the two-stage model works. In an example run, eight agents $\{a_1, \dots, a_8\}$ negotiate over one subject containing three issues (i.e. x , y and z). The target pre-coalition C_T consists of three agents, namely $\{a_6, a_7, a_8\}$. Agreement is reached by weighted majority voting. For an oversight of other relevant parameters, see Table 5.1 below^[15].

Table 5.1: The example run's most important input parameters.

	a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8
Power (p_i)	1	2	3	1	1	3	1	3
Risk behavior (β_i)	4	2	0.5	0.1	0.5	1	1	1
Disagreement utility (d_i)	0.25	0.4	0.1	0.25	0	0.25	0.15	0.45
Reservation value (r_i)	[40,60,20]	[40,60,5]	[90,20,50]	[50,50,50]	[80,30,10]	[30,30,40]	[80,80,15]	[1,1,20]
Target value (t_i)	[30,30,15]	[90,90,10]	[80,10,40]	[40,1,10]	[90,90,5]	[80,50,5]	[60,30,5]	[40,60,10]

As displayed in the top-left corner of Figure 5.10, the model initializes at $t = -1$, such that the first step (towards $t = 0$) equals the pre-negotiation phase. In this step, we observe the formation of C_T , without the formation of any other pre-coalition^[16]. This, because $\phi = 1$, which does not imply other pre-coalitions will not form, but only the chance decreases, especially if agents' ideological values differ heavily. After the pre-negotiation phase, four negotiation steps go by before a tentative agreement is reached between $\{a_2, a_5\}$. This cooperation could be expected from both agent's target values as they are relatively similar. Perhaps their difference in risk behavior or rather low power positions have caused not to cooperate more quickly. From $t = 4$ and onwards, two coalitions and three "free-riding" agents continue to negotiate. Until $t = 11$ no more agreement is made before $\{a_3, a_4\}$ decide to work together. At the next step, a_1 makes a fruitful offer to this new coalition as well. Although both steps may contradict at first glance – as these three agents' target values are relatively distinct – we see that their target values are *lower* than the reservation values. This implies these agents all aim for low values, in return implying they have declining utility functions. So, from $t = 12$ and onwards, only three coalitions are left negotiating with each-other, namely $\{a_1, a_3, a_4\}$, $\{a_2, a_5\}$ and target coalition $\{a_6, a_7, a_8\}$ with aggregated power $p = 5$, $p = 3$ and $p = 7$ respectively. At $t = 19$, C_T makes a good offer to $\{a_1, a_3, a_4\}$ after which the joint coalition becomes winning. As they consider themselves sufficiently large enough – as $p = 12$ is way more than $p = 8$ – the negotiation ends with outcomes [63,33,8] for the three issues respectively. The utility of agents then result in [0.63,0.39,0.68,0.43,0.36,0.69,0.89,0.82]. Although it looks C_T outperforms other agents, they may also have luck this simulation run.

Table 5.2: Part of the winning coalition (or not) in 100 simulation experiments.

	a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8
In winning (#)	52	54	52	18	29	65	15	73
In winning (%)	69.2	72.4	69.2	24.8	38.9	86.5	20.5	97.3

Therefore, to complete this example, 100 simulation experiments were conducted of which 75% resulted in an agreement. When considering the utilities displayed in Figure 5.11 (right), we notice the above run was actually an outlier. This, because a_7 has a median utility of 0.46, significantly lower than a_6 and a_8 . Moreover, as their reservation and target values distinct too much, from Table 5.2 we observe that a_7 actually leaves C_T the vast majority of simulation runs as they are part of the winning coalition approximately 20% of the times, compared to 86% and 97% for a_6 and a_8 respectively. In Figure 5.11 (left) the actual outcomes of the negotiations are also presented^[17]. It can be seen that both issue x and y often result high values, whereas z results in low values. This, because a_6 and a_8 are two dominant and powerful agents both aiming high for issue x and y and low for z . In addition, we observe that a_5 manages to get lucky from time to time, as this agent is rarely part of the winning coalition but has some high (tailed) outliers. The opposite holds for a_3 , if this agent is part of the winning coalition, his power is substantial enough to enforce good outcomes. However, if not, he has bad luck most of the time, resulting in low (tailed) outliers.

^[15]Input parameters not mentioned in this table are either constant or not important for this example.

^[16]In Figure 5.10, coalitions are visualized by lines getting closer to each other.

^[17]If discrete issues were used in this example, this figure becomes a histogram displaying number of times an option is chosen.

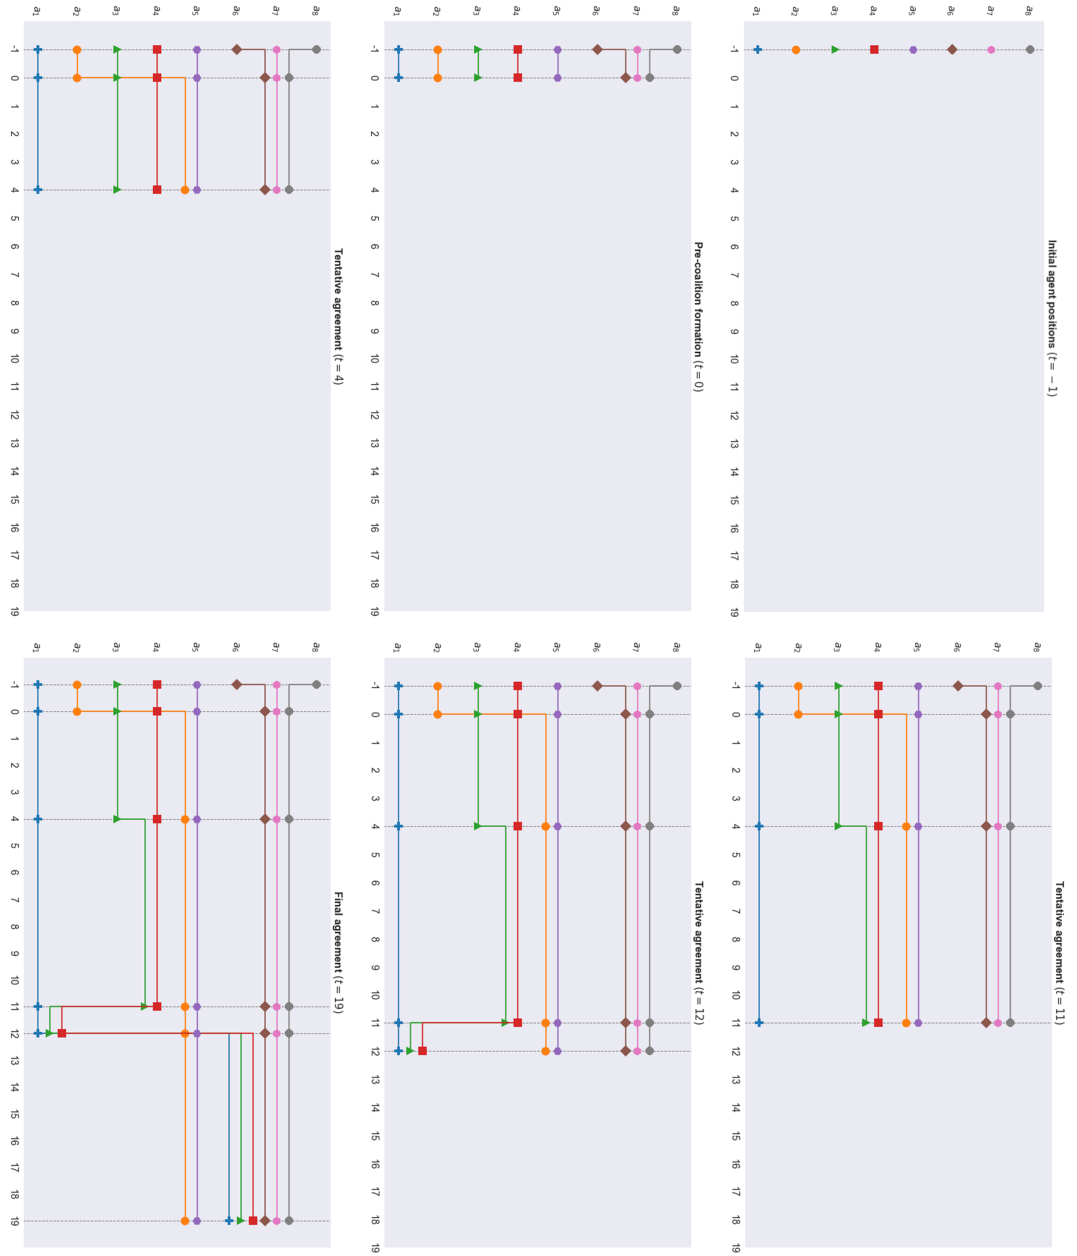


Figure 5.10: Example simulation run: different coalition formation steps in the negotiation process.




Figure 5.11: Boxplots of 100 simulations presenting negotiation outcomes (left) and derived utilities (right).




Figure 5.12: Some characteristics of 100 simulation experiments.

Finally, Figure 5.12 presents some other relevant output parameters of 100 simulations. Again, the run presented before is an outlier in terms of negotiation time as the majority of simulation runs result in agreement before $t = 11$. From Figure 5.12 (middle and left) we also observe that the final coalition is often slightly larger than necessary to become winning, both in terms of size and power.

5.6. Code quality and verification

In Appendix F, we have described a thorough evaluation of tool quality and verification. This section briefly summarizes the results. Let us first consider tool **quality**, for which it is good to remember the main purposes why the tool was developed: i) it must allow us to answer the research question; ii) it should allow others to review our work and iii) it should allow others to extend our work. A high quality tool contributes to each of these goals. But, as quality is ambiguous in many ways, we refer to tool quality by conforming to high-level requirements stated in ISO-9126^[18], including functionality, reliability, usability, efficiency, maintainability and portability (see Appendix F.1 for more detail). The *functionality* is evaluated by performing a functional requirement assessment which determines whether functional requirements are met. We show that:

- All 15/15 "Must haves" (M) functional requirements have been met;
- A considerable 8/11 "Should haves" (S) functional requirements have been met;
- A few or 3/10 "Could haves" (C) functional requirements have been met and;
- As expected, none or 0/4 of the "Won't haves" (W) functional requirement have been met.

As all "Must haves" (M) and a considerable amount of "Should haves" (S) and "Could haves" (C) functional requirements are met, we are confident that its (functional) quality is adequate enough to be used for answering the research questions. Furthermore, *non-functional* requirements are assessed by performing simulation experiments and applying Python's *Pylint* and *Radon* packages. The latter are used to determine, among others, coding standards, code complexity and maintainability. Simulations are used to test for errors (or even failures) and to determine performance and efficiency. Based on all these tests, we show that:

- All 2/2 "Must haves" (M) non-functional requirements have been met;

^[18]Or its more extensive replacement: the ISO-25010 model (see [69]).

- A considerable 3/5 "Should have" (S) functional requirements have been met;
- As expected, none or 0/1 of the "Won't have" (W) functional requirement have been met.

As all "Must have" (M) and a considerable amount of "Should have" (S) and "Could have" (C) non-functional requirements are met, we are again confident that its quality is adequate for answering our research question, and to be used or extended by others. However, there are still some improvements possible. First, maintainability can be improved by changing bad name conventions to adhere to Python's PEP8 guidelines. Secondly, usability improvements can be made by facilitating the input of negotiation domains by simple XML or JSON formats. Finally, especially for large negotiation environments and domains, there are also improvements to be made regarding tool efficiency. We show that for more than 20 agents and high-dimensional outcome spaces, the becomes infeasible as calculations take several minutes.

Testing the quality of our tool does not tell the complete story. We still need to make sure if the requirements shaping the conceptual model and tested before are correctly implemented in the computational tool in Python. This step is called **verification** – or more specifically *implementation verification* – and defined as "...assuring that the computer programming and implementation of the conceptual model is correct" [151, p. 108]. The following tests have been designed to verify our model:

- First, for every (class) method and function in our tool, we examine the input and output and verify whether or not it matches theoretical predictions from the conceptual model;
- Secondly, we setup a simple three-agent model for which we make theoretical predictions based upon the conceptual model. These predictions are then verified by comparing them to simulated results;
- Thirdly, we push our tool to the limit by entering extreme parameters for the simple three-agent model;
- Moreover, we repeat both aforementioned tests for a more complex, eight-agent model;
- Finally, variability tests have been performed to explain and verify unintended randomness in our tool.

In Appendix F, we demonstrate the verification of our tool and conclude the model passed all tests. However, similar to tool quality, we acknowledge that still some improvements can be made or additional analysis is in place. For example, we recommend to resolve an issue that occurs when entering grand coalition (i.e. all agents in negotiation environment) as target coalition C_T . Moreover, we recommend more thorough analysis of extreme risk behavior – either extremely risk averse or seeking – especially in those negotiations in which extreme risk behavior plays a dominant role for the involved parties.

5.7. Conclusion & summary

The formation of (pre-)coalitions is modeled in two stages. First, in the *pre-negotiation stage*, the formation of coalitions is considered to be a strategic choice, based upon negotiators' ideological values and power positions. This is modeled by a short bargaining game in which an agent offers partnership to another agent. The latter accepts this offer if both the ideological distance or presumed power position is not too large. When accepted, a "tentative" agreement between two (or more) agents is reached. Tentative, because in the next stage, agents can still decide to leave the coalition whenever the offers he gets are lower than his guaranteed payoff (i.e. outside option).

The second stage consists of the actual *automated negotiation environment*, in which (coalitions of) involved agents negotiate over various subjects and issues and, as a consequence, form greater (or even winning) coalitions. In this environment, the subjects are dealt with sequentially, whereas the issues allow for package-deals. Moreover, agents are considered to be the proposer according to a Markov Chain in which the initial transition matrix is based on different power positions and updated by the number of times an agent has already been the proposer. Offers are made and/or accepted by (coalitions of) agents according to their strategy profiles. Though, because of their cognitive limitations, they occasionally make mistakes.

The two environments together form the *two-stage model*, as they cover the first two stages of the international political negotiation process. This is implemented by creating a Python library named "negopol". This package is able to simulate multiple instances of the two-stage model, given certain negotiation domains and settings. It also enables users to define their own negotiation scenarios, create intelligent agents, adjust the way they interact and vary the level of uncertainty of all input parameters. Moreover, in order to evaluate the sufficiency of pre-coalitions, users can insert target coalition C_T . It is assumed this coalition always forms in the first stage, but collapse thereafter.

6.

Agent Objectives & Experimental Design

Keep your friends close, but your enemies closer.
Michael Corleone, in *The Godfather Part II* (1974)

This chapter answers the fifth supporting question: which objectives make pre-coalitions sufficient? Moreover, it discusses the experimental setup used for the simulation experiments in Chapter 7. But first, in Section 6.1, the concept of "sufficiency" is delineated more carefully. Thereafter, in Section 6.2 we explain how sufficiency is used to evaluate pre-coalitions in different negotiation domains and settings.

6.1. Defining sufficiency: when are pre-coalitions good enough?

So, how can we evaluate if Carl's decision – in the introductory holiday negotiation story – to ally with Alex was better than allying with Bob? Or, in general, how can we determine for agent i whether to join k or $k+1$, any other pre-coalition or stay alone prior to the negotiations? In this thesis, we assume that this depends on the quality of negotiation outcomes. This **quality** can be perceived collectively or for a single agent individually. As this thesis assumes political negotiators to be self-interested, we implicitly assume quality of negotiation outcomes mainly depends on personal gains. However, in the following section, we argue that even if negotiators are self-interested, occasionally they may still aim for collective or fair outcomes.

6.1.1. Maximum social welfare: utilitarian, egalitarian and Nash product

One can argue that the regimes^[1] who shaped international political negotiations for decades provide mechanisms supporting more fair outcomes. For example, the "(...) *WTO was born out of negotiations aimed at progressively reducing obstacles to trade (...) by discouraging 'unfair' practices*" [181, p. 4]. Moreover, these regimes also "forces" states to discuss and negotiate serious topics regularly. In Section 2.2.7, we stated this repeated interaction may lead to changing strategies [42], but we argue it may also gradually lead to collective objectives. Best examples of the latter are found in various conventions on climate change.

The study of fair negotiation outcomes is known as *welfare economics* [25, p. 247-248]. In contrast to individual utility discussed in Section 4.2, welfare economics focuses on maximizing collective utility functions, defined as any function $f: \mathbb{R}^n \rightarrow \mathbb{R}$. Three well-known functions f are described below (see [25, p. 251-253]):

- The **utilitarian** function computes social welfare by maximizing the sum of individual utilities. Thus, collective utility given outcome $\mathbf{x} = \{x_1, \dots, x_i, \dots, x_n\}$ is defined as $\sum_{i \in A} (u_i(x_i) - d_i)$, in which $\mathbf{x} \in O$ is the accepted offer and u_i and d_i agent i 's utility function and disagreement utility respectively. Note that in this function, the higher the average individual utility, the higher the social welfare. So, this measure may not be individually fair as "(...) an extra unit of utility awarded to the agent best off cannot be distinguished from an extra unit of utility to the agent currently worse off" [25, p. 251];

^[1]Such as the EU, UN and WTO (see Section 2.2.4).

- Secondly, the **egalitarian** function tries to encounter the utilitarian deficiency by defining social welfare as $\min_{i \in A} (u_i(x)_i - d_i)$. This implies that the social welfare is to be the utility of the agent worse off with final outcome \mathbf{x} . Thus, the egalitarian approach neglects increases in social welfare if individual utility increases for any other agent except the one currently worse off;
- Finally, the **Nash product** way of defining social welfare needs some introduction. The Nash bargaining solution, introduced by the famous mathematician John Nash in 1950 [114], provides a unique solution to simple bilateral bargaining problems under strict assumptions. For two agents i and $i+1$, the Nash bargaining solution is the unique pair (x_i^N, x_{i+1}^N) that solves the maximization problem:

$$\max_{(x_i, x_{i+1}) \in O} (u_i(x_i) - d_i) \cdot (u_{i+1}(x_{i+1}) - d_{i+1}),$$

The maximization objective is called the Nash product, which, for multilateral bargaining, can be extended to $\prod_{i \in A} (u_i(o_i) - d_i)$, see [20, p. 156-166]. This product is alternatively used to calculate social welfare. Like the utilitarian approach, Nash product increases with an overall increase in average individual utility, but more so for agents being worse off [25, p. 252].

All functions above are measures of social welfare and can be used to define the quality of the negotiation outcome collectively. As standardized utility functions are used, both the egalitarian and Nash product approach to social welfare are in $[0, 1]$. Obviously, the utilitarian method is not, but by averaging the sum of individual utilities, the average social welfare is obtained in the interval $[0, 1]$. Note that in order to calculate social welfare, theoretically we need *complete information* as all agents' utility functions, goals and disagreement levels must be known [11, p. 248]. As this is not assumed in this thesis (see Section 3.2.1), social welfare is based upon an agent's prior beliefs of his opponents' preferences.

6.1.2. Maximum self-interestedness: only personal gains?

Of course, for self-interested agents, the primary objective is maximizing their own personal gains given their individual utility functions and set of preferences. For agent i , this implies maximizing $u_i(x_i)$ given the final offer $\mathbf{x} = \{x_1, \dots, x_i, \dots, x_n\}$. However, we argue whether this is the only individual objective agents have. Below, we suggest some other possible objectives, partly adopted from literature:

- **Negotiation time** (adopted from [11, p. 116-117]) – Actors might want to trade in some individual gains if agreements are settled sooner. Of course, this is already partly reflected by their (possible) discounted utility μ_i , but might also be an individual objective on its own. Therefore, we suggest that in some cases agent i tries to minimize:

$$\frac{T}{\min(\sum_{k \in S} \delta_k^i, \sum_{k \in S} \theta_k, \Theta)}, \quad (6.1)$$

in which T is the simulated total negotiation time (in steps), δ_k^i is agent i 's private deadline for subject k , θ_k the public deadline for subject k and Θ the overall negotiation deadline;

- **Political support**^[2] – In situations with repeated interaction or to invigorate certain outcomes, negotiators may also prefer maximizing political support. Even if this political support exceeds the necessary voting rule. Thus, if we define C_w as the winning coalition, actors might want to maximize either:

$$\frac{|C_w|}{n} \quad \text{or} \quad \frac{\sum_{i \in C_w} p_i}{\sum_{i \in A} p_i}, \quad (6.2)$$

in which p_i is agent i 's (structural) power defined prior to the actual negotiation phase;

- **Likelihood of agreement** – Finally, actors might also want to secure their personal gains by increasing the likelihood of reaching an agreement. Naturally, this likelihood is defined by:

$$\frac{N_T}{N}, \quad (6.3)$$

in which N_T are the number of simulation runs resulting in agreement.

This finalizes an agent's individual set of objectives assumed in our tool. In return, they can be used to define the overall quality of outcome. Note that by definition, also these objectives are in the interval $[0, 1]$.

^[2]Without political support, ratification and implementation of international legislation (see Section 2.2.5) is often impossible as "(...) agreement may be reached, but the loser might end up wanting to sabotage its implementation." [108, p. 27-28].

6.1.3. Conclusion: sufficient pre-coalitions and quality of outcome

Every agent involved in the negotiation has its own definition of quality of outcome, depending on his set of objectives. Then, the quality is represented by weighted combinations of social welfare (i.e. collective) and self-interested (i.e. individual) objectives. Thus, the *quality of outcome* (QO) for agent i is defined by:

Quality of outcome (QO)

$$QO^i = \eta_1^i \cdot Q_1 + \eta_2^i \cdot Q_2 + \eta_3^i \cdot Q_3 + \eta_4^i \cdot Q_4 + \eta_5^i \cdot Q_5, \quad \text{in which:}$$

Q_1 = either *utilitarian*, *egalitarian* or *Nash product* approach to social welfare;

Q_2 = agent i 's utility $u_i(x)$ derived from the final offer x ;

Q_3 = agent i 's time objective as defined in Equation (6.1);

Q_4 = agent i 's political support objective as defined in Equation (6.2);

Q_5 = agent i 's likelihood of agreement objective as defined in Equation (6.3);

$\eta_1^i, \eta_2^i, \eta_3^i, \eta_4^i, \eta_5^i$ = agent i 's weights for Q_1, Q_2, Q_3, Q_4 and Q_5 , with $\sum_{k=1}^5 \eta_k^i = 1$.

As every quality measure $Q_x \in [0,1]$ and $\sum_{k=1}^5 \eta_k^i = 1$ also $QO \in [0,1]$. In the above definition, Q_x denote measures determining the quality of the objectives for agent i . Hence, *quality measures*, describing how well certain objectives are met. Sufficient outcomes are then defined as outcomes which substantially improve QO^i in specific negotiation settings and domains. As we have assumed the only strategic action agents make prior to actual negotiations is choosing pre-coalitions (see Section 2.2.5), we define:

Sufficient pre-coalitions as coalitions established prior to (actual) negotiation phase substantially improving the quality of outcome QO^i in specific negotiation settings and domains.

So, in order to find sufficient pre-coalitions – and, thereby, answer the main research question – we need to create negotiation settings and domains and determine their influence on the sufficiency pre-coalitions.

6.2. Experimental design

Settings and domains vary so widely in real-world international political negotiations that in this section 'examples' are introduced to answer the main research question. This type of theory building is called *exemplifying theory* and well-known in academics. Contrary to general theory, exemplifying theory "(...) does not tell us what must happen (...) [but] what can happen" [52, p. 117]. We justify its use by the vastness of real negotiation settings; it is impossible to derive grand assumptions upon which generalizing theory depends. Or if derived, they hamper practical application, which is the major criticism of many game theory models. Unfortunately, as a consequence, we lose generality.

The next two sections derive example negotiation settings (i.e. with whom/how to negotiate?) and domains (i.e. what is discussed?) by relating the balance-of-power theory to model parameters. Together, they describe *environment* factors for which pre-coalitions are evaluated. In Section 6.2.3, we define a handful of *decision* factors states face when considering pre-coalitions. Both influence the (added) quality of outcome and, therefore, the sufficiency of pre-coalitions.

6.2.1. Settings of negotiation environment

Remember a negotiation setting consists of the number and attributes of agents involved, strategies they deploy and protocol by which they interact. Then, example settings are derived from the balance-of-power theory in which states either balance or bandwagon. Choosing between the two is mainly motivated by power positions – weak, secondary and strong states. Linking this to model parameters, we define three types of agent power p , such that:

$$p_{\text{strong}} \gg p_{\text{secondary}} \gg p_{\text{weak}}$$

Though, in Section 3.1.1 we described that framing the two strategic choices solely in terms of power is seriously flawed. Therefore, the following environment factors are also included (from [175, p. 8-13]):

Table 6.1: Possible issues and outcome spaces in the simulation experiments.

Issue type	Possible outcomes O	p	Motivation
Binary	-	-	Not included in model, see #SC2 in Section F.1.
Indivisible resources	$ O = \{2, 3, 5, 8\}$ and $\sum_{j \in O} o_j = 100$	0.3	Represent either binary alternatives and high option issues from which agents derive utility.
Unrestricted resources	$O = \{[30, 100], [5, 25], [120, 150], [200, 400], [20, 30]\}$	0.5	These spaces represent different types of negotiations, from narrow low-outcome to high-outcome spaces and much in between.
Limited (shared) resources	$\pi = \{50, 200, 500\}$	0.2	Different pies representing scarcity and abundance.

- *Proximity* links to the ideological values Λ of negotiating agents;
- *Offensive capabilities* refer to cognitive limitations and incomplete information agents have and links to rationality parameter R ;
- *Offensive intentions* refer to the strategies agents deploy and the risk B they take to achieve their objectives and therefore links to their strategy profile(s) Ω . More aggressive agents will have higher probability to deploy competitive strategies, thus the function f changing their strategy profiles chooses complete profiles (i.e. "Competitor" or "Inverter"; see Section 5.3.3) more often;
- *Preferring status quo* is resembled by an agent's willingness to negotiate. The latter is associated with an agent's outside option or BATNA: the better his alternative, the less willing an agent is to negotiate [113, p. 103]. Thus, this factor links to an agents disagreement utility D in our tool.

The examples are completed by assuming agents make offers according to a Markov Chain ξ in which the initial transition matrix is based upon power, but updated by number of times agents made proposals already.

6.2.2. Domains in negotiation environment

Balance-of-power theory is applied in situations of immediate threat [174, p. 4]. Thus, the negotiation domain contains one subject, presumably dealing with the facing threat. Contrary, international political negotiations deal with many subjects and issues. Again, exemplifying theory is used illustrating such domains.

First, a "pool" of different types of issues is constructed, consisting of 4 distinct indivisible resource issues, 5 unrestricted resource issues and 3 limited resource issues. They are used to represent negotiation domains in which there are scarcity of resources, numerous of alternatives or complex ranges of possible outcomes (see Table 6.1). Following the linguistics of exemplifying theory, this pool is denoted as *event population*, whereas the different types of issues are *sub-populations* [185, p. 6-10]. Then, if we adjust the negotiation domain – hence, subjects S and issues I – we proportionally sample from the event population according to a specified distribution, as illustrated in Figure 6.1. To complete the domain, public deadlines are drawn from $N(20, 5)$ for each subject, after which the overall deadline is the accumulation of all subject deadlines.

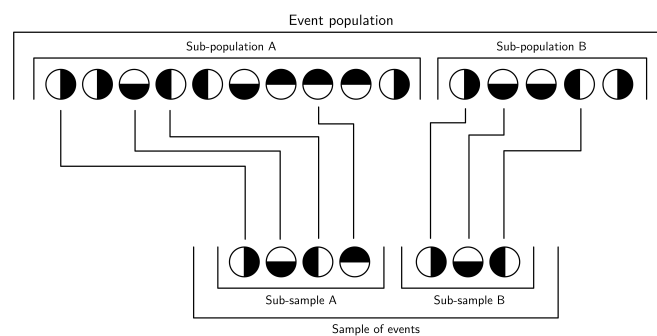


Figure 6.1: Conceptual approach to proportional sampling from known event population (adopted from [185, p. 9]).

6.2.3. Decisions: composition of pre-coalitions

When considering pre-coalitions, states potentially face many options. We have limited ourselves to the following *decisions*, describing the composition of the possible pre-coalition:

- *Relative size* as fraction of all agents in the negotiation environment;
- *Relative power* as fraction of all structural power in the negotiation environment;
- *Fraction of weak agents* given the power of the pre-coalition;
- *Fraction of secondary agents* given the power of the pre-coalition;
- *Fraction of strong agents* given the power of the pre-coalition;

It is this composition that makes pre-coalitions sufficient or not. Thus, the main parameters of interest are attributes of target pre-coalition C_T . Note that in literature, balancing and bandwagoning is unambiguous; either states fully balance with equally powered opponents, or they perfectly bandwagon with more powerful countries. Given the above decisions, we also consider in-between choices. For example, possibly a state's sufficient pre-coalition is to mainly bandwagon (e.g. 70%) but also to slightly balance (e.g. 30%).

Together with environment factors, these decisions complete the experimental setup^[3]. Subsequently, the experimental design is concluded by defining the relations between the environment and decisions (also denoted as *factors*), and the quality of outcome (denoted as *responses*). This is explained in the next section.

6.2.4. Relations between environment, decisions and quality of outcome

Figure 6.2 shows the presumed relations between factors and responses. We expect both the environment and decisions to affect the quality of outcome singly, and, therefore, the sufficiency of pre-coalitions. Furthermore, we are interested in which environments states should take different decisions. Therefore, we analyze interactions between negotiation settings and domains, and the composition of pre-coalitions. To achieve the aforementioned, four analysis steps have been taken:

- i) First, we examine the **added value** of pre-coalitions in general by comparing the difference between staying alone or forming (any) pre-coalition, given the same simulation experiment (see Figure 6.3a);
- ii) Then, we examine the **importance of each factor** on differences in outcome (see Figure 6.3b);
- iii) Subsequently, the **influence of each factor** on the quality of outcome is analyzed (see Figure 6.3c);
- iv) Finally, **interaction effects** between composition of pre-coalitions and the negotiation domain and setting are considered, combined with their influence on the quality of outcome (see Figure 6.3d).

The first analysis compares both simulation results for staying alone and forming pre-coalitions. Kernel density plots are provided to visualize both distributions of quality measures, as well as the differences between them. For the subsequent analyses, Random Forest regression models are fitted. These are used to calculate factor importances, (single) factor dependencies and interactions. To conclude, in Section 6.1.3 the quality of outcome QO is defined from the perspective of agent i . As our example settings include three types of agents, the above analyses are repeated from a weak, secondary and strong agent perspective.

6.3. Discussion: is this the right experimental design?

There are two components determining whether this is the right design. First, we should determine whether the experimental setup – i.e. the example negotiation domains and settings – represent real international political negotiations. This is part of *model validation* as discussed in Section 8.5. In brief, this section shows our results are valid in terms of social welfare Q_1 , personal gains Q_2 and political support Q_4 for the vast majority of negotiation domains. Unfortunately, both the time objective Q_3 and likelihood of outcome Q_5 are not independent of the domain and, therefore, the results should be interpreted within the boundaries of the adopted domain. In terms of negotiation settings, it is shown that our results are only valid within the 4:2:1 power ratio of strong, secondary and weak agents. Thus, in settings of true hegemony, results become insignificant.

^[3]See Appendix G for an oversight of all factors, used parameter ranges and derived negotiation scenarios.

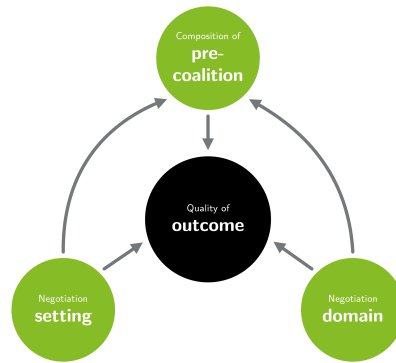


Figure 6.2: Presumed relations between factors (green) and responses (black).

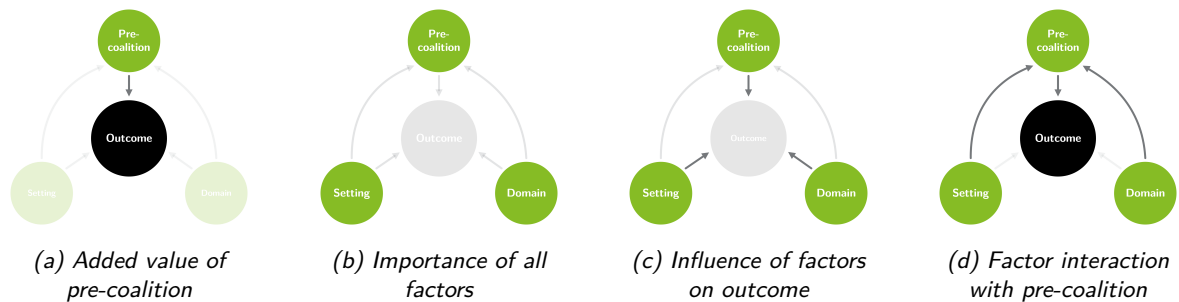


Figure 6.3: Focus of analyses to examine the presumed relations.

Secondly, the presumed relations should allow us to answer the main research question of this study. The four analysis steps of the previous question are designed to answer three intuitive questions states have: should they consider pre-coalitions in the first place? If so, in what negotiation environments are pre-coalitions sufficient? And, how should the sufficient pre-coalition be composed given the negotiation environment? By examining the added value (*i*) we answer the first question. Furthermore, the subsequent question is answered by determining both the importance and influence of the negotiation environment (*ii* and *iii*). By analyzing interactions between the composition of pre-coalitions and the negotiation domain and setting (*iv*) the last question is answered. To conclude, the assumed relations are adequate to answer the main research question. However, the results cannot be generalized to other domains and settings.

6.4. Conclusion & summary

This section answers the fifth supporting question: which objectives make pre-coalitions sufficient? We describe pre-coalitions are sufficient if they substantially improve the quality of outcome QO^i in specific negotiation settings and domains. This is determined by calculating quality measures – as function of an agent's personal objectives – and weight them accordingly. Agents might have the following objectives:

- Maximize *personal gains* (or utility) given preferences and goals of agent i for all subjects and issues;
- Maximize *social welfare* by increasing personal gains for all agents A involved in the negotiations;
- Minimize *time to agreement* given agent i 's own private or public deadlines for each subject;
- Maximize *political support* by increasing the number (or power) of agents agreeing with the outcomes;
- Maximize *likelihood of agreement* given an agent's willingness to negotiate.

For many agents, maximizing personal gains is the primary objective. However, we argue that in some situations, agents may have one or more secondary objectives as well. An experimental design is constructed to analyze the influence of pre-coalitions on agent i 's full set of objectives QO^i . This design lays the foundation for simulation experiments and, thereby, for answering the main research question.

In this thesis, *negotiation settings and domains* are described using exemplifying theory, as they vary widely in international political negotiations. Examples are derived from the balance-of-power theory and distinguish three types of hypothetical states based on structural power: weak, secondary and strong. These states vary in number, ideological differences, rationality, risk behavior and deployed negotiation strategies. The domain includes various subjects and issues, aiming to capture the negotiation's capriciousness.

To examine the influence of negotiation settings and domains on the (structure of) sufficient pre-coalitions, we presume certain relations between them. These are examined by i) determining the added value of the pre-coalition in general, compared to staying alone; ii) calculating the importance of each relationship; iii) determining influence of each relationship on the quality of outcome and iv) detect interactions between the negotiation setting, domain and structure of pre-coalition on quality of outcome. This is accomplished by considering the different power positions of agents.

7.

Simulation Results

A great man once said, everything is about sex. Except sex. Sex is about power.
Francis (Frank) Underwood, in House of Cards (2013)

In this chapter, a selection of simulation results are presented (see Appendix H for the complete oversight). Thereby, it lays the foundation for answering the last supporting question: how do pre-coalitions influence a state's set of objectives in different negotiation environments? As discussed, this question is answered by three intuitive questions agents may have: should pre-coalitions be considered in the first place? If so, in what negotiation environments are pre-coalitions sufficient? And, how should pre-coalitions be composed given their environment? Note that the four subsequent analyses provide the answers to these questions (see Section 6.2.4). Additionally, as the main research question is answered using the concepts of bandwagoning and balancing, in the results we are ultimately looking for evidence showing pre-coalitions are sufficient if:

- States **fully** balance with one or more other states;
- States **fully** bandwagon with one or more other states;
- States follow **mixed** balancing and bandwagoning strategies;
- States do none of the above (so, staying alone before the actual negotiations).

Now, we commence this chapter by presenting raw model results and linking them to the quality measures Q_x . Subsequently, in Section 7.2, the first intuitive question is answered by analyzing the added value of pre-coalitions. Furthermore, Section 7.3 commences with justifying the use of Random Forests, after which in Section 7.3.1 we analyze the negotiation environments in which pre-coalitions are sufficient (or not). Finally, in Section 7.3.2, the composition of pre-coalitions is discussed. Together, this should answer all three intuitive questions above.

7.1. Raw output: before quality measures are calculated

The model output includes agent *utilities* (as derived from negotiation outcomes), *time* (or steps) to agreement and, *size* and *power* of the final winning coalition, which are based upon 240,000 simulations. More specifically, for each type of agent, two times 2,000 samples are drawn from the parameter space (see Table G.1) using Latin Hypercube sampling (see Appendix G.1). For every sample, the quality of outcome QO and Q_x are calculated for either staying alone or forming pre-coalitions, based upon the raw model outputs. As QO itself is determined by 20 simulations to capture variety, 80,000 simulations are performed for every agent type, taking roughly 14, 17 and 15 hours of computing time respectively.

In Figure 7.1 the distribution of actual steps needed to reach agreements (if any) is plotted for all simulations. The spread is substantial – from approximately 1 to 120 – as the histogram shows negotiations with 1 to 15 subjects. Moreover, the tail is heavily skewed to the left due to the sequential simultaneous implementation of subjects, in which negotiation break down whenever a subject ends in no agreement. Moreover, in Figure 7.2 we have plotted the fraction of simulations that resulted in agreement and, the size and power of the winning coalition. We observe that overall the simulations have a low (median) fraction of




Figure 7.1: Number of simulation steps before agreement is reached for weak, secondary and strong agents.

negotiations in agreement. Again, this can be explained by the break down of negotiations similar to above. Furthermore, simulations of weak agents have a slightly lower fraction in agreement – i.e. approximately 0.23 versus 0.28 for both secondary and strong agents. Remark these boxplots show the raw output, they contain both the choice of forming pre-coalitions or staying alone. So, we also 'force' weak agents to stay alone which most likely hinders the likelihood of agreement due to reduced relative decisiveness.




Figure 7.2: Fraction of negotiation in agreement (left) size (middle) and power (right) of the winning coalition.

Additionally, the two other boxplots show both remarkable and expected results. First, in the right plot we see that on average the power of the winning coalition is the same for all simulations. Which can be expected, as the perspectives we take – i.e. weak, secondary and strong – should not affect the final winning coalition, probably containing various types of agents. What is striking is that from the perspective of strong agents, the size of the winning coalition is slightly greater. In turn, this implies the distribution of power within the winning coalition is rather asymmetric, as more agents have the same aggregated power and at least one agent is strong. Therefore, this is a first indication bandwagoning is slightly more fruitful than balancing for strong agents. Though, this will be supported with more evidence in subsequent sections.

In the following section, we split the simulation results for each agent for their two options of either forming pre-coalitions or staying alone. Moreover, we calculate the quality of outcome QO and its measures Q_x , among others based upon the raw model outputs displayed above. Finally, we calculated the *added value* of pre-coalitions^[1] by subtracting QO_{alone} from QO_{coal} for the exact same simulation experiments. This will be mainly used for the analyses in Section 7.3.1 and Section 7.3.2.

7.2. Forming pre-coalitions or staying alone?

First, let us compare the added value of pre-coalitions for the different types of agents, similar to the first relation in our experimental design (see Figure 6.3a). These values are presented in Figure 7.3 to 7.5. Note that all values *above* the dotted green line indicate true added value. In terms of overall QO , we notice weak agents benefit the most from forming pre-coalitions as nearly 70% of the simulations result in positive added value. Also the vast majority of simulations for secondary agents result in added QO . In contrary, for strong agents, this percentage is dropped to slightly above 50%.

Taking a look at which measures Q_x benefit from forming pre-coalitions, in Figure 7.3 to 7.5, for all agents we observe the likelihood of agreement Q_5 increases substantially. For weak agents, we also see personal gains Q_2 are likely to increase, closely followed by their time Q_3 and political support Q_4 objectives. Additionally, we observe a slight increase in Q_2 for secondary agents, whereas the median of their time and political support objectives are centered around 0. Finally, strong agents mainly benefit from forming

^[1]Although 'added value' has a positive connotation, whenever its value becomes negative, staying alone should be preferred.

pre-coalitions by slight increases in their time and political support objectives. In contrary, strong agents, seem to loose slightly in terms of personal gains for the majority of simulations. The same holds for their social welfare objectives Q_1 , whereas no changes are observed for weak and secondary agents.

Table 7.1: Bayesian probabilities $P(\mu_{coal} - \mu_{alone} > v)$ for weak, secondary and strong agents' QO and Q_x .

	Weak agents						Secondary agents						Strong agents					
	v	0	.01	.02	.05	.10	.15	0	.01	.02	.05	.10	.15	0	.01	.02	.05	.10
QO	1	.99	.99	.98	.50	.02	1	1	1	.99	.89	.23	.27	.15	.08	0	0	0
Q_1	.87	.09	.01	0	0	0	.89	.11	.02	0	0	0	.46	.01	0	0	0	0
Q_2	.97	.94	.86	.57	.01	0	.92	.91	.75	.45	.11	0	.14	.08	.02	0	0	0
Q_3	.99	.99	.99	.99	.85	.38	.89	.83	.68	.21	.12	.22	.99	.99	.99	.99	.91	.50
Q_4	1	1	1	1	.99	.81	.91	.81	.72	.42	.21	.25	.99	.99	.99	.96	.62	.12
Q_5	1	1	1	1	1	.98	1	1	1	1	1	.99	1	1	1	1	.99	.97

Bayesian inference. Table 7.1 supports the above observations quantitatively by performing Bayesian inference^[2]. We see that the probability the likelihood of agreement Q_5 increases by 15% is approximately 0.98 for all agents. Moreover, there is 0.57 chance personal gains of weak agents increases by 5%, whereas these chances are only 0.45 and 0.00 for secondary and strong agents respectively. In contrary, whenever strong agents want to 'get things done' – i.e. increase their time objective Q_4 , Table 2.1 shows that with $p = 0.91$ this increases with at least 10% (and, thus, the negotiation time decreases accordingly).

Kernel density estimation. The distribution of quality measures Q_x and quality of outcome QO for all types of agents are approximated by kernel density estimation in Figure 7.6, Figure 7.7, Figure 7.8 for weak, secondary and strong agents respectively. Besides (partly) increasing added value, from these plots, we also observe forming pre-coalitions affects the distribution of simulation outcomes as QO partly becomes 'dichotomous' for all agents. The outcomes for staying alone are in some way normally distributed with relatively high standard deviation, of which Q_2 is the major driving force. By forming pre-coalitions, the distribution of QO becomes double peaked. In turn, this is mainly caused by dichotomous distributions of Q_3 and Q_4 , in combination with higher and more secure – hence, lower standard deviation – expected personal gains Q_2 . In Chapter 8 we discuss a possible explanation for this dichotomous behavior.

Though, the dichotomous distributions of (among others) QO indicate that apparently in some negotiation settings, domains and compositions pre-coalitions become undoubtedly beneficial, whereas in others they truly do not. Therefore, in the next sections, we dive into the negotiation environment and composition of coalitions by applying Random Forests.




Figure 7.3: Added value of pre-coalitions for QO and its measures Q_x for weak agents.




Figure 7.4: Added value of pre-coalitions for QO and its measures Q_x for secondary agents.

^[2]We prefer Bayesian inference – as an alternative to 'traditional' frequentist hypothesis testing – as it is more fundamentally sound and does not draw hard significance boundaries [76, p. 4].




Figure 7.5: Added value of pre-coalitions for QO and its measures Q_x for strong agents.

7.3. Sufficiency of pre-coalitions

In this section, in total 36^[3] Random Forest models are fitted explaining the *added value* of pre-coalitions. This, to examine the impact of the negotiation environment – i.e. negotiation setting and domain – as well as to what extent agents can influence the added value by altering the pre-coalition’s composition. Note that this is in line with the second relation of our experimental design (see Figure 6.3b). First, let us justify the use of Random Forest models on simulated (instead of empiric) data.

Our experimental design consists in total of 16 factors – of which 11 environment and 5 coalition composition related – and 2×6 responses. Hence, leading to 28 relevant variables for each simulation experiment. Conclusions of this thesis, therefore, rely upon a 240,000 by 28 dataset, far too large to be analyzed by simple descriptive statistics and visualizations. Fitting Random Forest regression models not only allows us to inspect relationships between different variables (see next paragraph), it is also a convenient and efficient way of inspecting the simulated data *without* losing (too) much detail. Good examples include the factor importances and interaction plots in the next subsections. Additionally, in contrast to, for example, linear regression models, it deals with multiple data types – i.e. discrete and continuous – much better.

Furthermore, we are aware the constructed simulation model completely and utterly behaves as we have designed it to do so. Therefore, it may seem contradictory to analyze its outcomes as if they are empiric observations. Still, we have three reasonable explanations for doing so. First, one can argue we could predict model outcomes based on an axiomatic approach, similar to most bargaining models (see [113, p. 30-33]). However, contrary to these models, we have included substantially more factors increasing complexity tremendously and making the axiomatic approach cumbersome. Secondly, we intendedly added randomness to the model, describing the capricious and chaotic nature of negotiations and unpredictability of negotiators [35, 36, 79]. Of course, this could still be included in axiomatic approaches. In fact, randomness is explicitly assumed in, for example, Rubinstein’s alternating offers protocol [148]. However, various interrelated elements of our model are random. Together with increased complexity, this makes it extremely hard to make predictions and draw conclusions manually. Especially, given the large parameter space in our experimental design (see Table G.1). Random Forests can, in such situations, more effectively quantify complex relations between factors and responses, compute their importances and visualize non-linear dependencies.

Finally, when effectively embedded in our tool, Random Forests are also easier to use by others. Remember, one of the objectives of this thesis is allowing others to (re)use our tool for their own negotiations. Following an axiomatic approach may lead to understandable conclusions for game theorists and mathematicians, but decreases the adaptability of our tool for those who are less familiar with such approaches.

Fitness. Now their use is justified, let us consider fitness of the Random Forest models. Table 7.2 shows the R^2 values of the fitted models from the perspective of weak, secondary and strong agents in a clockwise direction. Note that ‘Env.’ and ‘Comp’ represent the models in which only environment or composition factors are used respectively. Overall, we observe high R^2 values, indicating most variation in outcomes is well explained by. Moreover, most models have higher R^2 values when factors describing the composition of

^[3]To be precise, $(2 \times 6 =) 12$ distinct models are fitted for both the environment and coalition composition factors, predicting 5 quality measures Q_x and quality of outcome QO respectively. Additionally, this is repeated for all three agent perspectives. Some important model parameters can be found in Appendix H.

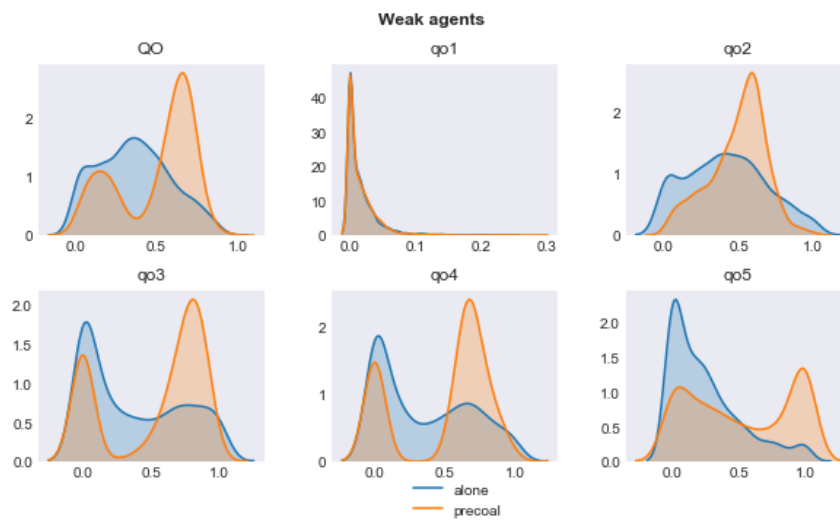


Figure 7.6: Kernel density estimation of distribution of QO and its measures Q_x for weak agents.

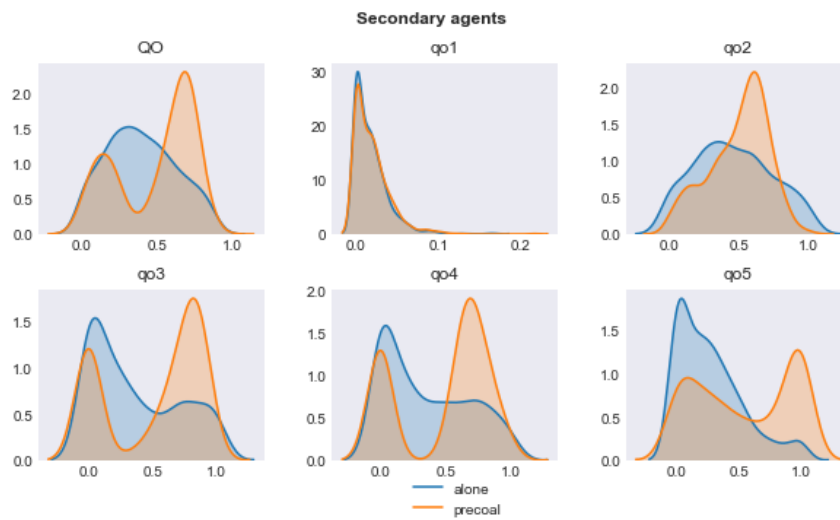


Figure 7.7: Kernel density estimation of distribution of QO and its measures Q_x for secondary agents.

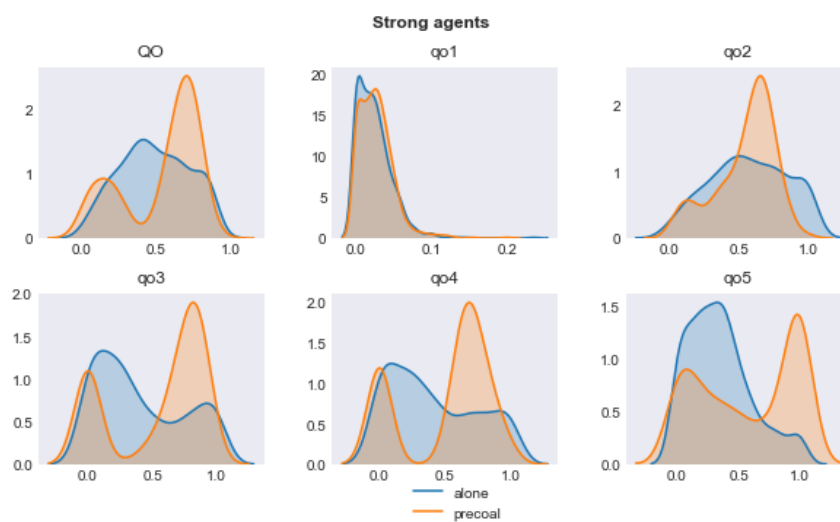


Figure 7.8: Kernel density estimation of distribution of QO and its Q_x for strong agents.

Table 7.2: R^2 of Random Forests for weak (top-left), secondary (top-right) and strong (bottom) agents.

Responses							Responses						
	Q0	Q1	Q2	Q3	Q4	Q5		Q0	Q1	Q2	Q3	Q4	Q5
Env.	0.80	0.78	0.82	0.76	0.68	0.81	Env.	0.80	0.80	0.80	0.77	0.79	0.74
Comp.	0.85	0.82	0.81	0.84	0.86	0.84	Comp.	0.87	0.87	0.87	0.87	0.87	0.87

Responses						
	Q0	Q1	Q2	Q3	Q4	Q5
Env.	0.81	0.80	0.77	0.80	0.80	0.80
Comp.	0.87	0.87	0.86	0.81	0.87	0.87

pre-coalitions are used as independent variables. This implies they explain the added value of pre-coalitions (slightly) better, indicating that changing the composition of pre-coalitions does effect its sufficiency.

7.3.1. In different negotiation environments

Factor importances. To examine what environment factors contribute to the overall fit, feature importances have been plotted in Figure 7.9, Figure 7.10 and Figure 7.11 for weak, secondary and strong agents respectively. These plots show the factors (y-axis), relative importances (x-axis) and standard deviation (black lines). Bars with high opacity have mean decrease accuracy or MDA > 0.10 ^[4].

The following can be concluded for *weak* agents:

- For Q0 and almost all Q_x the number of agents in the negotiation environment and how they behave – i.e. willingness to negotiate and rationality – affect the added value of pre-coalitions;
- Though, the average risk and aggressiveness other agents are willing to take is of less importance for all Q_x and Q0;
- The number of subjects discussed also influences the added value substantially for all Q_x and Q0;
- The distribution of power among secondary and strong agents has little influence on the added value;
- Number of weak agents do affect the added value, especially for Q_4 , Q_3 and Q_2 (in decreasing order);
- Both the number of initial ideological clusters and issues in per subject have the least impact.

Additionally, for *secondary* agents we conclude:

- Again, for Q0 and almost all Q_x the number of agents in the negotiation environment and how they behave – i.e. willingness to negotiate, rationality and risk behavior – affect the added value of pre-coalitions;
- The number of weak agents in the environment especially affects the added value for Q_1 , Q_3 and Q_4 ;
- The number of subjects in the negotiation domain is particularly important for Q_3 and Q_4 and to lesser extend also for Q_2 ;
- Both the power of other secondary and strong agents have little influence on the added value of pre-coalitions on Q0 or all Q_x . Except Q_4 , which is to some extent affected by the number of strong agents in the negotiation setting;
- Both the number of initial ideological clusters and issues in per subject have the least impact;

Finally, for *strong* agents we conclude:

- Added value of pre-coalitions for Q0 and almost all Q_x is still affected by the number of agents in the negotiation environment and how they behave. However, average risk behavior is of less importance, whereas average aggressiveness becomes increasingly important;

^[4]Plots only show importances relative to each other, whereas MDA shows the global importance of features independently [165]. Thus, $MDA > 0.10$ implies the factor contributes at least 10% to explain the variance in quality outcomes.

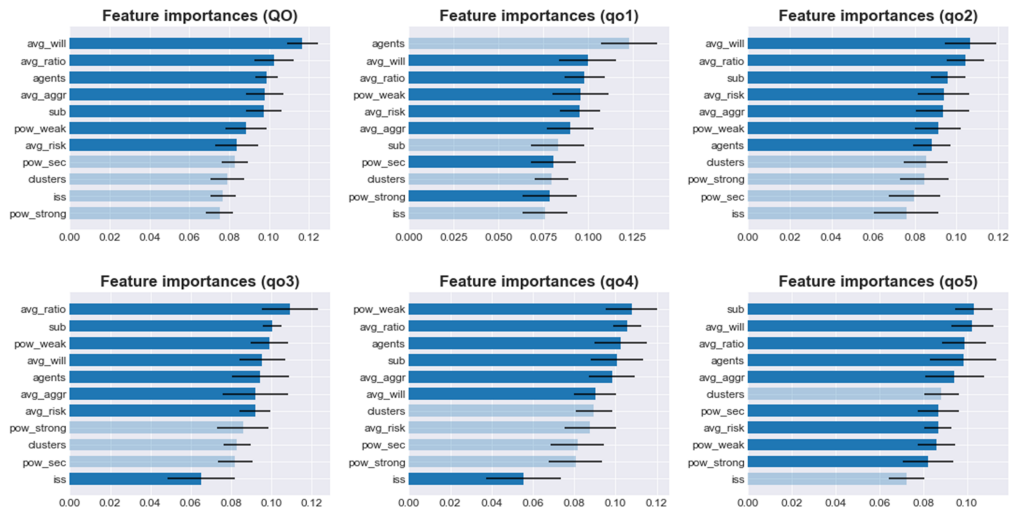


Figure 7.9: Environment factor importances for predicting the added value of pre-coalitions for weak agents.

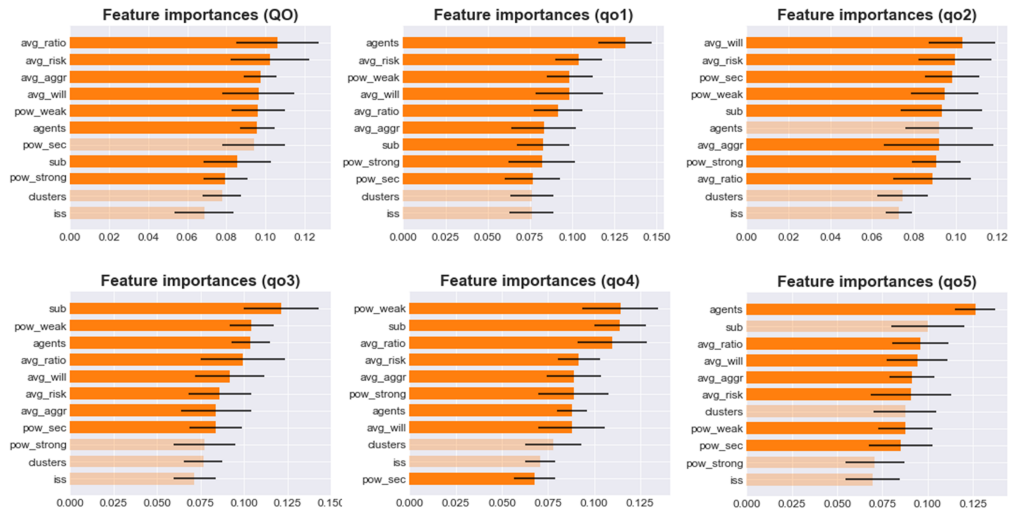


Figure 7.10: Environment factor importances for predicting the added value of pre-coalitions for secondary agents.

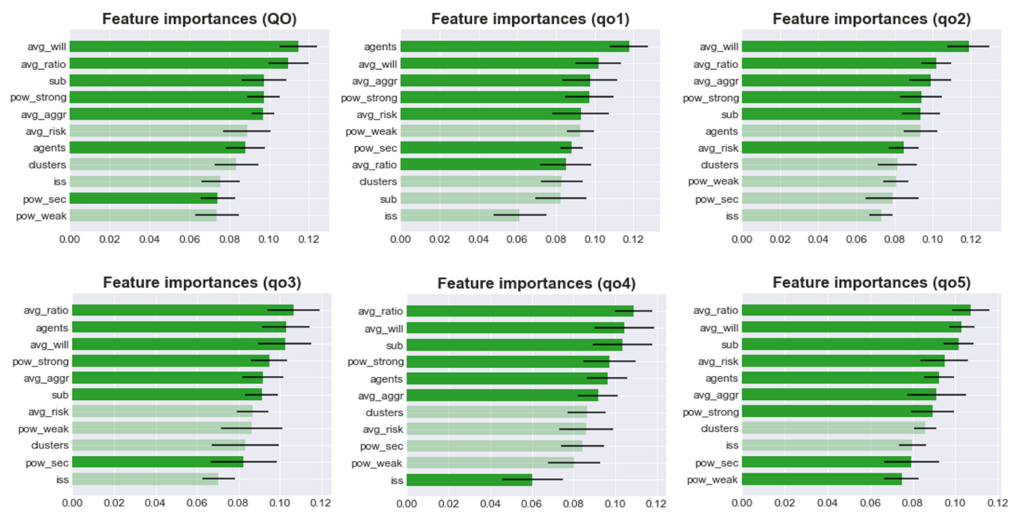


Figure 7.11: Environment factor importances for predicting the added value of pre-coalitions for strong agents.




Figure 7.12: Single environment factor dependencies on QO for weak agents.

Figure 7.13: Single environment factor dependencies on Q_1 for weak agents.

- Additionally, for personal gains Q_2 , number of agents is not significant, rather the way they behave;
- The number of other powerful agents in the negotiation environments is of substantial importance for the added value of pre-coalitions. Though, to a lesser extent for QO and likelihood of agreement Q_5 ;
- On the other hand, the number of weak and secondary agents is not important;
- The number of subjects in the negotiation domain is of great importance explaining the added value of pre-coalitions for QO , Q_4 and Q_5 , of some importance for Q_2 and no importance for Q_1 ;
- Both the number of initial ideological clusters and issues in per subject have the least impact;

So, to summarize, generally the number of agents in the negotiation environment and the way they behave is of substantial importance explaining the added value of pre-coalitions. Though, especially for weak and strong agents their average risk behavior is of less importance, whereas for secondary agents it does. Furthermore, the initial power distributions within the negotiation potentiometer is mainly of importance for secondary and strong agents. Especially the added value of pre-coalitions for strong agents is affected by the number of other equally powerful agents. Finally, for most quality measures and overall quality of outcome, the number of subjects in the negotiation domain are substantially important in explaining the added value, whereas the number of issues for each subject and ideological clusters are not.

Single factor dependencies. Unfortunately, the importances above not determine in what way the factors affect the added value – and, thus, sufficiency – of pre-coalitions for different negotiation environments. Therefore, in line with the third relationship in our experimental design (see Figure 6.3c), we visualized single factor partial dependencies. In brief, these plots show how and to what extent the added value of pre-coalition changes when increasing the value of a single factor. Together with importances, this answers the second intuitive question. As there are many plots, we only discuss most relevant results for our research [5]. First, for *weak* agents we can conclude the overall added value QO of pre-coalitions increases for more agents in the negotiation environment and becomes positive > 8 agents, see Figure 7.12 (top-right). This is mainly due to increases in their time objectives Q_3 and Q_4 who also become positive for approximately > 8 agents, see Figure 7.15 (middle) and Figure 7.16 (right) respectively. Moreover, from Figure 7.12 (top-left), it can be seen that low numbers of subjects negatively affect the overall added value QO of pre-coalitions, whereas numbers only slightly lead to greater added value. The same dependency is also observed for social welfare Q_1 , personal gains Q_2 and, to a lesser extent, also for Q_4 . The latter is shown in Figure 7.13 (left), Figure 7.14 (left) and Figure 7.16 (left) respectively.

In addition, we observe that the average rationality of agents – thus, the chance they might make a wrong or different decisions given the available information an agent has – has a moderate, yet capricious affect on the overall added value of pre-coalitions QO , see Figure 7.12 (bottom-left). We observe similar dependencies on Q_1 , though slightly negative. So, for when agents become more rational, the added value of pre-coalitions in terms of social welfare decreases slightly. The latter is shown in Figure 7.13 (bottom-left). The capriciousness is, however, mainly caused by the dependency for Q_3 , in which we see a slight overall – yet unsteady – increase of added value of pre-coalitions in Figure 7.15 (right).

Finally, for weak agents, we observe a (very) slight decrease in overall added value QO if the average willingness of other agents in the negotiation environment increases, as shown in Figure 7.12 (bottom right).

[5] See Appendix H for all results and additional analysis.




Figure 7.14: Single environment factor dependencies on Q_2 for weak agents.




Figure 7.15: Single environment factor dependencies on Q_3 for weak agents.




Figure 7.16: Single environment factor dependencies on Q_4 for weak agents.




Figure 7.17: Single environment factor dependencies on Q_0 for secondary agents.




Figure 7.18: Single environment factor dependencies on Q_1 for secondary agents.




Figure 7.19: Single environment factor dependencies on Q_2 for secondary agents.




Figure 7.20: Single environment factor dependencies on Q_3 for secondary agents.




Figure 7.21: Single environment factor dependencies on Q_4 for secondary agents.

In turn, this is mainly caused by the dependency for personal gains Q_2 , as can be seen in Figure 7.14 (right). There, we see for $\text{avg_will} > 0.3$ – thus, relatively high – the added value becomes negative. Note that if 'avg_' is high, this implies agents have high outside options and, are, therefore, less willing to negotiate. So, as they are less willing to negotiate, Figure 7.14 (right) shows pre-coalitions might become beneficial. To a lesser extent we also observe the same dependency for social welfare Q_1 in Figure 7.13 (bottom-right).

Secondly, for *secondary* agents we observe similar dependencies as before. So, the overall added value of pre-coalitions QO becomes positive if more subjects are discussed in the negotiation domain, see Figure 7.17 (left). In this figure, for 8 or more subjects, pre-coalitions become beneficial. If we take a look at its major drivers, we observe that especially personal gains Q_2 and time objective Q_3 increases by forming pre-coalitions, as shown in Figure 7.19 (left) and Figure 7.20 (left) respectively. In a lower degree, same dependencies are observed in Figure 7.18 (left) and Figure 7.21 (left) for social welfare Q_1 and political support Q_4 .

Furthermore, for secondary agents we also observe the overall added value of pre-coalitions QO seems stable for decreases^[6] in the average willingness to negotiate, as displayed in Figure 7.17 (middle). Though, we do see a strong dependency for Q_2 as we take a closer look at Figure 7.19 (middle). In this figure, lower willingness to negotiate lead to substantially less added value.

Also, we observe a slight dichotomous effect on the overall added value QO for increasing rationality, as displayed in Figure 7.17 (right). For low average rationality values (i.e. < 0.3), pre-coalitions become slightly beneficial, whereas for high values (i.e. > 0.75) they affect QO negatively. We see this dependency is partly caused by the added value of Q_2 . Moreover, in Figure 7.19 (right) we observe a sudden drop in added value in terms of personal gains for average rationality > 0.3 , after which the added value becomes nil. We have also identified environment factors which overall have little effect on the added value of pre-coalitions, but for individual quality measures they do. For example, in Figure 7.18 (middle), we observe the average aggressiveness affects the added value of pre-coalitions negatively in terms of social welfare Q_1 . In contrary, added value of Q_3 is slightly positively influenced by the aggressiveness, indicating that if others negotiate more aggressively, it might be beneficial for secondary agents to form pre-coalitions for meeting their time objectives. The latter is displayed in Figure 7.20 (middle). Moreover, we observe some minor dependencies on the average risk agents are willing to take for Q_3 and Q_4 . In Figure 7.20 (right), for example, we see that for more risky behavior – i.e. $\tilde{\beta} \rightarrow 0$ (!) – pre-coalitions become beneficial for secondary agents.

To conclude, for *strong* agents we observe slightly different dependency relationships. If we take a look at Figure 7.22, we see that even the most relevant environment dependencies are relatively stable. For example, the plot in the middle shows a slight increase in overall added value QO of pre-coalitions if the fraction of other strong agents increases. Additionally, the plot on the right shows a weak negative dependency on the average aggressiveness. Thus, whenever on average agents negotiate more aggressively, the QO decreases. More specifically, in terms of personal gains Q_2 , pre-coalitions become more beneficial for higher – i.e. > 0.5

^[6]Remember higher values for 'avg_will' indicate better outside options and, thus, a decreasing willingness to negotiate.




Figure 7.22: Single environment factor dependencies on Q_0 for strong agents.




Figure 7.23: Single environment factor dependencies on Q_2 for strong agents.




Figure 7.24: Single environment factor dependencies on Q_3 for strong agents.

fractions of secondary and strong agents in the negotiation environment. The latter is displayed in Figure 7.23 on the top-right and bottom-left respectively. In contrary, the fraction of weak agents have no dependency on the added value. Though, in the top-middle plot we observe that for very low fractions of weak agents, pre-coalitions are beneficial. However, very low fractions of weak agents implies higher fractions of secondary or other strong agents, which explains the slight peak. Additionally, Figure 7.23 (top-left) shows the added value of pre-coalitions for Q_2 increases if the average number of issues for each subject increases. This implies that strong agents only benefit from pre-coalitions in those environments in which package-deals are more common. Furthermore the last two plots of Figure 7.23 show that strong agents' gains might slightly benefit from forming pre-coalitions whenever on average agents act less risky or whenever the willingness to negotiate is relatively high. Additionally, in Figure 7.24 (left and middle), we observe that the added value of pre-coalitions in terms of time objective Q_3 slightly increases for either more issues or subjects in the negotiation domain. So, strong agents may benefit from pre-coalitions in terms of Q_3 if the negotiations are rather long or package-deals are more common. Moreover, we also see that strong agents are more likely to meet their time objective by forming pre-coalitions if there are more agents in the environment. Or, put differently, whenever the number of agents < 7 , strong agents can better stay alone.

Furthermore, in Figure 7.25 (left) we observe a rather strong dependency on the number of subjects in terms of political support Q_4 . So, whenever strong agents aim for high political support, forming pre-coalitions becomes beneficial for high (i.e. > 8) subjects in the negotiation domain. We also observe in this figure – in the plots in the middle and right – a weak negative dependency on the average risk $\tilde{\beta}$ and rationality. This implies that strong agents weakly benefit from forming pre-coalitions if, on average, other agents tend to negotiate more aggressive or if they become more rational. Finally, in Figure 7.26 (right) a very strong dependency is observed on the number of agents with respect to the likelihood of agreement Q_5 . This implies that whenever strong agents really want negotiations to succeed, they should form for pre-coalitions, especially if there are (way) more than 10 agents in the negotiation environment. The other two plots of Figure 7.26 show rather capricious dependencies. So, we cannot really conclude if pre-coalitions are beneficial in terms of Q_5 whenever agents are more/less rational or willing to negotiate.




Figure 7.25: Single environment factor dependencies on Q_4 for strong agents.




Figure 7.26: Single environment factor dependencies on Q_5 for strong agents.

Conclusions. We briefly conclude this section by answering the second intuitive question: in what negotiation environments are pre-coalitions sufficient? We have seen that both weak and secondary agents benefit from pre-coalitions in more or less the same negotiation environments. That is, whenever there are many agents at the table, they are willing to negotiate or act rational. Moreover, more complex negotiation domains increase weak agents' overall objectives QO substantially, whereas secondary agents mainly benefit from pre-coalitions in terms of personal gains Q_2 , time Q_3 and political support Q_4 objectives. In contrary, for strong agents, the distribution of power within the negotiation environment is of importance for the sufficiency pre-coalitions, especially for relatively high fractions of secondary or other strong agents. Moreover, they also benefit more – especially in terms of personal gains Q_2 and time Q_3 – if package-deals are common. Finally, strong agents also benefit from pre-coalitions if agents, on average, are willing to negotiate.

7.3.2. By changing the composition

Factor importances. This section presents the results of the analyses conducted to explore how pre-coalitions should look like given the negotiation environment. Or, more specifically – and in line with the final relationship in our experimental design (see Figure 6.3d) – the interactions between the composition of the pre-coalitions and the negotiation setting and domain. Again, we start with determining the importances of the 5 decision factors describing this composition, as presented in Figure 7.27, Figure 7.28 and Figure 7.29 for weak, secondary and strong agents respectively^[7]. As expected, from these figures, we observe that for all types of agents and quality measures Q_x and QO , the added value of pre-coalitions can be best explained by their size and power. However, we observe as well that for all types of agents the influence of the same size and power on Q_3 , Q_4 and Q_5 have substantial greater uncertainty. This is especially true for the likelihood of agreement Q_5 , indicating that the added value of pre-coalitions is strongly environment dependent.

Moreover, we also see that to a lesser extent the power distribution within the pre-coalition^[8] – i.e. the fraction of weak, secondary and strong agents – is important, particularly for all types of agents' personal gains Q_2 . Although substantially smaller, the power distribution is more stable in explaining the added value of pre-coalitions in terms of Q_2 . More specifically, it can be seen that of this power distribution, for weak agents the fraction of other weak agents is (globally) most important; for secondary agents, the fraction of other secondary agents explains the added value best and; for strong agents, the fraction of other strong agents is most relevant. Though, the aforementioned power distributions are quite close to one another.

Two-level factor interactions. To explore how each of the decision variables interact with the environment, we have plotted factor interaction plots as part of the Random Forest models. By means of these plots, we are trying to answer the third intuitive question: how should pre-coalitions be composed given their environment? Note that in this section we mainly discuss the overall added quality of outcome QO by

^[7]Similar to Section 7.3.1, those bars with high opacity have mean decrease accuracy or MDA > 0.10

^[8]Note that it is this power distribution that resembles (mixtures of) balancing and bandwagoning behavior.

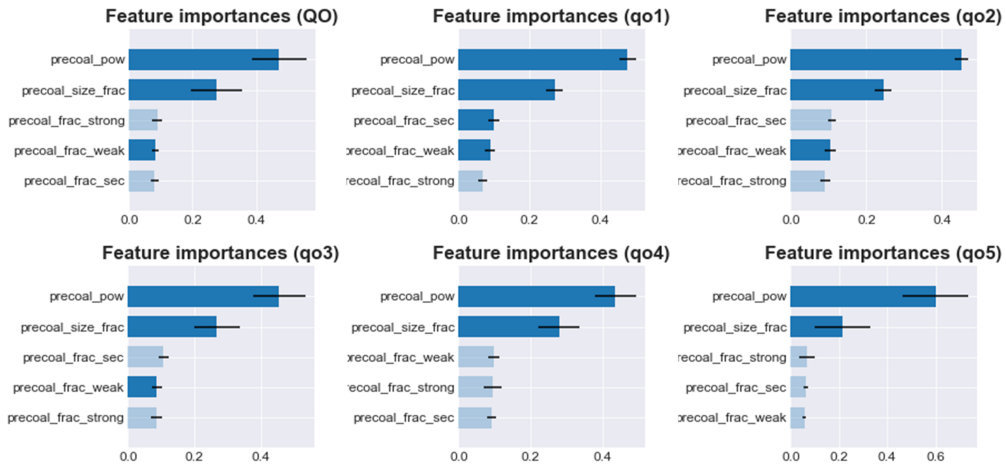


Figure 7.27: Pre-coalition composition importances for predicting added value of pre-coalitions for weak agents.

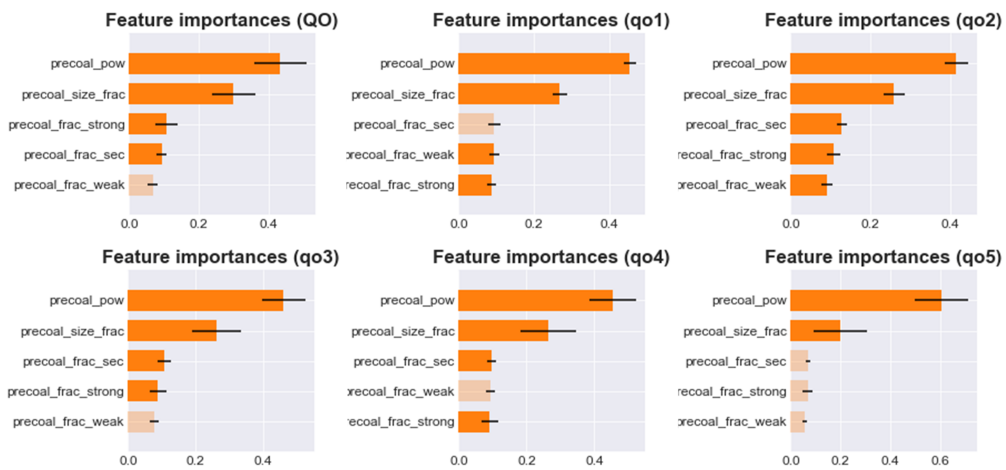


Figure 7.28: Pre-coalition composition importances for predicting added value of pre-coalitions for secondary agents.

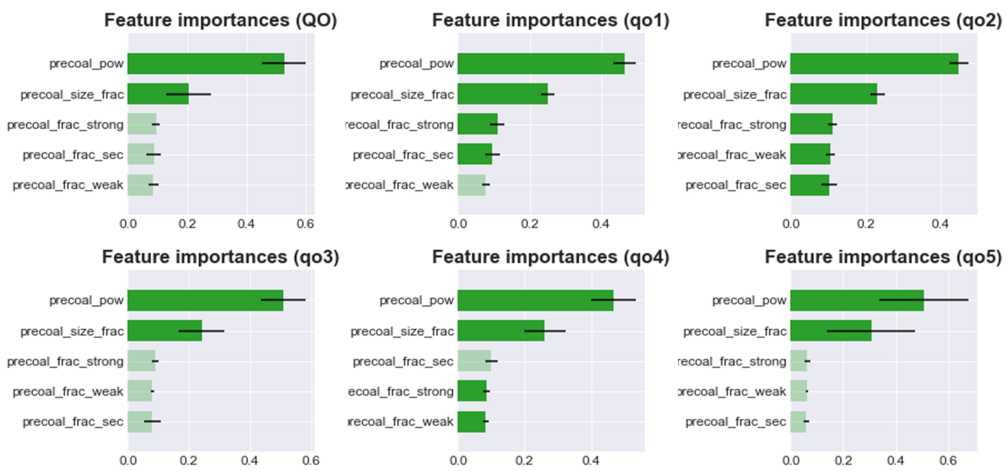


Figure 7.29: Pre-coalition composition importances for predicting added value of pre-coalitions for strong agents.




Figure 7.32: Interaction between environments and decisions for Q_2 and weak agents.

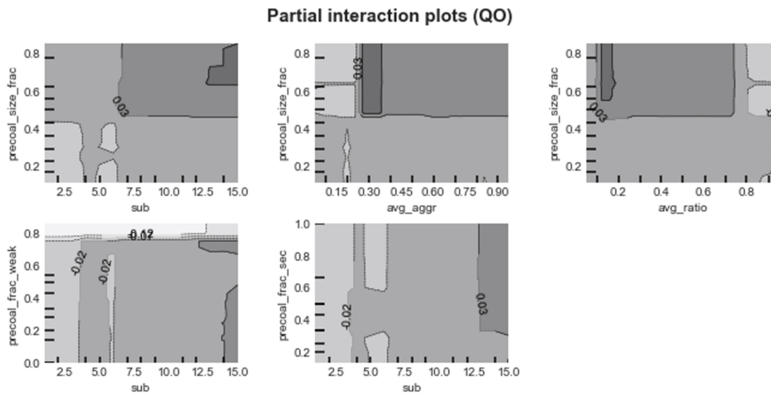


Figure 7.33: Interaction between environments and decisions for QO and secondary agents.

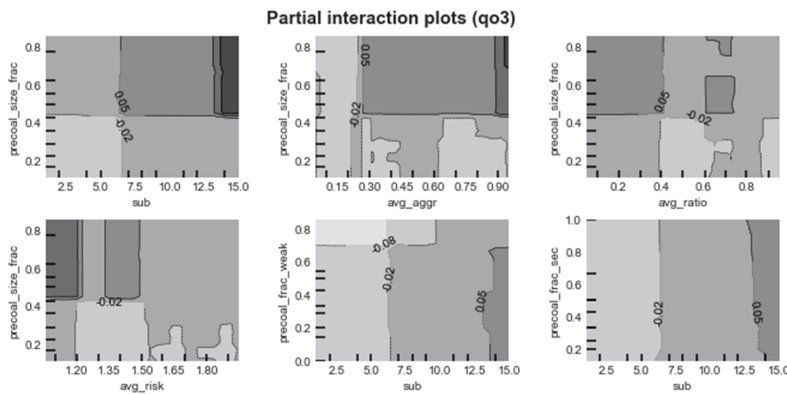


Figure 7.34: Interaction between environments and decisions for Q_3 and secondary agents.

power distribution of pre-coalitions, we see in more complex negotiation domains, high fractions of secondary agents (above 30%) contribute to more sufficient pre-coalitions. Though, we also see that up to 50% of weak agents in the pre-coalition increases the overall added value QO . The same interactions are observed for Q_2 (see Figure H.27).

Furthermore, in Figure 7.34, the interaction effects for a secondary agent's time objectives Q_3 are presented. We observe that large pre-coalitions > 0.3 positively affect added value in complex negotiation domains (> 10 subjects), settings with mildly/heavily aggressive agents (i.e. $\bar{p}(\Omega = \text{"Competitor"}) > 0.3$ and $\bar{p}(\Omega = \text{"Competitor"}) > 0.9$), risky agents ($\bar{\beta} < 1.2$) and on average irrational agents ($\bar{\rho} < 0.4$). Moreover, we observe in the same complex negotiation domains, time objective Q_3 benefits from having small to medium fractions of weak agents in the pre-coalition. In contrary, for such negotiation domains, it does benefit of having secondary agents in the pre-coalition, which indicates mixtures of weak and secondary agents is preferred. Moreover, in Figure 7.35 we see that higher fractions of weak or secondary agents in pre-coalitions ($> 30\%$) lead to positive added value for Q_4 in settings in which agents show risky behavior ($\bar{\beta} < 1.2$). In such settings, it increases the political support objective with at least 4%.

Finally, in Figure 7.36 we describe the interaction effects for *strong* agents in terms of overall added value QO . We observe strong agents marginally benefit from larger pre-coalition if it least covers 0.25 of all




Figure 7.35: Interaction between environments and decisions for Q_4 and secondary agents.

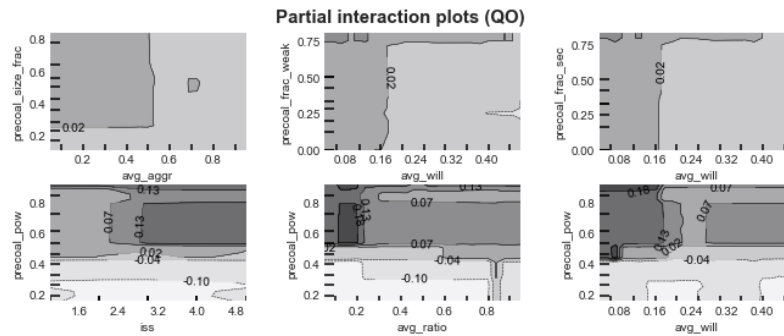


Figure 7.36: Interaction between environments and decisions for QO and strong agents.

agents in the negotiation domain, with limited aggressiveness (i.e. $\bar{p}(\Omega = \text{"Competitor"}) < 0.45$). Moreover, we see that for any fraction of weak or secondary agents in the pre-coalition, the coalition of strong agents benefit, as long as the willingness to negotiate is high (or, put differently, if $\bar{d} < 0.16$). It can also be stated that strong agents benefit substantially from more powerful pre-coalitions if the number issues per subject is high (> 3), if agents overall are relatively irrational ($\bar{\rho} < 0.2$) or if the average willingness of agents to negotiate is high $\bar{d} < 0.20$. In addition, in Figure 7.37 interactions with the composition of the pre-coalitions, the environment and strong agents' personal gains are presented. We see that the added value increases if they opt for lower fractions of weak agents (< 0.45) in their pre-coalition, when there are higher fractions of secondary agents in the negotiation environment (> 0.4). Also if average willingness to negotiate is high (or if $\bar{d} < 0.2$), strong agents' personal gains increase if they embrace around 50% of weak agents in their pre-coalitions. Moreover, if the willingness to negotiate is again high, they should opt for very low fractions of other strong agents in the pre-coalition ($\ll 0.25$ including themselves). Finally, strong agents benefit from more powerful pre-coalitions in terms of personal gains if the negotiation domain contains many issues (> 4), either the average rationality or risk is low ($\bar{\rho} < 0.2$ or $\bar{\beta} > 1.5$ respectively), fraction of power of secondary agents is relatively high (< 0.6) or if the willingness to negotiate is high ($\bar{d} < 0.2$).

7.4. Conclusion & summary

This chapter explores the simulation results by answering three intuitive questions negotiators may have: should pre-coalitions be considered in the first place? If so, in what negotiation environments are pre-coalitions sufficient? And, how should pre-coalitions be composed given their environment?

Both weak, secondary and strong agents may benefit from forming pre-coalitions. Though, the extent to which they contribute to their overall objectives differs. The vast majority of simulation results show a substantial increase in the likelihood of agreement Q_5 for all agents. Specifically for weak agents, the majority of simulations also show personal gains Q_2 are likely to increase, closely followed by their time Q_3 and political support Q_4 objectives. Additionally, we observe a slight increase in Q_2 for secondary agents, whereas their time and political support objectives seem not to benefit as much. Finally, strong agents mainly benefit from pre-coalitions by slight increases in their time and political support objectives. In contrary, strong agents loose slightly in terms of personal gains. The same holds for their social welfare objectives Q_1 , whereas no changes are observed for weak and secondary agents.

The above (overall) results do not imply that weak agents always benefit, neither does it suggest strong agents never gain from forming pre-coalitions. We have shown the negotiation environment is of great importance in explaining the overall added value of pre-coalitions. More specifically, both weak and secondary agents benefit from pre-coalitions in more or less the same negotiation environments. That is, whenever there are many agents at the table, they are willing to negotiate or act rational. Moreover, more complex

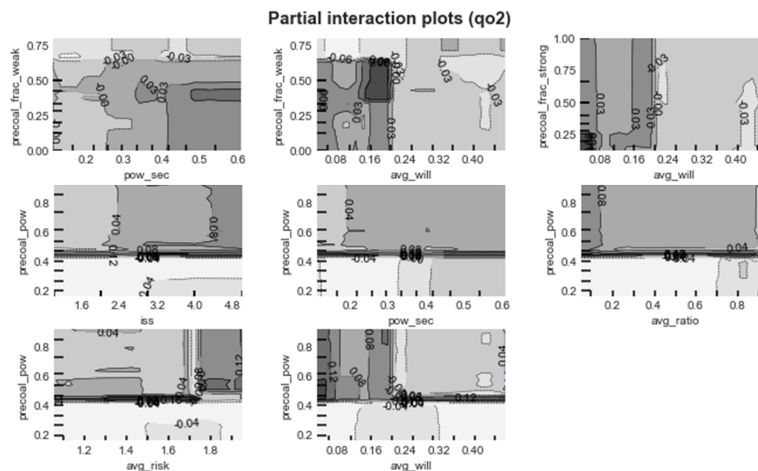


Figure 7.37: Interaction between environments and decisions for Q_2 and strong agents.

negotiation domains increase weak agents' overall objectives QO substantially, whereas secondary agents mainly benefit from pre-coalitions in terms of personal gains Q_2 , time Q_3 and political support Q_4 objectives. In contrary, for strong agents, the distribution of power within the negotiation environment is of importance for the sufficiency pre-coalitions, especially for relatively high fractions of secondary or other strong agents. Moreover, they also benefit more – especially in terms of personal gains Q_2 and time Q_3 – if package-deals are common. Finally, strong agents also benefit from pre-coalitions if agents, on average, are willing to negotiate.

To conclude, given the negotiation environment, agents may take certain decisions to make pre-coalitions more sufficient. The most obvious choice is trying to establish a more powerful pre-coalition. The vast majority of simulation results for all types of agents have shown the initial power of the pre-coalition without doubt of great influence on its sufficiency. Though, there are also negotiation environments for which we found potential evidence for either balancing, bandwagoning or mixed strategies. For *weak* agents:


- In negotiation settings with many agents (> 15), pre-coalitions become overall sufficient for high fractions of other weak agents ($> 80\%$);
- The overall added value QO increases by $> 60\%$ of secondary agents in the pre-coalition in settings with 10 or more agents;
- Between fraction of 30% and 50% strong agents in the pre-coalition increases the sufficiency for Q_2 in negotiation settings with many secondary agents ($> 50\%$);

Additionally, for *secondary* agents:

- In large negotiation domains (> 8 subjects), high fractions of secondary agents ($> 50\%$) contribute to more sufficient pre-coalitions in terms of QO ;
- Additionally, in the same large negotiation domains, up to 50% of weak agents in the pre-coalitions also makes the pre-coalition overall sufficient QO ;
- Personal gains Q_2 increases in the same manner when following the above strategies;

Finally, for *strong* agents:

- Their overall added value QO increases for any fraction of weak or secondary agents in the pre-coalition – and, thus, no other strong agents – as long as the willingness to negotiate is high ($\tilde{d} < 0.16$);
- In terms of Q_2 , added value increases if strong agents opt for low fractions of weak agents (< 0.45) in their pre-coalition, when there are high fractions of secondary agents in the environment (> 0.4).
- Also if average willingness to negotiate is high (or if $\tilde{d} < 0.2$), strong agents' personal gains increase if they embrace around 50% of weak agents in their pre-coalitions. This, while aiming for very low fractions of other strong agents in the pre-coalition ($\ll 0.25$ including themselves).



Discussion of Results

You can't handle the truth!

Colonel Nathan R. Jessup, in *A Few Good Men* (1992)

The simulation results of the previous chapter are still observations of agent behavior in the two-stage model. In this chapter, these results are aggregated to conclusions for international political negotiations. Thereby, the final supporting question is answered: how do pre-coalitions influence a state's set of objectives in different negotiation domains and settings? For answering this question, we first focus on the overall (i.e. given all objectives as defined in Section 6.1.3) sufficiency of pre-coalitions in Section 8.2. Then, in Section 8.3, we discuss those negotiation environments for which pre-coalitions might become sufficient. Finally, we discuss to what extent balancing or bandwagoning contribute to more sufficient pre-coalitions in Section 8.4. The latter is based upon the simulation evidence found in the previous chapter.

8.1. Fitted Random Forest models

For all models, the input factors are nearly uniformly distributed and most are also uncorrelated, which suggests Latin Hypercube Sampling worked (see Appendix H). Environment or decisions factors for which correlation occurred were expected as some are conceptually related. These correlated factors also behaved as expected, as they show high factor importances [165, p. 8-9]. Fortunately, this does not affect the model fit, as we know Random Forests' performance is little affected by strongly correlated predictors [165, p. 10] and as we have observed high R^2 values for all models (at least $R^2 > 0.75$).

In addition, we observe that all response variables, both the quality of outcome QO and its measures Q_x , are distributed over the $[0, 1]$ interval, which corresponds to their conceptual definitions (see Section 6.1.3). Moreover, both Q_3 , Q_4 and to a slightly lesser extent Q_5 are all "double-peaked". This is explained by the fact all measures result in 0 for no agreement – which justifies the first peak – whereas, if agreement is reached, they most likely result in outcomes > 0.5 . We also notice that Q_2 is almost never zero. This is in line with what was assumed: agents have disagreement utility d derived from their outside options. Whenever the negotiation ends in no agreement, agents fall back on their disagreement utility. This is comparable with states reverting the "status-quo", which they might prefer more [113, p. 99-105].

8.2. Overall sufficiency of pre-coalitions

The discussion of results commences with three general conclusions derived from the aggregated simulation results in Section 7.2. These conclusions describe if the initial power positions of the states assumed in this thesis affect the overall^[1] sufficiency of pre-coalitions. The conclusions are presented in Table 8.1, after which they are discussed.

1. In Figure 7.3, we have seen the median outcomes for all quality measures are substantially larger than zero, indicating weak states in general benefit from forming pre-coalitions. In addition, Bayesian

^[1]In this case, *overall* implies 'given all simulation outcome', whereas in subsequent sections we imply 'given all objectives'.

Table 8.1: General conclusions of forming pre-coalitions for weak, secondary and strong agents.

#	Overall conclusion
1.	For weak states, pre-coalitions are sufficient in the vast majority of negotiation environments as on average all their objectives are better met.
2.	For secondary states, pre-coalitions are sufficient in a large number of negotiation environments, especially when aiming for (slightly) higher personal gains and substantially higher likelihood of agreement and, to a lesser extent, for increased political support and better able to meet deadlines.
3.	For strong states, pre-coalitions are sufficient in a large number of negotiation environments, especially when aiming to meet hard deadlines, for more political support and substantially higher likelihood of agreement.

inference has showed there is almost a guaranteed gain of 5% for their time, political support and likelihood of agreement objectives, added to a 60% chance their personal goals are at least 5% better met (see Table 7.1). This is more or less completely in line with bandwagoning theory in which weak states ally with strong states to harvest some of its gains [175, p. 6-8]. That *harvesting* works for weak states is further supported by evidence found for bandwagoning (see Section 8.4). In addition, their lack of power during the actual bargaining process often causes them to be policy takers [74, 268-269]. So, presumably, by forming pre-coalitions weak states are still not policy shapers, but it does allow them to choose sufficient policy shapers for their personal objectives;

2. In Figure 7.4, we have seen the median outcomes for secondary states' personal gains and likelihood of agreement are above zero. This indicates overall they benefit from pre-coalitions whenever primarily aiming for these two objectives. These observations are again supported by calculating the Bayesian probabilities these outcomes are actually higher. We have observed there is a solid 45% chance secondary states' personal gains increase by 5% and an almost guaranteed outcome the likelihood of agreement is improved substantially as well. However, we have seen the median improvements of other objectives are less substantial. This is in line with [177, p. 12-13], who argues because secondary states are slightly more decisive in the negotiations compared to weak states, they already have better chances to meet their objectives Therefore, they should be more "picky" when forming alliances;
3. The above is observed for strong states as well. In Figure 7.5, we have seen the overall median outcomes for strong states' personal gains actually decrease when forming pre-coalitions. Also this seems reasonable and in line with, for example, [74] and [108]. Because of their structural power, strong states are policy shapers and benefit from negotiation [74, 269-274]. In addition, especially strong states should be extremely aware at what time they make concessions during the negotiations [108, p. 59-63]. So, if they benefit from the actual negotiation phase, they should perhaps not (always) make concessions by creating pre-coalitions beforehand. And when they do, they have the resources and power to be more picky [177, p. 12-13]. Though, we have seen a slight median increase in political support objectives. So, in those negotiations in which future ratification and implementation might become cumbersome [108, p. 27-28], also strong agents might want to form pre-coalitions.

These discussions raise the obvious questions in which negotiation environments pre-coalitions become (more) sufficient and to what extent states can influence this as well. Questions which are addressed in Section 8.3 and Section 8.4 respectively.

8.3. Sufficiency in different negotiation environments

In Section 7.3 we have shown many results of environment factors on the added value of pre-coalitions. Many of them are in line with other academic work stemming from both game theory and political sciences. For instance, pre-coalitions become more sufficient for all types of states by an increasing number of subjects in the negotiation domain (see, among others, Figure 7.12, Figure 7.17 and Figure 7.22). This is in line with [65, p. 22-26], who states that being member of a coalition increases someones decisiveness and, therefore, there chance of being "lucky". So, if states form pre-coalitions in negotiations with larger domains, they increase their chance of being "lucky" not once, but multiple times. Hence, most likely resulting in higher quality outcomes. Moreover, we also see that the total number of negotiating states in general is an important factor to decide to form pre-coalitions. Presumably, this is caused by the decreasing influence of single states in

larger negotiation settings. By forming pre-coalitions, decisiveness of a group of states with respect to others increases as their structural power aggregates, which is shown by some game theorists [164, p. 1138-1140]. In turn, the increased decisiveness improves their overall bargaining power in the actual negotiation phase (see Section 3.1.2), most likely resulting in better outcomes for them. In contrary, our simulations have also showed more interesting results, either supporting other studies, or providing new insights. These results have been translated to conclusions in Table 8.2. Each conclusion is discussed afterwards.

Table 8.2: Translating results to conclusions for environment factors in international political negotiations.

#	Result	#	Conclusion
a.	Almost all environment factor partial dependencies for weak states are negative for Q_4	1.	If weak states form pre-coalitions, this negatively affect the overall political support of negotiation outcomes.
b.	The way agents behave is increasingly ^[2] important to determine overall added value QO of pre-coalitions for secondary and strong states, especially in terms of willingness to negotiate, rationality and risk behavior.	2.	Compared to weak states, the way other states negotiate is increasingly important for the overall sufficiency of pre-coalitions for secondary and strong states. Especially in environments with irrational and risk taking states, pre-coalitions are sufficient in terms of political support.
c.	For secondary agents, added value of Q_3 becomes positive if the majority of states behave risky ($\beta < 1.3$) or if agents tend to negotiate less rational ($\rho < 0.5$).		
d.	For strong states, the fraction of power of other strong states has increased ^[3] explanatory power for the overall added value QO of pre-coalitions.	3.	Pre-coalitions become sufficient for strong states if there are one or more other strong states at the negotiation table. In such settings, pre-coalitions might especially increase a strong state's personal gains.
e.	For strong agents, overall added value QO become positive for high fractions of other strong agents in the negotiation environment ($> 70\%$)		
f.	For strong agents, added value of Q_2 increases substantially for high fractions of other strong agents in the negotiation environment.		
g.	For strong agents, overall added value Q_2 becomes negative if agents are willing to negotiate ($0 \leq \tilde{d} \leq 0.2$), but becomes positive if $\tilde{d} > 0.2$.	4.	For strong states to benefit from pre-coalitions, on average, other states should have good outside options.

- As shown in Figure H.8, all partial dependencies for Q_4 are negative or close to zero for the majority of environment factors and ranges, apart from the dependencies shown in Figure 7.16. This is an interesting insight, as it shows that whenever weak states are able to create a priori coalitions, this most likely negatively affect the political support of the negotiation outcomes. Perhaps this supports Martin and Mansbridge (2013, [101]) descriptions of, among others, political negotiations. They claim parties "(...) are merely trying to exercise power, exploit institutional advantages, and gain as much as possible at the expense of the other". Whenever weak states form pre-coalitions with more powerful allies, the latter may become disproportionately more decisive (see Section 3.1.2). Thus, according to [101], this forces others to compromise. We argue these compromises may be too severe, such that it negatively affects the overall political support. In turn, this may also have negative consequences for future ratification and implementation of the negotiation outcomes [108, p. 27-29], as well as for future negotiations [42, p. 568-575];
- This conclusion is based on two results. First, *b* refers to the overall added value of pre-coalitions stemming from negotiating behavior for weak and secondary agents respectively. Comparing Figure 7.9 to Figure 7.10 (both top-left) we observe negotiating behavior – i.e. willingness to negotiate, risk behavior, aggressiveness and rationality – explain approximately 36% of the added value of pre-coalitions for weak agents, compared to 43% for secondary agents. Secondly, *c* is obtained from Figure 7.34 in which we see slightly positive partial dependencies – and, thus, added value of pre-coalitions – if other agents act irrational or risky. This supports Walt's view the formation of alliances should not be solely seen in terms of power, but also in terms of aggressiveness, offensive capabilities and other opponent characteristics [175, p. 12-18]. Additionally, from [45, p. 14-15] we would expect with increasing risk behavior in the negotiation setting, in general negotiations take longer as states stay closer to their

^[2] Compared to weak agents.

^[3] Compared to both weak and secondary agents.

initial demands. Thus, forming pre-coalitions potentially allows secondary agents to 'win' negotiations while being less dependent on demands in the actual negotiation phase – i.e. with more a priori power, less votes are needed to win the subsequent negotiations. Moreover, this conclusion also shows pre-coalitions become sufficient for secondary states in negotiation settings in which state are unpredictable and might (re)act emotionally. Such settings might be those of imminent crises;

3. This conclusion is based on three results. First, d stems from a comparison between factor importances in Figure 7.9, Figure 7.10 and Figure 7.11 (all top-left). For strong agents, the fraction of other strong agents explains approximately 9.5% of the overall added value of pre-coalitions, compared to almost 8% for weak and secondary agents respectively. Secondly, d is derived from Figure 7.22 (middle) in which we see a positive partial dependency of the overall added value of pre-coalitions on higher fractions of other strong agents. Thirdly, in Figure 7.23 (top-right), we see an even stronger positive partial dependency, only now specifically for increased personal gains^[4]. High fractions of other strong states indicate a highly competitive negotiation table, as power and competitive (or hard) bargaining strategies are often related [42]. So, this conclusion shows that in such competitive environments, strong agents gain from forming pre-coalitions. Presumably because, then, they are can defend themselves from competitive tactics from equally powered opponents. Actually, this is exactly why weak agents would balance against strong states [174, p. 3-5]. Only now, a single strong state bandwagons *against* other strong states by forming pre-coalitions with other (less powerful^[5]) states. Hence, this demonstrates an alternative use of the balance-of-power theory, applied in dual, trial or other plural hegemonies. Until now, there is only little work on this type of 'great-power balancing'. Some work covering this niche includes Sweeney and Fritz (2004, [166]) and Levy and Thompson (2005, [93]).
4. As derived from Figure 7.23 (bottom-right), in which the added value of pre-coalition for strong states shows a negative dependency on the willingness to negotiate. As higher values indicate better outside options, this negative relationship thus implies pre-coalitions become sufficient – in terms of personal gains – if states are more reluctant to negotiate. This is an interesting conclusion, because of the following reasons. First, we expect stronger states to benefit from the actual negotiation phase [74, 268-269]. Moreover, in general, poor outside options result in larger outcome spaces O , as more becomes negotiable [106, p. 42]. So, if the opposite holds, the outcome space is smaller, making it harder for stronger states to benefit from the actual negotiations. We might rationalize this conclusion by assuming the pre-coalition that emerges is powerful enough to become minimal winning with other states who have worse outside options. One would expect in such setting the political support to be negatively effected. Unfortunately, we have neither evidence supporting, nor rejecting the aforementioned. This, because the obtained partial dependence shows no effect at all (see Figure H.38).

Now we have discussed some interesting conclusions defining in what environments pre-coalitions may be sufficient, we discuss decisions negotiators have, to make pre-coalitions (more) sufficient. Therefore, in the next section, we discuss evidence for either balancing, bandwagoning or mixed strategies.

8.4. Sufficiency of balancing or bandwagoning strategies

In Figure 7.27, 7.28 and 7.29, the simulation results have shown the power^[6] of pre-coalitions is without doubt of great influence on its sufficiency. Obviously, this is not a striking conclusion. Among others, many political scientists have described power is important to force issues into good outcomes, especially in negotiation settings with opposing interests [108, p. 59]. In addition, as described in Section 3.1.2, game theorists often assume direct relationships between power, decisiveness and negotiation outcomes [14, 113, 156, 163]. So, again nothing new, as long as pre-coalitions manage to stay together during the negotiation process. Then, the increased aggregated power position probably yield in better outcomes.

More interestingly are those simulation results in which negotiators can influence the sufficiency of pre-coalitions by strategically choosing the 'right' coalition partners in specific negotiation environments. In Table 8.3, we have summarized the most important results supporting either bandwagoning, balancing or mixed strategies. These will be discussed afterwards.

^[4]Note that to some extent this also holds for large fractions of secondary states, as shown in Figure 7.23 (bottom-left).

^[5]As we will see in the next section.

^[6]And size, as the two are interrelated.

Table 8.3: Translating simulation results to conclusions for weak states in international political negotiations.

#	Result	#	Conclusion
a.	In negotiation settings with many agents (> 15), pre-coalitions overall add value QO for high fractions of other weak agents (> 80%).	1.	With many states at the negotiation table, overall it becomes beneficial for weak states to <i>almost fully</i> balance with other weak or <i>mainly</i> balance with secondary states.
b.	Overall added value QO increases for weak agents by embracing around 60% or more secondary agents in the pre-coalition in settings with 10 or more agents.		
c.	In negotiation settings with many secondary agents (> 50%), weak agents should form pre-coalitions with 30% to 50% strong agents to increase the overall negotiation outcomes QO .	2.	With many secondary states at the negotiation table, pre-coalitions become sufficient if weak states <i>partly</i> bandwagon with stronger states.
d.	In larger negotiation domains (> 8 subjects), higher fractions of secondary agents (> 50%) contribute to more sufficient pre-coalitions in terms of QO .	3.	With many subjects to discuss at the negotiation table, pre-coalitions become sufficient for secondary states when <i>mainly</i> balancing with other secondary states and <i>only partly</i> balancing with weak states.
e.	In the same large negotiation domains, pre-coalition also overall add value if with no more than 50% of weak agents in the coalition.		
f.	Higher fractions of weak and secondary agents in pre-coalitions (> 40%) lead to positive added value for Q_4 in settings in which agents show risky behavior ($\tilde{\beta} < 1.2$).	4.	If secondary states aim for large political support, they should <i>mainly</i> balance with weak or secondary states in settings in which states appear to be more offensive (or risk taking).
g.	Overall added value QO increases for any fraction of weak or secondary agents in the pre-coalition – and, thus, no other strong agents – as long as the willingness to negotiate is high ($\tilde{d} < 0.16$).	5.	Whenever states have on average poor outside options, strong states should <i>never</i> balance with other strong states. Instead, they must bandwagon with weak or secondary states for pre-coalitions to become sufficient, especially in terms of personal gains.
h.	If willingness to negotiate is high ($\tilde{d} < 0.2$), strong agents' personal gains Q_2 increase if they embrace around 50% of weak agents in their pre-coalition, while aiming for low fractions of other strong agents ($\ll 0.25$).		
i.	In terms of Q_2 , added value increases if strong agents opt for significant fraction of weak agents (< 0.4) in their pre-coalition, when there are high fractions of secondary agents in the environment (> 0.4).	6.	Whenever there are many secondary states at the negotiation table, strong states should <i>only partly</i> bandwagon with weaker states to create pre-coalitions sufficiently increasing personal gains.

1. This conclusion is based on two main results (a and b), stemming from Figure 7.30 (top-right and bottom-left respectively). In both plots, we observe at least a 5% increase in the overall added value of pre-coalitions for weak states in negotiation environments with many other states, if they opt for i) mainly balancing with secondary or ii) almost fully balance with weak states. Obviously, in large negotiation settings, the number of possible pre-coalition partners increases tremendously. It is, therefore, easier to make a stance against more dominating states. This is, for example, currently observed in the region of South-East Asia, where numerous of smaller countries surrounding China – i.e. Vietnam, Taiwan and Philippines – are partly^[7] working together to avert China's rising hegemony in the region [63, 147].
2. This conclusion is based on Figure 7.30 (middle-right) in which we observe at least a 5% increase in the overall added value of pre-coalitions for weak states, if they partly bandwagon with strong opponents in negotiation environments with over 50% secondary states. Note that this implicitly assumes a mixed balancing and bandwagoning strategy, as 30% to 50% of strong states still allows for secondary states to join the (bandwagoning) coalition. This may support Panke (2011) in her conclusion why Luxembourg is sometimes more effective in European negotiations than some secondary member states – e.g. Lithuania – as it has (historic) ties with many strong and secondary countries [129, p. 311-314];
3. This conclusion is based on two main results (d and e), stemming from Figure 7.33 in which we observe a modest, though minimal 3% increase in overall added value of pre-coalitions for secondary agents when following mainly balancing strategies in large negotiation domains. This is partly in line

^[7]Partly, because there is also some conflict between these states themselves. For instance with respect to ongoing territorial disputes in the South China Sea.

with the vast majority of balance of power research, pointing that balancing has historically been the main alliance strategy for less strong states [174, p. 6]. It also supports the decisiveness argument in Section 3.1.2. A priori balancing with equally powered states does entail making concessions. However, as a result, the gained decisiveness increases the chance of being "lucky" [65, p. 22-26] – and, thus, (partly) winning the negotiations – not once, but multiple times, potentially averting more severe concessions from stronger states;

4. This conclusion is derived from both interaction plots in Figure 7.35 in which we observe at least 4% increase in added value of pre-coalitions for secondary state's political support objectives if they mainly balance with other secondary or weak states in settings in which on average states tend to negotiate more risky. This is partly in line with [175, p. 12], in which the author argues: "*(...) states that appear offensive [i.e. risk taking] are likely to provoke others to balance against them*". Weak and secondary states will most likely be the victim of such behavior, if it is conducted by strong opponents. It seems, therefore, obvious to balance against them, as there is a fair chance that they will take more risk. Though, this does not explicitly explain the increased political support. Thus, this needs some additional experiments;
5. This conclusion is based on two main results (*h* and *i*), stemming from Figure 7.36 (right-bottom) and Figure 7.37 (right-middle) respectively. Together, both results indicate strong states should never balance with other equally powered states in negotiation environments in which states in general have poor outside options, as this mainly negatively affect their personal gains. This stands in some contrast with point 4 in the previous section, in which we conclude that (on average!) strong states benefit from pre-coalitions whenever others have good outside options. However, the partial dependency of point 4 is an aggregation of all compositions of pre-coalitions in the experimental design. So, here, this conclusion is actually an exception if, and only if, strong states not balance with other strong states. Which makes sense, as by balancing they have to make strong a priori concessions (see Section 3.1.1), whereas because of the poor outside options of others they can benefit from the actual negotiation phase (see [106, p. 42] and [74, p. 268-269] respectively);
6. This conclusion is primarily derived from *j*, which in turn is based on Figure 7.37. Moreover, it implies a mixed balancing and bandwagoning strategy most likely results in sufficient pre-coalitions for strong agents when there are relatively many secondary states in the negotiation environment. This can be concluded as bandwagoning with 45% weak agents leaves enough room for at least another secondary state, but potentially also a strong state, hence balancing. This supports again the work on 'great-power balancing' (e.g. [93] and [166]).

Note that as power is a major driver of sufficiency of pre-coalitions, implicitly this suggests weak agents should in general opt for bandwagoning or large balancing pre-coalitions, whereas secondary agents might mainly balance whenever forming pre-coalitions. Indeed, we have found some direct simulation evidence supporting the latter (conclusion 3 and 4). Additionally, as we have assumed international political negotiations are multilateral in nature, we have also found few evidence supporting balancing pre-coalitions (conclusion 1). However, we have not examined the relationship between balancing and size of pre-coalitions, so we do not have the complete proof. We did find some evidence that in explicit, more competitive^[8] negotiation environments bandwagoning becomes the preferred strategy for weak agents (conclusion 2), which is in line with empiric evidence of the traditional balance-of-power theory [175, p. 6].

The raw output results in Figure 7.2 have also shown the size of winning coalitions of strong agents is slightly larger than of weak and secondary agents. Given the final power of winning coalitions is equal – see again Figure 7.2 (right) – strong agents are, thus, likely to opt for bandwagoning strategies with weak agents. More specifically, they are likely to accept bandwagoning strategies of weak agents. Indeed, we found simulation evidence supporting the latter, as a priori balancing with other strong states – especially in settings in which others are willing to negotiate – negatively affects personal gains (see conclusion 5). However, in some explicit, more competitive negotiation settings, strong agents can balance with equally powered states (conclusion 6), thus supporting the view even strong powers balance from time to time [93, 166]

However, overall, our simulation results have shown the actual negotiation environment is far more important for determining whether or not pre-coalitions become sufficient. This implies states cannot (always) influence the sufficiency of pre-coalitions, as the negotiation environment in which they act has a greater

^[8]Due to more higher fractions of more powerful secondary agents (see also conclusion 3 of the previous section).

influence on negotiation outcomes. To be precise, the only sufficient choices they have are either to stay alone or form pre-coalitions. Balancing or bandwagoning might, then, to some extent, explain the 'extra' added value, on top of forming pre-coalitions in the first place.

Finally, there are some environment factors in the balance-of-power theory – as introduced by [175, p. 8-13] – which are used to explain both balancing and bandwagoning, but of which no explicit simulation evidence is found. For example, we have not seen interactions between the (assumed) ideological proximity, certain fractions of states in the pre-coalition and its sufficiency. The same holds for offensive capabilities, as assumed relating to a state's rationality ρ ^[9].

8.5. Model validation: is the right model constructed?

The last important point to discuss is whether all conclusions above are valid, given the two-stage simulation model, experimental design and objectives of this thesis. In the following sections, we briefly discuss the conclusions of Appendix C and G.2, in which we demonstrate the impact of our assumptions on the validity of the two-stage model and experimental design respectively.

8.5.1. Two-stage model

In Appendix C, we analyzed the impact of every assumption on model validity by comparing them with existing research or providing quantitative analysis. Most important to note is that our results should *never* be generalized as the best pre-coalition strategy resulting in optimal negotiation outcomes. Instead, the conclusions indicate good strategies, most likely causing better negotiations outcomes for specific agents.

Moreover, we also show conclusions of this thesis solely apply to negotiation settings in which states have near stable preferences. So, in situations of domestic pressures – e.g. before tense national elections – our results may become irrelevant. Additionally, intergovernmental organizations need to be stable. This is often the case if they deal with global issues^[10], but not when a handful of governments work together on regional levels [108, p. 54]. Furthermore, we show that whenever negotiations in two or more distinct international arenas interact – for example, EU's refugee crisis and Brexit – our results may not be valid, at least not for those subjects most likely interacting. Moreover, if states have different currencies, the impact of exchange rates – especially in currency crises^[11] – also affect the validity of our conclusions, as negotiators face difficulties evaluating negotiation outcomes [163, p. 130]. So, results are only valid if negotiators adopt a common currency (e.g. Euro for EMU) or if a global currency is valued equally among all parties involved.

A discrepancy with real negotiations is that normally negotiators provide arguments why offers are rejected. This may lead to better counteroffers and, thus, more efficient negotiations [139, p. 343-345], which is especially true in situations in which parties are exploring their options by gathering information while negotiating [139, p. 346-347]. Therefore, conclusions in this thesis are only valid in those settings in which negotiators already have some idea about opponents' preferences. Hence, in a more mature stage of the negotiation process in which gathering of information has already silently been done by diplomats. Moreover, our conclusions need hard deadlines. Though, even Brexit is extended by a transition period postponing some decisions^[12]. So, results may become invalid if negotiators have the strong ability to delay negotiations.

To conclude, most assumptions have limited effect on model validity as it are simply decisions that need to be made, including the selection procedure and within coalition power. However, assuming final coalitions are larger than minimal winning is cumbersome as different empiric studies show conflicting results (e.g. [37, 99]). Moreover, these studies examine government formation on national or regional level. Therefore, we do not know the impact on validity of our (international) conclusions. Assuming stable preferences also affect model validity, as often preferences are the result of ongoing (domestic) negotiations [138]. Therefore, our model might be best applied in those situations in which negotiators are already part of a pre-coalition and want to explore other options. Or, the model can be applied during ongoing negotiations in which (suddenly) something in the negotiation setting changes. Three of such examples are described in Section 9.4.1.

^[9]Note that we have seen pre-coalitions are sufficient whenever states act less rational (see conclusion 2 in the previous section), but we have not seen whether this is established by following balancing or bandwagoning strategies.

^[10]To give an example, voting procedures in the Council of the European Union have changed only 5 times after the Treaty of Rome establishing the European Economic Community in 1958.

^[11]Examples include the Asian (1997) and Russian (1998) financial crises or the default of Argentina in 2006.

^[12]See Article 121 of the Draft Withdrawal Agreement [33, p. 68].

8.5.2. Experimental design

In Table 6.1 we briefly discuss why example negotiation domains resemble real-life negotiations. Unfortunately, it is hard to compare these domains with historic negotiations, as they differ for sure. Therefore, in Appendix G.2, we analyze whether our conclusions *directly* depend on example domains. We show an agent's utility is relatively independent of the chosen negotiation domain. Exceptions that occur are domains in which agents simply have low issue-specific power [59], among others caused by bad outside options. Moreover, size of the winning coalition is also not affected by the domain. However, different domains do influence negotiation time, power of the winning coalition and likelihood of agreement significantly. Therefore, conclusions of this thesis concerning the time objective Q_3 , political support Q_4 and likelihood of agreement Q_5 cannot be generalized. They are only valid in mixed issue domains, varying from scarcity to abundance and low or high number of possible outcomes.

Additionally, we examine whether the assumed power ratios between strong, secondary and weak states affect the conclusions of this thesis^[13]. We demonstrate that, apart from the political time objective Q_3 , all other objectives – and therewith the sufficiency of pre-coalitions – are dependent on the power ratios assumed for the negotiation setting. Thus, the conclusions of this thesis cannot be generalized for all power settings. In fact, they are only valid for those settings in which secondary states have some decisiveness compared to strong states, but are still clearly distinguishable from weak states. A good example of such setting is the European Union. In contrary, in settings of true hegemony, conclusions become insignificant.

8.5.3. Taking the next steps in validity

In general, there are three ways to carry out validation [91, p. 259-264]: comparison with historic data, face validation by experts or comparison with other models and studies. Above, we have mainly done the latter, whereas a comparison with data of real international political negotiations would potentially provide more concrete validity results. As this is extremely time consuming, it is beyond the scope of this research. However, we do provide an approach taking the next step in testing (and improving) model validity.

We suggest using negotiations between the EU, UK and Ireland discussing the North-Ireland border is a good case testing model validation. Although relating to the Brexit, it are clearly distinct negotiations at separate negotiation tables^[14], potentially not infringing the independent issues assumption. Moreover, all parties' demands look stable due to historic events in this region [33, p. 98-99]. The same events are also indicate the maturity of the negotiations. Finally, these (separate) negotiations are not part of the transition period [33, p. 68-69] and, therefore, have a very hard deadline.

Not only do these negotiations adhere to (the validity of) our model assumptions, they are also well documented and covered in media. Moreover, these negotiations have also the interesting property that model validity can be tested incrementally in order to reduce its complexity. First, negotiations could be modeled with the three primary negotiators, i.e. EU, UK and Ireland. Then, one might add Northern Ireland as a separate actor, as it is a devolved legislature of the UK with its own statutory powers^[15]. Additionally, one could gradually incorporate other relevant actors. For example, important trade partners of both Ireland and the UK, including the Netherlands, Germany and Belgium.

Finally, we would advise to focus on the model's primary output parameters as, among others, presented in Section 5.5.3. Negotiation outcomes are perhaps the most obvious to compare. But, as predicting actual outcomes is not the main purpose of our model, we mainly recommend to examine whether the simulated winning coalitions make sense given actual negotiation outcomes.

8.6. Conclusion & summary

This chapter discusses the most important simulation results to answer the final supporting question: how do pre-coalitions influence a state's set of objectives in different negotiation domains and settings? Many environment factors of the experimental design affecting the sufficiency of pre-coalitions were in line with other academic work. For instance, pre-coalitions become more sufficient for all states if negotiation domains are complex. This is in line with [65, p. 22-26], who states that being member of a coalition increases

[13] Which are assumed to be at least 4:2:1 and at most 24:10:1.

[14] See the Draft Withdrawal Agreement in which these negotiations are dealt with in stand-alone chapter, discussing subjects from fishery to free movements of goods [33, p. 98-105].

[15] See the Northern Ireland Act 1998 for more details at <https://www.legislation.gov.uk/ukpga/1998/47/contents>.

someones decisiveness and, therefore, there chance of being "lucky". So, if states form pre-coalitions in large domains, they increase their chance of being "lucky" not once, but multiple times. Moreover, we see the total number of negotiating states in general is an important factor to decide to form pre-coalitions. Presumably, this is caused by decreasing decisiveness of single states in larger negotiation settings. More interestingly, simulations also show negotiating behavior is increasingly important for the overall sufficiency of pre-coalitions for both **secondary and strong** states. Especially in environments with irrational and risk taking states, pre-coalitions are sufficient to increase their political support objectives. More specifically for **strong** states, pre-coalitions are sufficient if there are other strong states to negotiate with. In such settings, pre-coalitions might substantially increase personal gains. Additionally, if other negotiators have increasingly good outside options – and, are therefore less willing to negotiate – pre-coalitions become sufficient as well. In such settings, they might break the negotiation deadlock by quickly becoming minimal winning. Unfortunately, we have not found simulation evidence rationalizing the latter.

We have also been looking for results supporting either bandwagoning, balancing or mixed strategies to increase the sufficiency of pre-coalitions. Simulations have shown the aggregated power of pre-coalitions affects sufficiency without doubt, which implicitly suggests **weak agents** should in general opt for bandwagoning or large balancing pre-coalitions. Interaction dependencies only partly support this as we did not explicitly analyze the relationship between balancing and size of pre-coalitions. However, we did obtain that:

- With many states at the negotiation table, overall it becomes beneficial for weak states to almost *fully balance* with other weak or *mainly balance* with secondary states
- With many secondary states at the negotiation table, pre-coalitions become sufficient if weak states *partly bandwagon* with stronger states.

In addition, it also implies **secondary agents** might mainly balance. Indeed, we have found direct simulation evidence supporting the latter, as:

- With many subjects to discuss at the negotiation table, pre-coalitions become sufficient for secondary states when *mainly balancing* with other secondary states and *partly balancing* with weak states;
- When states appear to be more risk taking and secondary states aim for large political support, they should *mainly balance* with weak or secondary states for pre-coalitions to become sufficient.

Moreover, most simulations show **strong states** are likely to accept bandwagoning strategies of weak agents:

- Whenever states have on average poor outside options, strong states should *never balance* with other strong states. Instead, they must *bandwagon* with weak or secondary states for pre-coalitions to become sufficient, especially in terms of personal gains;
- However, whenever there are many secondary states at the negotiation table, strong states should *only partly bandwagon* with weaker states to create pre-coalitions sufficiently increasing personal gains.

The latter leaves room for *balancing* with other strong states, hence a mixed strategy. Thereby, it partly supports the scarce research on 'great-power balancing' (e.g. [93] and [166]). However, overall, our simulation results have shown the actual negotiation environment is far more important for determining whether or not pre-coalitions become sufficient. This implies states cannot (always) influence the sufficiency of pre-coalitions, as the negotiation environment in which they act has a greater influence on negotiation outcomes. To be precise, the only sufficient choices they have are either to stay alone or form pre-coalitions. Balancing or bandwagoning might, then, to some extent, explain the 'extra' added value, on top of forming pre-coalitions in the first place.

To conclude, we have discussed the validity of the above results given the two-stage model and our experimental design. Most important to note is that results of the two-stage model should never be generalized as the best pre-coalition strategy resulting in optimal negotiation outcomes. Instead, the conclusions indicate good strategies, most likely causing better negotiations outcomes for specific agents. Additionally, because of the hypothetical negotiation domain and setting of the experimental design, the above conclusions cannot be generalized in settings of true hegemony, or whenever negotiation consists of single issue types.

9.

Conclusion & Discussion

Well, nobody's perfect.
Osgood Fielding III, in *Some Like It Hot* (1959)

In the previous chapters, we put substantial effort in capturing most relevant elements of international political negotiations, abstracting a two-stage model, developing an experimental design and simulating various pre-coalition arrangements in different negotiation environments. This chapter puts all the pieces together by answering the main research question in Section 9.2. But first, we commence by providing a short recap in Section 9.1. Moreover, in Section 9.3 we give an overview of the main contributions. Subsequently in Section 9.4 we illustrate social relevance of our study, show potential use cases and describe how this study is embedded in the CoSEM (IA) master program. Finally, in Section 9.5 we reflect on this thesis by showing its limitations and describing where we believe future research effort should be directed.

9.1. Short recap

In everyday life, negotiations take place anywhere, anytime. They are also predominantly present in international politics: negotiations among governments in various intergovernmental organizations. Crucial in these negotiations is the formation of coalitions, both before and during the actual negotiations. While negotiating, these coalitions are often the result of ongoing discussions, whereas *ex ante*, they can be seen as strategic choices of negotiators. Political scientists provided the *balance-of-power* theory, rationalizing this choice in terms of "balancing" – allying against more powerful states to avert its hegemony – and "bandwagoning" – aligning with more powerful countries to harvest some of its gains [157, 175, 177]. Unfortunately, until now, they have only been applied in situations of immediate threat. Additionally, bargaining models are developed predicting coalitions under strict assumptions, suppressing many relevant context variables and neglecting strategic choices in the pre-negotiation phase. Automated negotiation might bring both fields together. However, currently, they lack automated environments supporting both multilateral multi-issue negotiation using alternating offers and the strategic choice of forming coalitions *ex ante*. Therefore, the objective of this study was to *develop* an automated negotiation environment supporting complex political contexts, able to *generate* pre-coalition arrangements and *evaluate* their sufficiency.

In the previous chapters, we have defined sufficient pre-coalitions as coalitions established prior to actual negotiation phase and substantially improving the quality of outcome of international political negotiations. In turn, the quality of outcome depends on how well states meet their own set of objectives. In this thesis, self-interested negotiators are assumed to primarily aim for maximizing (absolute) personal gains or utility. However, they may also have secondary objectives, including to maximize social welfare; minimize time to agreement; maximize political support; or maximize the likelihood of agreement. Without political support, for instance, ratification and implementation of outcomes is often impossible as "*(...) agreement may be reached, but the loser might end up wanting to sabotage its implementation.*" [108, p. 27-28].

Additionally, international political negotiations themselves can be decomposed in many components, ranging from the regimes structuring the negotiation process, voting procedures they impose and the actors they include. Following automated negotiation linguistics, this can be abstracted in both the negotiation

domain (i.e. what is discussed?) and the negotiation setting (i.e. with whom/how to negotiate?). As international political negotiations vary widely in domains and settings, we derived a hypothetical context from the balance of power theory, assuming states have three types of power positions – i.e. weak, secondary and strong – in which weak states are significantly less powerful than strong states, and, as introduced by Waltz, secondary states are: "(...) *major powers, rather than great powers*" [177, p. 44].

In this context, simulations have shown weak states benefit most from forming pre-coalitions as on average all their objectives increase. In addition, forming pre-coalitions is also sufficient for secondary states, although their personal gains improve less. In contrary, strong states mainly benefit by increased political support and time objectives, whereas their median personal gains actually decrease. Though, the variety of simulation results indicate in some negotiation environments – i.e. domains and settings – pre-coalitions may become (more) sufficient. Especially strong, and to a lesser extent secondary states should be extremely aware at what time they make concessions during the negotiation process [108, p. 59-63]. So, if they benefit from the actual negotiation phase, they should perhaps not make concessions by forming pre-coalitions beforehand. And when they do, they have the resources to be more picky [177, p. 12-13] and follow the best pre-coalition strategy. Obviously, this brings us to answering the main research question.

9.2. Answering the main research question

In the very beginning of this study, the following main research question was stated:

Research question

When are pre-coalitions sufficient in international political negotiations given a state's own set of objectives and specific power position?

To answer this question, we distinguish between the negotiation environment in which pre-coalitions may become (more) sufficient and what strategies states can follow to increase its sufficiency. In this thesis, we assume these strategies to be either balancing, bandwagoning, a combination of both or staying alone.

First, for the **environment**, simulations have shown many factors affect the sufficiency. For all types of states, pre-coalitions become more sufficient if the magnitude of negotiations increases and, hence, more subjects are discussed. Overall sufficiency increases as well when the number of states at the negotiation table increases. Both are in line with most political science theory [65, 175, 177]. More interestingly, simulations also show negotiating behavior of other states is increasingly important for the overall sufficiency of pre-coalitions for both secondary and strong states. Especially in environments in which other states tend to negotiate irrational and risky, pre-coalitions may increase the quality of outcome significantly by increasing secondary and strong states' political support objectives. Particularly for strong states, pre-coalitions are sufficient if there are (many) other strong states to negotiate with. As such settings are often more competitive [42, p. 568-575], pre-coalitions can substantially increase the personal gains of a single strong state. Additionally, if other states have good outside options – and, therefore, are less willing to negotiate – pre-coalitions become more sufficient as well. In such settings, a priori coalitions might break the deadlock by quickly becoming minimal winning.

In many other negotiation environments also weak states, and, to a lesser extent, secondary states, may take the 'gamble' by forming pre-coalitions. This, as the majority of simulation experiments have shown there is a substantial chance that pre-coalitions become sufficient and, thus, increase the overall quality of outcome. On the contrary, strong states should be cautious forming pre-coalitions, as there is a fair chance they benefit more from the actual negotiations. Unless, they opt for specific **strategies** rationalizing the pre-coalition choice in a certain negotiation environment. For strong states, these include:

- Whenever other states have on average poor outside options, strong states should *never balance* with other strong states. Instead, they must *bandwagon* with weak or secondary states for pre-coalitions to become sufficient, especially in terms of personal gains;
- Whenever there are many secondary states at the negotiation table, strong states should *only partly bandwagon* with weak states to create pre-coalitions sufficiently increasing personal gains.

The latter leaves room for *balancing* with other strong states, hence a mixed strategy. Thereby, it partly

supports the scarce research on 'great-power balancing' (e.g. [93] and [166]). In addition also **weak agents** may prefer distinct strategies:

- With many states at the negotiation table, overall it becomes beneficial for weak states to almost *fully balance* with other weak or *mainly balance* with secondary states;
- With many secondary states at the negotiation table, pre-coalitions become sufficient if weak states *partly bandwagon* with stronger states.

Finally, we have also found simulation evidence for two strategies **secondary agents** should follow:

- With many subjects to discuss at the negotiation table, pre-coalitions become sufficient for secondary states when *mainly balancing* with other secondary states and *partly balancing* with weak states;
- When states appear to be more risk taking and secondary states aim for large political support, they should *mainly balance* with weak or secondary states for pre-coalitions to become sufficient.

9.2.1. The balance-of-power theory: a reflection

So, did the concept of 'balancing' and 'bandwagoning' explain the a priori choice of coalition partners in international political negotiations? This, instead of threatening (near war) situations. Based upon the simulation results and the main conclusion above, we can only state the balance-of-power theory *partly* explains this choice. In the majority of simulation experiments, an (undefined^[1]) mix between both strategies is preferred for pre-coalitions to become sufficient. In some specific negotiation environments, it does explain the added value of pre-coalition. For example, in general, we obtained that weak states should mainly balance with other weak or secondary states, as international political negotiations are assumed to be multilateral in nature. This is in line with the majority of evidence on balancing, stating that balancing is more common than bandwagoning for weak states [175, p. 15-18].

However, overall, our simulation results have shown the actual negotiation environment is far more important for determining whether or not pre-coalitions become sufficient. The latter thus implies that states cannot (always) influence the sufficiency of pre-coalitions themselves, as the negotiation environment in which they act has a greater influence on negotiation outcomes. Or, to be precise, the only sufficient choices they have are either to stay alone or form pre-coalitions. Then, balancing or bandwagoning might, to some extent, only explain the 'extra' added value, on top of forming pre-coalitions in the first place.

9.3. Scientific contributions

The aim of this thesis was to develop an automated negotiation environment supporting complex international political contexts, and that is able to generate pre-coalition arrangements and evaluate their sufficiency. Therefore, the contributions of this thesis to scientific research are the following (in no distinct order):

- **Introducing an approach to align multilateral bargaining models and empiric political research (Chapter 2, 3 and 4)**^[2].

As stated by Bolton and Croson (2012, [23, p. 85-86]) and endorsed by ourselves, there is a need to diminish the gap between theory on multilateral bargaining and empiric and experimental research. This thesis can be seen as both an approach as well as first steps in aligning the two: we have used political case studies to sketch contours of a potential model (Chapter 2); derived rational choices with the bargaining theory in mind (Chapter 3) and; created contours of a seemingly unpredictable and capricious negotiation environment in which experiments are conducted (Chapter 4). Thus, starting from reality and abstracting to theory, contrary to extending theory back to reality;

- **Modeling the formation of pre-coalitions in an adaptable and open-source automated negotiation environment (Chapter 5).**

Most scientific work focuses on modeling the formation of coalitions during, or as a result of, actual negotiation processes. There are some models analyzing the strategic decision of choosing a priori

^[1]Undefined, as the overall power of the pre-coalition is more important than the states delivering that power.

^[2]Note that the actual content of these chapters do make very little contributions to current research. Instead, it is the overall approach to align multilateral bargaining models and actual political research.

coalition partners in elections, but these have been only applied to regional or national levels (e.g. [13, 111, 133]). Additionally, to our best knowledge, there are no automated negotiation models developed supporting the strategic decision of choosing coalition partners *ex ante*. In fact, most of them not even support coalition formation in general [159, p. 220-222]. Therefore, the two-stage simulation model in Chapter 5 can be seen as a major contribution, in which agents first 'bargain' over the formation of pre-coalition based on ideological values and power positions, after which multiple rounds of actual negotiations are held;

- **Developing evaluation methods for determining the sufficiency of pre-coalitions (Chapter 6).** Most (multilateral) bargaining models aim at finding equilibrium solutions in which agents rationally maximize their own or collective utility [73, p. 1-2]^[3]. In this thesis, we argue this is perhaps too blunt as real negotiators also have one or more secondary objectives. For instance, a single state may really want to obtain good results on short notice, because of pending national elections. Or, they want to increase political support, as without it, ratification and implementation of international legislation is often impossible [108, p. 27-28]. Therefore, in Chapter 6, we provided an evaluation method combining both primary and secondary objectives;
- **Drawing parallels between balance-of-power theory and diplomatic negotiations (Chapter 8).** In the beginning of this thesis, we have introduced the balance-of-power theory as a widely applied method to rationalize strategic alliance formation in situations of imminent *threat* [175, p. 3]. Moreover, we have stated this theory might explain a priori coalition choices in (luckily) less threatening international political negotiations. By doing so, in Chapter 8 we have drawn parallels between the balance-of-power theory in its traditional application and more diplomatic negotiations. Additionally, in Chapter 6, some important context factors – as introduced by Walt (see [175, p. 8-13]) – are linked to elements in international political negotiations. Combined together, both can be seen as first steps bringing the balance-of-power theory into the world of negotiations, instead of (near) war.

9.4. Social relevance: why should we bother?

Politics and coalitions are inseparable. Even dictators cannot rule alone and must ally with societal elites to impose his will [173, p. 1]. In broad sense, decades of research studying coalitions by political scientists have helped to enforce decisions, passing bills, winning elections, achieving (climate) goals and more. So, in best case, the social relevance of this research and the developed tool originates from its ability to expand the view of negotiators and find new or better partners to work with. More realistically, the answers given in this study might start, support, explore and guide discussions about *ex ante* coalition formation between one or more government representatives. In addition, the use of open-source technologies might enable other researchers to freely adapt and extend our tool for their work, thus potentially accelerating future research.

9.4.1. Who should use the two-stage simulation model?

To better understand why tools like ours might be relevant for real international political negotiations, we give some actual example settings in which they can be applied. For instance at the start of this year (2018), The Netherlands have changed seats with Italy for one year to become member of the primary international body for peace and security, UN's Security Council. Although The Netherlands have been member of the UN since its existence in 1945, it had never had such a 'powerful' place in an important council. Therefore, it has perhaps little awareness of strategic coalition choices to support their objectives. Our tool can provide insights and increasing awareness, by exploring different pre-coalition arrangements and calculate their added value. Of course, given the knowledge The Netherlands have of their counterparts.

Another good example is (or was) the ongoing enlargement of the European Union. Currently, Albania, the former Yugoslav Republic of Macedonia, Montenegro, Serbia and Turkey, are all candidate member states, possibly entering ongoing EU negotiations. Not only can this change the sufficiency of already established coalitions^[4], it also provides new strategic coalition options. Our tool is able to predict – naturally, within its own limits (see the following section) – the impact on already established coalitions by the entrance of new member states. Additionally, it can explore the new strategic options to get insight whether or not new

^[3]Own unpublished work.

^[4]Among others, currently, a liberal and conservative (pre-)coalition can be distinguished for negotiations concerning EU subsidies to maintain the agricultural production [129, p. 311-312].

member states might be helpful in achieving your own goals (and vice versa).

Finally, as one of the first things in office, the 45th president of the United States, Donald Trump, withdrew the US from the Trans-Pacific Partnership (TPP) on 23rd of January 2017. According to many news sources, this created much turmoil between its left-over members [62]. The withdrawal of such an important state from this trade agreement initiated re-negotiations between its original members, but it also triggered other (non-member) countries to announce interest or put forward alternatives (e.g. RCEP). Although perhaps slightly more far-fetched, in potential our tool could also explore the strategic coalitional options in this new situation by i) showing the impact of the US leaving the agreement on future re-negotiations and ii) exploring whether or not those parties who announced interest would potentially help you to achieve your own goals. Though, this last example is slightly in contrast with model validity, as we have stated model results are only valid in those settings in which negotiators already have some idea about opponents' preferences. Hence, in a more mature stage of the negotiation process. So, for using our tool in the last example, at least you should have some global idea about preferences of potential new members. This may be the case if you already have good diplomatic relations with those countries.

Additionally to the examples above – in which the actual negotiation environment changed – negotiators might also be interested in using our tool to 'think outside the box'. It is very natural to work together with those states you work together with for a long time [36, p. 5]. However, the most natural way might not always be a good way to proceed, especially as subjects to discuss changes. Our tool has the predictive power to show negotiators what other strategic coalition options are possible and, thereby, potentially expanding their views (and coalition strategies).

9.4.2. Link to master program

The research links to the CoSEM (IA) study program in the following ways. First, it has a clear and systematic design component as we start with the complex process of states negotiating with each other; deriving assumptions both in terms of this process as well as how actors (theoretically) behave; transform these assumptions to practical – software oriented – model requirements and; implementing them in a final software application. But, as taught in the IA track, it does not stop there as, while implementing, I kept in mind important implementation requirements such as scalability, modularity, maintainability to ensure future organizational and/or academic alignment. Furthermore, I have used data analyses techniques that has been taught both during the IA program and my specialization track abroad. In addition, the study is multidisciplinary in nature, considering and combining technical, institutional, economic and social knowledge and, therewith, clearly linking the technology to a public domain. Finally, a critique of this thesis could be that process management strategies and assessing the impact of my tool in organizations is underexposed. Although I agree on this, clear choices had to be made. This thesis already makes a great step in diminishing the gap between multilateral negotiation models and the institutional reality, both from a bargaining theory and automated negotiation perspective. Thus, embedding this work in actual organizations and assessing its impact would be a great next step.

9.5. Discussion: nobody's perfect

As Osgood Fielding III stated at the end of *Some Like It Hot* and at the beginning of this chapter: "Well, nobody's perfect". Not only did this apply to Jerry ("Daphne") the double bass player, it also concerns this work. Therefore, we conclude this thesis by addressing its limitations and show future research directions.

9.5.1. Limitations of this thesis

Of course, most limitations of this research are caused by the assumptions made in Chapter 2, 3 and 4. As shown, they affect the validity of the model, as well as of the experimental design. Affecting model validity, for instance, negotiators are assumed to be unitarian actors, representing the exact preferences of the states they represent. Because of this assumption, conclusions are only valid in negotiation settings in which states have near stable preferences. Naturally, this is a bold assumption, as presumably preferences result from ongoing (national) negotiations and discussions, and continuously change along the way.

Another example affecting model validity is disregarding the influence of the agenda both in the pre- and actual negotiation phases. We assumed the agenda is exogenously given, whereas in most real-life

negotiations it is endogenously part of the process. More specifically, negotiators may exploit the agenda in their advantage or bring new issues to the negotiation table [36, p. 49-54]. Thus, our model lacks flexibility in changing the agenda, both before and during negotiations. Another consequence of this limitation is that package-deals are only possible within the same subjects. So, compensating losses in the allocation of refugees within EU member states by, for example, infrastructural projects – two seemingly unrelated subjects – is not (really^[5]) possible.

Assumptions regarding the evaluation of outcomes by utility functions also lead to severe limitations. Not only have we assumed that subjects are evaluated independently – and, therefore, separate outcomes do not influence each other — there is also substantial body of knowledge contradicting states aim for high *absolute* gains. In contrary, when evaluation their gains, they compare their absolute gains with other states, hence strive for high *relative* gains [110, p. 471]. Though, in general, either striving for relative or absolute gains should not make (much) difference for cooperation among nations [110, p. 472].

Additionally, in Section 8.5, we have shown the assumed power positions – i.e. weak, secondary and strong – affect the validity of the experimental design and, thus, our conclusions with respect to balancing and bandwagoning in international political negotiations. This, as the assumed relations between each of these power positions directly affect negotiation outcomes. As a result, conclusions of this thesis cannot be generalized for all power settings. For example, in situations of true hegemony, conclusions become invalid. In addition, they also not apply to "micro-states"^[6], as, for example, introduced by Panke (2011, [129]).

Finally, the means of analysis (i.e. using Random Forests) to describe relations in the experimental simulations proceed in some limitations. Some argue that ensembles of decision trees are merely used predictively, rather than descriptively. Indeed, by using Random Forests, we lose the statistical power^[7] of, for example, linear regression models. As a consequence, it becomes hard to make relationships explicit, which is a serious limitation of the conclusions in this chapter. However, in return, Random Forests have enabled us to examine and visualize all kinds of non-linear dependencies and (high-order) interaction effects effortlessly.

9.5.2. Future research directions

In this section, we describe to where we believe future research effort should be directed. First, we provide directions *improving* – or further investigating – some conclusions derived in Chapter 8, which include:

- We have obtained contradictory results concerning balancing behavior of weak states in aggressive situations. The balance-of-power theory prescribes even weak states balance in settings with aggressive strong powers [174, p. 12-13]. Our simulation experiments have shown they do not. More simulations – including different factors – or linking the results to case studies is needed;
- To reach higher social welfare, our model has shown weak and secondary states need to balance with each other. We presume this is caused by the formation of a defensive block, forcing strong states to make concessions. However, further research should try to link the latter to practical case studies;
- Finally, results show strong states should partly balance with other strong states if there are many secondary states at the table. Again, it is interesting to link this result to empiric examples (see [93, 166]).

We have also directions *extending* our work and tool by changing assumptions, altering negotiation protocols and more. As there are many of these suggestions, we present some interesting ones:

- Expanding the model by forming the agenda endogenously in the pre-negotiation phase would be an interesting and realistic extension of the model. There exist some examples in literature, including [68] and [46]. Though, only for bilateral negotiations;
- Another interesting expansion, especially considering international political negotiations, would be to incorporate veto rights;

^[5]It is only possible if they have the same "overarching" subject.

^[6]States that are even weaker than weak states, such as Lichtenstein, Luxembourg or Cyprus.

^[7]One can argue if this statistical power could have been used in the first place, as we have created a model that intendedly creates randomness.

- Moreover, in our model we assume every negotiator is able to communicate (and make offers) to all other negotiators. Of course, real-world negotiators do not always have the means of communication to do so. Therefore, an interesting extension is bargaining in networks (e.g. see [1, 2, 100]);
- States are also often involved in many international negotiations at the same time. Thus, potentially they are quite likely to be part of more than one pre-coalition. As these negotiations partly overlap and evolve sequentially, one can imagine outcomes of one negotiation to influence the ex ante coalition choice of another. An interesting expansion is therefore to expand our model by bargaining in sequential networks (e.g. see [58, 102, 103]);
- Finally, in our model we have assumed agents negotiate by alternating offers. Subsequently, all the information they have about their opponent is derived from those offers. Especially in politics, *motivating* offers is extremely important. Therefore, another interesting extension would be to incorporate a "alternating arguments" protocol. Almost similar to arguments (or "chats") real negotiators make. This is considered to be the next frontier in agent-based automated negotiation (e.g. see [187]).

Bibliography

- [1] Dilip Abreu and Mihai Manea. Bargaining and efficiency in networks. *Journal of Economic Theory*, 147(1):43 – 70, 2012. ISSN 0022-0531. doi: <https://doi.org/10.1016/j.jet.2011.11.003>.
- [2] Dilip Abreu and Mihai Manea. Markov equilibria in a model of bargaining in networks. *Games and Economic Behavior*, 75(1):1 – 16, 2012. ISSN 0899-8256. doi: <https://doi.org/10.1016/j.geb.2011.09.004>.
- [3] Gregory Adams, Gordon Rausser, and Leo Simon. Modelling multilateral negotiations: An application to California water policy. *Journal of Economic Behavior & Organization*, 30(1):97–111, JUL 1996. ISSN 0167-2681. doi: {10.1016/S0167-2681(96)00844-X}.
- [4] Rebecca Adler-Nissen and Vincent Pouliot. Power in practice: Negotiating the international intervention in Libya. *European Journal of International Relations*, 20(4):889–911, DEC 2014. ISSN 1354-0661. doi: {10.1177/1354066113512702}.
- [5] Deniz Aksoy. Who gets what, when, and how revisited: Voting and proposal powers in the allocation of the EU budget. *European Union Politics*, 11(2):171–194, 2010. doi: 10.1177/1465116510363658.
- [6] José C.R. Alcantud. A measure of utility levels by euclidean distance. *Decisions in Economics and Finance*, 25 (1):65–69, May 2002. ISSN 1129-6569. doi: 10.1007/s102030200004.
- [7] Bedour Alrayes, Özgür Kafalı, and Kostas Stathis. Recon: A robust multi-agent environment for simulating concurrent negotiations. In Naoki Fukuta, Takayuki Ito, Minjie Zhang, Katsuhide Fujita, and Valentin Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 157–174. Springer International Publishing, Cham, 2016. ISBN 978-3-319-30307-9. doi: 10.1007/978-3-319-30307-9_10.
- [8] Dan Ariely. *Predictably Irrational, Revised and Expanded Edition: The Hidden Forces That Shape Our Decisions*. Harper Collins, NY: New York, 2010. ISBN 9780061353246.
- [9] R. J. Aumann and J. H. Dreze. Cooperative games with coalition structures. *International Journal of Game Theory*, 3(4):217–237, Dec 1974. ISSN 1432-1270. doi: 10.1007/BF01766876.
- [10] Robert M. Axelrod. *Conflict of interest: a theory of divergent goals with applications to politics*. Markham political science series. Markham Publishing Company, 1970.
- [11] Tim Baarslag. *What to Bid and When to Stop*. Dissertation, Delft University of Technology, Sep 2014. URL <http://dx.doi.org/10.4233/uuid:3df6e234-a7c1-4dbe-9eb9-baadabc04bca>.
- [12] Stefanie Bailer. Strategy in the climate change negotiations: do democracies negotiate differently? *Climate Policy*, 12(5):534–551, 2012. doi: 10.1080/14693062.2012.691224.
- [13] Siddhartha Bandyopadhyay, Kalyan Chatterjee, and Tomas Sjöström. Pre-electoral coalitions and post-election bargaining. *Quarterly Journal of Political Science*, 6(1):1–53, 2011. ISSN 1554-0626. doi: 10.1561/100.00010043.
- [14] John F. Banzhaf. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Review*, 19(2): 317–343, 1965. ISSN 0036-0465.
- [15] David P. Baron and John A. Ferejohn. Bargaining in legislatures. *The American Political Science Review*, 83 (4):1181–1206, 1989. ISSN 00030554, 15375943. URL <http://www.jstor.org/stable/1961664>.
- [16] Scott Barrett. *Environment and Statecraft: The Strategy of Environmental Treaty-Making: The Strategy of Environmental Treaty-Making*. Oxford University Press, New York, 2003. ISBN 9780191531446. doi: DOI:10.1093/0199286094.001.0001.
- [17] Brian Barry. Is it better to be powerful or lucky? part i. *Political Studies*, 28(2):183–194, 1980. ISSN 1467-9248. doi: 10.1111/j.1467-9248.1980.tb01244.x.
- [18] Richard E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.

- [19] Stefano Benati and Giuseppe Vittucci Marzetti. Probabilistic spatial power indexes. *Social Choice and Welfare*, 40(2):391–410, Feb 2013. ISSN 1432-217X. doi: 10.1007/s00355-011-0608-4.
- [20] Elaine Bennett. Multilateral bargaining problems. *Games and Economic Behavior*, 19(2):151 – 179, 1997. ISSN 0899-8256. doi: <https://doi.org/10.1006/game.1997.0547>.
- [21] G. Berridge, M. Keens-Soper, and T. Otte. *Diplomatic Theory from Machiavelli to Kissinger*. Studies in Diplomacy. Palgrave Macmillan UK, 2001. ISBN 9780230508309. URL <https://books.google.nl/books?id=DSqHDAAAQBAJ>.
- [22] Tom Blockmans and Marie-Anne Guerry. Probabilistic spatial power indexes: The impact of issue saliences and distance selection. *Group Decision and Negotiation*, 24(4):675–697, 2015. ISSN 1572-9907. doi: 10.1007/s10726-014-9408-4.
- [23] G.E. Bolton and R.T.A. Croson. *The Oxford Handbook of Economic Conflict Resolution*. Oxford Handbooks in Economics. Oxford University Press, 2012. ISBN 9780199730858.
- [24] Pierre Bourque and Richard E. Fairley, editors. *SWEBOK V3.0: Guide to the Software Engineering Body of Knowledge*. IEEE Computer Society, Los Alamitos, CA, 3rd edition, 2014. ISBN 978-0-7695-5166-1.
- [25] Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. In Gerhard Weiss, editor, *Multiagent Systems*, chapter 6, pages 213–259. MIT Press, 2013. ISBN 9780262018890.
- [26] Eric C. Browne and Peter Rice. A bargaining theory of coalition formation. *British Journal of Political Science*, 9(1):67–87, 1979. ISSN 14692112.
- [27] Barry C. Burden. Deterministic and probabilistic voting models. *American Journal of Political Science*, 41(4): 1150–1169, 1997. ISSN 00925853, 15405907.
- [28] Stephan Cass. The 2017 top programming languages. *IEEE Spectrum*, 2017. URL <https://spectrum.ieee.org/computing/software/the-2017-top-programming-languages>.
- [29] Suchan Chae and Jeong-Ae Yang. The unique perfect equilibrium of an n-person bargaining game. *Economics Letters*, 28(3):221 – 223, 1988. ISSN 0165-1765. doi: [https://doi.org/10.1016/0165-1765\(88\)90118-8](https://doi.org/10.1016/0165-1765(88)90118-8).
- [30] Pamela S. Chasek and Joanna Depledge. Raising the tempo: The escalating pace and intensity of environmental negotiations. In Pamela S. Chasek and Lynn M. Wagner, editors, *The Roads from Rio: 20 years of global environmental negotiations*, chapter 2, pages 150–172. RFF Press, 2012. ISBN 978-0-415-80975-7.
- [31] Kalyan Chatterjee, Bhaskar Dutta, Debraj Ray, and Kunal Sengupta. A noncooperative theory of coalitional bargaining. *The Review of Economic Studies*, 60(2):463–477, 1993. ISSN 00346527, 1467937X.
- [32] European Commission. European commission statement following the vote of the european parliament in favour of an emergency relocation mechanism for a further 120,000 refugees. Press Release, September 2015. STATEMENT/15/5664, http://europa.eu/rapid/press-release_STATEMENT-15-5664_en.pdf.
- [33] European Commission. Draft withdrawal agreement on the withdrawal of the united kingdom of great britain and northern ireland from the european union and the european atomic energy community. Draft document, February 2018. TF50 (2018) 33, https://ec.europa.eu/commission/sites/beta-political/files/draft_withdrawal_agreement.pdf.
- [34] Philip B. Crosby. *Quality is Free: The Art of Making Quality Certain*. Mentor book. McGraw-Hill, 1979. ISBN 9780070145122.
- [35] Hans de Bruijn and Ernst T. ten Heuvelhof. Policy analysis and decision making in a network: how to improve the quality of analysis and the impact on decision making. *Impact Assessment and Project Appraisal*, 20(4): 232–242, 2002. doi: 10.3152/147154602781766627.
- [36] Hans de Bruijn and Ernst T. ten Heuvelhof. *Management in Networks: On multi-actor decision making*. Taylor & Francis, 2012. ISBN 9781134042241.
- [37] Sarah L. de Lange. New alliances: Why mainstream parties govern with radical right-wing populist parties. *Political Studies*, 60(4):899–918, 2012. doi: 10.1111/j.1467-9248.2012.00947.x.
- [38] Joanna Depledge. *The Organization of Global Negotiations: Constructing the Climate Change Regime*. Taylor & Francis, 2013. ISBN 9781849773171.

- [39] Effrosyni Diamantoudi and Licun Xue. Coalitions, agreements and efficiency. *Journal of Economic Theory*, 136(1):105 – 125, 2007. ISSN 0022-0531. doi: <https://doi.org/10.1016/j.jet.2006.02.012>. URL <http://www.sciencedirect.com/science/article/pii/S0022053107000415>.
- [40] Christian Downie. *The politics of climate change negotiations : strategies and variables in prolonged international negotiations*. Edward Elgar Publishing, Cheltenham UK, 2014. ISBN 9781783472109.
- [41] Peter Drahos. When the weak bargain with the strong: Negotiations in the world trade organization. *International Negotiation*, 8(1):79–109, 2003. doi: 10.2139/ssrn.418480.
- [42] Andreas Dür and Gemma Mateo. Bargaining power and negotiation tactics: The negotiations on the eu's financial perspective, 2007–13. *JCMS: Journal of Common Market Studies*, 48(3):557–578, 2010. ISSN 1468-5965. doi: 10.1111/j.1468-5965.2010.02064.x.
- [43] Eithan Ephrati and Jeffrey S. Rosenschein. The Clarke Tax as a consensus mechanism among automated agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, volume 1, pages 173–178. AAAI Press, 1991. ISBN 0-262-51059-6.
- [44] European Medicines Agency. Ema to relocate to amsterdam, the netherlands. Press Release, November 2017. http://www.ema.europa.eu/docs/en_GB/document_library/Press_release/2017/11/WC500238862.pdf.
- [45] Peyman Faratin, Carles Sierra, and Nick R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3):159 – 182, 1998. ISSN 0921-8890. doi: [https://doi.org/10.1016/S0921-8890\(98\)00029-3](https://doi.org/10.1016/S0921-8890(98)00029-3). Multi-Agent Rationality.
- [46] Shaheen Fatima and Ahmed Kattan. Evolving optimal agendas for package deal negotiation. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation, GECCO '11*, pages 505–512, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0557-0. doi: 10.1145/2001576.2001646.
- [47] Shaheen S. Fatima, Michael Wooldridge, and Nicholas R. Jennings. Multi-issue negotiation under time constraints. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1, AAMAS '02*, pages 143–150, New York, NY, USA, 2002. ACM. ISBN 1-58113-480-0. doi: 10.1145/544741.544775.
- [48] Shaheen S. Fatima, Michael Wooldridge, and Nicholas R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1 – 45, 2004. ISSN 0004-3702. doi: [https://doi.org/10.1016/S0004-3702\(03\)00115-2](https://doi.org/10.1016/S0004-3702(03)00115-2).
- [49] Shaheen S. Fatima, Sarit Kraus, and Michael Wooldridge. *Principles of Automated Negotiation*. Cambridge University Press, 2015. doi: 10.1017/CBO9780511751691.
- [50] Dan S. Felsenthal and Mosché Machover. *The Measurement of Voting Power: Theory and Practice, Problems and Paradoxes*. Edward Elgar, 1998. ISBN 9781858988054.
- [51] Chaim Fershtman. A note on multi-issue two-sided bargaining: Bilateral procedures. *Games and Economic Behavior*, 30(2):216 – 227, 2000. ISSN 0899-8256. doi: <https://doi.org/10.1006/game.1999.0727>.
- [52] Franklin M. Fisher. Games economists play: A noncooperative view. *The RAND Journal of Economics*, 20(1): 113–124, 1989. ISSN 07416261.
- [53] Ian Fleming. Defining software quality characteristics to facilitate software quality control and software process improvement. In Ivan Mistrik, Richard Soley, Nour Ali, John Grundy, and Bedir Tekinerdogan, editors, *Software Quality Assurance*, pages 47 – 61. Morgan Kaufmann, Boston, 2016. ISBN 978-0-12-802301-3. doi: <https://doi.org/10.1016/B978-0-12-802301-3.00003-X>.
- [54] Pau Fonseca i Casas. *Formal Languages for Computer Simulation: Transdisciplinary Models and Applications: Transdisciplinary Models and Applications*. Research essentials. Information Science Reference, 2013. ISBN 9781466643703.
- [55] Y. Funaki and T. Yamato. Sequentially stable coalition structures. CentER Discussion Paper 2007-96, Tilburg University, 2007.
- [56] Pruitt Dean G. Indirect communication and the search for agreement in negotiation. *Journal of Applied Social Psychology*, 1(3):205–239, 1971. doi: 10.1111/j.1559-1816.1971.tb00363.x.

- [57] William A. Gamson. A theory of coalition formation. *American Sociological Review*, 26(3):373–382, 1961. ISSN 00031224.
- [58] Liang Guo and Ganesh Iyer. Multilateral bargaining and downstream competition. *Marketing Science*, 32(3): 411–430, 2013. doi: 10.1287/mksc.1120.0766.
- [59] William M. Habeeb. *Power and Tactics in International Negotiation: How Weak Nations Bargain with Strong Nations*. Johns Hopkins University Press, 1988. ISBN 9780801836206.
- [60] Emilie M. Hafner-Burton, Miles Kahler, and Alexander H. Montgomery. Network analysis for international relations. *International Organization*, 63(3):559–592, 2009. doi: 10.1017/S0020818309090195.
- [61] Maurice H. Halstead. *Elements of software science*. Operating and programming systems series. Elsevier, 1977. ISBN 9780444002150.
- [62] Robin Harding. Trump withdrawal from TPP creates high stakes turmoil, May 2017. URL <https://www.ft.com/content/37f1dcb8-0288-11e7-aa5b-6bb07f5c8e12>. Accessed on the 15th of May 2018.
- [63] Kai He. Institutional balancing and international relations theory: Economic interdependence and balance of power strategies in southeast asia. *European Journal of International Relations*, 14(3):489–518, 2008. doi: 10.1177/1354066108092310.
- [64] Jean-Jacques Herings and Arkadi Predtetchinski. One-dimensional bargaining with markov recognition probabilities. *Journal of Economic Theory*, 145(1):189 – 215, 2010. ISSN 0022-0531. doi: <https://doi.org/10.1016/j.jet.2009.10.002>.
- [65] Manfred J. Holler and Edward W. Packel. Power, luck and the right index. *Zeitschrift für Nationalökonomie / Journal of Economics*, 43(1):21–29, 1983. ISSN 00443158, 23048360.
- [66] Jo Hoogslag. How to win the game? strategic, simultaneous, many-to-many, non-cooperative negotiation. Master's thesis, University of Technology Delft, 2012.
- [67] Chen-Ying Huang and Tomas Sjöström. The recursive core for non-superadditive games. *Games*, 1(2):66–88, 2010. ISSN 2073-4336. doi: 10.3390/g1020066.
- [68] Younghwan In and Roberto Serrano. Agenda restrictions in multi-issue bargaining. *Journal of Economic Behavior & Organization*, 53(3):385 – 399, 2004. ISSN 0167-2681. doi: [https://doi.org/10.1016/S0167-2681\(03\)00087-8](https://doi.org/10.1016/S0167-2681(03)00087-8).
- [69] ISO/IEC 25010. Systems and software engineering—systems and software quality requirements and evaluation (SQuaRE)—system and software quality models. International Organization for Standardization, Geneva, Switzerland, 2011. URL <https://www.iso.org/standard/35733.html>.
- [70] ISO/IEC IS 9126. Software product evaluation: Quality characteristics and guidelines for their use. International Organization for Standardization, Geneva, Switzerland, 1991. URL <https://www.iso.org/standard/22749.html>.
- [71] Takayuki Ito, Mark Klein, and Hiromitsu Hattori. A multi-issue negotiation protocol among agents with non-linear utility functions. *Multiagent and Grid Systems - Negotiation and Scheduling Mechanisms for Multiagent Systems*, 4(1):67–83, January 2008. ISSN 1574-1702.
- [72] Tasos Kalandrakis. Equilibria in sequential bargaining games as solutions to systems of equations. *Economics Letters*, 84(3):407 – 411, 2004. ISSN 0165-1765. doi: <https://doi.org/10.1016/j.econlet.2004.03.012>.
- [73] Jesse Kaptein. A survey on multilateral negotiation models with alternating offers protocol. University of Technology in Delft, 2018.
- [74] Milana Karayanidi. Bargaining power in multilateral negotiations on intellectual property rules: Paradox of weakness. *The Journal of World Intellectual Property*, 14(3-4):265–275, 2011. ISSN 1747-1796. doi: 10.1111/j.1747-1796.2011.00418.x.
- [75] Heikki Kauppi and Mika Widgrén. Voting rules and budget allocation in the enlarged eu. *European Journal of Political Economy*, 23(3):693 – 706, 2007. ISSN 0176-2680. doi: <https://doi.org/10.1016/j.ejpoleco.2006.07.001>.
- [76] Maurice G. Kendall and Alan Stuart. *The Advanced Theory of Statistics: Distribution theory*. The Advanced Theory of Statistics. Macmillan, 1977. ISBN 9780028476308.

- [77] Robert O. Keohane. Comment: Norms, institutions, and cooperation. In Keith E. Maskus and Jerome H. Reichman, editors, *International Public Goods and Transfer of Technology Under a Globalized Intellectual Property Regime*, page 65–68. Cambridge University Press, 2005. doi: 10.1017/CBO9780511494529.004.
- [78] Robert O. Keohane. *International Institutions And State Power: Essays in International Relations Theory*. Westview Press, 1989. ISBN 9780813308371.
- [79] Gregory E. Kersten and Tomasz Szapiro. Generalized approach to modeling negotiations. *European Journal of Operational Research*, 26(1):142 – 149, 1986. ISSN 0377-2217. doi: [https://doi.org/10.1016/0377-2217\(86\)90166-9](https://doi.org/10.1016/0377-2217(86)90166-9). Second EURO Summer Institute.
- [80] Steven Ketchpel. Forming coalitions in the face of uncertain rewards. In *Proceedings of the Twelfth National Conference on Artificial Intelligence, AAAI '94*, pages 414–419, Menlo Park, CA, USA, 1994. American Association for Artificial Intelligence. ISBN 0-262-61102-3.
- [81] Barbara Kitchenham and Stuart Charters. Guidelines for performing systematic literature reviews in software engineering, 2007.
- [82] Janelle Knox-Hayes. Negotiating climate legislation: Policy path dependence and coalition stabilization. *Regulation & Governance*, 6(4):545–567, 2012. ISSN 1748-5991.
- [83] S.D. Krasner. *International Regimes*. Cornell paperbacks. Cornell University Press, 1983. ISBN 9780801492501.
- [84] Sarit Kraus. Automated negotiation and decision making in multiagent environments. In Michael Luck, Vladimír Mařík, Olga Štěpánková, and Robert Trappl, editors, *Multi-Agent Systems and Applications: 9th ECCAI Advanced Course, ACAI 2001 and Agent Link's 3rd European Agent Systems Summer School, EASSS 2001 Prague, Czech Republic, July 2–13, 2001 Selected Tutorial Papers*, pages 150–172. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001. ISBN 978-3-540-47745-7. doi: 10.1007/3-540-47745-4_7.
- [85] Sarit Kraus and Onn Shehory. Coalition formation among autonomous agents: Strategies and complexity (preliminary report). In C. Castelfranchi and J.P. Müller, editors, *From Reaction to Cognition: 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW '93, Neuchatel, Switzerland, August 25-27, 1993. Selected Papers*, pages 57–72. Springer, 1995. ISBN 9783540601555.
- [86] David M. Kreps. *Game Theory and Economic Modelling*. Clarendon Press, 1990. ISBN 9780198283812.
- [87] Vijay Krishna and Roberto Serrano. Multilateral bargaining. *The Review of Economic Studies*, 63(1):61–80, 1996.
- [88] Guoming Lai, Cuihong Li, Katia Sycara, and Joseph Andrew Giampapa. Literature review on multi-attribute negotiations. Technical Report CMU-RI-TR-04-66, Carnegie Mellon University Robotics Institute, Pittsburgh, PA, December 2004.
- [89] Sam Laird. Multilateral approaches to market access negotiations. WTO Staff Working Paper TPRD-98-02, World Trade Organization, Economic Research and Statistics Division, Geneva, 1998.
- [90] David A. Lake. Rightful rules: Authority, order, and the foundations of global governance. *International Studies Quarterly*, 54(3):587–613, 2010. ISSN 1468-2478. doi: 10.1111/j.1468-2478.2010.00601.x.
- [91] Averill Law. *Simulation Modeling and Analysis*. McGraw-Hill, Inc., 5th edition, 2014. ISBN 9780077595968.
- [92] M.P. Leffler. *A Preponderance of Power: National Security, the Truman Administration, and the Cold War*. Stanford nuclear age series. Stanford University Press, 1992. ISBN 9780804722186.
- [93] Jack S. Levy and William R. Thompson. Hegemonic threats and great-power balancing in europe, 1495-1999. *Security Studies*, 14(1):1–33, 2005.
- [94] Roy J. Lewicki. *Essentials of Negotiation*. McGraw-Hill/Irwin, 2004. ISBN 9780072545821.
- [95] Roy J. Lewicki and Robert J. Robinson. Ethical and unethical bargaining tactics: An empirical study. *Journal of Business Ethics*, 17(6):665–682, 1998. ISSN 1573-0697. doi: 10.1023/A:1005719122519.
- [96] Kevin Leyton-Brown and Yoav Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, 1st edition, 2008. ISBN 1598295934, 9781598295931.
- [97] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014. ISSN 1467-8640. doi: 10.1111/j.1467-8640.2012.00463.x.

- [98] Miguel A. Lopez-Carmona, Ivan Marsa-Maestre, Mark Klein, and Takayuki Ito. Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces. *Autonomous Agents and Multi-Agent Systems*, 24(3):485–535, May 2012. ISSN 1387-2532. doi: 10.1007/s10458-010-9159-9.
- [99] Anthony J. Madonna. Winning coalition formation in the u.s. senate: The effects of legislative decision rules and agenda change. *American Journal of Political Science*, 55(2):276–288, 2011. doi: 10.1111/j.1540-5907.2010.00491.x.
- [100] Mihai Manea. Bargaining in stationary networks. *The American Economic Review*, 101(5):2042–2080, 2011. ISSN 00028282.
- [101] Cathie Jo Martin and Jane Mansbridge. Negotiating political agreements. In Jane Mansbridge and Cathie Jo Martin, editors, *Negotiating Agreement in Politics*, pages 662–669, Washington, DC, USA, 2013. American Political Science Association. ISBN 978-1-878147-47-9.
- [102] Leslie M. Marx and Greg Shaffer. Rent shifting and the order of negotiations. *International Journal of Industrial Organization*, 25(5):1109 – 1125, 2007. ISSN 0167-7187. doi: <https://doi.org/10.1016/j.ijindorg.2007.03.001>.
- [103] Leslie M. Marx and Greg Shaffer. Break-up fees and bargaining power in sequential contracting. *International Journal of Industrial Organization*, 28(5):451 – 463, 2010. ISSN 0167-7187. doi: <https://doi.org/10.1016/j.ijindorg.2009.11.001>.
- [104] Thomas J. McCabe. A complexity measure. In *Proceedings of the 2Nd International Conference on Software Engineering*, ICSE '76, pages 407–, Los Alamitos, CA, USA, 1976. IEEE Computer Society Press.
- [105] Thomas R. McFaul and Al Brunsting. *God and Randomness*. Wipf & Stock, 2017. ISBN 9781532638985.
- [106] Heather E. McKibben. *State Strategies in International Bargaining: Play by the Rules or Change Them?* Cambridge Studies in International Relations. Cambridge University Press, 2015. ISBN 9781316240908.
- [107] John J. Mearsheimer. *The Tragedy of Great Power Politics*. W.W. Norton & Company, 2001. ISBN 9780393978391.
- [108] Paul W. Meerts. *Diplomatic Negotiation: Essence and Evolution*. Clingendael Institute, The Hague, 2nd edition, 2015. ISBN 978-94-610-8781-2.
- [109] Dale T Miller. The norm of self-interest. *American Psychologist*, 54(12):1053–1060, 1999.
- [110] Helen Milner. International theories of cooperation among nations: Strengths and weaknesses. *World Politics*, 44(3):466–496, 1992.
- [111] Massimo Morelli. Party formation and policy outcomes under different electoral systems. *The Review of Economic Studies*, 71(3):829–853, 2004. ISSN 00346527, 1467937X.
- [112] Hans J. Morgenthau. *Politics among nations: the struggle for power and peace*. Borzoi book. Knopf, 4th edition, 1967.
- [113] Abhinay Muthoo. *Bargaining Theory with Applications*. Cambridge University Press, 1999. doi: 10.1017/CBO9780511607950.
- [114] John Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [115] Margaret A. Neale and Max H. Bazerman. Negotiating rationally: The power and impact of the negotiator's frame. *The Executive*, 6(3):42–51, 1992. ISSN 19389779.
- [116] Muaz Niazi and Amir Hussain. Agent-based computing from multi-agent systems to agent-based models: a visual survey. *Scientometrics*, 89(2):479, August 2011. ISSN 1588-2861. doi: 10.1007/s11192-011-0468-9.
- [117] Emerson M. S. Niou and Peter C. Ordeshook. Stability in anarchic international systems. *The American Political Science Review*, 84(4):1207–1234, 1990. ISSN 00030554, 15375943.
- [118] John S. Odell. *Negotiating the World Economy*. Cornell paperbacks. Cornell University Press, 2000. ISBN 9780801486463.
- [119] John S. Odell. Three islands of knowledge about negotiation in international organizations. *Journal of European Public Policy*, 17(5):619–632, 2010. doi: 10.1080/13501761003748534.

- [120] Council of the European Union. European border and coast guard: final approval. Press Release, September 2016. STATEMENT/510/16, <http://www.consilium.europa.eu/en/press/press-releases/2016/09/14/european-border-coast-guard/pdf>.
- [121] Council of the European Union. Eu-turkey statement. Press Release, March 2016. STATEMENT/144/16, <http://www.consilium.europa.eu/en/press/press-releases/2016/03/18/eu-turkey-statement/pdf>.
- [122] Council of the European Union. Voting system. <http://www.consilium.europa.eu/en/council-eu/voting-system/>, 2017. Accessed: 23th of November 2017.
- [123] Akira Okada. Non-cooperative Bargaining and the Incomplete Information Core. Discussion Papers 2009-16, Graduate School of Economics, Hitotsubashi University, December 2009.
- [124] Jeffrey K. Olick. *The Politics of Regret: On Collective Memory and Historical Responsibility*. Routledge, 2013. ISBN 9781135909802.
- [125] Kayo Onishi. Interstate negotiation mechanisms for cooperation in the mekong river basin. *Water International*, 32(4):524–537, 2007. doi: 10.1080/02508060.2007.9709686.
- [126] Martin .J. Osborne and Ariel Rubinstein. *Bargaining and Markets*. Economic theory, econometrics, and mathematical economics. Academic Press, Inc., 1990. ISBN 0125286325.
- [127] Guillermo Owen. Evaluation of a presidential election game. *The American Political Science Review*, 69(3): 947–953, 1975. ISSN 00030554, 15375943.
- [128] Aseem Pahuja, Rangaraja P Sundarraj, and Venkataraghavan Krishnaswamy. Cloudneg: An autonomous multi issue negotiation system, with preference elicitation component, for trading cloud services. In *At the Vanguard of Design Science: First Impressions and Early Findings from Ongoing Research-in-Progress Papers and Poster Presentations from the 10th International Conference, DESRIST 2015. Dublin, Ireland, 20-22 May*. DESRIST 2015, 2015.
- [129] Diana Panke. Microstates in negotiations beyond the nation-state: Malta, cyprus and luxembourg as active and successful policy shapers? *International Negotiation*, 16(2):297–317, 2011. doi: <https://doi.org/10.1163/138234011X573057>.
- [130] Robert A. Pape. Soft balancing against the united states. *International Security*, 30(1):7–45, 2005. ISSN 01622889, 15314804.
- [131] Noam Peled, Ya'akov (Kobi) Gal, and Sarit Kraus. A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems*, 29(1):73–97, Jan 2015. ISSN 1573-7454. doi: 10.1007/s10458-014-9253-5.
- [132] Bezalel Peleg. Chapter 8 Game-theoretic analysis of voting in committees. In *Handbook of Social Choice and Welfare*, volume 1, pages 395 – 423. Elsevier, 2002. doi: [https://doi.org/10.1016/S1574-0110\(02\)80012-1](https://doi.org/10.1016/S1574-0110(02)80012-1).
- [133] Torsten Persson, Gerard Roland, and Guido Tabellini. Electoral rules and government spending in parliamentary democracies. *Quarterly Journal of Political Science*, 2(2):155–188, 2007. ISSN 1554-0626. doi: 10.1561/100.00006019.
- [134] Liga M. Pinto and Glenn W. Harrison. Multilateral negotiations over climate change policy. *Journal of Policy Modeling*, 25(9):911 – 930, 2003. ISSN 0161-8938. doi: <https://doi.org/10.1016/j.jpolmod.2003.08.002>.
- [135] Clara Ponsati and Joel Watson. Multiple-issue bargaining and axiomatic solutions. *International Journal of Game Theory*, 26(4):501–524, Dec 1997. ISSN 1432-1270. doi: 10.1007/BF01813888.
- [136] Robert Powell. Bargaining theory and international conflict. *Annual Review of Political Science*, 5(1):1–30, 2002. doi: 10.1146/annurev.polisci.5.092601.141138.
- [137] Dean G. Pruitt. *Negotiation Behavior*. Library and Information Science. Academic Press, 1981. ISBN 9780125662505.
- [138] Robert D. Putnam. Diplomacy and domestic politics: the logic of two-level games. *International Organization*, 42(3):427–460, 1988. doi: 10.1017/S0020818300027697.
- [139] Iyad Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, Peter Mcburney, Simon Parsons, and Liz Sonenberg. Argumentation-based negotiation. *Knowl. Eng. Rev.*, 18(4):343–375, December 2003. ISSN 0269-8889. doi: DOI:10.1017/S0269888904000098.

- [140] Howard Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982. ISBN 9780674048133.
- [141] Gordon Rausser and Leo Simon. *A non-cooperative model of collective decision making: a multilateral bargaining approach*. Department of Agricultural & Resource Economics, University of California Berkeley, 1992.
- [142] Debraj Ray. *A Game-Theoretic Perspective on Coalition Formation*. Oxford University Press, 2007. ISBN 9780191607240.
- [143] Debraj Ray and Rajiv Vohra. Equilibrium binding agreements. *Journal of Economic Theory*, 73(1):30 – 78, 1997. ISSN 0022-0531. doi: <https://doi.org/10.1006/jeth.1996.2236>.
- [144] William H. Riker. *The Theory of Political Coalitions*. Yale University Press, 1962.
- [145] Valentin Robu, D. J. A. (Koye) Somefun, and J. A. (Han) La Poutré. Modeling complex multi-issue negotiations using utility graphs. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '05*, pages 280–287, New York, NY, USA, 2005. ACM. ISBN 1-59593-093-0. doi: 10.1145/1082473.1082516.
- [146] James N. Rosenau. Change, complexity, and governance in globalizing space. In Jon Pierre, editor, *Debating Governance: Authority, Steering, and Democracy*, pages 167–200. Oxford: Oxford Univeristy Press, 2000. ISBN 9780191583315.
- [147] Robert S. Ross. Balance of power politics and the rise of china: Accommodation and balancing in east asia. *Security Studies*, 15(3):355–395, 2006. doi: 10.1080/09636410601028206.
- [148] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982. ISSN 00129682, 14680262.
- [149] Nicole J. Saam and David Sumpter. Coalition formation as emergent phenomenon in intergovernmental negotiations. Unpublished manuscript, 2004.
- [150] Thomas W. Sandholm and Victor R. Lesser. Coalition formation among bounded rational agents. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'95*, pages 662–669, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc. ISBN 1-55860-363-8, 978-1-558-60363-9.
- [151] Robert G. Sargent. Some approaches and paradigms for verifying and validating simulation models. In *Proceeding of the 2001 Winter Simulation Conference (Cat. No.01CH37304)*, volume 1, pages 106–114, 2001. doi: 10.1109/WSC.2001.977251.
- [152] Jurgen Scheffran. The dynamic interaction between economy and ecology: Cooperation, stability and sustainability for a dynamic-game model of resource conflicts. *Mathematics and Computers in Simulation*, 53(4):371 – 380, 2000. ISSN 0378-4754. doi: [https://doi.org/10.1016/S0378-4754\(00\)00229-9](https://doi.org/10.1016/S0378-4754(00)00229-9).
- [153] James K. Sebenius and David Lax. Thinking coalitionally: Party arithmetic, process opportunism, and strategic sequencing. In H. Peyton Young, editor, *Negotiation Analysis*, pages 153–193. University of Michigan Press, Ann Arbor, MI, 1992.
- [154] Alessandra Sgobbi and Carlo Carraro. A stochastic multiple players multi-issues bargaining model for the piave river basin. *Strategic Behavior and the Environment*, 1(2):119–150, 2011. ISSN 1944-012X. doi: 10.1561/102.00000006.
- [155] Lloyd S. Shapley. Simple games: An outline of the descriptive theory. *Behavioral Science*, 7(1):59–66, 1962. ISSN 1099-1743. doi: 10.1002/bs.3830070104.
- [156] Lloyd S. Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *Cambridge University Press*, 48(3):787–792, 1954. doi: doi:10.2307/1951053.
- [157] Michael Sheehan. *The Balance of Power: History and Theory*. Routledge, 1996. ISBN 0-415-11930-8.
- [158] Onn Shehory and Sarit Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1):165 – 200, 1998. ISSN 0004-3702. doi: [https://doi.org/10.1016/S0004-3702\(98\)00045-9](https://doi.org/10.1016/S0004-3702(98)00045-9).
- [159] Onn Shehory and Sarit Kraus. Feasible formation of coalitions among autonomous agents in nonsuperadditive environments. *Computational Intelligence*, 15(3):218–251, 1999. ISSN 1467-8640. doi: 10.1111/0824-7935.00092.

- [160] Leo Simon, Rachael Goodhue, Gordon Rausser, Sophie Thoyer, Sylvie Morardet, and Patrick Rio. *Structure and Power in Multilateral Negotiations: An Application to French Water Policy*. University of California, Davis, Giannini Foundation, 2007.
- [161] Leo K. Simon, Rachael E. Goodhue, Gordon C. Rausser, Sophie Thoyer, Sylvie Morardet, and Patrick Rio. *Structure and Power in Multilateral Negotiations: An Application to French Water Policy*. University of California, Davis, Giannini Foundation, 2007.
- [162] Israel Sofer, David Sarne, and Avinatan Hassidim. Negotiation in exploration-based environment. *Autonomous Agents and Multi-Agent Systems*, 30(4):724–764, Jul 2016. ISSN 1573-7454. doi: 10.1007/s10458-015-9303-7.
- [163] Philip D. Straffin. *Game Theory and Strategy*. Anneli Lax New Mathematical Library. Mathematical Association of America, 1993. ISBN 9780883856376.
- [164] Philip D. Straffin. Power and stability in politics. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2, chapter 32, pages 1127–1151. Elsevier, 1st edition, 1994.
- [165] Carolin Strobl, Anne-Laure Boulesteix, Thomas Kneib, Thomas Augustin, and Achim Zeileis. Conditional variable importance for random forests. *BMC Bioinformatics*, 9(1):307, Jul 2008. ISSN 1471-2105. doi: 10.1186/1471-2105-9-307.
- [166] Kevin Sweeney and Paul Fritz. Jumping on the bandwagon: An interest-based explanation for great power alliances. *The Journal of Politics*, 66(2):428–449, 2004.
- [167] Alan O. Sykes and Eric Posner. Voting rules in international organizations. Working paper, no. 458, University of Chicago Law School & Legal Theory, 2014.
- [168] Bedir Tekinerdogan, Nour Ali, John Grundy, Ivan Mistrik, and Richard Soley. Quality concerns in large-scale and complex software-intensive systems. In Ivan Mistrik, Richard Soley, Nour Ali, John Grundy, and Bedir Tekinerdogan, editors, *Software Quality Assurance*, pages 1 – 17. Morgan Kaufmann, Boston, 2016. ISBN 978-0-12-802301-3. doi: <https://doi.org/10.1016/B978-0-12-802301-3.00001-6>.
- [169] Sophie Thoyer, Sylvie Morardet, Patrick Rio, Leo Simon, Rachael Goodhue, and Gordon Rausser. A bargaining model to simulate negotiations between water users. *Journal of Artificial Societies and Social Simulation*, 4 (2), March 2001. ISSN 1460-7425.
- [170] Seth Tisue and Uri Wilensky. Netlogo: A simple environment for modeling complexity. In *International conference on complex systems*, volume 21, pages 16–21. Boston, MA, 2004.
- [171] James Tremewan and Christoph Vanberg. The dynamics of coalition formation – a multilateral bargaining experiment with free timing of moves. *Journal of Economic Behavior & Organization*, 130:33 – 46, 2016. ISSN 0167-2681. doi: <https://doi.org/10.1016/j.jebo.2016.07.011>.
- [172] George Tsebelis and Geoffrey Garrett. Agenda setting power, power indices, and decision making in the european union. *International Review of Law & Economics*, 16(3):345–361, 1996.
- [173] Ad M.A. Van Deemen. *Coalition Formation and Social Choice*. Theory and Decision Library. Kluwer Academic Publishers, 1997. ISBN 9781441947819. doi: 10.1007/978-1-4757-2578-0.
- [174] Stephen M. Walt. Alliance formation and the balance of world power. *International Security*, 9(4):3–43, 1985. ISSN 01622889, 15314804.
- [175] Stephen M. Walt. Explaining alliance formation. In *The Origins of Alliance*, pages 17–49. Cornell University Press, 1987. ISBN 9780801420542.
- [176] Richard E. Walton and Robert B. McKersie. *A Behavioral Theory of Labor Negotiations: An Analysis of a Social Interaction System*. Cornell University Press, 1965. ISBN 9780875461793.
- [177] Kenneth N. Waltz. *Theory of international politics*. Reading, Mass: Addison-Wesley Pub. Co., 1979.
- [178] Leo Wangler, Juan-Carlos Altamirano-Cabrera, and Hans-Peter Weikard. The political economy of international environmental agreements: a survey. *International Environmental Agreements: Politics, Law and Economics*, 13(3):387–403, Sep 2013. ISSN 1573-1553. doi: 10.1007/s10784-012-9196-8.
- [179] Uri Wilensky. *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL., 1999. URL <http://cc1.northwestern.edu/netlogo/>.

- [180] Herman Wold. A synthesis of pure demand analysis. *Scandinavian Actuarial Journal*, 1944(1-2):69–120, 1944. doi: 10.1080/03461238.1944.10404937.
- [181] World Trade Organization. Annual report 2017. Technical report, World Trade Organization, 2017. https://www.wto.org/english/res_e/booksp_e/anrep_e/anrep17_e.pdf.
- [182] James D. Wright. *International Encyclopedia of the Social Sciences*. Macmillan Reference USA, second edition, 2008. ISBN 9780028661179.
- [183] Alasdair R. Young. Not your parents' trade politics: the transatlantic trade and investment partnership negotiations. *Review of International Political Economy*, 23(3):345–378, 2016. doi: 10.1080/09692290.2016.1150316.
- [184] Nikolaos Zahariadis. Bargaining power and negotiation strategy: examining the greek bailouts, 2010–2015. *Journal of European Public Policy*, 24(5):675–694, 2017. doi: 10.1080/13501763.2016.1154977.
- [185] Dolf Zillmann and Hans B. Brosius. *Exemplification in Communication: the influence of Case Reports on the Perception of Issues*. Routledge, 2012. ISBN 9781135681609.
- [186] Taleh Ziyadov. Nagorno-karabakh negotiations: Though the prism of a multi-issue bargaining model. *International Negotiation*, 15:107–131, 2010.
- [187] Inon Zuckerman, Erel Segal-Halevi, Avi Rosenfeld, and Sarit Kraus. First steps in chat-based negotiating agents. In Katsuhide Fujita, Takayuki Ito, Minjie Zhang, and Valentin Robu, editors, *Next Frontier in Agent-Based Complex Automated Negotiation*, pages 89–109. Springer Japan, Tokyo, 2015. ISBN 978-4-431-55525-4. doi: 10.1007/978-4-431-55525-4_6.



Search Terms in Systematic Literature Review

Both Chapter 2 and Chapter 3 use a systematic literature review to answer the first two supporting questions of this research. Several steps are taken to ensure a such a review (stemming from [81]). The first step entails defining search terms used for querying different databases. As databases varied in size, specific and more general search terms were used (see Appendix A.1). Thereafter, different scientific databases were selected to search through (see Table 2.1). These databases were mainly searched for identifying highly-cited papers published in conferences and journals, but also books were found. By examining papers more closely, a handful of studies were selected by screening bibliography lists, categorized in Table 2.1 as "miscellaneous". In the final step, rules were developed to select *relevant* papers, contributing to answering the supporting question central to this chapter. Studies were first selected based upon *meta-data*. Although meta-data is defined differently across academic databases, generally it includes (sub)titles, abstracts, keywords and journal titles. Next, papers were selected more thoroughly after scanning titles. Subsequently, we have read the abstracts of the remaining papers to judge their usefulness. The remaining papers have been read entirely, after which some of them were still removed from further analysis. Finally, as databases partly overlap, duplicate articles were also taken out.

As part of a systematic literature review, we have carefully registered every search term in order to increase the reproducibility of our reviews. Depending on the database – i.e. IEEE Xplore, Web of Science, Scopus and ABI/INFORM Complete – dozens or even thousands of hits were found for each search term.

A.1. Characteristics of International Political Negotiations

Table A.1 shows the search terms used for the first systematic literature review, discussing the architecture of international political negotiations. For this review, Table 2.1 shows the aggregated results of the hits *after* selection based upon meta data. Note that for both Table A.1 and Table A.2 holds that the table below is not exhaustive as sometimes slight alterations of the same search terms are used.

Table A.1: Combinations of search terms used to perform literature review in Chapter 2.

Search term	Motivation
"politic*" AND "negotiation*"	Overview of political negotiation literature
"characteristic*" AND "politic*" AND "negotiation*"	Finding important elements of any political negotiations
"characteristic*" AND "politic*" AND "process*"	Finding important elements of political negotiation processes and phases

"characteristic*" AND "international*" AND "negotiation*"	Discover important elements of international negotiations (not specifically political)
"characteristic*" AND "diplomatic*" AND "negotiation*"	Discover important elements of diplomatic negotiations
"characteristic*" AND "inter-state*" AND "negotiation*"	Discover important elements of diplomatic negotiations between states
"characteristic*" AND "international*" AND "relation*"	Discover important elements of diplomatic negotiations between states
"international" AND "politic*" AND "process*"	Find papers about (the characteristics of) international political processes
"diplomatic" AND "characteristic*" AND "process*"	Find characteristics of diplomatic processes between states
"diplomatic" AND "power"	Find sources of diplomatic power (asymmetry)
"inter-state" AND "politic*" AND "process*"	Find something about political processes between states
"negotiation*" AND "international" AND "relation*"	Find something about negotiation in international relations
"negotiation*" AND "international" AND "affairs*"	Find something about negotiation in international affairs
"international" AND "politic*" AND "power"	Discover characteristics of power in world politics
"inter-state" AND "negotiation*" AND "power"	Discover characteristics of power in inter-state negotiation
"voting" AND "politic*" AND "negotiation*"	Find voting rules in political negotiations in general
"voting rule" AND "intergovernmental" AND "organization*"	Find voting rules in intergovernmental organizations
"voting rule" AND "international" AND "negotiation*"	Find voting rules in international negotiations
"negotiation*" AND "strateg*"	Overview of papers related to negotiation strategies in general
"negotiation*" AND "strateg*" AND "international"	Finding papers related to international negotiation strategies
"negotiation*" AND "strateg*" AND "diplomatic"	Finding papers related to diplomatic strategies in inter-state negotiations
"actors" AND "international" AND "negotiation*"	Find actors involved in any international negotiations
"actors" AND "diplomatic*" AND "negotiation*"	Find actors involved in international negotiations between one or more states


A.2. Coalition Formation in International Politics

Table A.2 presents the search terms used for the second systemic literature review, discussing different theories of coalition formation. Again, Table 3.1 shows the aggregated results of hits *after* selection based upon meta data.

Table A.2: Combinations of search terms used to perform literature review in Chapter 3.

Search term	Motivation
"coalition*" AND "formation"	Get a general impression of coalition formation literature (including highly cited articles)

"strateg*" AND "coalition*"	Discover strategies or strategic choices in coalition formation
"alliance*" AND "international" AND "relation*"	Find coalition formation in international relations research as it is often described as in terms of "alliances"
"coalition*" AND "formation" AND "game" AND "theory"	Find coalition formation literature described in game theory research
"explain*" AND "alliance*" AND "formation"	Find theories on coalition formation in international relations research as it is often described as in terms of "alliances"
"bandwagoning"	Find hypotheses and use cases of bandwagoning theory
"balance" AND "power" AND "theory"	Discover balance-of-power theory
"coalition*" AND "rationality"	Find game theoretic assumptions regarding rational behavior of coalitions (in contrary to individual agents)
"coalition*" AND "utility"	Find utility theories regarding coalition formation
"coalition*" AND "autonomous" AND "agent*"	Find agent-based theory to look into their assumptions regarding coalition formation behavior
"coalition*" AND "formation" AND "negotiation*"	Find coalition formation examples in bargaining theory
"Riker*" AND "size" AND "principle*"	Find more about Riker's size principle
"Riker*" AND "size" AND "principle*" AND "criticism*"	Find criticisms of Riker's size principle and empirical (counter) evidence
"political*" AND "coalition*" OR "formation"	Find political theory discussing coalitions and coalition formation
"dynamic*" AND "coalition*" AND "formation*"	Find research discussing the dynamics of coalition formation
"cooperative" AND "game" AND "theory*"	Get a general impression of (highly-cited) cooperative game theory literature
"politic*" AND "regret" OR "utility" AND "coalition*"	Find out if regret theory is used in political literature discussing coalitions
"spatial" AND "model*" OR "coalition*"	Find literature discussing spatial models of coalition formation
"coalition*" AND "formation" AND "power"	Find literature describing the influence of power on coalition formation
"core" AND "game" AND "theory*"	Find more (e.g. stability, uniqueness of solution, outcomes) about the core in game theoretic literature
"leav*" AND "political" AND "coalition*"	Find if leaving a coalition is described in any literature
"irrational" AND "coalition*" AND "formation"	Find if irrational behavior in coalition formation is described in any literature
"coalition*" AND "formation" AND "probability"	Find if (spatial-oriented) models are developed that determine the probability of certain coalitional arrangements
"spatial" AND "model*" OR "politic*"	Discover spatial models specifically focused on political coalitions



Issue Classification


Table B.1 provides an overview of the identified issue types examined in literature. Note that this table is by no means exhaustive. As explained in Section 2.2.3, *binary* issues are a special type of indivisible issues in which the "winner-takes-it-all". Moreover, *indivisible* issues have a fixed number of alternatives from which agents can choose. In addition, *unrestricted* issues can in theory take any value between a certain range, whereas *limited* issues can also take certain values between a range, but are dependent on the total number of available resources – i.e. divide the "pie" (see for a famous theoretical example [113, p. 10-16]).

Table B.1: Classification of different issues discussed in literature (and press releases).

Issue	Description	Source	Type
Location of IGO headquarters	Press release describing the relocation of the European Medicine Agency from London to Amsterdam after the Brexit. This was the result of a lengthy negotiation and voting session among the 26 remaining EU member states.	[44]	Binary
War (intervention)	The article describes in detail the negotiations between the United Nations, North Atlantic Treaty Organization and the European Union to either intervene in Libya's civil war or stay out of it [4, p. 890-891]. Note that the article also discusses different options for intervention, such as a "no-fly-zone" or "ground troops" [4, p. 900-908]. The latter has another classification in this table.	[4]	Binary
War (cease fire)	The article describes the negotiation attempts between Azerbaijan and Armenia by mediation of the Organization for Security and Co-operation in Europe (OSCE) and its Minsk Group (among others France, Russia, and the United States) for a cease of fire in 1992. After the cease of fire in 1994, more negotiations on different types of issues were conducted (see later on).	[186, p. 107-109]	Binary
Independence of a region	After the cease of fire, the Nagorno-Karabakh "(...) demanded either unification of NK with Armenia or full independence from Azerbaijan".	[186, p. 110]	Binary
Intellectual property rights	The article describes coalition formation between developed and developing countries during international negotiations on intellectual property rights. It describes different options for among others the violation of intellectual property, but also connections with education, culture and human rights.	[74]	Indivisible resources

Quality standards	The article describes Transatlantic Trade and Investment Partnership negotiations, which consists of various different issues. One of which includes negotiating the different (quality) standards among the countries involved, so that markets get accessible for each other's companies.	[183, p. 365]	Indivisible resources
Legislation	In this article, the adoption of new legislation is described in a model of accumulating amendments. A finite number of alternative amendments are defined, after which they are being negotiated. When one of these amendments passes, the process starts again until new legislation is mutual satisfactory.	[15, p. 1183-1186]	Indivisible resources
Means of infiltration	As stated earlier, this article also discusses different options for intervention, such as a "no-fly-zone" or "ground troops" [4, p. 900-908].	[4, p. 900-908]	Indivisible resources
Greenhouse gas emissions	This article discusses among others the property rights of the emission for trans-boundary pollutants and resources. It states that they are currently ill defined, which, as a consequence, results in unlimited emissions.	[178]	Unrestricted resources
CO ₂ reduction of greenhouse gas emissions	This article proposes a model to simulate the negotiations over carbon dioxide reductions. The model assumes that the emissions are in essence unlimited, with a striking example "(...) is the extent to which the European Union allowed differences in the implementation of their aggregate commitment within member states. Denmark and Germany agreed to targets that were less than 80% of their 1990 emissions, while Greece and Portugal were allowed emissions that were over 125% of their 1990 emissions".	[134, p. 911-913]	Unrestricted resources
Greenhouse gas emissions	This book describes numerous issues and options discussed during various climate negotiations (i.e. Kyoto and Copenhagen). Recurrent issue is the limitation of greenhouse gas emissions, started during the negotiators in Berlin	[40, p. 35-38]	Unrestricted resources
Import tariffs	This article describes multilateral market access negotiations by means of tariffs measures. In essence, every country can raise import tariffs to infinite high levels. As discussed, during and after the Uruguay negotiations, substantial reductions in import tariffs were succeeded.	[89]	Unrestricted resources
Budget allocation	This article describes the allocation of the EU budget between 1977-2003 over the various member states and different topics (i.e agriculture funds) as a result of various negotiations.	[5]	Limited (shared) resources
Rare materials	This article describes water as a rare resource in a specific part of China: the Mekong River Basin. As China shares this basin with amongst others Laos and Thailand, they gradually getting involved in negotiations with the downstream states, compromising little by little. There are, of course, many other examples of rare resources and materials.	[125]	Limited (shared) resources

Allocation of refugees	This press release announces the distribution of 120,000 refugees among all EU member states after the refugee crisis in 2015 after thorough negotiations. Because the total number of (to be allocated) refugees is restricted, it forms an (un)natural limit to the negotiations.	[32]	Limited (shared) resources
EU fishing quotas	This article describes a model in which well-defined fishery property right are divided among different players by means of negotiations. Because the total amount of fishery rights is well-defined, they form an (un)natural limit to the negotiations.	[152]	Limited (shared) resources
Territory	The section in this book describes some examples in which the issue of bargaining were borders or entire regions. These borders and regions provided "natural" limitations to the outcome of the bargaining process.	[108, p. 72-74]	Limited (shared) resources
Territory	As stated earlier, after the cease of fire in 1994, negotiations began regarding territory. Armenian forces occupied around 16% of Azerbaijan and, in order to remain peace, negotiations needed to be conducted to hand over some of the occupied regions back to Azerbaijan.	[186]	Limited (shared) resources



Overview of Assumptions

This appendix presents an overview of assumptions derived in Chapter 2, 3 and 4 of this thesis. Also the implications of assumptions are described, either by literature or supported quantitatively.

Table C.1: Overview of assumptions made in the first three literature review chapters.

#	Name	Description	Implication on validity
1	Self-interested actors	Actors involved in international political negotiations are self-interested and want to maximize their own (absolute) gains. If they engage in behavior that can be labeled as 'cooperation', then this will be in their personal best interest.	As a result, the unit of analysis changes. In non-cooperative coalitional bargaining agreements are made based on <i>individual</i> moves [142, p. 14]. Consequently, in contrary to cooperative game theory, it becomes difficult to avoid inefficient outcomes [142, p. 19]. Moreover, the notion of fairness is lost as negotiators aim for maximum personal gains. Note that in our model, we partly 'fight' this consequence by assuming individual negotiators may still partly aim for fair outcomes (see Section 6.1.1).
2	Omniscient negotiators	Negotiators (e.g. politicians or diplomats) on an international level represent the state in such a way that their preferences are already the results of ongoing domestic negotiations.	Impact of this assumption is examined by randomly changing preferences of one or more agents <i>during</i> the negotiations. Five settings are analyzed: the ' <i>base</i> ' setting as assumed in our model and four alternatives in which agents change their preferences with $p = \{0.2, 0.4, 0.6, 0.8\}$. From these experiments, we conclude that results of this thesis solely apply in negotiation settings in which states have more or less stable preferences for a certain period of time. So, in situations of (sudden) domestic pressures – e.g. before tense national elections – our results become meaningless. See point 2 for the supporting analysis.
3	Issue types	All to be negotiated issues can be represented by any of the identified issue types – i.e. <i>binary</i> , <i>indivisible</i> , <i>unrestricted</i> and <i>limited</i> – or be adjusted in such a way that they can.	In Chapter B we extensively surveyed many different types of issues and classified them accordingly. As this survey covers a wide range of real (international) political negotiations, the impact of this assumption is negligible. Note that in the final tool, binary issues are not incorporated because of time constraints. Therefore, negotiation domains discussing the location of IGO headquarters (see [44]), for example, cannot be adequately represented.
4	Strategically choose pre-coalitions	Regimes have already defined the bargaining table, voting and negotiation rules, issues under consideration and bargaining agenda. Pre-coalition formation is the only strategic action states can take <i>before</i> the actual negotiation takes place.	The impact of this assumption is two-fold. First, Fatima et al. (2004, [48]) have shown the agenda is not at all given. In fact, in many negotiations it is part of the pre-negotiation stage, such that negotiators can exploit it in their advantage. As a result, they might obtain higher personal gains or lower time to agreement [46]. In addition, from de Bruijn and ten Heuvelhof (2012) we know actual bargaining process is not fixed as well: actors may come and go [36, p. 27-28]; new issues can emerge [36, p. 49-54]; and (social) rules might change because of repetitive behavior [36, p. 114]. Again, agents use this as strategies to find new or better solutions [36, p. 49-54]. So, as a consequence, our model lacks some important strategies states adopt, which results in different negotiation outcomes or lower likelihood of agreement.

5	Defining power	Negotiators are all equally experienced and, therefore, have the same <i>behavioral power</i> . Moreover, <i>structural power</i> is derived from the voting rule in place (e.g. seats in the European Parliament). Finally, <i>issue-specific power</i> is defined implicitly from actors preferences and opponents' offers.	As both structural and issue-specific power are incorporated, potential impact of this assumption is (mainly) caused by neglecting behavioral power. Among others, [36] have shown that the experience of negotiators do impact outcomes. For example, in reality, negotiation takes place in networks [36, p. 35]. An experienced negotiator is likely to have a larger network. Consequently, he might have better information [36, p. 44-47] or more room to maneuver [36, p. 47-79] ^[1] . In addition, if we assume experience affects the likelihood offers are considered by others – e.g. similar to [141] – we can test the impact of this assumption quantitative. From the analysis in point 5 we conclude that by disregarding the experience of negotiators, we also neglect the possibility experienced negotiators bend negotiations to their will. Therefore, our results cannot be generalized to settings with great experience imbalances.
6	Negotiation strategy	Negotiators deploy both cooperative and conflictive strategies aimed at maximizing their goals, but adjusted for the amount of risk they want to take for ending in no agreement.	Traditionally, the two strategy classifications result in either 'sink or swim together' and 'when one swims, one sinks' [11, p. 150]. By creating dynamic mixtures between the two – and, therewith, following [11] – we have created more capricious settings. This matches reality better (e.g. see [79] and [35]), but as a consequence, it increases the outcome space. The latter is shown by example in point 6.
7	Stability of regimes	Regimes are fixed and stable during a single negotiation process, but may change thereafter.	This assumption is fairly common in most bargaining models as inter-governmental organizations tend to change very slowly (for a concise overview, see [108, p. 53-57]). To give an example, voting procedures in the Council of the European Union have changed only 5 times after the Treaty of Rome establishing the European Economic Community in 1958 ^[2] . So, assuming stable regimes seems valid. Though, as discussed in [108, p. 54], intergovernmental organization can also emerge on regional levels, which do have the tendency to change more rapidly. So, conclusions in this thesis generally apply for most regimes. Only in case of intergovernmental organization established regionally, one should validate whether or not it is likely to change before generalizing conclusions.
8	Negotiation memory	Past events <i>within</i> a single negotiation process influences a negotiator's future negotiation strategy. Past events in <i>previous</i> negotiation processes are already represented by a negotiator's preferences and strategies.	As a consequence, our model is not able to model sequential bargaining processes. States are often involved in many international negotiations at the same time. Thus, potentially they are quite likely to be part of more than one pre-coalition. As these negotiations partly overlap and evolve sequentially, one can imagine outcomes of one negotiation to influence the ex ante coalition choice of another [58, 102, 103]. So, if it is likely negotiations in two or more distinct international arenas will interact – for example, EU's refugee crisis and Brexit – our results may not be valid.
9	True power position	The a priori true power position of a state or coalition of states influences both its strategic demands and likelihood their respective demands are considered in the bargaining process.	For our experimental design, we have assumed states are aware of their decisiveness and, therefore, make demands accordingly (see Section 6.2.1). We have used to the Shapely-Shubik index to calculate the decisiveness and, thus, the respective demands. To test the impact of this assumption, we have run three simulation experiments in which we compare demands made by the 'base' (SS) situations, their structural 'power' position and the 'Banzhaf' index. From this analysis, we conclude our model is relatively robust for different ways of calculating a state's strategic demands. Therefore, this assumption has relatively little impact in the model, keeping in mind that all methods are based on structural power. This, implying that if negotiators are likely to make demands based on other (historic) reasons, conclusions of this thesis cannot be generalized. See point 10.
10	Power excess	The ability of a party to make proposals within a coalition is defined by his/her <i>power excess</i> , in which the power measure p_i is a priori defined according to some power index.	One can imagine that power excess of a state within a coalition is not always a reasonable assumption. For example, if some states have worked together for some time and a stronger country recently joins, then their history of working together might result in less (actual) power for the strong state than assumed. Therefore, we have tested the power excess assumption against two others; one in which states will always be equally powered (no matter their structural power) and one where the most power is hold by those who were part of the coalition from the very beginning. From these experiments, we can conclude, overall, changing the power excess assumption does not affect the results of this thesis much. So, conclusions may be generalized even if slightly different within coalition power positions hold. See also point 9.

^[1]For example, if a state does not like an offer, he simply goes to someone else.

^[2]For more information, please see https://europa.eu/european-union/about-eu/history_en.

11	Transferable utility	Utility is transferable between both <i>issues</i> and negotiating <i>parties</i> .	In many situations, this assumption is rather strong, though justifiable [163, p. 130]. It relies upon a common currency that is equally valued by all negotiators. In our day-to-day economic transactions, such medium of exchange is money. Of course, money is not always valued equally; think for example in differences between a millionaire and a homeless. Extending this analogy to diplomatic negotiations, differences may occur between developing and developed countries. Moreover, if states have different currencies, the impact of exchange rates – especially in currency crises ^[3] – may also affect the validity of this assumption. Therefore, the conclusions of this thesis are valid if negotiators all adopt a common currency (such as the Euro in the EMU) or if a global currency (e.g. the Dollar) is valued equally among all parties involved.
12	Joining and leaving rationality	Self-interested parties will only join or stay in a coalition if his gains are at least as much within the coalition as it could benefit by himself [159, p. 223], or in any other case they will leave.	As a consequence, agents loose the notion of <i>farsightedness</i> [142, p. 15-16]. This, because if agents or coalitions face an offer x^t at time t they only consider bids $\{x^0, \dots, x^t\}$ when joining or leaving. Thus, not looking beyond the immediate consequence of their actions by approximating $\{x^{t+1}, \dots, x^{t+n}\}$. In reality, it is easy to imagine negotiators sometimes may adopt more farseeing strategies; sticking to the coalition now by accepting a low offer, hoping future offers will increase. To partially mitigate the consequence of this assumption, additionally, we assume agents will not always join or leave immediately after facing a certain offer. Instead, they leave or join by probability p . Therefore, we expect the effect of this assumption on our conclusions to be minor.
13	Bounded rationality	Agents are bounded rational, which implies they make decisions given the (incomplete) information, cognitive limitations and available decision time.	An immediate consequence of bounded rationality is that agents might not make the optimal choices as they are incapable to do so. So, our results should never be interpreted as the 'optimal' ex ante coalition strategy resulting in optimal negotiation outcomes. Instead, our results indicate good pre-coalition strategies, most likely resulting better outcomes than staying alone before the start of actual negotiations.
14	Stopping criteria	Negotiations end if either the public deadline is passed or members of the winning coalition agree its size is adequately "large enough".	As a result of this assumption, winning coalitions may become larger than minimal winning (see [144]). Unfortunately, there is empiric evidence supporting our assumption – winning coalitions can be larger than minimal winning (e.g. see [99]) – and evidence showing Riker's size principle (e.g. see [37]). A quick look in some databases shows most of these empirical studies examines government formation on national or regional level. Therefore, we are not sure of the impact of this assumption on the validity of our conclusions. This should be further examined.
15	Public deadlines	Deadlines are hard, political targets set prior to the negotiation and, thus, public information.	Of course, we realize public deadlines are perhaps the least hard of all deadlines around. Even the 'hard' Brexit deadline of 29 th of March 2019 has been (sort of) extended by a transition period ^[4] . This also shows deadlines are often part of the (ex ante) negotiation process itself, as well as of a party's strategy to postpone large decisions. As this (currently) is not in scope of our model, conclusions of this thesis may be invalid for settings in which international negotiators are not effectively bound by hard deadlines and, therefore, have the ability to delay time after time.
16	Selecting initial proposer	Selecting an agent can either be done <i>randomly</i> , according to their <i>power position</i> or a combination of both in which the selection function is a <i>Markov chain</i> . These functions are used to select i) initial proposers and ii) opponents. The selection function for i) and ii) might differ for single negotiations.	This assumption basically entails making a choice which selection procedure is most valid given real international political negotiations. As discussed in Section 5.3.2 we think a Markov Process represents the selection procedure best as it acknowledges power positions without ignoring . To illustrate the impact of this choice, we have ran three simulation experiments in which we tested the ' <i>base</i> ' (Markov) procedure against ' <i>random</i> ' and ' <i>power</i> ' approaches (see point 16). From this analysis, we conclude as offers of more powerful negotiators are considered first in an international setting, the impact of this assumption on the conclusions of this thesis is minor.

^[3]Examples include the Asian (1997) and Russian (1998) financial crises or the default of Argentina in 2006.

^[4]See Article 121 of the Draft Withdrawal Agreement [33, p. 68].

17	Independent issues	Issues are independent and the combined utility function U can be defined by a (weighted) linear combination of the utilities u_j of all independent issues j .	A direct consequence of using linear additive utility functions is that they not represent interaction effects or dependencies between issues [88]. This seems a strong assumption, as many studies have shown discrepancies between model predictions and negotiation outcomes when using linear additive utility functions (e.g. [71, 98, 145]). Note that in our model, dependency does occur within a single package deal. So, we do capture some important dependencies. Though, only between related issues, whereas in international political negotiations, it is often assumed everything is interconnected [108, p.]. So, our conclusions may become invalid for multi-subject negotiations in which the subjects heavily differ. The Brexit in its entirety may be such an example.
18	Discounted utility with uncertainty	The utility function for agent i depends on the outcome x at time t is described by $U_i(x, t) = \mu_i^t \cdot u_i(x) + \epsilon$, in which u_i is the agents private utility function and $d_i \in (0, 1]$.	In their paper, Benati and Marzetti (2013, [19]) have shown for a European setting that the addition of a probabilistic component to utility functions increases model validity. In addition, in [88, p. 1-3] the authors state most actors face difficulties evaluating large numbers of issues. Consequently, for humans it is not trivial to create smooth utility functions. Therefore, also in international political negotiations, assuming some probabilistic components seems valid.
19	Ordering of alternatives	Agents have a distinct ordering of preferences for indivisible issues with discrete values, either on a <i>ordinal</i> (ranking) ^[5] or <i>interval</i> scale.	As we assume agents only need to be weakly preferred, implying that $u(x_A) \geq u(x_B)$ for alternatives A and B respectively. So, even if negotiators are not able to choose between A and B , they can still weakly prefer both over a third alternative C [11, p. 24-26]. Though, results of this thesis become invalid if both $u(x_A) < u(x_B)$ and $u(x_A) > u(x_B)$. This may be due to ongoing domestic negotiations. For these implications, see point 2.
20	Strategy profile	An agent's strategy profile is represented by i) bidding strategy which is a weighted combination of tactics (i.e. <i>time-dependent</i> , <i>resource-dependent</i> and <i>behavior-dependent</i>) and ii) acceptance strategy defining when to accept offers.	Consequently, in our model, an agent's strategy is solely related to making offers. Among other Jennings et al. (2003, [139]) show that most negotiators do more than that in real negotiations. They, for example, make threads, promises, rewards and other incentives, not captured by a single offer. They may also provide arguments why offers are rejected, potentially leading to better counteroffers and, thus, more efficient negotiations [139, p. 343-345]. This is especially true in situations in which negotiating parties are exploring their options by gathering information by means of negotiating [139, p. 346-347]. Therefore, conclusions in this thesis are only valid in those negotiation settings in which parties already have some idea about opponents' preferences. So, in a more mature stage of the negotiation process in which the gathering of information has already silently been done by diplomats.
21	Opening strategy	An agent's initial offer is always centered around its <i>target values</i> and, in our tool, drawn from a skewed normal distribution.	It is very common for negotiators to begin making high initial offers, even in excess of the target value [137, p. 27]. A reasonable explanation for this is that they are motivated by position loss and image loss [56]. Therefore, we expect this assumption to affect model validity positively. Though, as a consequence of using skewed normal distribution centered around an agent's target value, simulated negotiation outcomes might be slightly positively biased.

Supporting analyses^[6]

- From Figure C.1 we can conclude that changing preferences during the actual negotiations seriously affect the outcomes. Not only the (median) value of the outcomes are affected, but also in many simulations the spread of outcomes. The latter can be best seen for issue x in Figure C.1. Moreover, in Figure C.2 it can be observed that also the negotiation time, size and power of the winning coalition are affected by changing preferences to some extent. With regard to negotiation time, it seems that the time to agreement becomes unpredictable, as the spread of the data increases. Both size and power of the winning coalition is least affected, but still the results show differences, especially for high likelihoods agents' preferences may change. Consequently, results of this thesis solely apply in negotiation settings in which states have more or less stable preferences for a certain period of time. So, in situations of (sudden) domestic pressures – for example, before national elections – our results become meaningless;

^[5]For convenience, we use an inverted scale in which the highest number denotes the most preferred alternative.

^[6]Similar to the validation of the experimental design in Section G.2, we use the negotiation setting and domain of the example run of Section 5.5.3 to validate the assumptions. Note that for validation purposes, the example run may be altered.

5. In Figure C.3 we show the results of two simulation experiments. Again, the 'base' setting resembles the assumption in this thesis, whereas 'alt1' includes two negotiators – a_1 and a_2 — whose offers are considered with higher likelihood because of their experience. As expected, both negotiators obtain slightly higher utilities as they have a greater influence on the bargaining process. We also observe that spread of negotiation outcomes decreases, most likely because they tend to center around the preferences of a_1 and a_2 . More clearly, as their offers are considered the majority of times, they have a greater chance of being accepted (as a consequence of diminishing time) and, therefore, the final negotiation outcomes of 200 simulations are closer together. So, by disregarding the experience of negotiators, we also neglect the possibility more experienced negotiators bend negotiations to their will. Our results, therefore, cannot be generalized to settings with great experience imbalances. Though, with some simple modifications, the tool can be modified in such way that behavioral power is considered. However, incorporating 'network-like' negotiations – and derived behavioral power – is more cumbersome;
6. In Figure C.4 we present the negotiation outcomes for 200 simulations of our 'base' model – assuming dynamic mixtures between conflictive and cooperative strategies – and 'alt1'. In the latter, agents can only statically choose (ex ante) between either conflictive and cooperative strategies. As can be seen, the outcome space significantly reduces by assuming static strategies. Therefore, this assumption actually improves model behavior as it resembles real political negotiations better. As a consequence, we increase variability in outcomes. So, our tool should not be used to predict future negotiation outcomes. Rather, it must be used to analyze pre-coalition strategies in capricious negotiation settings;
9. In Figure C.5 the negotiation outcomes are presented in which we test the *power excess* assumption. In the 'base' experiment, parties make proposals by his/her power excess. Moreover, 'equal' describes the setting in which parties have equal power positions within the coalition – no matter what their actual structural power is. Finally, in 'longest', parties part of the initial coalition hold more power than those who recently joined. We observe that the latter affects the negotiation outcomes the most, whereas 'base' and 'equal' do not differ much. Considering negotiation time, size and power of winning coalition in Figure C.6, we observe fairly similar results for each experiment. Thus, overall, changing the power excess assumption does not affect the results of this thesis much. So, conclusions may be generalized even if slightly different within coalition power positions hold;
10. In Figure C.7 the negotiation outcomes and derived utility are presented in which we compare different methods to calculate a state's true power position. We observe that all simulation experiments – i.e. Shapley-Shubik (base), actual structural power and Banzhaf index – show fairly similar results regarding negotiation outcomes and derived utility. So, our model is relatively robust for different ways of calculating a state's strategic demands. Therefore, this assumption has relatively little impact in the model, keeping in mind that all methods are based on structural power. This, implying that if negotiators are likely to make demands based on other (historic) reasons, conclusions of this thesis cannot be generalized;
16. In Figure C.8 the negotiation outcomes and derived utility are presented in which we compare different selection procedures – i.e. Markov process, 'base', 'random' and 'power' approaches. Overall, both plots show the selection procedure has not a substantial impact on negotiation outcomes. Though, we observe that selecting negotiators randomly results in the most substantial differences. This is not surprising, as the relative influence of weak agents is equal to those of strong agents. Thus, outcomes are affected by weak agent's demands, which is also reflected in the derived utility (see a_4 and a_5). Moreover, the Markov and power-based approach lead to very similar outcomes. Again, not surprising, as both selection mechanisms are (initially) based on power, but differ slightly over time. As we expect offers of more powerful negotiators are considered first in an international setting, the impact of this assumption on the conclusions of this thesis is minor.




Figure C.1: Negotiation outcomes of 200 simulations in which preferences can change.




Figure C.2: Negotiation time, size and power of winning coalition in negotiations in which preferences can change.

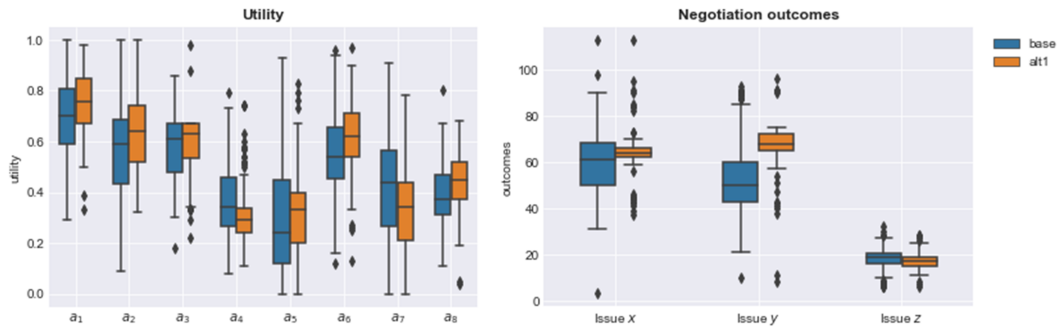


Figure C.3: Utility and negotiation outcomes of 200 simulations in which a_1 and a_2 have more experience.




Figure C.4: Negotiation outcomes of 200 simulations with dynamic ('base') and static ('alt1') strategy mixtures.




Figure C.5: Negotiation outcomes of 200 simulations testing the power excess assumption.




Figure C.6: Time, size and power of winning coalition of 200 simulations testing power excess assumption.

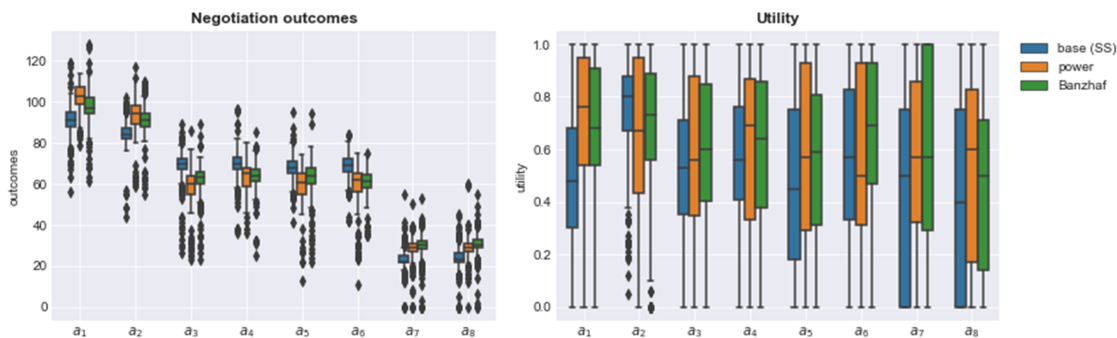


Figure C.7: Negotiation outcomes and derived utility of 200 simulations testing the true power position.




Figure C.8: Negotiation outcomes and derived utility of 200 simulations testing the selection procedure.



Model Requirements

As we are developing a *software* tool (i.e. model) to simulate coalition formation, there is a need to structure the design process in more detail. The first four chapters of this study have presented numerous methods, models, functions, theories and other solution concepts that can be used or have been developed to analyze coalition formation behavior. In addition, together with the research definition, they define the "contours" of the to be developed automated negotiation model. This appendix will define these rather vague "contours" more precisely by formulating *software requirements* for the automated negotiation model. These requirements express the needs and put constraints on the final model in order to develop a tool that is able to sufficiently answer the main research question.

In order to do so, we will use *SWEBOK V3.0: Guide to the Software Engineering Body of Knowledge* – further denoted with SWEBOK – of the IEEE computer society [24]. At most basic, SWEBOK defines a software requirement as [24, p. 1-1-1-2]: "(...) a property that must be exhibited by something in order to solve [analyze] some problem in the real world". An essential property of such requirements is that it needs to be *verifiable* as an individual feature of the software model. Sometimes it may be hard to verify certain requirements. Thus, they should be verified within the available resource constraints. Let us assume that it is required that the to be developed negotiation model should be used by more than 500 users in parallel. Then, for a study of this size, testing this requirement becomes infeasible and we should either lower our requirement or motivate why we think we have (not) accomplished it without testing. Another important property of software requirements is that they often have a *priority rating*. This enables us to make trade-offs during the design as our resources are limited [24, p. 1-2]. The SWEBOK distinguishes the following types of software requirements [24, p. 1-2-1-1-3]:

- *Product requirement*. This type defines a need or puts a constraint on the software as a whole – i.e. on the to be delivered software "product".
- *Process requirements*. This type of requirement puts a constraint on the development process of the software. These requirements are not in scope of this research;
- *System requirements*. Similar to software requirements, only than for the system as a whole. So, they are applicable to the interaction of different software tools together designed for a defined objective. The latter are not in scope of this research.

Each of the above requirements are divided in two dimensions, namely:

- *Functional requirements*. These requirements describe the functions software must be able to execute;
- *Non-functional requirements*. Also known as quality requirements act to constrain the solution. Many types exist, among others performance, maintainability, safety, reliability, interoperability and more. These requirements are only partly in scope of this research;

To the above dimensions quite often two other components are added, that is to say *priority* and *scope*. As described, the first enables us to make trade-offs during the development process between (functional)

requirements of higher or lower importance. As no prioritization method is specified in the SWEBOK, we choose to apply the *MoSCoW method* as it is simple to apply and easy to understand ^[1]. This method is an acronym for "Must have", "Should have", "Could have", and "Won't have", which are four prioritization categories in decreasing order of importance. The scope of a requirement refers to the extend software or software components are affected by the requirement. Especially non-functional requirements will affect the product as a whole (i.e. "globally"), while functional requirements often affect specific components. As our automated negotiation model has three major components (see Chapter 4), our scope is defined in terms of a *protocol*, *scenario* and *strategy*, accompanied with *global*, *agent* and *pre-negotiation*, the component that models the first stage in international political negotiations.

The table on the following pages set the (product) requirements for developing the automated negotiation tool for analyzing coalition formation behavior in international political negotiations. The requirements will be referred to by means of their *#ID*. For consistency, we refer to the *tool* if we indicate the entire development application – as displayed in Figure 5.1. In contrary, with *model* we indicate only the automated negotiation environment, whereas with a *run* we imply a single simulation of the negotiation process in scope – as displayed in Figure 2.2.

Table D.1: Requirements for the to be developed simulation tool.


#ID	Name	Description	Type	Scope	Priority
AG1	Number of agents	The software should be able to simulate international political negotiations with at least 20 agents (e.g. diplomats) involved.	Functional	Agent	M
AG2	Agent's strategy adjustment	Agents should be able to intelligently adjust their strategy profile over time, to represent their (concession) behavior and reflect opponents offers.	Functional	Agent	S
AG3	Custom utility functions	The model should be able to deal with custom utility functions that differ for each agent.	Functional	Agent	S
AG4	Power indexes	The model should accommodate multiple ways to calculate an agent's power (excess) at a specific moment in time.	Functional	Agent	S
AG5	Agent's character	Agents could reflect real negotiators characteristics, among others its risk behavior (both averse and postpone).	Functional	Agent	C
GL1	Sufficient coalition partners	The model must evaluate the sufficiency of coalitional arrangements based on yet to be determined measures.	Functional	Global	M
GL2	Simulations	The tool must be able to generate different coalitional arrangements computationally, given a certain scenario (and uncertain input parameters).	Functional	Global	M
GL3	Uncertainty	The tool must account for uncertainty in the input parameters (see GL2) and be able to represent it quantitatively.	Functional	Global	M
GL4	Open-source programming language	For re-usability purposes, we should a widely applied programming language that is freely available to researchers and others.	Non-functional	Global	M

^[1]For a concise description, see https://en.wikipedia.org/wiki/MoSCoW_method.

GL5	Reliability	The tool must at most have an error once every 10,000 runs.	Non-functional	Global	M
GL6	Visualization single run level	Visualizations displaying the outcomes and coalitions on a single run level should be made available.	Functional	Global	S
GL7	Visualizations for investigating sufficient outcomes	Visualizations comparing different coalitional arrangements should be made available for different sufficiency statistics.	Functional	Global	S
GL8	Dependencies	The model should have little dependencies with packages that are not frequently maintained/updated.	Non-functional	Global	S
GL9	Adhere to style guide principles of programming community.	During the code development, we should adhere to style guide principles set out by the particular programming community in order to ensure readability and reuse-ability.	Non-functional	Global	S
GL10	Documentation	The different modules of the programming tool should be thoroughly documented to ensure future use.	Non-functional	Global	S
GL11	Performance	Every model run (or negotiation) should take «1 second.	Non-functional	Global	S
GL12	Reporting	The most important simulation data (e.g. outcomes, utility, time, runtime, coalition members) should be made available for custom reporting/visualization purposes.	Non-functional	Global	S
GL13	Optimal coalitions	The software model could find optimal coalition partners, based on the negotiation scenario, protocol and agents involved.	Functional	Global	C
GL14	Visualization simulation level	Multidimensional visualizations for multiple runs over multiple issues could be made available.	Functional	Global	C
GL15	Multiple optimization algorithms	The tool could support multiple (or custom) optimization algorithms and compare the best one for a given scenario.	Functional	Global	C
GL16	User interface	The tool will not provide an interface for potential users in which they can insert new scenarios or in which outcomes are visualized.	Non-functional	Global	W
PN1	Pre-coalitions based on ideological values	The model must be able to find pre-coalitional arrangements based upon agents' ideological values.	Functional	Pre-negotiation	M
PN2	Pre-coalitions based on opponent's power	The model must be able to find pre-coalitional arrangements based upon agents' power asymmetry.	Functional	Pre-negotiation	S
PN3	Visualization of coalition formation	Visualizations continuously displaying the dynamic pre-negotiations and negotiations coalitions could be made available.	Functional	Pre-negotiation	C

PR1	Package-deals within subjects	Agents within the simulation software should be able to suggest different package-deals within a certain subject (i.e. combination of issues) and therefore allow trade-offs in one area against another.	Functional	Protocol	M
PR2	Voting rules	The model must offer the most basic voting rules (e.g. majority, unanimity, super majority).	Functional	Protocol	M
PR3	Select next agent	The protocol should support multiple ways of selecting initial proposers (and opponents) when negotiating.	Functional	Protocol	S
PR4	Coalition formation	The protocol should be able to allow agents to join, leave, stay or accept new members in the coalition.	Functional	Protocol	S
PR5	Package-deals across subjects	Agents within the simulation software should be able to suggest different package-deals across a certain subject (i.e. combination of issues).	Functional	Protocol	C
PR6	Agent interactions	The protocol could support multiple ways of agent interaction when dealing with different issues and subjects.	Functional	Protocol	C
PR7	Agenda setting	The software model should be able to allow agents setting up the negotiation agenda themselves by means of negotiation or during the negotiation process.	Functional	Protocol	W
PR8	New issues	The tool will not be able to allow agents bringing new issues to the bargaining agenda during the negotiation process.	Functional	Protocol	W
PR9	Rich proposals	Agents are not able to motivate their offers by any means (e.g. textual arguments).	Functional	Protocol	W
PR10	Re-negotiations	The model will not be able to deal with "re-negotiations". So when an outcome is reached, agents cannot go back in time.	Functional	Protocol	W
SC1	Number of subjects	The software should be able to simulate international political negotiations with at least 20 subjects on the negotiation agenda.	Functional	Scenario	M
SC2	Binary issue	The software should enable bargaining over issues of a binary type.	Functional	Scenario	M
SC3	Discrete issue	The software should enable bargaining over issues of a list (discrete) type.	Functional	Scenario	M
SC4	Unlimited range issue	The software should enable bargaining over issues of an unlimited range (continuous) type.	Functional	Scenario	M

SC5	Investigate influence of power asymmetry, ideological values and conflicting preferences	The model should be able to investigate the influence of power asymmetry, ideological values and conflicting preferences on finding good coalition partners, given a certain scenario.	Functional	Scenario	S
SC6	Custom input parameter distributions	The tool must be able to accept custom distributions for different input parameters (other than the normal distribution).	Functional	Scenario	S
SC7	New scenarios	Users of the model must be able to update or insert new scenarios within 15 minutes, assuming the protocol does not change.	Non-functional	Scenario	S
SC8	Limited range issue	The software should enable bargaining over issues of a limited range (continuous) type. The latter is fixed by a limited resource and cannot take any value.	Functional	Scenario	C
SC9	Formalized input	The tool could be able to read structured input (e.g. new scenarios) by means of, for example, XML or JSON standards.	Functional	Scenario	C
ST1	Time-dependent strategy	An (adjustable) time-dependent strategy should be constructed that allows agents to "feel" time pressure in their bidding behavior.	Functional	Strategy	M
ST2	Behavior-dependent strategy	An (adjustable) behavior-dependent strategy should be constructed that allows agents to mimic their opponents offers in order to capture real-life negotiation behavior .	Functional	Strategy	M
ST3	Resource-dependent strategy for coalition pressure	An (adjustable) resource-dependent strategy should be constructed that enables agents to make concessions as the number of "free-riders" (i.e. agents who are not part of a coalition) decreases.	Functional	Strategy	M
ST4	Resource-dependent strategy for limited assets pressure	An (adjustable) resource-dependent strategy should be constructed that enables agents to take more risk as a shared asset (or resource) reaches its limited.	Functional	Strategy	S
ST5	Intelligent acceptance strategy	Agents could be able to intelligently adjust their acceptance strategy, based upon previous offers.	Functional	Strategy	C
ST6	Custom strategies	The model could allow for custom (intelligent) agent strategies.	Functional	Strategy	C



Class Diagram of the Two-Stage Model

In Figure E.1 we present an overview of the most important relationships between the different classes in the automated negotiation environment and its entities – i.e. agents and coalitions. This is done by using UML class diagram^[1]. Note, that this diagram represents the environment in which the entities actually negotiate over subjects and issues, thus *without* the ability to form coalitions *ex ante*. Note too that this diagram is still a high-level (conceptual) representation, as the implemented environment consists of many more attributes and methods. Though, the shown classes are the actual implemented classes in the negotiation environment.

In the top of the diagram, we observe that a *negotiation* consists of one set of rules and an agenda. The *set of rules* guide the way in which agents interact and when they are allowed to make proposals. Moreover, it also defines when negotiations end in agreement (or not), thus it operates the decision rule in place. Together with one negotiation agenda, consisting of the order of subjects and their respective deadline, the set of rules form the negotiation **protocol**. Thus, a single negotiation always uses always one clearly defined protocol.

Furthermore, we observe that a negotiation consists of several *rounds*. Subsequently, each round deals with one *subject* in itself containing multiple *issues*. This implementation therefore indirectly shows that that we deal with subjects sequentially (in rounds), whereas issues are dealt with simultaneously (within the subjects) or in package-deals. Together, the subjects and issues conclude the negotiation **domain** (i.e. what is discussed?).

As this environment deals with multilateral negotiations singly, each round contains at least three *agents*. Of course, these agents have unique names, but they also have distinct power positions, ideology, private deadlines and more. Moreover, each agent can make or accept offers, form coalitions but also has a personal utility function. At least two or more agents can form *coalition* together, independently whether this happens before or during the actual negotiations. In a coalition together, agents make combined offers to opponents and accept potential offers together.

In addition, every agent has one unique **strategy profile**. This profile consists of one or more *bidding strategies*, which is an individual aggregation of tactics. In turn, these tactics are either time-, resource- or behavior dependent, as presented in Faratin et al. (1998, [45]). In our final model, we have assumed agents have four different strategies, depending on his opponent. These include matcher, concenter, inverter and competitor. Moreover, each agent has a distinct *acceptance strategy*, which determines whether to accept or reject offers from other agents. For sake of simplicity, in our final model, we have assumed agents accept offers if they are more advantageous than offers they would make themselves. However, this can, of course, simply be adjusted in the negotiation environment. Agents have also a certain *coalition strategy*, which basically entails how willing they are to be part of a certain coalition (or not).

Finally, the negotiation environment is completed by incorporating agents' preference profiles. Note that,

^[1]Unified Modeling Language.

technically, this is not part of the (general) negotiation environment, but only of a specific instance. This *preference profile* describes an agents preferences towards the possible outcomes of the negotiation domain. Thus, it includes certain weights for the various issues to discuss. But, more importantly, it includes an agents target values (i.e. what he aims for) and reservation values (i.e. what he deems just acceptable). Together with the negotiation domain, the preference profile finalizes the negotiation **scenario**.

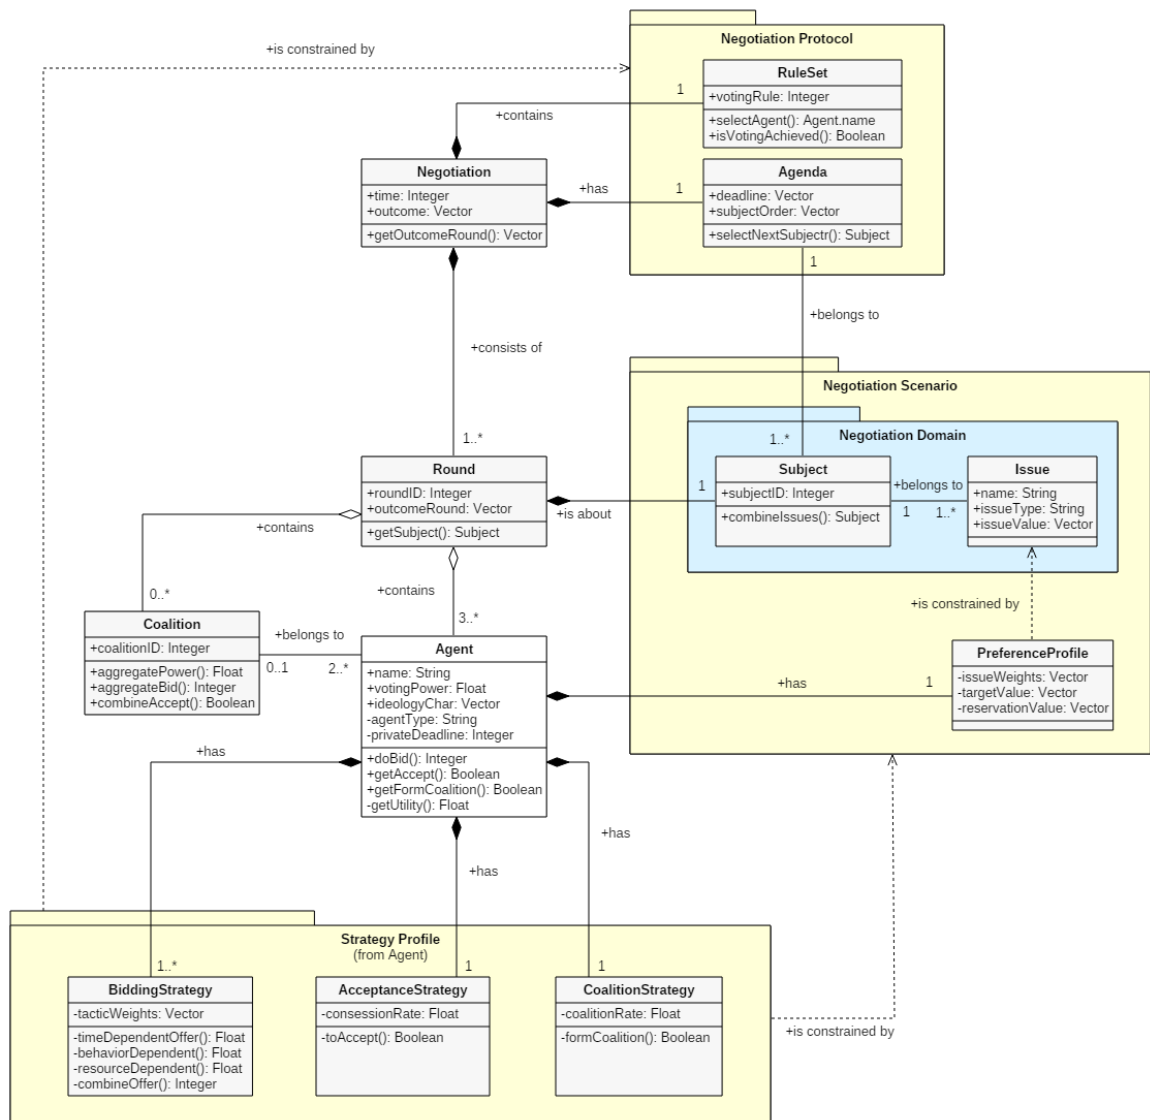


Figure E.1: High-level class diagram (UML) of the automated negotiation environment.



Code Testing and Verification

This appendix contains both the verification test results as well as tests to assess code quality. We will start by testing the quality of the Python code of our negotiation tool. Thus, Section F.1 focuses on the quality of the delivered software itself. In contrary, in Section F.2, we run several verification tests to ensure that the implemented (computational) model represents the formal (conceptual) model as defined in Chapter 5.

F.1. Testing tool quality: is it any good?

Of course in many ways, the term *quality* is a highly ambiguous. To some, quality may refer to the software itself – i.e. reliability, maintainability or scalability – while for others the process of developing the software is in focus, which implies they refer things as development time and costs [168, p. 1]. Another, much simpler definition was provided by Philip Crosby as he refers to quality as "...conformance to requirements" [34, p. 92]. As our tool is developed for academic purposes, we are basically the only relevant stakeholder in the development process. Therefore, we are not referring to process-oriented approaches of software quality (see among others [53, p. 48-49]). Moreover, remember that our tool – and therewith our code – serves three major purposes, namely:

- i) it must allow us to answer the research question(s) in this thesis;
- ii) it should allow others to review our work and;
- iii) it should allow others to extend our work.

To accomplish the first goal, all *functional* requirements prioritized as "(M)ust have" must be met. Moreover, for feasibility purposes, errors should be reduced as much as possible and calculations have to be done relatively efficient and quick – and thus we should aim for *reliability* and *efficiency*. To achieve the second point, *usability* of the code is vital, as reviewers should be able to understand and learn the code. Finally, to accomplish the third and last point, our code should be *maintainable* allowing potential users to make modifications. All the aforementioned prerequisites are described in the ISO-9126^[1] model for high-level software quality guidelines (see [70]). Therefore, we refer to software *quality* as conforming to the high-level requirements as stated in ISO-9126 model, with the exception of adhering to the *portability* requirement (see Table F.1). So, to assess the quality of our code, in the next subsection we first perform an assessment whether or not the functional requirements – as presented in Appendix D – have been met. Moreover, in Section F.1.2 we apply Python's *Pylint* package to check for coding standards (with Python's PEP8 style guide as a benchmark), error detection and refactoring. Moreover, we use the *Radon* package to determine code complexity and an overall maintainability score. Finally, we assess the code's time efficiency (or performance) by conducting several simulation experiments. By doing so, we assess non-functional requirements (i.e. #GL5, #GL10 and #SC7) on reliability, usability, maintainability and efficiency.

^[1]Or its more extensive replacement: the ISO-25010 model (see [69]).

Table F.1: ISO-9126 high-level model for quality guidelines (adopted from [53, p. 51]).

Quality	Description	Main focus	Means
<i>Functionality</i>	Capability of software to provide functions, which meet stated needs when software is used under specified conditions	All functional requirements	Functional requirement assessment
<i>Reliability</i>	Capability of software to maintain a specified level of performance when used under certain conditions	Stability/errors (e.g. #GL5)	<i>Pylint</i> and simulation experiments
<i>Usability</i>	Capability of software to be understood, learned, used, and attractive to the user	Documentation (e.g. #GL10 and #SC7)	<i>Pylint</i> and cyclomatic complexity (<i>Radon</i>)
<i>Efficiency</i>	Capability of software to provide appropriate performance, relative to the amount of resources used, under stated conditions	Time performance (e.g. #GL11)	Simulation experiments
<i>Maintainability</i>	Capability of software to be modified. Modifications may include corrections improvements or adaptation of the software to changes in environment, and in requirements and functional specifications	Coding standards (e.g. #GL9)	<i>Pylint</i> and maintainability index
<i>Portability</i>	Capability of software to be transferred from one environment to another	Not in scope	Not in scope

F.1.1. Functionality: assessment of requirements

In this section, we examine the functionality of our tool by performing a functional requirement assessment. Table F.2 presents this assessment, based upon the *functional* requirements set in Appendix D and accompanied with a brief description of how the assessment was conducted. As can be concluded from this table:

- All 15/15 "Must haves" (M) functional requirements have been met;
- A considerable 8/11 "Should haves" (S) functional requirements have been met;
- A few or 3/10 "Could haves" (C) functional requirements have been met and;
- As expected, none or 0/4 of the "Won't haves" (W) functional requirement have been met.

As our tool has met all the "Must haves" (M) functional requirements and a considerable amount of "Should haves" (S) and "Could haves" (C), we are confident that it's functional quality is adequate enough to meet our first goal. In the next section, we will discuss whether the quality of the tool is also abundant to meet the second and third goals respectively.

Table F.2: Functional requirement assessment.

#ID	Name	Priority	Assessment	Achieved
AG1	Number of agents	M	In the efficiency experiments (see Section F.1.2), we have effectively achieved multiple simulation runs with 50 agents.	Yes
AG2	Agent's strategy adjustment	S	The names of different strategy profiles (i.e. 'inverter', 'conceder', 'competitor', 'matcher') were printed out in a single simulation run. By (manual) observation, we have seen that agents reacted upon their opponents' strategy profiles and the correct names were printed.	Yes

AG3	Custom utility functions	S	As a compromise, (independent) Euclidean distance utility function is hard coded and cannot be changed.	No
AG4	Power indexes	S	The power positions of agents in the tool can be inserted manually using different power indexes. The latter can be computed using Python's <i>powerindex</i> package. Thus, different power indexes can be incorporated in the model, but it is not specifically facilitated.	No
AG5	Agent's character	C	There are two model parameters reflecting an agent's risk and rationality characteristics (i.e. ρ and β) and a third, hard coded function describing capriciousness. The latter is remained constant for all agents.	Yes
GL1	Sufficient coalition partners	M		
GL2	Simulations	M	There is a simulation function in the <i>simulation.py</i> module which automatically runs N negotiations and allows users to adjust some simulation parameters.	Yes
GL3	Uncertainty	M	The user can set a generic uncertainty parameter σ in the <i>simulation.py</i> module (see above).	Yes
GL6	Visualization single run level	S	The tool does not provide any custom visualizations. However, we have aggregated all relevant data and prepared them in such a way that they can be (easily) visualized using standard Python plot libraries.	Yes
GL7	Visualizations for investigating sufficient outcomes	S	See above.	Yes
GL13	Optimal coalitions	C		
GL14	Visualization simulation level	C	Although the data is available, we have not been able to present this multinomial data effectively.	No
GL15	Multiple optimization algorithms	C		
PN1	Pre-coalitions based on ideological values	M	The user is allowed to insert a spatial model Λ based upon all kinds of (quantifiable) differences between agent.	Yes
PN2	Pre-coalitions based on opponent's power	S	The above spatial model Λ can easily be extended with agents' power positions, if quantifiable.	Yes
PN3	Visualization of coalition formation	C	See #GL6, #GL7 and Figure 5.10 in which the coalition formation process is visualized for an example model run.	Yes
PR1	Package-deals within subjects	M	This is hard coded in the sequential (simultaneous) implementation of subjects and issues as described in Section 5.3.	Yes
PR2	Voting rules	M	The voting rule is implemented as an input parameter (i.e. χ) in which the user is allowed to select simple or super majority, (super) weighted majority and unanimity rule. Moreover, in case of super majority, the user is also able to insert the fraction of votes/agents needed to reach an agreement.	Yes
PR3	Select next agent	S	The selection rule of the model can be adjusted by changing parameter ζ . The rules incorporated and verified (see Section F.2) are based on <i>randomness</i> and <i>power</i> .	Yes

PR4	Coalition formation	S	This is hard coded in the protocol. Numerous of experiments have been ran and by examining the results we have indeed seen that agents join, leave, brake up or accept new coalitions.	Yes
PR5	Package-deals across subjects	C	This is not supported by the tool.	No
PR6	Agent interactions	C	As the protocol is hard coded, the only way of interaction is participating in multilateral negotiations is by performing series of bilateral negotiations.	No
PR7	Agenda setting	W	This is not supported by the tool. The inserted domain does also account for the negotiation agenda.	No
PR8	New issues	W	This is not supported by the tool. Agents can only negotiate on subjects known in advance.	
PR9	Rich proposals	W	This is not supported by the tool. The only information agents have (and give) is incorporated in their (historic) bids.	No
PR10	Re-negotiations	W	This is not supported by the tool. If an agreement is struck – i.e. the voting rule has surpassed – agents respect the outcomes.	No
SC1	Number of subjects	M	In the efficiency experiments (see Section F.1.2), we have effectively achieved multiple simulation runs with 20 subjects.	Yes
SC2	Binary issue	M	The tool allows for binary issues by inserting a (preference) dictionary with only two issue alternatives.	Yes
SC3	Discrete issue	M	The tool allows for discrete issues by inserting a (preference) dictionary with three or more issue alternatives.	Yes
SC4	Unlimited range issue	M	The tool allows for unlimited range issues by inserting a (preference) range with a target and reservation value.	Yes
SC5	Investigate influence of power asymmetry, ideological values and conflicting preferences	S	This is supported by the aggregation of simulation results (see #GL6 and #GL7).	Yes
SC6	Custom input parameter distributions	S	This is not supported by the tool. Only generic uncertainty can be inserted by altering uncertainty factor σ (see #GL3).	No
SC7	Limited range issue	M	The tool allows for limited issues by inserting a (preference) dictionary in which the keys represent the agents and values the ranges the agent is preferred to have. Moreover, a limit (or pie) can be set by inserting the negotiation domain. This limited has been verified (see Section F.2).	Yes
SC10	Formalized input	C	This is not supported by the tool. Input can only be given in Python code.	No
ST1	Time-dependent strategy	M	This is incorporated and verified in the model by multiple simulation experiments (see Section F.2).	Yes
ST2	Behavior-dependent strategy	M	This is incorporated and verified in the model by multiple simulation experiments (see Section F.2).	Yes

ST3	Resource-dependent strategy for coalition pressure	M	This is incorporated and verified in the model by multiple simulation experiments (see Section F.2).	Yes
ST4	Resource-dependent strategy for limited assets pressure	S	This is incorporated and verified in the model by multiple simulation experiments (see Section F.2).	Yes
ST5	Intelligent acceptance strategy	C	This is not incorporated in the tool.	No
ST6	Custom strategies	C	This is not supported by the tool.	No

F.1.2. Examining non-functional requirements

As maintainability, usability, reliability and efficiency are to some extent interrelated, we discuss them together in this section. We start with the Pylint package, which is widely applied to check the quality of source codes of different modules written in Python. Pylint's basic functionality generates among others an aggregation of different messages by the following categories^[2]:

- *[R]*efactor for a "good practice" metric violation;
- *[C]*onvention for coding standard violation;
- *[W]*arning for stylistic problems, or minor programming issues;
- *[E]*rror for important programming issues (i.e. most probably bug);
- *[F]*atal for errors which prevented further processing;

In which *[R][C][W]* are considered to represent coding standards for maintainability purposes, whereas *[E][F]* are considered to represent (serious) reliability issues. Table F.3 shows the results for the different modules in our tool including an overall automatically generated grading score^[3]. In addition, Table F.4 and Table F.5 present documentation and bad name convention statistics respectively. Let us first discuss the results regarding maintainability.

Maintainability

We observe that most modules in our tool face messages regarding code convention *[C]* in comparison to Python's PEP8 guidelines. This is can be expected, as code convention entails a comprehensive list of style "errors", such as indentation, comments, statements, doc strings white spaces and naming conventions. As code conventions are typically trade offs between someones own coding style, readability, development time and guidelines, compromises have to be made. We have prioritized bad name conventions (see Table F.4) above stylistic guidelines (e.g. white spaces), as good name conventions increases understandability and therewith maintainability of the code. As can be seen in Table F.4, the majority of modules have 0% bad name conventions. The ones who ha are considered to be compromises, as changing these name conventions takes considerable amount of times. The latter can thus be seen as future improvements to our tool. Moreover, we observe as well that many messages entail refactoring *[R]* notices. Code refactoring is the practice of restructuring existing code, without changing its behavior (e.g. output). It is, therefore, mainly done for readability purposes. If we look closer to the individual messages, we see that many refactoring messages concern too many (class, function or method) arguments. Again, we have compromised on restructuring our code to improve readability, given the development time. Thus, we take the remaining refactoring messages for granted which again can be seen as future improvements to our tool. Finally, in Table F.3, we observe that mainly the *agent.py* and *round.py* and to lesser extend *strategies.py* and *model.py* have some minor stylistic or programming warnings *[W]*. If we look at these messages closer, they appear to be minor programming warnings in which we have redefined local variable names identical to global names. As these warnings not influence the reliability nor the efficiency of our tool, we accepts these minor warnings for the time being. So, to conclude, we find the overall maintainability of our code acceptable and reasonable in line with PEP8

^[2]Retrieved from https://pylint.readthedocs.io/en/latest/user_guide/output.html on the 21st of February 2018.

^[3]Note that this score ranges from -10 to 10.

Table F.3: Output of Pylint by category for each module in our tool.

	<i>agent.py</i>	<i>coalition.py</i>	<i>protocol.py</i>	<i>round.py</i>	<i>scenario.py</i>	<i>strategies.py</i>	<i>model.py</i>	<i>reporting.py</i>
[R]	12	3	3	16	4	9	10	0
[C]	42	24	21	60	4	4	27	3
[W]	8	1	0	6	1	4	3	0
[E]	1	0	0	4	0	1	6	0
[F]	0	0	0	0	0	0	0	0
Grade	5.8/10	7.3/10	7.0/10	5.0/10	5.5/10	7.3/10	5.7/10	8.3/10

standard.

Usability

Secondly, we discuss the usability of our tool. There are many factors that can importer the usability. These include documentation, code complexity and user interfaces (e.g. [53, p. 52-58]). Table F.5 shows Pylint's results regarding the documentation of our model. This includes generic module descriptions, more detailed class attribute and method descriptions as well as function and argument descriptions. We see that all the necessary documentation is present and according to PEP8 standards. Moreover, we have measured code complexity by determine the cyclomatic complexity of our tool. This metric – introduced by Thomas McCabe in 1976 [104] and implemented by the *Radon*^[4] package in Python – defines complexity as the number of linearly independent paths within source code. Note that this is by far not the only measure of code complexity, but one of the first and widely used metrics in Python. In Table F.6, we present the output of the Radon package for our tool. In this output, complexity is classified in ranks from *A* to *F*, where *A* stands for the simplest and *F* the most complex code^[5]. In the final column of Table F.6, we show the number of linearly independent paths. From this analysis, we observe that only 7 out of 46 is classified as *C* (> 10 independent paths): moderate risk, slightly complex block. These are precisely the blocks of code covering the most complex parts of our model. So, the increased complexity of these parts does not come as a surprise. Moreover, they are still acceptable with respect to usability. Finally, there is much interaction between the maintainability and usability requirements of code. Therefore, Radon is also able to generate a maintainability score, which is a score between 0 and 100. This score is based upon a factored formula consisting of the total lines of code, cyclomatic complexity and Halstead volume^[6]. The latter is another measure of complexity – introduced by Maurice Halstead in 1977 [61] – in which many more measurable properties of software are included. The maintainability scores and accompanying ranks are displayed in Table F.7. As can be concluded from this table, all modules have more than acceptable scores for maintainability, which underlines our conclusion in the previous paragraph. To finalize the usability requirements, we consider user interfaces. Unfortunately, our tool does not facilitate them in any way. Partly, because a GUI was not a prerequisite to achieve our main goals (see the introduction of Section F.1), but also to compromise development time. Ideally, we could have facilitated the input of new negotiation domains by simple XML or JSON descriptions so that the tool can be used without (much) Python knowledge. The latter may hamper our second goal, especially if potential reviewers of our work are less skilled in Python. Altogether, we find the usability of our tool acceptable enough for answering the research question(s) in this thesis. However, there is still room for improvement if we consider the second and third goal of our tool: allow others to review and extend our work respectively.

^[4]See <http://radon.readthedocs.io/en/latest/>

^[5]For more information, please consult Radon's documentation at <http://radon.readthedocs.io/en/latest/api.html#radon.visitors.ComplexityVisitor>.

^[6]For more information, please consult Radon's documentation at <http://radon.readthedocs.io/en/latest/intro.html>.

Table F.4: Bad name convention for each module in our tool.

	<i>agent.py</i>	<i>coalition.py</i>	<i>protocol.py</i>	<i>round.py</i>	<i>scenario.py</i>	<i>strategies.py</i>	<i>model.py</i>	<i>reporting.py</i>
Module	0%	0%	0%	0%	0%	0%	0%	0%
Class	0%	0%	-	0%	0%	-	0%	-
Method	33%	27%	-	40%	0%	-	0%	-
Function	-	-	80%	-	-	0%	-	100%

Table F.5: Documentation for each module in our tool.

	<i>agent.py</i>	<i>coalition.py</i>	<i>protocol.py</i>	<i>round.py</i>	<i>scenario.py</i>	<i>strategies.py</i>	<i>model.py</i>	<i>reporting.py</i>
Module	100%	100%	100%	100%	100%	100%	100%	100%
Class	100%	100%	-	100%	100%	-	100%	-
Method	100%	100%	-	100%	100%	-	100%	-
Function	-	-	100%	-	-	100%	-	100%

Table F.6: Output of Radon package classifying cyclomatic code complexity.

Module	Block type	Block name	Complexity rank	Independent paths
<i>agent.py</i>	C	Agent	A	4
<i>agent.py</i>	F	opp_bids	B	10
<i>agent.py</i>	F	conv_bids	A	3
<i>agent.py</i>	M	Agent.bid	B	10
<i>agent.py</i>	M	Agent.utility	B	6
<i>agent.py</i>	M	Agent.opponentModel	B	6
<i>agent.py</i>	M	Agent.generate_bid	A	4
<i>agent.py</i>	M	Agent.__init__	A	3
<i>agent.py</i>	M	Agent.acceptFunction	A	3
<i>agent.py</i>	M	Agent.__repr__	A	1
<i>agent.py</i>	M	Agent.M	A	1
<i>agent.py</i>	M	Agent.strategySelection	A	1
<i>coalition.py</i>	C	Coalition	A	4
<i>coalition.py</i>	M	Coalition.bid	C	15
<i>coalition.py</i>	M	Coalition.target	B	7
<i>coalition.py</i>	M	Coalition.utility	A	5
<i>coalition.py</i>	M	Coalition.addMember	A	4
<i>coalition.py</i>	M	Coalition.delMember	A	4
<i>coalition.py</i>	M	Coalition.acceptance	A	4
<i>coalition.py</i>	M	Coalition.head	A	3

<i>coalition.py</i>	M	Coalition.__init__	A	2
<i>coalition.py</i>	M	Coalition.proposer	A	2
<i>coalition.py</i>	M	Coalition.uniqueID	A	2
<i>coalition.py</i>	M	Coalition.__repr__	A	1
<i>protocol.py</i>	F	majority	B	6
<i>protocol.py</i>	F	powerSelect	A	5
<i>protocol.py</i>	F	selectionFunction	A	1
<i>protocol.py</i>	F	randomSelect	A	1
<i>protocol.py</i>	F	votingRule	A	1
<i>round.py</i>	C	Round	B	7
<i>round.py</i>	M	Round.negotiate	C	15
<i>round.py</i>	M	Round.bid_zero	B	7
<i>round.py</i>	M	Round.myAgents	A	4
<i>round.py</i>	M	Round.ini_result	A	4
<i>round.py</i>	M	Round.__init__	A	3
<i>round.py</i>	M	Round.myCoalitions	A	3
<i>round.py</i>	M	Round.addCoalition	A	3
<i>round.py</i>	M	Round.generate_bid_zero	A	2
<i>round.py</i>	M	Round.T	A	1
<i>round.py</i>	M	Round.removeCoalition	A	1
<i>scenario.py</i>	C	Domain	A	1
<i>scenario.py</i>	C	Subject	A	1
<i>scenario.py</i>	C	Issue	A	1
<i>scenario.py</i>	M	Domain.__init__	A	1
<i>scenario.py</i>	M	Subject.__init__	A	1
<i>scenario.py</i>	M	Issue.__init__	A	1
<i>strategies.py</i>	F	time_dep	B	6
<i>strategies.py</i>	F	conv_bids	A	3
<i>strategies.py</i>	F	rand_tft	A	2
<i>strategies.py</i>	F	rel_tft	A	1
<i>strategies.py</i>	F	avg_tft	A	1
<i>model.py</i>	C	Model	B	9
<i>model.py</i>	M	Model.pre_coalitions	C	18
<i>model.py</i>	M	Model.draw_value	C	14
<i>model.py</i>	M	Model.simulate	C	13
<i>model.py</i>	M	Model.initialize	A	5
<i>model.py</i>	M	Model.dict_to_pandas	A	3

<i>model.py</i>	M	Model.__init__	A	1
<i>reporting.py</i>	F	updateCoalMatrix	C	14

Reliability

Thirdly, we discuss the reliability of our tool. We analyze the reliability by means of the $[E][F]$ outputs of the Pylint package and dive into external dependencies. First, let us discuss the Pylint results. In Table F.3 we do not observe any fatal errors (obviously), but we do see some important programming errors, both in the *model.py* and *round.py* modules. If we look more closely, the *round.py* module results four times in the same error: "Instance of 'list' has no 'keys' member (no-member)". This error is caused due to the fact that the instance – in this case *dummy_bid* – can be both a list and dictionary. The latter has the 'keys' member, whereas the first does not. As the instance is of type list, we skip the 'keys' member and thereby ignoring the error. Although we agree that this is slightly dodgy programming, it does not have any consequences for the model. In the *model.py* module, we see six times the same error: "Module 'numpy.random' has no 'normal' member (no-member)". This is a known and harmless error^[7], probably caused by conflicting versioning of different packages – in this case *numpy* and *random*. Thus, we can conclude that the errors discovered by Pylint in our tool do not affect the reliability model. Now, let us examine the external dependencies of the model, which are presented in Table F.8. Our tool depends on 4 external and 7 Python own packages. Apart from *bidict*, the external dependencies are not only widely used in Python's (data science) community, they are well established, mature and frequently updated. Therefore, we do not expect issues on short notice. In contrary, the *bidict* package is less mature as it is only founded since 2014. However, according to the documentation, the package is used by several teams at Google, Venmo and Bank of America Merrill Lynch. Of course, this does not guarantee stability in the near future for our tool, but it does give enough confidence that again we do not expect issues on short notice.

Efficiency

Finally, the efficiency is discussed, in which simulation performance is most crucial to achieve our first goal. As we expect that performance issues might occur in the more complex modules of our tool, (overnight) simulations have been run with varying numbers of agents (i.e. 3, 20 and 50), subjects (i.e. 1 and 20) and issues (i.e. 1 and 20). In total, 9 simulation experiments have been conducted with 1000 simulations each, thus in total 9000 model runs. The initialization of the model is similar to the example run in Section 5.5.3 whether or not extended with multiple issues, subjects or agents. The results of these runs are presented in

^[7]See, for example, <https://stackoverflow.com/questions/20553551/how-do-i-get-pylint-to-recognize-numpy-members>.

Table F.7: Output of Radon package indexing maintainability scores.

Module	Rank	Score
<i>agent.py</i>	A	52.9
<i>coalition.py</i>	A	59.6
<i>protocol.py</i>	A	71.5
<i>round.py</i>	A	50.0
<i>scenario.py</i>	A	99.9
<i>strategies.py</i>	A	58.2
<i>model.py</i>	A	53.9
<i>reporting.py</i>	A	86.9

Table F.8: External dependencies of our tool.

Dependency	Module	Created	Updated	Version
<i>bidict</i>	agent.py, round.py	≈ 2014	12/08/2017	0.14.2
<i>collections</i>	agent.py, coalition.py,	-	-	-
<i>copy</i>	round.py	-	-	-
<i>math</i>	strategies.py, model.py	-	-	-
<i>NumPy</i>	agent.py, coalition.py, round.py, strategies.py, model.py, reporting.py	≈ 2006	01/06/2018	1.6.1
<i>operator</i>	coalition.py	-	-	-
<i>pandas</i>	agent.py, round.py, scenario.py, model.py, reporting.py	≈ 2008	12/29/2017	0.22.0
<i>random</i>	agent.py, coalition.py, protocol.py, round.py, model.py	-	-	-
<i>SciPy</i>	model.py	≈ 2001	12/25/2017	1.0.0
<i>string</i>	coalition.py	-	-	-
<i>time</i>	model.py	-	-	-

Figure F.1^{[8][9]}. Overall, the results show a considerable spread in the simulation times of each experiment. In the simplest experiment – i.e. 1 subject, 1 issue and 3 agents – simulation times range from just above 0 to over 4 seconds, with the median around 0.6 seconds. Whereas the most time-consuming experiment – i.e. 1 subject, 20 issue and 50 agents – times range from 45 to over 260 seconds, with a median around 135 seconds. The spread in these results are explained by the time to reach an agreement. In some runs, an agreement is reached in just a few negotiation steps, whereas in others the overall deadline has surpassed without no agreement at all. In addition, we also observe that the number of issues a subject contains causes longer simulation times than negotiating over multiple subjects. This behavior was expected, as multiple issues leads to higher dimensional solution spaces, whereas multiple subjects lead to a series of lower dimensional solution spaces. Finally, we see that if we create twice as much agents in the negotiation environment, the simulation time roughly doubles. Again, this behavior is expected as the number of (potential!) interactions between agents are of $\mathcal{O}(n^2)$ ^[10]. Note that as agents make coalitions, the number of possible interactions decrease seriously^[11] and, thus, for higher numbers of n , we foresee the simulations to be of $\mathcal{O}(n \log n)$. Note too that the latter is an hypothesis based on simulation experiments shown in Figure F.1 and not proven theoretically. So, to conclude the efficiency experiments, for agents $n < 20$ and lower-dimensional issues, we are confident that our tool is feasible to answer the main research question(s). However, for $n \gg 20$ and/or (series of) high-dimensional solution spaces, the tool becomes infeasible as a single simulation run can take several minutes. The latter can somewhat hamper the extendability of the tool for more complex negotiation domains (third goal).

F.1.3. Assessment of non-functional requirements

The outcomes of the previous tests for the different high-level ISO-9126 quality standards allow us to assess the non-functional requirements as defined in D. The results of this assessment are presented in Table F.9.

^[8]The simulation experiments have been run on a laptop equipped with Intel(R) Core(TM) i5-6300U CPU 2.40GHz, 2496Mhz, 2 Cores processor and 16.0 GB physical memory (RAM).

^[9]Please be aware of the different scales on the y-axes.

^[10]The number of interactions between n agents can be determined by calculating the number of edges in a complete graph with n -vertices. Thus, the number of interactions between n agents equals $\frac{1}{2} \cdot n(n-1)$.

^[11]As coalitions are assumed to be a single entity.

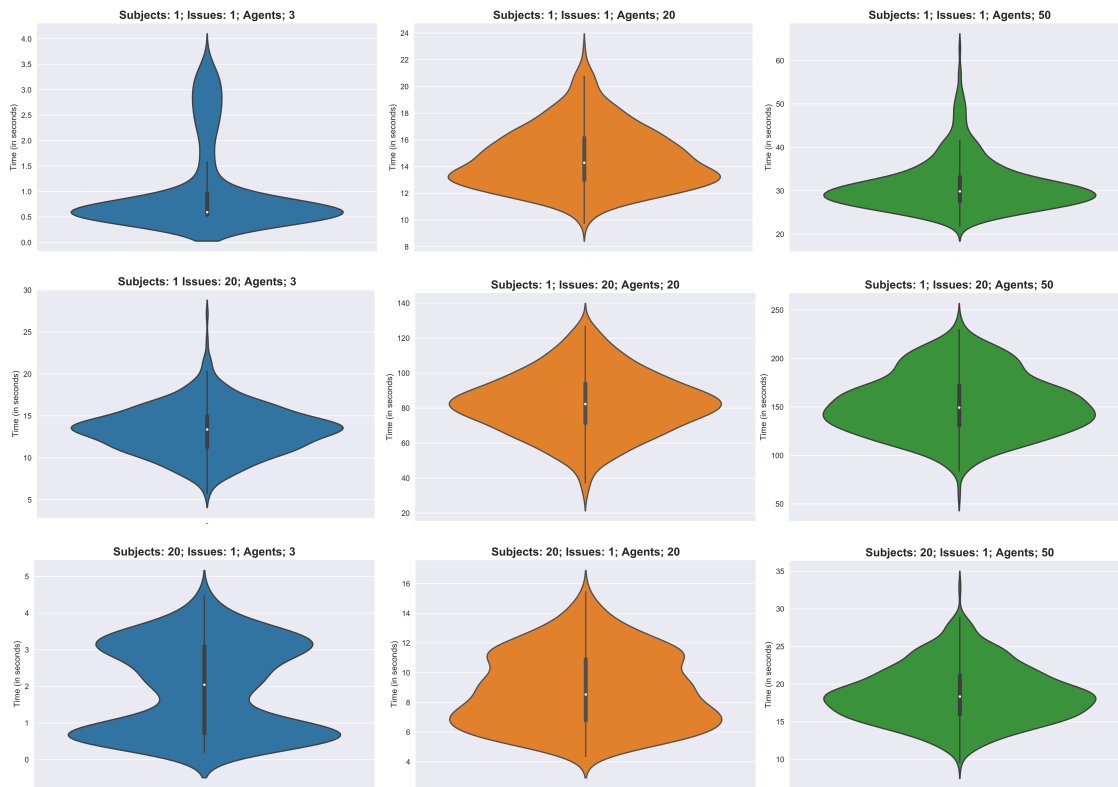


Figure F.1: Violin plots showing results of 9 simulation experiments examining tool performance.

As can be concluded from this table:

- All 2/2 "Must have" (M) non-functional requirements have been met;
- A considerable 3/5 "Should have" (S) functional requirements have been met;
- As expected, none or 0/1 of the "Won't have" (W) functional requirement have been met.

As our tool has met all the "Must have" (M) non-functional requirements and a considerable amount of "Should have" (S) and "Could have" (C), we are confident that its quality is adequate enough to meet our first goal and contributes to the realization of the second and third goal as well.

Table F.9: Non-functional requirement assessment.

#ID	Name	Priority	Assessment	Achieved
GL4	Open-source programming language	M	As described thoroughly in Section 5.5.1, Python is chosen as programming language.	Yes
GL5	Reliability	M	While examining the performance of our tool, we have run 9,000 simulations without any problems. This is just under 10,000 runs set as a prerequisite in Appendix D, but we are confident enough to state the tool will effortlessly achieve its reliability requirement. Moreover, efforts have been made to reduce tool complexity, as shown in the previous section.	Yes
GL8	Dependencies	S	We have shown that there are only four external dependencies in our tool. Those include dependencies with mature and stable packages that most likely not hamper the (long-term) reliability of the tool.	Yes

GL9	Adhere to style guide principles of programming community	S	As discussed in the usability and maintainability paragraphs in Section F.1.2, our tool adheres as much as possible to PEP8 guidelines. However, there is still (much) room for improvement, especially regarding name conventions.	No
GL10	Documentation	S	As shown in Table F.5, we have documented all modules, classes, methods and function conforming to Python's PEP8 guidelines.	Yes
GL11	Performance	S	As shown in Figure F.1, this requirement is achieved for simple, low-dimensional negotiation domains (with median runtime of 0.6 seconds). However, for more complex domains, the performance requirement is not met.	No
GL16	User interface	W	As explained, this is not supported by our tool.	No
SC7	New scenarios	S	Inserting negotiation scenarios for 9 simulation experiments in Figure F.1 took close to 65 minutes. This is roughly 7.2 minutes for each scenario. Although there is considerable overlap between those scenarios, they are still relatively complex and contain lots input parameters. Thus, we are confident that similarly complex (stand-alone) scenarios will not take more than 15 minutes to prepare.	Yes

F.1.4. Conclusion of tool quality

In the previous two subquestions, we have shown that the tool developed in this thesis meets all the necessary functional and non-functional requirements set in Appendix D to meet our definition of quality and (to some extent) support the three major development goals. As all *functional* requirements have been met, the tool should allow us to answer the research question(s) in this thesis. In addition, the usability and maintainability standards are in our view appropriate enough to allow others to review our work. However, there still exist areas for continued development, especially regarding extensibility for even more complex negotiation domains and necessary simulations.

F.2. Verification: have we computed the correct model?

Testing the quality of our tool does not tell the complete story. We still need to make sure if the requirements that shape the conceptual model presented in Chapter 5 are correctly translated to the computational tool in Python. This step is called verification – or more specifically *implementation verification* – and defined as "...assuring that the computer programming and implementation of the conceptual model is correct" [151, p. 108]. The following tests have been designed to verify our model:

- First, for every (class) method and function in our tool, we examine the input and output and verify whether or not it matches theoretical predictions from the conceptual model;
- Secondly, we setup a simple three-agent model for which we make theoretical predictions based upon the conceptual model. These predictions are then verified by comparing them to simulated results;
- Thirdly, we push our tool to the limit by entering extreme parameters for the simple three-agent model;
- Moreover, we repeat both aforementioned tests for a more complex, eight-agent model;
- Finally, variability test have been performed to explain and verify unintended randomness.

The model setup is again similar to the example run in Chapter 5, whether or not extended with multiple agents, subjects or issues. Note that if multiple issues are entered, we always make sure that at least 1 discrete, 1 unlimited continuous and 1 limited continuous issue is present. Note that both the first and second test are run in a simple three-agent model. In contrary, the full model contains 10 agents. Although our model is capable of handling more agents (see previous section), for the sake of time "only" 10 agents

are inserted. In addition, for sake of clarity, we will label the agents a with subscripts (i.e. from a_1 to a_{10}) the subjects with capitals and the issues with capitals followed numbers – thus, if subject A contains three issues, they are labeled $A1$, $A2$ and $A3$.

F.2.1. Examining inputs and outputs

The output of the *Radon* package in Table F.6 gives a clear oversight of the methods (M) and functions (F) in the different modules in our tool. Although the majority of functions or methods lead to outputs, not all of them do so. In some cases, they change class attributes for example. The latter are not taken into account for this verification test, as they only *use* (contrary to *define*) the output of functions who do have outputs. So, for this analysis, the following functions and methods are excluded: *Agent.M*, *Round.T* *Round.negotiate*, *Coalition.addMember*, *Coalition.delMember*, *Coalition.head*, *Coalition.proposer*, *Coalition.target* and *Coalition.uniqueID* Moreover, in this part we verify the two-stage model, as that model is defined conceptually and implemented computationally. Most in- and output tests are realized by adding "print" statements and observing the results. Some are realized by creating plots. Note that for sake of simplicity, we have described only the relevant input parameters. The results are discussed below.

agent.py – opp_bids

<i>Description</i>	This function aggregated and averages coalitional bids of different moments in time.
<i>Input</i>	Equally powered agents a_1 and a_2 , historic bids [30,34,36,36] and [32,34,37,37] at times $t = [1,2,4,6]$ and $t = [1,2,3,5]$.
<i>Prediction</i>	Should average times $t = [1,2]$ and aggregates other time steps, thus the bids [31,34,37,36,37,36].
<i>Output</i>	The function output is Pandas Time Series with values [31,34,37,36,37,36] at times (indexes) $t = [1,2,3,4,5,6]$.
<i>Note</i>	For discrete (string) bids, the historic bids are voted for by power.
<i>Verified</i>	Yes

agent.py – conv_bids

<i>Description</i>	Converts non-integer opponent's bids for specific issues into quantifiable bids according preference profiles.
<i>Input</i>	Discrete issue $A1$ with options $A1 = \{opt1,opt2,opt3\}$, single agent a_1 with preferences {5,7,11}, historic bids [opt3,opt3,opt2,opt3,opt2] at times $t = [1,3,4,10,14]$.
<i>Prediction</i>	Should be time series with values [11,11,7,11,7].
<i>Output</i>	The function output is Pandas Time Series with values [11,11,7,11,7] at times (indexes) $t = [1,3,4,10,14]$.
<i>Verified</i>	Yes

agent.py – Agent.bid and generate_bid

<i>Description</i>	Generate bid based upon tactics specified in an agent's strategy profile(s).
<i>Input</i>	Agent a_1 with strategy profiles $\{[1,0,0,0], [1,2,0,0], [1,1,1,0], [1,1,1,1], [1,2,1,2]\}$ for 'initial', 'inverter', 'conceder', 'matcher' and 'competitor' respectively. Agent a_2 is perceived by a_1 as 'inverter'. Tactics in strategy profile result in 15.9, 17.3, 19.4 and 15.0 respectively.
<i>Prediction</i>	According to Section 4.3: $x_{a_1 \rightarrow a_2}^t = \omega_{j_1} \tau_{j_1}(t)[j] + \omega_{j_2} \tau_{j_2}(t)[j] + \dots + \omega_{j_x} \tau_{j_x}(t)[j] = \frac{1}{1+2+0+0} \cdot 15.9 + \frac{2}{1+2+0+0} \cdot 17.3 + \frac{0}{1+2+0+0} \cdot 19.4 + \frac{0}{1+2+0+0} \cdot 15.0 \approx 16,8$
<i>Output</i>	Output is a float with value 16,833333.
<i>Note</i>	Agent.bid is an extension of generate_bid and, therefore, verified in the same manner. It determines the strategy, rounds bids or converts them to string bids. Be aware that entries in strategy profile correspond to tactics in Section 4.3.1 respectively.
<i>Verified</i>	Yes

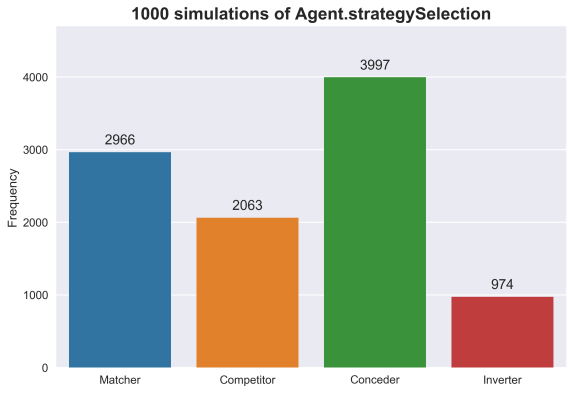
agent.py – Agent.utility

<i>Description</i>	Determines the utility of an agent given the bids for a subject (whether or not containing multiple issues).
<i>Input</i>	Subject A with two different types of issues (unlimited continuous and discrete). Agent a_1 with preference profiles $A1 = [90, \mathbf{140}]$ and $A2 = [20, \mathbf{40}, 10]$ (for $A2 = [opt1, opt2, opt3]$), Weights $[2,5]$, aiming for high bids each and opponent bid $[110, opt1]$.
<i>Prediction</i>	As utility is determined by the Euclidean distance, independently weighted for both issues and added with an error term (see Section 5.3.4), we predict the (standardized) utility to be: $\frac{2}{2+5} \cdot \frac{110-90}{140-90} + \frac{5}{2+5} \cdot \frac{20-10}{40-20} \approx 0,47$.
<i>Output</i>	Output is a float with value 0,54.
<i>Note</i>	The difference can be explained by the addition of the error term $\epsilon \sim N(0,0.1)$.
<i>Verified</i>	Yes

agent.py – Agent.acceptFunction

<i>Description</i>	Determines whether or not a bid is accepted, based upon an agent's utility (see above).
<i>Input</i>	Utilities from own and opponent bids. In this verification, $u_{own} = 0.87$ and $u_{opp} = 0.74$.
<i>Prediction</i>	As $u_{opp} < u_{own}$, the acceptance function must return <i>False</i> . See also Section 4.3.3.
<i>Output</i>	The method returns <i>False</i> .
<i>Verified</i>	Yes

agent.py – Agent.strategySelection

<i>Description</i>	Selects a strategy from an agent's strategy profile based upon opponent's historic bids and strategy selection function.										
<i>Input</i>	Historic bids result in opponent model strategy 'conceder' (see next test). Agent a_1 strategy selection function for 'conceder' is uniform distribution over different profiles with probabilities $P_{inv} = 0.1$, $P_{con} = 0.4$, $P_{com} = 0.2$ and $P_{mat} = 0.3$ (see Section 5.3.3).										
<i>Prediction</i>	If we run the strategy selection function 10,000 times, we expect it selects 'matcher' approximately 3000 times, 'competitor' 2000, 'conceder' 4000 and 'matcher' 1000.										
<i>Output</i>	The results in the simulation experiment below are close to what was expected.  <table border="1" style="display: none;"> <caption>1000 simulations of Agent.strategySelection</caption> <thead> <tr> <th>Strategy</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>Matcher</td> <td>2966</td> </tr> <tr> <td>Competitor</td> <td>2063</td> </tr> <tr> <td>Conceder</td> <td>3997</td> </tr> <tr> <td>Inverter</td> <td>974</td> </tr> </tbody> </table>	Strategy	Frequency	Matcher	2966	Competitor	2063	Conceder	3997	Inverter	974
Strategy	Frequency										
Matcher	2966										
Competitor	2063										
Conceder	3997										
Inverter	974										
<i>Note</i>	Note that in practice, this function returns the strategy profile and risk profile based upon the above output!										
<i>Verified</i>	Yes										

agent.py – Agent.opponentModel

<i>Description</i>	Tries to determine opponent strategy profile by fitting LSE regression model.
<i>Input</i>	Prior strategy profiles for $a_2 = \{[1, 2, 0, 0], [1, 1, 1, 0], [1, 1, 1, 1], [1, 2, 1, 2]\}$ for 'inverter', 'conceder', 'matcher' and 'competitor' respectively; a_2 's historic bids x_H for issue A1: $x_H = [20, 22, 25, 30, 40]$.
<i>Throughput</i>	Based upon prior for a_2 's strategy profiles, the following bids \tilde{x}_i are simulated: $x_{inv} = [21, 24, 25, 28, 42]$; $x_{con} = [16, 23, 23, 25, 38]$; $x_{mat} = [16, 23, 23, 26, 37]$ and $x_{com} = [24, 26, 28, 35, 34]$
<i>Prediction</i>	By applying Least Square Estimation (LSE), we try to minimize the sum of squares of residuals (SS) between historic bids and simulated bids in of prior strategy profile. Thus, $LSE = \min(SS_{inv}, SS_{con}, SS_{mat}, SS_{com}) = \min(13, 50, 46, 102) = 13$ which corresponds to opponent a_2 strategy of 'inverter'.
<i>Output</i>	The method returns a string with 'inverter'.
<i>Note</i>	In practice, this methods runs multiple simulations before determining the LSE. Demonstration above shows only one of such simulation run.
<i>Verified</i>	Yes

coalition.py – Coalition.bid

<i>Description</i>	Coalition bids are combinations of individual member agent bids (see before).
<i>Input</i>	Opponent a_3 , coalition members a_1 and a_2 with power $p_{a_1} = 3$ and $p_{a_2} = 10$, and continuous bids $x_{a_1 \rightarrow a_3} = 25.7$ and $x_{a_2 \rightarrow a_3} = 31.2$ respectively.
<i>Prediction</i>	If <i>weighted</i> is <i>True</i> , coalition bids are weighted combinations of individual member bids. If <i>False</i> , they are not (see Section 5.3.3). So we expect either a coalitional bid of $x = \frac{3}{3+10} \cdot 25.7 + \frac{10}{3+10} \cdot 31.2 \approx 28,45$ or of $x = \frac{1}{2} \cdot 25.7 + \frac{1}{2} \cdot 31.2 = 26,97$.
<i>Output</i>	The final output of this method is an integer valuing 29, so it meets our expectations for weighted bids.
<i>Note</i>	The output is rounded (and thus integers) as we assumed only integer bids are possible. Similar to <code>Agent.opp_bids</code> , if <i>weighted</i> , than for discrete (string) bids we vote for power.
<i>Verified</i>	Yes

coalition.py – Coalition.utility

<i>Description</i>	Coalitional utility is calculated by a (weighted) average of the utility of its member agents.
<i>Input</i>	Member agents a_1 and a_2 receive an offer from a_3 . Utilities derived from this offer are $u_{a_1} \approx 0.59$ and $u_{a_2} \approx 0.46$ respectively. Again, they have both $p_{a_1} = 3$ and $p_{a_2} = 10$.
<i>Prediction</i>	As coalitional utility U is simply (weighted) combination of individual member agent utilities (see Section 5.3.4), we expect either $U = \frac{3}{3+10} \cdot 0.59 + \frac{10}{3+10} \cdot 0.46 \approx 0,49$ or $U = \frac{1}{2} \cdot 0.59 + \frac{1}{2} \cdot 0.46 = 0.53$.
<i>Output</i>	The final output of this method is a float valuing 0.53. So, it meets the theoretical prediction from the conceptual model.
<i>Verified</i>	Yes


coalition.py – Coalition.acceptance

<i>Description</i>	Except new members in coalition. If accepted, every coalition member in the new member coalition decides for himself if the newly formed coalition is more beneficial than the old one (and thus if the individual rationality requirement is satisfied; see Section 3.2.1).						
<i>Input</i>	Coalitional utility from own (a_1 and a_2) and opponent (a_3) offer: $U_{own} = 0.51$ and $U_{opp} = 0.53$ respectively. Agent rationality behavior are 0.2, 0.3 and 0.1 for agent a_1 , a_2 and a_3 respectively.						
<i>Throughput</i>	The individual utilities are respectively $u_{a_1} = 0.59$, $u_{a_2} = 0.46$ and $u_{a_3} = 0.87$.						
<i>Prediction</i>	As $U_{own} = 0.51 < U_{opp} = 0.53$ we expect that the coalition accepts a_3 in his new coalition. However, $u_{a_2} < U_{own}$ and therefore, agent a_2 might leave the coalition. This should happen with probability 0.3.						
<i>Output</i>	10,000 simulations show that a_2 leaves the new coalition close to 3,000 times. <div style="text-align: center;"> <table border="1" style="margin: 10px auto;"> <caption>10000 simulations of Coalition.acceptance</caption> <thead> <tr> <th>Agent a_2 leaves coalition</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>False</td> <td>7074</td> </tr> <tr> <td>True</td> <td>2926</td> </tr> </tbody> </table> </div>	Agent a_2 leaves coalition	Frequency	False	7074	True	2926
Agent a_2 leaves coalition	Frequency						
False	7074						
True	2926						
<i>Note</i>	Of course, this method not only evaluates if a_2 leaves the new coalition, but does this for all members. However, the other members will not leave the new coalition as the new coalition is more beneficial.						
<i>Verified</i>	Yes						

protocol.py – votingRule and majority

<i>Description</i>	Function determining whether or not a large enough majority of votes is obtained.
<i>Input</i>	Total number of agents in environment ($N = 3$), a_1 and a_3 in winning coalition ($n = 2$), total power in environment $P = 20$ and total power in winning coalition $p = p_{a_1} + p_{a_3} = 3 + 7 = 10$.
<i>Prediction</i>	If <i>simple majority</i> is selected, this function should return <i>True</i> as $\frac{n}{N} = \frac{2}{3} > 0.5$. If <i>weighted majority</i> is selected, the function should return <i>False</i> as $\frac{p}{P} = \frac{10}{20} \not> 0.5$. For any (<i>weighed</i>) <i>super majority</i> greater than 50%, it should also return <i>False</i> (see Section 2.2.4).
<i>Output</i>	The function returns a boolean <i>True</i> if we select 'SiM' (= simple majority), return <i>False</i> if we select 'WM' or 'U' (weighted majority and unanimity respectively). Moreover, for weighted majority, it returns <i>True</i> if the majority rule is < 0.50 .
<i>Note</i>	As 'votingRule' is an extension of 'majority' it is verified in the same manner.
<i>Verified</i>	Yes

protocol.py – selectionFunction, randomSelect and powerSelect

<i>Description</i>	Selects function for (next) agent/coalition based upon randomness or power.
<i>Input</i>	Coalition $\{a_1, a_2\}$ and agent a_3 with (aggregated) powers $p_{col} = p_{a_1} + p_{a_2} = 3 + 10 = 13$ and $p_{a_3} = 7$ respectively.
<i>Prediction</i>	If random selection is preferred both coalition $\{a_1, a_2\}$ and a_3 should be selected equally often. If not, we expect the coalition to be selected with probability $\frac{13}{20}$ and the agent with probability $\frac{7}{20}$. Thus, with 10,000 simulation experiments, roughly 6,500 and 3,500 times respectively (see Section 5.3.2).
<i>Output</i>	Results below meet our expectations for power and random-based selections. 
<i>Verified</i>	Yes

round.py – Round.myAgents


<i>Description</i>	Initiates agents in the two-stage negotiation environment.
<i>Input</i>	Agents a_1 , a_2 and a_3 with their attributes as defined in Section 5.4. Thus P , B , R , Δ and D respectively.
<i>Prediction</i>	Should return dictionary containing keys $[a_1, a_2, a_3]$ and instances of Agent classes as values.
<i>Output</i>	Results in the following dictionary: $\{ 'a_1': \text{Agent}('a_1'), 'a_2': \text{Agent}('a_2'), 'a_3': \text{Agent}('a_3') \}$. The latter aligns with our expectations.
<i>Verified</i>	Yes

round.py – Round.myCoalitions

<i>Description</i>	Function transforming list of coalition members to coalitions of class Coalition with accompanying members.
<i>Input</i>	List of lists of prior coalitions (as result of pre-negotiations; see later on) is the following $[[a_1, a_2], [a_3]]$.
<i>Prediction</i>	Should return dictionary containing prior coalitions with unique keys and instances of Coalition class as values.
<i>Output</i>	Results in the following dictionary: $\{ 'A': \text{Coalition}('A'), 'B': \text{Coalition}('B') \}$ in which $\text{Coalition}('A').members$ returns a list with two instances of Agent class, namely $[\text{Agent}('a_1'), \text{Agent}('a_2')]$.
<i>Verified</i>	Yes


round.py – Round.generate_bid_zero and bid_zero

<i>Description</i>	Generates initial bids to start negotiations.
<i>Input</i>	Agent a_1 with the following preference profile $A_1 = [30, 100]$, thus aiming for high values. Moreover, $\sigma = 0.05$.
<i>Prediction</i>	Depending on whether or not agents aim for low or high values – i.e. utility increases/decreases with a larger part of the "pie" (see Section 4.2) – we assumed initial bids x_0 are "drawn" from a normal distribution, thus $x_0 \sim N(\mu, \sigma \cdot \mu)$. As the mean should represent target value $\mu = 100$, we expect $x_0 \sim N(100, 0.05 \cdot 100)$.
<i>Output</i>	10,000 simulation experiments of this method meet our expectations (see below)
<i>Verified</i>	Yes



strategies.py – time_dep

<i>Description</i>	Function describing time-dependent tactic [45, p. 14-15] and returning bids.
<i>Input</i>	Agent a_1 with similar preference profile $A_1 = [30, 100]$ (but now aiming for low values), private deadline $\delta_1 = 30$ and risk behavior $\beta_1 = [0.1, 0.5, 1, 5]$ for four different strategy profiles respectively and initial bid $x_0 = 31$.
<i>Prediction</i>	For polynomial tactics, the results should be similar to Figure 4.3 (left).
<i>Output</i>	The simulation is presented below and follows the theoretical assumptions.
<i>Verified</i>	Yes



strategies.py – rel_tft

<i>Description</i>	Relative tit-for-tat tactic as one of three behavior dependent tactics [45, p. 15-17].
<i>Input</i>	Inserted are agents a_1 and a_2 with historic bids till t_n equal to $x_{a_1 \rightarrow a_2} = [108, 108]$ and $x_{a_2 \rightarrow a_1} = [96, 88, 90]$. respectively. a_1 employs relative tit-for-tat tactic by looking one time step back ($w = 1$) with preference profile for issue in range [26, 112].
<i>Prediction</i>	From Equation (4.1) we know: $x_{a_1 \rightarrow a_2}^{t_{n+1}} = \min(\max(\frac{x_{a_2 \rightarrow a_1}^{t_{n-2w}}}{x_{a_2 \rightarrow a_1}^{t_{n-2w+2}}} x_{a_1 \rightarrow a_2}^{t_{n-1}}, \min^{a_1}), \max^{a_1})$ $= \min(\max(\frac{96}{90} \cdot 108, 26), 112)$ $= \min(\max(115.2, 26), 112) = 112$
<i>Output</i>	The function returns an integer valuing 112.
<i>Verified</i>	Yes

strategies.py – rand_tft

<i>Description</i>	Random tit-for-tat tactic as one of three behavior dependent tactics [45, p. 15-17].
<i>Input</i>	Same setup as previous experiment, only adding capriciousness parameter $M_{a_1} = 1$.
<i>Prediction</i>	From Equation (4.2) we know: $x_{a_1 \rightarrow a_2}^{t_{n+1}} = \min(\max(x_{a_1 \rightarrow a_2}^{t_{n-1}} + (x_{a_2 \rightarrow a_1}^{t_{n-2w}} - x_{a_2 \rightarrow a_1}^{t_{n-2w+2}}) + (-1)^s R(M), \min^{a_1}), \max^{a_1})$ $= \min(\max(108 + (96 - 90) + (-1)^0 \cdot R(1), 26), 112)$ $\leq \min(\max(108 + 6 + 1, 26), 112)$ $= \min(\max(115, 26), 112) = 112$
<i>Output</i>	The function returns an integer valuing 112.
<i>Verified</i>	Yes

strategies.py – avg_tft

<i>Description</i>	Average tit-for-tat tactic as one of three behavior dependent tactics [45, p. 15-17].
<i>Input</i>	Same setup as previous experiment, only changing step parameter with window $v = 1$.
<i>Prediction</i>	From Equation (4.3) we know: $x_{a_1 \rightarrow a_2}^{t_{n+1}} = \min(\max(\frac{x_{a_2 \rightarrow a_1}^{t_{n-2v}} [j]}{x_{a_2 \rightarrow a_1}^{t_n} [j]} x_{a_1 \rightarrow a_2}^{t_{n-1}} [j], \min_j^{a_1}), \max_j^{a_1})$ $= \min(\max(\frac{90 + 90 + 90}{96 + 88 + 90} \cdot 108, 26), 112)$ $\leq \min(\max(\frac{270}{274} \cdot 108, 26), 112)$ $= \min(\max(106.4, 26), 112) \approx 106.4$
<i>Output</i>	The function returns a float valuing approximately 106.4.
<i>Verified</i>	Yes

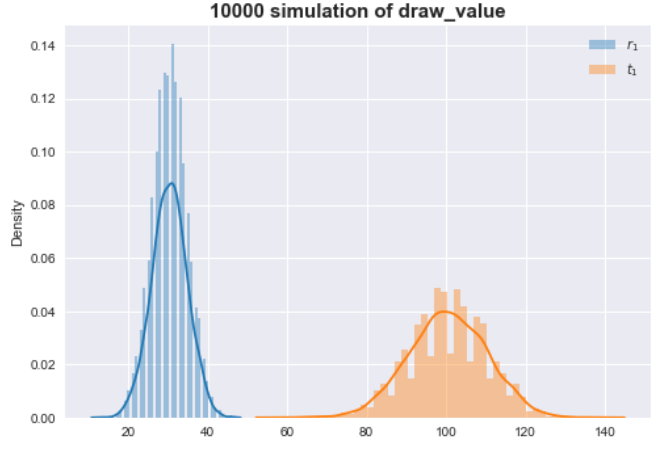
strategies.py – conv_bids

<i>Description</i>	Converts non-integer opponent's bids for specific issues into quantifiable bids according to an agent's own preference profile.
<i>Input</i>	Entered are agent a_1 historic (string) bids $x_H = [opt1, opt1, opt3, opt4]$ at different moments in time. Moreover, a_1 's preference profile for this issue is described by $\{10, 5, 8, 6\}$ for $A1 = \{opt1, opt2, opt3, opt4\}$ respectively.
<i>Prediction</i>	The function should result in a list containing the following values: $[10, 10, 8, 6]$.
<i>Output</i>	The function returns $[10, 10, 8, 6]$, which meets our expectations.
<i>Verified</i>	Yes

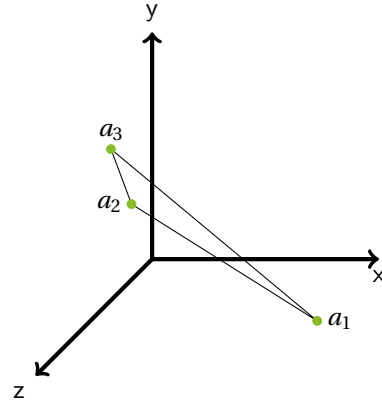
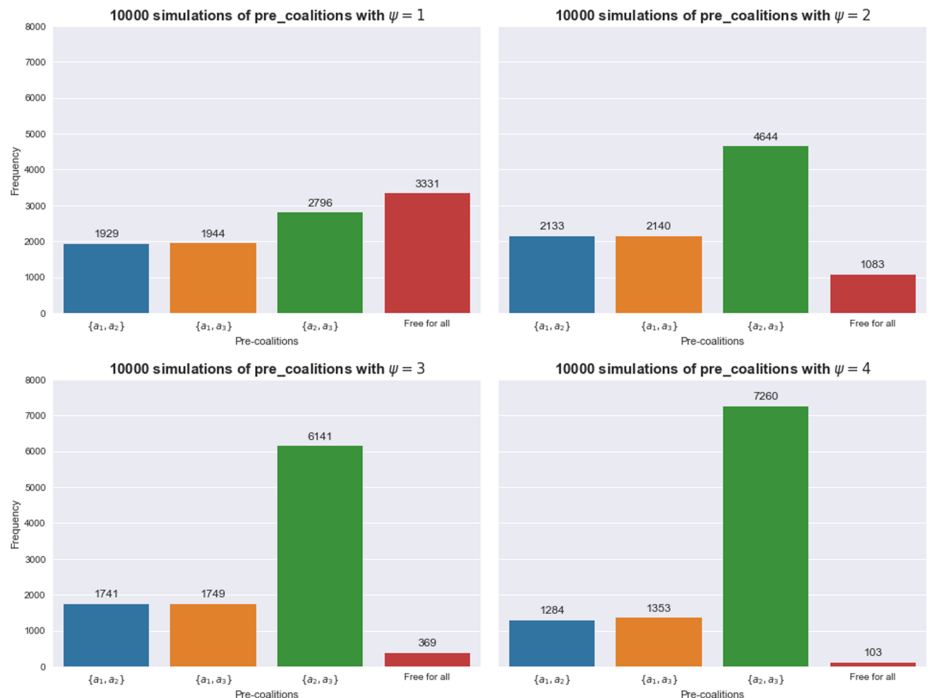
model.py – Model.initialize

<i>Description</i>	Initializes simulation run by initiating new negotiation rounds.
<i>Input</i>	All input parameters of the two-stage negotiation model.
<i>Prediction</i>	Should return instance of class Round, containing agents, scenario and pre-coalitions.
<i>Output</i>	Returns rnd , an instance of class Round.
<i>Verified</i>	Yes

model.py – Model.draw_value

<i>Description</i>	Draws values for uncertain scenario input parameters from normal distribution in order to generate multiple instances of (uncertain) negotiation scenarios.
<i>Input</i>	Uncertainty factor $\sigma = 0.10$ and agent a_1 's preference profile $[30, 100]$ for issue $A1$.
<i>Prediction</i>	Both a_1 's reservation value $r_1 = 30$ as well as target value $t_1 = 100$ should be drawn from a normal distribution in which $\tilde{r}_1 \sim N(r_1, \sigma \cdot r_1) = N(30, 3)$ and $\tilde{t}_1 \sim N(t_1, \sigma \cdot t_1) = N(100, 10)$, which results in preference profile $\tilde{A}_1 = [\tilde{r}_1, \tilde{t}_1]$ for single simulation run (see Section 5.4). Note that the results will always be non-negative.
<i>Output</i>	<p>The figure below shows 10,000 simulation experiments for \tilde{r}_1 (blue) and \tilde{t}_1 (orange).</p> 
<i>Note</i>	All agents' preferences, characteristics and strategy profiles are considered to be uncertain (see again Section 5.4).
<i>Verified</i>	Yes

model.py – Model.pre_coalitions

<p><i>Description</i></p>	<p>Generates possible pre-coalitions based upon ideological/intrinsic agent characteristics.</p>																									
<p><i>Input</i></p>	<p>Agents a_1, a_2 and a_3 with three-dimensional ideological/intrinsic values $\lambda_1 = [3,0,3]$, $\lambda_2 = [0,1,1]$ and $\lambda_3 = [0,2,2]$ respectively. For the verification experiments, we choose $\psi \in \{1,2,3,4\}$ Note that no we do not insert a target coalition C_T.</p>																									
<p><i>Throughput</i></p>	<p>If we plot the data, we observe that a_2 and a_3 are closer together than a_1.</p> 																									
<p><i>Prediction</i></p>	<p>Based upon the above plot, we would expect that a_2 and a_3 more often form a pre-coalition, especially if we increase the number of tries ψ.</p>																									
<p><i>Output</i></p>	<p>The figure below shows 10,000 simulation experiments for $\psi \in \{1,2,3,4\}$ each. We indeed see that coalition $\{a_2, a_3\}$ is more likely to form, especially for higher values of ψ</p>  <table border="1"> <caption>Frequency of Pre-coalitions for different values of ψ</caption> <thead> <tr> <th>ψ</th> <th>$\{a_1, a_2\}$</th> <th>$\{a_1, a_3\}$</th> <th>$\{a_2, a_3\}$</th> <th>Free for all</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1929</td> <td>1944</td> <td>2796</td> <td>3331</td> </tr> <tr> <td>2</td> <td>2133</td> <td>2140</td> <td>4644</td> <td>1083</td> </tr> <tr> <td>3</td> <td>1741</td> <td>1749</td> <td>6141</td> <td>369</td> </tr> <tr> <td>4</td> <td>1284</td> <td>1353</td> <td>7260</td> <td>103</td> </tr> </tbody> </table>	ψ	$\{a_1, a_2\}$	$\{a_1, a_3\}$	$\{a_2, a_3\}$	Free for all	1	1929	1944	2796	3331	2	2133	2140	4644	1083	3	1741	1749	6141	369	4	1284	1353	7260	103
ψ	$\{a_1, a_2\}$	$\{a_1, a_3\}$	$\{a_2, a_3\}$	Free for all																						
1	1929	1944	2796	3331																						
2	2133	2140	4644	1083																						
3	1741	1749	6141	369																						
4	1284	1353	7260	103																						
<p><i>Verified</i></p>	<p>Yes</p>																									

Conclusion – The above verification tests show that the outputs for all functions and methods in our tool matches the predicted theoretical output of the conceptual model, given certain input parameters. Thus, for all individual functions and methods in the conceptual model it is verified that they are correctly implemented in the computerized model. However, this still does not imply that altogether our tool is implemented correctly. Therefore, the coming sections will discuss and test the entire tool.

F.2.2. Theoretical prediction: simple three-agent model

In this section, we first setup a simple three-agent model. Then, based on the input parameters, we predict what most likely will happen – in terms of agent behavior – given the conceptual model discussed in Chapter 5. Finally, we simulate the simple model in our tool and compare these practical results with our theoretical presumptions. The agent setup and negotiation scenario are described in Table F.10.

Table F.10: Negotiation scenario setup for the simple three-agent model.

Component	Parameter	Value	Component	Parameter	Value
<i>Domain</i>	Deadline Θ	30	Agent a_2	Preference $A1$	[20, 50]
<i>Subject A</i>	Issue $A1$	[30, 100]		Preference $A2$	[15 , 10, 5]
	Issue $A2$	{'high', 'medium', 'low'}		Preference $B1$	{ a_1 : [110, 150], a_2 : [40, 50], a_3 : [40, 70]}
	Deadline θ_1	20		Weights w_2	[1, 1, 1]
<i>Subject B</i>	Issue $B1$	270		Power p_2	1
	Deadline θ_1	20		Ideology λ_2	[5, 5]
Agent a_1	Preference $A1$	[50, 80]	Agent a_3	Preference $A1$	[30 , 70]
	Preference $A2$	[10, 12, 18]		Preference $A2$	[10 , 12, 18]
	Preference $B1$	{ a_1 : [110, 150], a_2 : [30, 50], a_3 : [40, 70]}		Preference $B1$	{ a_1 : [70, 100], a_2 : [10, 30], a_3 : [80, 120]}
	Weights w_1	[1, 1, 1]		Weights w_3	[1, 1, 1]
	Power p_1	1		Power p_3	1
	Ideology λ_1	[4, 4]		Ideology λ_3	[0, 0]
Coalition	C_T	<i>None</i>			

Input parameters not defined in this table are resumed constant for all agents and, partly therefore, will not affect this verification test. This implies that all agents have the same strategy profiles, risk behavior and act completely rational upon each other. Moreover, negotiations will end if *weighted majority* is reached^[12]. We will simulate the three-agent model 100 times with uncertainty factor $\sigma = 0.10$.

Prediction – As a_2 and a_3 have similar preferences for subject A , we expect them to be in the winning coalition for this subject the majority of simulation runs. Moreover, for subject B , expect a_1 and a_2 to form winning coalition more often as they have similar preferences of how the pie should be divided. Regarding negotiation outcomes, we expect $A1$ to be in between a_2 and a_3 target values for the majority of simulations. Similarly, regarding $A2$, we predict the option 'high' to be selected more often, before 'medium' and 'low'. As we expect a_2 to be in the winning coalitions more often, we predict it's overall utility to be the highest. Similarly, we predict a_1 's utility to be the lowest, as they probably only win 1/3 issues. Finally, as in both negotiation rounds – one for each subject – there will be dominant winning coalitions, we predict that the majority of simulations will reach in an (overall) agreement. The aforementioned predictions are summarized in Table F.11 as well as the results after simulations.

^[12]Note that for the tests in this section there is no practical difference in choosing *simple majority* or *weighted majority* as $p = 1$ for all agents. However, this will be different for the tests in section Section F.2.3.

Table F.11: Predictions and results of simulation experiments for three-agent model.

Output parameter	Prediction	Simulation results	Verified
Agreement	> 90%	≈ 84%	No
Winning coalition <i>A</i>	$a_1 < 50\%$, $a_2 > 75\%$ and $a_3 > 75\%$	$a_1 \approx 22\%$, $a_2 \approx 100\%$ and $a_3 \approx 78\%$	Yes
Winning coalition <i>B</i>	$a_1 > 50\%$, $a_2 > 50\%$ and $a_3 < 50\%$	$a_1 \approx 66\%$, $a_2 \approx 88\%$ and $a_3 \approx 35\%$	Yes
Outcome <i>A1</i>	≈ [30, 50]	75% of outcomes are in [42, 52]	No
Outcome <i>A2</i>	'high' > 'medium' > 'low'	'high' = 60, 'medium' = 24 and 'low' = 0	Yes
Utility (median)	$u_{a_2} > u_{a_3} > u_{a_1}$	$u_{a_2} \approx 0.87$, $u_{a_3} \approx 0.6$ and $u_{a_1} \approx 0.37$	Yes

Conclusions – The results in Table F.11 show that the majority of simulations match the theoretical predictions derived from our conceptual model. In our opinion, the two dissimilarities that have occurred stem from uncertainty in the input parameters and randomness in agents' strategies. Even though the three agents have similar strategy profiles, the way they react upon each other still depends on stochastic behavior. Thus, to conclude, the model has *passed* this verification test.

F.2.3. Pushing the tool to the limit: simple three-agent model

We will now push the three-agent model to the limit, by changing remaining model parameters – that previously were assumed to be constant and/or non-affecting – to extreme values. If the selected model parameter is an agent's attribute, it is only changed for a_1 . Each model parameter is selected one at a time, accompanied with a brief prediction and followed by obtained result. This is presented in Table F.12. Note that for these extreme tests, we only perform 10 simulations.

Conclusion – As can be seen in Table F.12, all entered extreme values – either minimum or maximum – behaved as expected, except for entering the grand coalitions as target coalition C_T . Instead of running the simulation rounds, the tool returns an error. However, we can ignore this error, as it is both in theory as well as in practice not relevant to enter the grand coalition as target coalition C_T . To be specific, if in practice we would be able to form the grand coalition – containing all involved negotiation parties – it would be pointless to negotiate afterwards as all of them have agreed to work together in the first place. So, in practice, there is nothing to negotiate about any more. Thus, to conclude, the model has *passed* this verification test. However, for future work, we recommend to resolve this issue as it is not in line with other outputs.

F.2.4. Theoretical prediction: multi-agent model

This sections extends the simple three-agent model introduced in Section F.2.2 by adding 5 more agents in the negotiation environment. We will use the same negotiation domain, thus containing of two subjects *A* and *B* containing 2 and 1 issue respectively. Also the same outcome spaces are used, so, in the following verification experiments, the total pie in *B1* should now be divided by 8 agents, instead of 3. Table F.14 shows the attributes and preferences that have been assigned to the different agents for this negotiation scenario. Note that as we now have a multi-agent model, we do allow pre-coalitions but we do not insert a target coalition C_T . Agent have 3 tries to form pre-coalitions, thus $\psi = 3$. Similar to the three-agent model, the tool simulates this multi-agent model 100 times with uncertainty factor $\sigma = 0.10$.

Predictions – From Table F.14 we observe that the set of agents a_1, a_2, a_3, a_4 and a_5 have ideological values relatively close to each other, as well as the set a_6, a_7 and a_8 . Thus, we predict pre-coalitions between members of these two sets more often than cross set formed pre-coalitions. Moreover, with respect to subject *A*, the first set of agents also have more similar preferences. Therefore, we expect the majority of outcomes for *A1* to be between minimum and maximum target value of this set (i.e. [40, 50]). In addition, with respect to *A2*, we predict the outcome to be either 'medium' or 'low' with a slight advantage the first. Finally, regarding subject *B*, it can be anybodies game as they all want a large piece of the pie without giving in easily. However, as the set of agents a_1, a_2, a_3, a_4 and a_5 have a slightly higher probability to form

Table F.12: Predictions and results of simulation experiments for three-agent model.

Parameter	Prediction	Results	Verified
Power $p_1 = 10000$	Both negotiation rounds for each subject should end after a single time step, giving a_1 (near) maximum utility after forming winning coalition by himself.	Total negotiation time is indeed 2 – 1 for each round. The only member of the winning coalition is a_1 , with $u_{a_1} \approx 0.98$.	Yes
Risk behavior $\beta_1 = 10000$	Agent a_1 will give in extremely quickly, resulting in very low utilities for this agent. Moreover, most likely, the overall negotiation time decreases.	Overall negotiation times do not accede 5 steps and a_1 's utility ranges from $[0, 0.22]$	Yes
Risk behavior $\beta_1 = 0.0001$	Agent a_1 takes a lot of risk and therefore probably hesitates to accept offers. This results in longer negotiation times. In addition, we expect a_1 's utility to be similar to previous run.	Negotiation times range from $[6, 12]$ and a_1 's utility ranges from $[0, 0.35]$. Thus its utility seems slightly higher.	Yes
Rationality $\rho_1 = 0$	Will probably not have an impact in such simple negotiations.	No significant differences observed.	Yes
Disagreement point $d_1 = 1$	Agent a_1 is never part of the winning coalition as he most likely never obtains maximum utility levels.	In 10 simulation runs, a_1 was never part of the winning coalition, for neither of the subjects.	Yes
Disagreement point $d_1 = 0$	Agent a_1 more often be part of the winning coalition, but still not all the time.	For 10 simulations, a_1 was 40% and 88% part of the winning coalition, for A and B respectively.	Yes
Weights $w_1 = [1, 1, 10000]$ for issues A_1 , A_2 and B_1 .	As a_1 is more likely in winning coalition of subject B , putting extreme emphasis on B_1 probably increases overall utility. In addition, overall utility will be 0 more often, especially if a_1 is not in winning coalition of B .	Median utility u_{a_1} equals 0.7, significantly higher than other test. 3/10 simulation ended with 0 utility.	Yes.
Subject deadline $\theta_A = 0$ and $\theta_B = 0$	Tool results in error as all negotiations end in disagreement.	Tool returns "ValueError: None of the simulations did end in an agreement."	Yes
Subject deadline $\theta_A = 10000$ and $\theta_B = 10000$	Will most likely not really affect these simulations, as there is still an overall deadline $\Theta = 30$.	No significant differences.	Yes
Overall deadline $\Theta = 0$	Tool results in error as all negotiations end in disagreement.	Tool returns "ValueError: None of the simulations did end in an agreement."	Yes
Overall deadline $\Theta = 10000$	Will most likely not really affect these simulations, as there are still subject deadlines of $\theta = 20$ each.	No significant differences.	Yes
Acceptance rule $\chi = \text{unanimity}$	Majority (or all) of negotiations will not end in agreement.	Tool returns "ValueError: None of the simulations did end in an agreement."	Yes
Pre-coalition $C_T = \{a_1, a_2, a_3\}$ (all agents)	As pre-coalition is large enough to surpass the acceptance rule, both subjects end after a single time step.	The tool returns "ValueError: A 2-dimensional array must be passed."	No
Pre-negotiation tries $\psi = 0$	No pre-coalitions will form – as no C_T is entered.	Results in three, single agent coalitions that start negotiating at $t = 1$.	Yes
Pre-negotiation tries $\psi = 10000$	Agents have so many attempts to form pre-coalitions, probably the grand coalition will always form before entering the actual negotiation round. Thus, both subjects end in just one time step.	Returns the grand coalition 10 times and overall negotiations are finished in just 2 time steps	Yes




Figure F.2: Utilities (left) and Top-10 most frequent pre-coalitions (right) in 100 simulations of eight-agent model.

pre-coalitions, we expect them to perform slightly better in these negotiations. So, the overall utility of a_1 , a_2 , a_3 , a_4 and a_5 is expected to be higher compared to a_6 , a_7 and a_8 . As a result, we also assume a_1 , a_2 , a_3 , a_4 and a_5 to be more often part of the winning coalition.

Table F.13: Predictions and results of simulation experiments for three-agent model.

Output parameter	Prediction	Simulation results	Verified
Winning coalition A	Members of $\{a_1, a_2, a_3, a_4, a_5\}$ more likely to be part of winning coalition	$a_1 \approx 87\%$, $a_2 \approx 90\%$, $a_3 \approx 87\%$, $a_4 \approx 65\%$, $a_5 \approx 66\%$, $a_6 \approx 55\%$, $a_7 \approx 56\%$ and $a_8 \approx 53\%$	Yes
Winning coalition B	Members of $\{a_1, a_2, a_3, a_4, a_5\}$ more likely to be part of winning coalition	$a_1 \approx 67\%$, $a_2 \approx 71\%$, $a_3 \approx 72\%$, $a_4 \approx 72\%$, $a_5 \approx 75\%$, $a_6 \approx 54\%$, $a_7 \approx 67\%$ and $a_8 \approx 72\%$	No
Outcome A1	Majority of outcomes in range [40,50]	75% of outcomes are in [42,51]	No
Outcome A2	'medium' > 'low' and 'high' ≈ 0	'high' = 0, 'medium' = 25 and 'low' = 56	No
Utility (median)	Utilities of a_6 , a_7 and a_8 will be slightly lower than of remaining agents.	Figure F.2 (left) show utility results for all agents. It appears that utilities for a_6 , a_7 and a_8 are indeed lower as their median utilities barely meet the lower quartiles of remaining agents.	Yes
Pre-coalitions	Pre-coalitions are most likely subsets of either $\{a_1, a_2, a_3, a_4, a_5\}$ and $\{a_6, a_7, a_8\}$	Figure F.2 (right) shows top 10 most frequent pre-coalitions. These appear to be all subsets of either $\{a_1, a_2, a_3, a_4, a_5\}$ or $\{a_6, a_7, a_8\}$.	Yes

Conclusion – As the outcomes are slightly more complex in comparison to Section F.2.2, we present the results in Table F.13 supported by some visualizations. This table shows that some results match the theoretical predictions derived from our conceptual model, but some do not. As the complexity increases tremendously when adding more agents, it becomes harder to predict agent behavior. Though, those presumptions that do not (completely) match the predictions are still deemed reasonable as they fall within the uncertainty ranges. Thus, we conclude that the tool has *passed* this multi-agent verification test, even though for some observed behavior it is not entirely possible to be explained by the conceptual model due to increased uncertainty and agent interactions.

Table F.14: Negotiation scenario setup for the eight-agent model.

Agent	Parameter	Value	Agent	Parameter	Value
Agent a_1	Preference $A1$	[30, 50]	Agent a_5	Preference $A1$	[10, 40]
	Preference $A2$	[5, 12, 16]		Preference $A2$	[20 , 16, 10]
	Preference $B1$	$\{a_1 : [20, \mathbf{50}], a_2 : [20, 50], a_3 : [20, 50], a_4 : [20, 50], a_5 : [20, 50], a_6 : [20, 50], a_7 : [20, 50], a_8 : [20, 50]\}$		Preference $B1$	$\{a_1 : [20, 50], a_2 : [20, 50], a_3 : [20, 50], a_4 : [20, 50], a_5 : [20, \mathbf{50}], a_6 : [20, 50], a_7 : [20, 50], a_8 : [20, 50]\}$
	Weights w_1	[1, 1, 1]		Weights w_5	[1, 1, 1]
	Power p_1	1		Power p_5	1
	Ideology λ_1	[9, 8]		Ideology λ_5	[8, 8]
	Agent a_2	Preference $A1$		[30, 50]	Agent a_6
Preference $A2$		[5, 12, 16]	Preference $A2$	[613, 19]	
Preference $B1$		$\{a_1 : [20, 50], a_2 : [20, \mathbf{50}], a_3 : [20, 50], a_4 : [20, 50], a_5 : [20, 50], a_6 : [20, 50], a_7 : [20, 50], a_8 : [20, 50]\}$	Preference $B1$	$\{a_1 : [20, 50], a_2 : [20, 50], a_3 : [20, 50], a_4 : [20, 50], a_5 : [20, 50], a_6 : [20, \mathbf{50}], a_7 : [20, 50], a_8 : [20, 50]\}$	
Weights w_2		[1, 1, 1]	Weights w_6	[1, 1, 1]	
Power p_2		1	Power p_6	1	
Ideology λ_2		[8, 9]	Ideology λ_6	[-9, -5]	
Agent a_3		Preference $A1$	[30, 50]	Agent a_7	
	Preference $A2$	[5, 12, 16]	Preference $A2$		[6, 13, 19]
	Preference $B1$	$\{a_1 : [20, 50], a_2 : [20, 50], a_3 : [20, \mathbf{50}], a_4 : [20, 50], a_5 : [20, 50], a_6 : [20, 50], a_7 : [20, 50], a_8 : [20, 50]\}$	Preference $B1$		$\{a_1 : [20, 50], a_2 : [20, 50], a_3 : [20, 50], a_4 : [20, 50], a_5 : [20, 50], a_6 : [20, 50], a_7 : [20, \mathbf{50}], a_8 : [20, 50]\}$
	Weights w_3	[1, 1, 1]	Weights w_7		[1, 1, 1]
	Power p_3	1	Power p_7		1
	Ideology λ_3	[8, 7]	Ideology λ_7		[-4, -5]
	Agent a_4	Preference $A1$	[10, 40]		Agent a_8
Preference $A2$		[20 , 16, 10]	Preference $A2$	[6, 13, 19]	
Preference $B1$		$\{a_1 : [20, 50], a_2 : [20, 50], a_3 : [20, 50], a_4 : [20, \mathbf{50}], a_5 : [20, 50], a_6 : [20, 50], a_7 : [20, 50], a_8 : [20, 50]\}$	Preference $B1$	$\{a_1 : [20, 50], a_2 : [20, 50], a_3 : [20, 50], a_4 : [20, 50], a_5 : [20, 50], a_6 : [20, 50], a_7 : [20, 50], a_8 : [20, \mathbf{50}]\}$	
Weights w_4		[1, 1, 1]	Weights w_8	[1, 1, 1]	
Power p_4		1	Power p_8	1	
Ideology λ_4		[7, 8]	Ideology λ_8	[-1, -1]	

F.2.5. Pushing the tool to the limit: multi-agent model

This verification test is a duplication of Section F.2.3 only then using the eight-agent model in our tool. This time though, if the selected model parameter is an agent's attribute, it is changed for a_8 . The results of this test are presented in Table F.15, again after performing 10 simulations.

Table F.15: Predictions and results of simulation experiments for eight-agent model.

Parameter	Prediction	Results	Verified
Power $p_8 = 10000$	Both negotiation rounds should end after one time step, giving a_8 maximum utility. Though, as pre-coalitions are allowed, it is likely that a_6 and a_7 also benefit from a_8 's power.	Simulations end after 2 time steps. Agent a_8 always in winning coalition as well as a_6 and a_7 . Agent a_8 obtains near maximum utility and a_6 and a_7 have significantly higher utility compared to other agents.	Yes
Risk behavior $\beta_8 = 10000$	Agent a_8 gives in extremely quickly, thus possibly leading to significantly lower utilities.	Compared to the base case (see Section F.2.4) we observe a telling decrease in utility. However, we also see some (unexpected) increase in the variability of simulation utility outcomes for all agents.	Yes*
Risk behavior $\beta_8 = 0.0001$	As agent a_8 now plays an extremely bold strategy, we expect the utility to be slightly higher. Thus at the same level as a_6 and a_7 , as it is possibly a better strategy for their (minority) situation.	Utility – both in terms of outcome and variability – fairly similar to a_6 and a_7 .	Yes
Rationality $\rho_8 = 0$	We expect agent a_8 variability in utility to be higher, as he makes decisions more randomly.	No significant changes noticeable.	No
Disagreement point $d_8 = 1$	Agent a_8 will almost never be part of the winning coalition, as his disagreement point is so high that no offer is good enough.	Agent a_8 is part of winning coalition for subject A and B respectively 4% and 8%.	Yes
Disagreement point $d_8 = 0$	Once in a coalition, agent a_8 possibly will never leave as he will never disagree with coalitional offers and acceptance.	Agent a_8 is part of winning coalition for subject A and B respectively 60% and 82%. Significantly higher than the previous experiment and his like-minded opponents a_6 and a_7 .	Yes
Weights $w_8 = [1,1,10000]$ for issues $A1$, $A2$ and $B1$.	Agent a_8 utility decreases as it almost solely depends on $B1$ and as he is too weak to be decisive for $B1$.	Agent a_8 's median utility is less than lower quartile of like-minded opponent agents a_6 and a_7 . In addition, we also observe a tremendous increase in the variability of a_8 's outcomes.	Yes**
Subject deadline $\theta_A = 0$ and $\theta_B = 0$	Tool results in error as all negotiations end in disagreement.	Tool returns "ValueError: None of the simulations did end in an agreement."	Yes
Subject deadline $\theta_A = 10000$ and $\theta_B = 10000$	Will most likely not really affect these simulations, as there is still an overall deadline $\Theta = 30$.	No significant differences.	Yes
Overall deadline $\Theta = 0$	Tool results in error as all negotiations end in disagreement.	Tool returns "ValueError: None of the simulations did end in an agreement."	Yes
Overall deadline $\Theta = 10000$	Will most likely not really affect these simulations, as there are still subject deadlines of $\theta = 20$ each.	No significant differences, neither in time, nor in results.	Yes
Acceptance rule $\chi = \text{unanimity}$	Most likely none of the negotiations end in agreement as too many agents have conflicting preferences and the negotiation time is rather limited.	Tool returns "ValueError: None of the simulations did end in an agreement."	Yes

Pre-coalition $C_T = A$ (all agents)	As pre-coalition is large enough to surpass the acceptance rule, both subjects end after a single time step.	The tool returns <i>"ValueError: A 2-dimensional array must be passed."</i>	No
Pre-negotiation tries $\psi = 0$	No pre-coalitions will form – as no C_T is entered.	Every simulation starts with eight, single agent coalitions that start negotiating at $t = 1$.	Yes
Pre-negotiation tries $\psi = 10000$	Agents have so many attempts to form pre-coalitions, probably the grand coalition will always form before entering the actual negotiation round. Thus, both subjects end in just one time step.	Returns the grand coalition 10 times and overall negotiations are finished in just 2 time steps	Yes

Discussion – As we can see in Table F.15, two verified outcomes^[13] need some more thorough explanation since they resulted in both expected as well as unexpected behavior. First consider the test in which we increased $\beta_8 = 10000$. Besides the predicted decrease in utility, we also observed an increase in variability of outcomes. Therefore, we extended our experiment by performing 1000 simulations of the based model (as setup in Section F.2.4) and extreme model each. As we observed more variability in outcomes for all agents, we then randomly choose an agent (i.e. a_3) and compare utility distributions for both models. Since we assume utilities to be continuously distributed on the interval $[0,1]$, we compute the Kolmogorov-Smirnov statistic D ^[14] to test if the two datasets stem from the same distribution (null hypothesis H_0) or not (alternative hypothesis H_1). The results of the simulation experiment are shown in Figure F.3. Although both plots in this figure look very similar, the results of the Kolmogorov-Smirnov statistic – $D \approx 0.098$ with a p -value of ≈ 0.001 – suggest to reject the null hypothesis, even for very low values of α . Unfortunately, we cannot provide a theoretical explanation why these outcomes differ. Therefore, in our tool, we limit ourselves by not entering extreme input parameters for risk behavior.

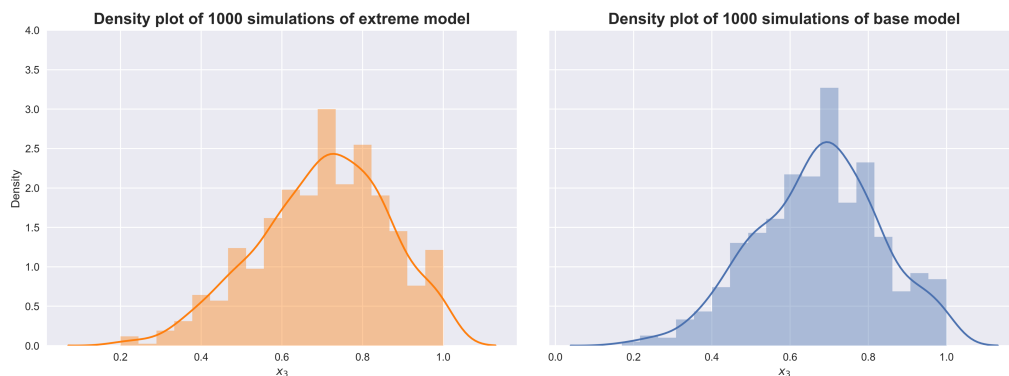


Figure F.3: Utility distributions for 1000 simulations of the base and extreme model.

Secondly, we consider the case in which we observed unexpected behavior after increasing $w_8 = [1, 1, 10000]$. Although the increased spread in outcomes was unexpected, we assume there is a plausible explanation. In the agent setup (see Table F.14), we have deliberately created two sets of near like-minded agents for subject A and B , namely $\{a_1, a_2, a_3, a_4, a_5\}$ and $\{a_6, a_7, a_8\}$. This implies that a_8 is not decisive, not even if he forms a pre-coalition with any member of $\{a_6, a_7, a_8\}$. By putting extreme emphasis on issue $B1$ – in which a_8 has no decisive ability to influence the outcome effectively – a_8 becomes extremely dependent on the other agents. Moreover, the contributions of $A1$ and $A2$ to a_8 overall utility result is negligible. As a consequence, a_8 's overall utility solely depends on the outcome of $B1$ in which a_8 himself has no major role. As a consequence, resulting utilities are more volatile.

[13] Indicated by '*' and '**' respectively

[14] Note that we are able to do so because the data is one-dimensional and assumed to be continuous.

Conclusion – With the discussion above taken into account, our tool *passes* this verification test. However, for future work, we do recommend a more thorough analysis of extreme risk behavior – either extremely risk averse or seeking – especially in those negotiations in which extreme risk behavior plays a dominant role for the involved parties. Moreover, we again recommend resolving the error when entering the grand coalition as target coalition.

F.2.6. Testing tool variability

We conclude our verification experiments by performing different variability tests. As we have seen in previous sections, variability plays an important role in the outcomes. However, it can be both intended – i.e. caused by uncertainty in input parameters – as well as unintended. The latter is often undesired, but does not necessarily have to be [105, p. 137-138]. If models gain complexity, they become more uncertain and, as a result, unintended or random outcomes increase [105, p. 139]. In our tool, the possibility for unintended randomness increases if more complex negotiation scenarios are inserted. Complexity of the negotiation scenario is primarily determined by an increase in negotiating agents – as number of possible interactions between agents are of $\mathcal{O}(n^2)$ (see Section F.1.2) – and/or an increase in issues per subject^[15], as the latter boosts the dimensionality of the outcome space.

In this section, we therefore perform 12 experiments – consisting of 1000 simulation runs each – in which we vary the number of negotiating agents and issues according to Table F.16. Agents are setup in a similar way as in Section F.2.2. Thus they have similar strategy profiles and risk behavior, whereas their ideological values are randomly distributed over the xy -plane in such a way that pre-coalition might (not) form. The number of tries agents are allowed to form pre-coalitions is set to $\psi = 3$. Moreover, the negotiation scenario solely consists of continuous, unlimited issues in the range $[0,100]$. Agents' preferences for these issues – i.e. their target and reservation values – are thus subsets of $[0,100]$. As we are interested in unintended randomness, we assume the uncertainty factor $\sigma \rightarrow 0$.

Discussion – The variability is determined for the following simulation outcomes: overall utility, aggregated issue outcomes^[16], negotiation time and size of winning coalition. Figure F.4 presents the results of the 12 simulation experiments. We will discuss them one at the time.

The first compelling, slightly counter-intuitive observation is that the variability of *utility* results does not depend on the number of agents, but primarily on the number of issues in the negotiation domain. A reasonable explanation for this behavior is that with only one or few issues, it becomes hard for agents to make package-deals. As a result, offers are either acceptable to agents or not, as there is no potential to balance unfavorable issue results for favorable ones. Therefore, we observe that when the number of issues for each subject increases, the more likely it is agents derive some – instead of none – utility from the final offer and, thus, shrinking the utility variability. In our experiments, this is the case for E_7 and beyond, which translates to 10 or more issues.

Secondly, we observe again an increase in variability due to an increase in issues, though this time for *outcomes*. Two distinct situations become apparent: down to E_6 or 5 issues, a variability in the range of $[10, 70]$ is realized, whereas from E_7 and beyond we this range increase to roughly $[10, 80]$. This is most likely caused by the increase in dimensionality of the subject. For E_7 to E_9 negotiation outcomes are 10-dimensional which increases to 20-dimensional for E_{10} to E_{12} . As the outcome space tremendously expands for higher dimensions and preferences for every agent and issue are randomly distributed within $[0,100]$, variability in outcomes must inevitably increase as well.

Thirdly, we see that variability in total negotiation time – steps needed to reach an agreement – mainly increases by adding more agents to the negotiation environment. For three agents (i.e. E_1 , E_4 , E_7 and E_{10}), in most simulations a majority agreement is reached after just a few time steps, whereas this can be any where between 1 and 20 time steps for the six and eight agent models. This increase in variability is, of course, not surprising, as the number of possible (winning) coalitions increases extensively. Though, it seems that the effect slightly diminishes each time agents are added. However, we think that this is mainly due to the overall negotiation deadline, which is set to $\Theta = 20$. The latter should be examined more carefully, but

^[15]Note that we do not expect unintended randomness caused by an increase in subjects as they do not affect the dimensionality of the outcome space.

^[16]As the experiments differ in the number of issues, it is not possible to compare the results of all individual issue outcomes. However, because issues vary over the same range, we can aggregate the results.

Table F.16: Experimental design with experiments E_i examining variability due to agents and/or issues.

		Agents		
		3	6	9
Issues	1	E_1	E_2	E_3
	5	E_4	E_5	E_6
	10	E_7	E_8	E_9
	20	E_{10}	E_{11}	E_{12}

Table F.17: Fraction of experiments that resulted in an agreement.

	E_1	E_2	E_3	E_4	E_5	E_6	E_7	E_8	E_9	E_{10}	E_{11}	E_{12}
In agreement	1.0	0.19	0.99	0.89	0.34	0.45	0.9	0.05	0.13	0.87	0.04	0.07

does not have consequences for the remaining of this thesis.

Finally, the variation in size of the minimal winning coalition is very limited. For experiments with only three agents (i.e. E_1 , E_2 , E_3 and E_4) the winning coalition almost always consists of two agents. In some cases, even the grand coalition establishes. This is also the case for experiments containing six agents, but not for nine. Obviously, it becomes harder to form the grand coalition in negotiation environments with lots of agents due to more possible conflicting preferences. Moreover, agents' willingness to form larger than minimal winning coalition potentially diminishes as well. Aforementioned provide reasonable explanations why there is limited variability in winning coalition sizes.

Conclusion – Altogether our tool *passes* the variability tests as the observed unintended randomness can be well explained by theory. However, the variability of negotiation time needs some closer attention. Note that the latter does not affect the remaining of our thesis.

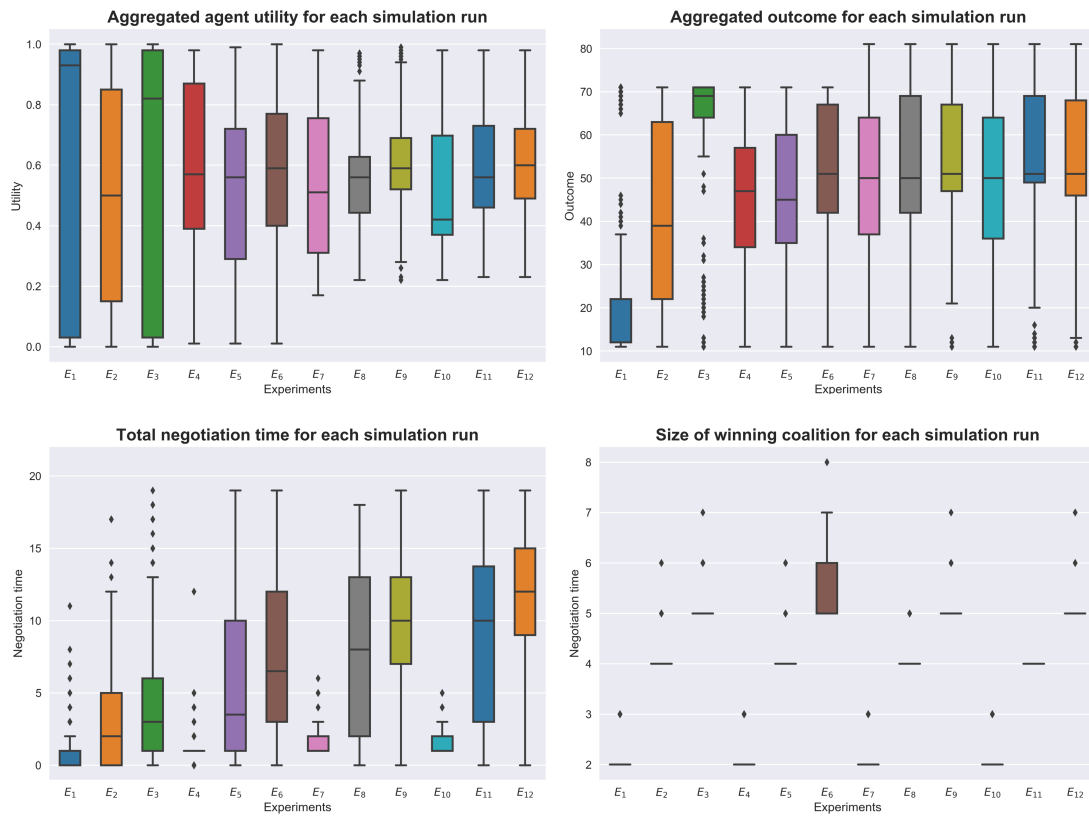


Figure F.4: Results of variability experiments for aggregated utility (top-left), outcomes (top-right), negotiation time (bottom-left) and size of winning coalition (bottom-right).



Experimental design

This appendix gives an oversight of all factors in the experimental design, including to what extent they are varied. Moreover, we test the validity of the experimental design by running several tests.

G.1. Factors in the experimental design

Table G.1 present all factors derived from Section 6.2.1, Section 6.2.2 and Section 6.2.3. In the negotiation setting, we have chosen the ratio between strong, secondary and weak states to be at least 4 : 2 : 1. But, in extreme cases, this changes to 24 : 10 : 1. Moreover, as the agents represent states, we should presume some ideological differences between them. For n groups of agents, these differences are represented by k -dimensional clusters drawn from (n, k) -dimensional multivariate normal distributions with $\mu \in \mathbb{R}^k$ and $\Sigma \in \mathbb{R}^{k \times k}$. So, if there are 10 agents in the negotiation environment, differing in three ideological dimensions ($k = 3$) and divided between three groups ($n = 3$), this results in the ideological space as presented in Figure G.1 (left). Note that both the number of agents in a cluster and its internal- and external distances are still randomly sampled. So, if agents are in the same cluster, this does not imply they have to form pre-coalitions, but the likelihood increases.

Furthermore, characteristics of the target pre-coalition C_T are not specific input parameters of our tool, as they are (aggregated) attributes of the member agents. Therefore, to create the target pre-coalition C_T , agents were randomly sampled from the negotiation environment such that they resemble the factors for each simulation experiment.

Finally, Table G.1 also contains the ranges over which the factors are varied. These ranges together form a 15-dimensional parameter space, containing endless of possible simulation experiments. Therefore, Latin Hypercube sampling has been chosen to sample experiments from this space. This is done such that we




Figure G.1: Two examples 10 agents distributed over a three-dimensional ideological space with (left) 3 clusters of ideological like-minded agents and (right) no clusters.

Table G.1: Factors included in the experimental design.

Factor	Notation	Part of	Ranges	Type	Note
Number of agents	$ A $	Setting	[3, 18]	Integer	-
Power of weak agents	p_{weak}	Setting	[1, 6]	Integer	$p_i \sim U(1,6)$
Power of secondary agents	$p_{\text{secondary}}$	Setting	[10, 15]	Integer	$p_i \sim U(10,15)$
Power of strong agents	p_{strong}	Setting	[19, 24]	Integer	$p_i \sim U(19,24)$
Number of ideological clusters	Λ	Setting	[0, 8]	Integer	-
Mean rationality of agents	μ_R	Setting	(0, 1)	Float	$\rho_i \sim N(\mu_R, 0.1)$
Mean risk seeking behavior	μ_B	Setting	(0, 2)	Float	$\beta_i \sim N(\mu_B, 1)$
Mean aggressiveness of agents	μ_a	Setting	(0, 1)	Float	$f_a \sim P(\Omega_{i \rightarrow i+1} = X) = a_i$ in which $a_i \sim N(\mu_a, 0.1)$ ^[3]
Mean willingness to negotiate	μ_D	Setting	(0,1)	Float	$d_i \sim N(\mu_D, 0.1)$
Number of subjects	$ S $	Domain	[1, 15]	Integer	-
Number of issues	$ I $	Domain	[1, 5]	Integer	-
Size of coalition	$\frac{ C_T }{ A }$	Pre-coalition	(0, 0.5)	Float	-
Power of coalition	$\frac{\sum_{i=1}^{C_T} p_i}{\sum_{i=1}^A p_i}$	Pre-coalition	(0, 0.5)	Float	-
Fraction weak	$\frac{\sum_{i=1}^{C_T} p_{\text{weak}}}{ C_T }$	Pre-coalition	(0, 1)	Float	-
Fraction secondary	$\frac{\sum_{i=1}^{C_T} p_{\text{secondary}}}{ C_T }$	Pre-coalition	(0,1)	Float	-
Fraction strong	$\frac{\sum_{i=1}^{C_T} p_{\text{strong}}}{ C_T }$	Pre-coalition	(0,1)	Float	-

maximize the minimum distance between the sample experiments^[1]. This completes the experimental setup as part of the design. Note that this setup is implemented in the *scenario.py*^[2] module which transforms the Latin Hypercube sample, to a single simulation experiments with a concrete negotiation scenario.

G.2. Validity of the experimental design

As described in Section 6.3, there are two components determining whether the experimental design as presented above is right. First, we should determine whether the experimental setup represent real international political negotiations. Secondly, the presumed relations should answer the main research question of this study. In this section, we give both quantitative and qualitative evidence supporting the experimental setup. We do so by testing two important assumptions in the example negotiation domains and settings.

G.2.1. Negotiation domain

Let us first validate the negotiation domain by having a fixed setting. The fundamental assumption made to construct the experimental design is creating a "pool" of different issue types from which domains are derived. We shall validate whether our simulation results are biased by creating 4 other pools and comparing them with the event population as introduced in Table 6.1. This comparison is done on negotiation *time*, obtained *utilities* and the *size/power* of winning coalition^[4] and the fraction of negotiation resulting in *agreement*. For sake of convenience, the same eight agents of the example run in Section 5.5.3 are used. Note that only their reservation and target values differ, as different issue types are used. Table G.2 presents the pool of issues for the 'base' pool, as well as for four alternative pools, denoted with 'alt1' to 'alt4'.

For every pool (or negotiation domain) 200 simulation experiments have been conducted. In Figure G.2 we compare agents' utilities derived from the negotiation outcomes in each of the domains. Of course,

^[1]The *pyDOE* package provides an efficient way to do so (see <https://pythonhosted.org/pyDOE/randomized.html>).

^[2]See also Section 5.5.2.

^[3]Note that $X = \text{"Competitor"}$.

^[4]Note that the actual negotiation *outcomes* are left out as they will differ anyway.

Table G.2: Possible issues and outcome spaces in the simulation experiments.

Issue type	Outcomes 'base'	Outcomes 'alt1'	Outcomes 'alt2'	Outcomes 'alt3'	Outcomes 'alt4'
Indivisible resources	$ O = \{2, 3, 5, 8\}$ and $\sum_{j \in O} o_j = 100$	$ O = \{3, 5, 15, 18\}$ and $\sum_{j \in O} o_j = 150$	$ O = \{2, 7, 13, 14\}$ and $\sum_{j \in O} o_j = 500$	$ O = \{2, 3, 10\}$ and $\sum_{j \in O} o_j = 50$	$ O = \{2, 3\}$ and $\sum_{j \in O} o_j = 100$
Unrestricted resources	$O = \{[30, 100], [5, 25], [120, 150], [200, 400], [20, 30]\}$	$O = \{[500, 1500], [50, 1300], [20, 100], [16000, 17000], [10, 50]\}$	$O = \{[15, 20], [5, 25], [7, 20], [800, 1400], [7, 290]\}$	$O = \{[10, 30], [10, 25], [10, 150], [10, 50], [10, 80]\}$	$O = \{[750, 1000], [325, 525], [20, 50], [52200, 53400], [200, 300]\}$
Limited resources	$\pi = \{50, 200, 500\}$	$\pi = \{1000, 1200, 50\}$	$\pi = \{500, 2000, 500\}$	$\pi = \{5000, 2000, 50000\}$	$\pi = \{30, 50, 510\}$

we observe differences in utility; as the negotiation outcomes changes, most likely the derived utility also differs to some extent. However, overall, the derived utility should be in line for every negotiation domain as the setting – and, thus, the power and negotiation behavior – is fixed. We observe that this holds for most agents. For the 'base' pool, only a_6 is a true exception. Such exceptions are not surprising, as in Section 2.2.6 we stated power can not only be derived from a state's structural resources, but also from the issues under consideration. Thus, if those issues do not particularly suit a certain negotiator – especially compared to others – despite his structural power, he may still achieve bad results.

Moreover, in Figure G.3 simulation results for negotiation time, size and power of the winning coalition are presented. We do observe great differences in time to agreement and power of the winning coalition. Both 'alt1' and 'alt4' take significantly less negotiation steps before reaching agreement. Whereas 'alt2' and 'alt3' are more or less in line with the 'base' domain upon which the results of this thesis are based. The same holds for the power of the winning coalition in which 'alt1' and 'alt4' are significantly more powerful than the others. In terms of the size of the winning coalition, all pools show comparable results. Finally, Table G.3 compares the likelihood of agreement of each negotiation domain. This table supports our conclusion that 'alt2' and 'alt3' significantly differ from other domains as they have high likelihoods of agreement.

Table G.3: Likelihood of agreement comparisons of 5 different negotiation domains.

	'base'	'alt1'	'alt2'	'alt3'	'alt4'
In agreement (#)	24	180	23	59	196
In agreement (%)	12.0	90.0	11.5	29.5	98.0

To conclude, an agent's utility is relatively independent of the chosen negotiation domain. The exceptions that do occur are domains in which agents simply have bad luck, as apparently that particular domain does not suit the agent well. In addition, we have observed the negotiation domain is insignificant for the size of the winning coalition as well. In contrary, different domains do affect negotiation time, power of the winning coalition and likelihood of agreement significantly. Therefore, conclusions of this thesis concerning the time objective Q_3 , political support Q_4 and likelihood of agreement Q_5 can definitely not be generalized.

G.2.2. Negotiation setting

Let us now validate the negotiation setting by having a fixed (base) domain. The fundamental assumption made to sketch the negotiation setting is the ratio of power between weak, secondary and strong states. In the 'base' setting upon which the results of this thesis are based a minimum ratio of 4:2:1 is adopted. To validate this assumption, the following four alternative ratios have been examined:

- 'base' adopts a minimum power ratio of 4:2:1;
- 'alt1' adopts a minimum power ratio of 3:2:1;
- 'alt2' adopts a minimum power ratio of 10:2:1;
- 'alt3' adopts a minimum power ratio of 10:3:1;
- 'alt4' adopts a minimum power ratio of 15:4:1;

Table G.4: Likelihood of agreement comparisons of 5 different power ratios.

	'base'	'alt1'	'alt2'	'alt3'	'alt4'
In agreement (#)	110	118	173	31	164
In agreement (%)	55.0	59.0	86.5	15.5	82.0

Other factors in the negotiation setting – as presented in Table G.1 are again equal to the example run in Section 5.5.3. Similar to the validation test before, we run 200 simulations for each alternative ratio. Note that in every simulation run $\{a_1, a_2\}$ are characterized as strong, $\{a_3, a_4, a_5, a_6\}$ as secondary and $\{a_7, a_8\}$ as weak. Figure G.4 compares the differences in utility for all adopted ratios. We observe that as strong nations become stronger, it is more likely that only one of them is part of the winning coalition and, thus, in this setting a_2 has lower utilities. For secondary and weak agents we observe more capricious results. First, it looks like the personal gains of weak agents are more dependent on others as the utility ranges are larger. Secondly, a higher ratio compared to weak agents and lower ratio compared to strong agents seems beneficial for secondary agents as in general 'alt4' results in higher utilities. Finally, one very (strong) strong agent is beneficial for weak agents, especially if the secondary states differ only little; presumably because in such settings it is more convenient to bandwagon with strong agents (and *against* secondary states).

Furthermore, Figure G.5 simulation results for negotiation time, size and power of the winning coalition are presented. Obviously, as the absolute power of strong states increases tremendously in these alternatives, the power of the winning coalition increases as well. Though, we see that its size is more or less stable for all ratios. Finally, we observe that both the likelihood of agreement in Table G.4 and the negotiation time are certainly not independent of the adopted power ratio. Especially 'alt2' and 'alt3' significantly differ.

To conclude, apart from the political time objective Q_3 , all agent's objectives – and therewith the sufficiency of pre-coalitions – are dependent on the power ratios. Thus, the conclusions of this thesis can definitely not be generalized for all power settings. It can only be generalized for those settings in which secondary states still have got some decisiveness compared to strong powers. So, in settings of true hegemony, the conclusions are insignificant.




Figure G.2: Agent utility comparison of 5 different negotiation domains.




Figure G.3: Negotiation time, size and power of the winning coalition comparisons of 5 different negotiation domains.




Figure G.4: Agent utility comparison of 5 different power ratios.





Figure G.5: Negotiation time, size and power of the winning coalition comparisons of 5 different power ratios.



Simulation results

This appendix contains all results and supporting analyses of the simulation experiments conducted in Chapter 7. These have been categorized according to the three different agent perspectives, namely weak (in Section H.1), secondary (in Section H.2) and finally strong agents (in Section H.3).

H.1. Perspective of weak agents

The results in this section are solely devoted to weak agents.

H.1.1. Correlation between factors

Figure H.1 shows a pairplot is shown of all (original) factors in the Random Forest regression model to examine correlation between them. From left to right and top to bottom, the following factors are paired: subjects, issues, fraction power of weak agents, fraction power of secondary agents, fraction power of strong agents, number of agents, average aggressiveness, average rationality, average risk behavior, average willingness to negotiate, number of ideological clusters of agents, size of pre-coalition, fraction weak members, fraction secondary members, fraction strong members and finally, the power of the pre-coalition. Besides, the correlations are also shown in Table H.2.

We observe that the vast majority of factors are (nearly) uncorrelated. The factors for which correlation occurs are mainly related to pre-coalitions. Which come as no surprise as, for example, the fractions of weak, secondary and powerful agents are interrelated. Also the size of the pre-coalition and its power are to some extent correlated. Finally, we notice some correlations between the number of agents in the negotiation and the pre-coalition's size and power. As the above factors are important to predict the influence of pre-coalitions on the quality of outcome, we will keep them as independent variables. However, a side note is placed when calculating factor importances and/or partial dependencies.

Finally, if we look at the distributions of each factor in the data (diagonals in Figure H.1, we see that Latin Hypercube sampling has done a good job in sampling uniformly over the parameter space. Only those factors related to fractions of agents are slightly skewed. However, this is related to the fact that we have chosen to represent them as fractions of power. Thus, the underlying factors – namely the number of weak, secondary and strong agents respectively – are still uniformly sampled over the parameter space.

Table H.1: Factor conversion table.

#	Factor	#	Factor	#	Factor	#	Factor
1	subject	5	p_{strong}	9	avg risk	13	fraction weak
2	issue	6	agents	10	avg willingness	14	fraction sec
3	p_{weak}	7	avg aggressiveness	11	clusters	15	fraction strong
4	p_{sec}	8	avg rationality	12	pre-coalition size	16	pre-coalition power




Figure H.1: Pairplot visualizing the correlation between factors in Random Forest model for weak agents simulations.

Table H.2: Correlations between factors explaining differences in outcome for weak agents.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	1.000	0.004	-0.013	0.015	-0.009	0.012	-0.003	-0.020	-0.011	0.045	0.020	-0.023	0.014	-0.006	-0.011	-0.021
2	0.004	1.000	-0.008	-0.004	0.007	0.010	0.035	-0.005	-0.013	0.023	-0.020	0.025	-0.053	0.010	0.056	0.040
3	-0.013	-0.008	1.000	-0.020	-0.349	-0.051	0.013	-0.031	-0.061	-0.006	0.033	0.034	0.198	-0.082	-0.168	0.089
4	0.015	-0.004	-0.020	1.000	-0.929	-0.042	-0.051	0.025	-0.030	0.000	0.010	0.020	-0.048	0.183	-0.126	0.0748
5	-0.009	0.007	-0.349	-0.929	1.000	0.058	0.043	-0.012	0.050	-0.005	-0.022	-0.032	-0.027	-0.141	0.180	-0.103
6	0.012	0.010	-0.051	-0.042	0.058	1.000	-0.002	0.017	0.032	0.050	0.018	-0.422	0.038	-0.004	-0.044	-0.334
7	-0.003	0.035	0.013	-0.051	0.043	-0.002	1.000	0.006	-0.012	-0.024	-0.020	0.006	-0.010	0.020	-0.008	0.003
8	-0.020	-0.005	-0.031	0.025	-0.012	0.017	0.006	1.000	0.012	-0.044	0.031	-0.027	-0.001	-0.015	0.017	-0.009
9	-0.010	-0.013	-0.061	-0.030	0.050	0.032	-0.012	0.012	1.000	-0.023	-0.042	-0.006	-0.035	0.023	0.021	0.009
10	0.004	0.023	-0.007	0.008	-0.005	0.050	-0.022	-0.044	-0.023	1.000	0.008	-0.054	0.023	-0.029	0.001	-0.036
11	0.020	-0.020	0.033	0.010	-0.022	0.018	-0.020	0.031	-0.042	0.008	1.000	-0.012	0.029	-0.003	-0.033	-0.012
12	-0.023	0.025	0.034	0.020	-0.032	-0.422	0.006	-0.027	-0.006	-0.054	-0.012	1.000	-0.590	0.355	0.387	0.917
13	0.014	-0.053	0.198	-0.045	-0.027	0.038	-0.010	-0.001	-0.035	0.023	0.029	-0.590	1.000	-0.641	-0.614	-0.720
14	-0.006	0.010	-0.082	0.183	-0.141	-0.004	0.020	-0.015	0.023	-0.029	-0.003	0.355	-0.641	1.000	-0.211	0.323
15	-0.011	0.056	-0.168	-0.126	0.180	-0.044	-0.008	0.017	0.021	0.001	-0.033	0.387	-0.614	-0.211	1.000	0.585
16	-0.022	0.040	0.089	0.074	-0.103	-0.334	0.003	-0.009	0.009	-0.036	-0.012	0.917	-0.720	0.323	0.587	1.000

H.1.2. Fitness of Random Forest models

The scatter plots in Figure H.2 and Figure H.3 show the predicted versus actual responses – i.e. quality of outcome QO and its quality measures Q_i – for the factors in "base" and "+pre-coalition" Random Forest models. Thus, these plots give us an insight how well the simulation results are explained by the factors in the Random Forest models.

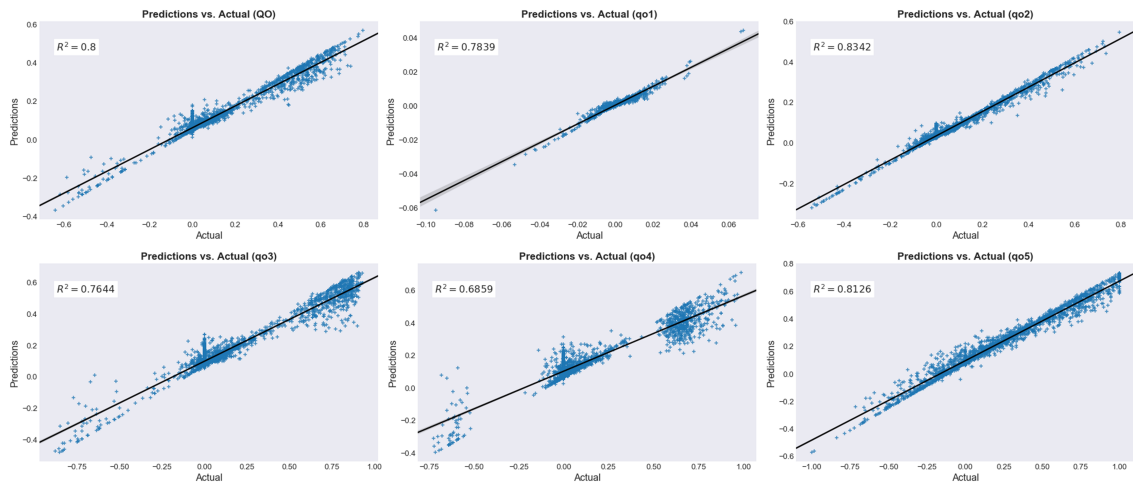


Figure H.2: Weak perspective: predictions vs. actual for the "base" model.

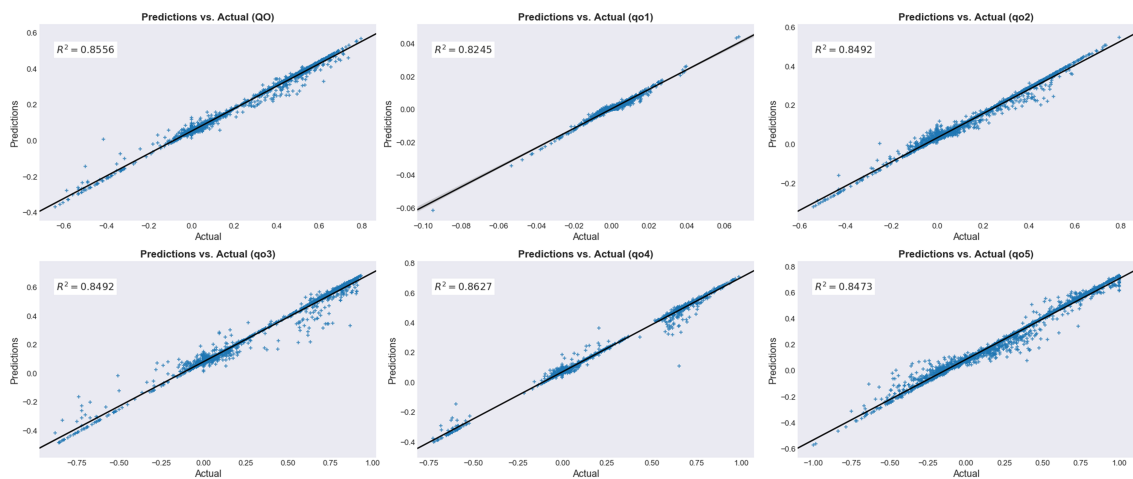


Figure H.3: Weak perspective: predictions vs. actual for the "+pre-coalition" model.

H.1.3. Partial dependency plots

Figure H.5 to Figure H.9 show the single factor partial dependence plots for the overall important factors in the Random Forest models (see Figure H.5). Positive values on the y -axis indicate pre-coalitions add value to weaker agents given the factor under consideration. Of course, the opposite holds for negative values. Note that these plots have been created using *scikit-learn*^[1].

Partial dependence QO

- In Figure H.4, we observe two strong (positive) relations, five neutral relations and three moderate relations, of which two positive;
- We observe that only in negotiation settings with more than 8 agents a positive difference is noticeable. Thus, overall, forming pre-coalitions for weak agents becomes fruitful if there are more agents in the negotiation environment. This effect is strong as the factor is highly important;
- We also notice that if the pre-coalition's power covers roughly 25% of the total power in the negotiations, the difference in outcome becomes positive. The latter increases tremendously for fractions over 40%. Also this effect is strong, as the factor is overall highly important;
- Moreover, we see that differences in quality of outcome become positive for larger negotiation domains, containing over 5 subjects or more;
- Finally, a moderate negative dependence is observed if the willingness to negotiate increases^[2]. Thus, the difference in outcome is negatively affected if agents refuse to negotiate;

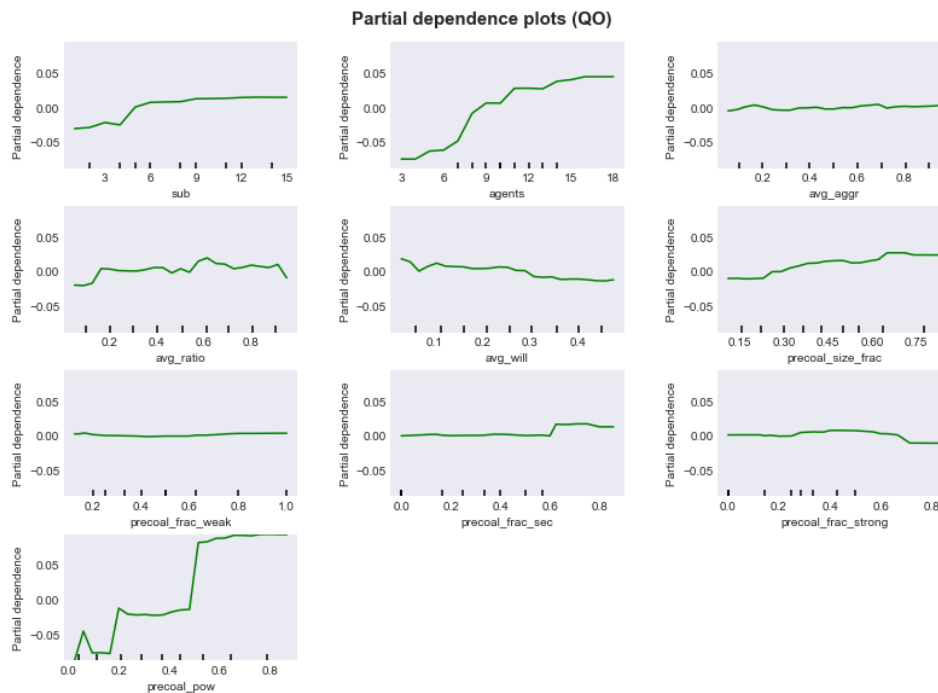


Figure H.4: Weak perspective: single factor partial dependencies for QO.

Partial dependence Q_1 ^[3]

- Judging the y -axis' scale – ranging from -0.001 to 0.001 – the contribution to the difference of outcome for Q_1 is very minor;
- Only increasing the fraction of secondary agents in the pre-coalition above 65% seem to have a modest effect on the difference in the outcome of the Nash product (as social welfare measure).

[1]For the partial dependence plot library see http://scikit-learn.org/stable/auto_examples/ensemble/plot_partial_dependence.html.

[2]Note that the higher the willingness, the less likely it is agents negotiate as their outside options become very good.

[3]See Figure H.5.

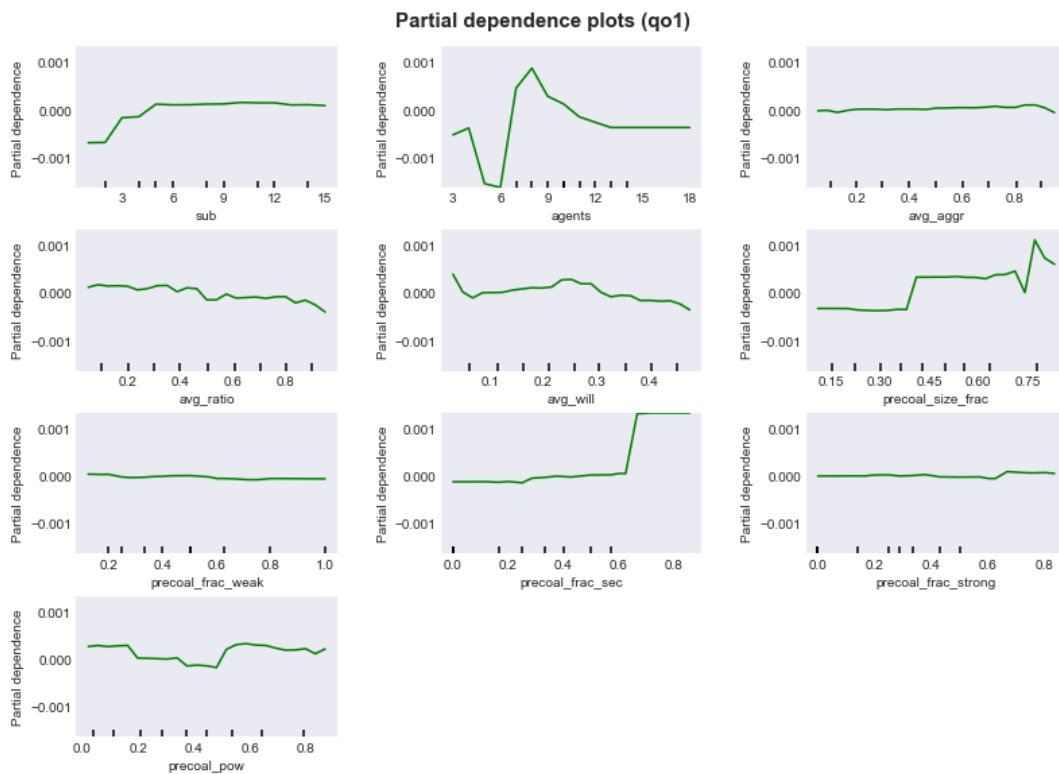


Figure H.5: Weak perspective: single factor partial dependencies for Q_1 .

Partial dependence Q_2 ^[4]

- Obviously, the dependence relations are very similar to Q_0 , as they account for 75% of its value;
- Though, we observe a less strong positive effect on the quality of outcome if the number of agents increases. Only from 9 or more agents the effect becomes positive, which implies that in terms of personal gains, forming pre-coalitions is beneficial in negotiation settings with high numbers of agents;
- Again, we notice that if the pre-coalition's power covers roughly 25% of the total power in the negotiations, the differences in personal gains become positive, which increases tremendously for higher fractions. Though, for lower fractions, it sticks slightly below zero, which might be best explained by the concessions the weaker agents have to do to form pre-coalitions;

Partial dependence Q_3 ^[5]

- Again, also for the weaker agents' time objective, two partial dependencies stand out whereas the other remain more or less neutral;
- First, we observe a strong dependence on agents. From 8 or more agents, it becomes beneficial for weaker agents to form pre-coalitions, as the difference is positive. Thus, pre-coalitions in negotiation settings with more agents contribute to achieving weaker agents' personal deadlines better;
- Obviously, the same holds for larger fractions of pre-coalitions. If the pre-coalition already exceeds nearly 50% of the power or more, chances weaker agents meet their time objectives increase as the difference in outcome between staying alone and forming a pre-coalition increases too;
- The above is also reflected in Q_5 , in which we observe the difference in likelihood of outcome increases tremendously if the pre-coalition already has nearly 50% of the power;

^[4]See Figure H.6.

^[5]See Figure H.7.

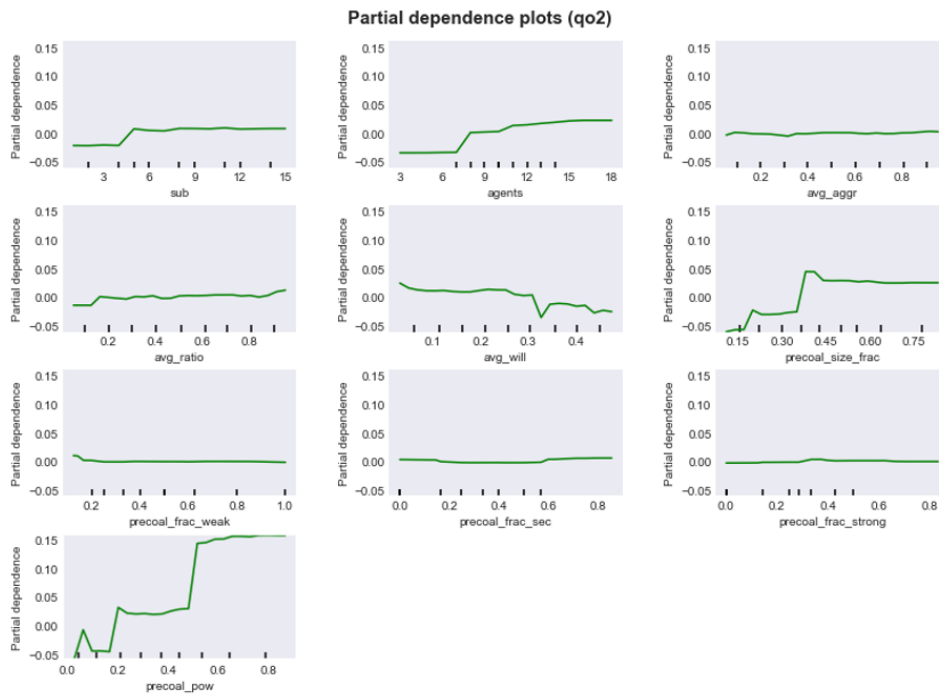


Figure H.6: Weak perspective: single factor partial dependencies for Q_2 .

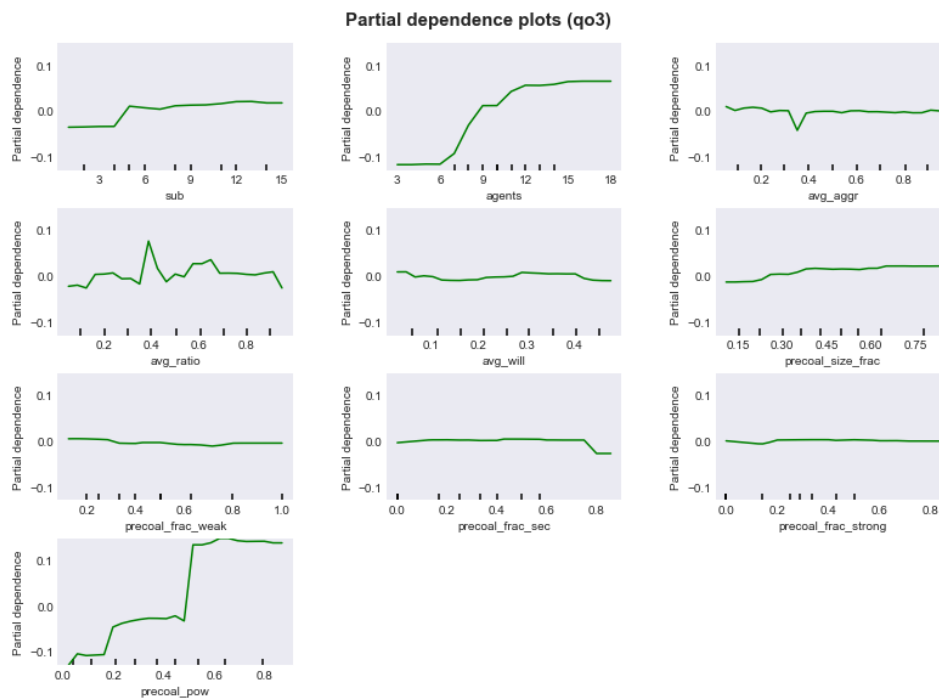


Figure H.7: Weak perspective: single factor partial dependencies for Q_3 .

Partial dependence Q_4 ^[6]

- First, we notice that almost no partial dependence plot exceed $y = 0$ (or x-axis). This implies that forming pre-coalitions almost always negatively affects the political support (as the added difference is zero or less most of the time);
- Increasing the fraction of power of the pre-coalition is the only clear exception. Evidently, as weaker agents manage to form pre-coalitions with the vast majority of (powerful) agents, political support increases. However, one can argue if this is realistically.

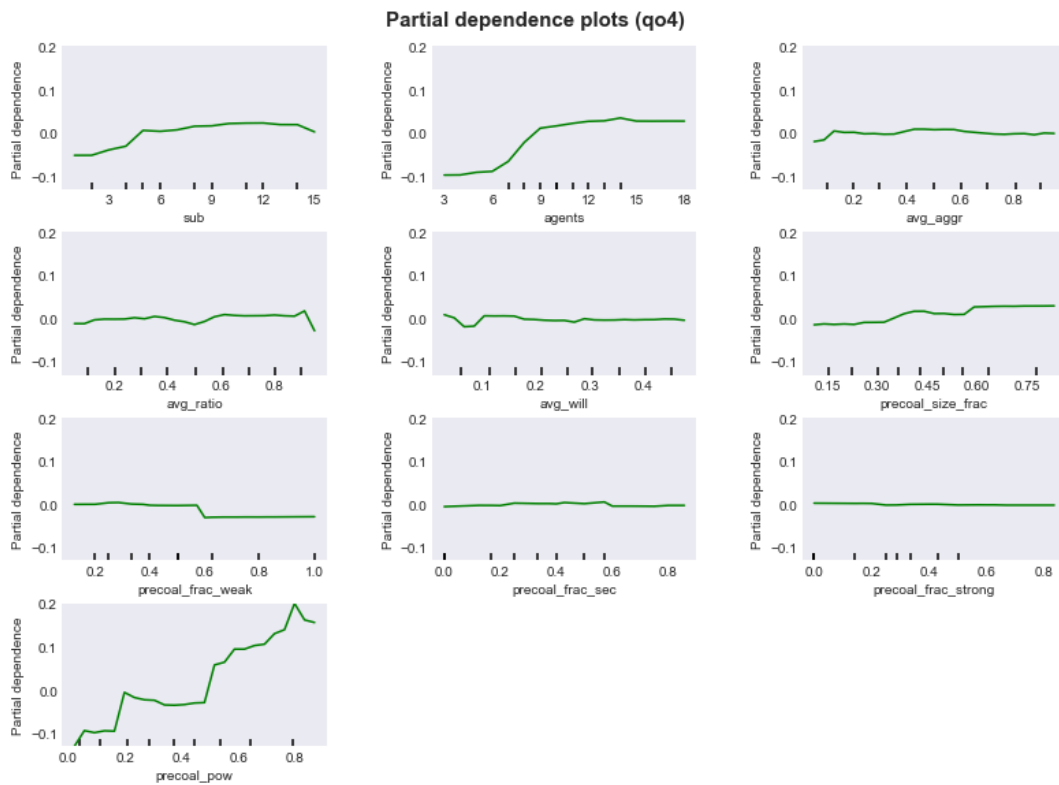


Figure H.8: Weak perspective: single factor partial dependencies for Q_4 .

Partial dependence Q_5 ^[7]

- First, we observe that if weak agents form pre-coalitions this does not necessary contribute to higher likelihood of outcome (or, to be precise, a greater difference). Especially up to 7 agents, if weaker agents form pre-coalitions with other agents, this might negatively affect the likelihood of agreement;
- Secondly, increasing the size of the pre-coalition contributes significantly to increasing the difference and thus increases the likelihood of agreement;
- Of course, the same holds for increased power of the pre-coalition, as discussed earlier in Q_3 .

^[6]See Figure H.8.

^[7]See Figure H.9.

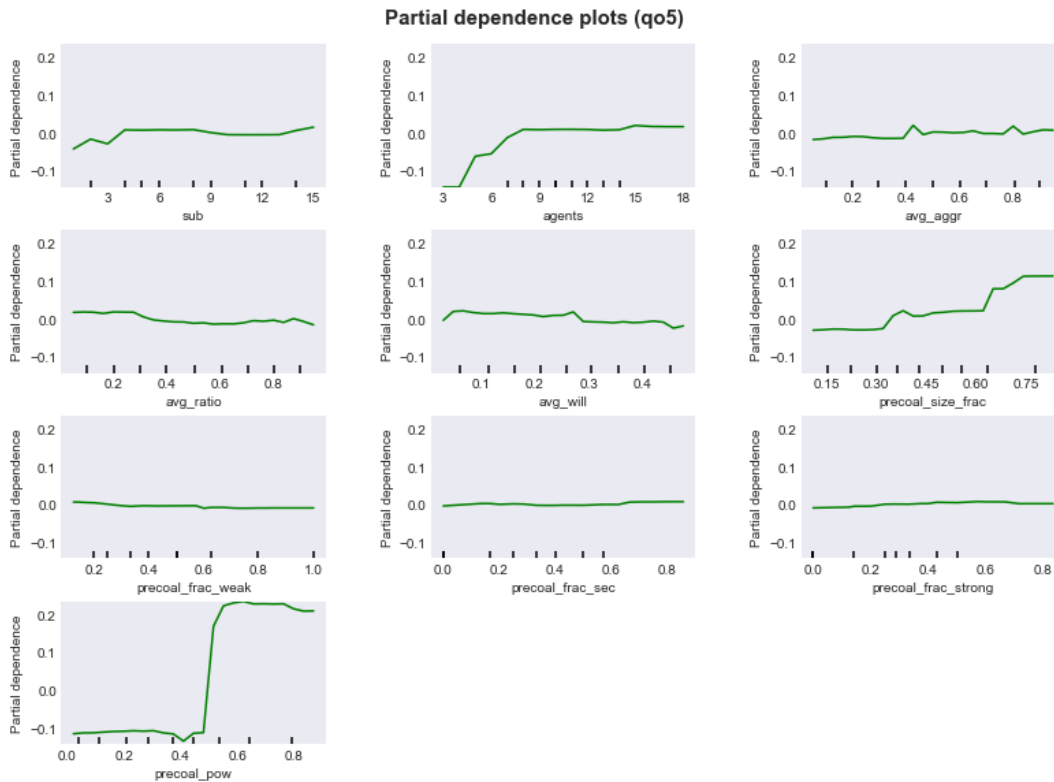


Figure H.9: Weak perspective: single factor partial dependencies for Q_5 .

H.1.4. Partial interactions plots

Figure H.11 to Figure H.15 show interaction plots between characteristics of the pre-coalitions – i.e. size, power and fraction of weak, secondary and strong agents – and overall important factors in the negotiation domain and setting. Again, these plots are created using Python’s *scikit-learn* package. Note that already many plots are excluded from this appendix as it only contains those plots important for further analysis. To illustrate why, note that for Q_2 6 vectors were overall important. Creating interactions with the 5 characteristics of pre-coalitions would already result in 30 different plots. Therefore, we have made a selection.

Factor interactions QO

- The first plot in Figure H.10 shows that increasing the size of the pre-coalition is only beneficial for weaker agents in negotiation settings with more than 11 agents. Then, the fraction of agents in the pre-coalition should at least be above 50%;
- Secondly, expanding the size also increases the difference in QO if on average agents act rational less than 45% of the times;
- Moreover, in situations with many agents (around 15), it becomes beneficial to balance with other weak agents as we observe fractions of weak agents of over 80%;
- However, we also see that from 10 and more agents, we can increase the difference in QO by mainly bandwagoning with strong agents (roughly 60% or more);
- Also, in settings with more than 50% secondary agents, difference increases if weak agents try to bandwagon with between fraction of 30% and 50% strong agents;
- Finally, we observe that increasing the pre-coalition power to above 40% tremendously increases the difference in QO in most negotiation settings and domains, but especially those with high number of subjects (> 5) or agents (> 8), or low willingness to negotiate (< 0.3) and rationality (< 0.5).

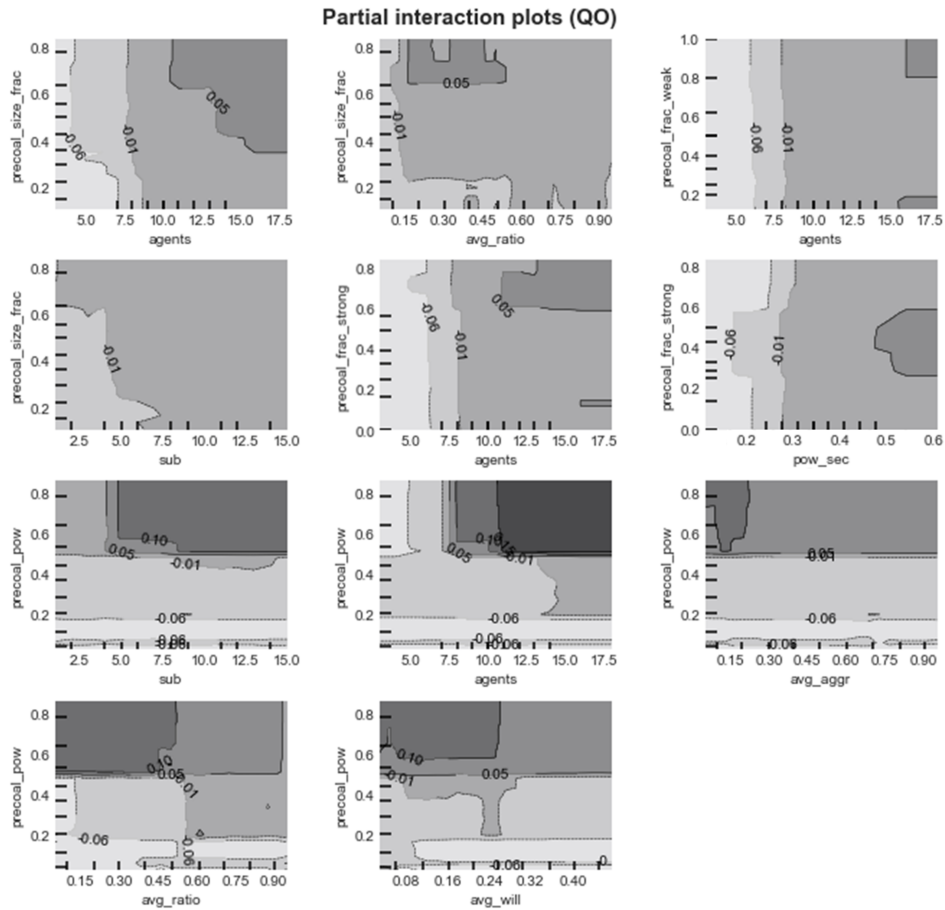


Figure H.10: Weak perspective: factor interaction dependencies for QO.

Factor interactions Q_1 ^[8]

- Observed is that weak agents can increase the social welfare by forming relatively smaller pre-coalitions in negotiations settings with low number of agents;
- The same holds for the power of the pre-coalition;
- We also see that differences in social welfare can be increased if the pre-coalition contains large fractions of secondary agents, especially in combination with low numbers of agents (< 8);
- However, for all these interactions hold that their contribution to the overall difference in QO is relatively small.




Figure H.11: Weak perspective: factor interaction dependencies for Q_1 .

^[8]See Figure H.11.

Factor interactions Q_2 ^[9]

- The (strong) interactions of the relative power of the pre-coalitions with factors in the negotiation setting and domain are very similar to those described in QO ;
- Though, we do notice changes in the interactions between subjects and the pre-coalition size. Only in more complex negotiation domains (subjects > 8) larger pre-coalitions ($> 35\%$ of all agents) become more beneficial for weak agents;
- Also in settings with more agents (> 8) weak agents may improve their personal gains by forming pre-coalitions if the fraction of agents within this coalition exceeds 35%;
- Finally, larger pre-coalitions become beneficial in situations in which agents are eager to negotiate. Presumably because then they do not have the incentive to leave the pre-coalition during the negotiations;

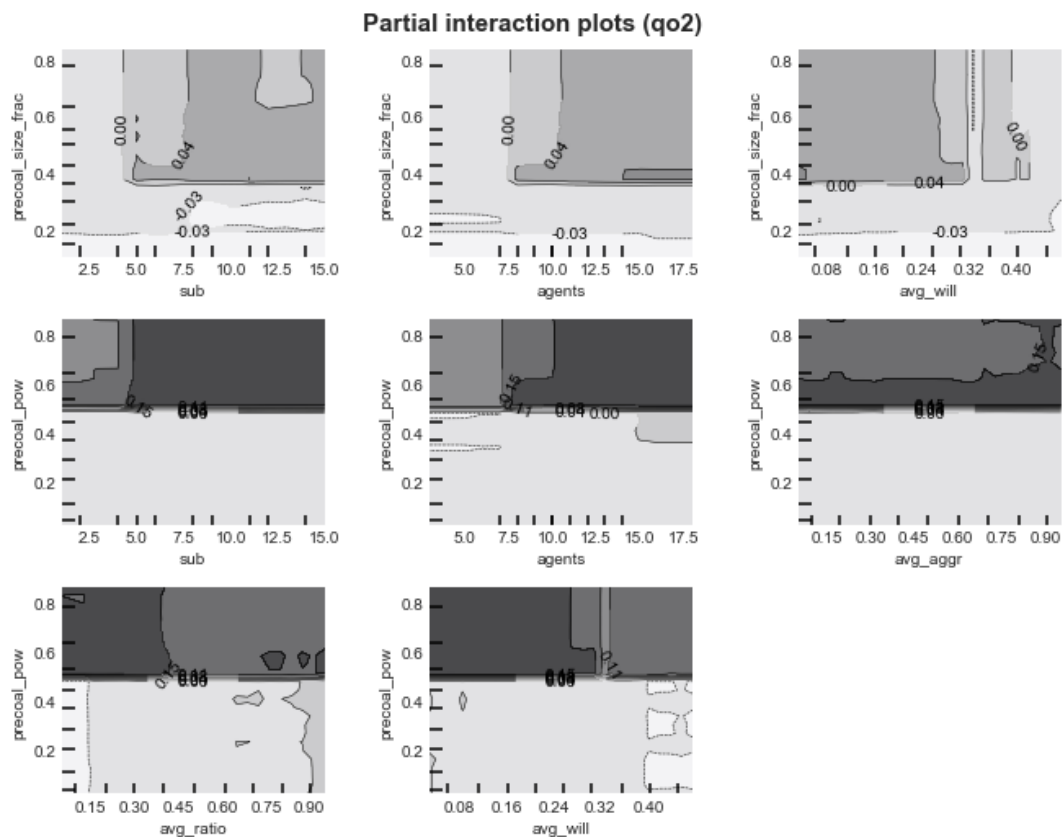


Figure H.12: Weak perspective: factor interaction dependencies for Q_2 .

Factor interactions Q_3 ^[10]

- No noticeable differences with QO for the interactions between the pre-coalition a priori power and negotiation domain and setting. Which implies that differences in the time objective relate to each other in the same way as those described in QO ;

Factor interactions Q_4 ^[11]

- No noticeable differences with QO for the interactions between the pre-coalition a priori power and negotiation domain and setting. Which implies that differences in the political support measure relate to each other in the same way as those described in QO ;

^[9]See Figure H.12.

^[10]See Figure H.13.

^[11]See Figure H.14.

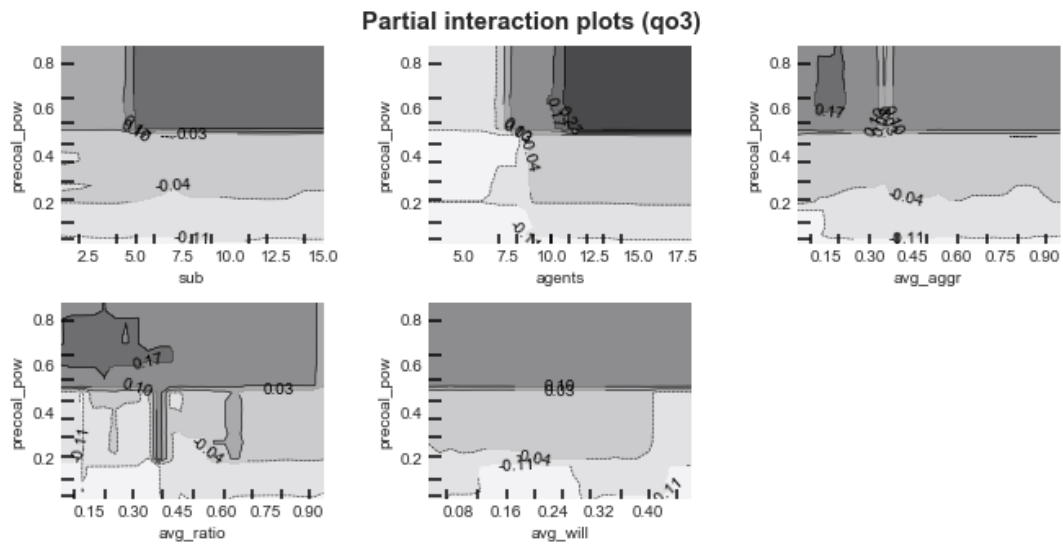


Figure H.13: Weak perspective: factor interaction dependencies for Q_3 .

Factor interactions Q_5 ^[12]

- Larger pre-coalitions increase the likelihood of agreement differences substantially for weaker agents in situations in which the majority of agents are not willing to negotiate;
- The same holds for situations in which the majority of agents do not act rational;

^[12]See Figure H.15.

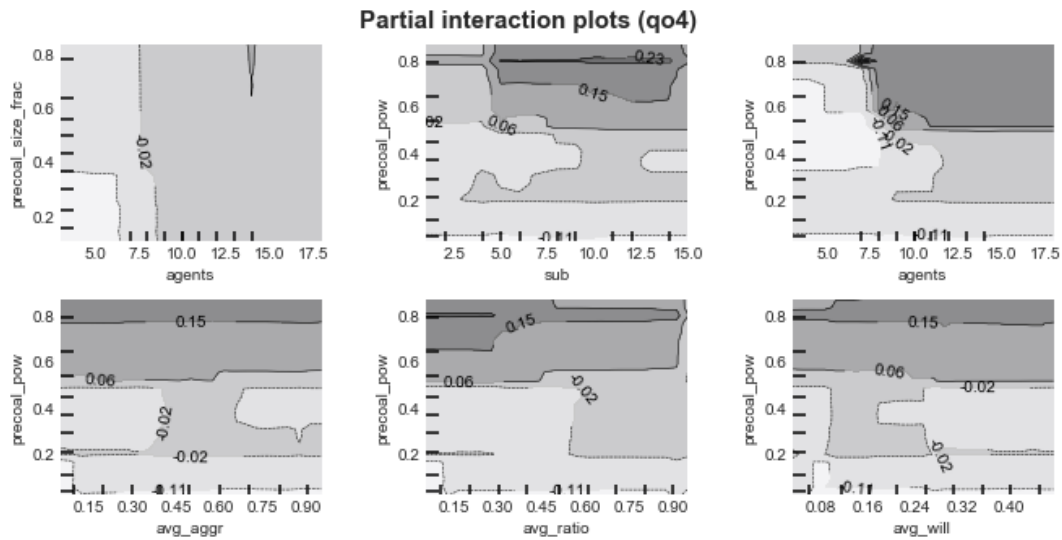


Figure H.14: Weak perspective: factor interaction dependencies for Q₄.

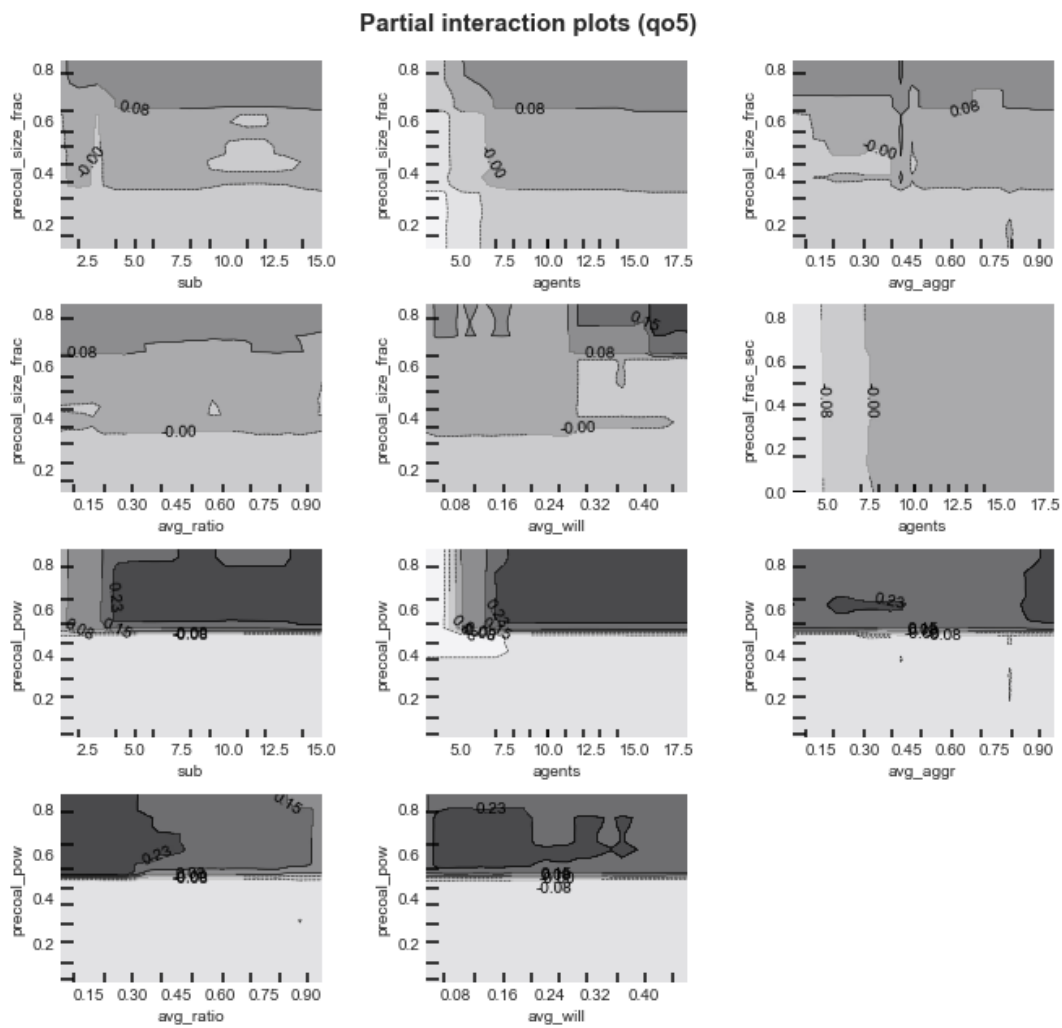


Figure H.15: Weak perspective: factor interaction dependencies for Q₅.

H.2. Perspective of secondary agents

The results in this section are solely devoted to secondary agents.

H.2.1. Correlation between factors

Following the perspective of weaker agents, also for the simulation experiments from the perspective of secondary agents correlations between factors were examined. Figure H.16 shows a pairplot visualizing paired correlation between the 14 factors in the (extended) model. From left to right and top to bottom, the following factors are paired: subjects, issues, fraction power of weak agents, fraction power of secondary agents, fraction power of strong agents, number of agents, average aggressiveness, average rationality, average risk behavior, average willingness to negotiate, number of ideological clusters of agents, size of pre-coalition, fraction weak members, fraction secondary members, fraction strong members and finally, the power of the pre-coalition. Besides, the correlations are also shown in Table H.2.

Again, we observe that Latin Hypercube sampling has done a good job in sampling as evenly as possible from the immense parameter space without (too) much correlation between the different factors. Only those factors in which we expect correlation (by definition) relatively high correlations are detected.

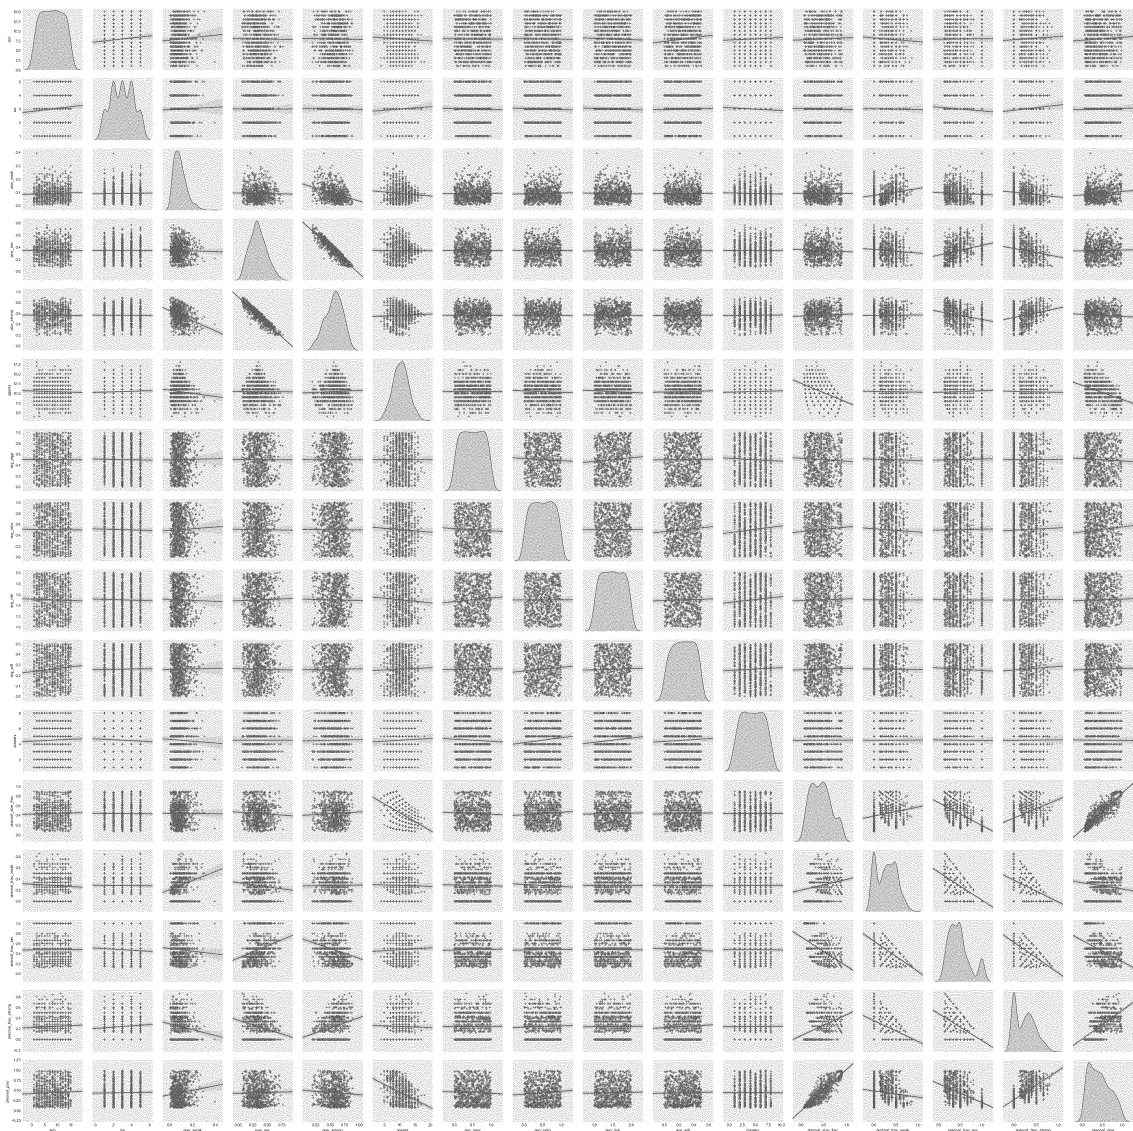


Figure H.16: Pairplot visualizing the correlation between factors in Random Forest model for secondary agents simulations.

Table H.3: Correlations between factors explaining differences in outcome for secondary agents (for conversions, see Table H.1).

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	1.0000	0.0768	0.0435	-0.0072	-0.0094	-0.0326	-0.0031	-0.0079	-0.0234	0.0862	0.0317	0.0280	-0.0514	0.0072	0.0433	0.0479
2	0.0768	1.0000	0.0189	-0.0005	-0.0066	0.0344	-0.0142	-0.0184	-0.0195	0.0099	-0.0337	-0.0088	-0.0082	-0.0535	0.0673	0.0160
3	0.0435	0.0189	1.0000	-0.0231	-0.3504	-0.0920	0.0088	0.0308	-0.0171	0.0085	-0.0318	-0.0043	0.2053	-0.0615	-0.1371	0.1214
4	-0.0072	-0.0005	-0.0231	1.0000	-0.9283	-0.0062	-0.0024	-0.0085	0.0210	-0.0129	-0.0002	-0.0518	-0.0922	0.2812	-0.2182	0.0218
5	-0.0094	-0.0066	-0.3504	-0.9283	1.0000	0.0401	-0.0010	-0.0035	-0.0134	0.0089	0.0120	0.0501	0.0100	-0.2406	0.2554	-0.0656
6	-0.0326	0.0344	-0.0920	-0.0062	0.0401	1.0000	-0.0002	-0.0407	-0.0489	0.0246	0.0278	-0.4248	-0.0191	0.0383	-0.0233	-0.3743
7	-0.0031	-0.0142	0.0088	-0.0024	-0.0010	-0.0002	1.0000	-0.0408	0.0659	0.0082	-0.0376	-0.0410	-0.0206	0.0163	0.0026	-0.0399
8	-0.0079	-0.0184	0.0308	-0.0085	-0.0035	-0.0407	-0.0408	1.0000	0.0025	0.0706	0.0880	0.0641	-0.0462	0.0132	0.0315	0.0795
9	-0.0234	-0.0195	-0.0171	0.0210	-0.0134	-0.0489	0.0659	0.0025	1.0000	0.0226	0.0770	0.0215	-0.0016	-0.0056	0.0077	0.0206
10	0.0862	0.0099	0.0085	-0.0129	0.0089	0.0246	0.0082	0.0706	0.0226	1.0000	0.0390	0.0204	-0.0079	-0.0411	0.0533	0.0217
11	0.0317	-0.0337	-0.0318	-0.0002	0.0120	0.0278	-0.0376	0.0880	0.0770	0.0390	1.0000	0.0018	0.0024	-0.0075	0.0059	-0.0062
12	0.0280	-0.0088	-0.0043	-0.0518	0.0501	-0.4248	-0.0410	0.0641	0.0215	0.0204	0.0018	1.0000	0.1908	-0.5251	0.3888	0.8793
13	-0.0514	-0.0082	0.2053	-0.0922	0.0100	-0.0191	-0.0206	-0.0462	-0.0016	-0.0079	0.0024	0.1908	1.0000	-0.5511	-0.3903	-0.1314
14	0.0072	-0.0535	-0.0615	0.2812	-0.2406	0.0383	0.0163	0.0132	-0.0056	-0.0411	-0.0075	-0.5251	-0.5511	1.0000	-0.5531	-0.4442
15	0.0433	0.0673	-0.1371	-0.2182	0.2554	-0.0233	0.0026	0.0315	0.0077	0.0533	0.0059	0.3888	-0.3903	-0.5531	1.0000	0.6213
16	0.0479	0.0160	0.1214	0.0218	-0.0656	-0.3743	-0.0399	0.0795	0.0206	0.0217	-0.0062	0.8793	-0.1314	-0.4442	0.6213	1.0000

H.2.2. Fitness of Random Forest models

The scatter plots in Figure H.17 and Figure H.18 show the predicted versus actual responses – i.e. quality of outcome QO and its quality measures Q_i – for the factors in "base" and "+pre-coalition" Random Forest models from the perspective of secondary agents. Thus, these plots give us an insight how well the simulation results are explained by the factors in the Random Forest models.

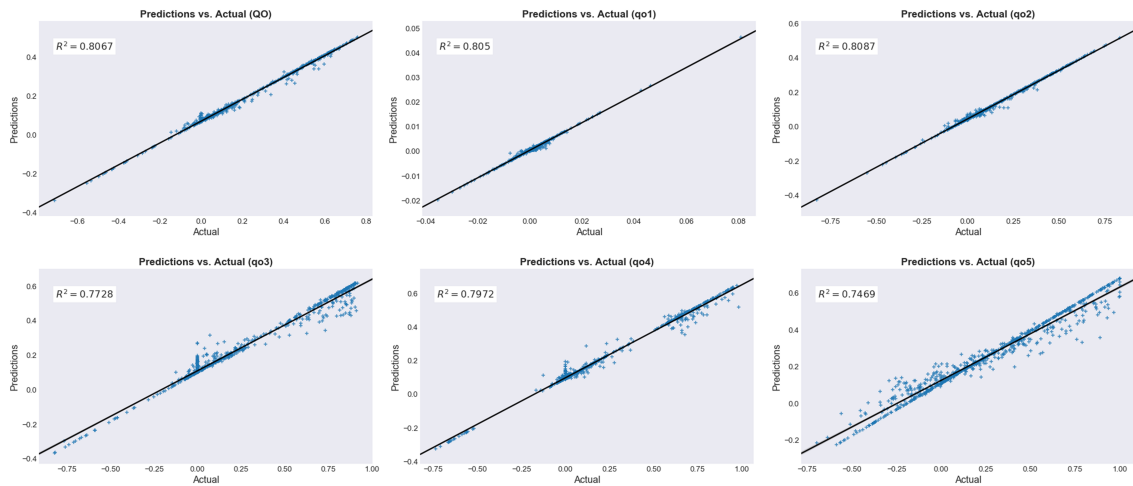


Figure H.17: Secondary perspective: predictions vs. actual for the "base" model.

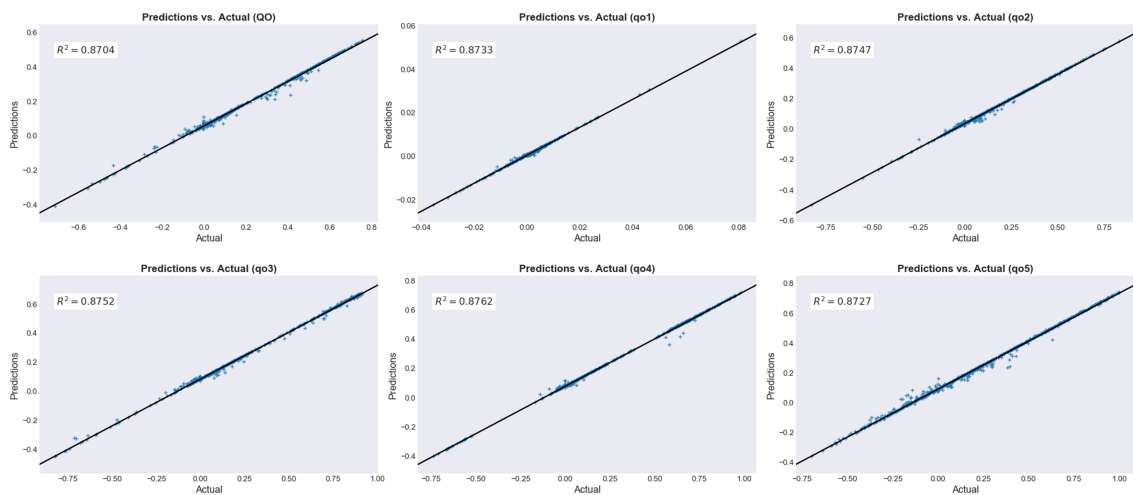


Figure H.18: Secondary perspective: predictions vs. actual for the "+pre-coalition" model.

H.2.3. Partial dependency plots

Figure H.20 to Figure H.24 show the single factor partial dependence plots for the overall important factors in the Random Forest models (see Figure H.5). Positive values on the y -axis indicate pre-coalitions add value to weaker agents given the factor under consideration. Of course, the opposite holds for negative values.

Partial dependence QO

- From Figure H.19 we observe that for secondary agents, forming pre-coalitions in negotiation domains with low numbers of subjects (< 5) negatively influences the difference in outcome. However, in more complex domains (> 9 subjects) it start to positively influence QO ;
- Moreover, pre-coalitions also become more beneficial if the majority of agents act irrational;
- In addition, pre-coalitions with high fractions $> 70\%$ of weak agents negatively influences the differences in outcome;
- Finally, if the power of pre-coalitions is too small ($< 30\%$) differences are negatively influenced. Whereas if the a priori power of pre-coalitions already exceed the majority, a strong positive effect is noticeable;

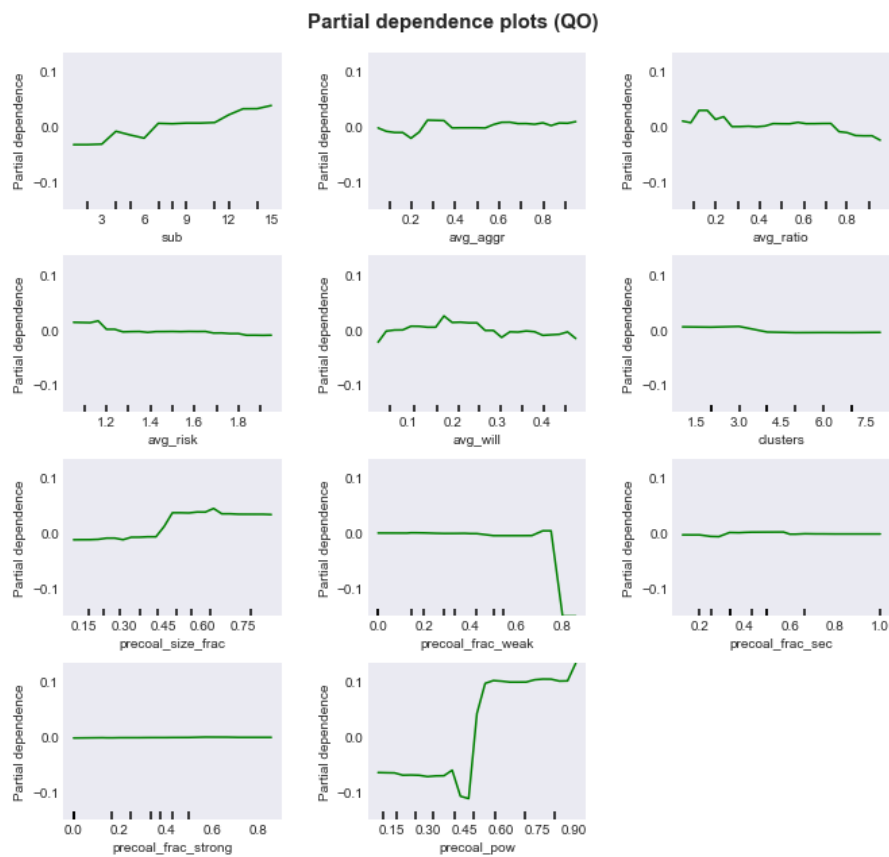


Figure H.19: Secondary perspective: single factor partial dependencies for QO .

Partial dependence Q_1 ^[13]

- Judging the y -scale, these partial dependencies do not explain the differences with great significance;
- Though, if secondary agents form pre-coalitions in negotiations with less than 9 subjects, it negatively (or neutrally) affects the social welfare of all agents;
- Remarkable is that greater sizes of the pre-coalition leads to positive differences in social welfare, whereas increased power has near neutral effects;




Figure H.20: Secondary perspective: single factor partial dependencies for Q_1 .

Partial dependence Q_2 ^[14]

- Obviously, much of the partial relations are similar to QO as Q_2 account for 75% of its explanation;
- Though, if the willingness to negotiate increases, it negatively influences the difference in personal benefits;
- Finally, we notice that if the fraction of weak agents in the pre-coalition increases – especially in the case of very high fractions $> 70\%$, differences are negatively affected.

Partial dependence Q_3 ^[15]

- Little noticeable different dependencies compared to QO . However, we do observe that in situations in which more agents behave risky $\tilde{\beta} < 1.3$, secondary agents benefit from pre-coalitions as the differences in time objective become positive;
- Also, we observe modest positive effects for time objective if on average agents in the negotiation environment tend to negotiate more aggressively $\tilde{p}(\Omega = \text{"Competitor"}) > 0.8$;
- Finally, the (positive) effect of subjects on the differences for time objective is slightly stronger. Especially for very complex domains containing more than 13 subjects;

Partial dependence Q_4 ^[16]

- Little noticeable different dependencies compared to Q_3 . Thus, differences in political support measure follow the same dependencies as time objective for secondary agents;

[13] See Figure H.20.

[14] See Figure H.21.

[15] See Figure H.22.

[16] See Figure H.23.

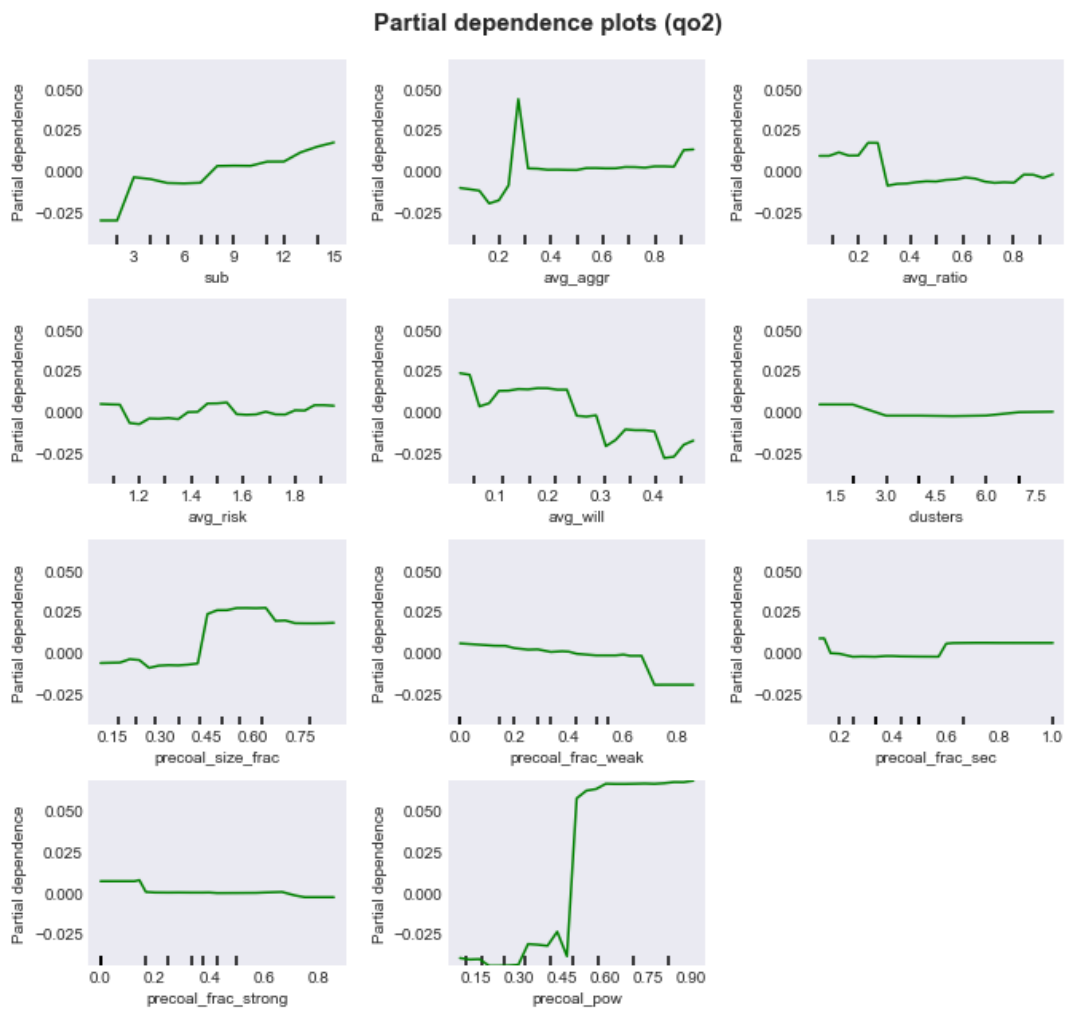


Figure H.21: Secondary perspective: single factor partial dependencies for Q_2 .

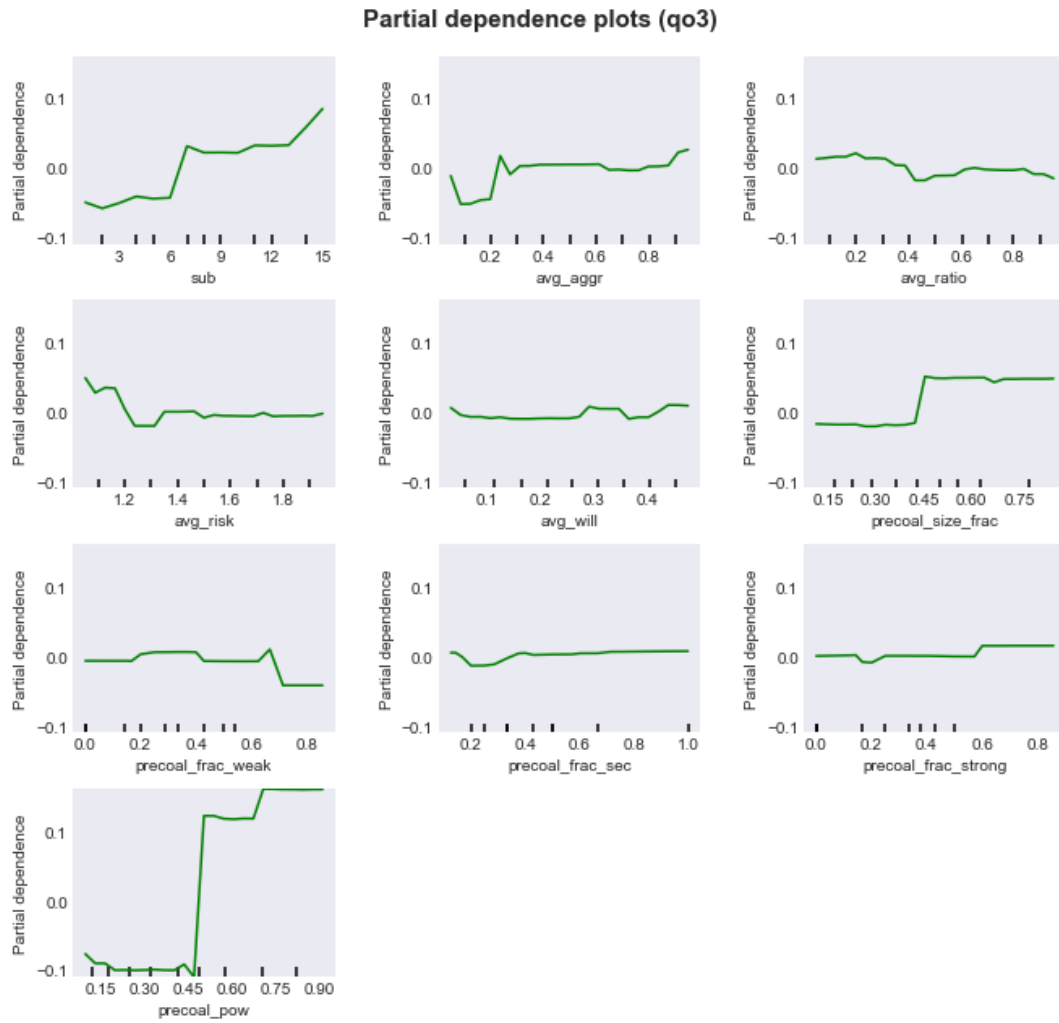


Figure H.22: Secondary perspective: single factor partial dependencies for Q_3 .

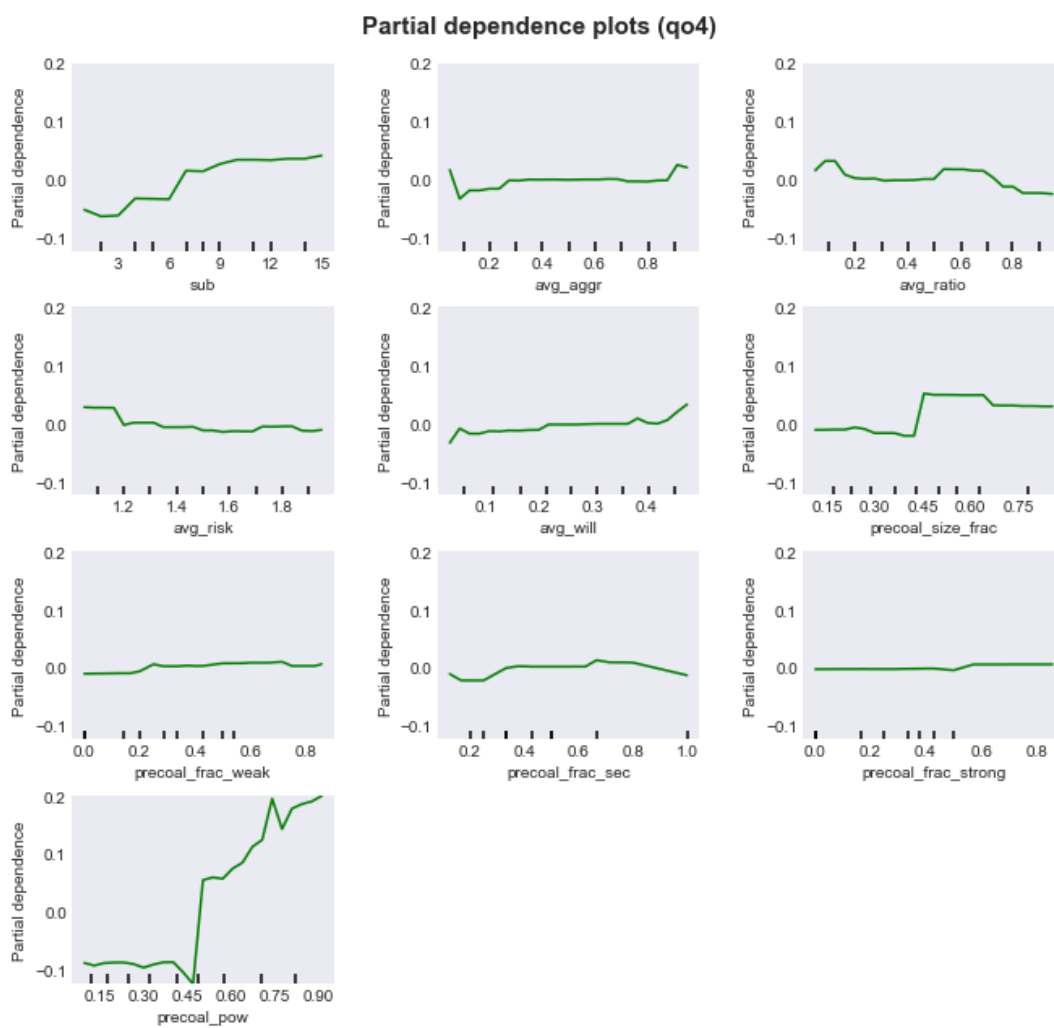


Figure H.23: Secondary perspective: single factor partial dependencies for Q_4 .

Partial dependence Q_5 ^[17]

- Differences in the likelihood of agreement for secondary agents almost solely dependent on a priori power of the pre-coalition. Thus, we see a strong increase in differences for fractions of power of over 45%.




Figure H.24: Secondary perspective: single factor partial dependencies for Q_5 .

H.2.4. Partial interactions plots

Figure H.26 to Figure H.30 show interaction plots between characteristics of the pre-coalitions – i.e. size, power and fraction of weak, secondary and strong agents – and overall important factors in the negotiation domain and setting. Note again that already lots of plots are excluded from this section in the appendix as they do not provide relevant insights.

Interaction dependence QO

- First, we observe interaction between the number of issues and size and pre-coalition size. Larger pre-coalitions positively affect the differences in outcome in more complex negotiation domains;
- In addition, larger sizes add more value in situations with more aggressive agents $\bar{p}(\Omega = \text{"Competitor"}) > 0.3$;
- The opposite is true for rationality. If agents make (irrational) mistakes 80% of the time^[18] larger pre-coalitions add more value to the difference in outcomes for secondary agents;
- Moreover, interactions is presented between number of subjects and the fraction of (other) secondary agents in the pre-coalition. In more complex domains, higher fractions of secondary agents contribute to greater differences;
- Finally, more powerful pre-coalitions add significantly more value in samples in which agents are relatively irrational ($\bar{p} < 0.3$) or reluctant to negotiate (thus if $\bar{d} > 0.2$).

^[17]See Figure H.24.

^[18]Thus have average rationality factor $\bar{p} = 0.2$.

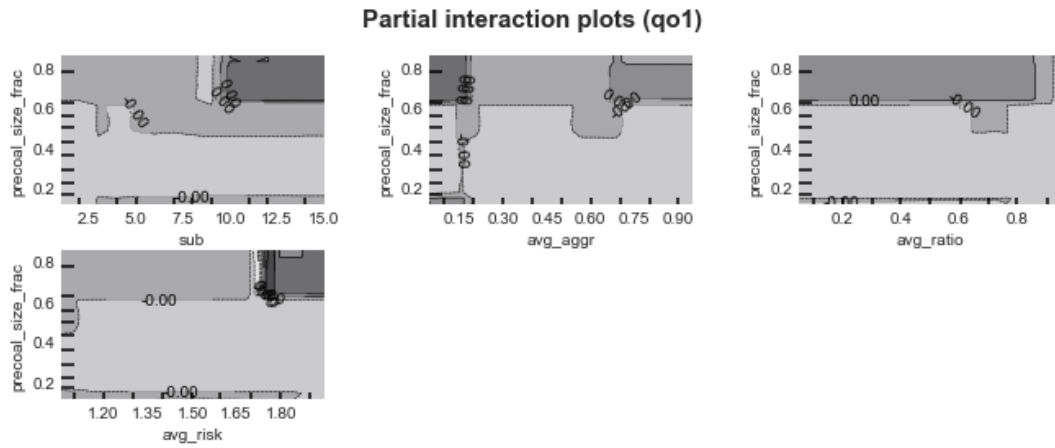


Figure H.26: Secondary perspective: factor interaction dependencies for Q_1 .

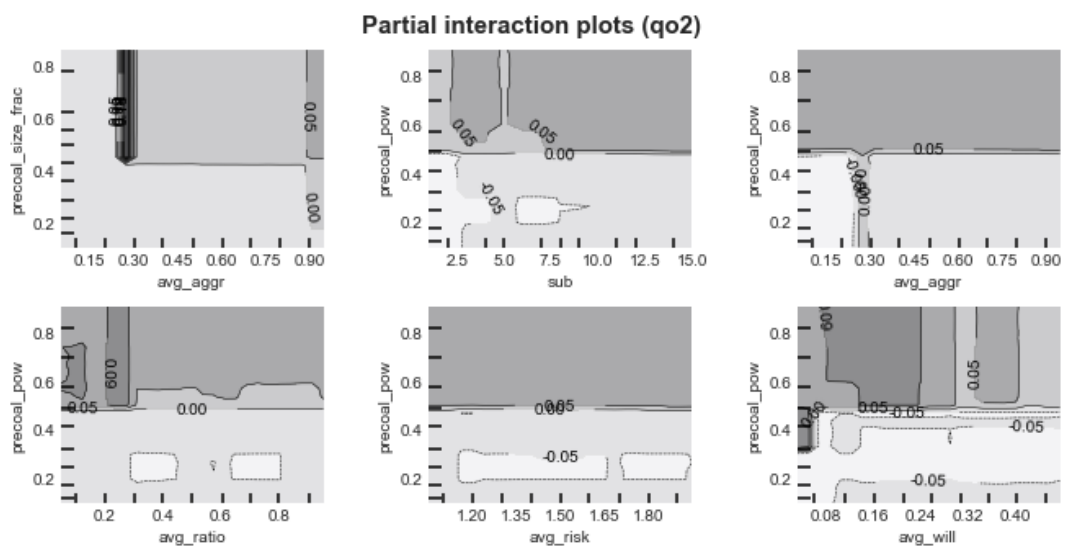


Figure H.27: Secondary perspective: factor interaction dependencies for Q_2 .

Interaction dependence Q_4 ^[22]

- Little noticeable differences compared to Q_3 . The most substantial interaction that differs is the fraction of weak agents interacting with average risk. From the perspective of secondary agents, to increase political support, it slightly benefit from adding more weak agents to their pre-coalition in negotiation settings in which the majority of agent show risky behavior (i.e. $\hat{\beta} < 1.2$).

Interaction dependence Q_5 ^[23]

- From the perspective of secondary agents, to increase the likelihood of agreement, one should not have high fractions of weak agents in the pre-coalition. Especially in negotiation domains with more than 7 subjects;

^[22]See Figure H.29.

^[23]See Figure H.30.

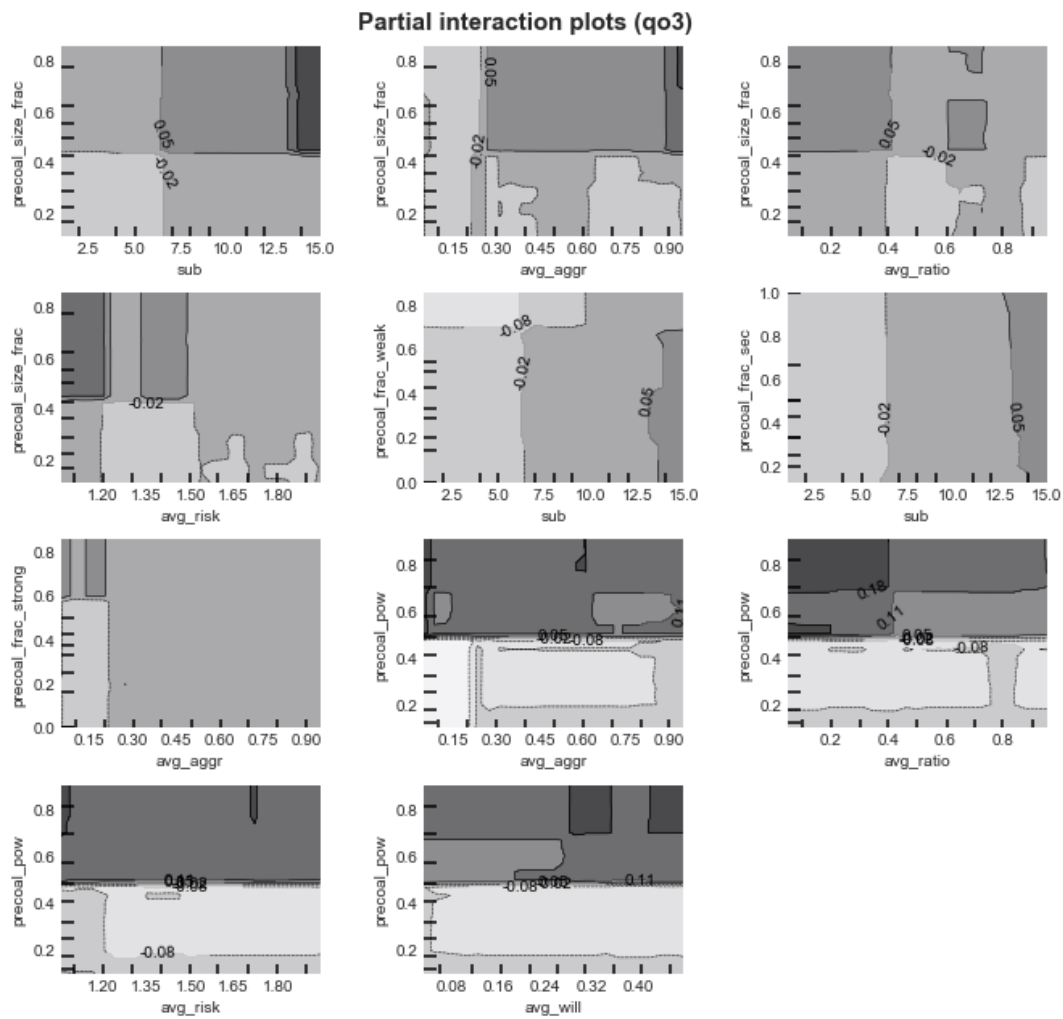


Figure H.28: Secondary perspective: factor interaction dependencies for Q_3 .

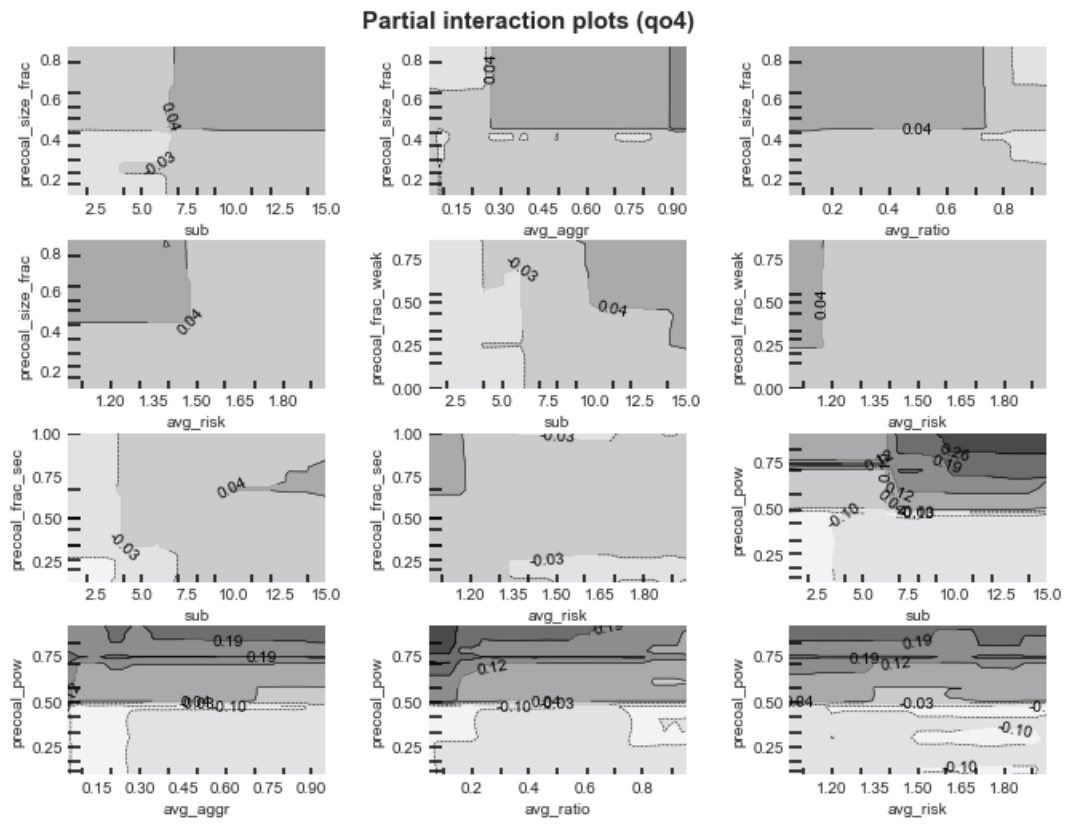


Figure H.29: Secondary perspective: factor interaction dependencies for Q_4 .




Figure H.30: Secondary perspective: factor interaction dependencies for Q_5 .

H.3. Perspective of strong agents

The results in this section are solely devoted to strong agents.

H.3.1. Correlation between factors

Following the perspective of weaker and secondary agents, also for the simulation experiments from the perspective of strong agents correlations between factors were examined. Figure H.16 shows a pairplot visualizing paired correlation between the 14 factors in the (extended) model. From left to right and top to bottom, the following factors are paired: subjects, issues, fraction power of weak agents, fraction power of secondary agents, fraction power of strong agents, number of agents, average aggressiveness, average rationality, average risk behavior, average willingness to negotiate, number of ideological clusters of agents, size of pre-coalition, fraction weak members, fraction secondary members, fraction strong members and finally, the power of the pre-coalition. Besides, the correlations are also shown in Table H.2. This analysis resulted in the same conclusions as for weak and secondary agents and, therefore, further explanation would be superfluous.

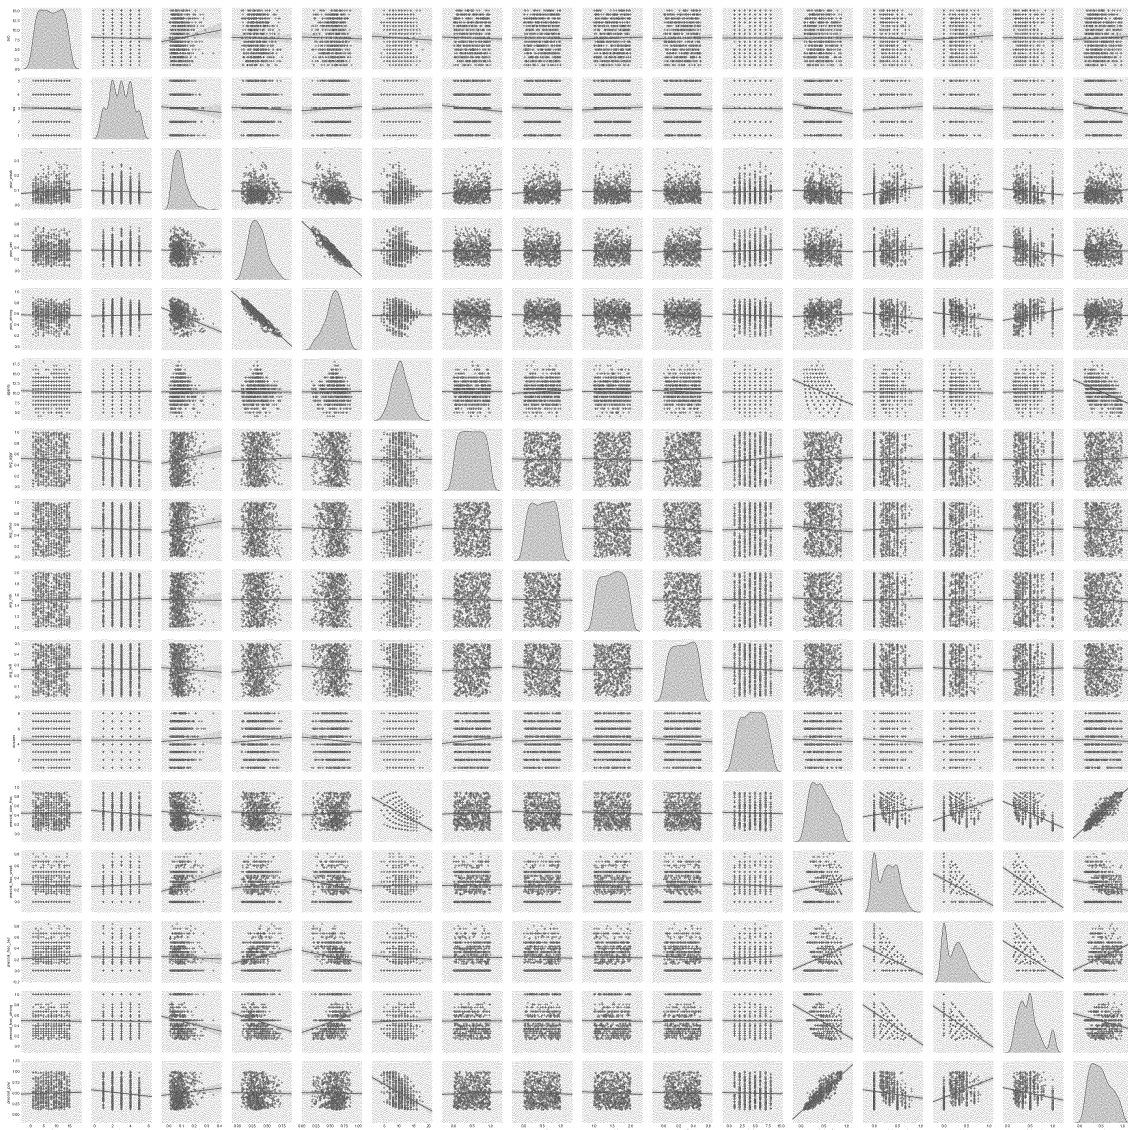


Figure H.31: Pairplot visualizing correlation between factors in Random Forest model for strong agents simulations.

Table H.4: Correlations between factors explaining differences in outcome for strong agents (for conversions, see Table H.1).

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	1.0000	-0.0176	0.0830	-0.0136	-0.0168	-0.0307	-0.0218	-0.0073	0.0188	0.0108	0.0001	0.0089	-0.0280	0.0399	-0.0108	0.0320
2	-0.0176	1.0000	-0.0381	-0.0231	0.0354	0.0133	-0.0628	-0.0196	0.0306	-0.0105	-0.0038	-0.0934	0.0359	-0.0136	-0.0204	-0.1213
3	0.0830	-0.0381	1.0000	-0.0241	-0.3338	0.0111	0.0894	0.0823	-0.0077	-0.0346	0.0265	-0.0541	0.1716	-0.0210	-0.1375	0.0835
4c	-0.0136	-0.0231	-0.0241	1.0000	-0.9343	-0.0134	0.0230	-0.0042	-0.0016	0.0666	0.0457	-0.0356	0.0708	0.1640	-0.2142	0.0835
5	-0.0168	0.0354	-0.3338	-0.9343	1.0000	0.0087	-0.0536	-0.0254	0.0043	-0.0505	-0.0457	0.0529	-0.1472	-0.0389	0.2510	-0.0177
6	-0.0307	0.0133	0.0111	-0.0134	0.0087	1.0000	0.0006	0.0651	-0.0041	-0.0410	0.0149	-0.0063	0.0141	-0.0389	-0.4265	0.0320
7	-0.0218	-0.0628	0.0894	0.0230	-0.0536	0.0006	1.0000	-0.0115	-0.0216	0.0412	0.0699	0.0366	0.0302	-0.0138	0.0226	-0.1213
8	-0.0073	-0.0196	0.0823	-0.0042	-0.0254	0.0651	-0.0115	1.0000	-0.0119	-0.0575	0.0191	-0.0594	0.0302	-0.0080	-0.0037	-0.0131
9	0.0188	0.0306	-0.0077	-0.0016	0.0043	-0.0041	-0.0216	-0.0119	1.0000	0.0122	-0.0043	-0.0399	0.0225	-0.0138	0.0025	-0.0488
10	0.0108	-0.0105	-0.0346	0.0666	-0.0505	-0.0410	0.0412	-0.0575	0.0122	1.0000	-0.0344	-0.0216	0.0195	-0.0445	0.0227	-0.0273
11	0.0001	-0.0038	0.0265	0.0457	-0.0525	0.0149	0.0699	0.0191	-0.0043	-0.0344	1.0000	-0.0246	-0.0376	0.0461	-0.0077	-0.0234
12	0.0089	-0.0934	-0.0541	-0.0356	0.0529	-0.4178	0.0366	-0.0594	-0.0399	-0.0216	-0.0246	1.0000	0.1596	-0.3482	-0.4631	0.8909
13	-0.0280	0.0359	0.1716	0.0708	-0.1280	0.0141	0.0036	0.0302	0.0225	0.0195	-0.0376	0.1596	1.0000	-0.3992	-0.5492	-0.1430
14	0.0399	-0.0136	-0.0210	0.1640	-0.1472	-0.0389	-0.0063	-0.0262	-0.0138	-0.0445	0.0461	0.3482	-0.3992	1.0000	-0.5470	0.3848
15	-0.0108	-0.0204	-0.1375	-0.2142	0.2510	0.0226	0.0025	-0.0037	-0.0080	0.0227	-0.0077	-0.4631	-0.5492	-0.5470	1.0000	-0.2202
16	0.0320	-0.1213	0.0835	-0.0177	-0.0131	-0.4265	0.0430	-0.0427	-0.0488	-0.0273	-0.0234	0.8909	-0.1430	0.3848	-0.2202	1.0000

H.3.2. Fitness of Random Forest models

Again, the scatter plots in Figure H.32 and Figure H.33 show the predicted versus actual responses – i.e. quality of outcome QO and its quality measures Q_i – for the factors in "base" and "+pre-coalition" Random Forest models from the perspective of secondary agents. Thus, these plots give us an insight how well the simulation results are explained by the factors in the Random Forest models.

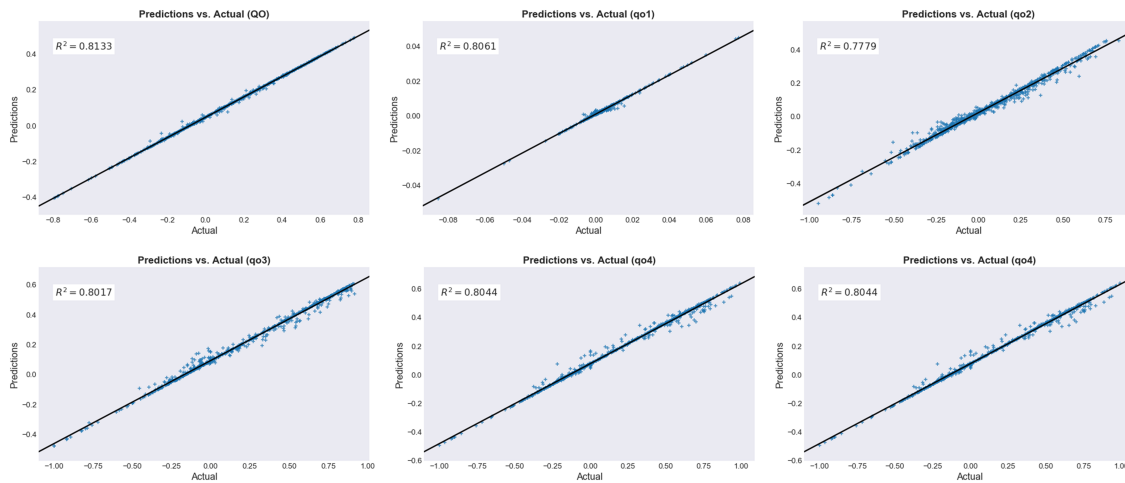


Figure H.32: Strong perspective: predictions vs. actual for the "base" model.

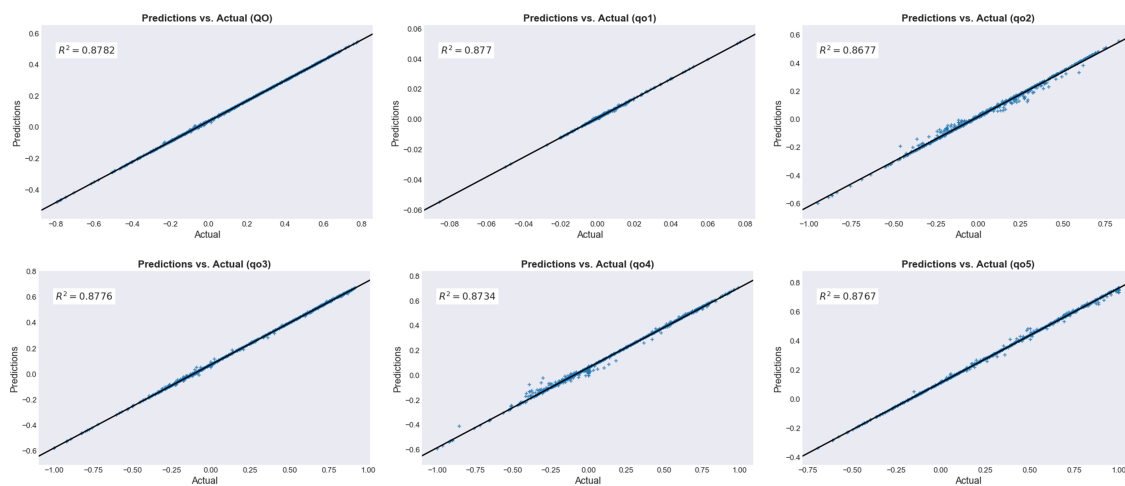


Figure H.33: Strong perspective: predictions vs. actual for the "+pre-coalition" model.

H.3.3. Partial dependency plots

Figure H.35 to Figure H.39 show the single factor partial dependence plots for the overall important factors in the Random Forest models (see Figure H.5). Positive values on the y -axis indicate pre-coalitions add value to weaker agents given the factor under consideration. Of course, the opposite holds for negative values.

Partial dependence QO

- With increasing number of issues in the negotiation domain, we observe a slight positive effect on the difference in QO ;
- We also see that if the fraction of strong agents in the negotiation environment increases, forming pre-coalitions also results in slightly positive differences. Especially for high > 0.7 fractions;
- Average aggressiveness gradually negatively affect the difference in QO ;
- Stronger negative dependencies are observed for the average willingness to negotiate. If agents are willing to negotiate (i.e. $\bar{d} < 0.2$), QO is positively influenced by forming pre-coalitions. However, if $\bar{d} > 0.2$, we observe negative effects;
- In addition, if strong agents form pre-coalitions than their QO is positively influenced if either they choose for high fractions of weak or secondary agents.




Figure H.34: Strong perspective: single factor partial dependencies for QO .

Partial dependence Q_1 ^[24]

- Again, judging the y -scale, partial dependencies for social welfare do not influence the differences in QO substantially;
- Though for small numbers of agents, difference in social welfare is positive from the perspective of strong agents;
- We also see that from social welfare perspective, strong agents should form pre-coalitions with others if the average rationality of agents is relatively low (i.e. $\bar{\rho} < 0.4$);

^[24]See Figure H.35.

- Other single factor dependencies for Q_1 are considered to be neutral, given the y -scale.

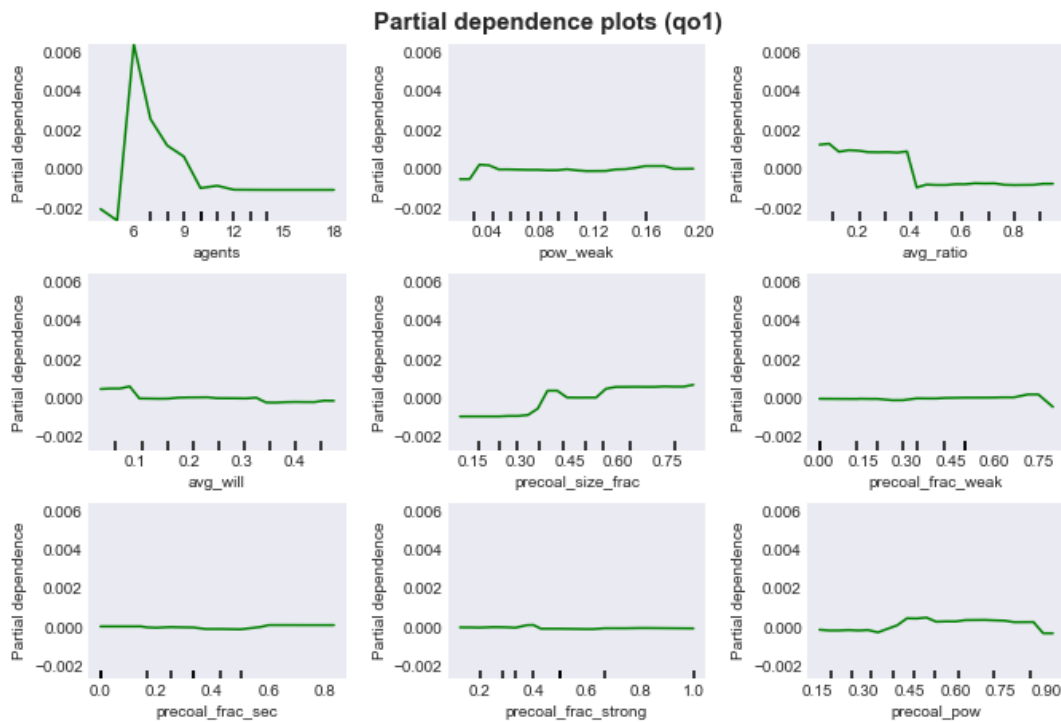


Figure H.35: Strong perspective: single factor partial dependencies for Q_1 .

Partial dependence Q_2 ^[25]

- Obviously, the partial dependencies for Q_2 are similar to Q_0 ;
- Though, we do observe that if either the initial fraction of secondary or (other) strong agents in the negotiation environment increases, it becomes very beneficial for strong agents to form pre-coalitions as the differences in personal gains Q_2 become positive;
- We also observe a limited effect on the differences for Q_2 if agents become less risk seeking. In that setting, forming pre-coalitions influences the difference positively;
- In addition, in terms of personal gains, increasing the fraction of weak agents in the pre-coalition has a slight negative effect;
- Finally, we observe that for strong agents in pre-coalitions covering over 40% of the power in the negotiation environment personal gains are strongly positively influenced;

Partial dependence Q_3 ^[26]

- For more complex negotiation domains, with both higher number of issues or subjects, forming pre-coalitions positively influences strong agents' time objectives;
- Also in negotiation environment with relatively more agents > 9 , the difference in time objective is positively influenced;
- Again, the higher the fraction of power of the pre-coalition for strong agents, the more likely it influences their time objective positively. Especially for fractions over 30%.

^[25]See Figure H.36.

^[26]See Figure H.37.

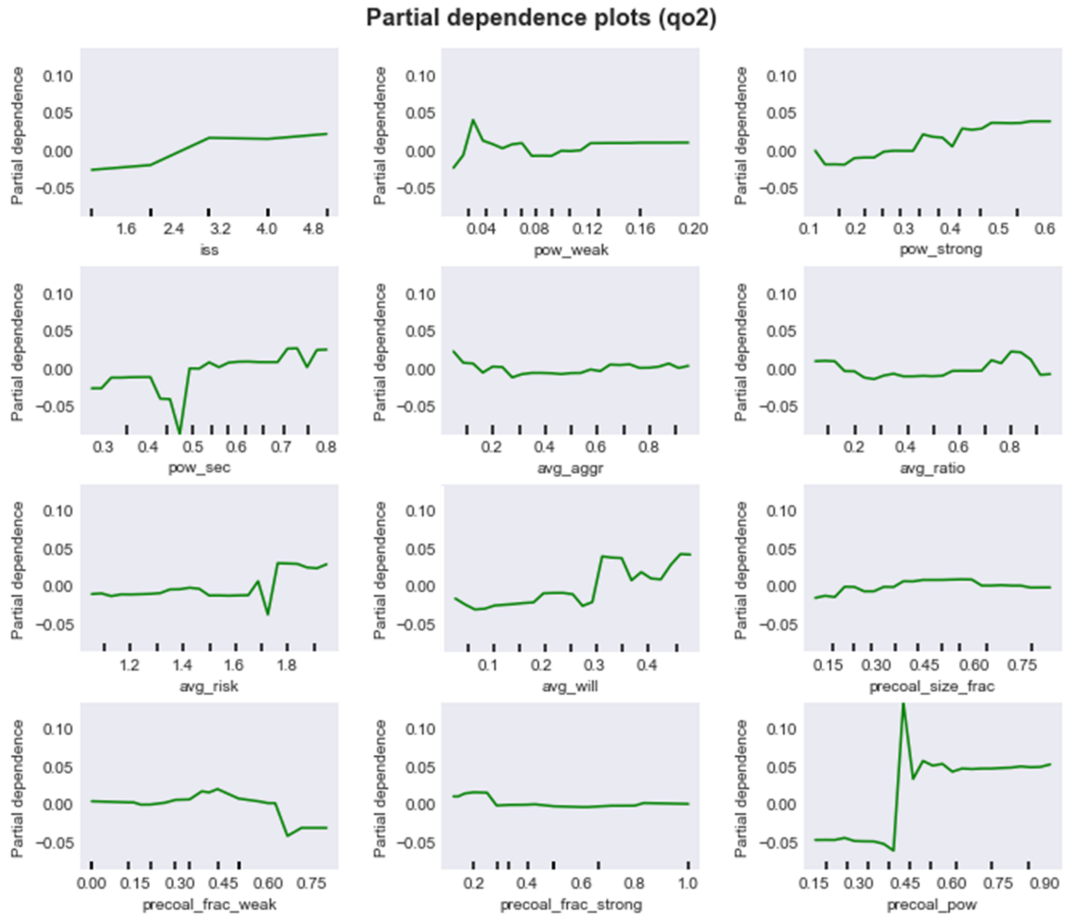


Figure H.36: Strong perspective: single factor partial dependencies for Q_2 .




Figure H.37: Strong perspective: single factor partial dependencies for Q_3 .

Partial dependence Q_4 ^[27]

- Little noticeable different dependencies compared to Q_3 . However, we do observe a negative dependency on average rationality. Apparently, the more rational other agents, the less beneficial it is for strong agents to form pre-coalitions from the perspective of political support;




Figure H.38: Strong perspective: single factor partial dependencies for Q_4 .

Partial dependence Q_5

- No noticeable differences with Q_0 .




Figure H.39: Strong perspective: single factor partial dependencies for Q_5 .

^[27] See Figure H.38.

H.3.4. Partial interactions plots

Figure H.41 to Figure H.45 show interaction plots between characteristics of the pre-coalitions – i.e. size, power and fraction of weak, secondary and strong agents – and overall important factors in the negotiation domain and setting. Note again that already lots of plots are excluded from this section in the appendix as they do not provide relevant insides.

Partial dependence QO

- From the interaction plots in Figure H.40, we observe that QO marginally increases if the size of the pre-coalition at least covers 0.25 of all agents in the negotiation domain if there is limited aggressiveness (i.e. $\tilde{p}(\Omega = \text{"Competitor"}) < 0.45$);
- Moreover, we see that for any fraction of weak or secondary agents in the pre-coalition, strong agents QO improves as the difference increases slightly, as long as the willingness to negotiate is high (or alternatively if $\tilde{d} < 0.16$);
- Finally, it can be stated that strong agents benefit substantially from more powerful pre-coalitions if the number issues per subject is high (> 3), if agents overall are relatively irrational ($\tilde{\rho} < 0.2$) or if the average willingness of agents to negotiate is high $\tilde{d} < 0.20$;

Partial dependence Q_1 ^[28]

- No noticeable interactions as they behave monotonically;

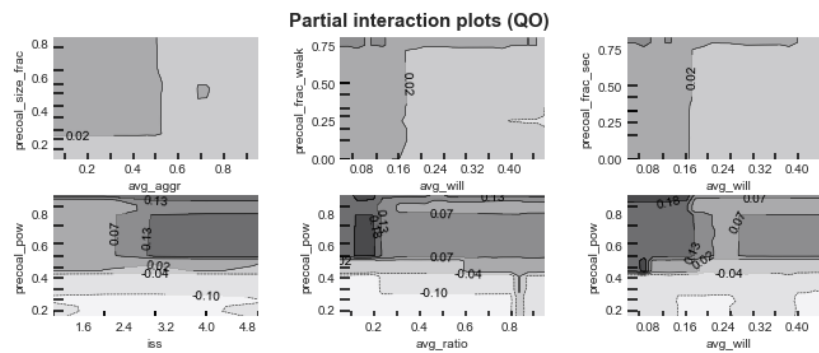


Figure H.40: Strong perspective: factor interaction dependencies for QO .




Figure H.41: Strong perspective: factor interaction dependencies for Q_1 .

[28] See Figure H.41.

Partial dependence Q_2 ^[29]

- Differences in personal gains for strong agents increase if they opt for lower fractions of weak agents (< 0.45) in their pre-coalition if there are higher fractions of secondary agents in the negotiation environment (> 0.4);
- If average willingness to negotiate is high (or alternatively if $\tilde{d} < 0.2$), strong agents' personal gains increase if they embrace around 50% of weak agents in their pre-coalitions;
- In addition, if the willingness to negotiate is again high, they should opt for very low fractions of other strong agents in the pre-coalition ($\ll 0.25$ including themselves);
- Finally, strong agents benefit from more powerful pre-coalitions in terms of personal gains if the negotiation domain contains many issues (> 4), either the average rationality or risk is low ($\bar{\rho} < 0.2$ or $\bar{\beta} > 1.5$ respectively), fraction of power of secondary agents is relatively high (< 0.6) or if the willingness to negotiate is high ($\tilde{d} < 0.2$).

Partial dependence Q_3 ^[30]

- Similar results for more powerful pre-coalitions are obtained for Q_3 . To be precise, in more complex negotiation domains with both high number of subjects or issues (> 8 and > 3 respectively) more powerful pre-coalitions are required for strong agents to result in positive differences;
- The same holds for more agents in the negotiation environment > 8 . Though, for small numbers of agents in the negotiation environment, the pre-coalition should cover less than 25% of the overall power to benefit strong agents in terms of time objectives;

Partial dependence Q_4 ^[31]

- The interaction results in terms of political support for the power of pre-coalitions are similar to Q_3 ;
- In addition, we see that differences in political support increase the size of the pre-coalition increases for larger fractions of secondary agents in the negotiation environment;

Partial dependence Q_5 ^[32]

- No noticeable differences with Q_0 .

^[29]See Figure H.42.

^[30]See Figure H.43.

^[31]See Figure H.44.

^[32]See Figure H.45.

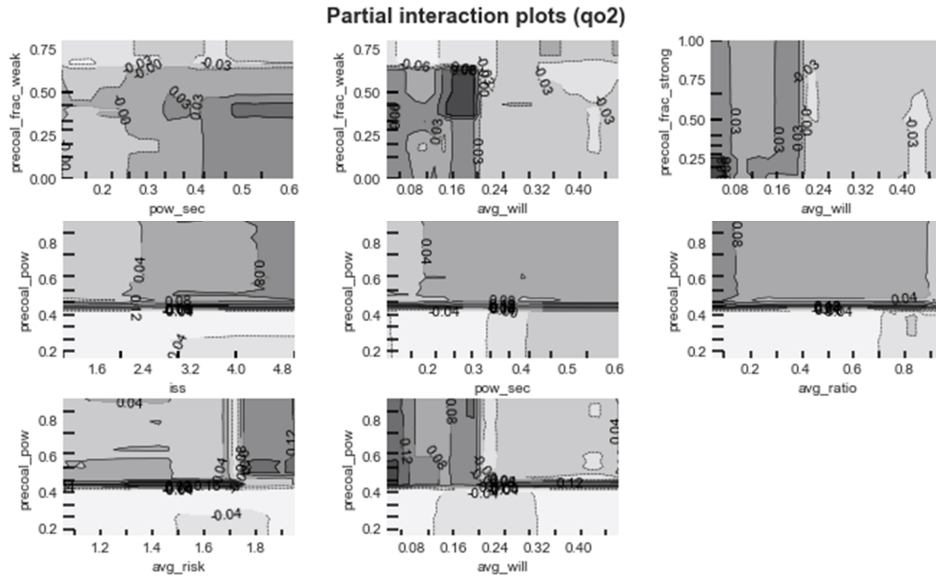


Figure H.42: Strong perspective: factor interaction dependencies for Q_2 .

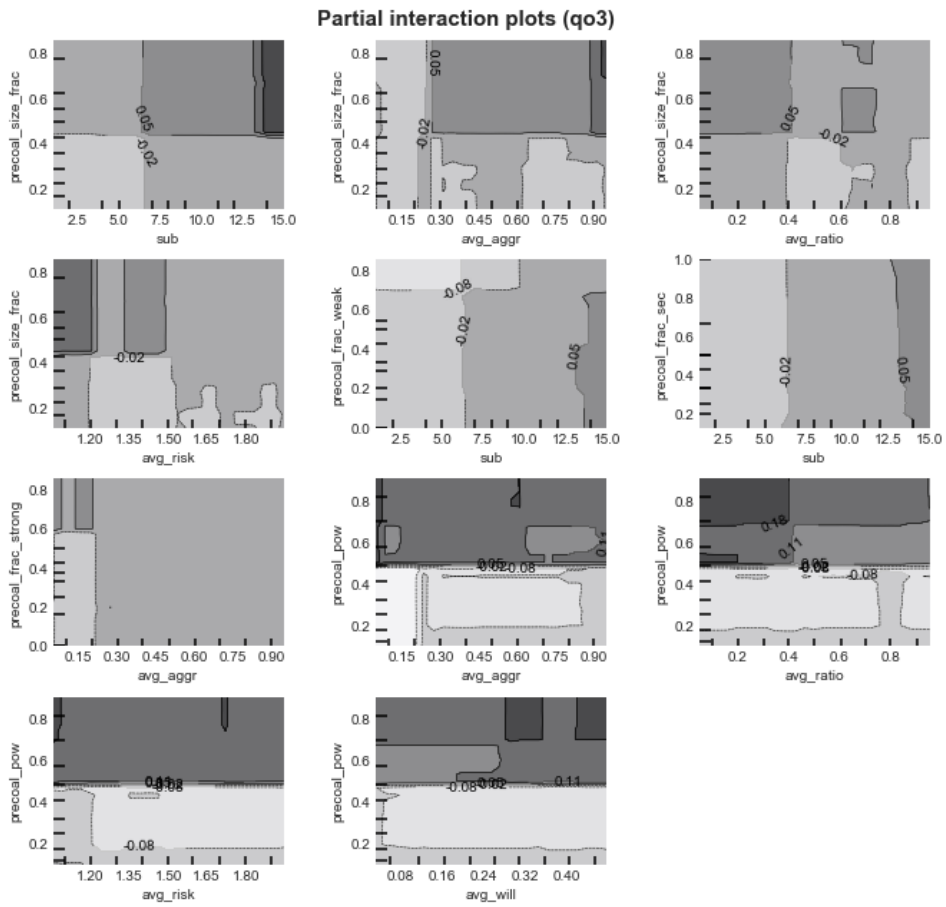


Figure H.43: Strong perspective: factor interaction dependencies for Q_3 .

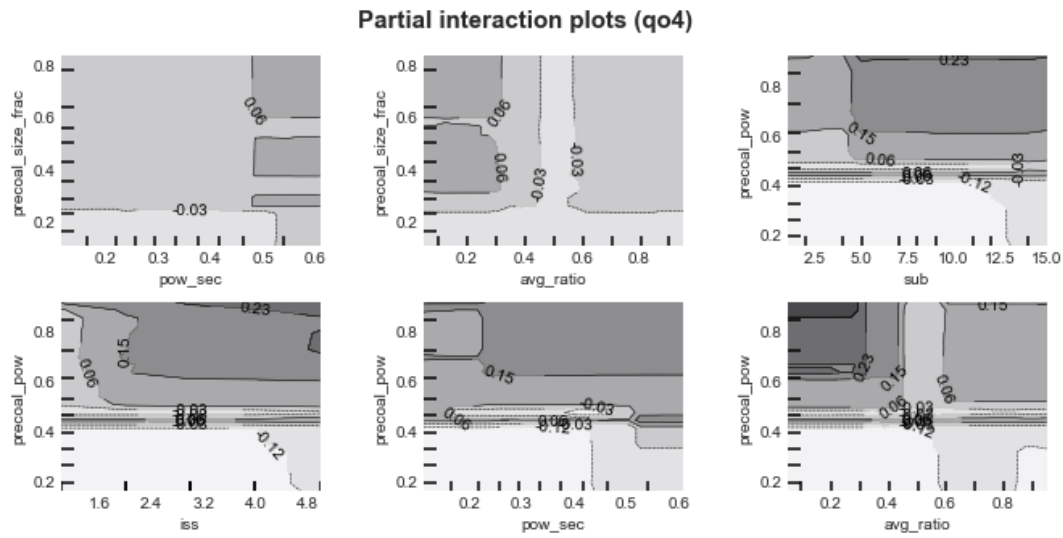


Figure H.44: Strong perspective: factor interaction dependencies for Q_4 .

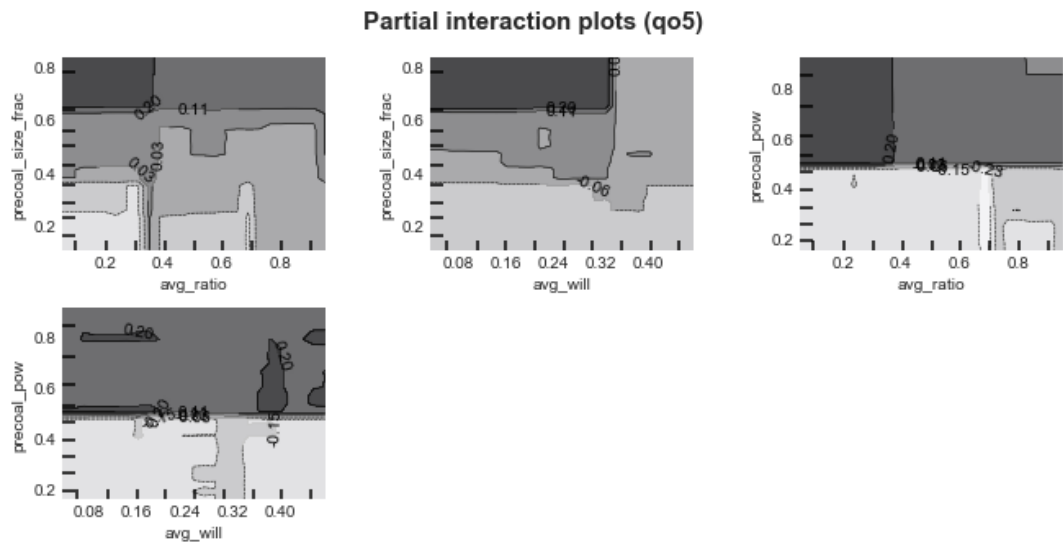


Figure H.45: Strong perspective: factor interaction dependencies for Q_5 .

Hasta la vista, baby.
Terminator 2: Judgment Day (1991)

 **TU Delft**
Deloitte.