# Guidance Under Uncertainty: Employing a Mediator Framework in Bilateral Incomplete-Information Negotiations

by

James Shew

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2008

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

# Abstract

Bilateral incomplete-information negotiations of multiple issues present a difficult yet common negotiation problem that is complicated to solve from a mechanism design perspective. Unlike multilateral situations, where the individual aspirations of multiple agents can potentially be used against one another to achieve socially desirable outcomes, bilateral negotiations only involve two agents; this makes the negotiations appear to be a zero-sum game pitting agent against agent. While this is essentially true, the gain of one agent is the loss of the other, with multiple issues, it is not unusual that issues are valued asymmetrically such that agents can gain on issues important to them but suffer losses on issues of less importance. Being able to make trade-offs amongst the issues to take advantage of this asymmetry allows both agents to experience overall benefit. The major complication is negotiating under the uncertainty of incomplete information, where agents do not know each other's preferences and neither agent wants to be taken advantage of by revealing its private information to the other agent, or by being too generous in its negotiating. This leaves agents stumbling in the dark trying to find appropriate trade-offs amongst issues.

In this work, we introduce the Bilateral Automated Mediation (BAM) framework. The BAM framework is aimed at helping agents alleviate the difficulties of negotiating under uncertainty by formulating a negotiation environment that is suitable for creating agreements that benefit both agents jointly. Our mediator is a composition of many different negotiation ideas and methods put together in a novel third-party framework that guides agents through the agreement space of the negotiation, but instead of arbitrating a final agreement, it allows the agents themselves to ratify the final agreement.

# Acknowledgments

First of all, I would like to thank my supervisor, Kate Larson. Her ability to stay calm under pressure, and more importantly, her ability to keep me calm under pressure, was invaluable in bolstering my confidence when I needed it the most. When I was searching for a research topic, the freedom and support she gave me to explore was astounding. As a supervisor, she is a spectacular mediator of the graduate student negotiation process. :)

I would also like to thank my readers, Robin Cohen and Chrysanne DiMarco. Their comments improved this thesis significantly, making it both easier to read and understand.

I would also like to thank the guys in the AI lab: Greg Hines, Jie Zhang, Nick Miller, Reid Kerr, Fred Kroon, and Tyrel Russell. Admittedly, I was quite the hermit while I was in Waterloo, but when I did venture into the AI lab it always led to interesting discussions and a good way to take a break.

Lastly, I would like to thank Annie Lee and the rest of the CCF hockey team. Hockey was a really important part of my graduate life and I am not sure how well I would have made it through the last year without it. We didn't win a lot of games, but I couldn't have asked for a better set of teammates, on and off the ice. In particular, Annie not only took the time to listen when I felt like complaining about graduate life, but she was also the one that said, "You know, I have a friend that runs a hockey team..."

# Dedication

First of all, I want to dedicate this work to my family; their love and support were always unwavering.

Saving the best for last, I also want to dedicate this thesis to my wife, Katherine. Katherine, you are my life and my world; I cannot remember what life was like without you and I cannot fathom going forward in it without being at your side. This thesis marks the end of a long time apart from each other, and in that time the patience, love, understanding, and support you gave me where beyond measure and definitely more than I deserved. This has been an experience infinitely enriched by your love, and I want to thank you for this and everything you have brought into my life.

# Contents

**List of Figures** **xi**

**List of Tables** **xv**

**1 Introduction** **1**

1.1 Our Philosophy . . . . . . . . . . . . . . . . . . . . . . . . . . 3

1.2 Contributions . . . . . . . . . . . . . . . . . . . . . . . . . . . 4

1.3 Outline of the Thesis . . . . . . . . . . . . . . . . . . . . . . . 5

**2 Background** **7**

2.1 Multi-agent Systems (MASs) . . . . . . . . . . . . . . . . . . . 7

2.1.1 MASs and Negotiation . . . . . . . . . . . . . . . . . . . 8

2.2 Preferences, Utility, and Uncertainty . . . . . . . . . . . . . . . 9

2.2.1 Utility Space . . . . . . . . . . . . . . . . . . . . . . . . 10

2.3 Mechanism Design and Game Theory . . . . . . . . . . . . . . . 11

2.3.1 Important Game-Theoretic Properties . . . . . . . . . . . 12

2.3.2 Myerson-Satterthwaite Bilateral Trading Theorem . . . . . . 14

2.4 Negotiation . . . . . . . . . . . . . . . . . . . . . . . . . . . . 16

2.4.1 Third Parties: Mediation and Arbitration . . . . . . . . . . 16

2.4.2 Single Negotiating Texts (SNTs) . . . . . . . . . . . . . . 17

2.4.3 Agendas . . . . . . . . . . . . . . . . . . . . . . . . . . 19

2.4.4 Solution Concepts . . . . . . . . . . . . . . . . . . . . . . 20

**3 Related Work** **31**

3.1 The Art and Science of Negotiation . . . . . . . . . . . . . . . . . . . 31

3.2 Arbitration and Mediation in Game Theory . . . . . . . . . . . . . . 32

3.2.1 Arbitrators in Games . . . . . . . . . . . . . . . . . . . . . . 32

3.2.2 Mediators in Games . . . . . . . . . . . . . . . . . . . . . . . 32

3.3 Agendas and Incremental Negotiation . . . . . . . . . . . . . . . . . . 35

3.3.1 Agendas . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 35

3.3.2 Incremental Negotiations . . . . . . . . . . . . . . . . . . . . 36

**4 Constraining the Solution Space** **37**

4.1 Constraints of Intermediary Agreements . . . . . . . . . . . . . . . . 37

4.2 Adding Stages . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 39

**5 The Bilateral Automated Mediator Negotiation Framework** **43**

5.1 Negotiation Setting . . . . . . . . . . . . . . . . . . . . . . . . . . . 44

5.2 BAM Negotiation Framework . . . . . . . . . . . . . . . . . . . . . . 48

5.3 Our Arbitration Module . . . . . . . . . . . . . . . . . . . . . . . . . 52

5.3.1 Greedy Punishment Solution (GPS) . . . . . . . . . . . . . . 52

5.3.2 Valuation Chips Mechanism . . . . . . . . . . . . . . . . . . . 59

5.3.3 Summary of Arbitration Module . . . . . . . . . . . . . . . . 67

5.4 Our Protocol Module . . . . . . . . . . . . . . . . . . . . . . . . . . 67

5.4.1 Translating Issue Orderings to Utility Space . . . . . . . . . . 69

5.4.2 Agenda Preliminaries . . . . . . . . . . . . . . . . . . . . . . 71

5.4.3 Agenda Creation . . . . . . . . . . . . . . . . . . . . . . . . . 76

5.4.4 Example of Choosing an Agenda . . . . . . . . . . . . . . . . 78

5.4.5 Alternative Considerations for Agendas . . . . . . . . . . . . 79

5.5 Example . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 80

5.5.1 Arbitration Module Example . . . . . . . . . . . . . . . . . . 82

5.5.2 Protocol Module Example . . . . . . . . . . . . . . . . . . . . 85

5.6 Conclusion . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 92

**6 Experimental Setup and Results** **93**
6.1 Negotiation Details . . . . . 93
6.1.1 Issue Setup . . . . . 93
6.1.2 Agenda Setup . . . . . 94
6.1.3 Offer Procedure . . . . . 94
6.2 Agents . . . . . 97
6.3 Results . . . . . 97
6.3.1 Metrics . . . . . 98
6.3.2 Data . . . . . 98
6.4 Discussion . . . . . 102

**7 Extensions of the System** **105**
7.1 Randomised Revelation Process . . . . . 105
7.1.1 The Method . . . . . 105
7.1.2 Questions . . . . . 111
7.1.3 Conclusion . . . . . 112
7.2 Valuation Chips Improvements . . . . . 113
7.3 Agendas . . . . . 115
7.3.1 Comprehensive Quality Measures: DRAB . . . . . 115
7.3.2 Further Agenda Extensions . . . . . 122

**8 Conclusion** **127**

**Appendices** **129**

**A The Number Of Agendas** **129**
A.1 Amount of Computation . . . . . 131
A.2 Computational Reductions . . . . . 132
A.2.1 Preliminary Discussion . . . . . 132
A.2.2 Dynamic Programming . . . . . 137

**Bibliography** **140**

# List of Figures


2.1 Examples of utility spaces . . . . . . . . . . . . . . . . . . . . . . . 11
2.2 Examples of efficient frontiers . . . . . . . . . . . . . . . . . . . . . 12
2.3 Advantages of negotiating with a Single Negotiating Text . . . . . . 18
2.4 Separating utility space according to an agenda . . . . . . . . . . . 20
2.5 Comparing agendas with and without renegotiation . . . . . . . . . 21
2.6 The invariance of affine transformation axiom . . . . . . . . . . . . 23
2.7 The independence of irrelevant alternatives (IIA) axiom . . . . . . . 25
2.8 The Kalai-Smorodinsky solution concept . . . . . . . . . . . . . . . 27

3.1 The Prisonner's Dilemma . . . . . . . . . . . . . . . . . . . . . . . . 33
3.2 The mediated Prisonner's Dilemma . . . . . . . . . . . . . . . . . . 34

4.1 Natural constraints in a utility space . . . . . . . . . . . . . . . . . 38
4.2 Constraining utility spaces with intermediary agreements . . . . . . 39
4.3 Constrained spaces with agendas and SNTs . . . . . . . . . . . . . . 41

5.1 An example of trade balances . . . . . . . . . . . . . . . . . . . . . 47
5.2 The BAM framework . . . . . . . . . . . . . . . . . . . . . . . . . . 49
5.3 The interactions of the arbitration module. . . . . . . . . . . . . . . 53
5.4 An example for explaining the continuous GPS mechanism . . . . . 55
5.5 Diagram for the proof of ex-ante IR of the GPS mechanism . . . . . 59
5.6 Relative valuations . . . . . . . . . . . . . . . . . . . . . . . . . . . 60
5.7 Summary of arbitration module properties . . . . . . . . . . . . . . 68

5.8 Converting preference orderings to utility space . . . . . . . . . . . 70
5.9 The solution line . . . . . . . . . . . . . . . . . . . . . . . . . . . 72
5.10 Explaining desired solutions . . . . . . . . . . . . . . . . . . . . . . 74
5.11 Explaining expected solutions . . . . . . . . . . . . . . . . . . . . . 75
5.12 The distance quality measure for agendas . . . . . . . . . . . . . . . 78
5.13 Choosing an agenda . . . . . . . . . . . . . . . . . . . . . . . . . . 79
5.14 Example with original pricing . . . . . . . . . . . . . . . . . . . . . 81
5.15 Example of converted prices . . . . . . . . . . . . . . . . . . . . . . 82
5.16 Declarations for the example . . . . . . . . . . . . . . . . . . . . . . 83
5.17 Estimated utility space . . . . . . . . . . . . . . . . . . . . . . . . . 87
5.18 Utility spaces for two possible agendas . . . . . . . . . . . . . . . . 88
5.19 Expected solutions in two agendas . . . . . . . . . . . . . . . . . . . 91

6.1 The offer procedure in our experimental setup . . . . . . . . . . . . 96
6.2 Benefits of getting further in the agenda . . . . . . . . . . . . . . . 100

7.1 Weak and strong declarations . . . . . . . . . . . . . . . . . . . . . 106
7.2 Creating declaration objects . . . . . . . . . . . . . . . . . . . . . . 107
7.3 An example of declaration objects for all the issues . . . . . . . . . 108
7.4 Declaration object pairing . . . . . . . . . . . . . . . . . . . . . . . 108
7.5 An example of inter-agent pairing of declaration objects . . . . . . . 109
7.6 Rewarding believed declarations . . . . . . . . . . . . . . . . . . . . 111
7.7 A simple valuation points system . . . . . . . . . . . . . . . . . . . 114
7.8 Assessing agenda risk by stage size . . . . . . . . . . . . . . . . . . 118
7.9 Examples of the contentiousness of issues . . . . . . . . . . . . . . . 119
7.10 Issue trade-offs . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 120
7.11 Calculating issue asymmetry and balance . . . . . . . . . . . . . . . 121
7.12 Adding a feedback loop in the BAM framework . . . . . . . . . . . 123
7.13 Using the current agreement in agenda re-evaluation . . . . . . . . . 124

7.14 Agenda trees . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 126

A.1 A graph of the number of agendas . . . . . . . . . . . . . . . . . . . . . 131
A.2 Building agendas by stages . . . . . . . . . . . . . . . . . . . . . . . . . 133
A.3 Agenda example without combining overlapping nodes . . . . . . . 134
A.4 Agenda example with common nodes combined . . . . . . . . . . . 135
A.5 The number of possible different stages in an agenda . . . . . . . . 136
A.6 The final agenda as a graph traversal of stage nodes . . . . . . . . . 139

# List of Tables


6.1 Experiment failure rates . . . . . . . . . . . . . . . . . . . . . . . . 99

6.2 Average stage reached . . . . . . . . . . . . . . . . . . . . . . . . 99

6.3 Percentage of maximum utility . . . . . . . . . . . . . . . . . . . . . 101

6.4 Percentage of NBS utility . . . . . . . . . . . . . . . . . . . . . . . 101

6.5 Improvement on the initial agreement . . . . . . . . . . . . . . . . . 102

7.1 The pairings of different belief types . . . . . . . . . . . . . . . . . 111

A.1 The number of possible agendas . . . . . . . . . . . . . . . . . . . . 130

A.2 The savings of combining common stage nodes . . . . . . . . . . . . 137

# Chapter 1

# Introduction

Negotiations are a cornerstone of human interaction, ranging from everyday interpersonal interaction to global interaction among nations. We are probably all familiar with trying to decide which movie or restaurant to go to amongst a group of friends, hearing about company mergers and takeovers, and at a larger scale, the negotiation processes of conflicting groups in areas such as the Middle East or the attempts of nations to come up with a global agreement for the reduction of greenhouse-gas emission. It is easy to see then that negotiations are central to human society, and the resolution of conflict ranges from simplistic to the impossible.

To understand why negotiations can be difficult, it helps to understand why negotiations exist at all. First of all, two or more parties realise that they can, or must, work together to come up with some resolution over a number of issues.[1] Parties negotiating an agreement on a set of issues have two major problems that we are concerned with: (1) they do not know the preferences of the other parties, and (2) the preferences of the parties are not aligned (in other words, they have conflicting interests). Let us consider an example:

> Two friends, S and B, need to decide on where to have dinner. S strongly prefers any sort of Japanese food and really wants to try the restaurant "Sushi Town", but is willing to go to an Italian restaurant if that is the only way an agreement can be made. B on the other hand, has no real preferences except that, unknown to S, B does not want to go to the restaurant "Sushi Town" because the last time he went there he was ill afterwards.

[1] "Issue" is used in a very broad sense here, representing any parameters that could be negotiated, tangible or intangible.

> At the start of their decision process they both have a general idea of what each other likes because they are friends and have been to dinner together many times. S having a strong preference and knowing that B is generally not very picky, suggests that they have Japanese food and B agrees to that. B, wanting to avoid “Sushi Town”, takes preemptive action in suggesting that they go to “Sushi Village”. However, since S does not know about B’s previous illness and wants to try “Sushi Town”, he makes a counter-offer by suggesting “Sushi Town”. B then strongly objects mentioning his illness and says if they want Japanese food it must be somewhere besides “Sushi Town”. S then agrees to go to “Sushi Village” since he would rather have some Japanese food then none at all.

What the example illustrates is a situation where the parties interests are not completely incompatible, but they must negotiate in order to bring about a solution that reasonably satisfies each of their private interests while still allowing them to collaborate together. The negotiation process allows them a method to search through the combined space of their preferences in order to find a solution to their problem of finding a place to eat. An example with greater conflict of interests could be made if B was averse to all Japanese food due to his illness. In this case, the friends probably would have gone to an Italian restaurant but the negotiation process may have lasted longer since S may continue to try to convince B to go to a Japanese restaurant (maybe enumerating several different restaurants) since that is what he strongly prefers.

In most real-world negotiations, we have a situation similar to the example; the parties have a general idea of what the other party prefers but do not know exactly what those preferences are, and at the same time there is some conflict between the preferences of the parties. Consider a labour negotiation with only the single issue of worker wages. The employer would prefer to pay as little as possible since that means more profits and the workers want as high a wage as possible. Naturally, both parties are aware of these preferences, but neither side knows the “breaking-point” of the other, the wages that would be absolutely unacceptable to a party; these points are known as the *reservation prices* of the parties. Not knowing the preferences of the other party, is where the great difficulty lies since both parties want to further their own interests but they do not want to be so greedy that a deal cannot be made, but at the same time, want to get as much as they can from the negotiation. Without knowing the preferences of the other party it is difficult for either party to understand what is going to be too greedy or too selfless.

However, much of the research in negotiations occurs with what is known as "complete information". This means that all the parties know all the preferences of the other parties. This is usually not very realistic, but makes the analysis of negotiations a lot more manageable. In our work, we try to maintain a negotiation environment that is realistic and that is why the concern of this work is bilateral, multi-issue negotiations of *incomplete information*. "Bilateral" means that there are two parties, as opposed to more than two (multi-lateral). "Multi-issue" is self-explanatory meaning that there are multiple issues over which the two parties negotiate. "Incomplete information" means that there is some information about preferences that are private to the parties that hold those preferences; it may be the case that a party knows some information about its counterpart, but it does not know everything.

To ameliorate the problems of negotiation under uncertainty, we employ the use of a third party. A third-party intervenor can operate as an unbiased shareholder in the negotiation process, able to do things that neither party trusts the other to do. Sometimes it takes an outside party to diffuse tense negotiations, make a negotiation proceed smoothly, or achieve an efficient outcome that is superlative to any that the parties themselves may produce. Our third-party framework will help the negotiating parties to help themselves by guiding them through the negotiation process, encouraging the production of results that are beneficial for both parties.

## 1.1 Our Philosophy

With as much conflict as there can be in some negotiation processes, negotiations are still a *collaborative* venture [22]. We would like to see parties working together to come up with a strong agreement for both of them instead of establishing an adversarial environment of getting the better of each other. If each party is too focused on what it can get out of the other, this can easily cause the breakdown of the negotiation or result in an agreement that allows joint gains to go unrealised.

Our goal in introducing a third-party is to direct the negotiating parties' attention towards building up agreements that are progressively better for both parties. By guiding the parties to explore areas of the solution space that contain likely candidates for joint gain, we hope to keep wasted opportunities to a minimum. The viewpoint we take is that parties need to be actively searching for these joint gains to get the full benefit from them.

## 1.2 Contributions

Our contribution in this work is two-fold: we create a novel third-party mediation framework for bilateral incomplete information scenarios, as well as provide a negotiation tool that can be used in real-life situations. The former contribution consists of a general framework that other researchers can instantiate to fit their needs. The latter contribution consists of a particular instantiation of the mediation framework that can be used to assist in real negotiations.

As a mediation framework, our model encourages new views on how to approach negotiation research. The framework embodies the idea of acting as a guide or facilitator for the negotiating parties and not a dictator of a solution. It is difficult to find quality solutions in incomplete information situations and our framework acknowledges this by trying its best to set up a negotiation environment that will be beneficial to the both parties, but ultimately re-involves them in the finalisation process. It is this paradigm that we provide in a concrete framework for others to explore and approach negotiation research from a different vantage point.

As a fully instantiated system, our model can be an aid in actual negotiations. Being a mediator that can handle general negotiations, it is applicable to many different situations. However, due to the generality of our scope and the complexity of real-world negotiating, we suggest that our system be used as a tool and not as a replacement for whatever current negotiation process parties may be using. As a tool, our system can be a simulation that negotiating parties can use to get a better grasp of how to analyse and understand their negotiation. The benefits of such a simulator are numerous, some examples being:

- Helping parties better understand their actual preferences for the issues (parties not being completely sure of their preferences can be a major problem in negotiations).
- Giving early warning of where potential roadblocks may exist and allowing exploration of possible solutions.
- Providing a starting agreement (the output of the simulation) that the parties may criticise and refine.

Even just as an "ice-breaker" to push through the initial uncertainty of how to negotiate is an important use of our system. Being a full simulation, our system can give parties a tutorial in the analysis of their negotiation before they participate in the real negotiation process.

## 1.3 Outline of the Thesis

**Chapter 2** Relates the necessary background material in negotiations and game theory needed throughout the thesis.

**Chapter 3** Discusses the related research that inspired our framework and contextually sets up other developments in later chapters.

**Chapter 4** Describes how the use of negotiation tools can be used to manipulate the negotiating parties' attention in the solution space.

**Chapter 5** Explores the heart of our work, the Bilateral Automated Mediator framework, or BAM for short; the general framework is described as well as our latest instantiation of it.

**Chapter 6** Describes the simulation that implements our instantiation of the BAM framework and some results from our experimentation.

**Chapter 7** Examines possible extensions of different areas of our framework that could be explored in the future.

**Chapter 8** Is the conclusion of the body of the thesis and reviews the highlights of the framework and our work.

**Appendix A** Goes into detail about some of the computational aspects of our use of agendas, giving an idea of how much work our system actually needs to perform when choosing an agenda.

# Chapter 2

# Background

This chapter outlines the preliminary knowledge that will be useful in later chapters. In Section 2.1, we introduce the idea of agents and multi-agent systems, then in Section 2.2, preferences of agents are discussed using utility theory, then Section 2.3 explores necessary ideas we will draw on from game theory and mechanism design, and the final section, Section 2.4, examines the necessary background knowledge we will need from the field of negotiations.

## 2.1 Multi-agent Systems (MASs)

From this point forward, when we discuss the parties of a negotiation, we will refer to them as *agents*. At its most basic, an agent is an entity that can perceive the world it is situated in and perform actions within that world [25]. This definition can encompass humans as agents, but also allows for agents to be artificially created entities, which will often be the case in electronic settings.[1] In the field of Computer Science, one often has the need to make representations of ideas, artifacts, and actors from the real world; agents can be thought of as the representation of the actors within a given setting.

When we need to consider the creation of agents and how they should act, we enter into the realm of Artificial Intelligence. With respect to agent design, artificial intelligence studies how to determine the details of implementing the behaviour of agents. It is not unusual that this behaviour be informed by the preferences of a user that is employing the agent as its representative to act on the user's behalf

[1] In our work, we allow that agents might be human or electronic.

in some electronic setting. Being able to choose actions that best implement the desired behaviour of an agent's creator is a property of *rational* agents. A rational agent is one that chooses the best actions to implement its goals based on the information it has. When having to deal with situations involving uncertainty, a rational agent should choose actions that maximise its expected benefit (however this benefit may be measured) [25].

As we move into a setting where an agent is interacting with other agents and not just an unintelligent environment, we begin to discuss *Multi-agent Systems* (MASs). A MAS is just that, a system consisting of multiple agents that need to interact with each other to satisfy their personal or collective goals. In our work, we are mainly interested in MASs where the agents are *self-interested*, meaning each agent is trying to satisfy its personal goals but is not concerned with whether or not other agents are satisfying their goals (unless that is a part of an agent's goals). Being self-interested does not mean that an agent acts with malice towards other agents, but it is a form of selfishness such that an agent is only concerned with its own goals; the presence of other agents only becomes a consideration insofar as those agents can help or hinder an agent's personal goals. *Game Theory* and *Mechanism Design* (Section 2.3) are fields of research often employed in artificial intelligence when agents have to start considering strategic interaction with other agents [25]. Negotiation is an essential part of the strategic interaction within a MAS [3].

### 2.1.1 MASs and Negotiation

Negotiation is one of the most important and critical interaction methods for multi-agent systems [10]. It has been well-studied and a few of the major areas of research have been: collective decision-making and consensus building, automation of negotiations, strategies for agent behaviour, and the role and consequences of different negotiation processes.

One focus in MAS negotiations centres around agents trying to influence other agents in order to come to a mutually acceptable agreement; this was particularly prevalent with respect to argumentation as negotiation [14, 10, 15]. In the earlier years of the research focused on consensus building, game theory was already becoming an important tool for negotiation research [4, 10]. There was also, and still is, a lot of focus on the strategies of agents within different negotiation settings [5, 6, 17], hence placing a greater focus on the self-interested aspect of agents.

Our work does not feature consensus building; that is an area of research typified by multilateral (more than two agent) societal settings and considerations. Our research focuses much more on negotiation instances between two agents, although these smaller individual interactions could be drawing agents from a multilateral society of agents. Our work is also not concerned with the strategies of agents. Instead, we are focused on the manipulation of negotiation processes to create desirable results (based on whatever arbitrary goals we may have) and on a form of negotiation automation.

We introduce a third-party into the negotiations that guides agents by developing a negotiation process that tries to help the agents achieve jointly beneficial agreements. Whereas much of the negotiation process research is concerned with looking at specific environments determining what the optimal behaviours will be [5], we keep our basic environment very general. However, based on the input we receive from each particular negotiation scenario, we create negotiation parameters that effectively change the environment for the agents. With our work, the negotiation environment changes a little bit between each separate negotiation.

The procedure of changing the negotiation environment is encapsulated within an automated mediator. Our automation differs from the automation research because we are automating our third-party to construct a negotiation procedure for the agents. In contrast, much existing automation research refers to automation as choosing correct protocols for agents using particular strategies or choosing agent strategies for particular negotiation protocols [5, 10, 6]. Although our automated third-party does, essentially, choose protocols and parameters for the agents to use, we introduce new types of protocols and ways of constructing them. Our protocols draw on novel ideas introduced in this work, but are also influenced by existing tools in negotiations research and existing ideas and theories from game theory and mechanism design (Section 2.3) that are already widely used in negotiation research.

## 2.2 Preferences, Utility, and Uncertainty

Agents will have preferences over different situations, and in order to quantify these preferences we use the idea of utility. Agents receive a certain amount of utility, that is based on their preferences, for outcomes or states in negotiations. Utility is a way to measure the value of an outcome for an agent; a higher utility agreement is more valued than a lower utility agreement.

**Notation 1 (Utility)** *The utility for an agent, $a$, for a particular state or outcome in the world, $S$, will be denoted $u_a(S)$. Abusing notation, we will refer to* *anything that renders utility to an agent by $u_a(i)$, where $i$ can be anything that renders utility, possibly indirectly (this could be something like taking a particular action that leads to a state that actually renders utility).*

Self-interested agents will want to maximise their utility by behaving in a manner that yields the highest utility outcomes or states in the world. In our negotiation setting, this means making offers that lead to the highest utility agreements possible. But what do agents do in the face of uncertainty (for example, if some element of chance exists that affects the outcome of a situation)? In this case, the rational agent should try to maximise its expected utility given the probabilities of the possible outcomes and the utility valuation the agent has for those outcomes. The expected utility is naturally calculated as follows:

**Definition 1 (Expected Utility)** *For an agent $a \in A$(the set of all agents), a set of outcomes $O$, and a probability distribution over the outcomes $P$, the expected utility of the agent is $E(u_a) = \sum_{o \in O} P(o)u_a(o)$.*

### 2.2.1 Utility Space

We will frequently represent the utility of agents geometrically:

**Definition 2 (Utility Space)** *For $n$ agents, a* utility space *is an $n$-dimensional space that consists of an $n$-dimensional point for each possible outcome, where each agent's utility is expressed along a different axis.*

As we are dealing with bilateral negotiations (two agents), our utility space will only be two dimensional, and in our particular case, is always convex.[2] A convex space is one where if you take any two points in the space, all the points between those two points are also in the space. We will generally only concern ourselves with utility spaces along the positive portions of the axes. See Figure 2.1 for some examples of both concave and non-concave utility spaces.

Of particular interest in a utility space is the boundary of the space, as this forms a frontier of efficiency.

[2]Having a convex utility space is useful because it means there are many troublesome situations that do not arise compared to a non-convex space. As such, we will not go over these situations in our work.

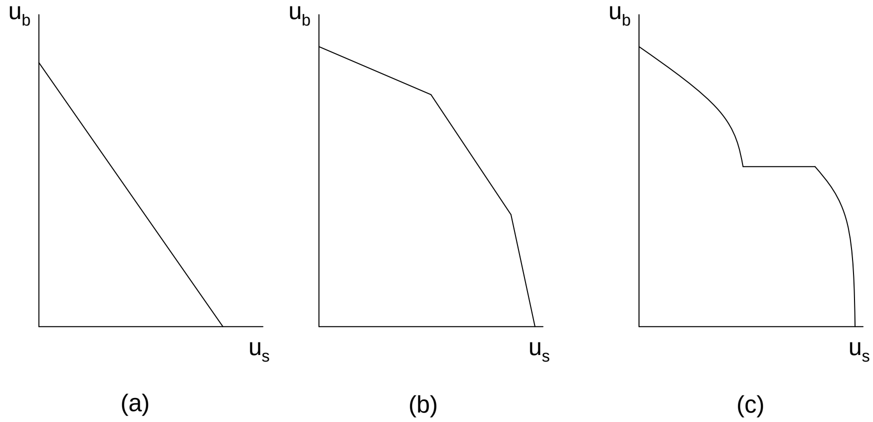


Figure 2.1: (a) A very simple convex utility space. (b) A slightly more complex convex utility space. (c) A non-convex utility space.

**Definition 3 (Efficient Frontier)** *The* efficient frontier*[22] consists of all the points in a utility space that have the following property: the point cannot be altered such that one agent receives a greater amount of utility (than in the original point), without another agent receiving less utility (than in the original point)* while *still being inside the utility space.*

Figure 2.2 highlights the efficient frontiers of the utility spaces introduced in Figure 2.1.

## 2.3 Mechanism Design and Game Theory

*Mechanism design* is a field in game theory that concerns itself with designing games that result in outcomes possessing desirable properties despite the existence of the self-interested agents playing the games. By designing the rules and protocols of a game correctly, the resulting mechanism produces a strategic situation where it is in the best interests of the agents to behave in a manner that results in a type of outcome desired by the mechanism designer. Mechanism design is often referred to as "inverse game theory" since *game theory* more traditionally has been about the study of the strategic aspects and properties of existing games and how to play them, whereas mechanism design works inversely creating new games with the aim of encouraging particular types of play and outcomes [16].

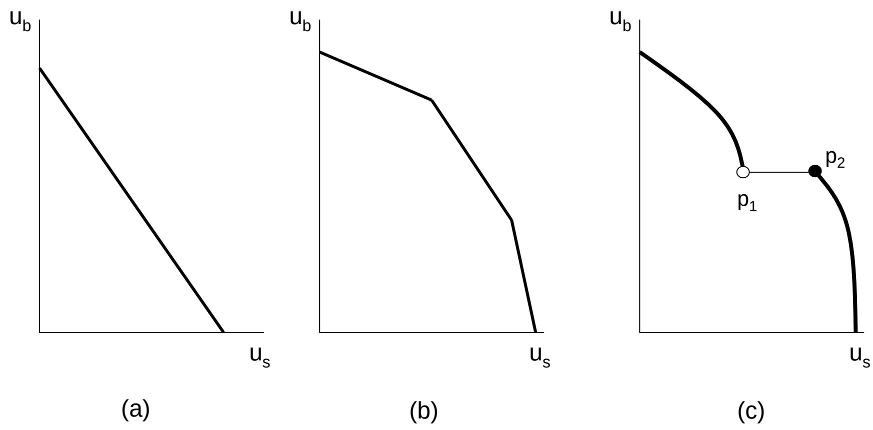


Figure 2.2: The same utility spaces as in Figure 2.1 with their efficient frontiers highlighted in bold. We see that in (c) the efficient frontier does not consist of the entire space boundary since all the points between p1 and p2 (including p1) can be replaced with point p2 to give agent $s$ a higher utility without reducing the utility of agent $b$. However, all the spaces we will deal with will be convex and so the efficient frontiers will be connected like in (a) and (b).

In this work, we will be acting as a designer of a negotiation mechanism that guarantees certain properties and attempts to achieve other desirable properties. The game theoretic properties of importance to us are outlined in the following section.

### 2.3.1 Important Game-Theoretic Properties

#### Individual Rationality (IR)

*Individual rationality (IR)* is a property portrayed by a mechanism if an agent participating in the mechanism is no worse off for participating. In terms of utility, this can be phrased as an agent participating in an individually rational mechanism will receive non-negative utility from the experience. This is a desirable property because agents, given the choice, are more likely to participate in an individually rational mechanism.

**(Individual Rationality)** A mechanism, M, is individually rational if for every agent $a \in A$, the utility of the agent is non-negative. $\boldsymbol{u_a \geq 0,\ \forall a \in A}$

There are three different types of IR that relate to the strength of the IR property with respect to time and knowledge: ex-ante, ex-interim, and ex-post. *Ex-ante IR* means that the expected utility for participating in the mechanism is greater than the utility for not participating, for each agent, *when the agents have no information about their own preferences yet.* *Ex-interim IR* is basically the same as ex-ante IR, except agents know their own preferences, but not the preferences of others. *Ex-post IR* means that the agents are guaranteed a better utility by participating in the mechanism than opting out, no matter what the personal preferences of the agents are.

### Incentive Compatibility (IC)

*Incentive Compatibility (IC)* is a property portrayed by a mechanism if agents participating in the mechanism have no incentive to be dishonest when asked to reveal private information to the mechanism. For example, if an agent is participating in an incentive compatible auction, then the agent will bid with its true reservation price rather than basing its bid on other agents' bids and on its personal beliefs about the other agents reservation prices. In terms of utility, incentive compatibility means that an agent's utility for declaring its private information honestly to the mechanism is never worse than making a dishonest declaration. This is a desirable property because it means that agents do not have to waste their time trying to figure out the best declaration strategy and the mechanism designer can trust the declarations it receives and actually make use of them.

**(Incentive Compatibility)** Given a mechanism, M, that requires agents to declare some private information, let each agent $a \in A$ have private information $I_a \in I$, and let the utility the agent receives for declaring $i$ be denoted $u_a(i)$. Then, M is incentive compatible if the utility each agent receives for declaring $I_a$ is at least as good as the utility for declaring any $i \in I$ where $i \neq I_a$. $\boldsymbol{\forall a \forall_{i \neq I_a} i \; u_a(I_a) \geq u_a(i)}$

We will also use a weaker form of IC called one-sided incentive compatibility [27]. Intuitively, one-sided IC means that an agent has no incentive to lie about their private information in a manner that would make it seem stronger, but it might still have incentives to declare information that makes its position seem weaker.

**(One-sided Incentive Compatibility)** Given a mechanism, M, that requires agents to declare some private information, let each agent $a \in A$ have private information $I_a \in I$, and let the utility the agent receives for declaring

$i$ be denoted $u_a(i)$. Let $I_{stronger} \subseteq I$ be the set of private information that would appear to make agent $a$'s position stronger than $I_a$, and let $I_{weaker} \subseteq I$ be the set of private information that would appear to make agent $a$'s position weaker than $I_a$. Then, M is one-sided incentive compatible if the utility each agent receives for declaring any value in $\{I_a\} \cup I_{weaker}$ is greater than or equal to the maximum utility of declaring any value in $I_{stronger}$. $\boldsymbol{\forall a \forall i \in (\{I_a\} \cup I_{weaker})\ u_a(i) \geq max_{j \in I_{stronger}}(u_a(j))}$

One might note that one-sided IC causes declarations that may have once produced a stronger position to now produce a weaker position. One-sided IC might also be considered as a psychological property as illustrated by first-price sealed bid auctions. In such an auction, bidders reveal their bids all at the same time and the highest bid gets the item and pays the highest bid amount. A bidder may appear to have a stronger bid position if it lies and bids higher than the item is actually worth to it; in this way, the bidder is more likely to get the item. However, if the bidder bids higher than its actual valuation and wins the item, it actually receives negative utility. In fact, if the bidder bids below its actual valuation, then it may still win the item and receive positive utility, even though it is declaring a "weaker" position.

#### Pareto Optimality

*Pareto optimality* is a property of a solution or outcome that says no agent can increase its utility over that of the solution without causing some other agent's utility to decrease. A mechanism is said to be pareto optimal if the outcomes/solutions it produces have this property. Note that the efficient frontier in a utility space is exactly the points in the space that are pareto optimal solutions (see Figure 2.2).

The value of a pareto optimal solution is that it acts as a measure of efficiency; if a solution is pareto optimal, then the solution is on the efficient frontier of the solution space. In other words, there are no further joint gains that can be had over a pareto solution; no amount of utility is going to waste.

### 2.3.2 Myerson-Satterthwaite Bilateral Trading Theorem

#### The Bilateral Trading Problem

The bilateral trading situation is where one agent has an item to sell to another agent, but the two agents do not know whether their respective reservation prices are compatible for an exchange to occur. More formally, we have two agents, the seller($s$) and the buyer($b$), and each agent has a reservation price (RP) for the item. The seller's RP, $p_s$ is drawn from an interval $[s_L, s_H]$ and the buyer's RP, $p_b$, is drawn from an interval $[b_L, b_H]$, each according to some probability distribution that fully supports each interval. In addition, there is a non-empty intersection between these intervals (i.e. $b_L < s_H$), which means it is possible for $p_s > p_b$. Agent $s$ is not willing to sell the item for less than $p_s$ and agent $b$ is not willing to pay more than $p_b$. If $p_s \leq p_b$, then an exchange should occur with the buyer paying the seller some price p, where $p_s \leq p \leq p_b$. If $p_s > p_b$ then an exchange should not occur.

The *bilateral trading problem* is the challenge of coming up with a mechanism that receives (possibly untruthful) declarations of reservation prices from the agents, and then decides whether or not a trade occurs and, if a trade occurs, at what the price the trade occurs. A useful mechanism for this problem should make sure that a trade occurs if $p_s \leq p_b$, and no trade occurs otherwise. If this is upheld, then no agent ever receives negative utility and the item is always exchanged when it should be. The heart of the problem is that the declarations from the agents cannot be trusted and the only other information available about the reservation prices are the intervals they were drawn from, but these intervals overlap and so we cannot know for certain whether $p_s \leq p_b$. To summarise, we cannot determine whether a trade should occur because we cannot determine the reservation prices from the intervals they are drawn from, nor can we determine them from the declarations of the agents. But can we design an incentive compatible mechanism where the agents declare their true reservation prices and make the mechanism designer's life easier? The following theorem from Myerson and Satterthwaite [19] responds in the negative.

**Theorem 1** *(Myerson-Satterthwaite Theorem) For the bilateral trading problem, no mechanism can be incentive-compatible, individually rational, and ex-post efficient.*[3]

This basically says that for the bilateral trading problem, we cannot get the following three properties to occur at the same time no matter what mechanism we use:

1. (incentive-compatible) Agents will declare their true reservation prices to the mechanism.

[3]For a formal proof of this theorem, please see [19].

2. (individually rational) Agents receive non-negative utility from the outcome of the mechanism.

3. (ex-post efficient) If the *actual* reservation price of the seller is less than the *actual* reservation price of the buyer, then a trade will occur, otherwise no trade occurs.[4]

This theorem is actually even more negative than it sounds, because it still holds when trying to work in expectation, which is a weaker setting. This theorem helps illustrate why negotiations are not nicely solvable when private information exists (which is almost always the case in real negotiation settings). Private information can create situations of uncertainty where no agent knows whether working together is a good thing for all agents, and no third party can extract this information from the agents truthfully.

In our work, we will be dealing with similar problems of incomplete information, except we will be dealing with multiple issues instead of just a single issue.

## 2.4 Negotiation

### 2.4.1 Third Parties: Mediation and Arbitration

A third party may engage in a negotiation in several ways: acting as a facilitator that brings parties together to negotiate, diffusing tensions by acting as an calming influence and messenger between the parties, suggesting useful procedures/ideas for the negotiation process[5], suggesting possible agreements, or even dictating an agreement. The actions open to a third party depend on how much power it has been imbued with.

A *mediator* can perform many of the functions listed, but cannot dictate an agreement as it does not have the power to enforce such an agreement. A mediator can suggest agreements it thinks are useful, but the agents are not required to accept such an agreement. The usefulness of a mediator comes from it being an unbiased

[4]This is a similar idea to pareto optimality/efficiency, but in the bilateral trading setting. However, technically any trade made in this setting is pareto optimal since one agent cannot receive more utility without the other receiving less, that is why we have the ex-post efficiency criteria instead.

[5]A relevant example is that a third-party might suggest playing some sort of negotiation game that helps the parties better understand their own personal preferences in the negotiation. They might even suggest using our system as a first step towards achieving a successful negotiation!

party in the negotiations which may give it the ability to see viable solutions or profitable areas of negotiation that a biased party might not realise is beneficial. Perhaps more importantly, the mediator has the credibility as a third party to suggest such solutions and have them seriously considered.

On the other hand, an *arbitrator* can usually do everything a mediator can with the addition of having the power to dictatorially enforce an agreement of its choosing, if it wishes to do so. The difference in enforcement powers can be illustrated by the following two institutions: a commission or tribunal that creates a report of recommendations on some situation that does not have to be followed (mediator) versus a judge that hands down a decision in a legal case (arbitrator).

Traditional mechanisms have the flavour of being an arbitrator: they tend to gather information from agents and then implement an outcome that agents must abide by. In our work we will act more as a mediator, and although we will be creating a negotiation mechanism that decides on many parameters of the negotiation, we will involve the agents in creating the final implemented agreement.

### 2.4.2 Single Negotiating Texts (SNTs)

A *single negotiating text* (SNT) is an agreement that allows all parties to focus on the same (intermediary) solution instead of focusing on their individual partisan offers [22]. Without this negotiation device, agents will typically propose independently developed solutions that will be widely disparate from the other agent's proposals. This leads to adversarial posturing where each side must start with an extravagant offer, as to not appear too weak, and then must come down towards a more reasonable position (which is probably closer to their actual position than what they first proposed). This type of politicised atmosphere puts too much emphasis on out-positioning an "opponent" and loses sight of the collaborative gains that can still be achieved. For example, consider Figure 2.3 (a), which illustrates a typical negotiation process against the backdrop of the utility space. Here the agents initially propose agreements $o_b$ and $o_s$, which are nowhere near each other. The most likely sequence of future offers will lead the two agents into the middle-ground of these two offers, the area A. However, even though the middle ground is more fair to all involved, there is still a whole area of joint gain that could still be achieved over any offer in A. This area, labeled W, is essentially wasted utility for both agents.

Using an SNT approach, the agents instead focus on a single agreement to build upon, only accepting a new SNT that all parties agree is better than the current agreement; this process continues iteratively until a final agreement is reached. The psychological shift in this approach is that now agents focus their attention and criticisms on the current SNT (that both agents came up with) as opposed to attaching criticism to the other agent's offer and the other agent itself. This incremental approach to the negotiation forces agents to think about how to improve the current agreement such that both parties can agree upon the new SNT; agents can no longer think in a purely egocentric way, producing offers that will just be rejected, but must focus themselves more realistically on mutually beneficial increments. The real shift in attention for the agents is that in order to increase their own utility, they must also focus on how the other agent can increase its utility at the same time in order to acquire that agent's approval.

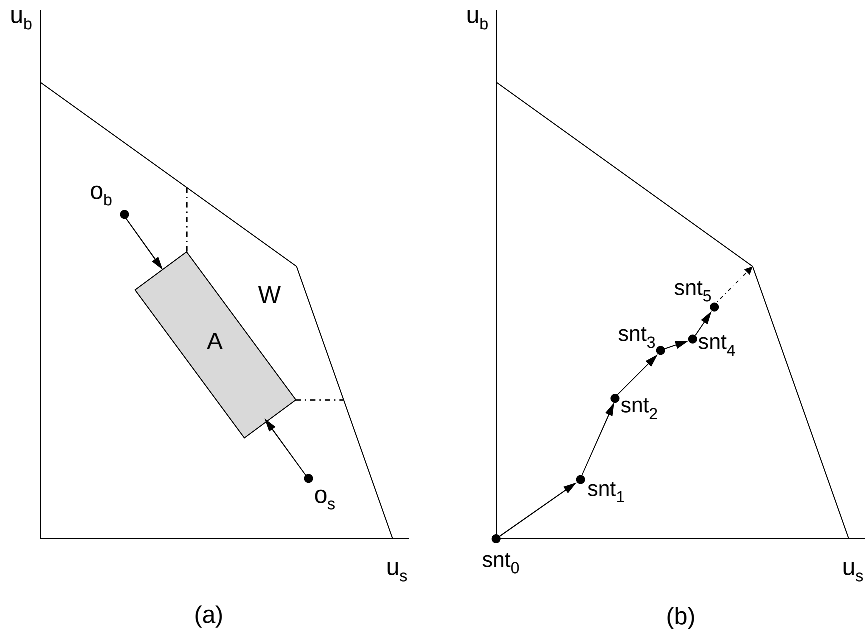


Figure 2.3: Illustration of the differences between negotiating with and without an SNT approach. (a) Without using an SNT, instead using an approach where agents take turns making offers. Offers tend to gravitate to somewhere in the middle ground (area labeled 'A') of the original offers until some agreement is reached. (b) Using an SNT. Agents proceed from the no agreement point (the origin) towards the efficient frontier, incrementally building upon previous agreements.

Looking at Figure 2.3 helps illustrate the approach of the typical negotiation process (Figure 2.3 (a)) versus that of an incremental SNT approach (Figure 2.3 (b)). Without using the SNT approach, offers from the two agents gravitate towards each other, moving somewhat parallel to the efficient frontier; there is little or no focus in regards to moving the agreement towards the efficient frontier, as most of the effort is spent on bridging the gap between the agents' initial offers. One benefit of this approach is that the ratified agreement is generally fair in the sense that it is a compromise between the two agents' initial offers.[6] The incremental SNT approach, on the other hand, starts from an unsatisfactory agreement but continually moves towards the efficient frontier, the focus always being "how can we squeeze more utility out of this agreement?". This approach also renders fair agreements since the intermediary agreements must continually satisfy both agents, but by moving orthogonally to the efficient frontier, it is more likely that the wasted utility space will be reduced.

### 2.4.3 Agendas

With regards to a multi-issue negotiation, a negotiating *agenda* is a schedule that specifies the order in which issues will be negotiated. Such a schedule specifies not only the order of issues, but separates issues into stages such that all the issues in one stage are negotiated before the issues in the succeeding stage in the schedule. Agendas allow negotiations to be broken down into more easily digestible stages (see Figure 2.4), but can also be used as a tool to manipulate negotiations depending on which issues are in which stages.

Much of the agenda research has only allowed the negotiation of a stage to consist purely of the issues in that stage, meaning that you are only negotiating the utility of that single stage. This means that we are throwing away utility at each stage if an efficient agreement is not made in the stage. The agendas that we will use in our work will allow for renegotiation of previous stages, which is based on the work of John and Raith [11, 12]. This means that at any given stage, the negotiation process covers the issues in the current stage as well as all issues in previous stages, the only caveat being that any prices already agreed upon for issues in previous stages now compose the default agreement for the current stage.[7] By allowing renegotiation of issues in previous stages, it is possible to recover utility

[6] However, this means that an agent must be very careful about how it chooses its initial offer, which can be very difficult.

[7] For example, if the agents have agreed that issue 5 in the last stage should have a price or value of 22, then if they do not renegotiate issue 5 in the current stage the agreement price will still be 22. However, if both agents agree to a renegotiated price for issue 5, then this new price is incorporated into the latest intermediary agreement.

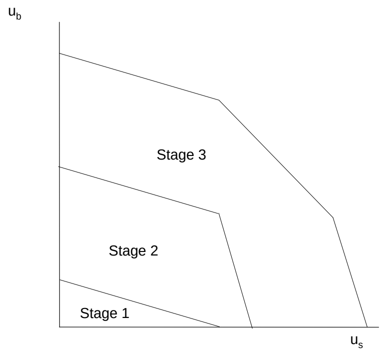


Figure 2.4: An agenda causes the utility space to be divided among its stages so that parts of the utility space are negotiated in each stage.

that would have just been wasted in a non-renegotiating agenda. Figure 2.5 helps explain the differences.

### 2.4.4 Solution Concepts

A solution concept [13] for a negotiation is essentially an agreement that has certain properties or comes about through the rational constraints that arise through analysis of a particular negotiation procedure. The first type of solution concept is called *axiomatic*, because the solution is derived from a set of axioms that must be satisfied. Most axiomatic solutions of interest define a unique solution and we consider the two most popular axiomatic solutions: the Nash bargaining solution and the Kalai-Smorodinsky bargaining solution. The second type of solution concept arises from the strategic analysis of a negotiation model or procedure. We will consider the Rubinstein-Ståhl alternating offers model and its analysis below.

In our work, we will make use of both an axiomatic approach as well as a more procedural approach. The axiomatic approach we use is the Nash bargaining solution, however, it is entirely possible to substitute another solution concept into our model as fits the needs or requirements of the particular negotiation situation. We use a similar analysis to the Rubinstein-Ståhl alternating offers model in order to determine how a rational agent would behave in our system.

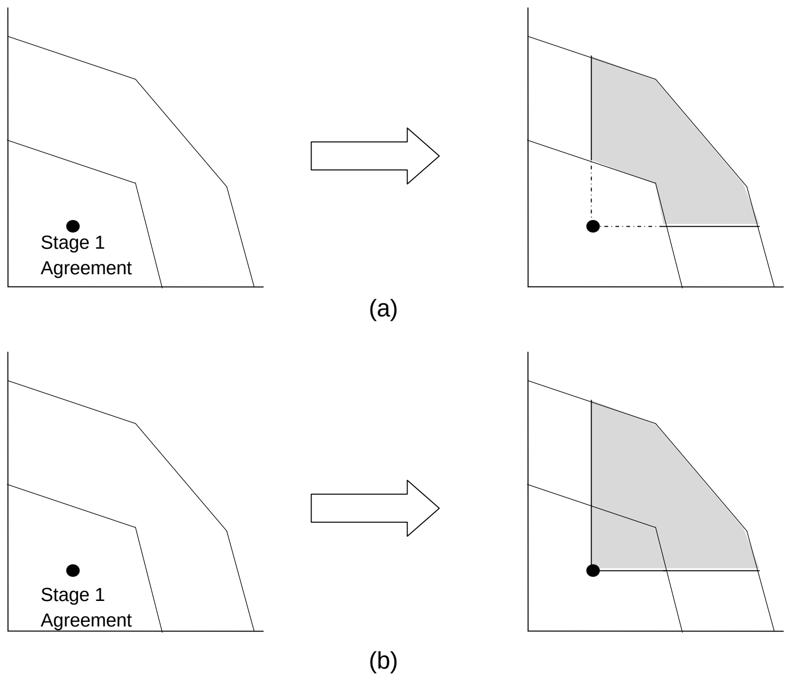


Figure 2.5: Showing the difference between agendas that do not allow renegotiation and those that do, in terms of the utility space. (a) Without renegotiation, when the stage one agreement is made, in stage two we can only negotiate the utility strictly represented by stage two (the shaded portion of the space). (b) With renegotiation, we can renegotiate the same space as in (a) but we are still open to renegotiating the utility that went to “waste” in the stage one agreement. This gives agents a chance to still have an overall efficient agreement.

#### Nash Bargaining Solution (NBS)

The *Nash Bargaining Solution* (NBS)[8] is probably the most well known axiomatic bargaining solution, and it stems from the following four axioms that the solution must satisify:

1. **Pareto Optimality/Efficiency -** The solution must not be improvable for one agent without harming another agent. This was already discussed in Section 2.3.1 on page 14.

2. **Symmetry -** The solution does not depend on the identities of the agents. If we switched the agents identities but kept all the utility information the same, then the solution should still be the same.

3. **Invariance to Affine Transformations -** If we already have a solution and then the utility information undergoes an affine transformation, the new solution for this new scenario is the same as the old solution having undergone the same affine transformation.

4. **Independence of Irrelevant Alternatives (IIA) -** If we have a utility space U, and the solution for this space is $s \in U$, then if we restrict our space to some subspace $V \subset U$ where $s \in V$, the solution for the utility space V should still be $s$.

The NBS is actually the only solution that satisfies these four properties [13]. But what do the properties mean? Pareto optimality should be fairly clear by now: it is an efficiency criteria. Symmetry just means that the solution is not biased by the identities of the agents, whoever the agents may be, whether they are powerful or weak entities, the solution remains the same.

The invariance of affine transformations needs a little bit of explaining. Consider the example in Figure 2.6. The utility space on the left has the solution $(\frac{x}{2}, \frac{y}{2})$. The utility space on the right is just the space on the left where the y-axis utilities are multiplied by 3 and the x-axis utilities by 2. According to the invariance of affine transformations axiom, the solution for the new space should just be $(2\frac{x}{2}, 3\frac{y}{2})$, or the old solution having undergone the affine transformation the utility space has undergone. This axiom is useful because it means that our solution does not depend on the agent's valuations of their own utilities, which can be operating on vastly different scales from agent to agent. For example, say that one agent is measuring

[8] First introduced by Nash in [20].

its utility in the range 0 to 10 and the other agent is measuring its utility in the range 0 to 1000. If we wanted to maximise the overall utility, we would be biased towards the latter agent, however the utility ranges of the agents are not, and should not, be comparable because these ranges are just abstract personal measurements of some measure of satisfaction or usefulness. If the first agent all of sudden decides to use a range from 50 to 100, has anything about the situation really changed? Not really, and neither should our solution, it should be based on other properties of the negotiation scenario.

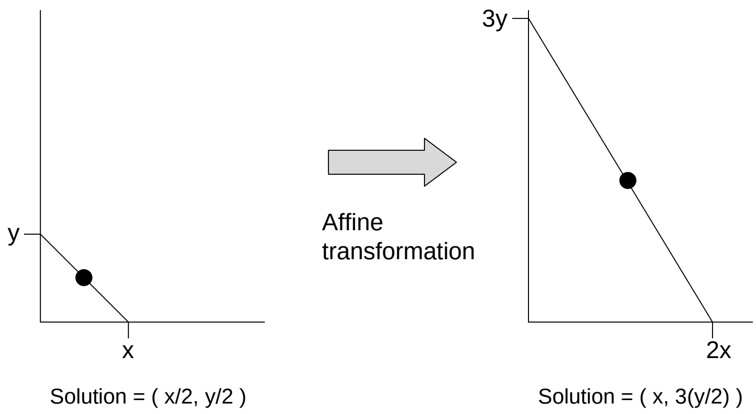


Figure 2.6: Example of *invariance of affine transformations*. If for the utility space on the left, the solution is $(\frac{x}{2}, \frac{y}{2})$ then when we move to the utility space on the right by the affine transformation that multiplies all utilities along the x-axis by 2 and the y-axis by 3, the solution for this new space is just the old solution under the affine transformation: $(2\frac{x}{2}, 3\frac{y}{2})$.

The last property, Independence of Irrelevant Alternatives (IIA), is the most controversial of the properties. To understand why, let us first give some examples of this property. Figure 2.7 shows two examples of IIA stemming from the same original space, U. In both of the derived subspaces, $V \subset U$ and $W \subset U$, the original solution $s$ still exists. The IIA property then says that if $s$ is the solution of the larger space $U$, then it should still be the solution in the smaller spaces $V$ and $W$. The rationale is that $s$ is originally the best solution, so why should removing solutions that were not considered before alter what is considered the best solution in the smaller spaces?

Raiffa [22] gives a very good example, that we paraphrase here:

> You go into a restaurant and peruse the menu, choosing the entree, braised beef. After making your decision, you notice several entrees have stars next to their names and you enquire as to why. The waiter tells you that those entrees are not available today. Your already chosen entree, braised beef, is not among these starred choices, so it is still available. Your choice of entree should still be the braised beef, should it not? After all, that was what you considered the best option from the full menu, shouldn't it be the best option for you on the restricted menu?

The story just mentioned makes it seem that IIA is a very reasonable axiom after all, so what is the criticism of IIA? The problem lies in the fact that the story just mentioned only considers a single agent's options, but a negotiation involves multiple agents. In Figure 2.7, the solution $s$ in the space $V$ gives the vertical axis agent its maximum utility, but only gives the horizontal axis agent about two-thirds of its utility. This seems more fair in the full space $U$, where each agent is getting a share of their maximum utility, but seems unfair in the space $V$. Why should one agent get its maximum utility while the other agent does not? This is the main criticism of the Nash solution and is addressed by the Kalai-Smorodinsky solution detailed in Section 2.4.4.

The question then becomes, why do we use the NBS in our system? The main answer is that the sorts of utility spaces that allow for this criticism do not really appear in our system. However, we could just as easily use the Kalai-Smorodinsky solution if we desire; one useful aspect of our system is that its functioning does not depend on the solution concept used. So the most appropriate solution concept for a particular negotiation scenario should be the one used, whether it is the NBS, the KSBS(Kalai-Smorodinsky), or some other solution concept. We just choose the NBS because it is the most well-known solution, has reasonable properties, and is easy to calculate.

We have just alluded to the fact that the NBS is easy to calculate, so how is it calculated? It would seem it might be complicated to satisfy all four of those axioms, but it turns out that the *unique* solution that satisfies those axioms is simply the solution that maximises the product of the agents' utilities [13]. Since the solution is known to be Pareto-optimal, all we have to do is find the point on the efficient frontier that maximises the product of the agents' utilities.

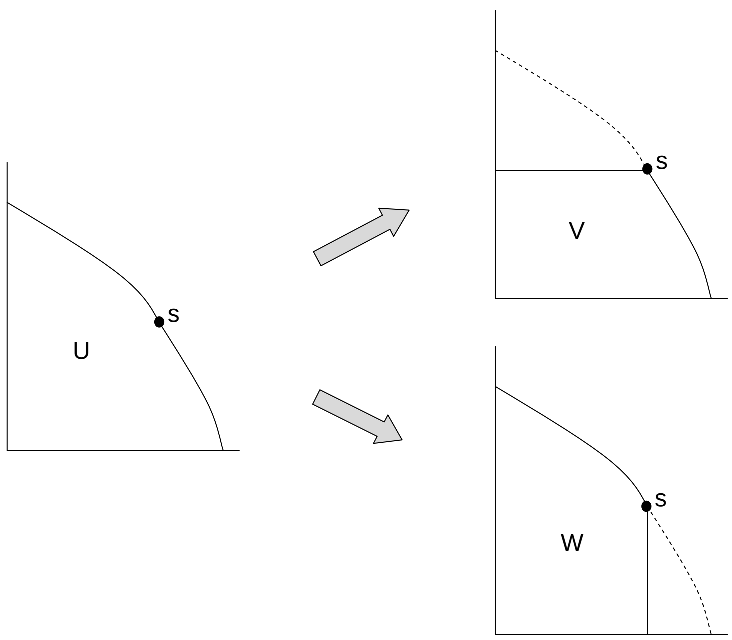


Figure 2.7: Example of the independence of irrelevant alternatives (IIA) axiom.

### Kalai-Smorodinsky Bargaining Solution (KSBS)

The *Kalai-Smorodinsky Bargaining Solution* (KSBS) is one of the most used alternatives to the NBS. The KSBS is another axiomatic solution that uses all the same axioms as the NBS except that the IIA axiom is replaced with a monotonicity condition. The monotonicity condition is a bit lengthy to explain and is unnecessary for our consideration (for the specific details see [13]).

The important aspect of the KSBS is that it is the solution that gives each agent the same percentage of its maximum utility. The percentage changes based on the shape of the efficient frontier, but is the same for each agent in a given solution. This seems very fair because each agent realises the same proportion of their maximum utility.

Figure 2.8 (a) illustrates how to calculate the the KSBS. First, an ideal point is defined, $u^*$ which is the point where each agent receives its maximum utility. A line is drawn between $u^*$ and the disagreement point (in most of our cases, the origin), and where this line intersects the efficient frontier is the KSBS solution. Figure 2.8 (b) shows where the KSBS solutions would be in the restricted spaces $V$ and $W$, at $s_v$ and $s_w$ respectively, which can be contrasted with the NBS solution $s_{nbs}$ that does not change from $V$ to $W$.

### Rubinstein / Ståhl Alternating Offers Model

The Rubinstein/Ståhl alternating offers model [24, 26] is a widely used negotiation procedure used in the literature. Its essential feature is that it gives one agent the opportunity to propose an agreement, and if that proposal is rejected by the other agent, then the other agent now gets to propose an agreement. This continues until an agreement is reached or some time limit expires or a discount factor eventually makes the negotiation worthless.

This model typically consists of dividing some item, usually a pie. The size of the pie starts at 1, but as time goes by the pie "shrinks": for each rejected proposal, the pie shrinks by a factor $\delta < 1$. At the beginning of the game the pie is of size 1, after one rejection, it is size $\delta$, after two rejections it is size $\delta^2$, and so on. If the current size of the pie is $\pi$[9], then when an agent proposes an offer, it offers $x \leq \pi$ to the other agent with $(\pi - x)$ left for itself.

To determine what an agent should offer in the first round, we use backwards induction, but intuitively, an agent wants to offer some amount that is between

[9] :)

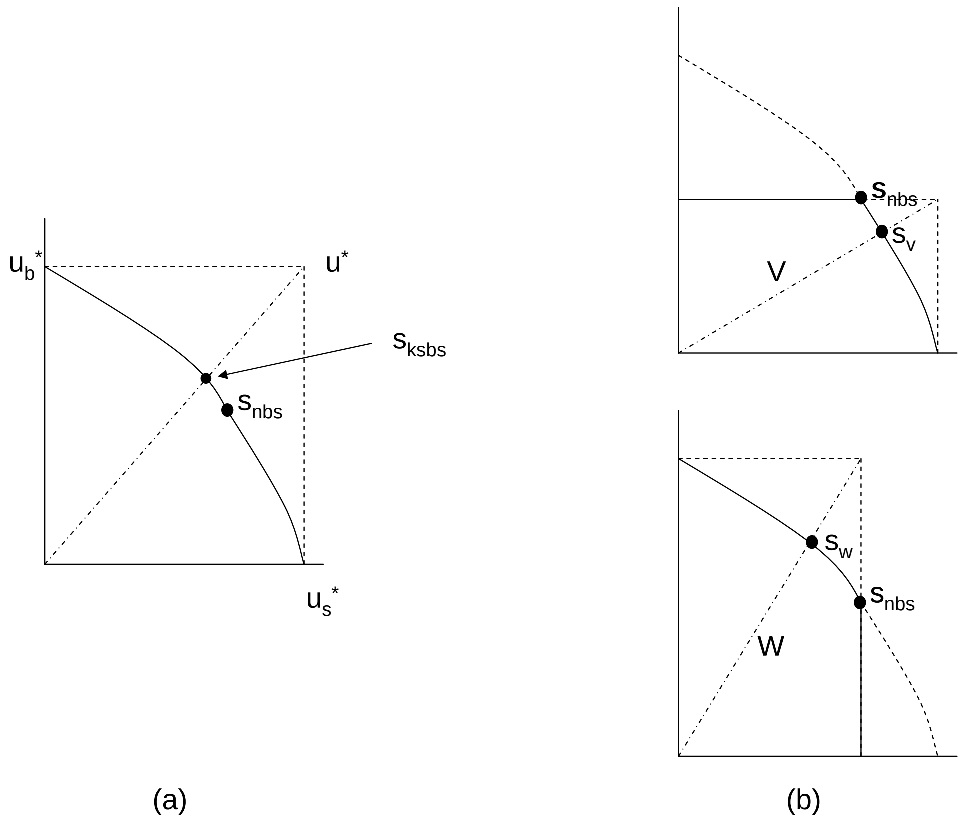


Figure 2.8: (a) An example of the Kalai-Smorodinsky solution compared to the Nash Bargaining solution. (b) Examples of how the Kalai-Smorodinsky solution changes as the utility space changes. The dotted lines of the old frontier show the area that has been removed from the utility space to give spaces $V$ and $W$, and the Kalai-Smorodinsky solutions are $s_V$ and $s_W$, respectively.

two values: the least amount that the other agent can expect in the next proposal round and the most that the other agent can expect in the next proposal round. The benefit that the proposer gets is that it gets an extra share due to the amount of the pie that would be discounted between this round and the next round, since the other agent cannot reasonably expect to get this portion in the next round (since it will have "shrunk" away by then). The backwards induction we outline next is from Gintis [9].

We assume that both agents are rational agents trying to maximise their personal utilities (or portions of the pie). Let the first proposer be agent 1, the second proposer agent 2, and let $z$ be the *most* that agent 1 can get in this game (the most agent 1 can get given that both agents are playing rationally). If this is the most agent 1 can get in the game, then this amount must be what agent 1 can get in the first round because its aspirations decrease as time goes by due to the discount factor and the fact that the later rounds are exactly the same as the earlier rounds except that the pie is smaller (meaning there is no strategic difference). The second time agent 1 gets to propose, the most it can get is $\delta^2 z$, therefore the *most* that agent 2 must offer, in its first chance to propose, to player 1 is this amount $\delta^2 z$.[10] This means that agent 2's share of the pie would be *at least* $\delta - \delta^2 z$, and therefore going back to the first round, agent 1 must give agent 2 at least $\delta - \delta^2 z$. Finally, this means that agent 1's share from the first round proposal is $1 - \delta + \delta^2 z$, but this must equal $z$ because agent 1 is offering agent 2 the least amount required and thus must be getting the most it can possibly get from a first round proposal, namely $z$.

$$\begin{aligned} z &= 1 - \delta + \delta^2 z \\ &= \frac{1-\delta}{1-\delta^2} \\ &= \frac{1}{1+\delta} \end{aligned}$$

Now let $a$ be the *least* agent 1 can possibly get in this game, then on agent 1's second turn to make an offer the least it can get is $\delta^2 a$. On player 2's first turn to offer, the *least* that needs to be offered to agent 1 is this amount $\delta^2 a$, and the *most* player 2 gets is $\delta - \delta^2 a$. Therefore in the first round, the *most* agent 1 must offer is

[10]Because then agent 1 is indifferent between accepting agent 2's proposal and going to the next proposal round. Technically, because $\delta^2 z$ is the most agent 1 can get in the next round, it may actually prefer accepting this proposal than going to the next round.

$\delta - \delta^2 a$, leaving itself $1 - \delta + \delta^2 a$, which must equal $a$ since now agent 1 is receiving the least it possibly could. And we get the following:

$$
\begin{aligned}
a &= 1 - \delta + \delta^2 a \\
&= \frac{1-\delta}{1-\delta^2} \\
&= \frac{1}{1+\delta}
\end{aligned}
$$

Therefore, $a = z$, the least agent 1 can get is the same as the most it can get and agent 1 will propose the following offer in the first round,

$$\left( \frac{1}{1+\delta}, \frac{\delta}{1+\delta} \right)$$

which agent 2 immediately accepts (since it is the most/least that agent 2 can get as well).

# Chapter 3

# Related Work

In this chapter we first discuss the philosophical influence of a major work in negotiations (Section 3.1), then move on to the use of arbitrators and mediators in game theory (Section 3.2), and finally, delve into research related to the use of agendas and incremental negotiation (Section 3.3).

## 3.1 The Art and Science of Negotiation

Published in 1982, *The Art and Science of Negotiation* by Howard Raiffa [22] is one of the most well-known books in the field of negotiation. Its compilation and exploration of every aspect of negotiations is both informative and inspiring to naive negotiators and seasoned professionals alike.

Raiffa had a particular bias towards being unbiased, considering the viewpoints of all parties and implicitly encouraging others to do the same in this book. Parties are involved in a negotiation because, for one reason or another, they all benefit from mutual participation and collaboration. Raiffa focused on this spirit of "we are all in this together", promoting exploration of joint gains over individual selfishness. When individual selfishness is the mindset of agents, an adversarial environment arises, often leading to inefficient agreements that "leave something on the table" that all parties could have jointly benefitted from, but none will (see for example, the space labelled $W$ in Figure 2.3).

The original inspiration for this entire thesis is the idea of this lost utility going to waste and finding a way to recover it. The result of this influence can be seen in our framework's guiding principle: direct agents to explore areas of the solution space that provide joint benefit.

## 3.2 Arbitration and Mediation in Game Theory

### 3.2.1 Arbitrators in Games

Arbitrators need to look at situations, acquire information, and determine an outcome or solution based upon this acquired knowledge combined with its principles for a good solution. When an arbitrator does not have the ability to gather information independently (as is often the case in game theoretic settings), it must rely on the agents for its data. Rosenthal [23] explains that if an arbitrator needs to ask for information, then the agents have incentives to lie if they can predict how their declaration may affect the arbitrator's solution. In bargaining, it is a difficult problem to get agents to tell the truth and maintain efficiency and individual-rationality as we saw with the work of Myerson and Satterthwaite [19] in Section 2.3.2. Pratt and Zeckhauser [21] studied the multi-issue setting and put it succinctly:

> There is no costless way out of the incentive compatibility problem.

Both of these groups of authors showed that if you want the truth you have to pay for it; you have to give the agents incentive to tell the truth. This problem of incentive compatibility is why our third-party framework does not act as an arbitrator. While our framework does have flavours of arbitration, we do not want to dictate a solution coming from "poisoned" data. As a third-party intervenor, we cannot avoid relying on such data, but we avoid dictating a final agreement from it, instead facilitating an agreement as a mediator.

### 3.2.2 Mediators in Games

Before discussing existing mediators in games, it will be instructive to use a running example, for which we use the classic Prisoner's Dilemma (PD) game shown in normal-form in Figure 3.1.

In this game, if both agents choose to cooperate then they each receive 4 utility units, if both defect they each receive 1 utility unit, and if they choose opposite strategies, the defector gets 6 and the cooperator 0 utility units. In this situation, game theory tells us that rational agents will both defect, which is an inefficient solution but guarantees each agent that they will not get 0 utility units. If only agents could trust each other enough to cooperate, they could both better their utility. This is where a third-party mediator can be of service.

| | cooperate | defect |
|---|---|---|
| cooperate | 4, 4 | 0, 6 |
| defect | 6, 0 | 1, 1 |

Figure 3.1: The classic prisoner's dilemma game.

One of the earliest forms of mediation in games can be found in Aumann's correlated equilibria [1], the essential idea of which is that each agent observes a public signal and chooses its action based on the outcome of that signal. For example, if a coin were flipped and an agent decided to cooperate if heads came up and defect if tails came up. An actual correlated equilibrium occurs if the strategies assigned to each signal are such that neither agent would want to deviate from its strategy. Unfortunately for the prisoner's dilemma, the only correlated equilibrium is to defect on every signal. One should observe that a correlated equilibrium requires agents to not deviate, although in theory they should not want to.

Monderer and Tennenholtz [18] take the idea of mediators even further. In a game like the prisoner's dilemma, they create a new game on top of the old game where a mediator action is now available. Agents have the option of choosing to defect, cooperate, or to let the mediator play for them. If an agent lets the mediator play on its behalf, the mediator behaves as follows: cooperate if the other agent also chooses the mediator and defect otherwise. The new game becomes the 3x3 normal-form game shown in Figure 3.2.

This mediator has given agents the best of both worlds! However, this type of mediator has added something to the original game that did not exist before, the ability of a party to act after a decision has been announced by an agent. That is, the mediator decides how it will play *after* the agents have announced their choice about using the mediator or not.

The types of mediators mentioned are layered on top of existing games as opposed to being an integrated part of the game. Our third-party mediation system is a bit of both: it is layered on top of an existing general negotiation setting, but at the same time will be integrated into the negotiation proceeding in a way that agents do not get to choose whether or not to use it. Monderer and Tennenholtz's

| | mediator | cooperate | defect |
|---|---|---|---|
| mediator | 4, 4 | 6, 0 | 1, 1 |
| cooperate | 0, 6 | 4, 4 | 0, 6 |
| defect | 1, 1 | 6, 0 | 1, 1 |

Figure 3.2: The prisoner's dilemma modified to allow for agents to choose to let a mediator play for them. In this game, the agents dominate strategy is to let the mediator play for them.

mediators are given permission to play on behalf of an agent, whereas our mediator will facilitate the negotiation for both agents but will not play for them.

Our mediation system is more akin to the market mechanism of Bartal et al. [2]. In their work, Bartal et al. have a set of buyers and sellers, where each seller is selling one good identical to each other seller's good, and the buyers and sellers have different valuations for this type of good. The mechanism developed manages to pair up buyers and sellers and come up with a negotiation range called a *Zone of Possible Agreements* (ZOPA) for each pair, within which the agents are guaranteed a positive utility. The mechanism provides the buyers with the high end of the ZOPA range and the sellers with the low end, and then lets the agents themselves negotiate the final price. The information provided to the agents is used as a guarantee of a beneficial agreement to both parties, where the seller does not agree to anything below the bottom of the ZOPA and the buyer does not agree to anything above the upper end of the ZOPA.

The truly remarkable thing about this work is that by pairing trading partners correctly and determining a negotiation range (as part of the pairing process) instead of just assigning a price to each trading pair, the mechanism becomes efficient, individually rational, incentive compatible and budget balanced. The authors are able to circumvent the results of Myerson and Satterthwaite using this idea of letting the agents themselves finalise the negotiations, acting merely as a mediator to the market process. This is the essence of our mediation process, create a particular negotiation scenario and allow agents to formalise the final agreement within the parameters we provide them. Our system will not be able to guarantee the same results however, and this is mainly due to the fact that we deal with the bilateral case, whereas Bartal et al. have multiple buyers and sellers to play off of each other.

## 3.3 Agendas and Incremental Negotiation

### 3.3.1 Agendas

Fershtman [8] was one of the early researchers to show how agendas can affect the final agreement of a negotiation. Like many researchers, he used a model of splitting pies; where a pie represents an issue. He used two pies and had agents negotiate the division of the pies one at a time. After both divisions were negotiated, the agents could "consume" their portion of the pies.

The catch in Fershtman's study was that the pies where valued asymmetrically by the agents such that each agent's favourite pie was different, but the amount they valued their favourites were the same. Additionally, there was a time discount factor in place that meant the longer negotiations continued, the smaller the pies became. Fershtman showed that in this setting, an agent wants to negotiate the pie it likes the most second, owing to the fact that the other agent will be impatient to eat its ever-diminshing share of the first pie, which happens to be its favourite pie.

Fershtman's setting was rather simplistic with only two issues and being complete information, but two important points were made: using different agendas changes the final agreement and agreements in this study were usually inefficient. The agreements were inefficient because each agent had a portion of pie that they did not really care about, when the efficient agreement would have each agent possessing all of its favourite pie. The problem that arose was that when negotiating the pies separately, agents can not make tradeoffs with respect to their asymmetric valuations of the pies.

Fatima et al. [7] performed a more thorough examination of the use of agendas under complete *and* incomplete information settings with many issues. They explored three situations in each setting, characterised by how many issues were negotiated at a time: all issues at once, one issue in each stage, and multiple issues in each stage. What they found was that the "package deal" (all issues at once) was the most able to reach Pareto-efficiency but also the most computationally expensive. The explanation for this is that the package deal allows the most flexibility in making tradeoffs amongst asymmetrically valued issues, although this ability to find all the tradeoffs to form pareto-optimal solutions comes at a computational cost.

Clearly, these results indicate the need for maneuverability for making tradeoffs amongst the issues. However, there are some problems we have with these studies. In our work, we try to maintain a general and realistic negotiation setting, supported by not assuming complete information and keeping our assumptions about the agents to a minimum. Fershtman's setting assumed complete information and Fatima et al.'s best results would require fully rational agents that can determine the optimal solutions across all the issues. We would like to reduce the threshold required to achieve optimal (or near optimal) agreements, while maintaining issue tradeoff flexibility. This leads us to discuss research regarding incremental negotiations.

### 3.3.2 Incremental Negotiations

The complexity in determining an optimal strategy in the package deal setting is due to everything needing to be figured out all at once. If we separate the negotiation into stages to make things simpler, we lose some of the tradeoff flexibility. The compromise is an idea espoused by John and Raith [11, 12] that suggests negotiating in stages but allow renegotiation of issues in stages that have already been negotiated; they refer to this as an "increasing-pie model". This means in the last stage the full range of issues is available to be negotiated, but at the same time negotiation is simpler because it is broken down into parts.

One might ask how such an incremental negotiation procedure is simpler since at each stage we are just adding more and more issues for the agents' consideration. The reason it is still simpler is that as we add new issues at each stage, we are also constrained by the partial agreement from the previous stage. Not all options for the issues need to be considered since we do not deal with the entire solution space for the given set of issues, but only a smaller space of viable solutions defined by the previous agreements. On the other hand though, this constraining factor can also be a disadvantage since it *could* direct attention away from the more desirable areas of the solution space similar to finding a local maximum in hill-climbing that is not the global maximum. Despite this, we use the agendas with renegotiation approach and will actually turn this "disadvantage" of a constrained space into a tool for guiding agents to desirable areas of the solution. We explore how this can be done in the next chapter.

# Chapter 4

# Constraining the Solution Space

The essential purpose of our mediator framework is going to be to guide agents through the solution space, the utility space of all possible agreements. If agents are left to negotiate an agreement on their own, then they face a search problem through an agreement space that neither agent knows the exact shape of, and contains many elements which are undesirable for one or both agents. Our mediator will reduce the agreement space to desirable agreement points, thus focusing the attention of the agents in an appropriate fashion. To this end, we propose using a mediator who will suggest intermediary agreements, with the understanding that these agreements are to be incrementally improved upon by the agents themselves. By doing this, we can control the subspace of solutions that arise from the intermediary agreement point and thus steer the direction of the final agreement.

## 4.1 Constraints of Intermediary Agreements

Consider Figure 4.1 (a), a solution space entirely unconstrained where the default agreement, or starting point, is the origin. Agents negotiating in this space will have a "reasonable" solution space that will look something like the subspace, $R$, in Figure 4.1 (b). This "reasonable" space arises from natural constraining factors: both agents want to gain some utility over the default agreement (moving away from the origin), but are unlikely to agree to something that gives them little utility while the other agent receives a lot of utility (moving away from the perpendicular axes of the space $R$). Within $R$, the negotiated agreement for this space might be anywhere in the space, some examples being the points $a$ and $b$ in Figure 4.1 (c).

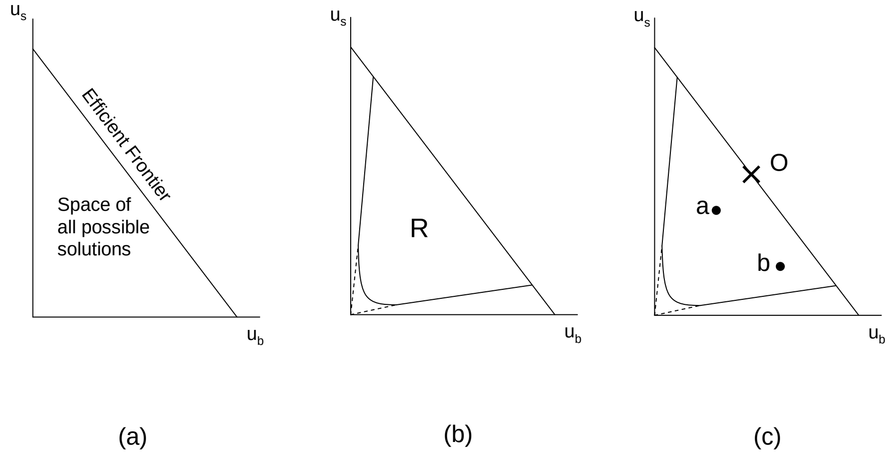


Figure 4.1: (a) An unconstrained utility space. (b) The subspace, R, of the unconstrained space that forms the set of "reasonable" solutions given the natural constraint that each agent wants more utility than the default agreement (the origin) but also does not want a very unbalanced agreement. (c) Examples of possible solutions within the reasonable space of solutions R. The solution marked O is a desirable solution for the negotiation, arbitrarily chosen.

Let point, O, in Figure 4.1 (c) be a solution that we would like the agents to attain. If there were a way to make it so that agents are only focused on the utility space near O, then it would be far more likely that the agreement would be closer to O than if the same agents negotiated in the unrestricted space of Figure 4.1 (a). This is the purpose of using intermediary agreements derived by the mediator.

If a mediator provides intermediary agreement $SNT_0$ (see Figure 4.2 (a)), then the new reasonable space of solutions becomes $R'$, since the same natural constraining factors apply and neither agent wants to receive less utility than it already has in the intermediary agreement. With this new space of consideration, $R'$, agreements $a$ and $b$ in $R$, become something more like $a'$ and $b'$ in $R'$, respectively (see Figure 4.2 (b)). The reduced space brings agreements closer to O because all the agents' attention is on the space $R'$ that is focused around O; by virtue of the restricted space, almost any reasonable agreement is close to O.

A further benefit of constraining the solution space with intermediary agreements is that we are more likely to achieve more efficient agreements since agents are starting from a more efficient default agreement ($SNT_0$ in Figure 4.2 (a)), by virtue of being closer to the efficient frontier. Agents will not waste their time going

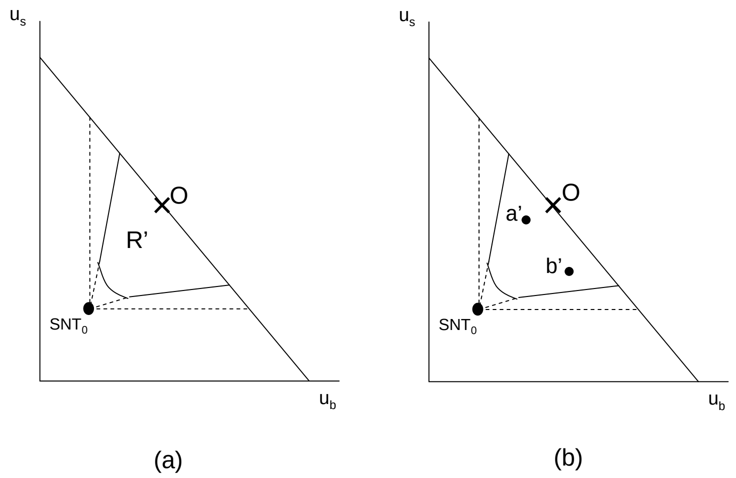


Figure 4.2: (a) The reasonable subspace of solutions, $R'$, when the utility space is constrained by an intermediary agreement, $SNT_0$. (b) Where the example solutions, a and b, from Figure 4.1 (c) would likely appear when agents are given an intermediary agreement to start with. The new solutions, $a'$ and $b'$, are much closer to the desired solution O than a and b were in the unconstrained space.

back and forth with offers that bring them up to the level of $SNT_0$, but instead they immediately focus their efforts advancing beyond an agreement like $SNT_0$. Even though the main goal of using intermediary agreements is guidance towards a desired area of the solution space, there is also an effect on efficiency since the constrained space an intermediary agreement creates is also a more efficient starting space.

## 4.2 Adding Stages

If we extend the idea of mediating intermediary agreements to mediating several intermediary agreements distributed through an incremental negotiation process, we obtain an incremental method with which to guide agents towards a specific point on the efficient frontier.[1] By introducing more than one stage to the process,

[1]Later when we start looking at the instantiation of our mediator framework, we are actually going to be doing the inverse of what is suggested here: given intermediary agreements, we will look for the agenda that best guides agents to a desired agreement on the efficient frontier, but not actually create the intermediary agreements specifically to guide agents to this chosen desired agreement.

the mediator has the ability to constrain the solution space at multiple stages and take advantage of the negotiations of the agents.

Consider a situation where agents negotiate half of the issues first and then the other half is negotiated incrementally off of the agreement from the first half. If a mediator is allowed to introduce intermediary agreements in each of these stages, then we get a situation like in Figure 4.3 (a). The mediator actually introduces only one intermediary agreement over all the issues, but when the agreement is broken down into stages, it becomes $SNT_0$ and $SNT_1$ in the figure, with $SNT_1$ including the issues from stage one as well.[2] In the first stage the mediator has proposed intermediary agreement $SNT_0$, restricting agents to the subspace A, and in the second stage, the mediator proposes $SNT_1$, which restricts the agents to the subspace B in stage 2. When the agents negotiate the first stage, coming up with agreement $z_0$, see how this affects the second stage in Figure 4.3 (b). The combination of $SNT_1$ and $z_0$ actually reduces the subspace under consideration in the second stage to an area even closer to the efficient frontier, and hopefully also closer to our desired solution; if we did not have stages, we could not take advantage of this combined effect of the intermediary agreement and the stage negotiations of the agents.

This example has shown how a third-party ability to intervene provides a significant measure of control over the location of the final agreement of the negotiation, while still allowing agents to be part of the negotiation process. The ideas in this chapter will be central to our instantiation of the mediator framework that is introduced in in Chapter 5.

[2]Since we are using an increasing-pie model.

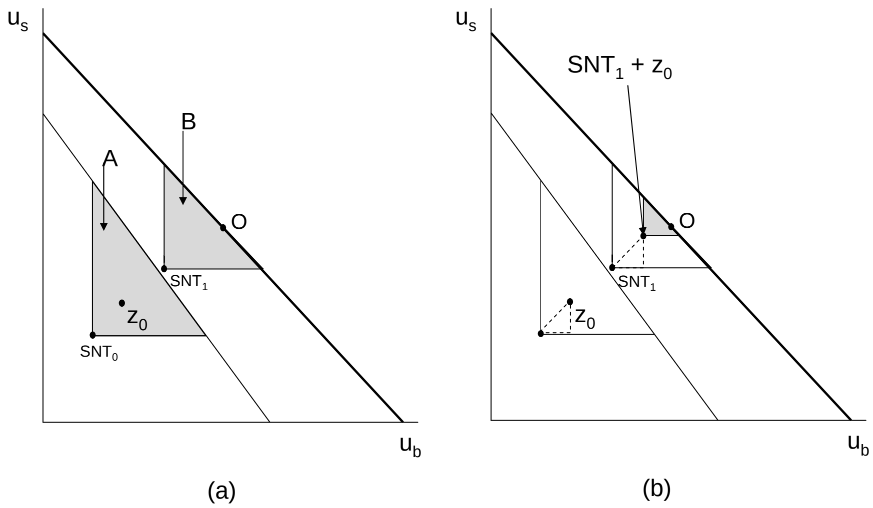


Figure 4.3: (a) A two stage negotiation with intermediary agreements constraining both stages. $SNT_0$ constrains the first stage to the space A, and $SNT_1$ constrains the second stage to the space B. An example of an agent-negotiated solution in the first stage, $z_0$, is also shown. (b) The shaded area shows the constrained solution space for the second stage when we take into account both the intermediary agreement for the stage, $SNT_1$, and the negotiated stage agreement from stage one, $z_0$. Notice how the point $SNT_1$ is pushed towards the frontier by the excess of $z_0$ over $SNT_0$ in the first stage (illustrated by the dashed triangles).

# Chapter 5

# The Bilateral Automated Mediator Negotiation Framework

The Bilateral Automated Mediator (BAM) negotiation framework approaches the problem of bilateral incomplete information negotiations by using a combination of methods and ideas to attack the problem from multiple angles. Due to the inherent difficulties of obtaining accurate information about agent information in the bilateral setting[1], one of the hallmarks of the BAM framework is that the agents are involved in the final negotiating stage in order to act as a self-correcting step in the system. By keeping the agents involved in the ratification of the final agreement, our system differs from a traditional mechanism that takes the input of agents and provides a solution. Instead we take input from the agents, create a temporary solution, and then try to re-utilise the agents' preferences (as expressed by the offers they make and accept or reject) as a secondary input phase that moulds the temporary solution our system provides into something more reflective of the specific negotiation scenario facing the agents. The BAM framework guides agents to a part of the solution space that it thinks is generally a good area to explore and then lets the agents explore it.

We will see later some of the more specific goals of the BAM framework and the particular instantiation of it that we have created, however the general framework has two major goals:

1. To encourage thought and research into hybrid approaches to negotiations, in

[1] Sometimes it is possible to obtain the agents' private information in particular multilateral settings by being able to "play agents off each other" as in Bartal et al. [2].

order to overcome the difficult realities of actual negotiations (that are often assumed not to exist in current negotiation research).

2. To manipulate the negotiation situation in a beneficial manner by creating a negotiation procedure for the agents to work within that guides them to desirable areas of the solution space of agreements.

The first goal is important because a lot of the literature tends to be focused on particular, very specific problems, or are forced to assume that the difficult realities in general problems do not exist in their problem setting. This is always a necessary approach in research to further the field, but we think it is time to re-explore the general *realistic* problem of negotiation, by utilising hybrid approaches that combine different negotiation tools to tackle the overall problem. Our work is a first step in encouraging this line of thinking, and we provide a framework within which these kinds of ideas can be explored.

The second goal is the topic of the rest of this chapter. We will first look at the particular negotiation setting that we will use in our model, then we examine the general BAM framework and how it structures the negotiation process through the application of multiple negotiation tools, then we explore a specific instantiation of the modules of the framework, and finally we mention some concluding remarks about the framework.

## 5.1 Negotiation Setting

The bilateral negotiating environment we study consists of the following:

- Two agents known as the buyer, $b$, and the seller, $s$,
- A set of issues on which to negotiate, $I = \{I_0,\ I_1, \ldots, I_n\}$

The agents are called the buyer and the seller because each issue can be thought of as a good for which a price is being negotiated.[2] For simplicity, we assume the range of prices that can be agreed upon is normalised to be $[0,\ V_j]$ for some constant $V_j$, for issue $I_j$, which is known to both agents.

[2] In reality, an issue can be anything that can be negotiated on. It could be, for example, how many weeks vacation employees get per year. Sometimes, the naming of agents as a buyer and seller confuses people into thinking issues are actually items being exchanged, when it really reflects the idea that the issues will be normalised such that the buyer always wants low prices and the seller always wants high prices for each issue.

**Definition 4 (Agreement)** *For a set of issues,* $I = \{I_0,\ I_1, \ldots, I_n\}$*, an* agreement *is a mutually agreed upon set of prices* $p = [p_0, p_1, \ldots, p_n]$*, one price per issue where* $0 \leq p_j \leq V_j$*. A* partial agreement *is just a subset of an agreement,* $p_{partial} \subset p$*.*

For each issue, $I_j$, each agent has a *private* reservation price, $r^s_j$ for the seller and $r^b_j$ for the buyer. The reservation prices represent the lowest and highest acceptable prices for the seller and buyer, respectively; sellers prefer higher prices on issues, while buyers prefer lower prices. If the agreed upon price for issue $I_j$ be $p_j$, then if $r^s_j \leq p_j$, the seller experiences non-negative utility on issue $I_j$, and if $p_j \leq r^b_j$, then the buyer experiences non-negative utility for this issue. Since agents in our model have the ability to declare what their reservation prices are (and they are not forced to declare their true reservation prices) we label $\hat{r}^s_j$ and $\hat{r}^b_j$ as the declared reservation prices of the seller and buyer, respectively.

The reservation prices of the agents for each issue are drawn from a distribution of possible values, which we will refer to as an agent's reservation price range, or RP range. On a particular issue $I_j$, there is a seller reservation price range of $RP^s_j = [0, x_j]$, and a buyer reservation price range of $RP^b_j = [y_j, V_j]$.[3] It is assumed, without loss of generality, that the prices on issues have been normalised to allow for the seller's range to start at 0 and the buyer's to end at $V_j$. The reservation price ranges are common knowledge to both agents and any third-party; $x_j$ and $y_j$ are known by all parties of the negotiation. The common-knowledge assumption of these distributions reflects the idea that the agents have some idea of what is an acceptable price to their negotiating partner for an issue. For simplicity, we will use uniform probability distributions over these ranges.

Abusing definitions a bit, we will refer to the range between the agents' reservation values as the *balance of trade or trade balance.*[4] The balance of trade is an idea of how much bargaining room there is for the price of a given issue; another way one could think about it is as a profit made on a joint venture. The negotiation process on the issue then becomes how to split the profits between the two agents. There are two forms of the balance of trade depending on whether we are considering the agents' RP ranges (which are common knowledge to all), or whether we are considering the actual reservation prices which are private information held by each agent. When we talk about RP ranges, we will refer to the *perceived* balance

[3]If we are only referring to a single issue, we may sometimes refer to the reservation price ranges using the following values: $RP^s_j = [s_L, s_H]$ and $RP^b_j = [b_L, b_H]$.

[4]We say we are abusing definitions here because the term balance of trade actually refers to the difference between the value of importation and exportation of goods and services. That is, it is a scalar value, not a range.

of trade, and when we talk about the actual reservation prices, we will refer to the *actual* balance of trade.

**Definition 5 (Perceived Balance of Trade)** *The range between the two agents' common-knowledge RP ranges. For issue $I_j$, the seller's RP range is $[0, x_j]$ and the buyer's RP range is $[y_j, V_j]$. When $x_j \leq y_j$, we will write the perceived balance of trade as $(x_j, y_j)$, and when $x_j > y_j$, we will write the perceived balance of trade as $-(y_j, x_j)$.*

**Definition 6 (Actual Balance of Trade)** *The range between the two agents' actual reservation prices. For issue $I_j$, when $r_j^s \leq r_j^b$, we will write the actual balance of trade as $(r_j^s, r_j^b)$, and when $r_j^s > r_j^b$, we will write the actual balance of trade as $-(r_j^b, r_j^s)$.*

For either type of trade balance, we will refer to the *size of the trade balance* as the absolute value difference of the two extreme ends of the range. When a trade balance is non-negative, we will refer to it as a *surplus* (perceived or actual), and when the trade balance is negative we will refer to it as a *deficit* (perceived or actual). Figure 5.1 illustrates the difference between the two definitions.

The perceived deficit is the main problem that occurred in the bilateral trading problem (Section 2.3.2) and results in the Myerson-Satterthwaite Theorem (see Theorem 1 page 15). When agents do not know if there is an actual surplus, it is difficult to create individually rational mechanisms for negotiation because there is no incentive-compatible way to guarantee non-negative utility for both agents; the multi-issue case is very similar. In order to guarantee some theoretical properties later on, we may need to make the assumption that all issues have a perceived (and thus actual) surplus. This is a relatively strong, and sometimes, unrealistic assumption because it means it is possible for both agents to experience non-negative utility on every issue. Not only that, this possibility is known beforehand as common-knowledge, which makes it easier to see where to focus price negotiations on every issue. In our experimental simulation, we drop this assumption, but we also lose some of our theoretical guarantees.

Lastly, agents have different preference weightings for the issues to reflect both the actual *private* preferences of the agents and to account for the normalisation of the issues' agreement ranges; we will refer to these weightings as the agents' valuations. The valuations represent how much utility one normalised unit in the price of an issue is worth to the agent; it is assumed that all valuations are positive (otherwise the seller might want low prices or the buyer high prices, but the issues

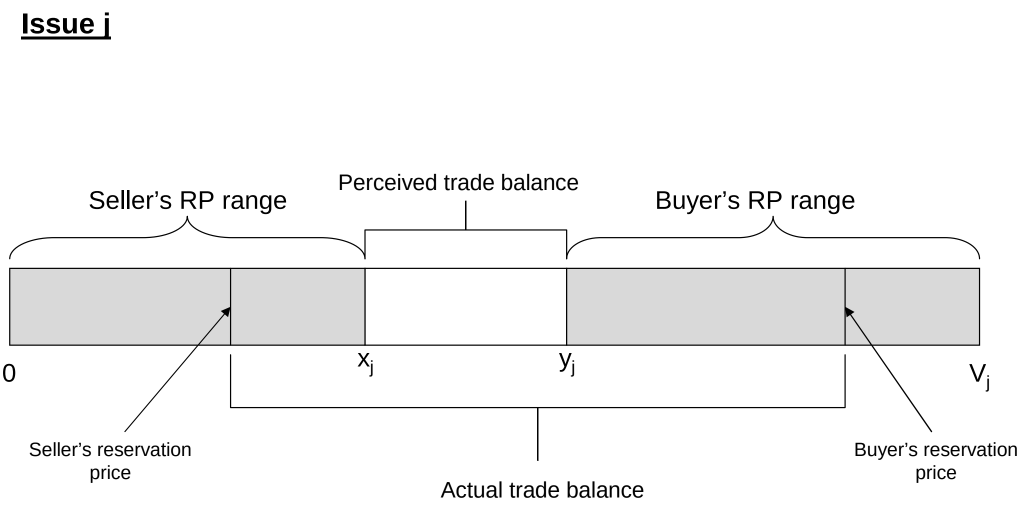


Figure 5.1: An example of trade balances. The reservation price (RP) ranges shaded in the figure represent the possible locations of each agent's respective reservation price. The RP ranges are common knowledge to all participants in the system, and so the *perceived trade balance* is also public knowledge. The reservation prices of the agents, shown as vertical lines within the RP ranges, are the private information of each agent. The *actual trade balance* which is based on these private values is not actually known to anyone in the system since no agent or third-party knows the private information of both agents.

have already been normalised such that neither of these is the case). We denote the seller's and buyer's valuations as $\lambda_j^s > 0$ and $\lambda_j^b > 0$, respectively. We do not, however, presuppose that the relative utilities of agents are comparable; these valuations only represent how an agent internally ranks the issues in importance against other issues. These valuations, like the reservation prices, are drawn from common knowledge distributions (that are uniformly distributed): $\lambda_j^s \in (0, \Lambda_j^s]$ and $\lambda_j^b \in (0, \Lambda_j^b]$. That these distributions are common-knowledge reflects the idea that agents have a general idea of which issues are more important to the other agent.

To summarise the negotiation environment, we define the utility for each agent:

**Definition 7 (Agent Utilities)** *The utility of each agent for an agreement $p = [p_0, p_1, \ldots, p_n]$ on the issues $I = [I_0, I_1, \ldots, I_n]$ is given by*

$$\begin{aligned} &(\textit{seller}) \qquad \sum_{I_J \in I} \lambda_j^s (p_j - r_j^s) \\ &(\textit{buyer}) \qquad \sum_{I_J \in I} \lambda_{I_j}^b (r_j^b - p_j) \end{aligned}$$

*where the valuations $\lambda_j^s$, $\lambda_j^b$ and the reservation prices $r_j^s$, $r_j^b$ are private information.*

Finally, we have one last requirement not usually made for multi-issue negotiations research: that every negotiation ends in agreement. This does not actually affect our model, but instead reflects the fact that there are many types of negotiations, particularly labour disputes, where this is a requirement. The actual effect this has is that our "no-agreement" state is a default agreement, possibly a very inefficient one.

## 5.2 BAM Negotiation Framework

The BAM framework is a negotiation framework that provides an outline for creating a third-party mediator that produces negotiation procedures for agents to follow. The framework is illustrated in Figure 5.2, and consists of the following four modules through which the negotiation flows: the setup, arbitration, protocol, and negotiation modules. The heart of the mediator system consists only of the

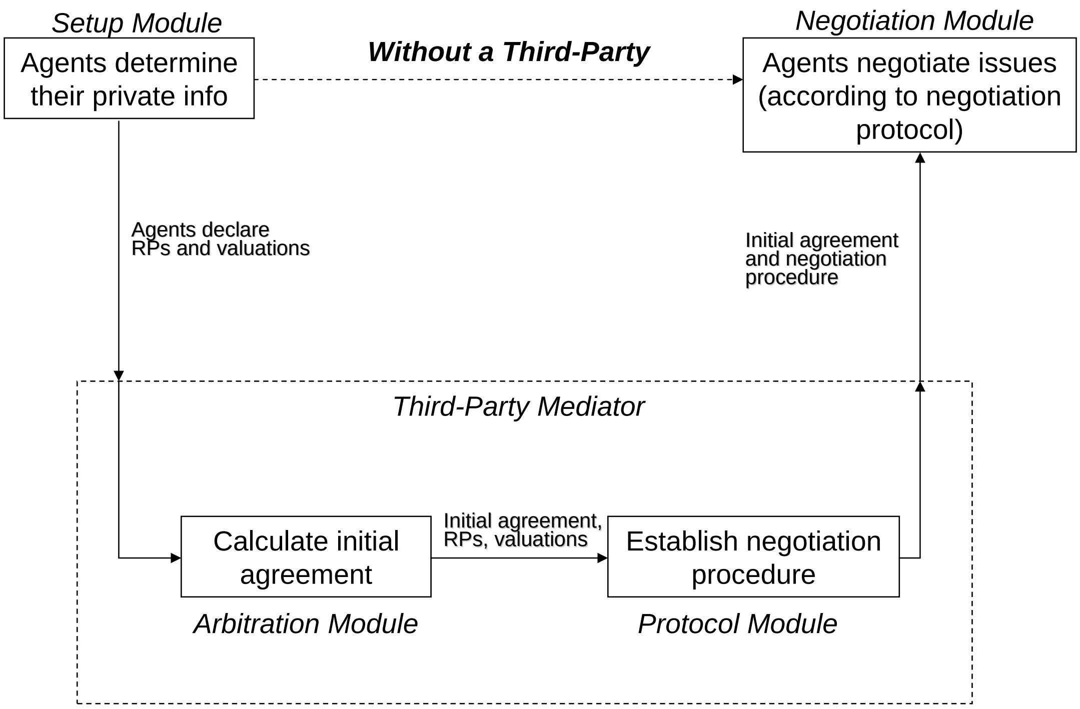


Figure 5.2: The BAM framework showing the information flow of the negotiation process through the mediator system. Agents declare their private information to the mediator which is received by the arbitration module. The arbitration module passes on this information as well as an initial agreement to the protocol module. The protocol module outputs the initial agreement and a negotiation procedure, which together, form the negotiation protocol that is the output of the mediator.

arbitration and protocol modules, the mediator is not actually involved in the other two. We will explain the framework by looking at the flow of the negotiation.

First, agents determine the public and private information of the system in the *Setup Module*. Then agents declare their private information (possibly untruthfully) to the *Arbitration Module*. The arbitration module uses the declared information to come up with a set of agreement prices, one per issue, that forms the initial agreement; the initial agreement forms the default agreement at the beginning of negotiations. The declaration information and the initial agreement are passed to the *Protocol Module*, which must determine the type of negotiation procedure the agents will use. Finally, in the *Negotiation Module*, the agents are left to negotiate on their own, starting with the initial agreement and following the negotiation procedure specified by the mediator.

The main purpose of bringing in a third-party is to ameliorate the difficulties involved in negotiating under uncertainty. Agents have a hard time knowing where to start negotiations and where to make tradeoffs amongst the various issues in order to achieve the best utility. The situation is exacerbated by the common occurrence of agents treating each other as adversaries, since efficiency is lost in the struggle to "win" the negotiation. It will be very difficult for the agents to come up with an efficient agreement if they do not work together to determine the best tradeoffs among the issues. Instead, they will waste all their effort just reaching an agreement and not focus on improving that agreement once it is found. A third-party, by virtue of being an impartial party, or at least being perceived as more impartial than either agent, can achieve things that neither agent trusts the other agent to do.[5] In our framework, the mediator is able to poll the information of the agents and form an initial agreement for the agents to build upon. An initial agreement posed by a mediator is easier for agents to accept as a starting point than an agreement posed unilaterally by either agent. Additionally, having a position to work from helps agents explore where issue tradeoffs exist.

Another major purpose of using a third-party, is to bring parties together and focus them on productive areas of discussion. Initial agreements are a first step in this regard, in fact the whole framework itself is designed for the purpose of guiding agents through the solution space. We do have to acknowledge, however, that much like the bilateral trading problem (Section 2.3.2), it may be impossible to come up with a revelation mechanism that guarantees the acquisition of honest information from the agents. This means that no matter how well we guide the agents and take advantage of our status as an impartial party, our efforts are based on uncertain

[5]In our work, the third party will actually be impartial, striving to be unbiased while trying to help both agents.

data and may need refinement; the negotiation module provides this refinement. As mentioned earlier, agents can refine our solution as a self-correcting step for our inability to know their true preferences and the crudely inefficient measures we need to take because of this lack of omniscience.

To summarise our framework, the mediator guides the agents through the solution space but does not take total control of the negotiation away from them. Instead, the framework allows the agents to modify the mediator's efforts so that they are more accurately aligned with the true preferences of the agents.

We would like to reiterate that this framework is meant for general use in incomplete information scenarios. The only requirements are that the arbitration module come up with some initial agreement and that the protocol module specify the negotiation procedure to be used. The only real constraint is that the negotiation procedure used must support using an initial agreement in some capacity, although this likely leads to an incremental style of negotiation. Negotiations can be quite complex. In separating the mediator module into two sub-modules, the task of mediating an incomplete information negotiation is broken down into smaller subtasks that are easier to manage. Although the most efficient mediation solution is probably a globally considered algorithm, the complexity of such methods make them elusive.

Although there are no other constraints on the framework, the approach we take is to obtain certain properties in the early stages of the mediation process and then maintain these properties while trying to achieve different properties later on. We guarantee individual rationality and one-sided incentive compatibility in the arbitration module's initial agreement, without concern for efficiency. Then, through the use of agendas developed in the protocol module, we try to guide the agents to an efficient and fair solution based on our solution that satisfies IR and one-sided IC. By controlling the format of negotiations, it is possible to encourage certain desirable areas of the solution space to be explored over less profitable areas. By separating our overall mediation goals into modular parts, it is possible that we sacrifice some efficiency that could be obtained through a globally-determined solution, but it also produces a more manageable mediator process that is easier to maintain and update. One of the major benefits of this framework is that local improvement can be made in one module without necessarily having to change other modules, allowing easy implementation of improvements to the mediator.

## 5.3 Our Arbitration Module

As mentioned earlier, the agents' declarations form the input to the arbitration module. The private information that agents declare consist of their reservation prices for each issue, their utility valuations for each issue, and the declaration of "valuation chips" that are used in the valuation chip mechanism described in Section 5.3.2.

The arbitration module will generally consist of some set of mechanisms that determine an initial agreement based on agent declarations. Ideally, we would like these mechanisms to be incentive-compatible (IC) and individually rational (IR), but as mentioned before this is difficult to achieve. Our arbitration module consists of two declaration mechanisms, one that takes reservation prices as input and the other which uses "valuation chips" as input. The reservation price mechanism is called the *Greedy Punishment Solution* (GPS) and the other is called the *Valuation Chips mechanism*. The GPS mechanism is one-sided IC and ex-post or ex-interim IR (depending on whether or not we make the assumption that the trade balance is always a surplus). The valuation chips mechanism is fully IC and ex-post IR.

The interaction of the arbitration module mechanisms is illustrated in Figure 5.3. First the reservation prices are fed into the GPS, which produces an agreement labeled "basis agreement" in the diagram. This agreement is like an intermediary initial agreement which may be modified by the valuation chips mechanism. Besides this intermediary initial agreement, the valuation chips mechanism also receives as input the placement of valuation chips by the agents (described later). The output from the valuation chips mechanism is the actual initial agreement used in the negotiation module.

### 5.3.1 Greedy Punishment Solution (GPS)

The greedy punishment solution is our way of determining the starting base prices for an initial agreement; it takes reservation price declarations from the agents and produces an intermediary agreement that will eventually turn into the initial agreement of the negotiations. We strive to equally divide any surplus or deficit between the two agents so that neither agent is helped or harmed more than the other. However, we do not know what the actual trade balance on any issue is because we cannot guarantee that agents reveal their true reservation prices to the mediator. The idea behind the GPS is that we do not want agents to make declarations that appear greedy. Our definition of being greedy is that an agent makes a reservation price declaration that makes it appear to be in a better bargaining position than

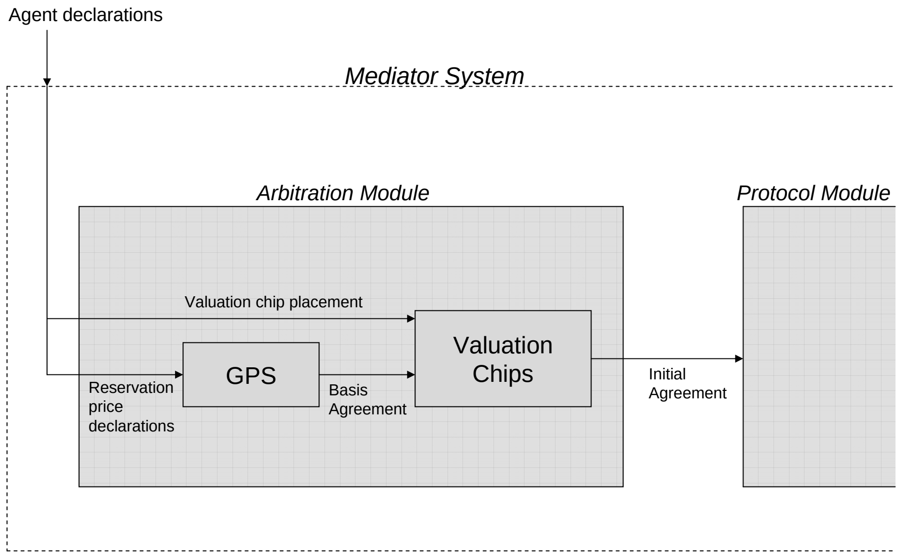


Figure 5.3: The interactions of the arbitration module.

it really is (i.e. its declaration would normally result in a more favourable price for that agent). Consider the following example:

> Typically in a negotiation, parties make initial offers which are unacceptable to each other and slowly move their offers towards each other until an agreement is reached. This sort of procedure often results in an agreement at the midpoint of the two initial offers, so if an agent makes a more extreme initial offer it is actually to its benefit.
>
> Say the seller and buyer are negotiating over a single issue with the seller's *actual* reservation price being 20 and the buyer's *actual* reservation price being 80 (that is $r^s = 20$ and $r^b = 80$). If each agent declares their reservation prices truthfully ($\hat{r}^s = r^s$ and $\hat{r}^b = r^b$), then the agreement price will be around $\frac{20+80}{2} = 50$. But if the buyer claims his reservation price is $\hat{r}^b = 50$, then the likely agreement price is going to be $\frac{20+50}{2} = 35$. And even if the price did not end up being as low as 35 with $\hat{r}^b = 50$, the price will probably still be lower than the true "fair" price of 50. The seller can also benefit by lying and saying its reservation price is higher than it actually is. Now the problem becomes that both agents have to strategically lie so that they are not cheated

by the other agent's lie.

Now consider what happens if the buyer actually declares a higher reservation price, say $\hat{r}^b = 100$. Now the agreement price becomes $\frac{20+100}{2} = 60$. This sort of declaration actually hurts the buyer because the buyer wants a lower price; the buyer is certainly not being greedy here.

**Definition 8 (Greedy)** *A buyer is being greedy when $\hat{r}^b < r^b$ and a seller is being greedy when $r^s < \hat{r}^s$. That is, a buyer declares a lower price than its actual reservation price and a seller declares a higher price than its actual price.*

How can we stop agents from acting strategically with their reservation price declarations? Ideally, we just want them to tell the truth, but perhaps we can settle for making it so that they are not acting "greedy". The GPS does this by making it painful for agents to be greedy, and beneficial for them to not be greedy. This is not the same as the agents being truthful however, instead it means that lying in the direction of being greedy will harm an agent and lying in the opposite direction will benefit an agent (this is the idea behind one-sided IC). Therefore it is in an agent's best interest to declare the "weakest" reservation price possible, which is at the extreme end of its reservation price range. We still will not know the true value of the reservation prices, but this type of declaration is sort of like the agent saying that they are flexible with the price and gives the mediator room to work with in the trade balance.

The Figure 5.4 will help us illustrate the specific implementation of the *continuous* GPS.

First observe that the starting price, $p$, is in the exact middle of the trade balance, S, which means each agent has received the benefit of half of the surplus $S = y_j - x_j$. Let the agents' share of the surplus be $s^s = \frac{S}{2}$ for the seller and $s^b = \frac{S}{2}$ for the buyer. The GPS gives at least this amount of the surplus to an agent if that agent does not act greedily, but if it does act greedily then portions of this surplus will instead be given to the other agent (meaning the price $p$ will shift). We declare two values that are sort of like greediness percentages for each agent:

$$G^s = \frac{\hat{r}^s - 0}{x_j - 0}$$

$$G^b = \frac{V_j - \hat{r}^b}{V_j - y_j}$$

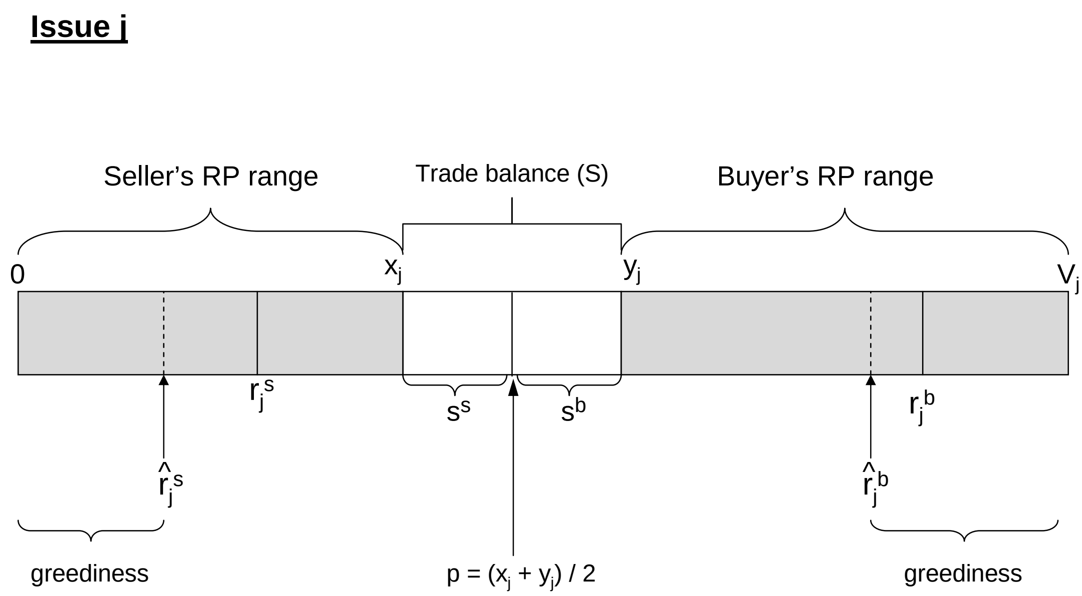


Figure 5.4: The RP range of the seller and buyer are shown and are $RP_j^s = [0, x_j]$ and $RP_j^b = [y_j, V_j]$, respectively. This makes the perceived balance of trade $(x_j, y_j)$, which is a surplus since $x_j < y_j$. The actual reservation prices are $r_j^s$ and $r_j^b$ and the declared reservation prices are $\hat{r}_j^s$ and $\hat{r}_j^b$. The "greediness" of each agent has also been labeled showing how far away from the extreme ends of the price range they declared as their reservation price.

These fractions represent how greedy an agent was with respect to the size of its surplus. These values also represent how much of its share of the surplus an agent will lose due to greediness; the higher these values are, the more an agent loses. Explicitly, the new price for issue $I_j$ using the GPS mechanism is:

**GPS Price Adjustment**

$$p_j = p - s^s G^s + s^b G^b \tag{5.1}$$

We see in the price adjustment equation that the greedier a seller is, the lower the price goes, and the greedier a buyer is, the higher the price goes.[6] This price adjustment is performed on every issue and this gives us the intermediary agreement that will be given to the valuation chips mechanism.

### Incentive Compatibility and Individual Rationality

We mentioned earlier that the GPS is one-sided IC and possesses different levels of IR depending on our assumptions. We have already seen what one-sided IC means for declaring reservation prices: it is in an agent's best interest not to appear greedy, and in fact being as modest as possible is the best thing to do. IR means that agents receive non-negative utility, and this occurs if the price for an issue satisfies $r_j^s \leq p \leq r_j^b$. We are always able to guarantee ex-ante IR, and we have ex-post IR if we assume that no issue has a *perceived* deficit (see definition 5).

Before we prove these results, there is one important thing to note about the GPS mechanism: it treats all issues separately. This means that in order to prove that the GPS satisfies the properties mentioned, we only need to prove it for a single issue since there is no interaction between issues with respect to price adjustments in the GPS.

**Theorem 2** *The GPS mechanism produces an agreement that is one-sided incentive-compatible.*

**Proof** We will only prove this for a single issue as that is all that is required due to the independence of the issues.

[6] Of course, these two adjustments counteract each other so that if each agent is equally greedy ($G^s = G^b$) then the price does not change. However, we will see later that there is no advantage in trying to strategically lie based on the other agent's greediness.

Consider the price adjustment equation for some issue $I_j$: $p_j = p - s^s G^s + s^b G^b$. The seller would like $p_j$ to be as high as possible, and since $p$, $s^s$, and $s^b$ are fixed, the only way that can happen is if $G^s$ is small and $G^b$ is large (the seller must not be greedy and the buyer must be greedy). The seller has no control over what $G^b$ will be, so it must focus on making $G^s$ as small as possible. $G^s = \frac{\hat{r}_j^s - 0}{x_j - 0}$ is minimised if $\hat{r}_j^s = 0$, or in other words, *if the seller lies as much as possible in the direction that gives it a weaker bargaining position*, which is the definition of one-side incentive compatibility in this situation. The seller has no incentive to be greedy since that would just raise $p_j$ and reduce the seller's utility.

The argument is analogous for the buyer and thus one-sided IC is satisfied by the GPS.

□

**Theorem 3** *If every issue has a perceived surplus, then the GPS mechanism is ex-post individually rational.*

**Proof** (We only prove this for an arbitrary issue $I_j$, as required.)

Because of the way the GPS performs its price adjustment for issue $I_j$ ($p_j = p - s^s G^s + s^b G^b$), the price $p_j$ will never be outside of the trade balance no matter what the agents declare. This means that $x_j \leq p_j \leq y_j$, noting that our assumption of a perceived surplus implies $x_j \leq y_j$. By definition of our negotiation environment, $r_j^s \leq x_j$ and $y_j \leq r_j^b$, and thus we have that $r_j^s \leq p_j \leq r_j^b$. Since $(p_j - r_j^s) \geq 0$ and $r_j^b - p_j \geq 0$, we can see from the definition of the agents' utilities (Definition 7 on page 48) that each agent will have non-negative utility for issue $I_j$ (and thus non-negative utility over all issues).

□

**Theorem 4** *The GPS mechanism is ex-ante individually rational.*

**Proof** (We only prove this for an arbitrary issue $I_j$, as required.)

We have two cases: (1) the issue has a *perceived* surplus and (2) the issue has a *perceived* deficit. If the issue has a surplus, then the proof is the same as in the Theorem 3 since ex-post IR is a stronger condition than ex-ante IR. If the issue has a deficit, then we have to consider the expected value for the price of the issue given by the GPS.

To assist in explaining the second case, we provide Figure 5.5 which is an issue with a perceived deficit. In the diagram we use the values $s_L$ and $b_L$ for the low end of the seller and buyer's reservation range, respectively, and similarly we use $s_H$ and $b_H$ for the high end of the ranges.

We will focus on the seller's expected utility for the issue, where the probability of any given reservation price in its RP range is $\frac{1}{s_H - s_L}$ (due to the uniform probability distribution of the seller's reservation price). Then the seller's expected utility is:

$$
\begin{aligned}
&\Rightarrow \int_0^p (p-x)\frac{\lambda}{(s_H - s_L)}dx + \int_p^{s_H}(x-p)\frac{\lambda}{(s_H - s_L)}dx \\
&\Rightarrow \frac{\lambda}{(s_H - s_L)}\left[\int_0^p (p-x)dx + \int_p^{s_H}(x-p)dx\right] \\
&\Rightarrow \frac{\lambda}{(s_H - s_L)}\left[\left(px - \frac{1}{2}x^2\right)|_0^p - \left(\frac{1}{2}x^2 - px\right)|_p^{s_H}\right] \\
&\Rightarrow \frac{\lambda}{(s_H - s_L)}\left[\left(p^2 - \frac{1}{2}p^2\right) - \left(\left(\frac{1}{2}s_H^2 - ps_H\right) - \left(\frac{1}{2}p^2 - p^2\right)\right)\right] \\
&\Rightarrow \frac{\lambda}{(s_H - s_L)}\left(ps_H - \frac{1}{2}s_H^2\right) \\
&\Rightarrow \lambda(p - \frac{1}{2}s_H)
\end{aligned}
$$

For the seller to have a non-negative expected utility we have:

$$
\begin{aligned}
\lambda(p - \frac{1}{2}s_H) &\geq 0 \\
&\Rightarrow p \geq \frac{1}{2}s_H
\end{aligned}
$$

An analogous calculation for the buyer shows that for ex-ante IR, $p \leq \frac{1}{2}(b_H + b_L)$. So our full constraint for satisfying ex-ante IR is:

$$\frac{1}{2}s_H \leq p \leq \frac{1}{2}(b_L + b_H)$$

This constraint is always satisfiable since $(b_L + b_H) - s_H \geq 0$, because $s_H \leq b_H$ and $0 \leq b_L$ $(b_H - s_H) + b_L \geq 0$. And it is satisfied by setting $p = \frac{1}{2}(s_H + b_L)$, which

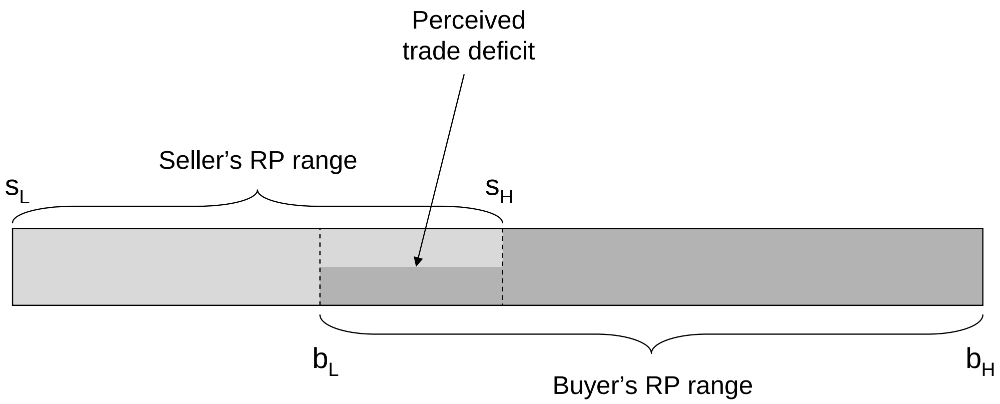


Figure 5.5: An example of an issue with a perceived deficit. This diagram is used to help explain the proof of Theorem 4. Instead of using [0, $x_j$] and [$y_j$, $V_j$] for the seller and buyer's reservation price ranges, they are labeled by the ranges [$s_L$, $s_H$] and [$b_L$, $b_H$], respectively. This is to indicate that they are low and high prices in the ranges, which makes the proof easier to read.

is the midpoint of the balance of trade. This is exactly where the GPS solution will place the price for this issue as long as agents act rationally and do not attempt to be greedy.

□

### 5.3.2 Valuation Chips Mechanism

In the GPS mechanism, we focused on the private information of reservation prices. Here we focus on the valuation of issues. Recall that a valuation of an issue represents how much utility an agent receives for one unit of price in that issue. Ideally, we want to know the relative importance of the issues for each agent, meaning how much more important is one issue over another to an agent. We do not care so much about the exact valuations, as the scales used have no meaning for the mediator, but instead we would like to obtain from each agent their ordering of issues by importance and the amount of separation of importance between issues in the ordering (this is illustrated in Figure 5.6). At the moment, the mechanism we use only achieves the former goal of determining the order of the issues for each agent,

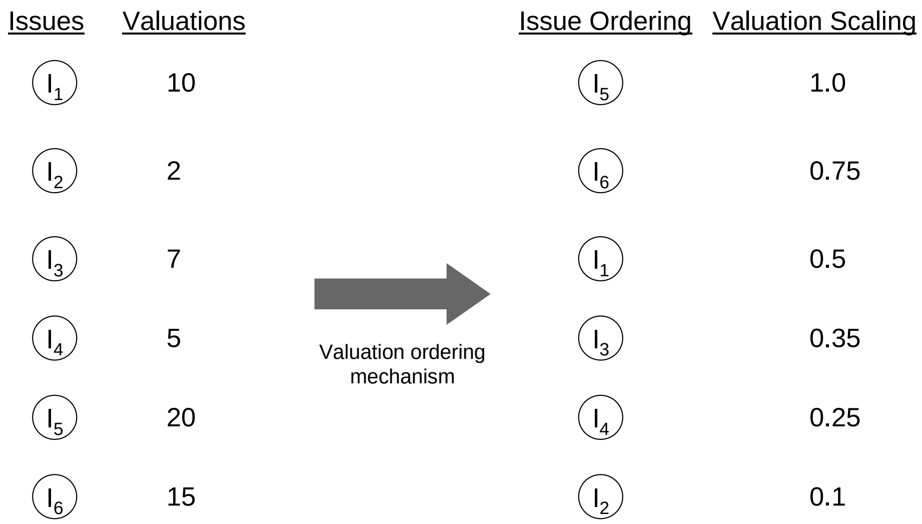


Figure 5.6: An example of the information we would like to acquire from an agent about its valuations. We see on the left, the set of issues with valuations for an agent. We would like to transform this information into the sorted ordering and scale of valuations shown on the right. Currently our valuation ordering mechanism determines the ordering but without the valuation scaling.

but without an idea of the relative importance of issues within the ordering; we know that one issue is more important to an agent than another, but not by how much.

How the valuation chips mechanism works is by giving each agent a set of chips equal to the number of issues. Each chip in a set has a different value on it and represents how much influence that chip has. The agents must assign one chip to each issue as part of their declaration to the mediator. The chips then affect the current prices of the agreement made by the GPS mechanism. The greater a value a chip has, the greater it influences the price of the issue it is assigned to in favour of the agent that assigned that chip. The assignment of chips then allows agents to express their preferences for the issues by giving them an opportunity to adjust the initial agreement in a way that reflects their preferences.

**Valuation Chips Mechanism** Given a negotiation as specified in Section 5.1 with $n$ issues, each agent is given a set of $n$ chips, each set being a duplicate of the other, with all the chips within a set having different values. Let the chips given to the seller be denoted by $C$ and the chips given to the buyer denoted by $D$, and we will denote by subscripts the $k^{th}$ ordering of chips (such that $c_1$ is the smallest chip and $c_n$ is the largest chip in set $C$).

An arbitrary value called the chip shift amount, $H$, is chosen by the mechanism that will be the largest amount an agent can shift the price of an issue using its valuation chips (the value of $H$ is the same for all issues and for each agent). Knowing $H$, each agent must assign its set of chips to the issues, one chip per issue. The declaration of chip placement is declared to the mediator privately so that neither agent can base their assignment on the other agent's actions. We will denote chips assigned to issue $I_j \in I$ with superscripts, such that the chips assigned to $I_j$ are $c^j$ and $d^j$ (subscripts will still be used to denote order within a chip set). The prices of the intermediary agreement from the GPS mechanism are then modified as follows:

> For each issue $I_j$, the price $p_j$ of the current agreement is adjusted by:
>
> $$\frac{c^j}{c_n}H - \frac{d^j}{d_n}H$$
> $$\Rightarrow \frac{H}{c_n}(c^j - d^j)$$
>
> Note that the seller's chip adds to the price of the issue and the buyer's chip subtracts from the price since a seller want a higher price and a buyer wants a lower price.

The resulting agreement that has been adjusted for all issues becomes the initial agreement that is the output of the arbitration module.

Before we prove some of the properties guaranteed by this mechanism, there are a few points to note. First, the largest chip in the chipset allows an agent to shift a price by the full shift amount, $H$, and the smaller chips allow shifts of smaller degrees. Second, the shift of price for an issue is based on the difference between the chips agents use for the issue; this reflects that an issue more important to one agent will be shifted in favour of that agent. For example, if the seller puts $c_n$ on an issue and the buyer $d_1$, then the issue is much more important to the seller than the buyer and the price on the issue will be shifted to reflect this. Third, we have not specified how to choose $H$ or whether there are any restrictions on it. The only real restriction on $H$ is the following:

$$H \leq \min_{I_j \in I}(\min(p_j, V_j - p_j)) \tag{5.2}$$

This says that $H$ can be no larger than the smallest distance of the current price to either end of the price range over all issues. The rationale for this constraint is that we do not want the price on any issue to be able to be shifted outside of the price range $[0, V_j]$. The method we actually use for choosing $H$ is to choose the smallest positive distance of the current price to either end of the *trade balance* over all issues. If the trade balance for issue $I_j$ is $(t_L, t_H)$, then our $H$ is:

$$H = \min_{I_j \in I}(\min(p_j - t_L, t_H - p_j)) \ where\ H \ > \ 0 \tag{5.3}$$

The rationale for choosing the chip shift amount like this is so that the shifted price stays within the trade balance on as many issues as possible while still having the shift amount be positive. In peculiar circumstances where many of the issues have prices on the edge of the trade balance after the GPS (meaning their smallest distance from their current price to the end of the trade balance is 0), it may be desirable to also put a maximum value on H in order to maintain as many prices as possible staying in their respective trade balances.

We mentioned earlier that our valuation chips mechanism is fully incentive-compatible and ex-post individually rational. In this setting, this means that an agent has no reason to lie about its ordering (i.e. it has no reason not to put its largest chip on its most valued issue, and then its next largest chip on its next most valued issue and so on) and an agent will not experience negative utility from participating in this mechanism *no matter what the other agent does.* In proving these results we show that an agent can affect its utility most positively by telling the truth about its ordering and that nothing the other agent can do can make the utility change resulting from this mechanism be negative *if the agent assigns its chips truthfully.*

Just before we prove our results about the valuation chip mechanism, we want to introduce the following useful definition:

**Definition 9 ((Valuation Chip) Utility Influence on issue $I_j$)** *For an issue $I_j$, the* (valuation chip) utility unfluence *is:*

$$\frac{H}{c_n}\lambda_j^s(c^j - d^j)$$

*and*

$$\frac{H}{c_n}\lambda_j^b(c^j - d^j)$$

*for the seller and buyer, respectively.*

*We will refer to the summation of these equations over all issues as just the* utility influence *for the seller and buyer, respectively.*

**Theorem 5** *The valuation chips mechanism is incentive compatible, providing the issue ordering of each agent, which is unique except for when there are issues that are valued equally.*

**Proof** WLOG we will assume we are considering the seller as our agent; the buyer's proof is analogous. Assume that the seller has assigned its chips such that its largest chip is on its most important issue, its second largest chip is on its second most important issue, and so on. Now, by way of contradiction assume that the seller can increase its overall utility by swapping the chips for two issues that it does not value equally.

Of the two issues having their chips swapped, let the higher valued issue be issue $I_j$ and the lower valued issue be $I_k$, and let the chips assigned to these issue in the original ordering be $c^j$ and $c^k$, respectively. This means that $c^j > c^k$ and $\lambda_j > \lambda_k$[7].

Now consider how much these two issues influence the agent's utility in the original chip assignment:

$$\frac{c^j}{c_n}H\lambda_j + \frac{c^k}{c_n}H\lambda_k \tag{5.4}$$

$$\Rightarrow \frac{H}{c_n}(c^j\lambda_j + c^k\lambda_k) \tag{5.5}$$

And then the utility influence of the swapped chip assignment:

$$\frac{c^k}{c_n}H\lambda_j + \frac{c^j}{c_n}H\lambda_k \tag{5.6}$$

$$\Rightarrow \frac{H}{c_n}(c^k\lambda_j + c^j\lambda_k) \tag{5.7}$$

[7]Recall that $\lambda$ represents the agent's valuation for an issue, the units of utility per unit price on that issue.

We only need to consider the influence on these two issues since the chip assignment is the same for all the other issues. For our assumption about the seller being able to increase its utility above that of the original assignment, we must have that (5.6) > (5.4):

$$\begin{aligned}
\frac{H}{c_n}(c^k\lambda_j + c^j\lambda_k) &> \frac{H}{c_n}(c^j\lambda_j + c^k\lambda_k) \\
\Rightarrow (c^k\lambda_j + c^j\lambda_k) &> (c^j\lambda_j + c^k\lambda_k) \\
\Rightarrow c^k(\lambda_j - \lambda_k) &> c^j(\lambda_j - \lambda_k) \\
\Rightarrow c^k &> c^j
\end{aligned}$$

However, this results in a contradiction because we know that $c^j > c^k$, and therefore our assumption that the seller can increase its overall utility by swapping the chips of two unequally valued issues must be impossible. Since we did this with two arbitrary unequally valued issues, we see that it is not possible for the seller to swap its chip assignment from a truthful assignment and receive more utility.

□

**Theorem 6** *The valuation chips mechanism is ex-post individually rational if the agents assign their chips in a truthful manner with respect to their preference ordering of the issues.*

**Proof** The key to proving this result is to show for one agent, as long as it assigns its chips truthfully revealing its true issue preference ordering, that there is no chip assignment that can be made by the other agent that will result in the first agent receiving less utility in the agreement after the valuation chip mechanism than in the agreement before it (the agreement that comes from the GPS).

WLOG let us assume we are considering the seller as the agent under consideration, and let the seller assigns its chips truthfully to the issues. We will show that even if the buyer maximises its negative utility influence on the seller's utility, the seller will still not have negative utility overall. In this case the buyer wants the seller's utility influence to be minimised:

$$\min(\sum_{I_j \in I} \frac{H}{c_n}\lambda_j^s(c^j - d^j)) \tag{5.8}$$

The only control the buyer has over this summation is what values $d^j$ have, therefore in order to minimise Equation 5.8, the buyer really only needs to focus on maximising the following:

$$\sum_{I_j \in I} d^j \lambda_j \tag{5.9}$$

But this is the same thing as the seller trying to maximise its share of the its own utility influence! And by our proof of Theorem 5, we saw that the seller maximises its utility influence by assigning its chips truthfully. We see then that the chip assignment the buyer must make to minimise Equation 5.9 is the same chip assignment(in terms of chip values) that the seller makes when making a truthful assignment. Since this theorem assumes that the seller is playing truthfully, this means that for every issue $I_j$ we have $c^j = d^j$ and Equation 5.8 equals 0. But this resulted from the buyer minimising the seller's utility influence as much as possible and thus we see that if the seller assigns its chips truthfully, it is impossible for it to experience less utility in the agreement after the valuation chip mechanism has run than the agreement before the mechanism has run.

□

Now that we have explained the mechanism and proved its properties, we should really ask ourselves why do we want the issue ordering in the first place? What good is this ordering when we do not know the amount of spacing of importance between issues? First of all, when we proceed to the protocol module, we need to have some idea of relative utilities for issues in order to create agendas. Although it would be ideal to know the spacing between issues, it is still useful to know just the preference orderings of the agents. A second use that we have not yet exploited, is that this information can be used to adjust the initial agreement prices even more to reflect the issue orderings by performing price adjustments across multiple issues, trading off on issue importance asymmetries so that the starting initial agreement is even closer to what each agent wants. However, we must be careful taking any actions based on the issue orderings retrieved from the valuation chips mechanism because we can jeopardise its ability to satisfy IC and ex-post IR.

### Constraints Spawning from the Valuation Chips Mechanism

For the valuation chips mechanism to work, the chips in a chipset must maintain their ordering throughout the processing of the mediator system. That is, we must maintain the following for the values of the chips:

$$c_i < c_{i+1}, \qquad 1 \leq i \leq n \tag{5.10}$$
$$d_i < d_{i+1}, \qquad 1 \leq i \leq n \tag{5.11}$$

It may seem odd to specify this constraint since the values of the chips do not actually change, however, their effective value can change if we modify the agreement after the valuation chips mechanism runs. Say we wanted to change some of the prices in the agreement such that they are better aligned with the issue orderings we get from the valuation chips mechanism. By changing the prices it is the same thing as changing the values of the chips that are assigned to those issues. For example, say we have an issue $I_j$ with a price adjustment (from the valuation chip assignments) that is worth

$$\frac{H}{c_n}(c^j - d^j)$$

and now we want to change the price on this issue by one positive unit. This effectively makes the perceived price adjustment

$$\begin{aligned}
&\frac{H}{c_n}(c^j - d^j) + 1 \\
&\qquad\qquad \Rightarrow \quad \tfrac{H}{c_n}(c^j - d^j) + \tfrac{Hc_n}{Hc_n} \\
&\qquad\qquad \Rightarrow \quad \tfrac{H}{c_n}(c^j - d^j + \tfrac{c_n}{H})
\end{aligned}$$

By changing the price by one unit, we have changed the difference between the buyer's and seller's chips by $\frac{c_n}{H}$ on this issue. Although this change in chip value should be distributed across both agent's chips, from either agent's point of view they just know the price changed by one unit and the value of their chip (either $c^j$ or $d^j$) has effectively changed by the full amount $\frac{c_n}{H}$. To summarise, if we raise the price of an issue in the agreement we make the seller's chip on that issue worth more and we devalue the buyer's chip according to the following:

$$\begin{aligned}
c_{i_{(new)}} &= c_{i_{(old)}} + \frac{c_n}{H}t_i \\
d_{k_{(new)}} &= d_{k_{(old)}} + \frac{c_n}{H}t_k
\end{aligned}$$

where $t_i$ is the amount of the price shift on the issue that chip $c_i$ has been assigned to, and $t_j$ is similar but for the buyer's chip.

Now, we turn to the question of why is it important to maintain the constraint in Equations 5.10 and 5.11? The whole mechanism hinges on the idea that the agents know the order of their chipsets, and if the order changes after the agents have assigned their chips then they may not want to assign them in the same way. If agents think that the ordering of the chipsets may effectively change later on, then they are uncertain how they should act: they no longer know what is the rational course of action. This jeopardises our guarantees of IC and ex-post IR, which is something we do not want to occur, and this is why we want to maintain the constraint in 5.10.

### 5.3.3 Summary of Arbitration Module

The table in Figure 5.7 summarises the properties of the arbitration module and the mechanism that compose it. Note that although both the GPS and valuation chip mechanism affect the prices in the initial agreement output by the module, each mechanism actually takes different declarations as input. The GPS takes the reservation price declarations as input, whereas the valuation chips mechanism takes the assignment of valuation chips to issues as its declaration input.

## 5.4 Our Protocol Module

Once the arbitration module has created an initial agreement, this agreement and all other information the system has is fed into the protocol module. The protocol module is responsible for coming up with the negotiation procedure that will be used in the negotiation module. The instantiation of this module that we use will create an agenda that will guide the negotiations. The actual negotiation procedure, beyond the agenda, that we use is not as important and is explained in Chapter 6.

The method we use to create our agenda is geometrical, using utility space representations of the negotiation environment to produce an agenda that fits our desires of coming up with a fair and efficient final agreement. There are many ways to extend the use of agendas in our system and some of these are explained in Section 7.3, but here we introduce the foundational concepts in our approach to using agendas.

| Mechanism | Incentive Compatible? | Individually Rational? |
|---|---|---|
| Greedy Punishment Solution (GPS) | One-sided IC, w.r.t reservation price declarations | Ex-ante * |
| Valuation Chip Mechanism | Yes, w.r.t valuation chip declarations | Ex-post |
| **Arbitration Module** | **Yes, cumulative properties of above entries** | **Ex-ante *** |

* = ex-post if every issue has a trade surplus

Figure 5.7: A summary of the game-theoretic properties of the arbitration module and its inner mechanisms.

### 5.4.1 Translating Issue Orderings to Utility Space

To use utility space in creating agendas, we need some way to convert the issues into their respective utility representations. This amounts to deciding for each issue $I_j$, how much utility that issue *could* represent to each of the buyer and seller. The problem here comes back to the fact that we do not have exact information about the valuation of issues for the agents nor what their reservation prices are, so we will have to make some sort of estimation. Our estimation is based on two pieces of information:

1. The *perceived* trade balance on the issue.
2. The issue preference orderings of each agent (from the valuation chips mechanism).

We use the perceived trade balance as if it where the whole price range such that if the range is $(x_j, y_j)$, then a price $p_j = x_j$ would give the seller 0 utility and a price $p_j = y_j$ would give the buyer 0 utility. Note however, that we do not actually need to know what $p_j$ is in the initial agreement to do this, we are just using the trade balance as part of the basis for how much the issue is worth to the agents. If an issue has a large trade balance, then it is likely worth a lot to the agents.

To obtain an estimation for the agents' valuations of the issues, we need to convert the preference orderings into valuations for each agent. We only have the issue ordering for each agent, but now we have to assign some value for a utility per price unit for each agent for each issue. In order to create these valuation estimations, we need to have some separation between the issues in the ordering and because we have no information about what the actual separations might be, we will just evenly separate all the valuations. For example, if we have three issues in an agent's issue ordering, we might assign the highest issue a valuation of 10 utility/unit, the next highest issue 6 utility/unit, and the last issue 2 utility/unit.

Determining the amount of utility an issue will be worth in our utility space is now simply an issue of taking the size of the trade balance and multiplying it by the valuation estimations we have given to each agent for that issue. One thing to note, that may not be obvious yet, is that the total amount of utility an issue is worth to an agent is not the same as the agent's estimated valuation; when ordering the issues by total utility value to an agent, this ordering may not be the same as the agent's issue preference ordering that is based on its actual valuations. To see an example of this and a general example of converting issues into utility space, consider Figure 5.8.

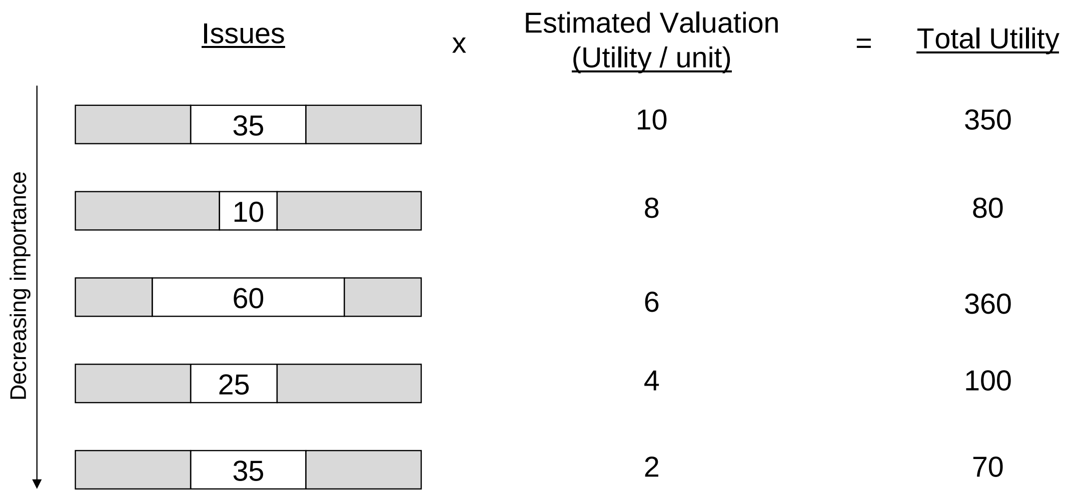


Figure 5.8: An example of converting preference orderings into utility for creating a utility space. The numbers in the issues are the sizes of the perceived trade balance (see Definition 5 for an explanation of how the perceived trade balance is determined for an issue). The estimated utility starts by assigning an arbitrary valuation to the most important issue and then decrementing this for each successive issue. Note that due to the different sizes of the trade balances, it is possible for an issue of low importance to actually be worth more, in total utility, than an issue of higher importance. This can be seen with the first issue being worth 350 total utility, which is less than the 360 total utility that the third issue is worth, even though the first issue has a higher valuation.

Before we move on to elements of the agenda creation, we should specify our actual transformation for turning issue preference orderings into estimated valuations. Let us consider one agent's issue ordering and for convenience let us say it is $\{I_0, I_1, \ldots, I_{n-1}\}$, where issue $I_0$ is the agent's most important issue. We will denote our estimated valuation for issue $I_j$ by $\lambda'_j$, and the size of the perceived trade balance as $|S_j|$. Finally, let $u_0$ be some arbitrary amount of utility that we are going to say issue $I_0$ is worth to the agent. Then the estimated valuations for the issues are:

$$\begin{aligned}
\lambda'_0 &= \frac{u_0}{|S_0|} \\
\lambda'_i &= \lambda'_0 - \left(\frac{\lambda'_0}{n}\right) i, \quad 1 \leq i \leq n-1
\end{aligned}$$

We can see from this that we just make sure that each issue in the ordering has a smaller valuation than the one previous to it. We might also note from this that we assign each agent the same estimated valuations.

### 5.4.2 Agenda Preliminaries

As mentioned earlier, we would like our agendas to guide agents to useful areas of the solution space so they are more likely to negotiate an agreement that is fair and efficient. What we will do is choose a solution in the utility space that we think is a good solution, and then focus our efforts on making an agenda that tries to direct the agents to that solution; our chosen solution will be the Nash Bargaining Solution (Section 2.4.4). We would like to re-iterate however, that our agendas are based on estimated information for the utility space and thus our chosen solution in our estimated space really is a guide to the area where a good solution probably exists.

In order to understand how we will create our agendas in a way that leads agents to particular areas of the solution space, we need to consider a few concepts first.

#### The Yellow Brick Road: The Solution Line

The negotiation protocol that we are going to use will separate the negotiations over several stages as in Figure 2.4. To be able to guide agents to a particular area of the utility space over the course of these stages we need some way to focus the agents' attention throughout this process, and not just near the end of the

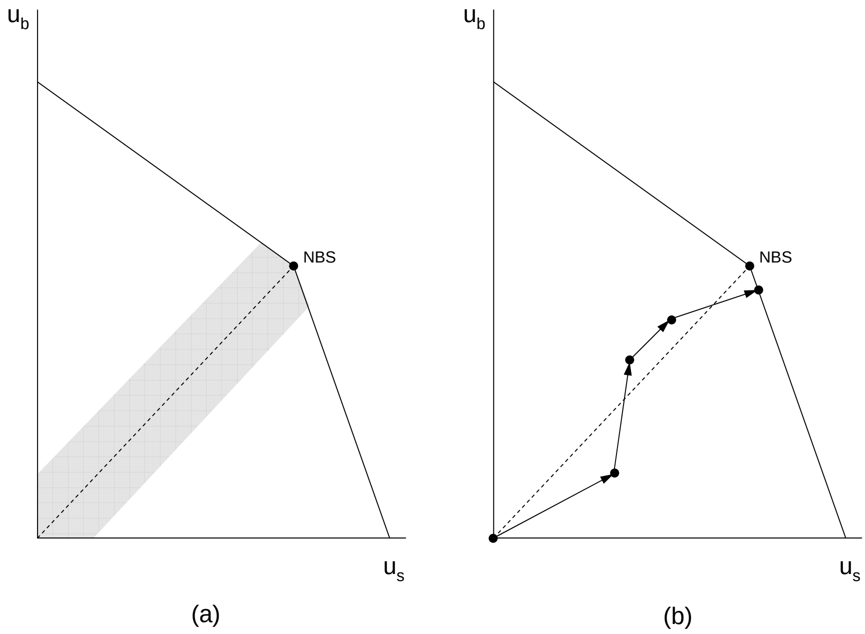


Figure 5.9: The solution line is just the line between the origin and our chosen solution concept, in this case the NBS. (a) We want to guide agents along the solution line, the shaded region acts as a "focus space" that we want to try and keep the intermediary agreements within throughout the negotiation process. (b) An illustration of how a negotiation might proceed if agents are guided along the solution line. The intermediary agreements will typically zig-zag around the solution line, ending up at a solution near both the efficient frontier and NBS solution.

negotiations. If agents stray too far of course early on in the negotiations, it may be impossible to reach the desired solution area later on.

The idea we use for keeping agents focused is simple. We consider a part of the utility space from the beginning of the negotiations to the end of the negotiations that acts as a path towards the area of the space that we want agents to end up in. This path is based on the solution concept we use for the entire space, in our case the NBS. The *solution line* is just the line from the no agreement point (the origin) to our chosen solution, the NBS. Throughout the negotiation process, through the various stages, we want to continually guide agents along the solution line towards the efficient frontier. This idea is shown in Figure 5.9.

### Desired Solutions

Ultimately, our desired solution is the global solution concept towards which we want to guide the agents. To achieve this we need to focus agents along the solution line, but how do we do this throughout the agenda? We break down this problem into subproblems at each stage, such that we have a desired solution at each stage and we create our agenda in a manner that it will guide agents to the desired solution at each stage.

**Definition 10 (Desired Solution)** *The* desired solution *at each stage is the local agreement for the stage that we, as a mediator, would like agents to achieve in order to be aligned with our desired global solution. The desired solutions for the stages are simply the intersection points of the solution line with the efficient frontier of each stage's utility space.*

Since we use the NBS as our global solution concept, we will denote the desired solution for a stage $i$ to be $NBS_i$, and the global solution $NBS_n$ may also be denoted by $NBS*$. The idea of desired solutions is represented in Figure 5.10.

### Expected Solutions

We now have a pretty good idea of our ideal solution concept and how it can be envisioned within each stage of an agenda, however, we cannot expect agents to come to an agreement at each stage that is the same as the desired solution. This is true for a number of reasons, but the main reason we will focus on is that the utility subspace agents are considering at a particular stage is not the same as the global utility space.

If we look at the utility space defined by stage k in Figure 5.11(a), even if we are considering the NBS for this subspace which is labeled as $NBS'_k$, it is not the same as the desired solution for the stage at $NBS_k$. Not every stage is going to have a utility space that results in a local solution in line with our global solution, even when using the same solution concept for the local space. This becomes even more evident when we have an existing agreement that constrains our local space, as in Figure 5.11(b). In that figure, the current agreement for the issues in stage k is constraining the local space to only the small shaded rectangle, which contains neither the desired solution $NBS_k$, nor the local solution from part (a). In this case the local solution, given just this smaller subspace, might be something like what is labeled $NBS'_k$ in Figure 5.11(b), which is even further away from the desired solution that derives from the globally desired solution.

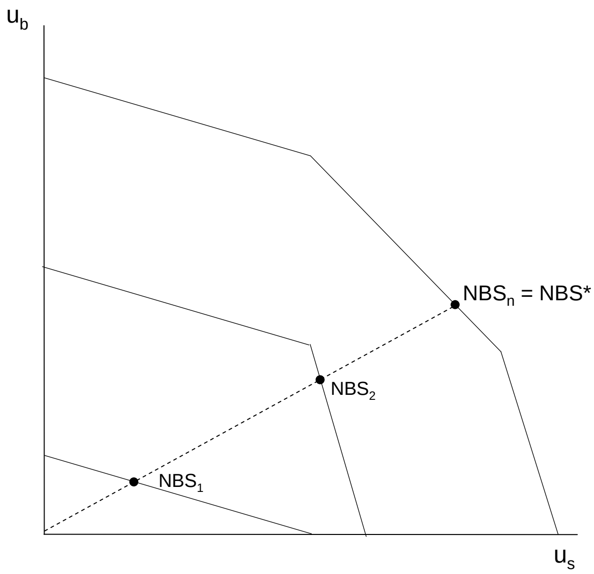


Figure 5.10: The desired solutions for the stages are the intersections of the solution line with the efficient frontier at each stage. The desired solutions are denoted $NBS_i$ for stage $i$.

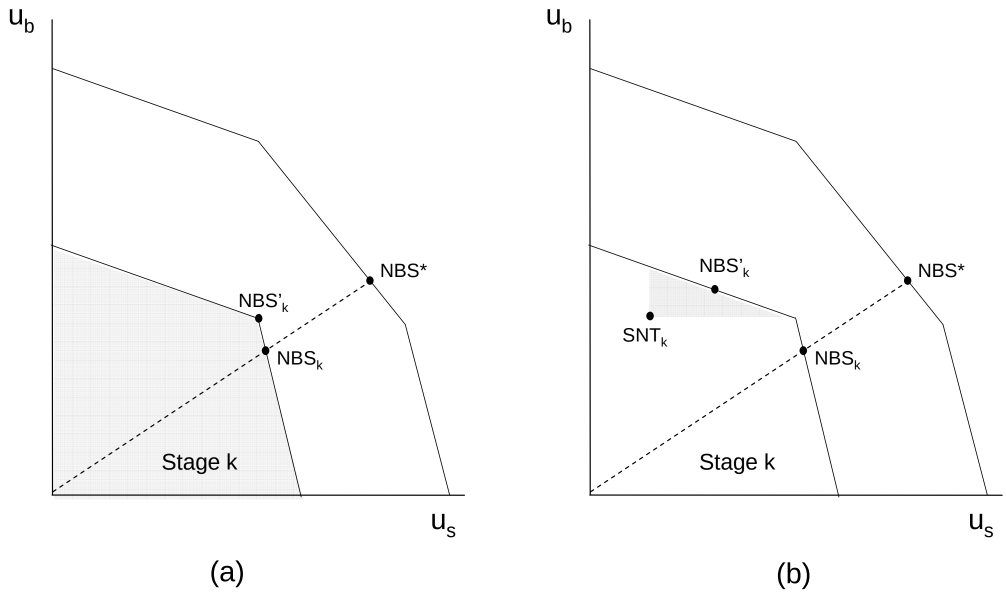


Figure 5.11: (a) In stage k, the utility space is the shaded region. The NBS for this local space for the stage is located at $NBS'_k$ which is not the same as the desired solution $NBS_k$. (b) The local utility space in stage k is constrained to an even smaller subspace of reasonable solutions by the current agreement, $SNT_k$. In this more constrained space, the local solution, $NBS'_k$ is even farther away from the desired solution of $NBS_k$.

The expected solutions that were calculated in the local (possible constrained) subspace for a stage were done using an NBS solution. However, it is possible to use any other solution concept to model the expected solution, it really depends on what is the best solution concept to model the expected behaviour of the agents, and this may vary between different types of agents. For example, if the mediator had some information that one agent was a stronger bargainer than the other, we might bias our expected solutions towards the stronger agent since we would expect them to come out with a better than fair agreement most of the time. Still, there are many other reasons we might use a different solution concept for locally expected solutions.

**Definition 11 (Expected Solution)** *The* expected solution *at each stage is the local agreement that agents are expected to reach for the stage. The calculation of the expected solution for each stage depends on the solution concept chosen to model agents' expected behaviour.*

It is not even necessary to use the same solution concept for the global and expected local solutions. Indeed, it may not even be desirable since the expected behaviour may not be aligned well with the desired global solution. From our previous example where one agent is a stronger bargainer, we may want a global solution that is more fair than the biased expected solution concept and therefore we use a Nash solution globally, but a biased solution locally. Nevertheless, for lack of a more certain way to model general agents, we will be using the NBS for our global solution as well as our expected local solution concept.

When we create our agendas, our goal will be to create an agenda in such a way that the expected solution for each stage is relatively close to the desired solution for the stage. In this way, we will guide the expected behaviour of agents towards the desired global solution.

### 5.4.3 Agenda Creation

Creating an agenda is really just a matter of choosing a schedule for the issues to be negotiated. The choice of an agenda involves answering a few questions:

How many stages will there be?

How many issues will be in each stage?

Which issues will be in which stages?

For simplicity, we use a constant number of issues per stage, typically around two to four issues. Let this constant be $z$, then there will be $f = \lceil \frac{n}{z} \rceil$ stages if we have $n$ issues. The only issue left to resolve is what order we will negotiate the issues, or how we place the issues into the stages. Our basic algorithm for deciding this is to look at every possible agenda, evaluate by some quality measure, and then choose the agenda that has the best quality measure:

```
bestAgendaSoFar = null
for (each agenda, A in all possible agendas) do
    if ( Q(A) > Q(bestAgendaSoFar) ) then
        bestAgendaSoFar = A
    end if
end for
return bestAgendaSoFar
```

This introduces two areas of inquiry, the first is in regard to how many agendas there are and how much work is involved in looking at every agenda, and the second is what is our quality measure. For a discussion about how many agendas we actually have to consider using this method and how much computation needs to be done to choose an agenda, see Appendix A. In this section we will focus on how we measure the quality of a single agenda and what our quality measure means.

The quality measure has to reflect the idea that a higher quality agenda is one that is more likely to lead agents along the solution line, as we established that this will be our method for directing the agents to jointly profitable solutions. We use a simple measure of the distance in the utility space from the expected solutions to the desired solutions throughout the stages of an agenda. The summations of these measurements from each stage form our quality measure, that is, if $d_k$ = the distance from the expected solution to the desired solution in stage k, then the definition for our quality measure is:

**Definition 12 (Agenda Quality Measure)** *The agenda quality measure, Q, for an agenda, A, is*

$$Q(A) = \sum_{k \in Stages_A} d_k$$

*The closer Q is to 0, the higher the quality of the agenda, A.*

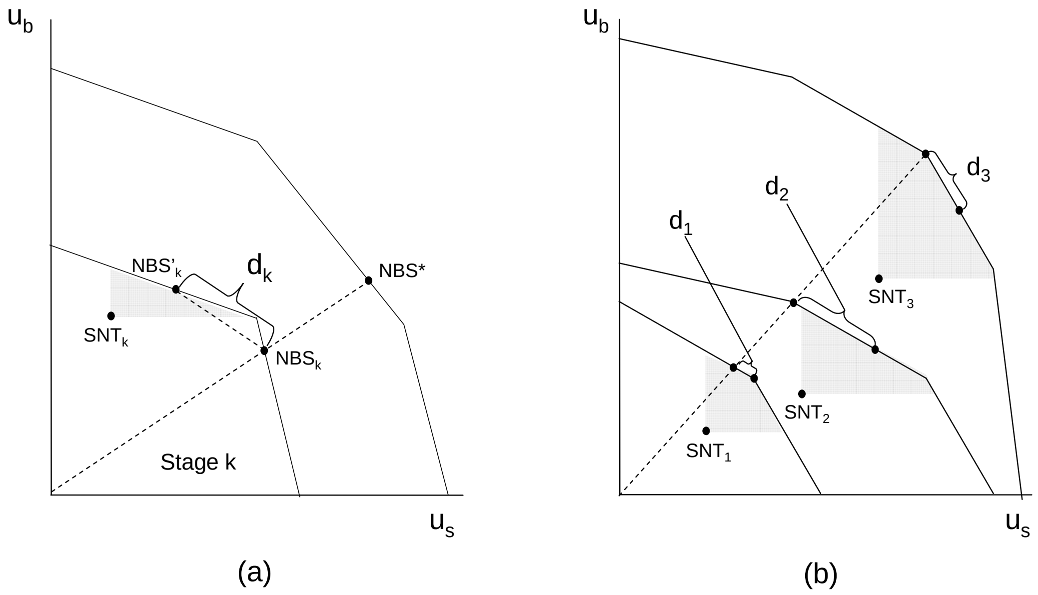


Figure 5.12: (a) Showing the distance measure for stage k from the example shown in Figure 5.11(b). (b) Showing the quality/distance measure for an agenda of 3 stages. The shaded region for each stage represents the space that the current agreement restricts that stage to. The desired and expected solutions are not labeled, but the desired solutions can be distinguished since they are on the solution line.

Our goal, when determining the best agenda, is to find the agenda, $A^*$, such that $A^* = \arg\min_A Q(A)$. Some examples of the distance measure for several stages is shown in Figure 5.12. When we look at Figure 5.12(b) we see that each stage's utility space is being restricted by the current agreement so far, $SNT_k$ (for the issues in that stage and lower). When we create our agendas, we only have the initial agreement to work from, so our $SNT_k$ agreement values will just be the partial agreement formed from the prices in the initial agreement for the issues in stage k and lower.

### 5.4.4 Example of Choosing an Agenda

Figure 5.13 illustrates how we choose between two agendas based on our quality measure (the desired solutions are shown as diamonds, the expected solutions as Xs). In these two agendas, the last two stages are actually composed of the same issues, so their distance measures for these two stages are the same. This means the

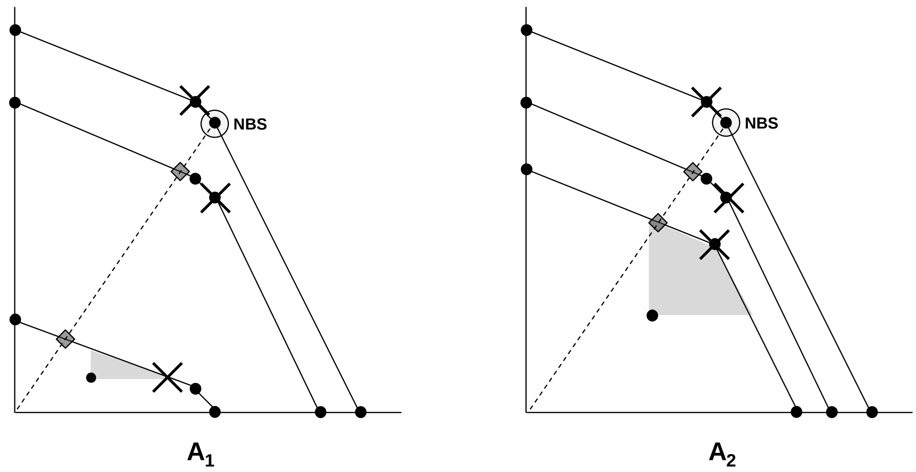


Figure 5.13: An example of choosing between two agendas. The expected solutions are shown with Xs and the desired solutions as diamonds. Each agenda is the same except for the first stage, which is why only the first stage shows the restricted subspace shaded representing where the initial agreement is for that stage. If we are using our measure Q(A) from Definition 12, then we want to minimise Q(A) and we can see that $d_1$ for agenda $A_2$ is smaller than for agenda $A_1$. Therefore, $Q(A_2) < Q(A_1)$, and $A_2$ is the better agenda.

choice of agenda comes down to which agenda has a smaller distance measure for the first stage. It is clear that the expected solution is much closer to the desired solution in the first stage for agenda $A_2$, and therefore if we had to choose between these two agendas, we would choose agenda $A_2$.

### 5.4.5 Alternative Considerations for Agendas

The heart of our agenda creation, or our method of choosing an agenda out of the possible agendas, was the simple quality measure we used. However, there are many alternative ways to measure agendas and different types of agendas altogether that could be used. We outline some extensions to our use, and creation, of agendas in Section 7.3.

## 5.5 Example

In this section, we are going to go through a simple example of using our instantiation of the BAM framework; this should make the workings of the framework clearer.

Our negotiation will consist of four issues that will actually be real-world items. The items and their associated reservation prices and valuations are shown in Figure 5.14. The reservation prices are in dollars and the valuations represent the amount of utility an agent receives per one dollar price shift in its favour. In the case of actual items, the valuations for a single agent might be the same for each item (since a dollar on one issue is the same as a dollar on another), however, for the sake of the example we give the agents different valuations for each item. Note that the valuations for a given agent are drawn from the same valuation range for each issue; this does not have to be the case, as each issue can have its own valuation range.

The first thing we do now is convert these values into our standard normalised representation where the price ranges are $[0, 100]$. This conversion also requires conversion of the valuations since we want the valuations to be per unit of our normalised price, not per unit dollar. Figure 5.15 shows the converted issues with all the old prices shown as well.

To explain how to convert a price, let $D$ be some price in dollars, and let $[L, H]$ be the total price range, in dollars, for the issue we are converting (the range that covers both agents reservation price ranges). The new price is then:

$$price_{new} = \frac{D - L}{H - L} * 100 \tag{5.12}$$

Consider the price \$5, which is the actual reservation price of the seller for issue number two. It becomes the new price of 12.5 by the following conversion:

$$price_{new} = \frac{5 - 4}{12 - 4} * 100 = 12.5$$

The conversion for valuations is slightly different. First we need to determine how much one dollar is worth in the normalised price range (this can be determined by setting $(D - L) = 1$ in Equation 5.12. Dividing the old valuation by this amount gives us the new valuation. So for example, the sellers valuation on issue two of 6 becomes $\frac{6}{12.5} = 0.48$, since one dollar is worth 12.5 units in the new price range.


Figure 5.14: An example of a 4 item negotiation with all the reservation prices and ranges shown, as well as all the valuations. Note that in this example, we are drawing all valuations for an agent from a single range.

Now that all our values have been converted into the new units, we can begin the declaration process. The agents will need to declare reservation prices and the placement of valuation chips on the issues. Let the set of valuation chips provided to the agents be $\{10, 5, 3, 1\}$, and the arbitrary chip shift amount be $H = 10$. The declarations made be the agents are illustrated in Figure 5.16. The seller's reservation price declarations are shown with dotted lines and the buyer's with solid lines.

Once agents have made their declarations, the mediator takes over. As shown in Figure 5.2, first the arbitration module runs, creating an initial agreement, and then the protocol module creates an agenda. Within the arbitration module (see Figure 5.3), first the GPS mechanism creates an initial agreement and then the valuation chips mechanism makes modifications to this agreement.

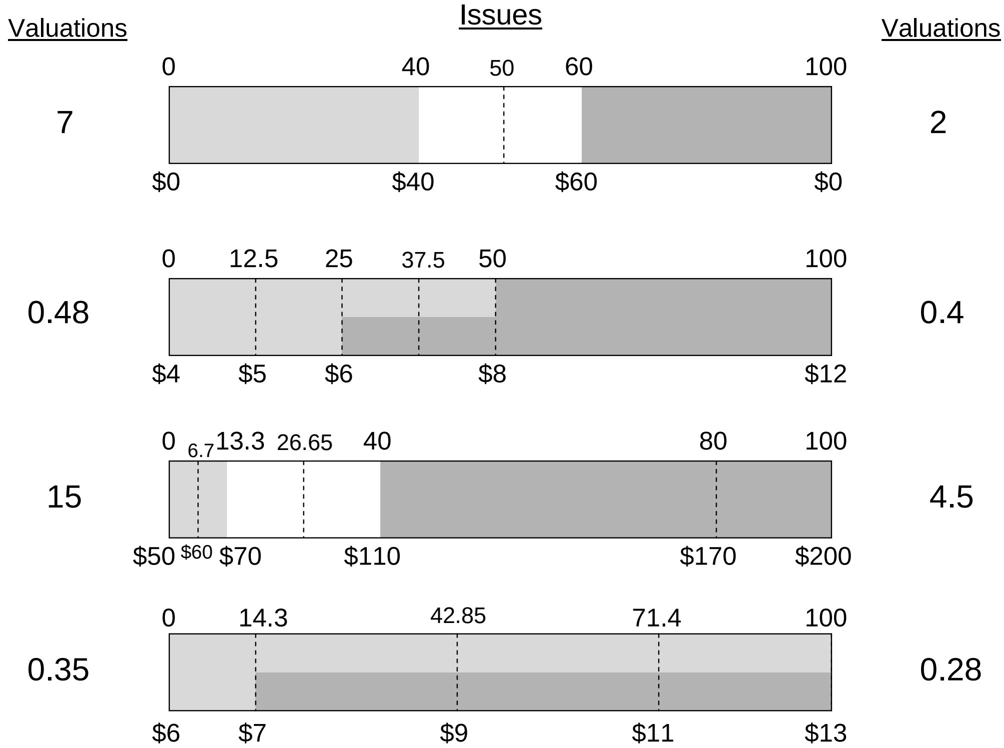


Figure 5.15: The issues converted into our normalised price ranges of $[0, 100]$ and valuations converted to from utilities per dollar to utility per price unit. All the dollar values for the original reservation prices are shown with the normalised values as well.

### 5.5.1 Arbitration Module Example

#### The Greedy Punishment Solution

The GPS creates an initial agreement by applying the GPS price adjustment for every issue:

$$p_j = p - s^s G^s + s^b G^b$$

where the starting default price for each issue is the midpoint of the trade balance. Recall that $s^s = s^b$, where each of these values is equal to half of the trade balance of the issue they refer to. The GPS price adjustments for the issues in our

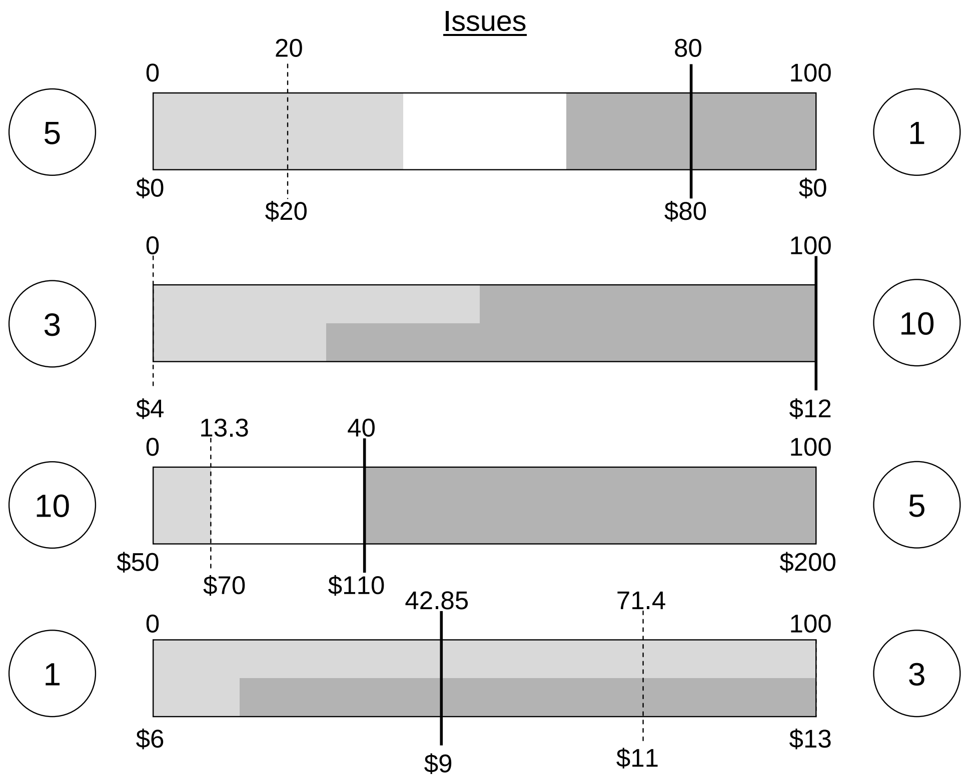


Figure 5.16: The agent declarations for the example. The chip placements are shown at the sides of each issue, and the reservation price declarations are shown within the issues graphical representation. The dotted lines within the issues are the seller's reservation price declarations, and the solid lines are the buyer's declarations.

negotiation are:

$$
\begin{aligned}
p_1 &= 50 - (10 * \frac{20}{40}) + (10 * \frac{20}{40}) \\
&= 50 \\
p_2 &= 37.5 - (12.5 * 0) + (12.5 * 0) \\
&= 37.5 \\
p_3 &= 26.65 - (13.35 * \frac{20}{20}) + (13.35 * \frac{90}{90}) \\
&= 26.65 \\
p_4 &= 42.85 - (28.55 * \frac{5}{7}) + (28.55 * \frac{4}{6}) \\
&= 41.49
\end{aligned}
$$

The initial agreement produced by the GPS is $\{p_1 = 50, p_2 = 37.5, p_3 = 26.65, p_4 = 41.49\}$. This agreement is provided as input to the valuation chips mechanism.

### Valuation Chips Mechanism

The valuation chips mechanism will modify the agreement that the GPS created. The current price on each issue is modified by the valuation chips price adjustment:

$$
\frac{H}{c_n}(c^j - d^j) \tag{5.13}
$$

where H = 10, $c_n = 10$ (since 10 is the largest chip value). $c^j$ and $d^j$ are the value of the chips placed on issue $j$ by the seller and buyer, respectively. According to the declarations in Figure 5.16, we have:

$$
\{c^1 = 5, c^2 = 3, c^3 = 10, c^4 = 1\}
$$
$$
\{d^1 = 1, d^2 = 10, d^3 = 5, d^4 = 3\}
$$

This means the prices in the current initial agreement now become:

$$\begin{aligned}
p_1 &= 50 + \frac{10}{10}(5-1) \\
&= 54 \\
p_2 &= 37.5 + \frac{10}{10}(3-10) \\
&= 30.5 \\
p_3 &= 26.65 + \frac{10}{10}(10-5) \\
&= 31.65 \\
p_4 &= 41.49 + \frac{10}{10}(1-3) \\
&= 39.49
\end{aligned}$$

Thus the finished initial agreement is $\{p_1 = 54, p_2 = 30.5, p_3 = 31.65, p_4 = 39.49\}$. This is the final output of both the valuation chips mechanism and the arbitration module.

### 5.5.2 Protocol Module Example

The protocol module is in charge of creating the agenda for negotiations. In order to evaluate agendas according to Section 5.4, we need to have a utility space representation of the negotiation. To produce this space, we need valuation information, but all we have are issue orderings. The seller's issue ordering is $\{3, 1, 2, 4\}$ and the buyer's is $\{2, 3, 4, 1\}$. We utilise these orderings and the perceived trade balances in order to convert the issues into the maximum amount of utility they could be worth to each agent (which is what we need to create a utility space).

The first thing to do is to produce valuation estimations as described on page 71. In those equations, if we set $u_0 = 100$, we have the following valuation estimations for the seller and buyer (where superscript $s$ or $b$ refers to which agent the valuation is for and the subscript indicates the issue):

$$\begin{aligned}
\lambda_3^{s'} &= \frac{100}{26.7} \\
\lambda_1^{s'} &= \lambda_3^{s'} - (\frac{\lambda_3^{s'}}{4})1 \\
\lambda_2^{s'} &= \lambda_3^{s'} - (\frac{\lambda_3^{s'}}{4})2 \\
\lambda_4^{s'} &= \lambda_3^{s'} - (\frac{\lambda_3^{s'}}{4})3
\end{aligned}$$

$$\begin{aligned}
\lambda_2^{b'} &= \frac{100}{25} = 4 \\
\lambda_3^{b'} &= \lambda_2^{b'} - (\frac{\lambda_2^{b'}}{4})1 = 3 \\
\lambda_4^{b'} &= \lambda_2^{b'} - (\frac{\lambda_2^{b'}}{4})2 = 2 \\
\lambda_1^{b'} &= \lambda_2^{b'} - (\frac{\lambda_2^{b'}}{4})3 = 1
\end{aligned}$$

Using our estimated valuations, we can now get the estimated utility that each issue can be worth for each agent. To do this, we multiply the perceived trade balance by our estimated valuations for each agent. This makes the estimated utilities for the seller:

$$\begin{aligned}
u_1^{s'} &= \lambda_1^{s'} * 20 = 56.2 \\
u_2^{s'} &= \lambda_2^{s'} * 25 = 46.8 \\
u_3^{s'} &= \lambda_3^{s'} * 26.7 = 100 \\
u_4^{s'} &= \lambda_4^{s'} * 85.7 = 80.2
\end{aligned}$$

And the estimated utilities for the buyer:

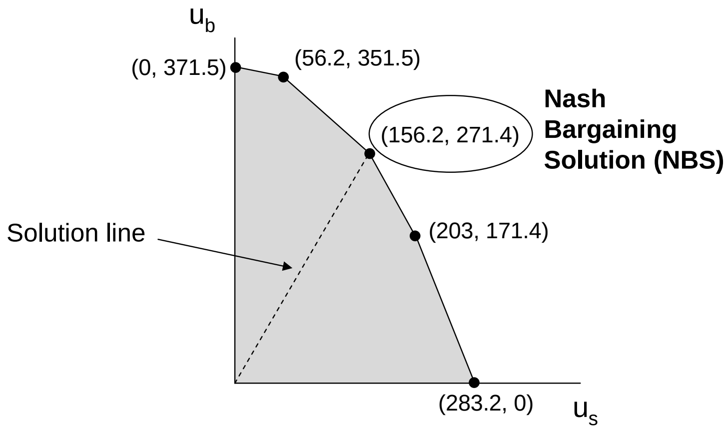


Figure 5.17: The utility space of our 4 item negotiation based on the estimations of the valuations and the size of the actual trade balance. The Nash Bargaining solution is shown on the efficient frontier as well as the solution line. The frontier is constructed from left to right by the following ordering of issues $\{I_1, I_3, I_2, I_4\}$.

$$
\begin{aligned}
u_1^{b'} &= \lambda_1^{b'} * 20 = 20 \\
u_2^{b'} &= \lambda_2^{b'} * 25 = 100 \\
u_3^{b'} &= \lambda_3^{b'} * 26.7 = 80.1 \\
u_4^{b'} &= \lambda_4^{b'} * 85.7 = 171.4
\end{aligned}
$$

Summing these values, this means the total (estimated) possible utility is 283.2 and 371.5 for the seller and buyer, respectively. These estimated utility values produce the utility space shown in Figure 5.17. The figure also identifies the Nash Bargaining solution (NBS) and the solution line.[8]

The last thing left for the protocol module to do in order to choose an agenda for the negotiations is to compare all the possible agendas and see which one has the best quality measurement. To make things simpler, we are going to just choose two agendas and compare them in the same manner the protocol module would.

[8]The Nash Bargaining solution is determined by finding the point on the efficient frontier that maximises the product of both agents' utilities.

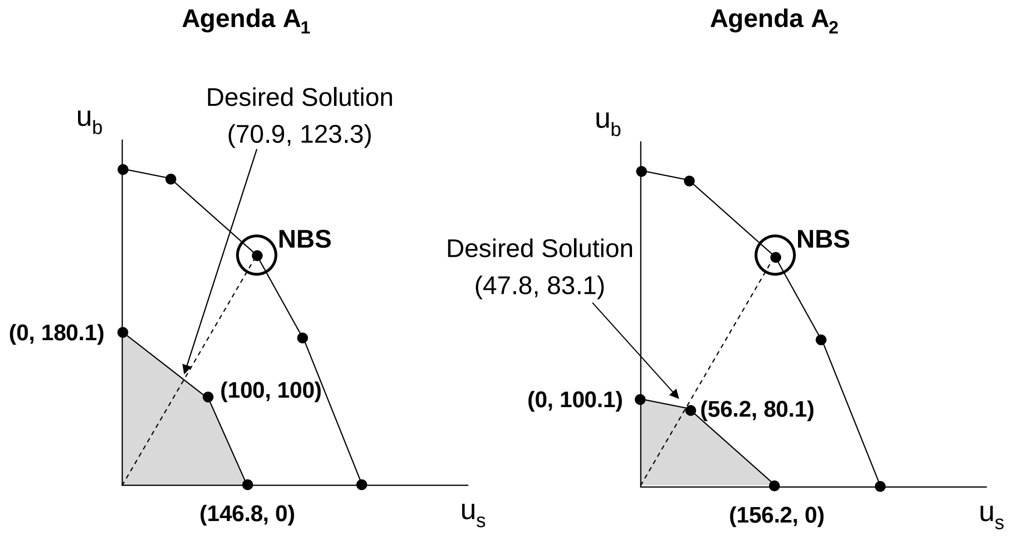


Figure 5.18: The utility spaces for the two agendas $A_1$ and $A_2$, with the stage utility frontiers shown. The solution line is also plotted along with the desired solution in the first stage, which is the intersection of the solution line with the first stage utility frontier.

The two agendas we are going to look at are: $A_1 = (\{I_2, I_3\}, \{I_1, I_4\})$ and $A_2 = (\{I_1, I_3\}, \{I_2, I_4\})$. Figure 5.18 shows these two agendas represented in utility space along with the solution line in each of them. The intersection of the solution line with the stage frontiers are marked showing where the desired solutions exist.

Before we compare the agendas, we should notice that the second stage is the same for both agendas in terms of the utility space (because the second stage comprises the entire utility space). This means all the calculations for that stage between the two agendas is the same, so we need only worry about the calculations in the first stage. This is similar to the agenda example in Section 5.4.3.

In order to compare the two agendas to see which is superior (according to our measures), we need to determine the expected solutions in the first stage of the agendas. We will need to plot where the initial agreement places us within the first stage's utility space. This is accomplished by calculating how much utility each agent receives in the first stage from the initial agreement; essentially we calculate the utility for each agent according to Definition 7 (on page 48), but restrict the calculation to the issues in the first stage. In place of the actual valuation and reservation price, we use our estimated valuation and the boundaries of the trade balance as substitutes, respectively. Letting the perceived trade balance on issue $I_j$ be $[L_j, H_j]$, the utility we will calculate for each agent, in the first stage of an agenda, is:

$$\begin{array}{ll}(seller) & \sum_{I_J \in I_{stage_1}} \lambda_j^{s'}(p_j - L_j) \\ (buyer) & \sum_{I_J \in I_{stage_1}} \lambda_{I_j}^{b'}(H_j - p_j)\end{array}$$

Recall that the initial agreement is $\{p_1 = 54, p_2 = 30.5, p_3 = 31.65, p_4 = 39.49\}$. Then the utility for each agent for agenda $A_1$ is:

$$\begin{aligned}(seller)u^s_{A_1} &= \lambda_2^{s'}(30.5 - 25) + \lambda_3^{s'}(31.65 - 13.3) \\ &= 79.03 \\ (buyer)u^b_{A_1} &= \lambda_2^{b'}(50 - 30.5) + \lambda_3^{b'}(40 - 31.65) \\ &= 103.05\end{aligned}$$

And the utility for each agent for agenda $A_2$ is:

$$\begin{aligned}(seller)u^s_{A_2} &= \lambda_1^{s'}(54 - 40) + \lambda_3^{s'}(31.65 - 13.3) \\ &= 108.05 \\ (buyer)u^b_{A_2} &= \lambda_1^{b'}(60 - 54) + \lambda_3^{b'}(40 - 31.65) \\ &= 31.05\end{aligned}$$

If we plot the $stage_1$ initial agreement utility points within our utility spaces (Figure 5.19), we see how the initial agreement restricts the agents to a subspace of the $stage_1$ utility space. Figure 5.19 also shows where the expected solutions exist for these subspaces, which is just the NBS calculated for the restricted subspaces. Now that we have the expected and desired solutions, we can calculate the agenda quality measure (Definition 12 on page 77) that depends on the distance between these two types of solutions.

For the first stage of agenda $A_1$, the expected solution is $(96.19, 103.05)$ and the desired solution is $(70.95, 123.27)$. This means the agenda quality measure for the first stage of agenda $A_1$ (the distance between the expected and desired solution) is 28.75.

For the first stage of agenda $A_2$, the expected solution is $(108.05, 38.57)$ and the desired solution is $(47.82, 83.08)$. This means the agenda quality measure for the first stage of agenda $A_2$ is 74.90.

Remember that we only need to compare the agenda quality measures of the two agendas for the first stage because all stages after the first are the same in the agenda. We can see that $Q(A_1) = 28.75 < 74.90 = Q(A_2)$. Since we want to minimise $Q$, this tells us that agenda $A_1$ is better than agenda $A_2$. Therefore, our mediator will choose agenda $A_1$ in the protocol module.

Our mediator process is now finished with its part in the negotiations, having chosen an initial agreement ($\{p_1 = 54, p_2 = 30.5, p_3 = 31.65, p_4 = 39.49\}$) and an agenda for the agents to use ($A_1 = (\{I_2, I_3\}, \{I_1, I_4\})$). The agents will now negotiate starting with our initial agreement and following the agenda we provided them, ratifying the final agreement themselves without further intervention from our mediator.

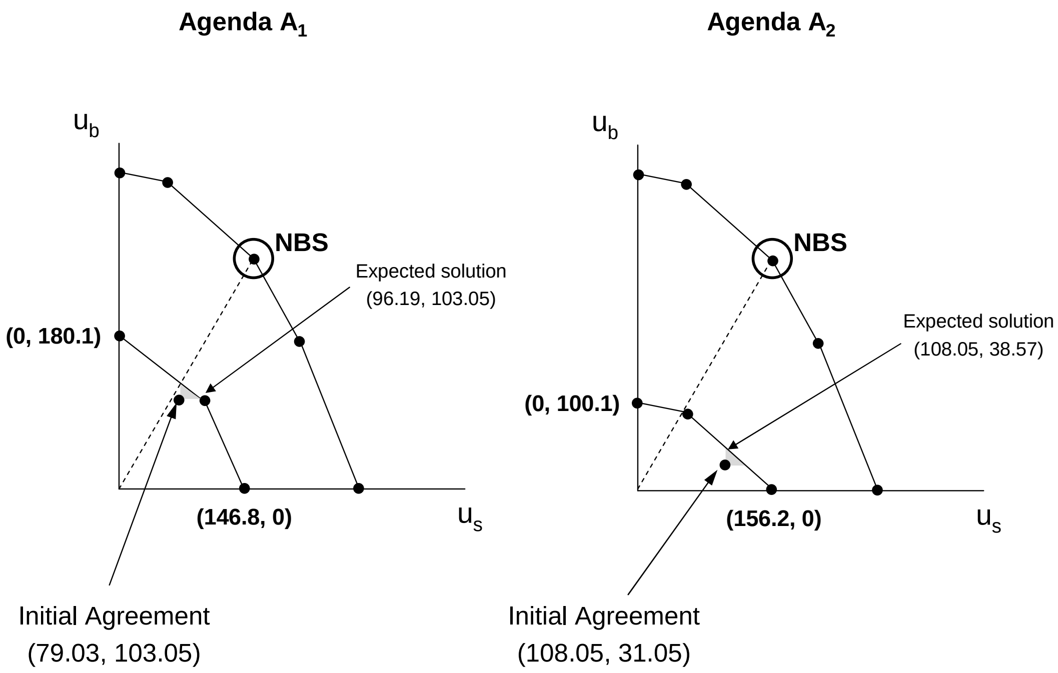


Figure 5.19: Showing the subspaces in the first stage of each of the two agendas, $A_1$ and $A_2$, that agents are constrained to due to the initial agreement. We also see where the expected agreements for each agenda are, being the Nash Bargaining solution of these restricted subspaces. Intuitively, we can see that the expected and desired solutions for the first stage are much closer together in agenda $A_1$ than they are in agenda $A_2$. This indicates to us that agenda $A_1$ is probably going to be the more desirable agenda to use.

## 5.6 Conclusion

The BAM framework was built with the intention of assisting agents to achieve efficient and fair agreements in incomplete information scenarios. In particular, our mediator's purpose was to steer agents away from adversarial and selfish thinking that may let agents lose sight of joint gains. However, the BAM framework cannot do anything that the agents could not do for themselves, but as a third-party it changes the negotiation environment in ways the agents are unlikely to do themselves. The framework changes what agents pay attention to and focuses them on joint efforts. By using an incremental agreement process, the framework shifts attention from developing individually created partisan offers to a common agreement developed mutually. Instead of agents having to bring their partisan offers down to a level acceptable to the other agent, now they are focused on improving a single agreement for both agents; after all, both agents need to agree to any changes to the common agreement so there must always been benefits for both agents in any update to the common agreement. In the same spirit, our mediator provided an agenda that keeps agents focused on solution subspaces that work well for both agents and represent efficient areas of the utility space.

In order to help the agents, our mediator not only had to understand negotiations, but additionally had to get the agents to trust it. We saw this in the arbitration module were the mechanisms strived for incentive compatibility and individual rationality. The mediator we created had the arbitration module focus on guaranteeing properties to make the agents want to participate, without worrying too much about overall efficiency. Then the protocol module took over and led agents through the solution space in a manner that was centred on wringing out as much efficiency as possible, while paying attention to fairness too. This was how our particular mediator delegated its overall goals into tasks within the modules, but the framework can support mediators with different methodologies structured within the arbitration and protocol modules.

The nice thing about having these separate modules is that they allow a "plug-and-play" incremental development. For example, if we came up with a way to create agendas that could guarantee fair pareto optimal outcomes, we good just plug this new algorithm into the protocol module without having to change the arbitration module. The inspiration for this separation of the mediator into modules is due to the fact that incomplete information negotiations can be difficult and may benefit from having multiple approaches to progressively tackle the problem; multiple modules encourages this hybrid approach to problem solving.

# Chapter 6

# Experimental Setup and Results

In order to examine the BAM framework more thoroughly, we have implemented a simulation environment from the framework. The simulation framework implements the version of the BAM framework as described in the instantiation detailed in Sections 5.3 and 5.4.

In Section 6.1, we discuss how data was setup for the experiment and the parameters of the negotiation procedure used. Section 6.2 describes the agents that participated in testing the framework and the behaviour they used during negotiation. Section 6.3 discusses the metrics we recorded during simulation and the specific results of our simulation trials. Lastly, we discuss the results of our experiment in Section 6.4.

## 6.1 Negotiation Details

The following sections describe the parameters and process of a single negotiation between two agents, which constitutes a single trial in our experiment.

### 6.1.1 Issue Setup

For each individual negotiation, ten issues were randomly generated based on a set of parameters that described each aspect of the issues. All the issues were generated using the same parameters and ranges. The generation properties were:

- The price range for each issue was from $[0, 100]$.

- The high end of the seller's reservation price range was randomly chosen within the range $[10, 55]$. So the seller's reservation price range would be $[0, x]$, where $10 \leq x \leq 55$.
- The low end of the buyer's reservation price range was randomly chosen within the range $[45, 90]$. So the buyer's reservation price range would be $[y, 100]$, where $45 \leq y \leq 100$.
- The agents' actual reservation values were randomly chosen within the reservation price ranges generated.
- Each agent also had a random valuation range generated for them which was a subset of the range $[1, 20]$.
- The agents' actual valuations for the issues were randomly drawn from their valuation ranges.

As can be seen, there was no particular bias towards either agent; the ranges chosen are mirrored identically. In future work, it may be useful to run experiments in more asymmetric settings to compare with our experiments here.

### 6.1.2 Agenda Setup

The agendas generated for our negotiations used a fixed stage size of two issues per stage. With ten issues total, this means all our agendas are going to be five stage agendas. This can be somewhat limiting, as variable stage sizes would allow for the best tradeoffs amongst issues and probably the highest ability to guide agents more precisely through the solution space. Again, these are options to consider in future explorations of the framework.

### 6.1.3 Offer Procedure

The offer procedure defines the heart of a negotiation procedure, since a negotiation is just a compilation of offers and their resolutions. Our particular offer procedure is similar to the Rubinstein / Ståhl alternating offers model (Section 2.4.4), but differs from it in that we choose an agent to make an offer randomly and instead of a discount factor we use an exogenous probability of breakdown. If the negotiation breaks down, then agents have to accept the current agreement, which is the initial agreement plus any modifications to it that have been already been agreed upon. This negotiation procedure is the one used by John and Raith [11, 12].

The probability of choosing the seller as the offerer, $P$, was equal to 0.5, so each agent was equally likely to be the offerer at each time step. The exogenous probability of breakdown, $1-\delta$, is implemented as a chance of breakdown that occurs before every offer is made. In our particular experiment, $1-\delta = 0.04$, or 4%. The result of using this offer procedure is that every single offer in a given stage has the exact same analysis illustrated by Figure 6.1. This means we do not have to resort to the backwards induction used in the alternating offers model (Section 2.4.4) when discussing rationality, but instead just consider the amount of utility each agent expects to get based on the probabilities of being chosen as the offerer and of exogenous breakdown.

The expectations of each agent are:

$$u_b^s = \delta[Pu_s^s + (1-P)u_b^s] + (1-\delta)u_0^s \tag{6.1}$$

$$u_s^b = \delta[Pu_s^b + (1-P)u_b^b] + (1-\delta)u_0^b \tag{6.2}$$

Subscripts in the above equations indicate the identity of the offerer, superscripts indicate the recipient of the utility. The interpretation of something like $u_b^s$ then means, the amount of utility the seller would receive from the buyer's offer. The subscript 0 indicates the "offerer" is the initial agreement. To fully understand the equations, let us dissect Equation 6.1.

$Pu_s^s$ is the probability of the seller being chosen multiplied by the utility the seller would receive if it were to make an offer

$(1-P)u_b^s$ is the probability of the buyer being chosen multiplied by the utility the seller would receive if the buyer were to make an offer

$Pu_s^s + (1-P)u_b^s$ is the expected utility the seller would receive if exogenous breakdown does not occur

$u_0^s$ is the utility the seller receives from the initial agreement for the stage

Therefore Equation 6.1 becomes the utility the seller expects overall based on whether or not exogenous breakdown occurs. Equation 6.2 is similar but for the buyer.

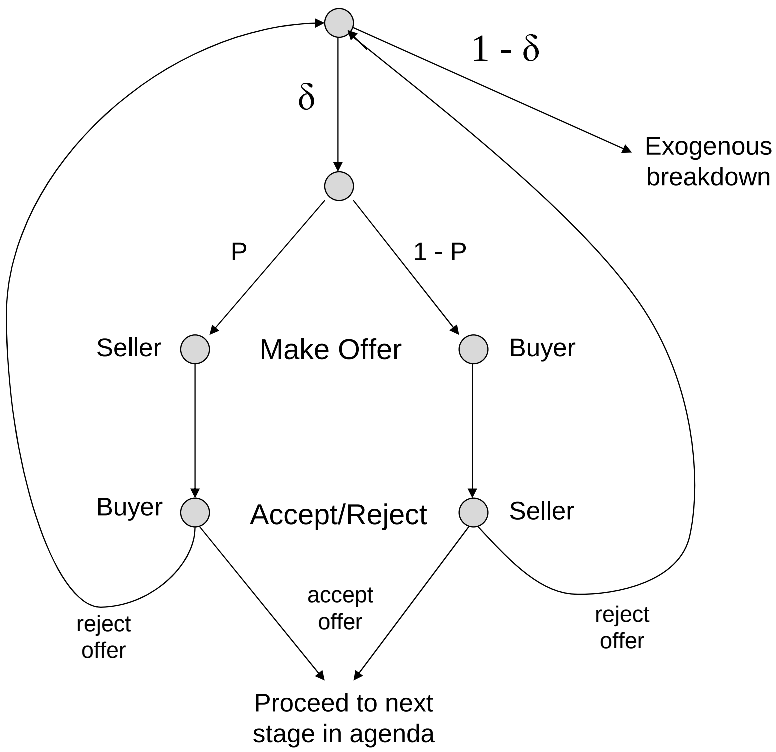


Figure 6.1: The diagram illustrates how a single offer in the negotiation procedure works. The top two nodes in the diagram start the offer and represent random elements in the offer, so the actor for these nodes can be considered to be Nature. First, before anything else, there is a $1 - \delta$ probability of exogenous breakdown. Then an agent is randomly chosen to make an offer to the other agent. If the offer is accepted, then we move on to the next stage were the same procedure is followed. If the offer is rejected we run through this exact same procedure within the same stage starting with the chance of exogenous breakdown.

## 6.2 Agents

The agents used in our simulation experiments and their behaviour are described below:

**Default Agent** This agent accepts any agreement, no matter how bad it is, and offers the stage's initial agreement as its offer.

**Smart Default Agent** This agent accepts any agreement that gives it at least as much utility as the stage's initial agreement and makes offers that give it very nominal utility increases (i.e. a 5% increase in utility over the initial agreement).

**Typical Agent** This agent starts out with a high desired utility which guides its offers and acceptance behaviour, but as the stage continues without successful offers, this agent will lower its desired utility all the way down to accepting the initial agreement for the stage. Additionally, this agent maintains an automatic acceptance utility level that is lower than its desired utility level, but it will accept any agreement offered that gives it that much utility. The idea is that the agent believes that this amount is better than it can honestly expect, even if it is not as much as it desires at the moment.

For example, at the start of a stage, a Typical Agent might initially want 80% of its maximum possible increase in utility, with an automatic acceptance utility of 50% of its maximum possible increase in utility. With each offer that goes unaccepted (by either agent), the agent lowers its desired amount by between 10%-20%. It will keep doing this until an agreement is accepted or it has bottomed out and is willing to accept the same utility as offered by the initial agreement (to try and prevent exogenous breakdown occurring).

## 6.3 Results

We ran our negotiation experiment with each possible pairing of agent types, first running 1000 trials, then switching the agent's identities (as the buyer and seller) and then running another 1000 trials. The results that appear below are the averages over these 2000 trials per agent pairing.

### 6.3.1 Metrics

We consider the following metrics in evaluating the trials in our experiment:

M1. The utility received by each agent as a percentage of its maximum possible utility. The maximum possible utility is not realistically possible in a negotiation, however, because it would require the other agent to give in completely on every issue. This is just used as one measure that can compare the differences between different agent type pairings.

M2. The utility received by each agent as a percentage of its utility if the Nash Bargaining Solution had been achieved. The NBS in this situation represents an ideal solution, although it is possible for an agent to receive more than it would in the NBS at the expense of the other agent. This measure can be used more successfully as an efficiency measure, as well as a measure of how well our system estimated the NBS with incomplete information to work from.

M3. The utility received by each agent as a percentage increase over the utility of the initial agreement proposed by the mediator. This allows us to see if agents actually negotiated a better deal than what they started with.

M4. The exogenous failure rate between different pairs of agents. This measure allows us to see if agents had difficulty reaching an agreement since a higher failure rate will generally indicate longer negotiating times.

M5. The average stage reached in negotiations between a pair of agents. This is a measure of how long negotiations were and how close agents came to reaching a successful agreement in the case of exogenous failure.

M6. The average utility received for negotiations, separated into categories according to the stage reached in the negotiations. This measure will allow us to see how negotiations improve as agents move further through the agenda.

### 6.3.2 Data

The first metric we consider are the failure rates for different agent pairings. The failure rate is just the percentage of the negotiation trials that resulted in exogenous breakdown (any trial that was successful meant that the agents managed to negotiate through the entire agenda). Table 6.1 shows the failure rates amongst different agent pairings. Note that the table should be symmetric, which is why some entries are not listed.

| | default | smart default | typical |
|---|---|---|---|
| default | 5.1 | 5.3 | 8.5 |
| smart default | - | 98.5 | 35.5 |
| typical | - | - | 37.1 |

Table 6.1: Showing the various rates of failure for different agent pairings.

The low failure rates with the default agent are to be expected since it accepts any offer proposed to it, and these failure rates form a baseline for the amount of minimum failure that occurs due to exogenous causes. The typical agent has a higher rate of failure, around the mid thirties, and this is due to the fact that these agents take longer to negotiate stages and thus opens up the chance for exogenous failure. The smart default agent negotiating with another smart default agent has a startling high failure rate, almost failing at every opportunity! This is due to the fact that these agents do not modify their offers as time goes by, instead stubbornly sticking to them.

Related to the failure rates is the average stage reached for various pairings. The average stage reached gives us an idea of how far through the agenda different agent pairings got to. Table 6.2 displays this metric.

| | default | smart default | typical |
|---|---|---|---|
| default | 4.849 | 4.843 | 4.7805 |
| smart default | - | 1.058 | 3.917 |
| typical | - | - | 3.81 |

Table 6.2: Showing how far each agent pairing got in the agenda, on average. Note that the agendas only had 5 stages, so this was the maximum possible value.

These values are generally what we would expect, especially considering the failure rates above. The default agents almost always reach later stages in the agenda since they are willing to accept anything and advance through the stages quickly. Again, we see, as we start having more complicated agents playing against each other, that they tend to take longer in their negotiations, and thus exogenous breakdown becomes more likely. Again we see the difficulties the smart default agent has negotiating against a similar agent.

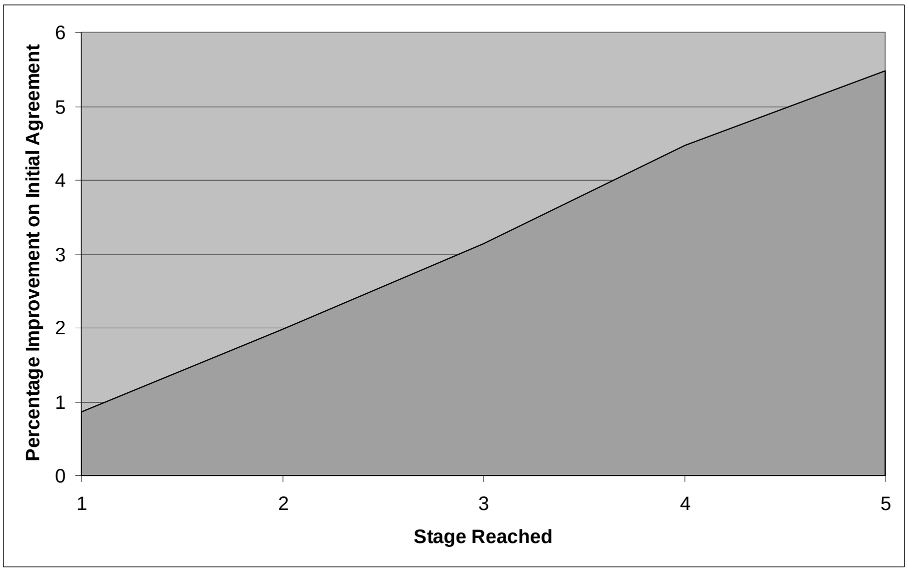


Figure 6.2: The graph shows how well agents do compared to how far they get in the agenda. The measurements represent the percentage increase in utility for the agents, above that of what was provided by the initial agreement decided by the mediator. It is clear the general trend is that agents get larger utility the further they go through the agenda. This is natural since it means more issues have been negotiated and more utility gains are possible to be found. The results here are averaged over all the experimental trials we performed and just shows the average utility an agent can expect if its negotiation only gets to a certain stage.

Next we show how well agents do the further they get in the agenda. In Figure 6.2, we find a graph illustrating how much agents bettered the initial agreement based on how far they got through the agenda. This graph averages over all agent pairings possible, and is used to show the general trend that as agents get further through the agenda, they receive more utility.

We now focus a little more on the specific utility values agents achieved during negotiations. Table 6.3 displays what percentage of the maximum utility each agent received. The entries listed show what percentage of the maximum utility the row agent received on average. Note that these values appear very low since the maximum utility achievable is not realistically achievable at all, since it would require one agent to basically give in completely on every possible issue.

We see that except for when playing against the default agent, there is very

| | default | smart default | typical |
|---|---|---|---|
| default | 40.2 | 38.4 | 28.0 |
| smart default | 43.0 | 41.0 | 41.2 |
| typical | 54.9 | 41.7 | 41.2 |

Table 6.3: This table lists the percentage of the maximum possible utility each agent received as an average over all the trials. The entries list these values for the row agent.

little movement towards the Nash bargaining solution in any of the agent pairing when compared to the baseline value of the top left entry.

The next measurement (Table 6.4) compares the agents' utilities received to the Nash Bargaining Solution, which is our desired solution. The values shown represent the percentage of the NBS utility that the row agent received, meaning if an entry says 78.5, then the row agent for that entry received 78.5% of its NBS utility. The entry for the default agent paired with the default agent acts as the baseline for the other values, since that entry represents where the initial agreement is with respect to the NBS.

| | default | smart default | typical |
|---|---|---|---|
| default | 77.0 | 73.8 | 53.5 |
| smart default | 82.8 | 78.5 | 79.7 |
| typical | 105.2 | 79.9 | 79.1 |

Table 6.4: This table lists the percentage of the Nash Bargaining Solution utility each agent received as an average over all the trials. The entries list these values for the row agent.

Again, we do not see a lot of improvement on the baseline entry for the default agent paired with another default agent, except in the case of the typical agent taking advantage of the default agent.

The last measurement shows how much improvement agents made over the initial agreement in terms of utility; it is expressed as the percentage of improvement over the initial agreement. This is shown in Table 6.5, and again the entries are listed for the row agent only.

Again, we see a similar trend of exploitation of the default agent by the typical agent, but when paired against the smart default agent or another typical agent, the improvements over the initial agreement are much more modest.

| | default | smart default | typical |
|---|---|---|---|
| default | 0 | -4.55 | -30.12 |
| smart default | 7.26 | 2.02 | 2.04 |
| typical | 39.2173 | 5.48 | 2.74 |

Table 6.5: This table lists the percentage improvement agents made over the utility they received in the initial agreement laid out by the mediator. The entries list these values for the row agent.

## 6.4 Discussion

As we saw in the last section, the experiments showed that little improvement is being made beyond the initial agreement that was provided by the mediator. One positive aspect to this is that the initial agreement brings agents up to a level of utility that provides them with a solid basis agreement. However, agents were having difficulty obtaining utility beyond this point except in situations of exploitation resulting from very unevenly matched agents.

The philosophy of our mediation framework is to encourage agents to explore areas of joint gain based on where they were being directed by the initial agreement and agenda. The failure of this to happen is believed to be due to two things: we have not yet examined enough agents with different search behaviour, and the implementation details of how offers are created might not be explorative enough of the solution space.

In order to simplify the simulation framework, the way offers were generated may have not focused enough on both agents, but only on the agent making the offer. In this way, desired utility increases in offers may not have been well-balanced enough and the concessions made to entice the other agent may not have been strong enough to allow agreements to be made early on; this would have resulted in offers having to be pulled back significantly to achieve agreement and opened up the possibility of exogenous breakdowns. In particular, the typical agent initially made offers close to the efficient frontier, but then scaled back away from the frontier almost to the point of the initial agreement in order to reach an agreement. This is why the typical agent's gains, while they are positive, are fairly small.

In future work we will have to explore offer creation and agreement behaviour that focuses more on starting from the initial agreement and moving towards the frontier, especially considering this is what our framework would seem to suggest is the best behaviour to follow. Unfortunately, advancements in this direction are still swimming against the tide of incomplete information, as agents still have little idea of the preferences of the other agent. It may be useful to explore using different types of learning agents within our system to see if this helps ameliorate this problem.

# Chapter 7

# Extensions of the System

This chapter is our future work section where we explore various extensions to the BAM framework including updates to the Greedy Punishment Solution, the valuation chips mechanism, and agendas.

## 7.1 Randomised Revelation Process

The Greedy Punishment Solution (GPS) of Section 5.3.1 guaranteed one-sided IC, but gave us little information about the actual reservation prices (RP) of the agents. It was essentially a declaration, under duress, from the agents that they will allow the mediator maximum flexibility on issue pricing by declaring the weakest RP value. Figure 7.1 reminds us of the definitions of weakest and strongest RP values.

It would be desirable to have a mechanism that obtains more information than the GPS, even though Myerson and Satterthwaite (Theorem 1) suggest we will not be able to obtain everything in an incomplete information scenario. But perhaps it is possible to obtain more than what we have now, but short of everything we desire. The method proposed here utilises randomisation in the hopes of achieving desirable goals in expectation: we would like agents to be truthful in general.

### 7.1.1 The Method

Imagine each declaration an agent makes for each issue is considered an object: for every issue we have two objects, one declaration for the seller and one for the buyer (Figure 7.2). Then for $n$ issues, we have $2n$ declaration objects (Figure 7.3).

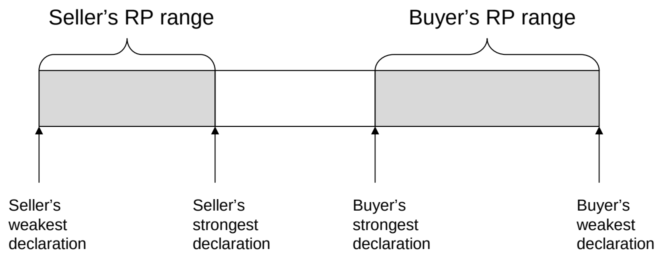


Figure 7.1: An issue showing the RP ranges of the agents and the weakest and strongest declarations of the agents.

Each declaration object represents the RP range of the agent it belongs to and the line within the rectangle represents where the agent's declaration occurs within this range.

The Randomised Revelation method takes the declaration objects and randomly pairs them together under two types of pairing relationships: intra-agent and inter-agent. Intra-agent pairings pair declarations from the same agent together, whereas inter-agent pairings form pairs between declarations from the two agents. Every declaration object will participate in one inter-agent pairing and one intra-agent pairing (If there are an odd number of issues, then one declaration from each agent will not be part of an intra-agent pairing). A full set of both pairings are illustrated in Figure 7.4.

### The Meaning of Intra-agent Pairings

When two declarations from the same agent are paired, we randomly choose one of these declarations that we will "believe" and the other that we will not believe.[1] The meaning of belief is that we will generally accept believed declarations as truthful. This will be explained in more detail when we discuss inter-agent pairings.

The reason for random intra-agent pairs is to make the agents uncertain which declarations will be believed by the mediator and which will not be believed. The agents will know that approximately half of their declarations will be believed and half will not, but the placement of declarations into these two groups will be unknown to the agents when the declarations are made. This uncertainty about

[1] If there is an unpaired declaration for each agent, we randomly choose to believe or disbelieve both of them.

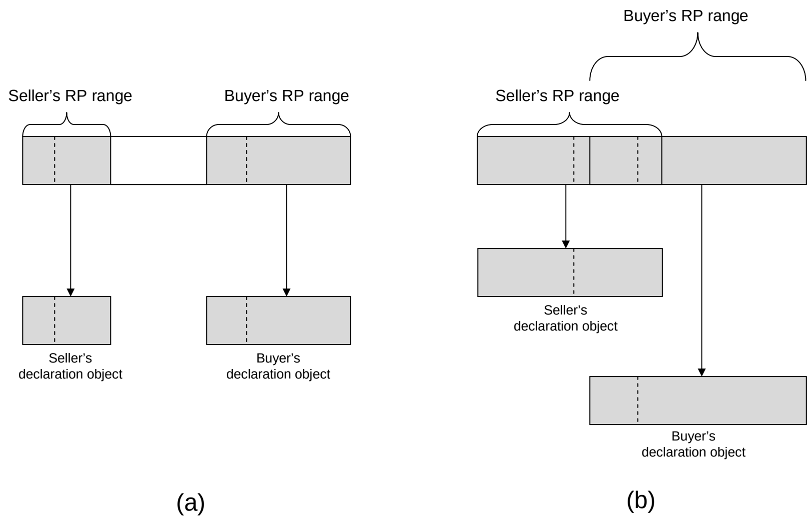


Figure 7.2: The derivation of declaration objects from an issue. Each declaration represents the agents' RP ranges and the lines within this range are the declarations they made to the mediator. (a) An example of deriving the declaration objects from an issue with a surplus. (b) An example of deriving the declaration objects from an issue with a deficit.

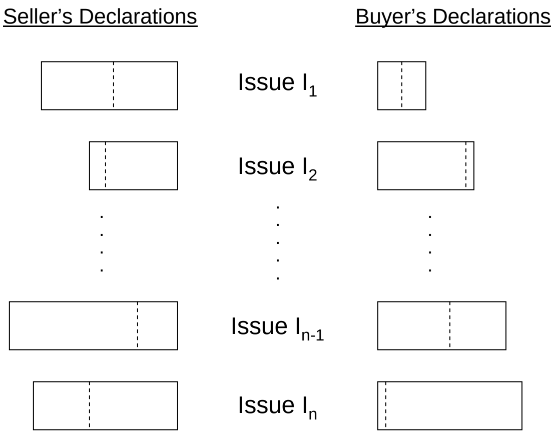


Figure 7.3: An example of the declaration objects for all n issues. Remember each object represents the RP range of the agent it belongs to, and the line within is where the agents declaration within this range occurred.

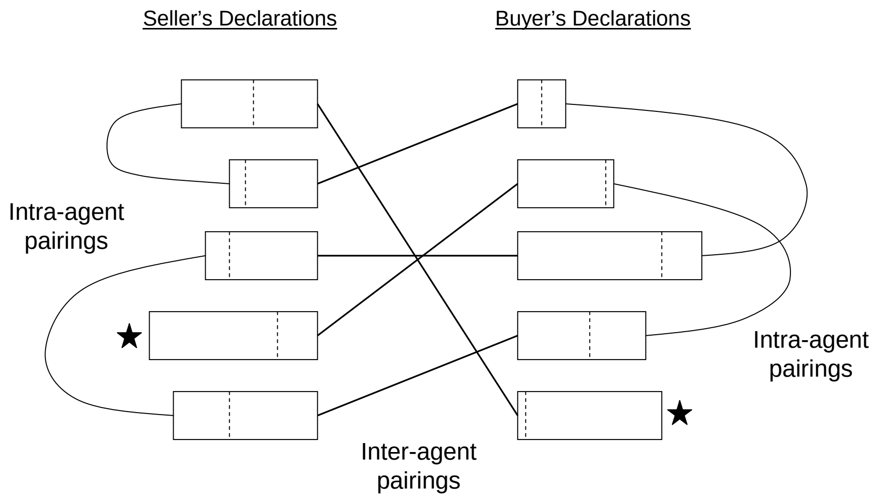


Figure 7.4: An example of pairings between declaration objects. Notice the declaration objects that are starred. Since the number of issues is odd, these are the randomly chosen issues that do not have intra-agent pairings.

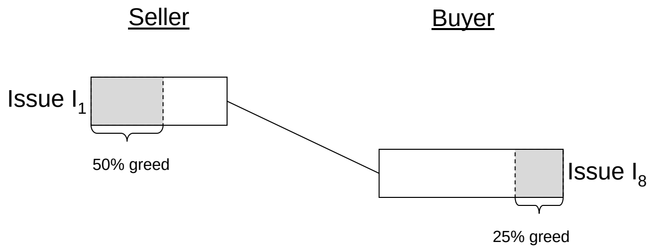


Figure 7.5: An example illustrating basic inter-agent pairing. In this case the seller is going to be punished on issue $I_1$ by the difference of the two agents' greed, $25\%$. The punishment is translated into lowering the default price on issue $I_1$ in the initial agreement.

which declarations will be believed makes it more difficult for agents to exploit a mechanism that uses these types of pairings.

### The Meaning of Inter-agent Pairings

When two declarations from different agents are paired, this puts the two declarations into a game similar to that of the GPS mechanism. The general idea is that the mediator looks at the level of "greediness" (see Definition 8) in each declaration and favours the less greedy agent by performing price adjustments on the issues represented by these declarations. For the rest of the chapter we will refer to an agent's level of greediness on an issue as how far away its actual declaration of its reservation price is from its weakest declaration, as a percentage of its entire reservation price range (in Figure 7.5 the level of greediness is shown as the shaded part of the RP range, making the declaration on the left 50% greedy and the one on the right 25% greedy).

For example, consider Figure 7.5, here we have a seller declaration for issue $I_1$ that is 50% greedy and a buyer declaration for issue $I_8$ that is 25% greedy. In this case, the mediator will punish the seller by the difference of the greediness, 25%, on issue $I_1$ (recall that this means the price on issue $I_1$ will be lowered by 25% of the seller's share of the trade balance). Note that every issue has two declarations associated with it, one from each agent, and so the single inter-agent pairing in this example does not fully define the price adjustments on issues $I_1$ and $I_8$.

The example shows that inter-agent pairings, as described so far, are very similar to the GPS mechanism. In fact, if declarations for the same issue were paired together, this would result in the original GPS mechanism as described in Chapter 5. The point of not doing this, but instead randomly choosing the inter-agent pairings is to reduce the strategic advantage of knowledge of the other agent when making declarations; we would like agents' declarations to not be based upon their counterparts, but instead on their own reservation prices. If an agent knows the importance of an issue to the other agent, this may affect the agent's own declaration on that issue. But if the agent does not know which issue its declaration will be paired with, it cannot take advantage of this knowledge and may be better off just not being greedy as in the GPS.

Agents not being greedy at all is no good, however. This was why the GPS solution did not provide any useful information in the first place, because the agents had no incentive to declare anything but their weakest position. We want agents to be greedy sometimes, to reflect their actual RP values. This is where belief from the intra-agent pairings will come into play; it will be used as a way to provide rewards to the agents as a counterbalance to the existing punishments. This will lead to agents having to strike a balance between modesty and opportunism that will lead to declarations that are neither very weak nor very strong.

### Combining the Pairings

Once the pairings are combined, there are four possible inter-agent pairings as shown in Table 7.1. All that needs to be decided is how each of these possibilities are to be handled. In the case of NB-NB, where neither agent is believed, nothing needs to be done; the current GPS like solution will suffice where the more greedy agent is punished by the difference of the amounts of greediness. The remaining symmetric case (B-B) and the two asymmetric cases need to be handled differently since belief is now involved in these pairings, and as mentioned earlier, this belief factor needs to factor into the ability to be rewarded.

The symmetric case of belief (B-B) is problematic since each agent should be believed, and it seems unfair to adjust the prices since any adjustment of prices counts as a punishment towards one agent. In this case, it is best to do nothing at all and leave the prices as they are with respect to these declarations. If we handle the symmetric cases as mentioned, then there is still no incentive for an agent to declare anything but its weakest declaration since no rewards are provided; rewards will have to be encapsulated in the asymmetric cases then.

The simplest way to reward an agent in the asymmetric case is to just reward

| Seller | Buyer | Probability($\sim$) | Type |
|---|---|---|---|
| B | B | $\frac{1}{4}$ | symmetric |
| NB | NB | $\frac{1}{4}$ | symmetric |
| B | NB | $\frac{1}{4}$ | asymmetric |
| NB | B | $\frac{1}{4}$ | asymmetric |

Table 7.1: The possible pairings of belief types in the inter-agent pairing of declarations. B = b

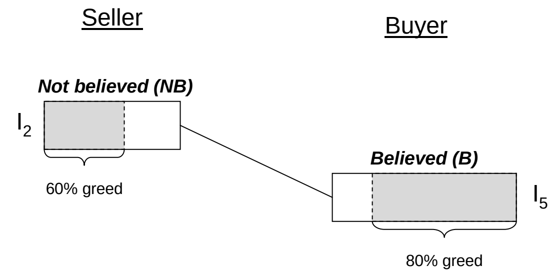


Figure 7.6: An example for illustrating how rewards might work. In this pairing the seller is not believed and the buyer is believed, the buyer will be rewarded on issue $I_5$ by its level of greed, $80\%$. No price adjustment will be made on issue $I_2$, however, observe that a reward for the buyer on issue $I_5$ is a punishment for the seller.

them the amount of their greed on the declaration issue. For example, in Figure 7.6, the buyer is believed for its issue $I_5$ declaration and has been 80% greedy so we will adjust the price on $I_5$ in favour of the buyer by 80% (which is actually the same as punishing the seller on issue $I_5$ if the seller had made a declaration that was 80% greedy). Doing this, agents will now have an incentive to declare strong reservation prices in addition to weak ones.

### 7.1.2 Questions

Now that we have explained the approach of this randomised method, it will be helpful to answer a few questions to understand how this method is different from the GPS and how it brings us closer to obtaining better information about the agents' reservation prices.

#### Why won't agents always declare their weakest reservation price?

Without loss of generality, say the seller declares its weakest price on every issue. Now, all the buyer has to do is declare something a little stronger than its weakest price on every issue and it should come out ahead. After all, the way the punishments are meted out, the buyer will only suffer on $\frac{1}{4}$ of the issues, but will gain or do the same on $\frac{3}{4}$ of the issues. If the buyer makes its declarations appropriately, it should expect to gain utility a majority of the time.

Each agent acts as the baseline for the other agent; an agent is only considered greedy if it is more greedy than the other agent. This was true in the GPS as well, but the difference here is that agents can gain a reward for being greedy, this was not true of the GPS. An agent can benefit by just being a little more greedy than the other agent, so neither agent will want to declare too weakly anymore, but at the same time they do not know how strongly to declare either.

#### Why won't agents always declare their strongest reservation price?

The answer to this question is similar to the last question. Say the seller declares its strongest price on every issue, then the buyer can declare prices that are just slightly more modest (weaker) without having to worry about punishment since no matter what the inter-agent pairings are, the seller will always be greedier and thus be the one punished in (NB-NB) situations. So the buyer could enjoy large gains when it is believed without any risk when it is not believed.

Again, each agent is the baseline for the other agent, and the best strategy is to exhibit the same level of greediness as the other agent, although there are benefits to being just above and just below the other agent's level of greediness. However, it will not necessarily be known to either agent how greedy the other agent is in general, and specific levels of greediness on particular issues (i.e. if an issue is known to be important it may cause higher or lower levels of greed) are complicated by the randomised pairing of declarations.

### 7.1.3 Conclusion

The randomised method presented here should be an improvement on the GPS solution we have been using, providing more information about the agents' private information. There is a balancing act introduced where each agent wants to be about as greedy as the other agent and now there is incentive to not be too greedy or too modest. Whether these incentives can be aligned with the actual reservation prices still needs to be determined. These ideas are preliminary, but the hope is that some randomised mechanism like the one here can be used to allow an agent to be truthful and still do well in expectation.

The idea is similar to having a balanced stock portfolio: some stocks will go up and some down, but overall a gain is experienced. This approach can be rather complicated, since we are randomising the pairings and since the agents do not know how they will be paired, they need to balance out their declarations between opportunism and honesty in order to maintain a balanced portfolio of declarations (i.e. they will stand to lose on some issues and gain on others). Due to this randomness, in expectation an agent is better off just telling the truth, or at least telling the truth is an easy option computationally that is not detrimental in general. From this viewpoint, there is hope that we can have more confidence trusting the agents.

## 7.2 Valuation Chips Improvements

The valuation chips mechanism (Section 5.3.2) being used currently provides an ordering of agents' preferences for issues with respect to their valuations. The ultimate goal is to not only acquire this ranking but to obtain a representation of the actual valuations agents have for the issues. This requires knowing how much more important one issue is over another issue. But since the value of the chips are fixed, it is not possible for agents to inform the mechanism beyond just a ranking of the issues. Clearly, some variable value mechanism is needed where the agents themselves assign points to each issue based on their private valuations. In fact, this was the original approach to this problem.

A naïve point system, that has no restrictions on how points can be assigned, reveals little useful information. If a maximum number of points per issue is set, then this will reveal a group of issues that is more important than the group of issues that have no points assigned to them. For example, if the agents are given 450 points and a 100 points/issue limit, then each agent will assign 100 points to each of its top four issues and 50 points to its fifth most important issue (see Figure 7.7). All this reveals is the top four issues for an agent (without any ordering information amongst these four issues) as well as specific identification of the fifth most important issue; no valuation information is revealed about any other issue

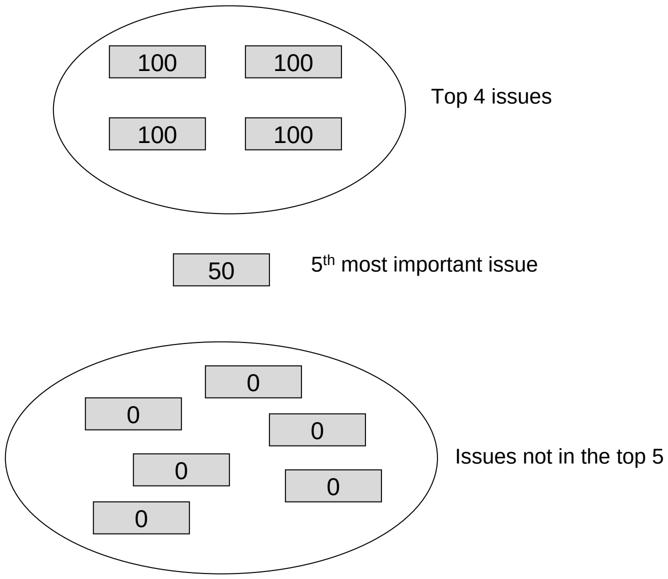


Figure 7.7: An agent is given 450 points to distribute with a 100 points/issue maximum. The result will be that 100 points are assigned to each of the top five issues, 50 points to the fifth most important issue, and no points to any other issue. This provides very little information except the identity of the fifth most important issue for the agent.

except that they are not in the top five important issues for an agent. Obviously, a better set of restrictions is needed.

In our example, the fifth most important issue was identified and this shows how to obtain rankings amongst the issues: put different maximum point limits on the issues. The way this would work would be to allow a maximum of, say, 100 points on only one issue, then 80 points maximum on another issue, then 60 on one other issue, and so on, providing the agent enough points that it can max out each of these limits. The rational agent behaviour in this case would be to put 100 points on its most important issue, 80 points on its second most important, 60 on its third most important, and so on, following the point maximums per issue as they descend. This will give an agent's preference ordering on the issues, and this is exactly our current valuation chips mechanism!

To obtain more information, further modifications to the point system need to be made. The obvious extension is to create a mechanism that uses the separation in assigned points of issues to alter the issues' prices in the initial agreement. It is not clear at the moment if this will be possible while balancing out the workings of the existing mechanism; further study is needed. However, what appears needed is some mechanism that effectively redistributes the assigned points (after they have been assigned by the agents) based on the point separation of issues.

## 7.3 Agendas

The agenda process we dealt with in Section 5.4 was relatively uncomplicated overall, but it illustrated a novel general method for choosing an agenda in a geometric manner for an incomplete information environment. However, the method we introduced can be made more complicated in order to address a greater range of criteria and situations. A more comprehensive set of quality measures are introduced in Section 7.3.1 and a further extension of agenda ideas in general is introduced in Section 7.3.2.

### 7.3.1 Comprehensive Quality Measures: DRAB

The DRAB quality measures are essentially an expansion of the simple distance measure we used in Section 5.4.3. DRAB is an acronym for Distance, Risk, Asymmetry, and Balance; four different metrics for evaluating an agenda. These various metrics are discussed below, but the idea would be to extend our quality measure for an agenda so that it is a weighted sum of these four metrics. Each of the individual metrics could be specialised to ascertain different desirable properties of an agenda, and the weighting could be used to decide how important each type of measurement is to the overall goal for the agenda.

**Definition 13 (DRAB Agenda Quality Measure)** *The expanded agenda quality measure, Q, for an agenda, A, is*

$$Q(A) \;=\; w_D D(A) \;+\; w_R R(A) \;+\; w_{Asym} Asym(A) \;+\; w_B B(A)$$

$$\begin{aligned} \textit{where} \quad w_D + w_R + w_{Asym} + w_B \;&=\; 1 \quad \textit{and} \\ w_D, w_R, w_{Asym}, w_B &\geq 0 \end{aligned}$$

*D, R, Asym, and B are individual measures that rate the quality of an agenda on different properties and are weighted by the w values.*

By making the quality measure more complex, the amount of computation time for performing an agenda assessment will likely increase significantly since the current basic distance measure (see Section 5.4.3 and Definition 12) is composed mostly of simple constant time computations. For more information about how much computation this might be with respect to how many agendas need to be evaluated, see Appendix A.

The quality measure functions used in Definition 13 are not actually specified in this section, but are descriptions of what the measures might be with respect to the goals of each measure as described below. Before going into this description though, it is helpful to consider the different classifications of agenda measurements.

Agenda quality measures can be classified into three general categories: stage dependent, agent dependent, and agenda dependent. Stage dependent measures focus on taking a measurement from the properties of individual stages; these stage measurements may be aggregated together to form the overall measurement, but the specific measurements focus mainly on single stages. Agent dependent measures are those that take into account elements of the agents' preferences, for example, their preference orderings. Agenda dependent measures are those that require looking at the entire agenda to make a measurement. This differs from aggregated stage dependent measures because a stage dependent measure can be made within a stage and still have meaning, whereas an agenda dependent measure has to consider trends across the entire agenda in order to make an assessment; a single stage provides little or no information to an agenda dependent measure. Many of the DRAB measures combine elements of some or all of the three classifications.

**Distance (D)**

Distance is the measure that is already being used in Section 5.4.3 that measures the distance between the expected solution (Definition 11) and the desired solution (Definition 10) in a stage. This measure is stage dependent.

**Risk (R)**

The risk measure is meant to assess qualities of an agenda that will make a negotiation more likely to breakdown. Besides measuring individual high risk aspects of an agenda, the measure also evaluates risk as a composite property over the entire agenda. In particular, as risk decreases as agents proceed through the agenda, their expected utilities increase.[2] Risk can be a complicated measure that considers many facets of an agenda and thus falls into all three categories of quality measures; risk is stage, agent, and agenda dependent. First we discuss stage dependent measurements and then move on to measurements that are agent and agenda dependent.

There are two forms of breakdown that we focus on in this work: endogenous and exogenous. *Endogenous breakdown* occurs if agents just cannot reach an agreement at some point in negotiations for whatever reason. The likelihood of this occurring increases when agents are too far apart to make a compromise or when things become too confusing to realise there is a compromise. If the former case is true, there is little that we can do as a mediator to relieve the situation, but if the latter is a problem, we may be able to arrange the agenda to be less complicated. *Exogenous breakdown* occurs due to factors that are outside of the negotiation that result in its deterioration, although these factors may be affected by events within the negotiation. For example, many outside factors put pressure on a negotiation the longer the negotiation continues without an agreement. This is seen in many labour disputes that become more bitter and polarising as the length of the dispute increases. It is difficult to know, in general, what these outside factors may be; however, it is likely that exacerbation can be avoided by reducing the length of negotiation.

Assessing risk then becomes a measurement of an agenda's contribution to the length and complexity of a negotiation. Fortunately, both of these things can generally be assessed by examining the size of the stages in the agenda. The size of a stage can be measured by how much utility it represents or how many issues are involved (or both). A large stage size increases the likelihood that that stage will be difficult to negotiate because there is a lot involved in that stage, making it harder for agents to determine a compromise *and* to do so in a timely manner.[3] The actual measurement of stage size can count the number of issues and might, for example, measure utility by looking at the distance between efficient frontiers of the stages in the agenda (see Figure 7.8).

To increase the quality of an agenda, risk should go down as the agenda progresses. This can be measured by taking the risk measurements just discussed and giving some weight to how well these measurements decrease from the early stages to the later stages. However, there are other considerations that need to

[2] The ramifications of risk over time within agendas is studied by John and Raith in an environment similar to the one we use [11, 12].

[3] In complete information scenarios, this is generally not true unless dealing with bounded rational agents, since all agents will know the "correct" agreement from the first time step.

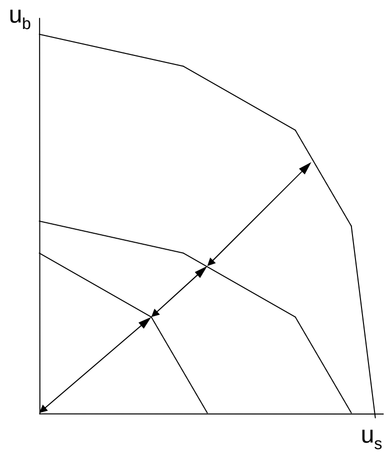


Figure 7.8: An example of how the utility size of an agenda may be measured. Some distance measure between frontiers in the stages of the agenda assesses how much utility there is per stage. A measure like this combines both agents' utility sizes into one measurement.

be balanced into the overall risk measurement that can be at odds with the specific measurements mentioned thus far. The model of negotiation we use says that agents can reach incremental agreements through the stages of the agenda, building more and more utility into the current default agreement at each successfully negotiated stage. In light of this, risk is considered to decrease over time as more and more utility is incorporated into the current agreement, since that means there is less and less utility to be lost due to breakdown. This is contradictory to the measures discussed thus far because it implies having large stage sizes early in the agenda; a balance must be found between these two ideas.

The latter viewpoint of measuring risk over time, motivates the idea that large non-contentious issues should occur early in the agenda, while small contentious issues occur later on. This way, it is more likely that large amounts of utility are easily secured within the current agreement as quickly as possible. The size of an issue can be estimated by the issue orderings of the agents and the surplus size on the issue (as we did in Section 5.4.1 when we needed to convert issues into utility space). The level of contention can be measured by how evenly agents measure an issue: if both agents think the issue is of similar importance in terms of issue ordering then the contention is higher (see Figure 7.9). This idea of contention is similar to ideas expressed in the next two sections of Issue Asymmetry and Issue

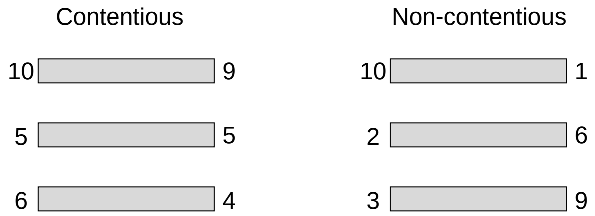


Figure 7.9: Examples of issues that are contentious and not very contentious. Issues are represented by the rectangles and the numbers on each side represent the issue ranking for each of the agents. Contention is measure by how similar the agents rank an issue in their issue orderings that are derived from the Valuation Chips mechanism (Section 5.3.2). With close rankings, contention is high, with rankings farther apart contention is low.

Balance.

### Issue Asymmetry (A)

When agents negotiate within a stage, it is desirable that each agent have something more desirable to the other agent than what the other agent currently has. When this happens, trade-offs are much easier because it is clear that the agents want to trade their share in an issue they do not care that much about for a larger share in an issue that they do care a lot about (see Figure 7.10). Typically, when the initial agreement is created for the agents, each agent gets a share of each issue as we try to distribute the trade balance on that issue. In other words, each agent is given a stake in every issue, including issues that the other agent cares more about.

The main concern of issue asymmetry is measuring the amount by which agents value issues differently within a stage. To measure this, one can sum the difference in issue ranking for each issue in a stage (taken from the issue orderings from the Valuation Chips mechanism in the arbitration module). The average issue asymmetry over all the stages in the agenda can then be used as the issue asymmetry measurement for the agenda. Issue asymmetry is both a stage and agent dependent measure.

See Figure 7.11 for an example of how to calculate this measure for an agenda. Notice that we do not use signed differences, just absolute differences, for the issue rankings.

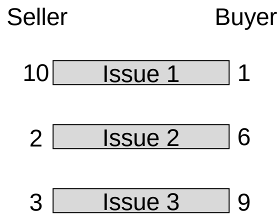


Figure 7.10: An example of trade-offs within a stage. The numbers on the sides represent the rankings in importance of the issues, with low numbers meaning higher importance. If these issues represent a stage in the agenda, then the seller wants to use any share it has in the Issue 1 (which the seller does not care about) as leverage to get a larger share in Issues 2 and 3. This may be possible since the buyer does not care too much about these issues, but values Issue 1 very highly.

### Issue Balance (B)

Issue balance has a lot in common with issue asymmetry and might be combined with the previous measure since their differences can be subtle. In issue asymmetry, the focus was on how differently agents ranked the issues in a stage, but not actually whether this asymmetry was balanced. That is, if all the issues in a stage where actually valued higher by the buyer than the seller, this still gave a good asymmetry score for that stage. But if this is the case, then the seller has no reason to budge from the current agreement in that stage since the buyer has nothing the seller wants. Thus in order to make negotiations proceed smoothly and have agents capitalise on joint gains, it is desirable to have asymmetry but also to have a balance between the asymmetry within a stage so that both agents have something to offer.

Issue balance, then, is the measure of the balance of issues within a stage. This can be measured similarly to asymmetry by looking at the difference between rankings of the issues, but instead the *signed* difference is used. Summing these signed differences of individual issue rankings within a stage means that a score close to 0 is a good balance score for a stage. Taking the average of the *absolute values* of these sums[4] over all stages forms the overall balance score of the agenda (see Figure 7.11 for an example of calculating the issue balance of an agenda). Similar to asymmetry, issue balance is both a stage and agent dependent measure.

[4]If we just take the average of the signed sums over all stages, it will always be zero since the sums will always add up to zero since each agent has the same set of rankings when we consider all the issues together.


Figure 7.11: An example of calculating issue asymmetry and balance. The nine issues are separated into three stages, with the issues rankings for each agent shown on the sides of the issues. The individual asymmetry and balance for each issue are shown with stage totals shown just to the right of the issues. The overall agenda asymmetry and balance are shown at the far right and are averages of the stage scores (except the balance measure first takes the absolute value of the stage scores). A high asymmetry score and a low balance score are desirable, however the particular scores shown have no meaning until they are compared against other agendas.

Note that a good balance score is achieved by asymmetric issues that "cancel" each other out in terms of their direction of asymmetry, *but also* can be achieved by sets of issues that are very symmetrical. Since the desire of the issue measures is to create stages that give each agent something to offer the other agent, making trading easier, it is necessary to have *both* good balance and good asymmetry measures together to attain this goal.

### 7.3.2 Further Agenda Extensions

In the previous section we focused on extending our particular agenda creation method, but here we introduce more general agenda extensions that can be used with or without our framework. By adding dynamic elements to agendas, these extensions grant further levels of control to the mediator over the negotiation process occurring in the Negotiation Module.

#### Blind Agendas

Normally, when the negotiation module begins both agents can see the entire agenda that was created by the mediator, allowing agents to lookahead while planning their stage negotiations. In *blind agendas*, the agents are only allowed to see the agenda up to, and including, the current stage. This is a pseudo-dynamism because the mediator cannot change the agenda once the negotiation module takes over. There are probably lots of ways to use this, and we discuss a few here. However, the appeal of blind agendas seems limited and the main reason for introducing it is for motivation for the next two topics.

By not letting agents know how far away particular issues are, it makes it more likely that agents need to "play nice" to get to unresolved issues that are important. This is because they do not know when they will be negotiating these issues and do not want negotiations to breakdown prematurely. In addition, a blind agenda also focuses agents on the current stage, which may be useful to more limited agents that may get ahead of themselves trying to plan according to the whole agenda.

#### Dynamic Agendas

Currently, once the negotiation module begins, the agenda is static and can no longer be changed. This means that if the agents negotiate in such a way that the agenda created for them is no longer assisting them, nothing can be done. It would

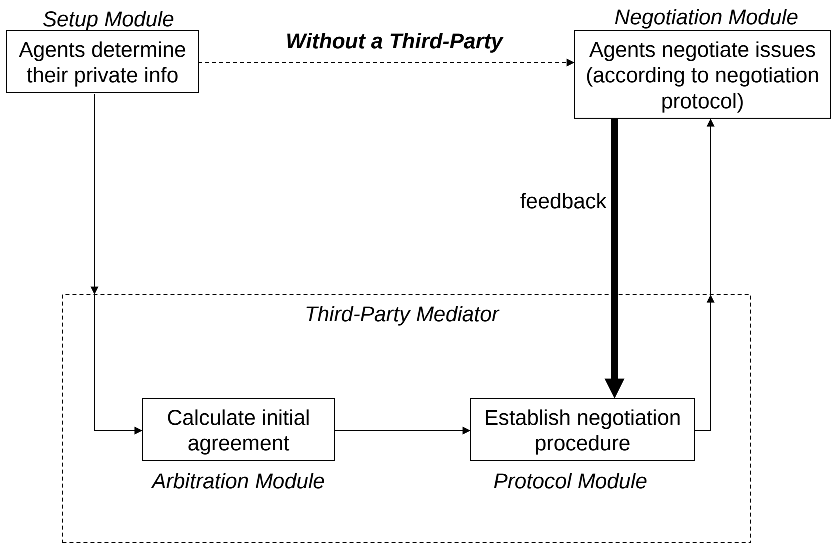


Figure 7.12: Updating Figure 5.2 with a feedback loop between the negotiation and protocol module (highlighted in this figure), that allows the use of dynamic agendas. Dynamic agendas allow modification of the agenda based on how the actual negotiations proceed.

be desirable to change the agenda to reflect the actual behaviour of the agents as they progress through the stages. A *dynamic agenda* creates a feedback loop from the negotiation module to the protocol module (see Figure 7.12), that allows the agenda to be changed dynamically based on the stage negotiations of the agents. Obviously, only future stages can be modified in this way. Since future stages may change, it may be best to also use a blind agenda approach so that agents are not misled by seeing future stages that may not be used; effectively for the agents, a dynamic agenda would appear the same as a blind agenda.

One way this technique can be used is for re-evaluating agendas after every stage's negotiations: the mediator recalculates the best agenda using the same criteria but only assessing the agendas based on the stages that have not yet been negotiated. In addition, the knowledge of the current agreement can be used in place of the initial agreement that was used to create the original agenda, the current agreement having the additional information of how agents negotiated all the previous stages. Figure 7.13 gives an example of how using the current agreement differs from using the initial agreement.

This sort of dynamic re-evaluation of agendas means that if agents stray too

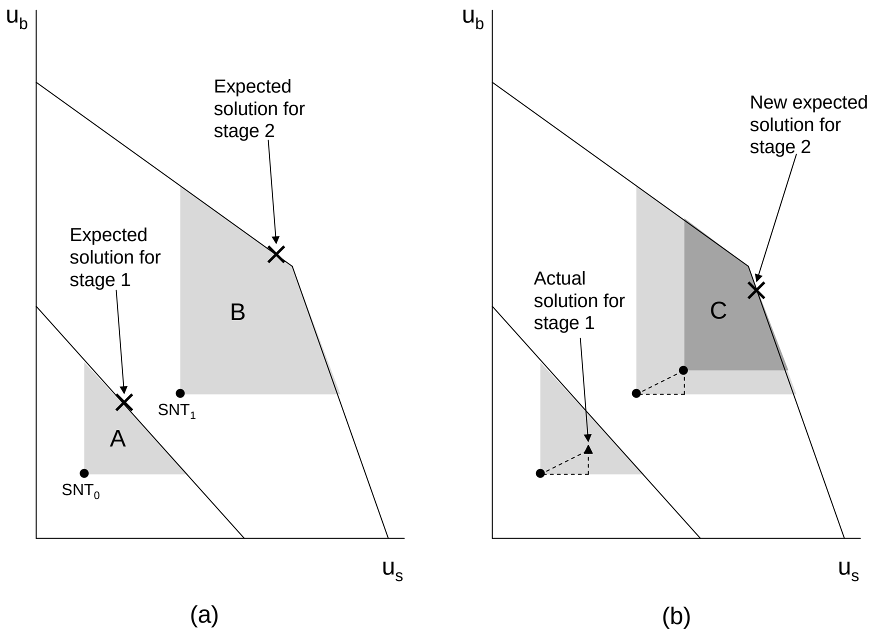


Figure 7.13: (a) The expected solutions using a regular agenda. (b) The expected solutions using a dynamic re-evaluation of the agenda to include the current agreement after stage 1. Note that the space of viable solutions in the second stage is now smaller than in part (a).

much from the direction they should be traveling in the solution space, the agenda can be modified to correct for this deviation. However, it may not always be desirable to do this since this deviation may actually be the agents fitting the initial agreement to be more accurately aligned with their true preferences (of which the mediator is not privy to).

### Agenda Trees

*Agenda trees* are a combination of both blind and dynamic agendas, but they are combined in a way that allows static creation of the agenda tree such that a feedback loop from the negotiation module is unnecessary. Static creation can be desirable since it can be costly and lengthy to re-evaluate agendas while negotiations are proceeding, causing delays in the negotiation process when the agents are involved.

The idea behind an agenda tree is to precalculate backup plans for when agents deviate from following along the path we are trying to lead them. If agents generally follow the path that is considered desirable, then the original agenda is used unaltered. But if agents stray too far off the path, then the current agenda is replaced with a different agenda that is more likely to help get agents back on track. This is done by separating the solution space of each stage into different areas (the shape of which are arbitrary and decided by the mediator framework), and for each area assigning an agenda. If agents negotiating in that stage produce an agreement within one of these areas, then the agenda associated with that area is used in the next stage, this is illustrated in Figure 7.14.

Agenda trees actually do require information from the negotiation module, but this information only needs to be used within the finite state machine that is the agenda tree, and does not have to go back to the mediator framework for a decision to be made. This allows the mediator to create a flexible agenda plan that accounts for contingencies without having to constantly recalculate the agenda during the negotiation process between the agents. Although dynamic agendas allow a greater level of control and accuracy for guiding agents back to desirable areas of the solution space, agenda trees provide a similar functionality statically. One thing to keep mind, however, is that a lot of work goes into generating an agenda tree since each agenda featured in the tree has to have its stages' utility spaces divided up and choose an agenda for each division. Figure 7.14 is a bit deceptive because it makes it look like a single agenda's utility space is divided for creating the tree, but every agenda can start a similar tree.

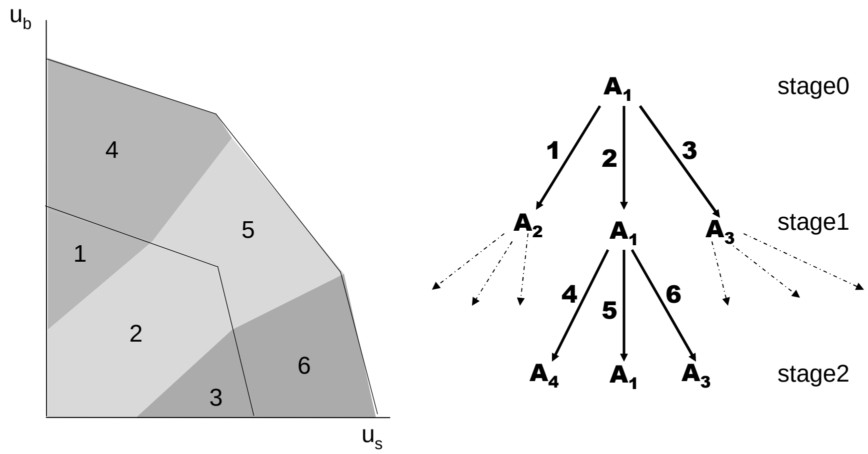


Figure 7.14: Example of agenda trees. The mediator chooses an initial agenda, in this case agenda $A_1$, and creates an agenda tree as shown on the right; the utility space of the first two stages of agenda $A_1$ are also shown. These spaces are divided up into the sections 1 through 6, and for each section an agenda is assigned to it. In the first stage, agents negotiate the new current agreement, depending on where this agreement lies in the utility space will determine which agenda will be used in stage 2. If the agreement is in section 1, then the tree tells us that agenda $A_2$ will be used in stage 2. Similarly, agreements in sections 2 and 3 will result in using agendas $A_1$ and $A_3$ in stage 2, respectively. As the agenda tree shows, if the agents keep their solutions in the centre sections (sections 2 and 5), which is likely where the mediator is trying to guide the agents, then the agenda tree will stay with the original agenda $A_1$. If agents stray to either side, then other agendas will come into use to account for this deviance.

# Chapter 8

# Conclusion

In the BAM framework we brought together many different concepts, techniques, and tools, in order to approach a problem from many different angles but integrated into one endeavour. Many of these ideas and artifacts were taken from different contexts and settings, but have been utilised in new ways as they got woven into our system; some are familiar stalwarts still very useful in complimenting our system. Some highlights of the methods and ideas we used are:

- Single Negotiating Texts (SNTs) were used as a way to focus agents on a single agreement to incrementally improve and move through the solution space.
- Agendas have been geometrically re-imagined as a tool for guiding agents along desirable negotiation pathways.
- With SNTs and agendas combined, we have conceived ways of constraining the solution space, creating a potent tool for illuminating profitable solution subspaces for the agents to explore.
- The increasing-pie model of John and Raith [11] is embedded in renegotiating agendas that achieve a balance of flexibility and simplicity, while still allowing efficient agreements to emerge.
- Game-theoretic principles have been incorporated in our arbitration module mechanisms to create initial agreements with provable properties.

All of these ideas have been brought together under one umbrella tasked with facilitating a positive negotiation environment for the agents, being a mediator that guides them to jointly beneficial agreements. The instantiation of the BAM framework developed in this thesis introduced an interesting new way to combine these ideas and tools into a working mediator that helps guide agents in an environment of uncertainty created by an incomplete information setting.

The assumption that all agents are rational is not realistic and undesirable in a framework that deals with very general situations, which is why the BAM framework did not assume that agents have a particular level of sophistication. However, the framework does keep in mind that agents are an incredible asset, being the only stakeholders in the negotiation that (collectively) know the true global information state. The framework expanded Bartal et al.'s [2] *multilateral* market mechanism idea of having agents finalise the agreement to a broader way of approaching *bilateral* negotiations. This opens up a whole new avenue for bilateral negotiation research where agents are key actors in finalising the third-party aspirations of the researchers. This is not necessarily an easy research methodology, but neither are incomplete information negotiations easy to "solve". We have already seen a number of possible advancements of our particular system in the previous chapter, however it would be very interesting to see the various ideas the research community could explore within the general BAM framework, which embodies the relatively new paradigm of using a mediator mechanism that allows agents to finalise the final agreement themselves.

# Appendix A

# The Number Of Agendas

In Section 5.4.3, it was mentioned that an agenda is chosen by examining every possible agenda. Naturally, one might wonder how much work this entails when the naïve number of agendas to consider is $n!$. This factorial number of agendas, however, is based on using a stage size of one issue for every stage. If we consider a fixed stage size of $z$, let $f = \lfloor \frac{n}{z} \rfloor$ and $F = \lceil \frac{n}{z} \rceil$ ($F$ is the number of stages in the agenda). Then we have:

- $\binom{n}{z}$ possible choices for the first stage
- $\binom{n-z}{z}$ possible choices for the second stage

$$\vdots$$

- $\binom{n-(f-1)z}{z}$ possible choices for the second to last stage
- $\binom{n-fz}{z}$ choices for the last stage. Although this value will always equal 1 because $n - fz \leq z$.

This means the number of agendas we have with a fixed stage size of $z$ is:

$$\binom{n}{z}\binom{n-z}{z} \quad \cdots \quad \binom{n-(f-1)z}{z}\binom{n-fz}{z}$$

$$\Rightarrow \underbrace{\frac{n!}{z!\cancel{(n-z)!}} \cdot \frac{\cancel{(n-z)!}}{z!\cancel{(n-2z)!}} \cdot \quad \cdots \quad \cdot \frac{\cancel{(n-(f-1)z)!}}{z!(n-fz)!}}_{f terms} \cdot 1 \tag{A.1}$$

$$\Rightarrow \frac{n!}{(z!)^f(n-fz)!} \tag{A.2}$$

Thus we have $\frac{n!}{(z!)^f(n-fz)!}$ agendas, which can be a significant improvement over the naïve $n!$ agendas, due to the savings factor, $(z!)^f(n-fz)!$. Table A.1 gives a comparison for the total number of agendas per stage size when we have 10 issues, and Figure A.1 shows the same results graphically. It is easier to see in the graph (Figure A.1) that as we exceed half the number of issues per stage, the decrease in the number of agendas is reduced.

| z = issues per stage | $\frac{n!}{(z!)^f(n-fz)!}$ | Savings factor $(z!)^f(n-fz)!$ |
|---|---|---|
| 1 | 3628800 | 1 |
| 2 | 113400 | 32 |
| 3 | 16800 | 216 |
| 4 | 3150 | 1152 |
| 5 | 252 | 14400 |
| 6 | 210 | 17280 |
| 7 | 120 | 30240 |
| 8 | 45 | 80640 |
| 9 | 10 | 362880 |

Table A.1: The number of possible agendas when we have 10 issues and varying numbers of issues per stage. Also shown is the savings factor achieved over using $z = 1$.

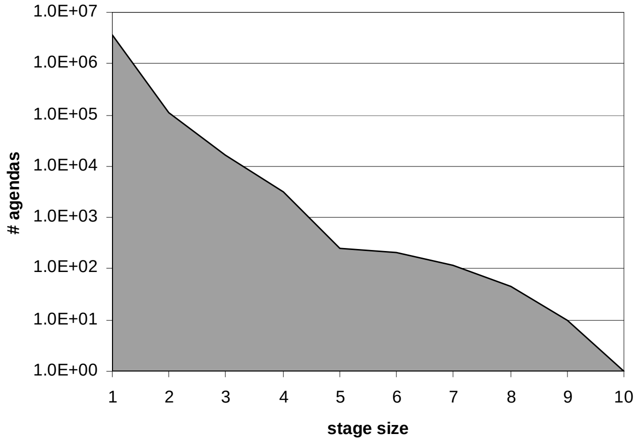


Figure A.1: The number of agendas for a negotiation of 10 issues with varying stage sizes. This is the graphical result of Table A.1. Note that the y-axis is logarithmic and thus this graph shows an exponential decline in the number of agendas with increasing stage size. After 5 issues per stage, the rate of decrease declines because $f = 1$ in $(z!)^f(n - fz)!$ since the stage size is now more than half the number of issues, thus we lose that exponential factor in $(z!)^f$. However, there is still a fairly steady decline as stage size increases.

## A.1 Amount of Computation

We have determined how many agendas we actually have to look at when choosing an agenda, but how much computation do we need to perform on each agenda? With the relatively simple quality measure we defined in Definition 12(Section 5.4.3), the actual computation per agenda is quadratic in the number of issues, but consists of very simple constant time calculations. It is quadratic because performing the necessary line intersections for determining the desired and expected solutions for each stage requires looking at each edge in each stage's efficient frontier. Since the number of edges in a frontier is the same as the number of issues in the stage, we have an approximately linear calculation over an approximately linear number of stages. This is really overestimating quite a bit especially as the stage sizes increase (and thus there are fewer stages), but quadratic is the upper bound on the work we perform for a single agenda.

However, the point we really want to make is that even despite being quadratic and having a potentially large number of agendas, the computation for a realistic number of issues, anywhere from 4 to 20, is still easily computable time-wise[1] As we get into large numbers of issues, say 30, things can become a little unwieldy if we do not perform any reductions to make the computation for choosing an agenda simpler. The intent of this system is not necessarily to be a real-time mediator, thus larger numbers of issues are still not that much of a problem as computations can be performed offline. However, it is always nice to find more efficient algorithms.

## A.2 Computational Reductions

The computational reductions we are going to achieve come about from how agendas are constructed and evaluated. We show that there are a lot of overlapping subproblems in creating agendas because many efficient frontiers appear in several agendas and the calculations performed on such frontiers need only be done once per frontier. In addition to overlapping subproblems, it is the case that we evaluate an agenda by breaking it down into its individual stages and evaluating each stage and then summing the results to determine the overall evaluation of the agenda. What we are alluding to here is that when considering many agendas that have overlaps in their stages, we are able to take advantage of an optimal substructure property. Together, these two properties allow us to take advantage of dynamic programming to achieve significant savings in choosing the best agenda to use in our negotiations.

### A.2.1 Preliminary Discussion

Consider counting the number of agendas as the leaf nodes of a tree being built from a root of no issues, while each branch represents adding a new stage of issues to the agenda. We end up with a tree like in Figure A.2 where the branching factor at each level follows the terms in Equation A.1. In this tree, each node represents a partial agenda composed of $d \cdot z$ issues where $d$ is the depth of the node.

Evaluating each agenda involves assessing the quality at each stage of the agenda. The quality measure is made for a stage by looking at its efficient frontier and calculating the desired solution and the expected solution.[2] It is important to

[1] Less than a second on a 1.8 GHz first-generation dual-core processor with 2GB of RAM.

[2] In order to calculate the desired solutions, we do have to determine the final negotiation frontier and determine the global Nash Bargaining Solution. However, this frontier and NBS solution are the same for all possible agendas, so this only needs to be done once. To calculate the expected solutions requires the initial agreement, which also remains constant for all the agendas.

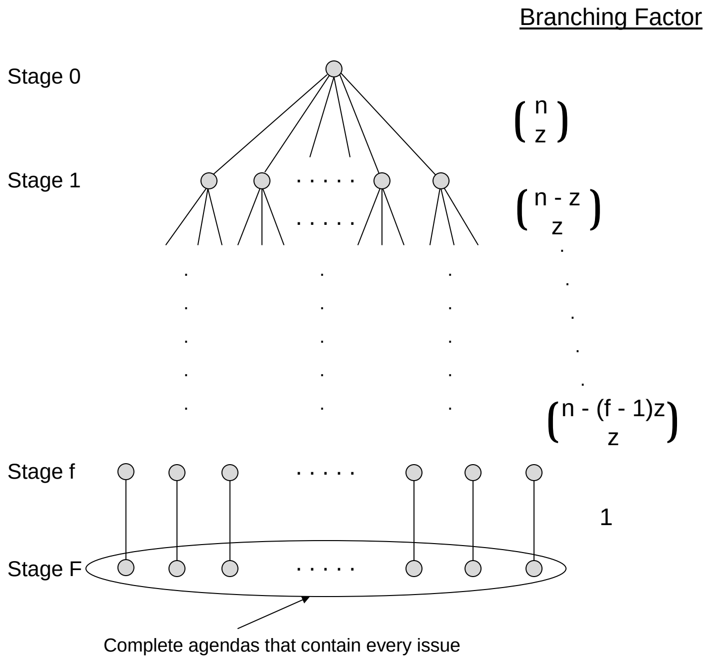


Figure A.2: The tree shows how we build the agendas by adding stages as our branches and the agendas are the leaf nodes. The tree uses n issues and a stage size of z.

realise that if two agendas have a stage with the same efficient frontier in the utility space, all the calculations are the same for both of these stages, which means we should really only do them once and store them.

Now let us consider a small example of four issues $I = \{a, b, x, y\}$ with a stage size of 2 issues, meaning we have 6 possible agendas.[3] If we paid no attention to overlapping subproblems (meaning stages that have the same frontier) within the agendas we would have a tree as in Figure A.3 which is an instantiation of the tree in Figure A.2. Whereas if we did observe the overlap we would have something like that shown in Figure A.4.

In the latter figure (Figure A.4), at stage 0, we have no issues in our agenda. At stage 1 we have six partial agendas consisting of the issues *ab, ax, ay, bx, by*, and *xy*. At stage 2, we will actually have included all the issues in the negotiation (remember that our negotiation process allows negotiation on all issues in the current stage and all issues in previous stages, as discussed in Section 2.4.3). Every one of the

[3] $\frac{4!}{(2!)^2(4-2\cdot 2)!} = 6$

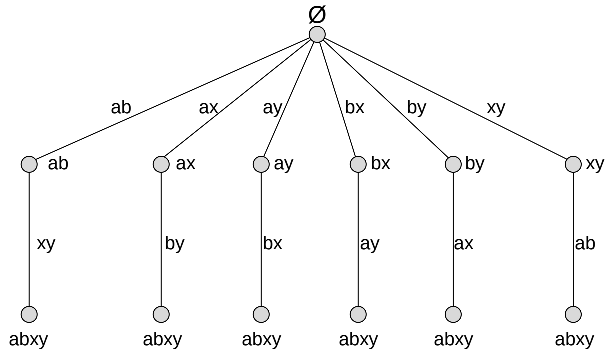


Figure A.3: A small example of an agenda that does not combine overlapping nodes.

partial agendas from stage 1 will branch to the one single agenda that includes all the issues, *abxy*. Of course, this makes sense because at stage 2 every agenda includes all issues and thus the efficient frontier is the same at this stage for every possible agenda. What this means is that at stage 2, we really only need to perform one set of calculations instead of six.

Although the example just mentioned is very small it illustrates how overlapping subproblems can occur: when nodes at the same depth of the tree actually contain the same issues, the nodes can be combined into one node (this is seen in Figure A.3, where the six *abxy* nodes have been collapsed into the single *abxy* node in Figure A.4). If this is kept in mind, then it is easy to see that there are lots of overlapping subproblems.

At the first stage of the "agenda graph" there will be $\binom{n}{z}$ nodes because there are that many ways to choose $z$ issues from $n$ issues. At the second stage, the naive graph shown in Figure A.2 will have $\binom{n}{z}\binom{n-z}{z}$ nodes. *However*, it is not possible for all of those nodes to represent unique sets of issues because there are only $\binom{n}{2z}$ possible unique sets of issues at stage 2, since there are only $\binom{n}{2z}$ ways to choose $2z$ issues from $n$ issues ($2z$ being two stages worth of issues). We see below that $\binom{n}{2z}$ is less than $\binom{n}{z}\binom{n-z}{z}$, proving that there must be some overlap.

$$\begin{aligned}\binom{n}{z}\binom{n-z}{z} &= \frac{n!}{(s!)^2(n-2s)!} \\ \binom{n}{2z} &= \frac{n!}{(2s)!(n-2s)!}\end{aligned}$$

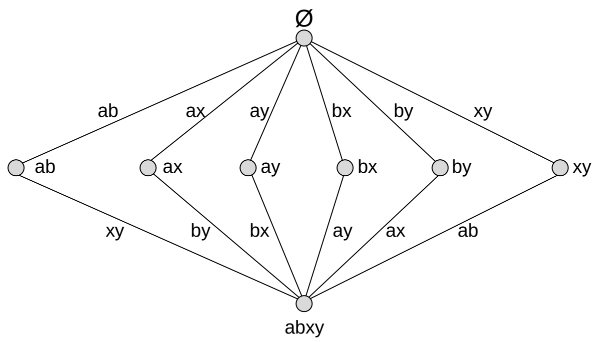


Figure A.4: The same example from Figure A.3 but combining the final nodes into one node since every agenda uses it.

The only difference in these two equations are the terms $(s!)^2$ and $(2s)!$.

$$\begin{aligned}(s!)^2 &= (1 \cdot 2 \cdots (s-1) \cdot s)(1 \cdot 2 \cdots (s-1) \cdot s) \\ (2s)! &= (1 \cdot 2 \cdots (s-1) \cdot s)((s+1) \cdot (s+2) \cdots (2s-1) \cdot 2s)\end{aligned}$$

$$\begin{aligned}(1 \cdot 2 \cdots (s-1) \cdot s) &< ((s+1) \cdot (s+2) \cdots (2s-1) \cdot 2s) \\ \Rightarrow (s!)^2 &< (2s)! \\ \Rightarrow \binom{n}{z}\binom{n-z}{z} &> \binom{n}{2z}\end{aligned}$$

This result is easily generalisable to other stages deeper in the graph. At a depth of $d > 0$ in the graph, the number of nodes is $\frac{n!}{(s!)^d(n-ds)!}$, but the maximum number of nodes is $\binom{n}{dz} = \frac{n!}{(ds)!(n-ds)!}$, and $(s!)^d < (ds)!$. There is potentially a lot of overlap, and the tree in Figure A.2, with nodes combined to account for overlap, actually becomes the graph in Figure A.5.

In conclusion, we have shown there are $\sum_{k=0}^{\lfloor \frac{n}{s} \rfloor} \binom{n}{kz} \leq 2^n$ nodes or efficient frontiers to evaluate. Thus the number of *nodes* needed to be evaluated is down to exponential, and these nodes can be used to represent all possible agendas. This is

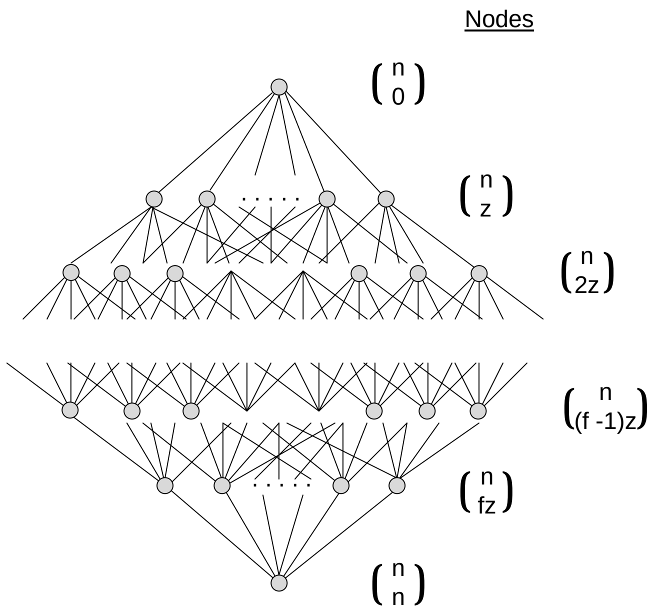


Figure A.5: Graph showing the minimum number of nodes(or efficient frontiers) that need to be evaluated in order to examine all possible agendas using a stage size of $z$ on a negotiation of $n$ issues. We can see that compared to the expanding tree shape in Figure A.2 there are a lot of overlapping nodes which results in a diamond shape to the graph.

a significant savings of computation as illustrated in Table A.2 where we recreate Table A.1 but add in the number of nodes needed when we collapse common nodes and when we do not.

| z<br># issues per stage | # agendas<br><br>$\frac{n!}{(z!)^f(n-fz)!}$ | # nodes<br>(without overlap) | # nodes<br>(with overlap)<br>$\sum_{k=0}^{f}\binom{n}{kz}$ |
|---|---|---|---|
| 1 | 3628800 | 36288002 | 1024 |
| 2 | 113400 | 567002 | 512 |
| 3 | 16800 | 50402 | 342 |
| 4 | 3150 | 6302 | 257 |
| 5 | 252 | 506 | 254 |
| 6 | 210 | 212 | 212 |
| 7 | 120 | 122 | 122 |
| 8 | 45 | 47 | 47 |
| 9 | 10 | 12 | 12 |

Table A.2: Showing the difference between collapsing common nodes in our agenda counting graphs and not collapsing them ("with overlap" and "without overlap", respectively). Up until stage sizes of one half the number of issues we see a savings in nodes. Note that $f = \lfloor \frac{n}{z} \rfloor$, and also we have added two nodes to each entry to account for the root node and the bottom node that represents the final frontier for all issues.

### A.2.2 Dynamic Programming

Our algorithm for choosing the best agenda according to the quality measure used is simply a depth-first traversal of the graph in Figure A.5, where at each node we do the following:

1. Calculate the efficient frontier line segments for the current node. This is usually based off of the frontier that is provided by the parent node, which would then just require adding a few new issues to that existing frontier.

2. Calculate the desired solution from the current frontier and the solution line (the solution line is calculated before the traversal of the graph).

3. Calculate the expected solution from the current frontier and initial agreement.

4. Calculate the quality measure from the expected and desired solutions and record its as this node's quality value.

5. Recursively call the depth first node-evaluation traversal on all possible child nodes.

6. Create a reference to the child node that had the best quality measure and add that quality value to our current node's quality value (this step here is where we store results for later memoization).

7. Return the current node's quality value to the parent node.

In doing this traversal, each time we evaluate a node, we record its quality value (which is the best quality value from this stage and onward that can be achieved by using this node) and a link to the node at the next stage that will lead to the best possible agenda from this node. When the traversal encounters a node (frontier) more than once, this recorded quality value can be used instead of evaluating the node again; this is why the savings in Table A.2 are achieved.

At the end of this traversal, the graph in Figure A.5 becomes a DAG (directed acyclic graph) as in Figure A.6. Determining the best agenda is then just a matter of traversing the DAG from the root to the node at the bottom.

Since our quality measure (see Definition 12) is a minimising measure, we are essentially just performing a shortest path dynamic programming algorithm for our graph.

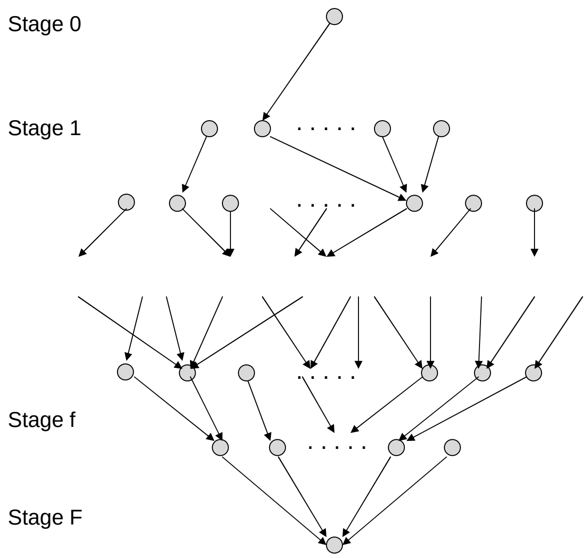


Figure A.6: Directed acyclic graph (DAG) that is the result of performing our node-evaluation traversal on a graph like in Figure A.5. The directed edges indicate the node/frontier to use in the next stage to achieve the best quality agenda from that point onward. Some nodes in the graph will not have incoming edges to them, but all nodes will have one outgoing edge to the next stage nodes.

# Bibliography


[1] Robert Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

[2] Yair Bartal, Rica Gonen, and Pierfrancesco La Mura. Negotiation-range mechanisms: Exploring the limits of truthful efficient markets. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 1–8, 2004.

[3] M. Beer, M. d'Inverno, M. Luck, N. R. Jennings, C. Preist, and M. Schroeder. Negotiation in multi-agent systems. *The Knowledge Engineering Review*, 14(3):285–289, 1999.

[4] Eithan Ephrati and Jeffrey S. Rosenschein. Deriving consensus in multi-agent systems. *Journal of Artificial Intelligence*, 87(1–2):21–74, November 1996.

[5] S. S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation strategies for agents with incomplete information. In *Proc. 8th Int. Workshop on Agent Theories, Architectures and Languages (ATAL)*, pages 53–68, 2001.

[6] S. S. Fatima, M. Wooldridge, and N. R. Jennings. On efficient procedures for multi-issue negotiation. In *Proc. 8th Int Workshop on Agent-Mediated Electronic Commerce*, pages 71–84, 2006.

[7] Shaheen S. Fatima, Michael Wooldridge, and Nicholas R. Jennings. Multi-issue negotiation with deadlines. *Artificial Intelligence Research*, 27:381–417, 2006.

[8] Chaim Fershtman. The importance of the agenda in bargaining. *Games and Economic Behavior*, 2:224–238, 1990.

[9] Herbert Gintis. *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton University Press, Princeton, New Jersey, USA, 2000.

[10] N. R. Jennings, S. Parsons, C. Sierra, and P. Faratin. Automated negotiation. In *Proc. 5th Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Systems*, pages 23–30, 2000.

[11] Reinhard John and Matthias G. Raith. Strategic step-by-step negotiation. *Economics*, 70(2):127–154, June 1999.

[12] Reinhard John and Matthias G. Raith. Optimizing multi-stage negotiations. *Economic Behavior and Organization*, 45(2):155–173, June 2001.

[13] Ehud Kalai. Solutions to the bargaining problem. In Leonid Hurwicz, David Schmeidler, and Hugo Sonnenschein, editors, *Social Goals and Social Organization: Esssays in Memory of Elisha A. Pazner*, pages 77–106. Cambridge University Press, Cambridge, UK, 1985.

[14] N. Karunatillake, N. R. Jennings, I. Rahwan, and S. D. Ramchurn. Managing social influences through argumentation-based negotiation. In *Proc. 3rd Int Workshop on Argumentation in Multi-Agent Systems*, pages 35–52, 2006.

[15] S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence Journal*, 102(1–2):1–69, 1998.

[16] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory.* Oxford University Press, New York NY, USA, 1995.

[17] N. Matos, C. Sierra, and N. R. Jennings. Determining successful negotiation strategies: an evolutionary approach. In *Proc. 3rd Int. Conf. on Multi-Agent Systems (ICMAS)*, pages 182–189, 2006.

[18] Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. In *AAAI-06: The Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 2006.

[19] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Economic Theory*, 29:265–281, 1983.

[20] John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, April 1950.

[21] John W. Pratt and Richard Zeckhauser. Multidimensional bargains and the desirability of ex post inefficiency. *Risk and Uncertainty*, 5:205–216, 1992.

[22] Howard Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, Massachusetts, USA, 1982.

[23] Robert W. Rosenthal. Arbitration of two-party disputes under uncertainty. *The Review of Economic Studies*, 45(3):595–604, Oct. 1978.

[24] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109, 1982.

[25] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach.* Pearson Education, Upper Saddle River, New Jersey, USA, 2nd edition, 2003.

[26] I. Ståhl. *Bargaining Theory.* Stockholm School of Economics, Stockholm, 1971.

[27] Thomas Weber and Abhishek Bapna. Bayesian incentive compatible parameterization of mechanisms. *Journal of Mathematical Economics*, 44(3-4):394–403, 2008.