Autonomous Agents in Bargaining Games

An Evolutionary Investigation of Fundamentals, Strategies, and Business Applications

Autonomous Agents in Bargaining Games

An Evolutionary Investigation of Fundamentals, Strategies, and Business Applications

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de Technische Universiteit Eindhoven, op het gezag van de Rector Magnificus, prof.dr. R.A. van Santen, voor een commissie aangewezen door het College voor Promoties in het openbaar te verdedigen op dinsdag 6 juli 2004 om 16.00 uur

door

Enrico Harm Gerding

geboren te Haarlem

Dit proefschrift is goedgekeurd door de promotoren:

prof.dr.ir. J.A. La Poutré
en
prof.dr. H.M. Amman

The research reported in this thesis has been carried out at the Centrum voor Wiskunde en Informatica (CWI) under the auspices of the Instituut voor Programmatuurkunde en Algoritmiek (IPA). The research is part of the project "Autonomous Systems of Trade Agents", funded by the Telematica Instituut.

CIP-DATA LIBRARY TECHNISCHE UNIVERSITEIT EINDHOVEN

Gerding, Enrico Harm

Autonomous agents in bargaining games : an evolutionary investigation of fundamentals, strategies, and business applications / by Enrico Harm Gerding. - Eindhoven : Technische Universiteit Eindhoven, 2004. - Proefschrift. -
ISBN 90-386-2047-0
NUR 788
Keywords: Bargaining / Evolutionary algorithms / Electronic commerce / Multi-agent systems

IPA Dissertation Series 2004-08

Cover by Tobias Baanders

Acknowledgement

I am grateful to many people for their support and collaboration in the course of my Ph.D. research. First of all, I thank by supervisor and first promotor Han La Poutré. Not only have the many discussions been inspiring and motivating, but I have also learnt a lot about research in general. I also thank my colleagues at the CWI with whom I've had the pleasure to co-author papers: David van Bragt, Koye Somefun, Sander Bohte, and Pieter Jan 't Hoen. Their input has been invaluable. A special thanks goes to my CWI roommates and former roommates Floortje Alkemade, Michiel de Jong, and Ivan Vermeulen for their friendship, discussions, and pleasurable talks. I also thank my other colleagues and former colleagues who have contributed through journal clubs, seminars, and other ways: Tomas Klos, Jano van Hemert, Rob van Stee, Valentin Robu and Erich Kutschinski.

Furthermore, I am grateful to my second promotor, Hans Amman, for his comments and support, and to the other people in the kernel committee, Nick Jennings and Rudolf Müller, for reading my thesis and providing valuable comments.

I would also like to thank my many friends who have helped me to keep my mind off the research! Last but not least, a very special thanks to my parents and to Bethany for your love and support.

Contents

Chapter 1

Introduction

Autonomous software agents are considered by many as the next step in computer automation. Given a set of goals and tasks, an autonomous agent will try to maximally satisfy the interests of its owner. These agents should be capable of autonomously performing certain tasks which are currently done manually, like searching for information on the Internet, planning, booking a holiday, and buying and selling goods and services.

Especially in the field of electronic commerce, an increased use of autonomous agents is expected [30, 56, 65, 70, 92, 119, 144]. Such agents should be able to autonomously negotiate with other agents about the price and other relevant aspects of a product or service, such as delivery time, quality, quantity, payment methods, and return policies. Furthermore, the agents should be adaptive in order to cope with diverse and changing environments. Current electronic markets are becoming increasingly transparent with low search costs. From a business perspective, this potentially results in strong price competition and low margins, with a negative effect on aspects such as quality and service. Through automated bargaining about a multitude of aspects, a business can go beyond price competition and gain a competitive advantage by personalising products and services to the needs of individual customers.

In such a setting, where multiple self-interested adaptive agents perform complex negotiations, the key question is how they will behave in a given environment and with specific rules of interaction. Moreover, an important challenge is to find effective bargaining strategies for the agents, and, if the rules can be changed, to determine the set of rules that achieves the best results. These are the main issues addressed in this thesis.

Game theory is a field that studies the behaviour of interacting agents and can be used to address the above issues through mathematical analysis. The limitation of game theory, however, is that many restrictive assumptions need to be made in order for a mathematical analysis to be feasible. Commonly made assumptions are, for example, that the agents act rationally and are completely informed. This means that the agents completely understand the rules of the game, have infinite reasoning capabilities, make no mistakes, and know all that needs to be known about the world and other agents' preferences to derive optimal outcomes. If such agents really existed, games like chess would no longer be a challenge. In reality, both humans and computational agents have only limited forward looking capabilities and information; instead, many tasks are learned through experience, by a process of trial and error. To analyse such settings with so-called boundedly rational agents, computer simulations are a helpful addition to the set of game-theoretic tools.

In this thesis we consider the setting where agents are adaptive to their environment, and learn effective bargaining policies by trial and error. We apply learning techniques from the field of artificial intelligence, specifically evolutionary algorithms, to model the adaptive nature of bargaining agents in practical settings. In the first part of the thesis, we consider fundamental aspects of bilateral bargaining between a buyer and a seller. We first validate the evolutionary model for bilateral bargaining by comparing the outcomes with game-theoretic results of relatively simple bargaining settings. We then investigate several extensions of game-theoretical bargaining games, which are more complex and closer to real-world settings than traditional models. Such settings are difficult to analyse game-theoretically, but can be approached using computational techniques.

In the second part, a number of business applications of automated bargaining are introduced and investigated using computational simulations. The focus here lies on one-to-many bargaining, where for example a seller negotiates with many buyers simultaneously. Either an auction or a bilateral bargaining protocol is applied to the one-to-many setting, depending on the application. Auctions can be an effective way to allocate scarce resources efficiently, or in other words, to ensure that goods are awarded to whoever values them the most. If resources are flexible, however, and negotiation involves multiple aspects, bilateral bargaining can again be the preferred way to reach an agreement. For the first case, we investigate the effectiveness of various auction rules using an evolutionary simulation for problems which are unwieldy to analyse mathematically. For the latter case, we present novel bargaining strategies for the agents that can be used in practical applications. These strategies are able to cope with complex goods and can maximise the gains of trade (i.e., the joint gains that results from an agreement) by adjusting different aspects of the goods to individual needs. We furthermore combine auctions with bilateral bargaining and propose strategies which benefit from the fact that the setting is one-to-many, even though the actual bargaining is bilateral. The performance of the strategies is evaluated using computational simulations.

1.1 Terms and definitions

This Section introduces the general terminology used throughout this thesis. A more detailed explanation of game-theoretic concepts related to bargaining is presented in Chapter 2, particularly Sections 2.1 and Sections 2.3. Furthermore, additional local definitions are provided in the corresponding chapters. Some definitions are numbered in order to facilitate the lookup. Note that the numbers contain the *page number* where the definition is introduced, plus an additional index number.

1.1.1 General economic concepts

In order to analyse the choices that people make, such as in bargaining, it is important to consider the preferences of decision makers for different outcomes. Within economics and in this thesis the notion of *utility* is used to quantify individuals' preferences. *Utility* can be considered as an individual's measure of goal achievement and is usually expressed in real numbers. In general, this measure is subjective and cannot be compared to the utility of other individuals. For many real-world applications, however, utility corresponds to a monetary value, in which case comparison is possible. A *utility function* describes an individual's preferences over possible outcomes in terms of utility.

In many cases, outcomes depend not only on choices made by individuals, but can also be affected by unpredictable events or *lotteries*. When such uncertainty exists, the notion of *expected utility* is used. *Expected utility* specifies the preferences over lotteries, and is computed by multiplying the utility of an event by the probability that this event occurs, and adding across all events (see [72, Ch.6] for further details).

Often, people have several goals and trade-offs between these goals. For example, when buying a house, trade-offs exist between the location, size, and price of the house. A *multi-attribute utility function* [10, 101] can be used in order to represent preferences in case of several (often independent) goals:

Definition 3.1 ***Multi-Attribute Utility Function*** A *multi-attribute utility function* defines the utility over multiple weighted *attributes*, where each attribute corresponds to a goal, and the weight indicates the relative importance of the corresponding attribute. An attribute is also called a *dimension* or an *issue*. In general, the attributes are assumed to be *preferentially independent* or *additive*. In that case, the utility is calculated by multiplying each attribute by its weight and adding across the attributes.

1.1.2 Game-theoretic concepts related to bargaining

Game theory [11, 90][72, Ch.8+9] is a collection of mathematical tools designed to analyse situations where decision-makers interact, for instance when bargaining.

The decision-makers are usually assumed to be *fully rational* (utility maximising) and to be completely informed of the circumstances in which the game is played[1] [11, Ch.10+11]. These assumptions are far from realistic, but are often necessary in order to make mathematical analysis feasible. We will elaborate on these assumptions in Chapter 2 (see Section 2.1).

A decision maker in a game is henceforth called a *player*. We often use the term *agent* instead of player, especially in a computational context.

Game theory is used in this thesis to investigate situations of *bargaining*. In a *bargaining situation* two or more players have the option to make a joint choice from a set of possible outcomes. The players may benefit from an agreement, but they have different preferences for the various outcomes. In economic terms, the players can jointly produce some type of *bargaining surplus*, provided that they agree on how to divide it [81]. Examples include bargaining over the price of a house, but also choosing a restaurant together; in both cases, all parties involved benefit from an agreement, but might have conflicting preferences for the different outcomes. The *bargaining surplus* or just *surplus* is the joint gains that can be achieved through cooperation. For example, if a seller wants to sell a house for at least $100000, and buyer is willing to pay up to $150000, then the bargaining surplus that is jointly produced equals $50000. We define *bargaining* as the corresponding attempt to resolve a bargaining situation, i.e., to determine the particular form of cooperation and the corresponding division of the bargaining surplus. Bargaining is *bilateral* when it concerns two players. We use the term *negotiation* interchangeably with the term bargaining.

The interaction between negotiating agents is usually restricted by certain rules. For instance, in the alternating-offers game (discussed in Section 2.3.2), the players are restricted to making offers and counter offers in a sequential order. The rules are set by the so-called *bargaining protocol*:

Definition 4.1 ***Bargaining Protocol*** A *bargaining protocol* (also called *negotiation protocol*) specifies the rules that govern the negotiation process [5].

The outcomes of a bargaining game have two desirable features: individual rationality and Pareto-efficient [11, Ch. 5]:

Definition 4.2 ***Individually Rational*** A bargaining outcome is *individually rational* [2] if the utility assigned to each player is at least as large as a player can achieve by himself without cooperation.

[1] Complete information does not rule out uncertainty (e.g. about the preferences of other players). In case of uncertainty, however, it is assumed that the *probabilities* are known to the players. This topic is further discussed in Section 2.1 of the next chapter.

[2] Individual rationality is also used to denote a property of a *mechanism* (see Def.5.1). In short, a *mechanism* is individually rational if it induces voluntary participation.

Definition 4.3 ***Pareto-Efficient, Pareto-Efficient Frontier*** A bargaining outcome is *Pareto-efficient* if no outcome exists that is strictly preferred by one player and not less preferred by any other player. The *Pareto-efficient frontier* connects all the Pareto-efficient points in an N-dimensional space, where each dimension corresponds to the utility level of a player (see Fig. 2.1 on page 21 for an example in a 2-dimensional space).

Loosely put, individual rationality of the bargaining outcome ensures that an agent benefits from the agreement. In most cases, a utility of zero is set as the agent's status quo (i.e., the agent's utility for not participating). Any positive outcome is then individually rational. A Pareto-efficient outcome is desirable since there is then no waste in the allocation of the resources [72, p. 313]. If outcomes are *not* Pareto efficient, another deal could have been made which was at least better for one player (and equally good for the other player), or even better for both.

The players are endowed with *strategies* that determine how the bargaining proceeds. In general, a player's *strategy* is a plan which lays out a course of action for each possible state or history [90]. In a bargaining setting, a strategy determines the bids of a player, given the history of the game. Moreover, the strategy decides how the player responds to the bids received by other player(s) in the game. In the alternating-offers game (see Section 2.3.2), for example, a player can respond by accepting or refusing the bid received by the opponent.

Mechanism design An important application area of game theory is setting up the rules of the games, such as voting procedures or auctions rules, as to induce a certain outcome, given that players act rationally and in their own best interest. For example, game theory can help to understand what type of penalties, rewards or tax system are most effective to induce industrial companies to apply environmentally friendly production methods. In the context of bargaining, common goals are maximising social welfare (i.e., the sum of utilities of the players) or maximising revenue. Choosing the right rules in order to achieve desired outcomes is known in economics as the problem of *mechanism design* [133][72, Ch. 23]. First, we define the notion of *mechanism*.

Definition 5.1 ***Mechanism*** A mechanism is a set of decision rules that map the strategies of the agents to a collective outcome.

A mechanism can be viewed as an institution with rules governing the procedure for making the collective choice [72, p. 866]. In a *direct* mechanism, the agents are asked to state their preferences directly (either truthfully or not). An agent's preferences or type is represented by a utility function, expressing the valuation of the possible outcomes or allocations. In an *indirect* mechanism, players do not communicate an entire utility function, but for instance bids in an iterative auction such as the English auction.[3]

Mechanism design deals with the problem of finding a mechanism that results in a desired collective outcome, given that the agents maximise their individual utility, and given that the institution that governs the rules does not know the *preferences* or *types* of the agents beforehand (i.e., we are in a setting characterized by incomplete information, see [72, Ch. 23.B] and Section 2.2). In other words, mechanism design tries to answer whether or not, and if so how, a desired social outcome can be materialised in a world of selfish agents.

A mechanism is called *incentive compatible* if it induces the agents to reveal their preferences truthfully. An interesting theorem is the *revelation principle* [11, Ch. 11][72, Ch. 23], which states that if a desired social outcome can be realised by an indirect mechanism, there exists an incentive compatible direct mechanism that also reaches the desired outcome.

1.1.3 Concepts from computer science

We describe *software agents* [144] in this thesis that fully or partially automate the task of negotiation. We define a *software agent* as an autonomous software program which operates on behalf of its owner. Software agents have a certain goal, which in this thesis is to maximise a given utility function. The software agents described here can usually *learn* from experience and *adapt* their behaviour given feedback from the environment, without any human intervention. When multiple software agents interact, the entire system is called a *multi-agent system.* Note that in a multi-agent system the agents can reside on different platforms, in which case communication occurs via a physical network. We also use the term *evolutionary agent* to denote an agent who's strategy is adapted using an evolutionary algorithm.

1.2 Evolutionary algorithms

Evolutionary algorithms (EAs) are powerful search algorithms from the field of artificial intelligence that are based on the principles of natural evolution [8, 45, 51, 75, 103, 115]. EAs are originally applied to solve optimisation problems, such as the travelling salesman problem and the knapsack problem [29], but are now increasingly being used to model societies of learning agents, especially within the field of agent-based computational economics (ACE) [4, 29, 104, 124, 127, 139]. Throughout this thesis EAs are applied to model adaptive agents that can learn to bargain effectively by means of trial and error. This section first briefly explains the basic

[3]In an English auction players call out increasingly higher bids until no more increases are made. The winner is the last bidder.

principles of EAs. Then it motivates and explains the use of EAs in the context of bargaining. Furthermore, Section 1.2.3 describes in more detail the actual algorithm used in this thesis. The basic approach is the same in all chapters that apply evolutionary algorithms.

1.2.1 Principles of evolutionary algorithms

The cornerstones of evolution in nature are "survival of the fittest" together with the transfer (with some variation) of genetic material from one generation to the next. EAs apply these aspects of biology to evolve an artificial population of individuals. These individuals are not living organisms in this case, but for instance solutions to a optimisation problem or bargaining strategies of an agent. The solutions are encoded on a *chromosome* of an individual, often consisting of a string of real or binary values.

As in natural ecosystems, the survival of these individuals depends on their *fitness*. A suitable fitness measure in artificial ecosystems depends on the problem domain. It can for instance be an objective function in case of an optimisation problem, or the mean utility obtained by a strategy in a game. Using the example of the well-known prisoner's dilemma[4] [90, p.16], an individual's chromosome encodes a player's (binary) strategy: confess or not confess. The fitness is determined by the final payoff (or utility) obtained when the game is played.

By reproduction new individuals are generated that inherit genetic material from the existing individuals in a population. Natural selection then removes individuals with a relatively low fitness from the population. This process of evolution causes good traits (i.e., that contribute to a higher fitness) to remain and bad traits to die out in the long run. Additionally, variation or "errors" in the transfer of genetic material creates new type of individuals or solutions.

1.2.2 Modelling adaptive bargaining agents

Traditional game-theoretic studies of bargaining rely on strong assumptions such as full rationality of the agents and common knowledge of beliefs and preferences (for details see Chapter 2). In reality it is rare that these criteria are met. Even in the case of computational autonomous agents, which are capable of performing calculations much faster than humans, optimal or "rational" solutions cannot always be found. More importantly, since agents can be programmed by different parties, it is better to avoid strict assumptions on other agents' behaviour, in particular concerning their rationality. Rather than fully rational, we assume that bargaining

[4]In this game, two suspects in a crime can choose either to confess or not to confess, without knowing the strategy of the other player. The *payoff* or final utility of a player depends on both his choice and of the choice made by the other player.

agents have little a-prori knowledge and gradually adapt and search for optimal solutions by a process of trial and error. Such agents are called *boundedly rational*.

In this thesis we apply an EA to model this learning aspect of bargaining agents and to develop effective strategies for these agents. EAs are frequently used for modelling (adaptive) behaviour of human societies and societies of computational agents from the bottom up, especially within the field of agent-based computational economics (ACE).[5] EAs are also increasingly being used to study situations of bargaining that are difficult to analyse game-theoretically, as in [31, 34, 73, 88, 126] (see also Section 2.4.1). The advantage of EAs is that they make no explicit assumptions or use of rationality; basically, the *fitness* of the individual agents is used to determine whether a strategy will be used in future situations. Nonetheless, surprisingly rational behaviour often emerges from such "low-rational" agents [146] (as we will also show in this thesis).

There are several ways of modelling adaptive agents using EAs. In the approach used in this thesis, agents select their bargaining strategies from a pool of strategies. A separate pool of strategies exists for each agent *type*, where a *type* is defined by the preferences (i.e., utility function) of the agent and/or the agent's role (e.g. buyer or seller). Agents of the same type select their strategies from the same pool, as these agents are likely to have similar behaviour. On the other hand, agents of different types will usually prefer different strategies, hence the use of separate pools. The pools then evolve independently, i.e. no genetic material is exchanged between the different pools. Note that if there is only a single agent of a certain type, all strategies in a pool belong to that agent. This is also called a model of *individual* learning. If there are several agents of the same type, this is called *population* learning, since a population of agents (of the same type) learns as a whole. Below, the implementation of the EA is explained in more detail.

1.2.3 Implementation

The term "evolutionary algorithm" refers to a broad class of algorithms. The implementation used in this thesis is based on a branch within EAs called *evolution strategies* (ES) [8], originally developed by Rechenberg [103] and Schwefel [115]. The ES were developed independently from the well known genetic algorithms (GAs) [45, 75], introduced by Holland [51]. Whereas GAs are more tailored toward binary-coded search spaces, ES are originally designed for real-encoded representations, the latter being a more natural encoding for the type of bargaining strategies we employ in the simulations. Other classes of evolutionary algorithms include genetic programming, evolution strategies, and evolutionary programming. For an interesting overview of the various approaches within evolutionary computation, see [7].

[5] For an on-line survey of the field of ACE, see [125].

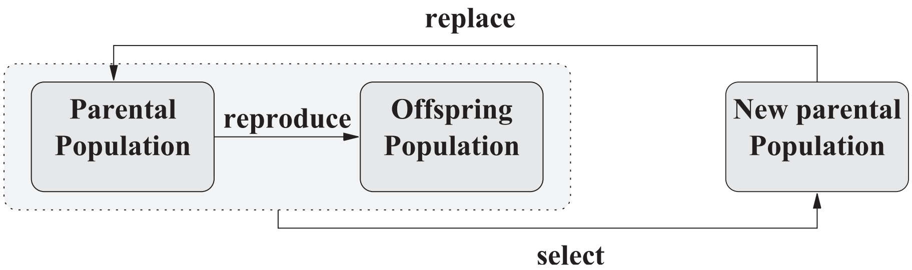

Figure 1.1: Iteration loop of the evolutionary algorithm.

An outline of the EA is given in Figure 1.1. The EA starts with a randomly initialised *parental* population of individuals. Each individual contains a bargaining strategy which is encoded on the chromosome, a fixed-size string $[x_0, ..., x_{l-1}]$ of length l and real values $x_i \in [0, 1]$. Subsequently, *offspring* individuals are created (see Figure 1.1) by first (randomly, with replacement) selecting an agent in the parental population, and then mutating his chromosome to create a new offspring (the mutation operator is described below). Figure 1.2 depicts the chromosomes of a parent individual and a corresponding (mutated) offspring individual. This process is repeated until the offspring population reaches the required size.

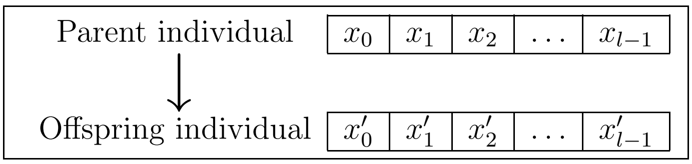

Figure 1.2: The chromosome of a parent individual and of an associated offspring individual. Each chromosome consists of l real values $x_i, x'_i \in [0, 1]$. The offspring individual is created by mutating the chromosome of the selected parent individual.

In the next stage, the *fitness* or performance of both the offspring and parent individuals is determined by a process of negotiation. The way in which this is achieved depends on the negotiation setup. Details are provided in the corresponding chapters.

In the final stage of the iteration (see Fig. 1.1), the fittest agents are selected as the new "parents" for the next iteration. Selection is performed using the deterministic $(\mu + \lambda)$-ES selection scheme [7, 8], where μ is the number of parents and λ is the number of generated offspring. The μ survivors with the highest fitness are selected (deterministically) from the union of parental and offspring agents. This final step completes one iteration or *generation* of the EA.

Mutation and Recombination

Mutation and recombination are the most commonly used EA operators for reproduction. Recombination exchanges parts of the parental chromosomes, whereas mutation produces random changes in a chromosome. In case of an ES, it is common to use mutation-based models without recombination, especially because the mutation operator (explained below) is much more advanced compared to the standard operator used in e.g. genetic algorithms. Moreover, for many computational experiments of the kind discussed in this thesis, the effects of recombination seemed to be negligible when using an ES (see also [126]). We therefore focus on mutation-based models in this thesis.

The mutation operator of an ES implementation works as follows. Each real value x_i of a parent chromosome (see Figure 1.2) is mutated by adding a zero-mean Gaussian variable with a standard deviation σ_i [8, 126], thereby producing a new value x'_i for the chromosome of the offspring:

$$x'_i := x_i + \sigma_i N_i(0,1). \tag{1.1}$$

All resulting values larger than unity (or smaller than zero) are set to unity (respectively zero).

In our simulations, we use two mutation models: a mutation model with self-adaptive control of the standard deviations σ_i [8, pp. 71-73][126], and a model with exponential decay of the standard deviations, which we describe below.

Self-Adaptive Control This model allows the evolution of both the strategy and the corresponding standard deviations at the same time. More formally, an agent consists of strategy variables $[x_0, ..., x_{l-1}]$ and ES-parameters $[\sigma_0, ..., \sigma_{l-1}]$, where l is the length of the chromosome.

The mutation operator first updates an agent's ES-parameters σ_i in the following way:

$$\sigma_i := \sigma_i \exp[\tau' N(0,1) + \tau N_i(0,1)], \tag{1.2}$$

where τ' and τ are the so-called learning rates [8, p. 72], and $N(0,1)$ denotes a normally distributed random variable having expectation zero and standard deviation one. The index i in N_i indicates that the variable is sampled anew for each value of i. We use commonly recommended settings for these parameters (see [8, p. 72]).[6] After the strategy parameters have been modified, the strategy variables are mutated as indicated in Eq. 1.1.

Note that, since selection works on the σ_i's as well as on the strategy variables, the σ_i's are part of the evolutionary process. The particular initial value chosen for σ_i is therefore typically not crucial for this model, as the self-adaptation process rapidly

[6] Namely, $\tau' = (\sqrt{2l})^{-1}$ and $\tau = (\sqrt{2\sqrt{l}})^{-1}$, where l is the length of the chromosome.

scales σ_i into the proper range. For example, if solutions are far from the optimal value, the σ_i can increase as a result of the evolutionary process. On the other hand, if good solutions are found, the σ_i's can converge to smaller values in order to maintain these solutions. To prevent complete convergence of the population, we force all standard deviations to remain larger than a small value ε_σ [8, pp. 72–73].

Exponential Decay Using this model, the standard deviations σ_i decay exponentially such that every t generations their value is reduced to half the size. We call t the half-life parameter. This model is similar to the simulated annealing mechanism, where a temperature parameter is slowly lowered to reduce variation in the exploration space. Using this model, the EA always converges if the simulation is run for a sufficient number of generations.

1.3 Organisation of the thesis

Readers that are new to the field of game theory and bargaining are recommended to read the introduction to this topic in Chapter 2. Specific topics include the ultimatum game, the alternating-offers game, bargaining with incomplete information, multi-issue bargaining, and one-to-many bargaining. Chapter 2 also contains a survey of approaches using techniques from artificial intelligence and are in that way related to the general topic of the thesis. Chapter 8 concludes the thesis with a discussion and an overview of the the main results.

The remaining chapters of the thesis are grouped into two parts: **Part A** considers fundamental aspects of bilateral bargaining systems using both game-theoretical and computational techniques. **Part B** investigates two business applications of automated bargaining, and introduces a number of effective bargaining strategies. Additionally, in the Appendix a game-theoretic analysis is provided for the games described in Chapter 3. Each chapter of parts A and B can, in principle, be read independently. Where necessary, cross-references are indicated within the chapters. A recurring theme is the application of evolutionary algorithms for simulating the strategic behaviour of the agents. The evolutionary algorithm is therefore treated separately in Section 1.2. Parts A and B are organised as follows:

Part A: Fundamental aspects of bargaining systems

Chapter 3 describes a system for bilateral negotiations in which artificial agents are generated by an evolutionary algorithm. The negotiations are governed by a finite-horizon version of the alternating-offers protocol. Several issues are negotiated simultaneously. This can reduce the competitive nature of the game since trade-offs can be made to obtain mutually beneficial solutions. These so-called Pareto-efficient solutions are indeed found by the evolutionary agents. The outcomes of the evolutionary system are also analysed and validated using the game-theoretic subgame-perfect equilibrium as a benchmark. We furthermore present and investigate an extended model in which the agents take into account the fairness of the obtained payoff. The concept of fairness plays an important role in real-life negotiations and experimental economics. We find that when the fairness norm is consistently applied during the negotiation, the evolving agents reach symmetric outcomes which are robust and rather insensitive to the actual fairness settings.

Chapter 4 extends the above game by allowing both agents to negotiate with other opponents in case of a disagreement. This way the basics of a competitive market are modelled where for instance a buyer can try several sellers before making a purchase decision. Negotiations are limited to a single round, which corresponds to the so-called ultimatum game. Whereas in the regular ultimatum game the proposer demands the entire surplus, responding agents can now choose to refuse unacceptable take-it-or-leave-it deals and negotiate with another opponent. As before, the game is investigated using an evolutionary simulation. The outcomes appear to depend largely on the information available to the agents. We find that if the agents' number of future bargaining opportunities is commonly known, the proposer has the advantage. If this information is held private, however, the responder can obtain a larger share of the pie, even if the initial number of bargaining opportunities is equal for both agents. For the first case, a game-theoretic analysis of the game is also presented and compared to the evolutionary results. Although a theoretical analysis is hard for the incomplete information case, the evolutionary simulation is very suitable for analysing both settings. The game is further extended to allow several issues to be negotiated simultaneously. Furthermore, effects of search costs are investigated and the case where uncertainty exists about future opportunities and a new opponent cannot always be found.

Part B: Bargaining systems for business applications

Chapter 5 considers a business application of automated negotiation, where several supplier agents of goods and services compete for banner space or "consumer attention space" by bidding in an auction. Bidding occurs based on information about the consumers, their so-called profile. As a result of the auction, a small selection of banners is short-listed and presented to the consumer, for instance on a web site. The supplier agents are simulated using an evolutionary algorithm, and can learn, given feedback from the consumers and whether or not they were short-listed, the type of consumers to target and the amount to bid. A number of consumer behaviour models are investigated that simulate the consumer's response to the presented banners. In a relatively simple model, the response is independent of other banners displayed concurrently. In other models, the response contains dependencies between the banners. The auctioneer can select the auction rules or *mechanism* that generates the best advertisements for the consumers, but at the same time provides the suppliers with sufficient profits. Several mechanisms are investigated using the simulation environment.

Chapter 6 applies automated negotiation to buy and sell bundles of information goods. A single information provider agent or seller agent negotiates with a number of buyer agents simultaneously. Whereas in Chapter 5 an auction is used for a one-to-many setting, a bilateral negotiation protocol is applied in this case, where the seller negotiates with each buyer by alternating offers and counter offers, as described in Chapter 3. A bilateral protocol is more suitable here because information goods have no constraints on the supply and different buyers can be interested in very diverse bundles of goods. A personalisation of bundles is achieved by bargaining over multiple issues. Bargaining in this setting essentially has a double purpose: (1) division of the surplus, and (2) maximising the joint gains that can be achieved by finding win-win or Pareto-efficient (see Def. 4.3) outcomes. This chapter focuses on the latter part and introduces negotiation strategies for multi-issue negotiations which can approximate Pareto-efficient solutions.

Chapter 7 also considers the one-to-many bargaining setting using a bilateral bargaining protocol, but focuses on the division of the surplus. Although the buyers perceive bargaining as bilateral, the seller can actually benefit from the fact that bargaining occurs with many buyers simultaneously. This is especially the case if buyers have time pressure and prefer early agreements. Several bargaining strategies for the seller are investigated and compared using an evolutionary simulation. A class of strategies are introduced which are based on the first-price auction. These strategies can especially benefit from competition arising from the time pressure. The seller's bargaining strategies also take into account a notion of fairness, which should ensure that buyers are treated fairly and do not feel discriminated based on their individual bargaining behaviour or preferences.

1.3.1 Publications

Chapters 3-6 are based on published work and/or work that has been accepted for publication but has yet to appear. **Chapters 2 and 7** are based on technical reports.

- **Chapter 2** is based on [41]: E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Scientific approaches and techniques for negotiation: A game theoretic and artificial intelligence perspective. Technical Report SEN-R0005, CWI, Amsterdam, 2000.

- **Chapter 3** is based on [42]: E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. *Computational Economics*, 22:39–63, 2003.
- **Chapter 4** is based on [38]: E.H. Gerding and J.A. La Poutré. Bargaining with posterior opportunities: An evolutionary social simulation. In M. Gallegati, A. Kirman, and M. Marsili, editors, *The Complex Dynamics of Economic Interaction*, Springer Lecture Notes in Economics and Mathematical Systems (LNEMS), Vol. 531, pages 241–256. Springer-Verlag, 2004.
- **Chapter 5** is based on [17]: S.M. Bohte, E.H. Gerding, and J.A. La Poutré. Market-based recommendation: Agents that compete for consumer attention. *ACM Transactions on Internet Technology*, August 2004 (to appear). A shorter version appeared earlier as [16]: S. M. Bohte, E. H. Gerding, and H. La Poutré. Competitive market-based allocation of consumer attention space. In M. Wellman, editor, *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC-01)*, pages 202–206. The ACM Press, 2001.
- **Chapter 6** is based on [120]: K. Somefun, E.H. Gerding, S. Bohte, and J.A. La Poutré. Automated negotiation and bundling of information goods. In *Agent-Mediated Electronic Commerce V*, Springer Lecture Notes in Artificial Intelligence (LNAI). Springer-Verlag, Berlin, to appear.
- **Chapter 7** is based on [40]: E.H. Gerding, K. Somefun, and J.A. La Poutré. Bilateral bargaining in a one-to-many bargaining setting. Technical Report, CWI, Amsterdam, to appear. A shorter version has been accepted for publication as [39]: E.H. Gerding, K. Somefun, and J.A. La Poutré. Bilateral bargaining in a one-to-many bargaining setting. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS2004), New York City, New York*. IEEE Computer Society Press, 2004.

Chapter 2

Bargaining: an overview

This chapter contains an overview of approaches and techniques concerned with bargaining. We here focus on the large body of literature that has been published in the fields of economics (in particular game theory) and artificial intelligence (AI). To give a brief impression of the rapid developments in this field, we first highlight some important breakthroughs in economic bargaining theory in Section 2.1. Section 2.2 discusses assumptions frequently made in game theory to make mathematical analysis feasible, and motivates the use of computational techniques. Details on game-theoretic bargaining approaches follow in Section 2.3. Bargaining approaches using computational techniques from the field of artificial intelligence are the topic of Section 2.4. Finally, Section 2.5 concludes this chapter with a short discussion.

2.1 A brief history of bargaining

Perhaps surprisingly, the bargaining problem has challenged economists for decades. Yet the bargaining problem is stated very easily [110]:

> Two individuals have before them several possible contractual agreements. Both have interests in reaching agreement but their interests are not entirely identical. What "will be" the agreed contract, assuming that both parties behave rationally?

Before the path-breaking work of Nash [82] and, much later, Rubinstein [110] the bargaining problem was considered to be indeterminate. For example, in their influential work Von Neumann and Morgenstern [137] argued that the most one can say is that the agreed contract will lie in the so-called *bargaining set*. The *bargaining set* is the set of all feasible outcomes (an outcome is feasible if it can be jointly achieved by the players involved) that are individually rational (see Def. 4.2) and Pareto-efficient (see Def. 4.3), i.e., it is no worse than disagreement and there is no agreement that both parties would prefer. But because this bargaining set consists in general of an infinite number of different agreements this requirement does not yield a unique bargaining solution. A unique solution can be found, however, if the agreed contract satisfies additional axioms such as those proposed by Nash [82]. This solution is called the Nash bargaining solution and is discussed in Section 2.3.1. Because one can argue about which axioms are *reasonable* and which are not, Nash suggested to complement this axiomatic approach with a strategic game. This route was followed by Rubinstein [110] who proved that an important bargaining game (the *alternating-offers* game) has a unique solution (see Section 2.3.2). Binmore [12] then connected the fields of axiomatic and strategic bargaining by proving that the solution of Rubinstein's bargaining model coincides with the Nash bargaining solution under special circumstances.

2.2 Game theory and artificial intelligence

Game theory frequently makes simplifying assumptions to facilitate the mathematical analysis. Common assumptions are for instance: (1) complete knowledge of the circumstances in which the game is played and (2) full rationality of the players. The first assumption implies that the rules of the game and the preferences (i.e., the utility functions) and *beliefs*[1] of the players are *common knowledge*.[2] A game has *incomplete information* if something about the circumstances in which the game is played, such as the preferences of other players, is not known to the players. Game theorists traditionally model incomplete information of other player's preferences and beliefs by specifying a limited number of player *types* (see also Section 2.4.3). Each type is then uniquely determined by a set of preferences and beliefs. Players are not completely certain about the exact type of their opponent. However, the *probability* that an opponent is of a certain type is, again, common knowledge for all players. In this manner, a game of incomplete information can be transformed in a game of *imperfect* information.[3]

The second assumption relates to the need for common knowledge on how players reason. It is assumed that players maximise their expected utility given their beliefs. Players have infinite computational capacity to pursue statements like "if I think that he thinks that I think..." ad infinitum. Furthermore, players are assumed to have a perfect memory.[4] These assumptions limit the practical applicability of game-

[1]Beliefs are subjective probability of events occurring about which the player is uncertain.

[2]Common knowledge means that the players know what the other players know, etc., in an infinite regress.

[3]In a game with imperfect information uncertainty exists about the state of the world. A game is said to have *perfect* information if (i) there are no simultaneous moves and (ii) at each decision point it is known which choices have previously been made [131, Ch. 1].

[4]Lately, much research in game theory focuses on the field of "bounded" rationality, in which players have limited computational power and/or limited hindsight. An overview of recent work

theoretic results. In the field of AI, however, assumptions like complete knowledge or full rationality are not necessary because the behaviour of individual agents can be modelled directly.[5] This gives the AI approach an important advantage over more rigorous (but at the same time more simplified) game-theoretical models.

Researchers in the field of AI are currently developing software agents (see Section 1.1.3) which should be able (in the near future) to negotiate in an intelligent way on behalf of their users. A survey of the potential of automated negotiation is given in [144, Ch. 9]. The state-of-the-art of agent technology is reviewed in [70]. In future applications for e-commerce, multi-agent systems will need to be flexible, especially for trading, brokering, and profiling applications [128]. In particular, it is important for the negotiating (software) agents to be able to adapt their strategies to deal with changing opponents, changing topics and concerns, and changing user preferences. Multi-agent learning, (the ability of the agents to learn how to communicate, cooperate and compete) becomes crucial in such domains [70, p.23]. This should lead to much more advanced and universal systems.

Nevertheless, due to this rapidly increasing complexity, the connection between the AI approach and a game-theoretic analysis remains important. Game theory may aid in the difficult task of choosing a suitable bargaining protocol [14] (see Def. 4.1). Tools and techniques from AI can be used to develop software applications, bargaining strategies, protocols and mechanisms which are currently beyond the reach of classical game theory.

2.3 Game-theoretic approaches to bargaining

Traditionally, game theory can be divided into two branches: cooperative and non-cooperative game theory. In *cooperative* game theory, groups of players are taken as primitives and binding agreements can be made. Cooperative game theory abstracts away from the rules of the game and is mainly concerned with finding a solution given a set of feasible outcomes.[6] A topic like coalition forming is typically analysed using cooperative game theory. Often, in real life, companies can gain profits by working together, for example by securing a larger market share or by reducing direct competition with the competitors. In such games, a surplus (see Section 1.1.2) is created when two or more players cooperate and form a coalition. Cooperative

in this field can be found in [112]. Binmore also gives a short discussion of this topic in [11, pp. 478-488].

[5] For example, agents can be programmed with a certain strategy and use for instance reinforcement learning to improve this strategy. These agents are not explicitly rational or fully informed. Nevertheless, after a period of learning, the agents could exhibit behaviour that resembles that of rational and fully informed agents.

[6] Recall from above that an outcome is feasible if it can be jointly achieved by the players involved.

game theory can then determine how the surplus is to be divided, given a coalition and a set of assumptions (called *axioms*). Likewise, cooperative bargaining theory determines how the surplus is to be divided which results from an agreement, given the set of axioms (an example of such axioms resulting in a unique solution, the so-called Nash bargaining solution, is discussed in Section 2.3.1).

Non-cooperative game theory, on the other hand, is concerned with specific games with a well defined set of rules, game *strategies*, and *payoffs* rather than axioms. All strategies, rules and payoffs are known beforehand by the players. A player's *strategy* is a plan which lays out a course of action for each possible state or history. Strategies can be *pure* or *mixed*. A *pure* strategy determines the actions for a given state deterministically. A *mixed* strategy requires a player to randomise between his pure strategies. *Payoffs* are the final returns (expressed in utility) to the players when the game is concluded.

Non-cooperative game theory uses the notion of a *strategic equilibrium* or just *equilibrium* to determine rational outcomes of a game. Numerous equilibrium concepts have been proposed in the literature (see [131] for an overview). Some widely-used concepts are *dominant* strategies, *Nash* equilibrium and *subgame perfect* equilibrium. We define these notions below.

Definition 18.1 ***Dominant Strategy*** A dominant strategy is optimal in all circumstances, that is, the strategy achieves the highest payoff no matter what the strategies of the other players are.

This is obviously a very strong notion of an equilibrium strategy. A slightly weaker, but still very powerful, equilibrium concept is the so-called Nash equilibrium [83, 84]:

Definition 18.2 ***Nash Equilibrium*** Strategies chosen by all players are said to be in *Nash equilibrium* if no player can benefit by *unilaterally* changing his strategy.

Nash proved that every finite game has at least one equilibrium point (in pure or mixed strategies [83, 84]). The concept of dominant strategies is a *refinement* of the Nash equilibrium. That is, if strategies are dominant, they also constitute a Nash equilibrium. The reverse is not necessarily true, however. Another important refinement of a Nash equilibrium is Selten's subgame-perfect equilibrium [116, 117] for extensive-form games. Extensive-form games are games with a tree structure, i.e., where players can make decisions sequentially and at various stages of the game (by contrast, in *strategic-form* games, players are required to make decisions once and simultaneously). Subgame-perfect equilibrium is defined as follows:

Definition 18.3 ***Subgame-Perfect Equilibrium*** Strategies in an extensive-form game are in subgame-perfect equilibrium if the strategies constitute a a Nash equilibrium at every decision point.

An overview of the main bargaining literature from the field of cooperative game theory is given in Section 2.3.1. We note that the concepts from cooperative game theory are not necessary to understand the remainder of the thesis, and are intended for the interested reader. In Section 2.3.2 several non-cooperative bargaining games are discussed. Particular attention is paid to the most important bargaining protocol: the *alternating-offers* game. In Section 2.3.2 bargaining over a single issue is assumed. Section 2.3.3 covers work on multiple-issue negotiations.

As we mentioned before, traditional game theory assumes complete information, implying that the player's preferences and beliefs are common knowledge. However, lately many researchers in game theory have focused on the consequences of players having private information. Among other things, incomplete information could explain why inefficient deals are reached or why no deal is reached at all. For instance, the occurrence of strikes and bargaining impasses, but also the occurrence of delays in negotiations can theoretically be addressed when complete information is no longer assumed. Literature related to this topic is discussed in Section 2.3.4. We also consider one-to-many bargaining, i.e., where one player interacts with multiple opponents simultaneously. Auctions are the most common approach for such a setting, and will be the topic of Section 2.3.5 (an alternative approach, using bilateral bargaining, is discussed in Chapters 6 and 7).

2.3.1 Cooperative bargaining theory

Cooperative game theory considers the space of possible outcomes of a game, without specifying the game itself in detail. In case of bargaining, the outcomes are often denoted in terms of *utilities* (see Section 1.1.1). In case of two-player games, the outcomes are then represented by utility pairs. Cooperative bargaining theory is concerned with the question of which outcome will eventually prevail, given the set of all possible utility pairs. A particular set of possible outcomes is also referred to as a *bargaining problem*.

A function which maps a bargaining problem to a single outcome is called a *solution concept*. Usually, a solution concept is only valid for a certain subset of all possible bargaining problems. For instance, the first and most famous solution concept, the Nash bargaining solution (see below) only applies to convex and compact bargaining sets (see [11, pp. 180–181]). Only if these requirements are satisfied the bargaining problem can properly be called a Nash bargaining problem.

An alternative bargaining solution has been proposed by Kalai and Smorodinsky [57]. Their approach is discussed below. Both the Nash and the Kalai and Smorodinsky bargaining solutions are invariant with respect to the calibration of the players' utility scales. The *utilitarian* solution concept differs in that respect and does actually depend on how the functions are scaled. For this reason, its application is limited to those situations where inter-personal utility comparison makes any sense. Cooperative theories of bargaining are discussed in more detail in [106].

The Nash bargaining solution

Nash proposed four properties, now called the *Nash axioms*, which should be satisfied by rational bargainers [82],[11, p. 184]:

1. The final outcome should not depend on how the players' utility scales are calibrated. This means the following. A utility function specifies a player's preferences. However, different utility functions can be used to model the same preferences. Specifically, any strictly increasing affine transformation of a utility function models the same preferences as the original function, and should therefore yield the same outcome.

2. The agreed payoff pair should always be individually rational (see Def. 4.2) and Pareto-efficient (see Def. 4.3)

3. The outcome should be independent of irrelevant alternatives. Stated otherwise, if the players sometimes agree on the utility pair s when t is also a feasible agreement, they never agree on t when s is a feasible agreement.

4. In symmetric situations, both players get the same.

The solution which satisfies these four properties is characterised by the payoff pair $s = (x_1, x_2)$ which maximises the so-called *Nash product* $(x_1 - d_1)(x_2 - d_2)$, where d_1 and d_2 are player 1's and player 2's outcomes in case of a disagreement. Nash proved that this is the *only* solution which satisfies all four axioms [82]. Given a Nash bargaining problem where the set of individually rational agreements is not empty, the Nash bargaining solution then leads to a unique outcome. Figure 2.1 illustrates how to construct the Nash bargaining solution for a given bargaining problem.

Due to the fourth axiom, both players are treated symmetrically if the bargaining problem is symmetric as well. In other words, if the players' labels are reversed, each one will still receive the same payoff. A more general solution attributes so-called *bargaining powers* α and β to player 1 and player 2, respectively. In this generalised or asymmetric Nash bargaining solution, the fourth axiom is abandoned and the bargaining solution comes to depend on the bargaining powers of the two players.[7] The generalised Nash bargaining solution corresponding to the bargaining powers α and β can be characterised as above as the pair s which maximises the product $(x_1 - d_1)^\alpha (x_2 - d_2)^\beta$ [11, p. 189].

[7] What these bargaining powers represent depends on the actual (non-cooperative) game played. For example, in case of negotiating companies the bargaining powers could be determined by the strength of their respective market positions. It should be clear however, that the bargaining powers have nothing to do with the bargaining skills of the players, since perfect rationality is assumed.

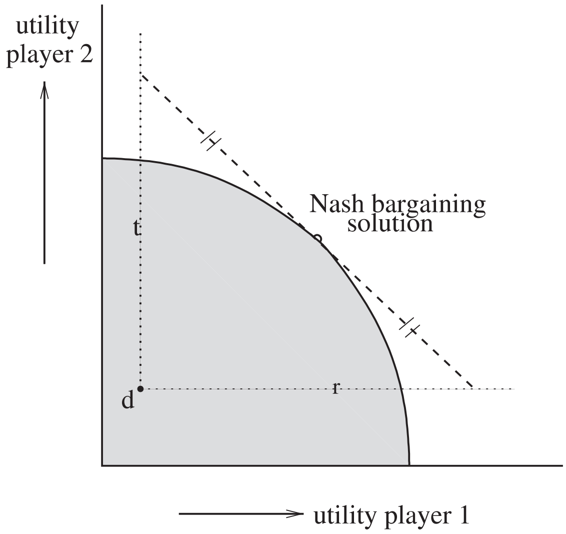

Figure 2.1: Construction of the Nash bargaining solution. This figure shows the Pareto-efficient frontier (denoted by the solid line, see also Def. 4.3) and the Nash bargaining solution for a specific bargaining problem. The bargaining problem is defined by the set of feasible utility pairs (denoted by the grey area) and the disagreement point d which specifies the players' payoffs in case of a disagreement. To find the (symmetric) Nash bargaining solution, one needs to find a supporting line on the Pareto-efficient frontier which is bounded by lines r and t such that the Nash bargaining solution is exactly halfway between these lines. The lines r and t are respectively the horizontal and the vertical lines drawn from the disagreement point d.

The Kalai-Smorodinsky bargaining solution

The third of the Nash axioms (independence of irrelevant alternatives) has been the source of great controversy (follow the discussion in [69]). Kalai and Smorodinsky therefore proposed an alternative to this axiom, which they refer to as the *axiom of monotonicity* [57][72, p. 844]. For a set S of individually-rational and Pareto-efficient points, let $m_i(S) = max\{s_i \mid s \in S\}$ be the maximum utility value that player i could attain (for $i = 1, 2$), given that the players are individually rational. The Kalai-Smorodinsky solution then selects the maximum element in S on the line that joins the disagreement point (d_1, d_2) with the point $(m_1(S), m_2(S))$. An example is given in figure 2.2.

Utilitarianism

A utilitarian policy in philosophy is one which prefers an outcome which maximises the total welfare of the individuals in a society [80]. Any bargaining solution which

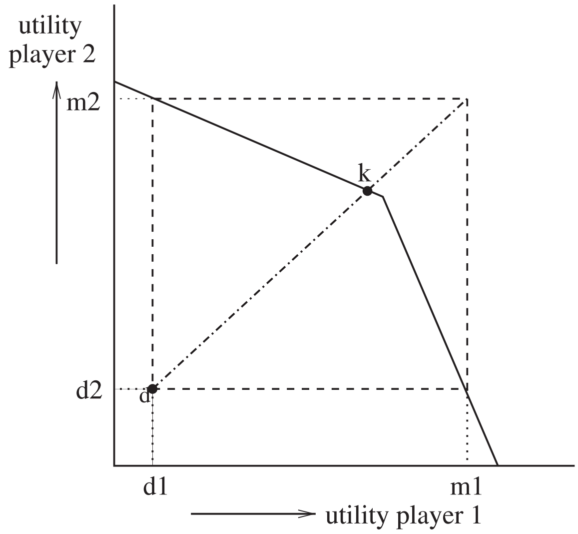

Figure 2.2: Construction of the Kalai-Smorodinsky solution. m_1 and m_2 are the maximum utilities for players 1 and 2 respectively, given that the players are individually rational. Point k is the unique solution which satisfies the four axioms proposed by Kalai and Smorodinsky [57].

maximises the sum of utilities is therefore called a utilitarian solution concept. Stated less formally, the utilitarian principle asserts that "you should do something for me if it will hurt you less than it will help me". Clearly, a utilitarian solution concept assumes that interpersonal utility comparisons are possible. Therefore, Nash's first axiom (independence of utility calibrations) no longer holds in utilitarian models.[8]

Concluding remarks

Apparently, many different types of solutions to the bargaining problem exist in cooperative game theory. The choice of a specific solution is of course based on norms existing in a society, or, more specifically, on which axioms seem to be "reasonable" in a specific bargaining context. Certain outcomes might be for instance be considered as "unfair". An example is given in [101, pp. 235–250].

Additionally, it is important to consider for which classes of non-cooperative games the solution concepts from cooperative game theory are appropriate. For instance, if no non-cooperative game can be found which results in a solution specified by cooperative game theory, then the results from cooperative game theory have little bearing. Fortunately, such a connection between cooperative and non-

[8]Note that the Pareto-efficiency axiom still holds. The other axioms depend on the specific solution concept.

cooperative game theory has been observed under special circumstances [12]. More details are given in the next section.

2.3.2 Bargaining over a single issue

Four different negotiation games or *protocols* (see Def. 4.1) are described in this section. These protocols can be used by two bargainers to divide a given *bargaining surplus* (see Section 1.1.2), that is, the mutual benefit resulting when the players reach an agreement. Without loss of generality, we assume that the bargaining surplus is of size unity in the following.

The following protocols are considered below: (1) the Nash demand game, (2) the ultimatum game, (3) the alternating-offers game and (4) the monotonic concession protocol. The first three games are well-known and widely-used. The fourth game is described in [105] and is an attempt to model a more realistic negotiation scenario. However, in all games described here analytical solutions are obtained using the strong assumption of common knowledge. The extrapolation of results obtained here to real-world cases is therefore a non-trivial step.

The protocols described in this section have been applied mainly to evaluate negotiations over a single issue. In real life, this issue is often the price of a good to be negotiated. Although this keeps matters simple, important value-added services such as delivery time, warranty or service are left out. Both the supplier and the consumer could for instance benefit if negotiations involve multiple issues. Moreover, multiple-issue negotiations can be less competitive because solutions can be sought which satisfy both parties. Multiple-issue negotiations are studied in more detail in Section 2.3.3.

The Nash demand game

Both players simultaneously demand a certain fraction of the bargaining surplus in this game, without any knowledge of the other player's demand [11, pp. 299-304]. In case the sum of demands exceeds the surplus, both players only receive their disagreement payoff. Otherwise, the demands are said to be compatible, and both players get what they requested. This game has an *infinite* number of Nash equilibria: all deals which are Pareto-efficient, but also deals where both players receive their disagreement payoff. For example, if both players ask more than the entire surplus, no player could ever gain by unilaterally changing his strategy.

The concept of a Nash equilibrium thus places few restrictions on the nature of the outcome. Nash therefore suggested a refinement for this game which does result in a *unique* solution. This refinement of the demand game is called the *perturbed* demand game [89, pp. 77-81]. In this perturbed game the players are not completely certain about which outcomes are within the bargaining set (i.e., the set of compatible demands) and which outcomes are not. When the degree of uncertainty approaches zero, the Nash equilibrium of the perturbed game approaches the Nash bargaining solution of the regular demand game (without uncertainty).[9] The reader is referred to [131] for technical details on this subject. A more introductory overview is given by Binmore [11].

The ultimatum game

Playing Nash's demand game, both players could easily receive nothing, or it could occur that some of the surplus is "thrown away". Players would do better by choosing a somewhat less competitive game. If they are unable to reach an agreement using this alternative game, the demand game still remains an option.

A very simple alternative is the so-called *ultimatum* game. In this game, one of the players proposes a split of the surplus and the other player has only two options: accept or refuse. In case of a refusal, both players get nothing (or the demand game is played). Although the game again has an infinite number of Nash equilibria, it has only one subgame perfect equilibrium (in case the bargaining surplus can be divided with arbitrary precision) where the first player demands the whole surplus and the second player accepts this deal [11, pp. 197-200].

The alternating-offers game

Basically a multiple-stage extension of the ultimatum game, the *alternating-offers* game is probably the most elegant bargaining model. As in the ultimatum game, player 1 starts by offering a fraction x of the surplus to player 2. If player 2 accepts player 1's offer, he receives x and player 1 receives $1 - x$. Otherwise, player 2 needs to make a counter offer in the next round, which player 1 then accepts or rejects (sending the game to the next round). This process is repeated until one of the players agrees or until a finite deadline is reached.

Bargaining over a single issue in an alternating fashion has been pioneered by Ingolf Ståhl [121]. A taxonomy and survey of economic literature on bargaining before 1972 is given in this reference. Ståhl analyzes bargaining games with a finite number of alternatives. Both games of finite and of infinite length are studied, but he primarily evaluates games of a finite length. Ståhl uses an assumption of "good-faith bargaining" to simplify the theoretical analysis. Good-faith bargaining prevents players from increasing their demands during play. He then identifies optimal strategies for rational players with perfect information by starting at the last stage of the game and then inductively working backwards until the beginning of play. This procedure yields those equilibria which can be found with dynamic programming methods.

[9] Note that only the Nash equilibria which result in solutions within the bargaining set are considered. Nash equilibria in which no agreement is reached still remain [89, p.79].

A straightforward dynamic programming approach can fail in case of imperfect information [131, Ch. 1]. Sensible strategies can then be found by requiring that each player's optimal strategy for the entire game also prescribes an optimal strategy in every subgame. As mentioned before, this concept of a subgame-perfect equilibrium (SPE, see Def. 18.3) is due to Selten [116, 117]. Rubinstein [110] successfully applied this equilibrium concept to identify a unique solution in his variant of the alternating-offers game. Rubinstein's game [110] has an infinite length and there is a continuum of alternatives. To simplify the analysis, Rubinstein made several assumptions with regard to the players' preferences. An important difference with Ståhl's model is that time preferences are assumed to be stationary (this means that the preferences of getting a part x of the surplus at time t over getting y at $t+1$ is independent of t).

Rubinstein analyses two specific stationary models: one in which each player has a fixed bargaining cost for each period (c_1 and c_2) and one in which each player has fixed *discount factors* (δ_1 and δ_2). Discount factors are used to relate the utility of future consumption to the utility of consuming immediately. In other words, discount factors model how impatient the player is [11, p. 202]. We provide a formal definition of a *discount factor*:

Definition 25.1 ***Discount Factor*** The *discount factor* is used to translate expected utility or costs in any given future into present value terms.

Player i's utility for getting a fraction x of the surplus at time t is equal to $x(\delta_i)^t$. If the discount factor is smaller than 1, a deal is therefore worth less if the agreement is reached in the future than if a deal is reached immediately.

Using stationarity and other assumptions, Rubinstein first demonstrated that the Nash equilibrium concept is too weak to identify a unique solution by proving that every partitioning of the surplus can be supported as the outcome of Nash equilibrium play. To overcome this difficulty, Rubinstein then applied the concept of a SPE and proved that there exists a *unique* SPE in the alternating-offers bargaining model. For example, if both players have a fixed discounting factor (δ_1 and δ_2) the only SPE is one in which player 1 gets $(1-\delta_2)/(1-\delta_1\delta_2)$ and player 2 the remainder (of a surplus of size 1). Furthermore, if both players use their SPE strategy, agreement will be reached in the first round of the game. Notice that Rubinstein's proof assumes that both players have perfect information about the other player's preferences (i.e., their bargaining cost or discount factor). Bargaining with imperfect information (i.e., where uncertainty plays a crucial role) is discussed further in Section 2.3.4.

Rubinstein's paper has been very influential in bargaining theory. At the moment, a vast body of literature exists on infinite-horizon games. An overview is given in [79, 89]. Many pointers to the literature are given in these references. We will conclude this section by discussing a few key papers in this field.

An particularly important paper is [12]. In this paper a relation between the SPE outcome of the alternating-offers game and the Nash bargaining solution is identified in case of weak player preferences (e.g., discount factors close to unity or small time intervals between rounds). This establishes a link between non-cooperative and cooperative bargaining theory and justifies the use of the Nash bargaining solution to resolve negotiation problems (at least in case of complete information).

Van Damme et al. [132] have investigated the role of a smallest monetary unit (i.e., a finite number of alternatives) in the alternating-offers game with payoff discounting. They show that in case of a finite number of alternatives, any partition of the surplus can be supported as the result of a subgame-perfect equilibrium if the time interval between successive rounds becomes very small. This means that Rubinstein's assumption of a continuous spectrum of bids is essential in deriving a unique solution of the alternating-offers game under these conditions.

Monotonic concession protocol

A more restricted protocol, compared to the alternating-offers game, is described in [105]. In this *monotonic concession* protocol the two players announce their proposals simultaneously. If the offers of both agents match or exceed the other agent's demand, an agreement is reached. A coin is tossed to choose one of the offers in case they are dissimilar.

If no agreement is reached, the players need to make new offers in the next round. The offers need to be monotonic, that is, the players are not allowed to make offers which have a lower utility for their counter player compared to the last offer. Hence, a player can either make the same offer (to stand firm) or concede. Negotiations end if both agents stand firm in the same round. The players receive their disagreement payoffs in this case. Because each round at least one of the players has to make a concession (or a disagreement occurs), the protocol has a finite execution time if the minimum concession per round is fixed and larger than zero.

Note that in order to make a (monotonic) concession possible, a player needs to have some knowledge about the other players' preferences. This knowledge is crucial when several issues are negotiated at the same time. In this case not only the sign of the utility function, but also the relative importance of the issues becomes important.

Rosenschein and Zlotkin discuss which kinds of strategies are *stable* and *efficient* when using this protocol (in negotiations over a single issue). A strategy pair is *efficient* in this case if an agreement is always reached. *Stability* is defined using the notion of symmetric Nash equilibrium: A strategy s constitutes a symmetric Nash equilibrium (and is stable) if player 1 can do no better than playing s, given that player 2 also uses s. Note that a strategy s in which both players make a concession in the same round is not stable: one of the players could do better by standing firm. On the other hand, a strategy where a player tosses a coin to determine whether to

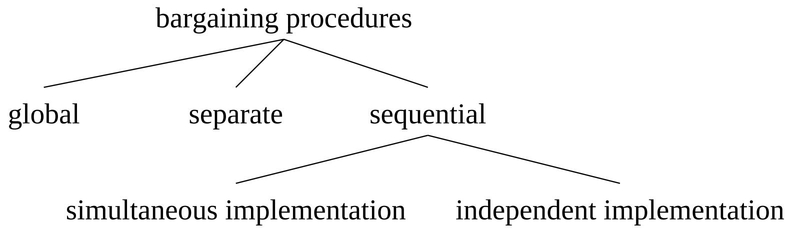

Figure 2.3: Four different bargaining procedures used in multiple-issue bargaining [97].

concede or stand still is not efficient (nor stable): a disagreement will occur with a probability of one fourth. The interested reader is referred to [105] for more details on the characteristics of this mechanism.

2.3.3 Bargaining over multiple issues

The above situations can be described as negotiations about how to divide a surplus. This means that the negotiations are distributive: a gain for one player always creates a loss for the other player. These kinds of negotiations are also referred to as *competitive* [48]. When more than a single issue is involved, and players attach different importance to these issues, tradeoffs become an option and negotiations may become *integrative*. The latter kind of negotiations is the main topic of this section. Results from cooperative game theory are discussed first, followed by a overview of results from non-cooperative game theory.

Cooperative game theory

An additive scoring system or an additive multi-attribute utility function (see Def. 3.1) can be used to represent the relationships or trade-offs between the issues if several issues are involved.[10] However, these methods are appropriate only if the issues are preferentially independent, that is, if the contribution of one issue is independent of the values of the other issues.

Once the preferences are mapped, for instance onto an additive multi-attribute utility function, the bargaining set can be determined. The main goal is again to reach a Pareto-efficient outcome (see Def. 4.3). Previously introduced solution concepts such as the Nash bargaining solution or the Kalai-Smorodinsky solution can be used for this purpose. Several practical considerations (concerning for example fairness of the outcome) and some instructive real-world examples are given by Raiffa in [101].

[10]See [101, pp.154-155] for a discussion of the differences between these methods.

Non-cooperative game theory

Four different bargaining procedures can be distinguished for multiple-issue bargaining [97] (see figure 2.3). In case of *global* or simultaneous bargaining all issues are negotiated at once. The second procedure is called *separate* bargaining. In this protocol the issues are negotiated independently. The final two procedures fall under the header of sequential bargaining and are distinguished by their *rules of implementation*. These rules specify when the players can start enjoying the benefits of the issues which have been agreed on.[11] Three possibilities are considered in [35]. Here, however, we will only mention the most important two. Using the so-called *independent implementation* rule, an agreement on an individual issue takes effect immediately, that is, the agreed upon issues are no longer discounted. In the *simultaneous implementation* on the other hand, the players have to wait until agreement is reached on all issues before they can enjoy the benefits of it. The time it takes to agree on the remaining issues also influences the profits gained on the already agreed upon issues.

When bargaining is sequential an agenda needs to be determined to set the order in which the issues will be negotiated. Agenda setting is of course only relevant if the issues are of different importance. Another concern is whether the players attach the same importance to each issue or whether different players have different evaluations regarding the importance of the issues. The latter is the most interesting case since this allows for integrative negotiations. Unfortunately, however, only a limited literature exists on this topic in game theory. Usually, either the issues are of equal importance (as in [6]) or the players have identical preferences (as in [19]). In [97] the assumption is made that preferences are additive over issues, implying that the multi-issue bargaining problem is equal to the sum of the bargaining problems over the separate issues.

One of the few papers in game theory on integrative bargaining is [35]. Fershtman considers sequential bargaining over two issues. He states that, when using Rubinstein's alternating-offers protocol for each issue in a sequential order, each player prefers an agenda in which the first issue to bargain on is the one which is the least important for him but the most important for his opponent. Notably, it is shown in [35] that the subgame-perfect equilibrium outcome for this problem does not need to be Pareto-efficient.

2.3.4 Bargaining with private information

Private information such as reservation values (i.e., limit values on what the players find acceptable), preferences amongst issues, attitudes towards risk or time preferences are often hidden from the opponent in real-life negotiations. In bargaining it

[11] This is relevant in case the payoff is discounted in the course of time.

might for example be beneficial to be dishonest about one's attitudes towards risk in order to get a greater share of the surplus (as would be the case in Rubinstein's alternating-offers game). Sometimes, however, a mechanism (see Def. 5.1) can be designed which gives agents a compelling incentive to be honest to the opponent. Such mechanisms are called *incentive compatible* (see Section 1.1.2, p. 6).

The Vickrey auction [136] is an example of such an incentive-compatible mechanism (this auction and other incentive-compatible mechanisms are discussed in Section 2.3.5). Unfortunately, however, a suitable mechanism does not always exist. Moreover, such mechanisms are static and mediated (e.g. by an auctioneer) [5]. In practice, bargaining is often dynamic and involves a sequence of offers and counter offers between two or more players. Therefore, it is necessary to analyse dynamic or extensive-form bargaining games with incomplete information. As mentioned in Section 2.2, game theory frequently assumes that the players have complete information. However, in order to analyse situations in which players are unsure of the opponent's type, the notion of imperfect information needs to be applied (see Section 2.2).

Imperfect information enables us to address important issues as reputation building, signalling and self-selection mechanisms [111]. For example, the fact that players are unsure of the other player's type might explain the occurrence of (inefficient) delays in reaching an agreement [89, Ch. 5]. Using such inefficient strategies may be the only way to signal for instance one's strength (an example is the outbreak of strikes during wage bargaining situations). Any utterance which is not backed up by actions can be considered as being cheap talk.[12] Delays may therefore be required to convey private information credible [58].

In a wage negotiation problem, for example, the union is often unsure about the actual value of its workers for a firm. If this value is high, the firm will be more eager to sign an agreement. In case of a low value however, the firm will behave credible by bearing the costs of a strike [58]. A firm could try to "bluff" by ignoring a strike even in case of a high valuation, and use this strategy to signal a lower valuation of the union workers than actually is the case. However, such a strategy can potentially be very harmful.

An overview of bargaining with incomplete information is given in [5]. More introductory texts on bargaining with private information can be found in [58] and [11, Ch. 11].

[12] In non-cooperative games, nothing anyone says constrains its future behaviour. If a player chooses to honour an agreement or threat that has been made, this will only be because it is optimal to do so.

2.3.5 One-to-many bargaining

In a one-to-many bargaining setting, one player negotiates contractual agreements with two or more opponents. A typical example is when a seller has one or more items for sale, and several buyers wish to purchase an item (or a bundle of items). Auctions are the most common *mechanism* (see Def. 5.1) to solve the one-to-many bargaining problem. An alternative approach, using bilateral bargaining, is discussed in Chapters 6 and 7. This section explains the most common auctions or mechanisms and discusses optimal bidding behaviour in these auctions.

We focus here on sealed-bid auctions, where buyers submit positive bids to an auctioneer and the auctioneer selects the winners and the amount that they have to pay.[13] Note that the amount that the winners pay in such auctions does not always correspond to the actual bid, which will become clear below. The auction is called sealed because a buyer's bid is hidden from the other buyers and is only revealed to the auctioneer. Often, the role of the auctioneer is taken by the seller.

Auctions for a single good are discussed first, followed by auctions for more complex cases. We assume in the following that buyers have *independent valuations*. In this context, a the buyer's *valuation* is the highest price that she[14] is willing to pay, such that she is indifferent between paying the highest price and not obtaining the good(s) at all (i.e., both options have equal utility). A player's valuation is *independent* if it does not depend on information available about the preferences of other players, nor on the allocation of the goods to other players.

Single unit

Perhaps the most common sealed-bid auction for selling a single item is the *first-price* auction. In this auction, the item is awarded to the highest bidder, and she pays the price equal to the submitted bid. We can use game theory to derive optimal strategies for the buyers in this auction. Take for example the case where two buyers compete for the good and have different valuations for the good. If a buyer knows the valuation of the other buyer, it is optimal to bid slightly above the valuation of the other buyer if she has the highest valuation, and to bid her valuation otherwise. This strategy constitutes a Nash equilibrium. In case the other buyer's valuation is not known, but is independently drawn from a distribution, the optimal response can again be calculated (we refer the interested reader to [72, p.865] for details). Clearly, the buyer's bid depends on a buyer's speculation about the valuations of other bidders. In general, the buyer will then bid below her valuation.

An interesting alternative auction is the aforementioned *Vickrey* or *second-price* auction [136]. In this auction the highest bidder wins as before, but pays the price

[13]Note that such auctions can be considered *direct* mechanisms (see Section 1.1.2, page 5), in which the players are asked to submit their preferences directly.

[14]In the following, we use *she* for a buyer and *he* to refer to a seller.

bid by the second-highest bidder.[15] In contrast to the previous auction, the optimal strategy in this case is to bid the true valuation for the good, irrespective of the valuations and bids of the other buyers [27, 136].[16] This is in fact a *dominant strategy* (see Def. 18.1). This auction is also called *incentive compatible* (see Section 1.1.2, p.6) because it provides the players with the incentive to reveal their preferences truthfully. Intuitively, this is because a buyer's payment is independent from her bid, and therefore she does not benefit by bidding lower than her valuation. Bidding a higher value is also not beneficial since it can result in paying more than the valuation. In fact, it appears that an auction is incentive compatible if and only if the auction is *bid-independent* [44], i.e., if the bid value of a bidder i does not determine bidder i's payment (but only determines if she wins or not).

The Vickrey or second-price auction has several advantages compared to the first-price auction. First of all, since the second-price auction is incentive compatible, calculating the optimal strategy for the buyers is straightforward. The auction is also robust, since the choices of buyers do not depend on the behaviour of others. Another advantage is that the second-price auction is an *efficient* auction; *efficient* auctions put goods into the hands of the buyers who value them the most [27]. Efficiency is a very desirable property, as it maximises the total gains of trade (i.e., the bargaining surplus). In [27] it is shown that any incentive compatible auction is efficient. By contrast, the first-price is not, in general, efficient. In case of uncertainty about other buyers' valuations and thus speculating buyers, *in*efficient outcomes can occur (see [27] for an example). Below, we consider incentive compatible (and thus efficient) auctions for the more general case of multiple units.

Multiple units

In case multiple goods are traded, the *Generalised Vickrey Auction* (GVA) [133] can be used to allocate the goods efficiently. Like the Vickrey auction, the GVA is also incentive-compatible, that is, truth-telling is a dominant strategy. In this section, we apply the GVA in case multiple (homogeneous) units of the same good are sold (for other applications, see e.g. [133]). The GVA then works as follows.

In the initial stage, each buyer i reports a utility function $u_i(\vec{x})$ to the auctioneer, which may or may not be the true utility function. The vector $\vec{x}$ specifies the number of units allocated to each buyer i.[17] For this application, the utility function expresses the amount of money a buyer is willing to spend for a given allocation $\vec{x}$. The auctioneer then calculates the allocation of units $\vec{x}^*$ that maximises the sum of

[15] In case of a single bidder, this bidder gets the good for free.

[16] This holds assuming independent valuations, as stated before.

[17] For the case described here, we assume that buyers only care about the units they receive, and not about the units received by others (which is part of the valuation independence assumption described earlier), i.e., $u_i(\vec{x}) = u_i(x^i)$; there are no so-called *allocative externalities* [55]. We note, however, that the GVA can also be applied to the case of allocative externalities, see e.g. [134].

utilities, under the constraint that the number of allocated units equals the number of available units. The auctioneer also calculates the allocation that maximises the sum of utilities *other* than that of buyer i. This allocation is denoted by $\vec{x}^*_{\sim i}$. Each buyer i then receives the bundle according to the allocation $\vec{x}^*$ and has to pay the following amount to the auctioneer: $\sum_{j \neq i} u_j(x^*_{\sim i}) - \sum_{j \neq i} u_j(x^*)$. In words, a buyer pays the other buyers' "losses" as a consequence of obtaining the bundle. Note that since the payment of a buyer i does not depend on the utility reported by buyer i, but only on the utilities reported by the other buyers, it follows that this mechanism is incentive compatible. Below we show the application of this mechanism for two examples.

Example 1 In case of a single unit, this mechanism is equivalent to the second-price auction. We show this in the following. We assume (without loss of generality) that a buyer's utility equals zero if no units are allocated to this player. In case buyer i is *not* the highest bidder (i.e., does not report the highest utility value for the good), the allocation is not affected by buyer i (i.e., $\vec{x}^*_{\sim i} = \vec{x}^*$), and the payment $\sum_{j \neq i} u_j(\vec{x}^*_{\sim i}) - \sum_{j \neq i} u_j(\vec{x}^*) = 0$. On the other hand, if buyer i is the highest bidder, then the second part of the equation $[\sum_{j \neq i} u_j(\vec{x}^*)]$ equals zero, since nobody else gets anything. The first part $[\sum_{j \neq i} u_j(\vec{x}^*_{\sim i})]$, however, equals the reported valuation (i.e., bid) of the second-highest bidder, since this would be the (reported) valuation of the winner if buyer i would not participate. The payment therefore equals the reported valuation (i.e., bid) of the second-highest bidder.

Example 2 In case of N units, and if each bidder is allocated up to one unit, the GVA mechanism reduces to an $(N+1)$-price auction, i.e., where each winner pays the price of the $(N+1)$-highest bidder.[18] To see this, consider first the case where buyer i is not a winner. As before, buyer i does not affect the allocation, and therefore pays zero. In the other case, i.e., when buyer i is one of the winners, then $\sum_{j \neq i} \vec{x}^*$ equals the total bids (reported valuations) of the remaining winners. Furthermore, since the unit would go to the $(N+1)$-highest bidder if buyer i would not participate (assuming there are at least $N+1$ participants), $\sum_{j \neq i} \vec{x}^*_{\sim i}$ equals the total valuation of the remaining winners of the actual allocation, plus the valuation of the $(N+1)$-highest bidder. The payment is then exactly the valuation (or bid) of the $(N+1)$-highest bidder. This holds for each winner, assuming there are at least $N+1$ bidders. Note that if there are less than $N+1$ bidders, all bidders receive the good for free.

2.4 Computational approaches to bargaining

Simplifying assumptions frequently made in game-theoretical analyses, such as assumptions of perfect rationality and common knowledge, do not need to be made

[18]This auction is applied in Chapter 5.

if the behaviour of boundedly-rational negotiating agents is modelled directly, for instance using techniques from the field of artificial intelligence (AI). This section provides an overview of the key research related to this thesis, where AI techniques such as evolutionary algorithms, reinforcement learning (specifically Q-learning) and Bayesian beliefs are applied to develop a negotiation environment consisting of intelligent agents. In addition, we shortly review the relatively new field of argumentation-based negotiation. Note that the evolutionary approach is the main focus of this thesis, and therefore the most relevant. The other techniques mentioned are intended for the interested reader.

Using the above-mentioned techniques, agents are able to learn from experience and adapt to changing environments. This learning aspect is essential for automated negotiation settings (where software agents, see Section 1.1.3, bargain on behalf of their owners), especially when the behaviour of competitors and the payoffs are not known in advance. Several aspects of learning are potentially important during the negotiation processes. First, a bargaining agent needs to have a strategy which specifies his actions during the course of play. On the basis of the agent's experiences in previous bargaining games, he can learn that it might be profitable to adjust his strategy in order to achieve better deals. Second, it might even be useful to update a strategy during play. This may be the case if the agent is initially unsure about the *type* of his opponent. After playing a bargaining game for a number of rounds, the agent may form a belief about his opponent's type and fine-tune his behaviour accordingly. Third, an agent might need to learn the preferences of his owner first. Here, attention is focussed on the first two kinds of learning.

This section is organised as follows. Section 2.4.1 discusses the main related research where bargaining agents adapt using evolutionary algorithms (EAs). Q-learning and an application hereof for bargaining is described in Section 2.4.2. Section 2.4.3 approaches learning during the negotiation process using Bayesian beliefs. Section 2.4.4 considers an alternative approach where negotiation is viewed as a dialogue game, and the parties attempt to reach consensus using argumentation.

2.4.1 The evolutionary approach

Oliver [88] was the first to demonstrate that a system of adaptive agents can learn effective negotiation strategies using evolutionary algorithms. Computer simulations of both distributive (i.e., single issue) and integrative (i.e., multiple issue) *alternating-offers* negotiations are presented in [88]. Binary coded strings represent the agents' strategies. Two parameters are encoded for each negotiation round: a threshold which determines whether an offer should be accepted or not and a counter offer in case the opponent's offer is rejected (and the deadline has not yet been reached). These elementary strategies were then updated in successive generations by a genetic algorithm (GA). Similar models are also investigated in this thesis.

In [126], a related model was investigated. Here, a systematic comparison between game-theoretic and evolutionary bargaining models is also made, in case negotiations concern a single issue. Chapters 3 and 4 of this thesis extend similar negotiation models even further by considering multiple issues and cases that are unwieldy to analyse mathematically.

More elaborate strategy representations are proposed in [73]. Offers and counter offers are generated in this model by a linear combination of simple bargaining tactics (time-dependent, resource-dependent, or behaviour-dependent tactics). As in [88], the parameters of these different negotiation tactics and their relative importance weightings are encoded in a string of numbers. Competitions were then held between two separate populations of agents, which were simultaneously evolved by a GA. The time-dependent tactics are further investigated in [34] using GAs, for the case that negotiating agents have different time preferences.

Dworman et. al [31] studied negotiations between three players. If two players decide to form a coalition, a surplus is created which needs to be divided among them. The third party gets nothing. Of course, all three players want to be part of the coalition in this case. Moreover, they also want to receive the largest share of the bargaining surplus. Genetic programming was used in this paper to adapt the offers and to decide whether to form a coalition or not. A comparison with game theoretic predictions and human experiments was made.

Evolutionary algorithms have recently been used not only to generate strategies but also to design auction *mechanisms* (see Def. 5.1 and Section 2.3.5), notably by Cliff [24] and Phelps et al. [96]. Especially for double auctions, where analytical solutions are typically intractable, the evolutionary approach has been successfully applied. Double auctions allow for many buyers and many sellers to exchange goods or services. In this type of auction, sellers and buyers submit bids (offered quantity and price) and asks (demanded quantity and price) respectively, which are then matched by the auctioneer. The auctioneer also determines the trading price for each match. In [96] genetic programming (GP) is used to evolve both the strategies of the traders and the auction mechanism. In this first endeavour towards automated design of auction mechanisms from scratch, GP is used to determine the rule for setting the trading price, while having a fixed matching algorithm. The goal is to optimise *market efficiency*, that is the total profits of both buyers and sellers as a fraction of the theoretical maximum, given that buyers and sellers are only concerned about maximising their individual profits. In a related approach by Cliff [24], a genetic algorithm is used to evolve both the traders and an additional parameter that selects between a continuum of auctions.

2.4.2 Using Q-Learning

Many learning techniques require feedback each time an action is performed. However, in many practical cases feedback is only received at the end of a (long) sequence of actions. A good example is a game like chess: only at the end of play the players know with certainty how well their strategy performs. In learning models like Q-learning, agents also try to evaluate the effect of intermediate actions. Q-learning is a reinforcement learning algorithm [113, p. 528] which learns an action-value function yielding the *expected* utility (see Section 1.1.1) of a given action in a given state [113, p. 599].

This algorithm maintains a list of so-called Q-values $Q(a, i)$, which denote the expected utility of performing an action a at state i. The action which maximises the expected utility is selected, and the system moves to a new state j. The Q-value is then updated depending on the Q-value of the new state and the received reward (if available). The following equation can be used [113, p. 613] for updating the Q-value in case of a transition from state i to j by taking action a:

$$Q(a, i) \leftarrow Q(a, i) + \alpha(R(i) + \max_{a'} Q(a', j) - Q(a, i)), \tag{2.1}$$

where $R(i)$ is the actual reward received in state i and α is the learning rate. The value $\max_{a'} Q(a', j)$ represents the expected utility of state j. For example, if the current state i has a relatively low expected utility and the next state j has a high expected utility, the Q-value $Q(a, i)$ is updated in such a way that the difference between these states is reduced. In this way rewards which are given at the terminal state are passed to the other states in the sequence.

As we mentioned before, selecting an action in the current state depends on the expected utility of each action. Hence, a trade-off needs to be made between "exploitation" and "exploration". In other words, should an action be chosen which has already proven itself or do we prefer to try out new actions which might produce even better results? This question of finding an optimal exploration policy has been studied extensively in the subfield of statistical theory that deals with so-called "bandit" problems [113, pp. 610-611].

The Q-learning approach was applied by Oliveira and Rocha [87] for the formation of virtual organisations in an e-commerce environment. The idea is that in order to satisfy some user's need, often a combination of services is needed, which is provided by different companies. The agent representing the user (called the "market agent") negotiates with several organisation agents, after which a selection of these organisations is made and a virtual organisation is created. The protocol used during the negotiation phase is as follows. First, each participating organisation generates a bid, based on previous experience, and sends this bid to the market agent. A Q-learning technique is then used to determine which bid to make. The actions (i.e., the bids) made are then evaluated using the feedback given by the market agent. The market agent compares the bids using a multi-criteria evaluation method based on qualitative measures (in which only the preference ordering is assumed to be important). The market agent selects the organisation which either proposes a satisfactory evaluation, or he chooses the highest evaluation when a deadline is reached. Organisations not selected are given feedback as to which attributes were not satisfactory. Negotiations take several rounds, and each round an organisation is selected.

2.4.3 Using Bayesian beliefs

Bayesian beliefs are used to model an agent's (probabilistic) knowledge of an uncertain environment. Suppose the agent has some a priori knowledge about the likelihood of a set of hypotheses H_i, with $i = 1, ..., n$. Furthermore, the agent has some conditional knowledge about the probability that an event e will occur, given that one of the hypotheses is true. If event e then occurs, the beliefs about the hypotheses are updated using the Bayesian update rule [148]:

$$P(H_i|e) = \frac{P(H_i)P(e|H_i)}{\sum_{k=1}^{n} P(e|H_k)P(H_k)}, \tag{2.2}$$

where $P(H_i|e)$ is the a posteriori probability of H_i and $P(H_i)$ the a priori probability. $P(e|H_i)$ is the conditional probability that event e occurs given hypothesis H_i.

When agents have incomplete information about one another, it becomes important to learn about the other agent by observing his behaviour during the negotiation process. Bayesian beliefs are often used to make assumptions about the opponent such as his *type* [64] or his *reservation price* [147],[148] (where the reservation price is defined here as an agent's threshold of offer acceptability). These beliefs are updated depending on the opponent's moves.

However, once both agents use beliefs to determine their strategies, they also need beliefs about their opponent's beliefs, and so on. This is known as the problem of *outguessing regress* [148]. In game theory this problem is solved by having a limited number of different types of players. The beliefs and preferences of each type are common knowledge, but there is uncertainty about which player is of which type. This theory, suggested by Harsanyi, is a technique for transforming a game of incomplete information into a game of imperfect (but complete) information (see also Section 2.2). In reality however, the number of different types is usually very large, and, moreover, it is not always realistic to assume that the preferences and beliefs of the different types are common knowledge. In more practical applications (such as [64] and [147]), the problem of outguessing regress is circumvented by assuming limited reasoning capabilities. In [147], for instance, a player has beliefs about e.g. the payoff function and reservation price of the other player, but not about the beliefs of the other player.

2.4.4 Argumentation-based negotiation

An alternative approach to automated negotiation is the use of *dialogues* or argumentation to resolve conflicts. In recent years, this field has received increasing interest within the agent community [71, 74, 94, 99, 100]. We therefore relate some of the main concepts and highlight some of the research in this field. A more extensive overview of the state-of-the-art on argumentation-based negotiation can be found in [99].

Argumentation can be useful when, for example, negotiations involve several issues and a mutually beneficially situation can be achieved (as described in Section 2.3.3). When agents have incomplete information about each others' preferences negotiations, inefficient deals are often obtained (see Section 2.3.4). This problem can be resolved using argumentation. The idea is that the agents are able to provide meta-information on *why* they have a particular objection to a proposal. This way, information is exchanged, but without fully disclosing each others' preferences.

A negotiation architecture using this kind of meta-information is described in [94]. This approach was also used in MIT's Tête-à-Tête system, a bilateral integrative negotiation system for online shopping [71]. Agents within this framework can: (1) make a new proposal, (2) accept the proposal of the counter agent, (3) criticise a proposal or (4) withdraw from the negotiations. This system uses the notion of a *critique* to enable agents to criticise a particular proposal. A critique is a comment of an agent specifying which part of the proposal he dislikes. In case of a new proposal or critique, the agent can also send additional information. For instance, a proposal may include conditions under which it holds (e.g., I will provide you with X if you provide me with Y).

Argumentation can also be used to influence the preferences, beliefs and/or goals of other players. In general, preferences are assumed to be fixed. In reality, however, it is often true that a player's preferences are not completely formed or that uncertainty exists about the environment. In that case, a player's preferences and beliefs can be influenced upon receipt of new information. The negotiation process then not only consists of dividing the surplus, but also of gathering information. An interesting approach is described in [100], where one player may influence another player's preferences by discussing the underlying motivations and interests behind adopting certain (sub)goals. For example, a buyer may want to negotiate a flight ticket with a travel agent for the more fundamental goal of travelling to Paris. If the fundamental goal is known to the travel agent, she can suggest a train ticket as an alternative means to satisfy the same goal. Another way of influencing a player's behaviour is by means of persuasion, for example by using threats, rewards or appeals [102].

2.5 Discussion

The first part of this chapter reviews, in broad lines, literature on bargaining from the field of game theory. This overview shows that game theory is a very useful tool to analyse bargaining situations in a mathematical fashion. Such a rigorous analysis is only tractable, however, if many details of human interaction, for instance emotions or irrational behaviour, are abstracted away. This may undermine the capability of game-theoretical models to explain or predict human behaviour.

This aspect may be less problematic when we consider systems in which *artificial* agents interact with each other, because these agents are often designed to behave (in good approximation) in a rational fashion. Game theory may therefore yield fundamental insights in the design of efficient negotiation protocols for automated trading. Furthermore, given a negotiation protocol and under certain assumptions, optimal strategies can sometimes be derived.

Nevertheless, game-theoretical assumptions like common knowledge and perfect rationality often appear to be too strong in modelling practical situations. The issue of common knowledge has been solved only partially in game theory by introducing a theory for players with *imperfect* information. The development of game-theoretic models for boundedly-rational players is a relatively young research direction. Our survey shows that techniques from the field of artificial intelligence are potentially very powerful in situations of incomplete information and boundedly-rational players. Learning techniques developed within the AI community can for instance be used to adapt the agents' behaviour in complex environments and to construct accurate models of the other agents' preferences.

Part A

Fundamental aspects of bargaining systems

Chapter 3

Multi-issue bargaining by alternating offers

Automated negotiations have received increasing attention in the last years, especially from the field of electronic trading [14, 56, 65, 71, 73, 88, 128]. In the near future, an increasing use of bargaining agents in electronic market places is expected. Ideally, these agents should not only bargain over the price of a product, but also take into account aspects like the delivery time, quality, payment methods, return policies, or specific product properties. In such *multi-issue* negotiations, the agents should be able to negotiate outcomes that are beneficial for both parties. The complexity of the bargaining problem increases rapidly, however, if the number of issues becomes larger than one. This explains the need for "intelligent" agents, which should be capable of negotiating successfully over multiple issues at the same time.

In this chapter,[1] we consider negotiations that are governed by a finite-stage version of the Rubinstein-Ståhl multi-round bargaining game with alternating offers (see Section 2.3.2 and [110, 121]). We investigate the computation of strategies of the agents by evolutionary algorithms (EAs) in case negotiations involve multiple issues. We first assess the efficiency of the agreements reached by the *evolutionary agents* (see Section 1.1.3). We then analyse to what extent the evolutionary outcomes match with game-theoretic results. We study models in which time plays no role and models in which there is a pressure to reach agreements early (because a risk of breakdown in negotiations exists after each round).

Furthermore, we present and study a more realistic negotiation model, where agents take into account the fairness of the obtained payoff. This use of fairness is based on the following observation. When no time pressure is present, extreme divisions of the payoff occur in the computational experiments, due to a powerful

[1]The results in this chapter have are published as [42]: E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. *Computational Economics*, 22:39–63, 2003.

'take-it-or-leave-it' position for one of the negotiating agents in the last round of the negotiation. Although such extreme outcomes are in agreement with game-theoretic results, they are usually not observed in real-life situations, where social norms such as fairness play an important role [13, 67, 107, 141]. We therefore introduce a fairness norm and incorporate this in the agents' behaviour. We perform computational experiments with various fairness settings, and show that, depending on the actual settings, "fair" deals indeed evolve.

A number of related paper demonstrate that, using an EA, artificial agents can learn effective negotiation strategies [34, 73, 88, 126] (see also Section 2.4.1). In [126], a systematic comparison between game-theoretic and evolutionary bargaining models is made, in case negotiations concern a single issue. In [34] single-issue negotiations are also studied using a genetic algorithm, when agents can select between a number of pre-specified strategies. The multi-issue problem is considered in [73, 88]. The main contribution in this chapter lies in the validation of the evolutionary model for multi-issue negotiations with possible breakdown, using game-theoretic subgame-perfect equilibrium (see Def. 18.3), and the introduction of a fairness norm in such negotiations. Especially the latter is a first attempt to study complex bargaining situations which are more likely to occur in practical settings. A rigorous game-theoretic analysis is typically much more involved or may even be intractable under these conditions.

The chapter is organised as follows. The alternating-offers negotiation protocol for multiple issues is described in Section 3.1. Section 3.2 gives an outline of the evolutionary simulation environment and how the strategies of the agents are represented. A comparison of the computational results with game-theoretic results is presented in Section 3.3. The extension with fairness is the topic of Section 3.4. Section 3.5 summarises the main results and concludes.

3.1 Description of the bargaining game

We consider negotiations that are governed by a finite-stage version of the Rubinstein-Ståhl multi-round bargaining game with alternating offers (see Section 2.3.2 for details). During the negotiation process, the agents exchange offers and counter offers in an alternating fashion at discrete time steps (rounds). In the following, the agent starting the negotiations is called "agent 1", whereas his opponent is called "agent 2".

Bargaining takes place over m issues simultaneously, where m is the total number of issues. We assume that mutual gains are possible for each issue by reaching an agreement, i.e., that a positive *bargaining surplus* is available (see also Section 1.1.2) for each issue. We further assume (without loss of generality) that the total bargaining surplus available per issue is equal to unity. We express an offer as a vector $\vec{o}$, where the i-th component o^i specifies the share that agent 1 receives of the bargaining surplus for issue i if the offer is accepted. Agent 2 then receives $1 - o^i$ for issue i. The index i ranges from 1 to m. Note that an offer always specifies the share obtained by agent 1.

The agents evaluate multi-issue offers using an additive multi-attribute utility function (see Def. 3.1 and [73, 88, 101]). Agent 1's utility function is $\vec{w_1} \cdot \vec{o_j}(r) = \sum_{i=1}^{m} w_1^i \cdot o_j^i(r)$, where $j = 1$ if the offer is proposed by agent 1 and $j = 2$ otherwise. Agent 2's utility function is $\vec{w_2} \cdot [\vec{1} - \vec{o_j}(r)]$. Here, $\vec{w_j}$ is a vector containing agent j's weights w_j^i for each issue i. The weights are normalised and larger than zero, i.e., $\sum_{i=1}^{m} w_j^i = 1$ and $w_j^i \geq 0$. Because we assume that $0 \leq o_j^i(r) \leq 1$ for all i, the utilities are real numbers in $[0, 1]$.

As stated above, agent 1 makes the initial offer. If agent 2 accepts this offer, an agreement is reached and the negotiations stop. Otherwise, play continues to the next round with a certain continuation probability p $(0 \leq p \leq 1)$. When a negotiation is broken off prematurely, both agents receive a utility of zero.

If negotiations proceed to the next round, agent 2 needs to propose a counter offer, which agent 1 can then either accept or refuse. This process of alternating bidding continues for a limited number of n rounds. When this deadline is reached without an agreement, the negotiations end in a disagreement, and both players receive nothing.

3.2 The evolutionary system

We use an EA to evolve the negotiation strategies of the agents. Implementation details of the EA are discussed in Section 1.2.3. Each strategy in the EA is associated with either an agent of type 1 (i.e., initiating the negotiation) or of type 2. The strategies of competing agents evolve in separate populations[2]: the strategies of the agents of type 1 evolve in population 1, and of type 2 in population 2. This way, the EA populations co-evolve since the performance of a strategy depends on the strategies in the opponent's population. An overview of the evolutionary system with separate populations for the strategies of the two agent types is depicted in Figure 3.1.

The fitness of the parents is determined by negotiation between the agents in the two parental populations (as shown in Fig. 3.1). Each agent negotiates with all agents in the population of the opponent. The utility functions are the same for agents within the same population (i.e., the weight settings are equal). The average utility obtained in all negotiations is an agent's fitness value. The fitness of the

[2]It is also possible to use a single population with strategies for both agent types on a single chromosome. The outcomes, however, are then affected by so-called *hitchhiking* [75], where relatively poor genes are selected because other genes on the chromosome yield a good performance.

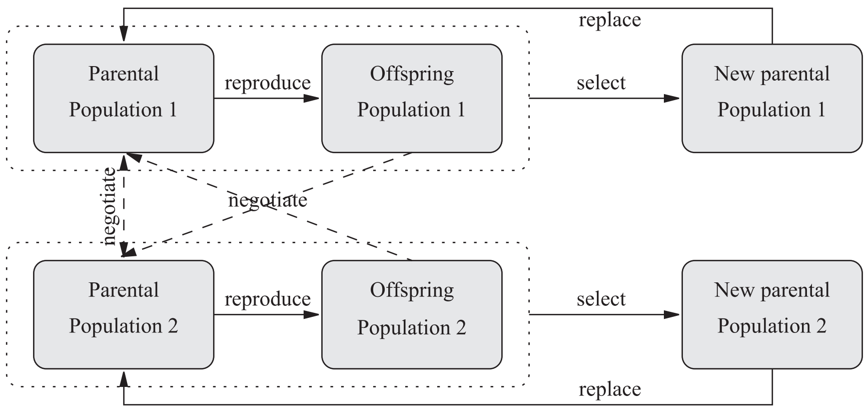

Figure 3.1: Iteration loop of the evolutionary algorithm where strategies for competing agents evolve in separate populations.

EA Parameters	Parental population size (μ)	25
	Offspring population size (λ)	25
	Selection scheme	$(\mu + \lambda)$-ES
	Mutation model	self-adaptive
	Initial standard deviations ($\sigma_i(0)$)	0.1
	Minimum standard deviation (ϵ_σ)	0.025
Negotiation parameters	Number of issues (m)	2
	Number of rounds (n)	10
	Weights of agents in population 1 ($\vec{w_1}$)	$(0.7, 0.3)^T$
	Weights of agents in population 2 ($\vec{w_2}$)	$(0.3, 0.7)^T$

Table 3.1: Default settings of the evolutionary system.

new offspring is evaluated by negotiation with the parental agents.[3] A social or economic interpretation of this parent-offspring interaction is that new agents can only be evaluated by competing against existing or "proven" strategies.

3.2.1 Representation of the strategies

An agent's strategy specifies the offers and counter offers proposed during the process of negotiation. In a game-theoretic context, a strategy is a plan which specifies an action for each history [11]. In our model, the agent's strategy specifies the offers $\vec{o}_j(r)$ and thresholds $t_j(r)$ for each round r in the negotiation process for agents $j \in \{1, 2\}$.

The threshold determines whether an offer of the other party is accepted or

[3]In an alternative model, not only the parental agents are used as opponents, but also the newly-formed offspring. Similar dynamics have been observed in this alternative model.

Agent 1	$\vec{o}_1(1)$	$t_1(2)$	$\vec{o}_1(3)$	$t_1(4)$	...
Agent 2	$t_2(1)$	$\vec{o}_2(2)$	$t_2(3)$	$\vec{o}_2(4)$	...

Figure 3.2: The strategies for agent $j \in \{1, 2\}$ specify a sequence of offers $\vec{o}_j(r)$ and thresholds $t_j(r)$ for rounds $r \in \{1, 2, \ldots, n\}$ of the negotiation.

rejected: If the value of the offer (see below) falls below the threshold the offer is refused; otherwise an agreement is reached.[4] This strategy representation is depicted in Fig. 3.2. Notice that in each round, the strategy of an agent specifies either an offer or a threshold, depending on whether the agent proposes or receives an offer in that round. Note that in odd rounds, agent 1 makes an offer and agent 2 either accepts or rejects, and visa versa in even rounds.

The strategy, consisting of offers and thresholds, is encoded on the chromosome using real values in the unit interval (one offer or threshold for each negotiation round). We use x_i to denote the (real) value at location i of the chromosome. The agents' strategies are initialised at the beginning of each EA run by drawing random numbers in the unit interval (from a flat distribution).

3.3 Validation and interpretation of the evolutionary experiments

Experimental results obtained with the evolutionary system are presented in this section. All relevant settings of the evolutionary system are listed in Table 3.1 (further explanation is provided in Section 1.2.3). A comparison with game-theoretic results is made to validate the evolutionary approach. Section 3.3.1 addresses the evolution of efficient negotiation results. Section 3.3.2 further analyses the results and compares the experimental results with predictions from game theory. In the following, we refer to the agents in the evolutionary system as *evolutionary agents* (see Section 1.1.3).

3.3.1 Efficiency

First, we investigate the experimental results w.r.t. disagreements. Without breakdown ($p = 1$), disagreements can only occur when the deadline is reached. The experiments show that the percentage of disagreements is then very small (around 0.1% after 1000 generations if $n = 10$). With a risk of breakdown of 30% ($p = 0.7$),

[4] A similar approach was used in [88, 126].

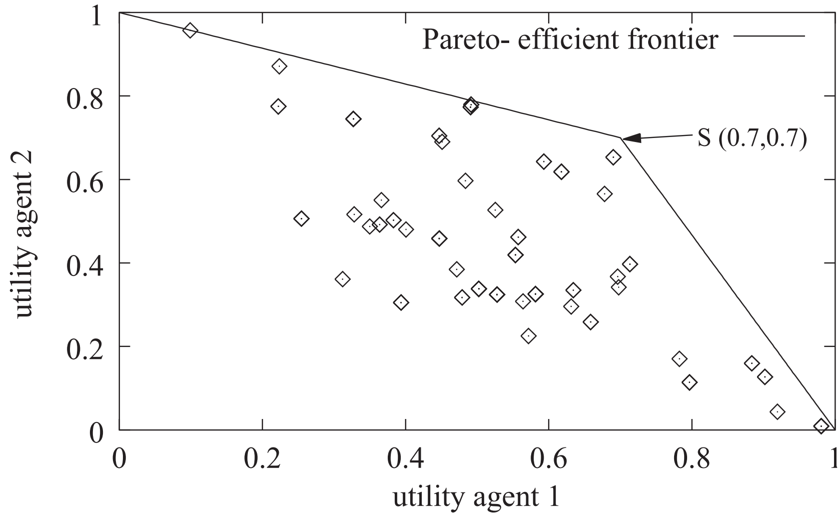

(a)

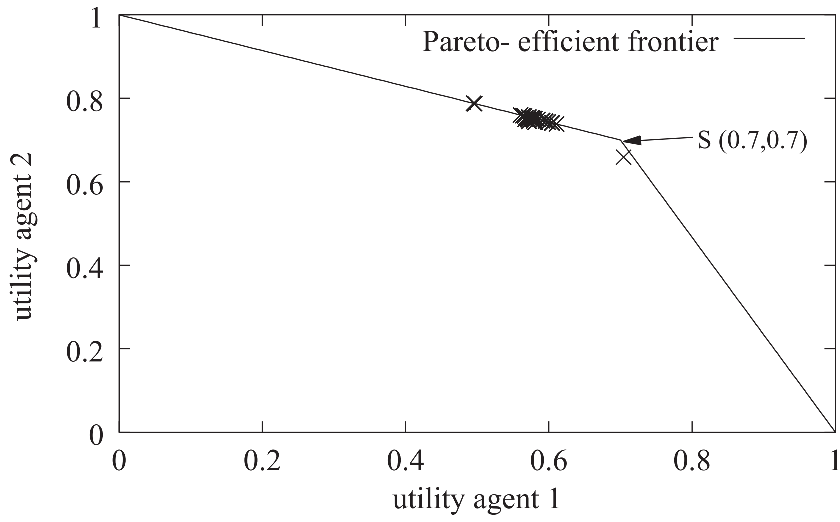

(b)

Figure 3.3: Agreements reached by the evolutionary agents at (a) the start of a typical EA run and (b) after 100 generations. The negotiation settings are $p = 0.7$ and $n = 10$. Each agreement is indicated by a point in these two-dimensional spaces. The Pareto-efficient frontier is indicated with a solid line. In point S [at $(0.7, 0.7)$] both agents obtain the maximum share for their most important issue, and receive nothing for the other issue.

this percentage is between 1% and 10%. Timing is now important for efficiency. The evolutionary agents avoid disagreements by reaching agreements early: after 1000 generations, approximately 75% is reached in the first round.

Next, we study the efficiency of the agreements reached in the experiments. The agreements are depicted in Fig. 3.3. This figure shows the utilities for both agents of the deals reached. Also depicted in Fig. 3.3 is the so-called "Pareto-efficient frontier". An agreement is located on the Pareto-efficient frontier when an increase of utility for one agent necessarily results in a decrease of utility for the other agent. Agreements can therefore never be located above the Pareto-efficient frontier. A special point is the symmetric point S [at $(0.7, 0.7)$], where both agents obtain the maximum share of the issue they value the most, and receive nothing of the less important issue.

Figure 3.3 shows that initially, many agreements are located far from the Pareto-efficient frontier. After 100 generations, however, the agreements are chiefly Pareto-efficient. We note that, even in the long run, the agents keep exploring the search space, resulting in a continuing moving "cloud" of agreements along the frontier.

Conclusion. Results in this section thus show that the evolutionary agents reach efficient agreements, viz. on the Pareto-efficient frontier, and that disagreements are avoided. The next section studies the actual outcomes more closely, using results from game theory as a benchmark.

3.3.2 Further Analysis

The computational results are analysed in more detail in this section and compared with game-theoretic results, and in particular the subgame perfect equilibrium (SPE) predictions (see Def. 18.3). Rubinstein and (much earlier) Ståhl applied this notion to the alternating-offers bargaining game [110, 121]. Our experimental setup differs in two respects from their model, however. First, the agents bargain over multiple issues instead of a single issue. Second, the evolutionary agents are "myopic": they do not apply any explicit rationality principles in the negotiation process, nor do they maintain any history. Actually, they only experience the profit of their interactions with other agents. The SPE behaviour of rational agents with complete information will nevertheless serve as a useful theoretical benchmark. The equations for deriving the SPE outcomes in case of multiple issues are presented in Appendix 1.

We distinguish between three classes of experiments w.r.t. the breakdown probability: (1) no risk of breakdown ($p = 1$), (2) a low breakdown probability ($0.8 \leq p < 1.0$) and (3) a high breakdown probability ($p < 0.8$). For each of these classes we consider the role of n on the outcomes.

We found that in our experiments, when $p = 1$, in the long run almost all agreements are delayed until the last round (about 80% after 1000 generations). Furthermore, the last offering agent makes a take-it-or-leave-it deal and demands almost the entire surplus (on each issue), which is accepted by the opponent. This extreme division of the surplus agrees with game-theory (see Appendix 1.1); it is rational for the responder to accept any positive amount in the last round. Note, however, that rational agents are indifferent about the actual round in which the agreement is reached. The deadline-approaching behaviour in our experiments corresponds better to "real-world" behaviour [108], however.

The EA results and SPE outcomes for different values of n (game length) are compared in Fig 3.4a. To guide the eye, the SPE outcomes for successive values of n are connected. Notice that the fitness of agents in population 1 converges to unity if n is odd, and to zero if n is even (the opposite holds for the agents in population 2). Figure 3.4b shows the results for $p = 0.95$. Note that the partitioning becomes less extreme with a low breakdown probability compared to no breakdown. This holds for both SPE outcomes and EA results, although the effect is much stronger in the evolutionary system (see Fig. 3.4b). These differences with SPE are due to the myopic properties of the agents in the EA. The evolutionary agents do not reason backwards from the deadline (as in SPE), since most agreements are reached in the first few rounds (if $p < 1$). As a result, the deadline is not perceived accurately by the evolving agents. In fact, the game length is strongly overestimated. Furthermore, in SPE *all* agreements are reached without delay (see [126]). The EA, on the other hand, also continues to explore other strategies, which results in a remaining small number of disagreements (see Section 3.3.1).

As p becomes smaller, the influence of the game length on the SPE outcome also decreases (see [126]). Instead, the first-mover advantage becomes more important. Therefore, if p becomes sufficiently small (e.g., $p < 0.8$), the computational results automatically show a much better match with SPE outcomes than if p is large: the match is almost perfect, although a small number of disagreements occur due to a continuing exploration of new strategies. This is clearly visible in Figure 3.5, which shows long-term results for $n = 5$ and different breakdown settings p.

Interestingly, in the limit of $n \to \infty$, game theory predicts that the agents in population 1 have a fitness of ≈ 0.71 when $p = 0.95$, whereas the agents in population 2 have a fitness of ≈ 0.68. This corresponds to a point in the vicinity of the symmetric point S, indicated in Fig. 3.3. The results reported in Fig. 3.4b show that the behaviour of the agents corresponds much better to an infinite-horizon model than the finite-horizon model for $n \geq 5$ (see Fig. 3.4b). The same behaviour was observed for other EA settings (e.g., larger population size) and other negotiation situations (e.g., other weight settings).

We also studied the performance of the EA in case the number of issues m is increased to 8.[5] We observe that, for $p = 1$, the long-term outcomes of the EA

[5] The 8-dimensional weight vector for agents in population 1 is set to $\frac{1}{3.9}(0.7, 0.3, 0.5, 0.2, 0.3, 0.4, 0.5, 1.0)^T$ and equal to $\frac{1}{3.9}(0.3, 0.7, 0.5, 1.0, 0.5, 0.5, 0.2, 0.2)^T$ for agents in population 2. These settings are such that they contain both "competitive" issues (e.g.,

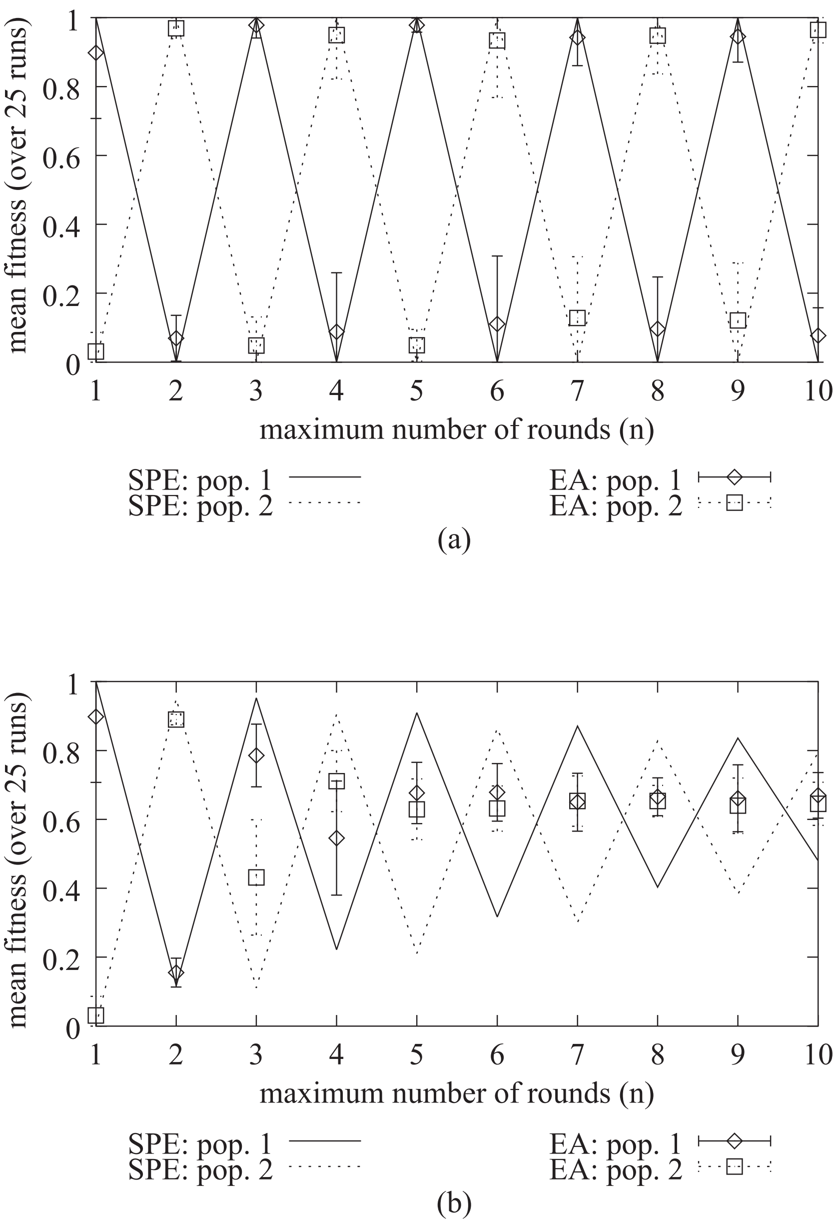

Figure 3.4: Comparison of the long-term evolutionary results with SPE results for (a) $p = 1$ (time indifference) and (b) $p = 0.95$. The error bars indicate the standard deviations across 25 runs.

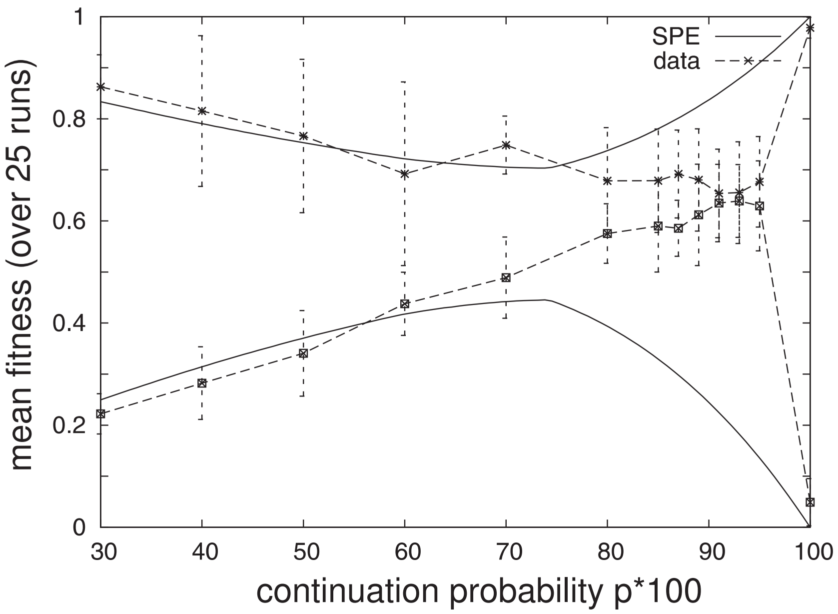

Figure 3.5: Average long-term results using 2 issues for different values of p, where $n = 5$.

are unstable and do not converge to the extreme partitioning. When we increase the population size for the EA from 25 to 100 agents,[6] the extreme partitioning reappears. Results are shown in Figure 3.6. Thus, for more complicated bargaining problems, the EA parameters must be adjusted. For $m = 8$ and $p < 1$, similar observations are found as reported in Section 3.3.2 (like Fig. 3.4) when using the adjusted population size.

Conclusion. Game-theoretic (SPE) results appear to be a very useful benchmark to investigate the results of the evolutionary simulations. In computational simulations without a risk of breakdown (case 1), agreements are predominantly reached in the final round. This deadline effect is consistent with human behaviour [108]. Furthermore, the last agent in turn successfully exploits his advantage and claims a take-it-or-leave-it deal (as in SPE). In case of a small risk of breakdown (case 2), the deadline is not accurately perceived by the evolving agents, and the last-mover advantage is smaller than predicted by game theory. In fact, if the finite game becomes long enough, results match the SPE outcomes for the infinite-horizon game. With a high risk of breakdown (case 3), however, this deviation from SPE becomes negligible. Finally, it appears to be important to adjust the EA parameter settings (e.g., by increasing population sizes) for more complex bargaining problems.

issue 3) and issues where compromises can be made (e.g., issue 8).

[6]To avoid a (quadratic) increase in the number of fitness evaluations, each agent negotiates with 25 (random) opponents.

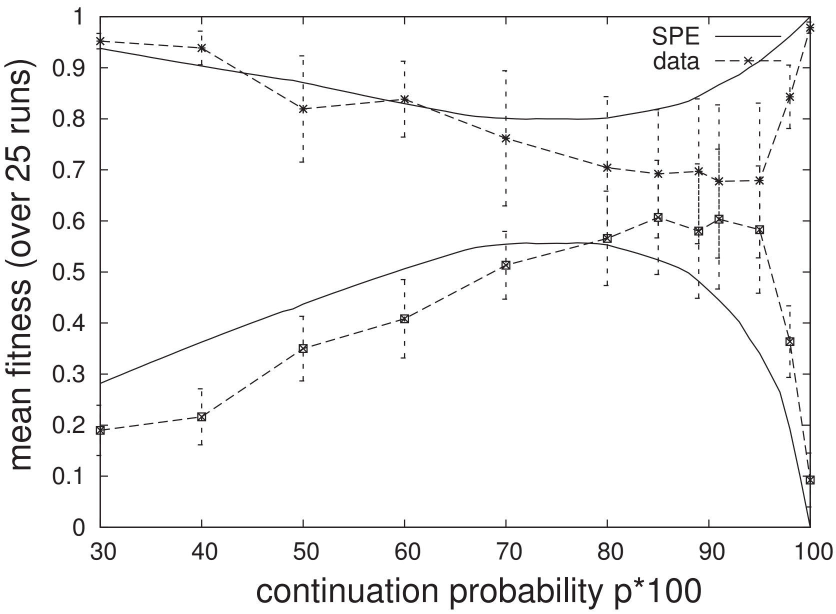

Figure 3.6: Average long-term results using 8 issues for different values of p, where $n = 5$. These results are obtained using a population size of 100.

3.4 Social extension: fairness

We extend the agent model within our evolutionary system in this section to study the influence of "fairness", an important aspect of real-life bargaining situations. The motivation and description of this fairness model is given in Section 3.4.1. In the fairness model studied in Section 3.4.2 the evolving agents only take the fairness of a proposed deal into account when the deadline is reached. Section 3.4.3 presents results obtained when agents perform a "fairness check" in each round. Section 3.4.4 further analyses the model in Section 3.4.3 for a simple case.

3.4.1 Motivation and description: the fairness model

Game-theoretic models for rational agents often predict the occurrence of very asymmetric outcomes for the two parties. We showed in Section 3.3.2 (see Fig. 3.4a) that such "unfair" behaviour can also emerge in a system of evolving agents, in particular when $p = 1$ or n is small (see Fig. 3.4). Large discrepancies between human behaviour in laboratory experiments and game-theoretic outcomes are found, however, both for ultimatum (a single round) and multi-stage (several rounds) games [13, 25, 67, 107, 109, 141]. A possible explanation for the occurrence of these discrepancies between theory and practice is the strong influence of social or cultural norms on the individual decision-making process. In [107, p. 264] and [50], for example, it is argued that responders tend to reject unfair or "insultingly low" proposals. There-

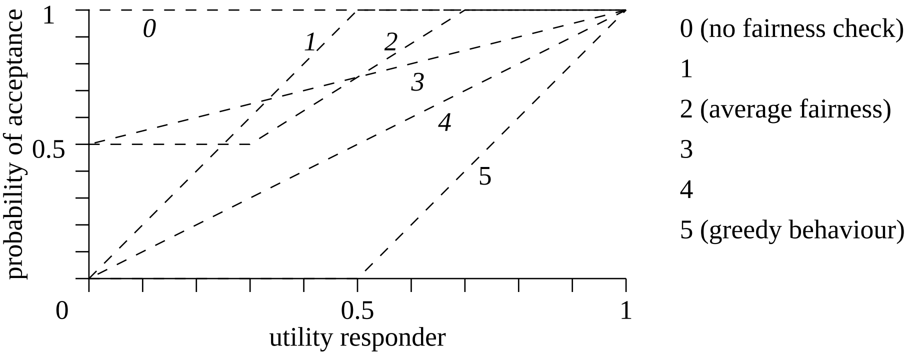

Figure 3.7: Fairness functions used by the agents in the EA.

fore, an anticipating agent should lower his demand in order to avoid a disagreement, this way taking into account the expectations about his opponent's behaviour.

In [67] a model is proposed in line with this hypothesis. In their model, the probability of acceptance of an offer increases with the amount offered to the responder. Such a model, making more realistic assumptions about the agents' behaviour, appears to organise the data from experiments with humans better than the SPE model [67].

Following [67], we introduce a fairness model in our evolutionary system. The agent model is extended as follows. If the value of an offer exceeds the responder's threshold, the agent has the opportunity to re-evaluate his decision. The probability that he finally accepts the agreement is then a function of the acquired utility. This so-called "fairness function" is assumed to be piece-wise linear (with up to three segments).[7] The instances that we use are shown in Fig. 3.7.[8] We now further distinguish between two different extended agent models. In the first model, the fairness function is used at the deadline only. This situation is studied in Section 3.4.2. In the second model, the fairness function is effective at any moment. This case is studied in Section 3.4.3. The first case is motivated by the deadline-effect observed in the experiments without a risk of breakdown (see Section 3.3.2), where most agreements are reached in the last round. The second case, however, is more likely to be an appropriate model of human behaviour.

3.4.2 Fairness check at the deadline

In this section, fairness is applied in the last round. We study the case in which $p = 1$ and $n = 3$. Figure 3.8 shows that if the evolving agents in population 2 use fairness function 1 (i.e., a "weak" fairness model), the partitioning is much less extreme than

[7]Piece-wise linear functions nicely fit the experimental data reported in [67, 109].

[8]We want to remark here that, although the fairness function is the same for all agents, the actual fairness function can depend on cultural norms in the real world [67].

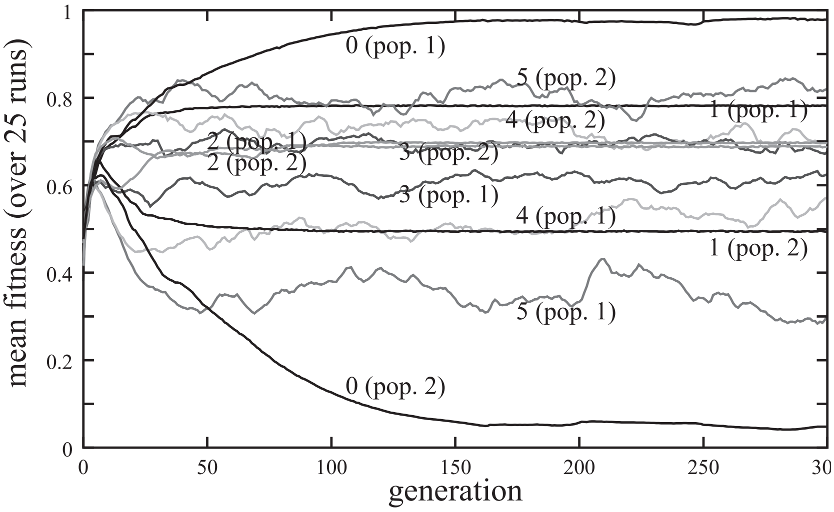

(a)

Figure 3.8: Mean fitness when fairness functions 0-5 are applied at the deadline.

in case of no fairness check (function 0). However, the agents in population 1 still reach a relatively high fitness (utility) level. Fair agreements evolve, on the other hand, when the agents in population 2 use function 2 (a case with average fairness). In this case the mean long-term fitness is approximately equal to 0.7 for all agents (most agreements are thus located close to the symmetric point S in Fig. 3.3).

When stronger fairness functions (e.g., functions 3 through 5) are used by the agents the roles reverse, and the agents in population 2 reach a *higher* fitness level than their opponents in population 1 (see Fig. 3.8). Because of the strong fairness check, many last-round agreements are rejected in this case and agents in population 2 can demand a larger share of the surplus in the round before last. As a result, the deadline is effectively reached one round earlier. This effect indeed occurs in our experiments.

Conclusion. Our results show that fair outcomes can evolve in an evolutionary system with a fairness model in the last round. However, there is a rather large sensitivity to the actual fairness function that is used by the evolved agents; an "average" fairness function yields symmetric results, whereas more extreme fairness functions yield more asymmetric outcomes.

3.4.3 Fairness check in each round

This section studies the second fairness model, in which the responding agent re-evaluates all potential agreements. The EA settings are the same as in the previous section.

The results in Fig. 3.9 for fairness functions 1 are similar to the previous case

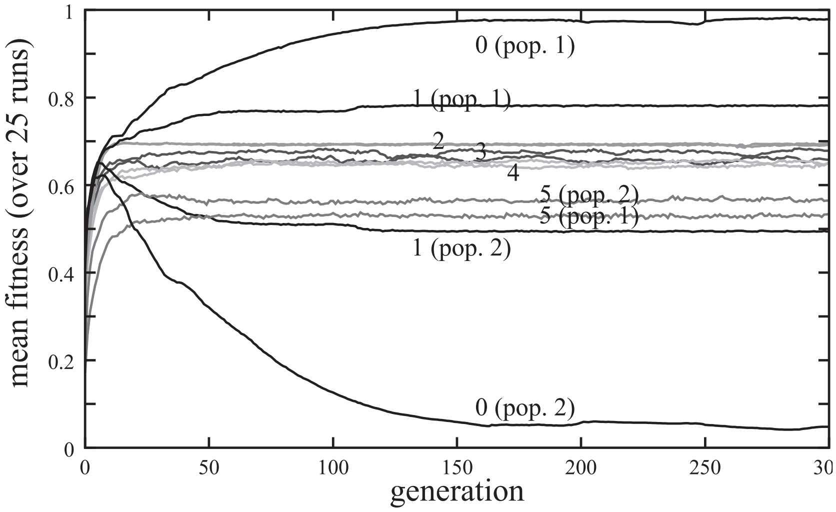

(b)

Figure 3.9: Mean fitness when fairness functions 0-5 are applied each round.

(see Fig. 3.8). However, when fairness functions 2 through 5 are used, the agents in both populations reach almost identical fitness levels. Most agreements now occur in the vicinity of point S in Fig. 3.3. Note that the agents have no explicit knowledge about the location of this point, and that this knowledge is also not incorporated within the fairness functions. We also observe that agreements are now reached in different rounds, whereas in earlier experiments without fairness most agreements occur at the very end of the game.

Fig. 3.9 thus shows that the agents' long-term behaviour is much less sensitive to the shape of the fairness function: the various "stronger" fairness functions all yield similar results. Figure 3.9 however indicates that when the agents use fairness function 5, the mean fitness of both agents decreases. This is due to the increasing number of disagreements which are a result of the strong fairness check.

We furthermore studied a 2-issue negotiation problem with an asymmetric Pareto-efficient frontier, as shown in Fig. 3.10. In this case, agent 1 values both issues equally important, whereas agent 2 has different valuations for each issue (his weights are 0.2 and 0.8 for issues 1 and 2 respectively). If each agent obtains the whole surplus on his most important issue, agent 1 obtains 0.5, whereas agent 2 gets 0.8. This outcome corresponds to the Nash bargaining solution (NBS), see Section 2.3.1. The symmetric point (S), on the other hand, is located at $(\frac{8}{13}, \frac{8}{13})$.[9]

Both solutions can be considered to be fair outcomes in different ways: the first solution maximises the product of the agents' utilities and also splits the surplus equally, whereas in the second case equal utility levels are obtained for both agents

[9]This outcome corresponds to the Kalai-Smorodinsky solution, see Section 2.3.1.

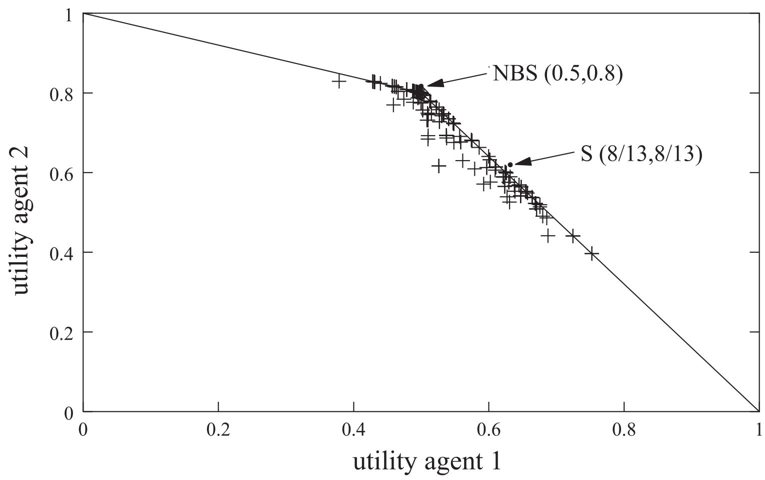

Figure 3.10: Resulting agreements in a single generation when the Pareto-efficient frontier is asymmetric and fairness function 4 is used.

(see [101, Ch. 16] for a related discussion). In the computational results, we observe that, when fairness functions 2-5 are applied, the agreements are divided and are usually concentrated in two separate clusters ("clouds"), see Fig. 3.10. The issue of the choice of and distribution over multiple "fair" agreement points seems an important issue for further research, both in a computational setting as well as in experimental economics.

We also experimented with different weight vectors and with $m > 2$. A general finding is that extreme outcomes do not occur in the evolutionary process if the agents apply a fairness check.

Conclusion. We have shown that fair agreements can evolve if fairness is evaluated each round, even with strong fairness norms: the fairness of the deals is much more stable w.r.t. the actual choice of the fairness function. Of course, the number of actual agreements drops if a very strong fairness function is used, resulting in a lower fitness for both parties. In case of two-issue negotiations with a symmetric Pareto-efficient frontier, most agreements are reached in the vicinity of the symmetric point. In the asymmetric case, fair solutions can also be obtained. The solutions are then distributed over various possible outcomes, which can all be considered fair in different ways.

In the following, we first derive the game-theoretic subgame-perfect equilibrium for a relatively simple game (with only a single issue and using fairness function 4), and then compare the results with evolutionary outcomes for this game.

	Payoff agent 1	Payoff agent 2
SPE	0.419	0.391
EA	0.391 (±0.022)	0.412 (±0.014)

Table 3.2: Comparison of the agents' payoffs in the EA with SPE results.

Round	Offer (SPE)	Offer (EA)	Threshold (SPE)	Threshold (EA)
1	0.609	0.58 ± 0.06	0.391	0.23 ± 0.21
2	0.375	0.39 ± 0.07	0.250	0.14 ± 0.13
3	0.500	0.48 ± 0.09	0.000	0.13 ± 0.13

Table 3.3: Comparison of the evolved strategies with game-theoretic (SPE) results for each round.

3.4.4 Validation and strategy analysis

Although our incorporation of fairness aspects makes a game-theoretic analysis much more complicated, SPE strategies can again be derived for a very simple version: the game with only a single issue ($m = 1$) and fairness function 4. These settings were chosen because of mathematical feasibility. The general equations are presented in Appendix 2.1. A derivation for $m = 1$, $n = 3$, $p = 1$, and fairness function 4 is given in Appendix 2.2.

Table 3.2 shows both the SPE results and the payoffs obtained by the evolving agents (in the long run) in the a with $m = 1$, $n = 3$, $p = 1$, and with the (rather strong) fairness function 4. Note that since $m = 1$, an agent's payoff equals the share obtained for issue 1. Results for the EA are obtained after 300 generations (averaged over 25 runs). Notice that the SPE payoffs are in good agreement with the outcome of the evolutionary experiments. However, in SPE agent 1's payoff is slightly larger than agent 2's payoff. In the EA this is reversed, although Table 3.2 shows that differences between theory and experiment are very small. We will further analyse the evolving strategies below.

Table 3.3 compares the offers of the evolving agents (for each round) with SPE results, showing a good match. From Table 3.3, it can be derived that agreements are reached in *all* rounds, with some emphasis on the first round.[10]

Table 3.3 also shows the acceptance thresholds (the thresholds are calculated based on the payoff which an agent expects to receive if he rejects the current offer, see Appendix 2). Because the thresholds in rounds 2 and 3 are much lower than the obtained utility, the thresholds in these rounds are not really relevant in SPE. This explains the large variance of the thresholds in the EA and why these thresholds can

[10] Acceptance rates are approximately 39%, 22%, 20% in SPE in rounds 1-3, and 36±4%, 25±3%, $20 \pm 2\%$ for the EA in rounds 1-3.

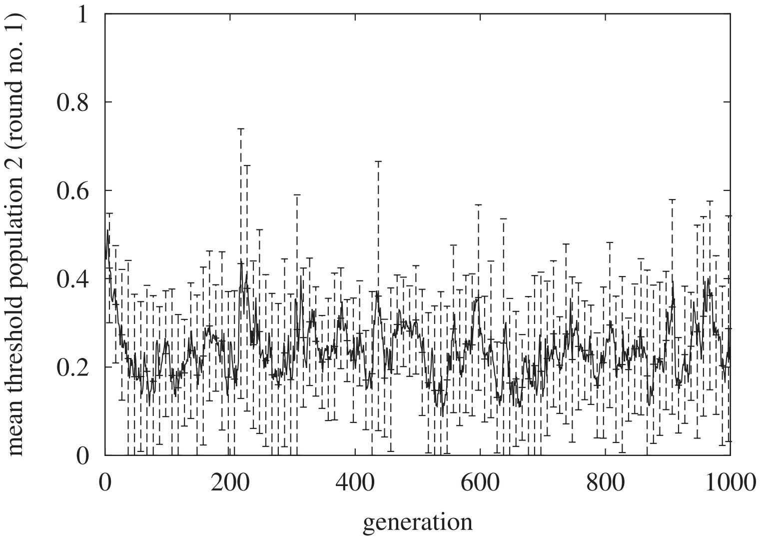

Figure 3.11: Average threshold values of the agent strategies in the EA in the first round.

deviate from SPE predictions in these rounds. In round 1, the threshold is important in SPE and influences the offer made. The experiments show a much lower average threshold value than the SPE (see Table 3.3). Nevertheless, the thresholds influences the offers made in the EA due to a high variance of the threshold values. We analyse this more closely.

Figure 3.11 shows the evolution of the threshold value for the first round for a single experiment. The indicate the variance in the population. Notice that this variance and the volatility of the mean threshold is rather high. This forces the offers in population 1 to be similar as in SPE.

In order to obtain an even better match with SPE results, we reduced the occurrence of frequent peaks by using a decreasing mutation step-size in the EA (instead of self-adaptive mutation step-sizes, see Section 1.2.3). With this approach, the mutation step sizes σ_i are gradually decreased in the course of evolution.[11]

At the beginning of each EA run, σ_i is set to 0.1 for all i (as before, see Table 3.1) and then exponentially decrease until $\sigma_i = 0.01$ after 1000 generations. This procedure indeed reduces the fluctuations in the threshold values and the offers in the long run. Results for experiments with this EA setting appear to be in excellent agreement with SPE results, see Table 3.4. We found no significant effect of the new mutation scheme on the evolutionary outcomes for $m = 2$, however. We suspect that this is due to the integrative nature of the negotiation problem, where the results obtained are already beneficial for both parties.

Conclusion. This relatively simple bargaining situation shows a good match

[11] A similar approach was applied in [3] for a genetic algorithm.

	Payoff agent 1	Payoff agent 2
SPE	0.419	0.391
EA with decreasing σ_i	0.416 ± 0.012	0.395 ± 0.009

Table 3.4: Comparison of the evolutionary agents' payoffs after 1000 generations (using exponentially decreasing mutation step-sizes) with SPE results

between theoretical (SPE) and experimental results. Furthermore, when fairness norms are applied, the outcome of the negotiation process comes to depend on the actual round in which an agreement is finally reached, while thresholds play an important role in some of the rounds. We also showed that EA parameters can be fine-tuned for a more stable situation if needed. This rendered an excellent match with the SPE for $m = 1$.

3.5 Concluding remarks

We have investigated a system for negotiations, in which agents learn effective negotiation strategies using evolutionary algorithms (EAs). Negotiations are governed by a finite-horizon version of the alternating-offers game. Several issues are negotiated simultaneously. Both negotiations with and without a risk of breakdown have been studied. Our approach facilitates the study of cases for which a rigorous mathematical approach is unwieldy or even intractable. We presented computational results for several difficult bargaining problems in this chapter.

We first validated the long-term evolutionary behaviour using the game-theoretic concept of subgame-perfect equilibrium (SPE). When no risk of breakdown exists, the last agent in turn proposes a take-it-or-leave-it deal in the last round and demands most of the surplus for each issue. This extreme division is consistent with SPE predictions. When a risk of breakdown exists, most agreements in the EA are reached in the first round. If the finite game becomes long enough, the deadline is therefore no longer perceived by the evolutionary agents and results actually match SPE predictions for the infinite-horizon game.

We also modelled and studied the concept of "fairness", where a responding agent carries out a fairness check before an agreement is definitely accepted. This fairness check was modelled in two ways: a responding agent considers fairness only at the deadline or all the time, for any potential agreement. In both cases, fair outcomes can be obtained but the outcomes in the second case are much less sensitive to the actual choice of the fairness function. In case of an asymmetric bargaining situation (where the players have asymmetric preferences), multiple outcomes then exist which can be considered "fair" in different ways. We also found a good match between the EA results and game-theoretic SPE predictions for a simple bargaining game (concerning a single issue).

An interesting line of research is to further explore the notion of fairness and to compare the computational outcomes with results from experimental studies with human subjects. Of particular interest is the study of asymmetric multi-issue bargaining situations, where more than one outcome can be considered "fair". This raises several new research questions for experimental economics as well as computational sciences.

Chapter 4

Bargaining with multiple opportunities

In the advent of ubiquitous application of agent technology, bargaining agents are expected to play an essential role in electronic market places. The agents in a competitive market are self-interested and can be equipped with the ability to autonomously search for products and services and negotiate the terms of an agreement. In this chapter,[1] we focus on strategic aspects of bilateral bargaining within a market-like setting.

We use the one-shot ultimatum game as the basic bargaining procedure for our model, a well-known approach within the field of game theory. In this game (see also Section 2.3.2), two players, a proposer and a responder, negotiate about the division of a *bargaining surplus* (see Section 1.1.2). The proposer makes an offer and the responder can only choose to accept or reject this offer. The ultimatum game has been extensively researched, both theoretically and by experiments using human subjects [67, 90, 107].

The ultimatum game models a negotiation between an isolated pair of players. In a market setting, however, an agent's behaviour can change if future opportunities are taken into account. This chapter introduces a natural extension of the basic ultimatum game in which fall-back opportunities are explicitly modelled. Both the proposing and the responding agents have several bargaining opportunities with different opponents before their final payoff is determined. In this way a market place is modelled where several sellers and buyers are available.

The game is further extended to allow several issues to be negotiated simultaneously, as in the previous chapter; not only the price, but also other important

[1] This chapter is based on [38]: E.H. Gerding and J.A. La Poutré. Bargaining with posterior opportunities: An evolutionary social simulation. In M. Gallegati, A. Kirman, and M. Marsili, editors, *The Complex Dynamics of Economic Interaction*, Springer Lecture Notes in Economics and Mathematical Systems (LNEMS), Vol. 531, pages 241–256. Springer-Verlag, 2004.

attributes such as delivery time, package deals, warranty, and other product-related aspects can be taken into account. This can reduce the competitive nature of the game since trade-offs can be made to obtain win-win solutions. Furthermore, we study the impact of search costs if an offer is refused and a new opponent needs to be found. In addition, we consider the case where uncertainty exists about future opportunities and a new opponent cannot always be found.

An important aspect within this setting is the information available to the agents about their opponents. We distinguish between the complete information case, where an agent's current number of future bargaining opportunities is common knowledge, and the incomplete information case, where this information is known to the protagonist but hidden from the opponent. The complete information case can be approached theoretically using game theoretic subgame-perfect equilibrium (see Def. 18.3) given reasonable assumptions. The subgame-perfect results show an extreme split of the surplus, similar to the ultimatum game: the proposer claims the entire surplus, and the responder accepts this deal.

The incomplete information case, on the other hand, seems much more difficult to analyse theoretically. We therefore apply an evolutionary simulation as described in Section 1.2 to investigate this setting. We also compare the evolutionary and the theoretical approach in the complete information case. The evolutionary outcomes show a good match with the game-theoretic results. Moreover, the simulation shows that results differ significantly if information about the opponent's future bargaining opportunities is not available: if the number of bargaining opportunities is sufficiently high, the responder now obtains the largest share.

The outcomes in the incomplete information case, however, also depend on the existence of positive search costs. Search costs stimulate agents to reach agreements early and discourage both players to exploit the additional opportunities. In the evolutionary simulation, the agreements are then similar to the one-shot ultimatum game. A similar effect is observed if bargaining is terminated with a small probability because no new opponent can be found.

This chapter is organised as follows. In Section 4.1 the bargaining game with multiple bargaining opportunities is described. Section 4.2 provides a game-theoretic analysis of the game in case of complete information. Section 4.3 outlines the evolutionary simulation and Section 4.4 discusses the obtained results from the simulation. Lastly, Section 4.5 concludes.

4.1 Description of the bargaining game

The modelled market consists of buyers and sellers who exchange a single good through bilateral negotiations. At each bargaining opportunity, an ultimatum-like game is played, where the proposer makes an offer and the responder can reject or

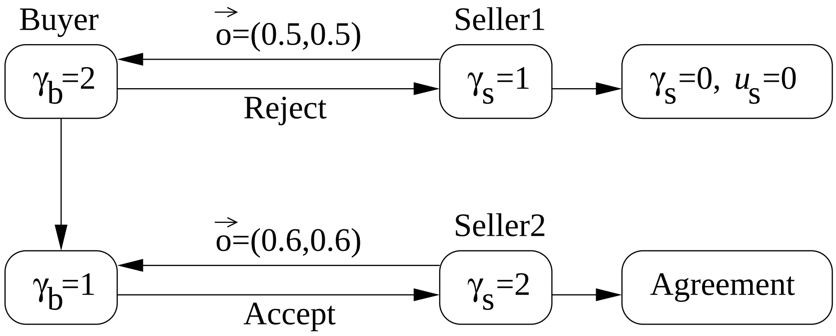

Figure 4.1: A two-issue negotiation example in a market where each agent has two initial bargaining opportunities ($n = 2$).

accept this offer.[2] If an agreement is reached, both agents obtain a payoff equal to their utility of the offer. For convenience, we use *seller* and *buyer* to denote a proposer and responder respectively in the following (although we previously used the terms agent 1 and agent 2, this is not suitable here since several buyers and sellers can participate in a single "market" game).

In our model an offer consists of one or more issues. The utility is calculated as in Chapter 3 (cf. Section 3.1): the seller's utility u_s for an offer $\vec{o}$ can be written as $\vec{w_s} \cdot \vec{o} = \sum_{i=1}^{m} w_s^i \cdot o^i$, where $\vec{w_s}$ is a vector containing the seller's weights for each issue and m is the number of issues. Similarly, the buyer's utility function $u_b = \vec{w_b} \cdot [\vec{1} - \vec{o}]$, where $\vec{w_b}$ represents the buyer's weights. The utilities of the agents are normalised between 0 and 1. The differences in weights of the two players determine the degree of *competitiveness* of the negotiations (i.e., to what extend trade-offs can be beneficial). We formalise the notion of competitiveness and address this issue further in Section 4.4.3.

Each buyer and seller initially has up to n bargaining opportunities to reach an agreement. In case of a disagreement the agents are newly matched with randomly selected opponents, until no more bargaining opportunities remain. The number of remaining bargaining opportunities we call an agent's *bargaining state*, denoted by $\gamma_s \in \{0, 1, \ldots, n\}$ for a seller and $\gamma_b \in \{0, 1, \ldots, n\}$ for a buyer. If an agent's bargaining state reaches zero, the agent obtains a disagreement payoff which is set to zero.

An example for a two-issue negotiation is shown in Figure 4.1 from a buyer's perspective. The buyer, whose initial bargaining state is $\gamma_b = 2$, first encounters a seller, seller 1, with bargaining state $\gamma_s = 1$. The seller proposes an offer $\vec{o} = (0.5, 0.5)$ and the buyer refuses this offer. Because the seller has no more bargaining

[2]Alternatively, the multi-round alternating-offers game (e.g. see chapter 3) can be used. As shown in chapter 3, however, outcomes are equivalent to the ultimatum game, if no time pressure exists; agreements are delayed until in the final round a take-it-or-leave-it offer is made.

opportunities his bargaining game ends and he obtains the disagreement payoff. The buyer, on the other hand, can continue bargaining when matched with another opponent, seller 2. In the example this opponent with $\gamma_s = 2$ offers $(0.6, 0.6)$. The buyer now accepts and the bargaining game ends for both agents.

Note that even though the agents initially have equal bargaining opportunities, the matched agents can have different bargaining states. Having agents with different states is an important aspect of the market game, particularly when agents are unaware of their opponent's remaining opportunities. We assume that, once an offer is rejected, agents cannot go back on a previous offer.[3] We also assume that there are an equal number of buyers and sellers in the market. This in contrast to the work in e.g. [89], where markets are studied with unequal number of buyers and sellers.

4.2 Game-theoretical approach

This section considers the game-theoretic subgame-perfect equilibrium (SPE) of the above game where the agents' bargaining states are common knowledge. A game-theoretical analysis seems to be very difficult if the agents have incomplete information of their opponent's bargaining state. We will, however, drop the complete information assumption in the evolutionary approach (Section 4.3). In the following analysis we assume all agents of a specific type (i.e., buyer or seller) apply the same negotiation strategy. This assumption is reasonable since the preferences are identical for a given type.

In case of a single opportunity, the bargaining game is reduced to the ultimatum game. The ultimatum game has a unique SPE where the seller (here the proposer) claims the total share for each issue, and the buyer (the responder) accepts this take-it-or-leave-it deal [90]. This result can be obtained by applying backward induction. Intuitively, a rational buyer will accept any positive amount, which is always better than obtaining the zero payoff in case of a disagreement. The SPE is precisely the point where the buyer is indifferent between accepting and refusing.

We argue that the game with multiple bargaining opportunities and complete information has an SPE with the same outcome as the ultimatum game: the seller obtains the entire share, and the buyer receives the disagreement payoff, which is set to zero.[4] Consider a buyer with $\gamma_b = 1$, i.e. with a final bargaining opportunity remaining. The buyer will then accept any positive amount offered by the seller. An anticipating seller will then claim the entire share, as in the ultimatum game, independent of γ_s. In SPE, the buyer's payoff for $\gamma_b = 1$ therefore equals zero. Note that this only holds if the seller is informed about the buyer's bargaining state.

[3] Agents are said to have no recall [149].

[4] This holds for continuous divisions of the surplus.

If $\gamma_b = 2$, the buyer has two bargaining opportunities. Using the above, we can replace the payoff for refusing the seller's offer when $\gamma_b = 2$ by the disagreement payoff. The situation for $\gamma_b = 2$ is now equal to $\gamma_b = 1$: the buyer is indifferent between accepting and refusing a value of zero and in SPE the buyer accepts this deal, independent of γ_s. By backward induction the same holds for $\gamma_b = n$.

We note that, because the agents are indifferent to the bargaining state in which the agreement is reached, actually several subgame-perfect equilibria can exist. In all cases, however, the divisions are the same. Note furthermore that the above argument only holds if the seller is informed about the buyer's number of remaining bargaining opportunities. If this information is not available, a game-theoretic analysis seems much more difficult. An evolutionary simulation, however, is very apt to analyse the case of incomplete information. We analyse both the completely informed and the uninformed case in Section 4.4. First, the evolutionary system is described in detail.

4.3 Evolutionary approach

We use an evolutionary algorithm to evolve the strategies of the agents. The evolutionary simulation is depicted in Figure 4.2. The evolutionary algorithm is based on the implementation described in Section 1.2.3. As in Chapter 3, each strategy in the EA corresponds to an agent of a certain type (buyer or seller), and we use separate populations to evolve the strategies of the two types of agents. The way in which the fitness of the agents is determined, however, differs from the approach described in Chapter 3. In the previous model, each agent was evaluated against all agents in the opponent's population. In this case, however, all agents together constitute a market-like setting, where buyers and sellers can bargain several times with different opponents before their final fitness is determined. Also because the interactions determine the bargaining states of the agents, another approach is required here.

The fitness of the agents is determined as follows. The parental and offspring populations are first combined to form a group of sellers and a group of buyers. The agents are then evaluated by a sequence of pair-wise matches. For each match, two agents are randomly selected (with replacement) and play the one-shot game. An agent obtains a payoff in case an agreement is reached or the disagreement payoff (which is zero) if no more opportunities are available for this agent. If an agent still has opportunities remaining, his fitness remains undetermined. Note that, since both agents can be in different bargaining states, the consequences of a disagreement may be different for the individual agents. Because an outcome depends on many random factors, each strategy is evaluated a number of times and the fitness is the average of r payoff values. The parameter r is called the *evaluation frequency.* This way the fitness becomes a more accurate measure of the expected payoff. The bargaining games continue until all agents have obtained at least r payoff values.

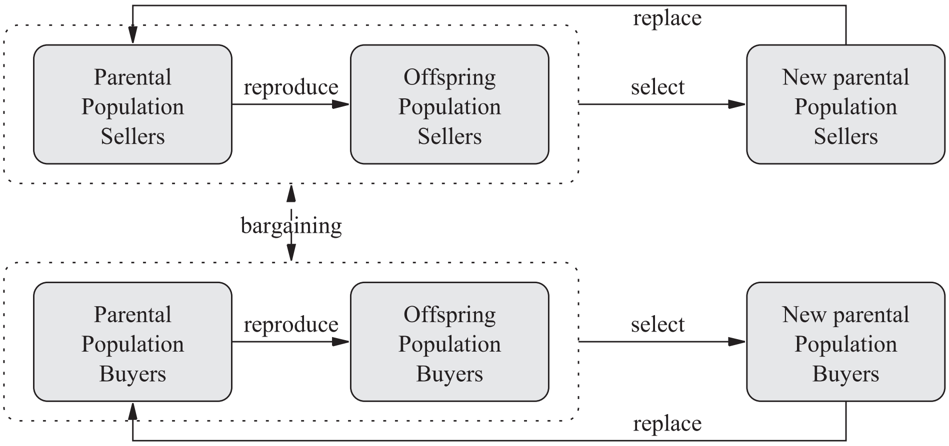

Figure 4.2: Iteration loop of the evolutionary algorithm.

Since both buyers and sellers start with the same bargaining state, in the first periods the opponent's bargaining states do not represent an ongoing bargaining society. To prevent so-called initiatory effects and to model an on-going bargaining society, a strategy's fitness is only measured after the first payoff is determined. A strategy is thus evaluated at least $r + 1$ times. Furthermore, we model a market situation where the number of agents remains constant over time, also called a steady-state market in [89]. Therefore, once the fitness of a strategy has been established, the strategy can still be selected to play again but its fitness is no longer affected by the outcome. The bargaining games are continued until the fitness for each strategy has been established.

4.3.1 Strategy Encoding

The strategy, encoded on the chromosome, specifies either an offer or a threshold for each bargaining state, depending on the type of the agent (i.e., sellers only have offers and buyers only have thresholds). The threshold determines whether an offer of the opponent is accepted or rejected: if the utility falls below the threshold the offer is refused; otherwise an agreement is reached. A similar representation was used in Chapter 3 for the alternating-offers game, although in that game all strategies contain both offers and thresholds.

We distinguish between the complete information setting and the incomplete information setting (see Section 4.1). The strategy representation depends on this setting and is schematically depicted in Figures 4.3 and 4.4 for the complete and incomplete information case respectively. In the incomplete information case (Figure 4.4), an offer or threshold is specified for each bargaining states of the agent. In case of complete information (Figure 4.3), an offer or threshold is also conditional on the opponent's bargaining state.

Seller Strategy	$\vec{o}(1\|1)$	$\vec{o}(2\|1)$	...	$\vec{o}(n\|1)$
	$\vec{o}(1\|2)$	$\vec{o}(2\|2)$	...	$\vec{o}(n\|2)$
	...	...	...	...
	$\vec{o}(1\|n)$	$\vec{o}(2\|n)$	...	$\vec{o}(n\|n)$
Buyer Strategy	$t(1\|1)$	$t(2\|1)$	...	$t(n\|1)$
	$t(1\|2)$	$t(2\|2)$	...	$t(n\|2)$
	...	...	...	...
	$t(1\|n)$	$t(2\|n)$	...	$t(n\|n)$

Figure 4.3: The strategies of a seller and a buyer for the market game with complete information about the opponent's bargaining state. The offers $\vec{o}(\gamma_s|\gamma_b)$ and thresholds $t(\gamma_b|\gamma_s)$ are conditional on the bargaining state of the opponent, where $\gamma_s, \gamma_b \in \{1, \ldots, n\}$.

Seller Strategy	$\vec{o}(1)$	$\vec{o}(2)$	...	$\vec{o}(n)$
Buyer Strategy	$t(1)$	$t(2)$	...	$t(n)$

Figure 4.4: The strategies of a seller and a buyer for the market game, where the players are uninformed about the opponent's bargaining state. An offer $\vec{o}(\gamma_s)$ or threshold $t(\gamma_b)$ is only determined by an agent's own bargaining state, since more information is not available.

4.3.2 Mutation Operator

Although several mutation models were tried, the mutation model with exponential decay showed a closest match with game-theoretic benchmark cases. We therefore only report the results using the exponential decay model. This mutation operator is explained in Section 1.2.3.

4.4 Evolutionary simulation results

The results are organised as follows. First, the game with complete information is studied in Subsection 4.4.1 and the results are compared to the game-theoretic (SPE) predictions. Subsection 4.4.2 studies the incomplete information case. Subsection 4.4.3 introduces a measure of competitiveness for multi-issue negotiations and compares results for different levels of integrative negotiations. Finally, in Subsection 4.4.4 considers the effects of fixed search costs in the market game and

Parental population size (μ)	30
Offspring population size (λ)	30
Initial standard deviations (σ)	0.1
Mutation model	exponential decay
Standard deviation half-life (t)	400
Number of generations	4000
Number of runs per experiment	30
Strategy evaluation frequency (r)	20

Table 4.1: Default settings of the evolutionary simulation.

uncertainty about future opportunities.

4.4.1 Game-Theoretic Validation

This section considers a competitive (i.e., single-issue) scenario with complete information of the agents' bargaining opportunities and compares the evolutionary algorithm (EA) outcomes to SPE predictions. Default parameter settings for the EA are shown in Table 4.1. Note that because of random fluctuations, the EA results are averaged over 30 runs using the same settings.

In SPE the share of the buyers is zero and the sellers obtain the whole surplus in case the initial number of bargaining opportunities of the players is finite, and the bargaining state of the opponent is common knowledge (see also Section 4.2). Figure 4.5 shows the EA outcomes for different values of n (initial bargaining opportunities). The results indicate an almost perfect match between evolutionary outcomes after 4000 generations and game-theoretic outcomes, particularly when n is small.

For larger values of n we find that, using the same EA parameter settings, the evolutionary outcomes become somewhat less extreme. See also Figure 4.6, which shows the long-term EA outcomes (after 4000 generations) for n up to 10. This is because as n becomes larger, the complexity of the problem increases due to a larger search space, making learning by an EA more difficult. However, a better match for larger values of n also appears by adjusting EA parameters, such as the evaluation frequency and the population size, to handle the increased complexity. Details on tuning the EA are not treated here. Instead, we refer the interested reader to previous research [126], in which different EA settings are systematically studied for an alternating-offers bargaining game. Henceforth, we present only experiments using uniform EA settings here.

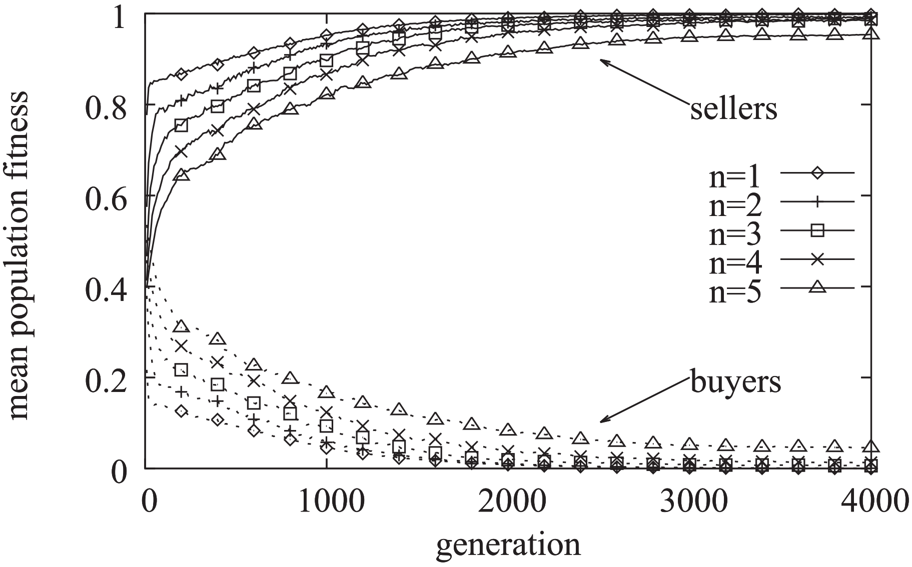

Figure 4.5: Development of the mean fitness (averaged over 30 runs) for complete information setting with varying initial number of bargaining opportunities.

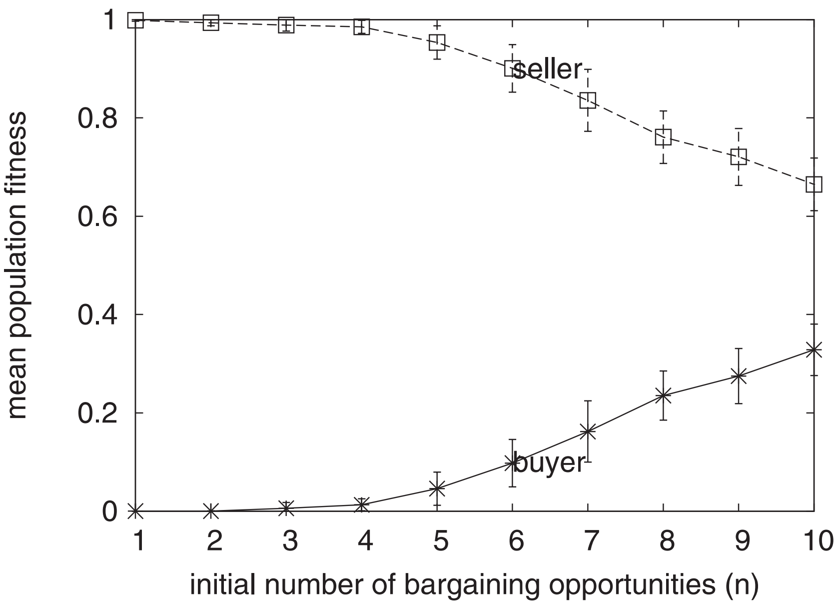

Figure 4.6: Results after 4000 generations (averaged over 30 runs) in case of complete information.

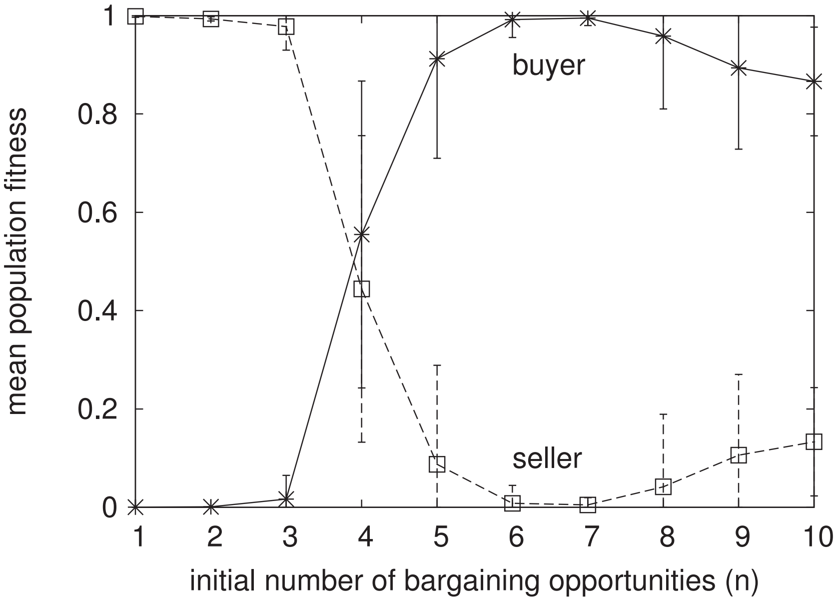

Figure 4.7: Results after 4000 generations (averaged over 30 runs) for incomplete information settings with various n. The error bars indicate the standard deviation of the averaged results.

4.4.2 Incomplete Information

We now examine the results when the agents do not know their opponent's bargaining states; the agents only know their own bargaining states. Although no explicit information is available, the agents implicitly learn the distribution of the bargaining states in the opponent's population. This distribution is endogenously determined by the strategies of the agents. The strategies, in turn, adapt to the distribution of the bargaining states. This complex interaction is one reason why theoretical analysis is difficult. An EA, on the other hand, is well suited to find outcomes that emerge from such local interactions.

Results produced after 4000 generations of the EA for the incomplete information case are shown in Figure 4.7, for different values of n (the initial number of bargaining opportunities). These results are averaged over 30 runs. The error bars indicate the standard deviation. Whereas in the complete information case the seller obtains almost the entire surplus, the responder (i.e., buyer) has the best bargaining position in the incomplete information case (see Figure 4.7). This holds as long as the initial number of bargaining opportunities are sufficiently large (i.e., ≥ 5). Note that these results are obtained even though the buyers' and sellers' initial settings are equal.

The results can be explained as follows. If the buyer is in her final state, she will accept any deal (as in the ultimatum game). In other states, however, the buyer can try to find a better deal elsewhere. Consider a seller in his last bargaining state. Because he does not know the buyer's bargaining state, he can no longer anticipate the buyer's behaviour. In order to prevent a disagreement, the sellers will then concede in the last bargaining state. The expected payoff in case of a disagreement and the offers in earlier bargaining states will then also decrease. After many generations, the simulation converges to an outcome where the seller concedes almost his entire surplus in each bargaining state. We also observe that the seller concedes slightly less if he has more bargaining opportunities remaining, resulting in less extreme deals if n becomes large, as shown in Figure 4.7.

In the incomplete information setting the first-mover (here the seller), has no information about his opponent. The responder, on the other hand, can make a relatively more informative decision based on the seller's offer. Whereas in the ultimatum game the proposer seems to dominate the outcome, a more competitive setting allows the responder to obtain a considerable advantage. This result, however, holds only if the number of bargaining opportunities is finite and equal for both players. Furthermore, the players incur no costs for refusing a deal. As we will show in Section 4.4.4, even slight costs completely change these results.

When the number of initial bargaining opportunities is set higher than three, a sudden transition in the long-term outcomes can be observed in Figure 4.7: up to $n = 3$, the seller obtains almost all, whereas the buyer obtains the largest share if $n > 3$. By increasing n, the number of possible states increases, making the buyer's behaviour less predictable for the seller. The value for which the transition occurs depends on game parameters such as r and the competitiveness of the negotiation. The latter will be discussed further in the next section.

4.4.3 Integrative Negotiations

An advantage of bilateral negotiation is the ability to negotiate complex contracts with several issues. When mutually beneficial solutions are available, negotiations are called *integrative* (see Section 2.3.3). We consider integrative two-issue negotiations in this section and introduce the notion of competitiveness. We show that the information in the integrative case has a very similar impact as in the competitive case. Due to increased complexity, however, the evolutionary results are less extreme when the number of bargaining opportunities is large.

The utility of an offer is an additive, weighted function of the share obtained for each issue (see also Section 4.1). The weights for sellers and buyers for the two issues are $\vec{w_s} = (0.5 - \alpha, 0.5 + \alpha)$ and $\vec{w_b} = (0.5 + \alpha, 0.5 - \alpha)^T$ respectively, where $\alpha \in [0.0, 0.5]$ is the so-called degree of *competitiveness*. When the parameter α is set equal to 0, negotiations are purely competitive; if $\alpha = 0.5$ there is no competition at all. Note that the maximum social welfare, i.e. the maximum total utility that a seller and a buyer can achieve together equals $2 \cdot (0.5 + \alpha)$, where each agent obtains $(0.5 + \alpha)$.

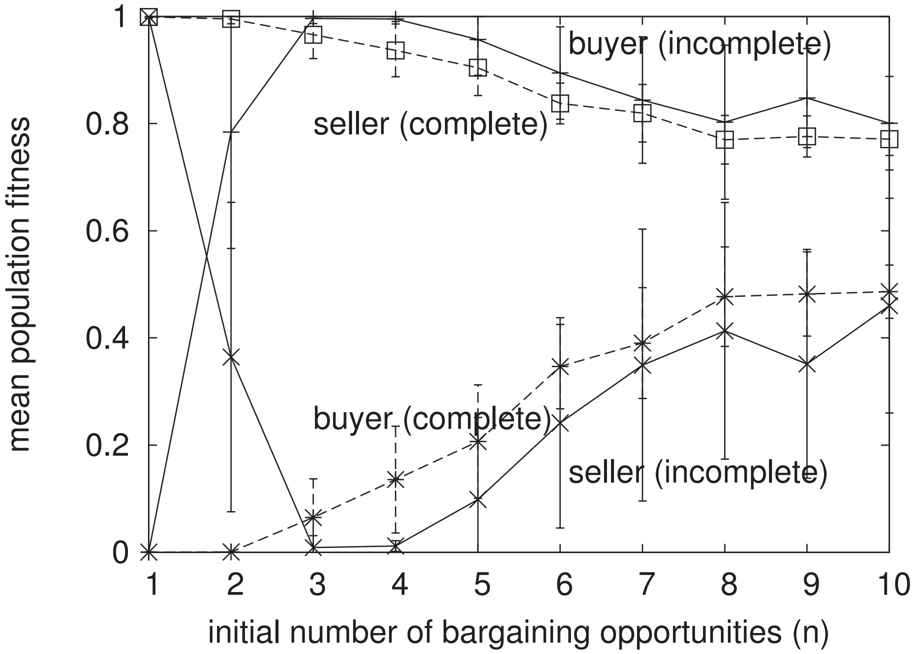

Figure 4.8: Mean long-term outcomes for two-issue negotiations and $\alpha = 0.2$.

Results for $\alpha = 0.2$ are visualised in Figure 4.8. The results show that, as in the competitive case, a transition occurs to a buyer-dominated outcome for sufficiently large n and incomplete information. We find, however, that this transition already occurs when $n = 2$ (see Figure 4.8). Only two bargaining opportunities are needed to obtain an advantage for the responder, as supposed to four in the single-issue game (Figure 4.7).

Figure 4.8 also shows a less extreme split compared to competitive negotiations, particularly for large n. This occurs firstly since the strategy search space is increased (a value for each issue needs to be learned), making learning more difficult. Moreover, the win-win possibilities are fully exploited: if one of the agents slightly concedes, the other agent can obtain a relatively large gain by negotiating a Pareto-efficient deal. As shown in Figure 4.9, this effect becomes stronger as α increases. In the extreme case, where $\alpha = 0.5$, both agents can obtain the full surplus without any concession.

Note that the EA parameters are fixed for the various game settings. As we mentioned in Section 4.4.1 we can adjust the parameters to handle more complex bargaining settings as a result of a larger n and an increased number of issues. By increasing the population size and adjusting other parameters of the EA, we obtain results which are closer to game-theoretic predictions.

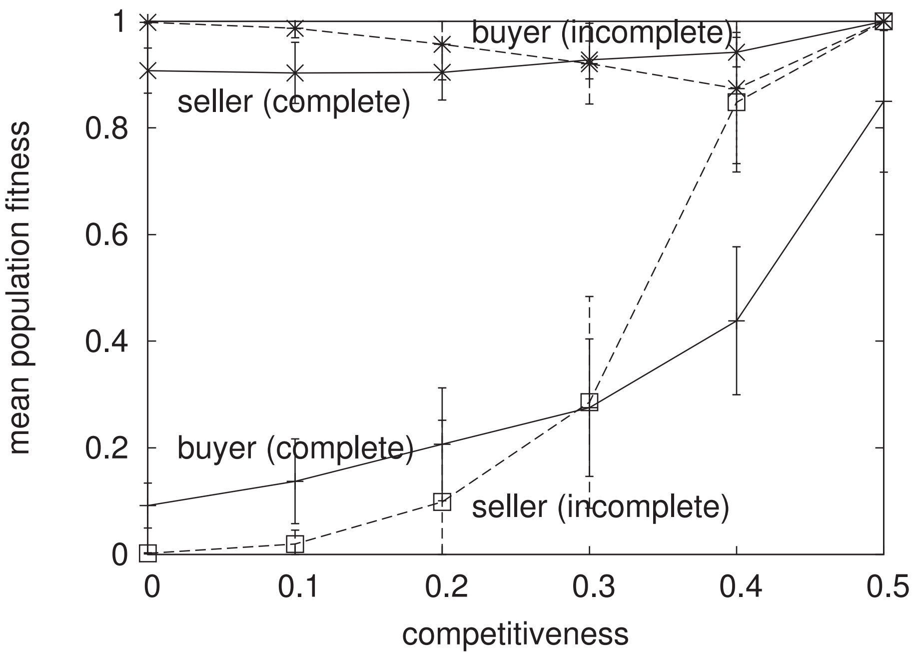

Figure 4.9: Mean long-term outcomes for $n = 5$ and different values for the competitiveness (α).

4.4.4 Search Costs and Premature Termination

We further extend the bargaining game in two ways. First, we introduce search or negotiation costs each time an offer is refused and agents engage in a new negotiation. Subsequently, we consider the case where there exists uncertainty about whether a new bargaining opponent can be found. Whereas we have assumed until now that the number of bargaining opportunities remains fixed, there can be external factors which influence the number of opportunities (e.g., if a seller has in the meanwhile sold the good to another buyer). This is modelled as a probability that negotiations terminate prematurely, i.e., before the final number of bargaining opportunities is completely exhausted.

Search costs can represent the amount of money, time, or effort that an agent may incur for finding a new opponent. It is shown theoretically that if buyers have search costs, the sellers charge monopolistic prices in equilibrium [22, Ch.7]. We consider the impact of search costs on the bargaining game where both buyers and sellers have equal search costs β. The final utility is reduced by fixed search costs β for each new bargaining opportunity. Only the first bargaining opportunity has no costs.

Evolutionary outcomes for the complete and incomplete information settings with different search costs are depicted in Fig 4.10. Negotiations are competitive and buyers and sellers each have 5 initial bargaining opportunities. Search costs seem

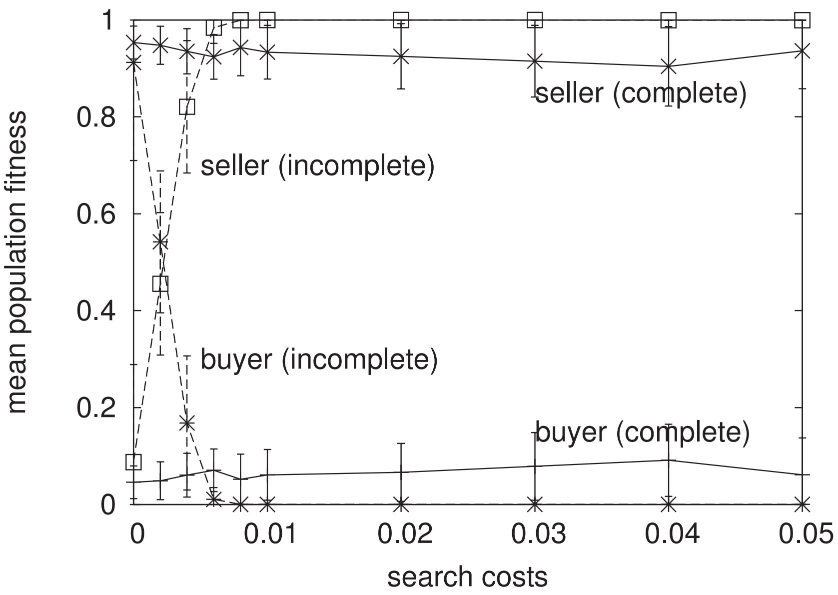

Figure 4.10: Mean long-term results as a function of the search costs (β) for $n = 5$.

to have little impact on the fitness in the complete information case; variations are not statistically relevant. Although the fitness does not change, the actual behaviour of the agents does: most agreements are now reached immediately. Without search costs, agreements reached are distributed over the various bargaining states.

In the incomplete information case, on the other hand, even small search costs have a drastic impact on the fitness of the agents, see Figure 4.10. The sellers claim almost the entire share even if search costs are very small (e.g. 0.01) and equal for both agents. Results are robust for different settings of the EA. These outcomes are consistent with economic theory, which states that prices become monopolistic even if search costs are infinitely small.

As in the complete information case, both buyers and sellers are stimulated to reach agreements early in case of search costs. The final opportunity of the seller is therefore almost never reached, removing the advantage for the buyer. The game changes from a game with incomplete information, to a game where almost all players complete a deal in their first bargaining opportunity. Now the seller can again claim the entire surplus as in the one-shot game.

Similar outcomes are observed when bargaining for a buyer and/or a seller is discontinued with a certain probability after each disagreement.[5] Figure 4.11 shows the long-term outcomes for different probabilities of premature termination after each bargaining opportunity. The probability is set equal for buyers and sellers, and

[5] This is analogous to discount factors or a probability of break down in case of multi-round bargaining, as used in e.g. Chapter 3.

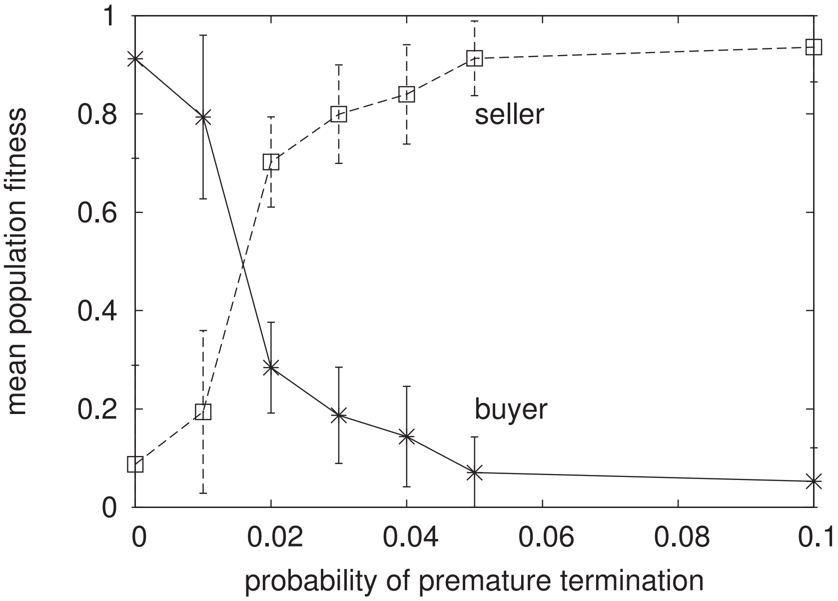

Figure 4.11: Long-term fitness values for $n = 5$ and incomplete information, when negotiations are discontinued with a certain probability after each disagreement.

for each bargaining opportunity, but drawn independently. As with search costs, the seller obtains the largest share if the probability is sufficiently high.

Note that the effect of premature termination is less extreme, however. This is because search costs also affect the utility if an agreement is *not* reached, providing an additional incentive to reach agreements (otherwise, a negative utility is obtained). In case of premature termination, on the other hand, an agent is indifferent between termination after the first bargaining opportunity and a disagreement in the last bargaining opportunity.

4.5 Concluding remarks

We study the evolutionary dynamics of a market-like game in this chapter, where a seller sells a single good and has several opportunities to do so. At the same time, a buyer wishes to buy an item by trying several sellers. The terms of an agreement are negotiated using an ultimatum-like game, where the seller proposes an offer and the buyer can choose to accept or reject the offer. The game is extended to allow for multiple opportunities for both the seller and the buyer if the deal is rejected. This way a competitive market is modelled. We furthermore investigate multi-issue integrative negotiations and the effects of search costs and premature termination if a disagreement occurs.

The game-theoretic outcome using subgame-perfect equilibrium (SPE) for the one-shot ultimatum game predicts an extreme split of the surplus: the seller obtains the whole surplus whereas the buyer obtains her disagreement payoff. We extend the analysis for multiple bargaining opportunities with complete information of the opponent's bargaining state and find an equivalent outcome. A theoretical analysis seems to be very difficult, however, if the bargaining states of the agents are not common knowledge. An evolutionary simulation, on the other hand, is very well suited to investigate such games with incomplete information.

We first compare the evolutionary results with the game-theoretical outcomes for the game with complete information to validate the evolutionary approach. If the initial number of bargaining opportunities is small, a very good match is found. In larger games or when the negotiations become less competitive, the EA shows somewhat deviating outcomes due to larger search space and the limited computational capacity of the EA. We note that we mainly report experiments using uniform EA settings in this paper. However, adjusting EA settings appear to improve results even further for more complex games.

The evolutionary simulation shows a large impact of the additional bargaining opportunities if the agents have no information on their opponent's number of future opportunities. Whereas in the complete information game the seller dominates the market, the buyer is better off in the incomplete information setting, as long as the number of bargaining opportunities is sufficiently high. By increasing the initial number of bargaining opportunities a sudden transition is observed where the buyer obtains the largest share instead of the seller. This occurs because the seller can then no longer anticipate the buyer's response and gives in to avoid a disagreement.

Similar outcomes are found for two-issue integrative negotiations. At the same time, integrative negotiations produce less extreme evolutionary outcomes, both in the game with complete and incomplete information, particularly if the number of initial bargaining opportunities is large. This mainly occurs since the space of possible deals increases. Moreover, the agents find win-win situations which benefit one agent without affecting the payoff obtained by the opponent.

An integrative setting also already affects small games with incomplete information: we find that for certain settings, a transition from a seller to a buyer dominated payoff occurs even in case both agents merely have two initial bargaining opportunities, whereas in the competitive case more bargaining opportunities are needed to achieve the same result.

We also study the effect of search or negotiation costs in case a negotiation fails and the agent needs to find a new opponent. Search costs induce players to reach an agreement in the very first bargaining opportunity. This changes an incomplete information game into an ultimatum-like game with only a single bargaining opportunity. Even very small search costs result in an extreme split where the seller obtains almost the entire share, similar to the ultimatum game outcome. This is consistent with economic theory which states that even infinitely small search costs produce monopolistic prices. The outcomes are similar but less extreme if search costs are replaced by a probability that bargaining is discontinued after a disagreement. This models the situation where uncertainty exists about future opportunities.

In this chapter we have shown that evolutionary simulations are extremely useful to investigate negotiations with incomplete information, which are unwieldy to analyse theoretically. Using evolutionary algorithms, we can simulate complex interactions involving a large number of agents, as is the case in bargaining with multiple opportunities. It is interesting to further refine the model to specific real-world settings, where for instance agents have incomplete information about their own future number of bargaining opportunities. Another interesting extension is allowing agents to return to previously encountered opponents.

Part B

Bargaining systems for business applications

Chapter 5

Competitive market-based allocation of consumer attention space

In this chapter,[1] we consider an e-business application of automated negotiation using software agents. We present a framework for a distributed Competitive Attention-space System, CASy, to allocate the scarce resource that is consumer attention via the techniques of dynamic market-based control [20, 23, 43] and adaptive software agents (see see Section 1.1.3 and [47, 60, 144]). In the example of an electronic shopping mall, CASy recommends shops to a consumer: the task of matching a consumer to a set of suitable shops is delegated to the individual shops, each of which evaluates the information that is available about the consumer and his or her interests (the consumer's interests and other information which the consumer is willing to provide; e.g. keywords, product queries, and available parts of a profile). Based on this information and on their domain knowledge, shops can make a monetary bid in an auction where a limited amount of consumer attention space, or banners, for the particular consumer is sold.

To facilitate CASy, the system is designed as a multi-agent system (see Section 1.1.3) where each shop is represented by a software agent that executes the task of bidding for the attention of each individual consumer. The use of learning software agents allows shops to rapidly adapt their bidding strategy such that they only bid for consumers that are likely to be interested in their offerings. Further-

[1]The results of this chapter have been published in [17]: S.M. Bohte, E.H. Gerding, and J.A. La Poutré. Market-based recommendation: Agents that compete for consumer attention. *ACM Transactions on Internet Technology*, Special Issue on Machine Learning in the Internet, August 2004 (to appear). A shorter version appeared as [16]: S. M. Bohte, E. H. Gerding, and H. La Poutré. Competitive market-based allocation of consumer attention space. In M. Wellman, editor, *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC-01)*, pages 202–206. The ACM Press, 2001.

more, efficient bidding for each customer is only feasible when automated: hence the use of software agents. These agents allow a shop to process a large number of small transactions, and enable them to make a deliberated bid for every customer entering the shopping mall.

In CASy, shops react to consumer behaviour and to behaviours of other shops, yielding various interdependencies in the commercial effects related to being displayed together with competitors. For various basic and simple models for on-line consumers, shops, and profiles, we demonstrate the feasibility of our system, i.e., that proper matchings of consumers with shops are achieved, and that shops can learn their niche in the market, even in the case of such interdependencies. Especially, to validate the economical concept of the market mechanism underlying CASy, we develop an evolutionary system for bidding supplier agents. In this approach, the agent system is investigated like an (adaptive) economic market, as in agent-based computational economics (ACE) (see also Section 1.2, and Chapters 3 and 4).

Furthermore, in this chapter we reflect on the merits of the system, and assess the advantages and issues that need further attention, from both the technological and the economical point of view. In [17] we extend this work and also develop adaptive software agents that learn bidding strategies based on neural networks and strategy exploration heuristics.

We note that the mechanism we describe is not limited to the example of the electronic shopping mall, but can easily be extended to other domains where (pre) selection of possibilities has to be guided, like banners on more general websites, attention spaces on mobile devices, or other types of marketplaces.

This chapter is organised as follows. First, Section 5.1 motivates the decentralised, agent-based approach for allocating attention space, and discusses related approaches. In Section 5.2, the design of CASy is presented. The evolutionary simulation is explained in Section 5.3, whereas Section 5.4 contains the results. Section 5.5 reflects on practical implementation issues such as privacy and the communication overhead of the mechanism. Finally, Section 5.6 concludes.

5.1 Motivation and related research

Before describing CASy in more detail, we first elaborate on the merits of such a system, and the motivation for using software agents. Also, we discuss related work. In Section 5.1.1 we compare the decentralised approach with the more commonly used centralised approach. In Section 5.1.2 we comment on the use of software agents. Section 5.1.3 gives an overview of related work.

5.1.1 Centralised vs. decentralised recommendation

With the advent of electronic marketplaces, scale limitations as encountered in the brick-and-mortar world no longer apply: the supply side of the market is no longer restricted by geographical considerations or lack of physical (shelf) space. At the same time, novel problems are encountered, like how consumers can find their way in a large marketplace where very many suppliers offer their products.

To this end, a mechanism provided by a trusted third party is desired to propose relevant shops and products to a consumer in e.g. a virtual shopping mall. A central filtering scheme is a feasible solution for several different business areas. For such an approach, knowledge on both the user and of the shops, as well as knowledge on the product domain needs to be stored in a central location in order to determine appropriate matches. This approach is used in recommender systems like *Amazon* and *eBay* [114] to recommend goods on specific domains such as books and CD's, and in shopbots or pricebots [46], as for instance BargainFinder [66]. Keyword profiling is also a popular method for ranking online sites in search engines. This amounts to contracts for charging monetary amounts for increased visibility, given specific keyword entries, e.g. [52–54].

A central or personal filtering system works well in the case of suitable and well-demarcated domains, as for instance for a book and music store. However, for a large heterogeneous marketplace with many participating shops and consumers, several complexity difficulties arise. This is due to the amount of relevant information that has to be tracked and processed by the filtering mechanism in the form of relevant up-to-date knowledge of e.g.: the consumer's interest in different product domains and shop categories; the shops' products, ways of doing business, and business interest; and ontologies and domain knowledge for various product categories. Also, the weighing of multiple issues like service, quality, price, and product diversity (add-ons and customisation of products) can be important.

Besides the computational complexity problems for information processing, this requires the transfer of business information of shops towards the central system as a trusted third party. Such a practice encounters many objections by businesses, even if only product catalogues are concerned [78, 135]. In addition, a central mechanism still needs to make decisions about what to display in which order to a consumer, in a way that is reasonable to all parties: all the suppliers and consumers. A fair and general ' of interests (utilities) of different market parties is usually not possible, however, and concepts like Pareto-efficiency (see Def. 4.3) are used instead.

Thus, central filtering mechanisms may suffer from increasing (computational) complexity as well as serious objections and obstructions from commercial parties in various sorts of business areas.

5.1.2 Use of adaptive software agents

We believe that the system as presented is the natural evolution of auction-based allocation systems like those currently employed by internet companies like Google (for sponsored keywords, [52]) and Overture (for banner targeting, [53]). Whereas these pre-cursor systems rely on the human factor to set essentially static prices for particular goods, the use of software agents in our system in principle allows a market-party to assess the value of each individual prospect, if desired at a very detailed level, as well as take into account real-time business-related domain knowledge and strategies. The implementation of adaptivity into the software agents allows the "market" for consumer attention to function more efficiently, where the targeting of potential prospects can be more precise, and changing buyer behaviour can be tracked and followed. As such, agent-assisted recommendation in competitive markets represents the next logical step for current auction-based allocation systems.

5.1.3 Related research

Our work relates to the large body of research concerning market-based control [20, 23, 43]. This paradigm is essentially about controlling complex systems using a (distributed) market mechanisms for allocating scarce resources. A large number of applications exist such as the allocation of computational resources [23, 43], load balancing and climate control [145]. Our work applies the paradigm of market-based control to generating recommendations in a distributed fashion using software agents.

Related to our approach for generating recommendations is a prototype called MATE [91] (Multi-Agent Trading Environment) that performs market-matching using agent technology. In [91], merchant agents receive the profile of the consumer, and each suggests one or more products to a personal consumer agent. The personal consumer agent then filters the appropriate products and ranks the remaining products according to the customer's preferences. In this approach, selection is done on the consumer side, and significant knowledge on a product domain should be incorporated in the personal consumer agent, being a task of a central party to provide.

A more recent approach by Wei et al. [142, 143] has a number of characteristics similar to CASy; they also apply a central auctioneer to shortlist the recommendations based on bids made by information providers (called recommending agents). In their approach, a reward agent determines the reward or feedback for the recommending agents based on the quality of the recommendations as perceived by the user. The rule used to calculate the reward is shown to be Pareto-efficient (i.e., maximise the social welfare) [142]. Based on this feedback, the bidding (recommending) agents update their strategy using heuristic rules. The bidding strategy proposed here, on the other hand, is more general and adapted by machine learning algorithms.[2] Also, the feedback is directly obtained via the consumers, and it is up to the supplier agent to determine the value of this feedback.

5.2 The design of CASy

In this section, we present the framework of CASy (Competitive Attention-space System) for matching consumers with relevant suppliers in the case of an electronic shopping mall. We note that the framework we describe is not limited to the example of the electronic shopping mall, but can easily be extended to other domains where (pre) selection of possibilities has to be guided, like banners on more general websites, attention spaces on mobile devices, or other types of marketplaces. Instead of addressing to the case of "shops" only, we henceforth mainly use the more general term "supplier" to refer to the suppliers of goods or services.

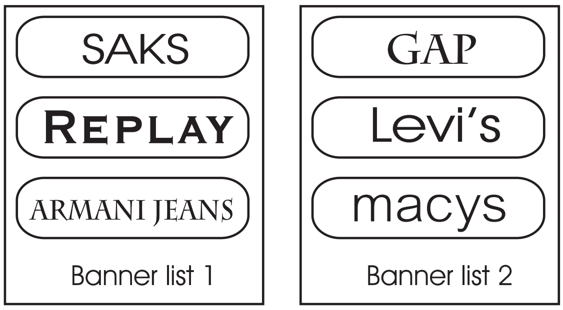

Figure 5.1: Advertisements are shown in the form of banners. The banner list is tailored towards a consumer's characteristics.

When a consumer enters a shopping mall, he[3] expresses his interest for certain products and selects the business sector of his interest. The information about his interest, possibly augmented by additional knowledge, is passed on to potential suppliers in the sector. The suppliers subsequently compete against each other in an auction by placing bids to "purchase" one of a limited number of entries of attention space for this specific consumer. Finally, the consumer is shown the list of winning suppliers, using for instance banner advertisements. An example is depicted in figure 5.1.

Software agents (see Section 1.1.3)) are used to facilitate the fine grain of interaction, bidding, and selection in CASy. For our mechanism, we have software agents for the suppliers and for the enabling intermediary: the mall manager. The model

[2]We discuss results using evolutionary algorithms in this chapter. For an approach using neural networks, see [17].

[3]"he" stands for "he or she."

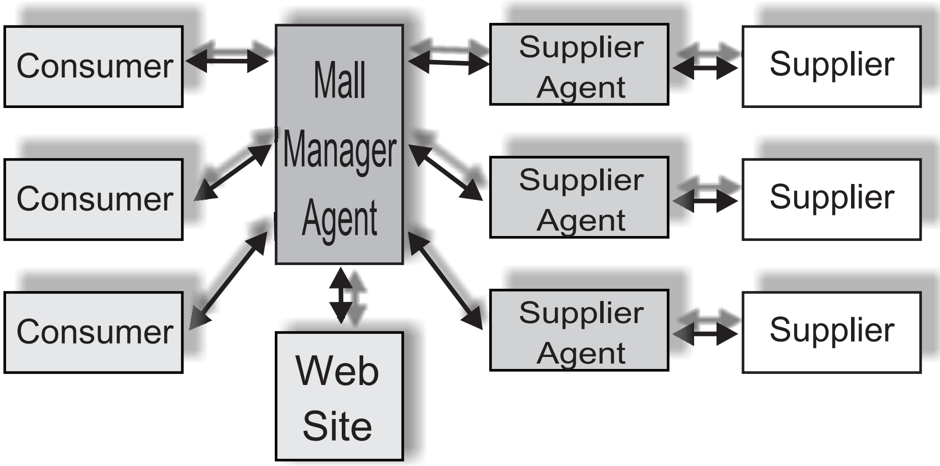

Figure 5.2: Components of the shopping mall and their interactions.

of the electronic shopping mall is depicted in figure 5.2, showing both the software agents and the actual economic players in the shopping mall: the consumers and the suppliers. The participants within the shopping mall and their roles are discussed in more detail in the sections that follow.

5.2.1 Mall manager agent

The Mall Manager Agent (MMA) acts as an intermediary between consumers and supplier agents. The task of the MMA is to facilitate bidding and information dissemination processes by providing the auctions and additional customer profiling services to the suppliers. Given privacy concerns, the consumer profile will not automatically be communicated in full to the suppliers, as e.g. described in Subsection 5.2.2. Information on the consumers could be stored within the MMA for revisiting consumers, leaving open consumers who wish to remain anonymous. The MMA applies the auction: it collects the bids of the supplier agents, selects the winners, charges the selected suppliers, and enables their display. In Section 5.2.4 we address the auction in more detail.

5.2.2 Consumers

In the model of figure 5.2, the consumer directly communicates its interest and preferences to the MMA, e.g. via a web page. Note, however, that the assistance of a personal software agent for the consumer is conceivable. Preferences include the product that is being searched after and various values for the attributes of the product. The MMA can also consider information on a consumer's profile. The consumer profile consists of more generic information on the consumer. This could include regular personal information like general interests, previous acquisitions, as well as age or zip code; but also general sales-related information like style or the interest in issues as price, quality, and service. The consumer can either be queried directly for this information, or the MMA can derive the information from previous interactions. The consumer can restrict or disable the dissemination of his profile information. E.g., distribution of such information can be limited to for specific or anonymised parts, or to general sales-related information that is derived from the private profile.

5.2.3 Suppliers and supplier agents

Each supplier "owns" an agent that acts on the supplier's behalf, called a *supplier agent*. The main task of a supplier agent is to effectively purchase attention space. The agent will do this by bidding on attention spaces that are to be displayed to consumers it deems interesting, thus maximizing the supplier's profits. To this end, it has to evaluate (information about) consumers. The valuation of a consumer by a supplier agent is closely linked to its bidding strategy: the bid should not outweigh the expected profit (if the supplier is to break even) or percentage thereof. This task can be complicated: the variety of consumers can be great, and the competitive environment can change rapidly. Also, the supplier's conception of the targeted audience may deviate from its actual audience.

The agent can learn this targeting by for instance using the push-back information from individual customers, e.g. the knowledge whether or not its advertisement was selected by the customer (click-through), subsequent buying actions, or, to be provided by the mall manager, (selected) click-stream information (e.g. time spent on pages, mouse actions). Additionally, the agent can use supplier-specific knowledge and (adaptive) rules for accurate targeting.

Along with a strategy for bidding on customers, a supplier agent is also equipped with knowledge about the supplier. Such knowledge can contain amongst others relevant business information on the supplier that is needed for the matching process. This information should determine the supplier's conception of its "niche" in the market, and hence the type of preferred consumer. Typical business information could be the products carried and the intended audience. Furthermore, the goals and limitations of the supplier can be taken into account, such as the current quantity of a certain product in stock or the service level.

5.2.4 Auctions

In this Section we address the auctions protocol and the payment procedure of the MMA. A payment procedure specifies *what* should be charged and *when*. See also Section 2.3.5, where various auction mechanisms are discussed.

Auction protocol

The actual choice of the auction protocol can depend on many factors. We focus on the single-bid sealed auction, being a communication-efficient auction. With this procedure, each supplier submits a single sealed bid for a particular consumer. The MMA allocates the available positions to the highest bidders, where the first position is allocated to the highest bidder, the second position to the second highest bidder, and so on. In some environments the ranking is not important, whereas in other cases the profits for the supplier depend on the position obtained. For this reason, the choice of payment scheme matters, and is discussed below. Note that, since the MMA executes the auction for each arriving consumer, suppliers losing an auction could increase their bid in the next auction for a similar consumer.

Payment procedure

Several different payment schemes are possible for various auction procedures. In the Vickrey auction, the winner pays the price of the second-highest bid (see also Section 2.3.5). This is a prominent and widely-used auction type, which has been shown to be efficient for independent valuations of the item [27, 133, 136]. The auction is also robust, since revealing ones true preferences is the dominant strategy in case of independent valuations.

For the case where multiple banners are shown concurrently, we apply an extension of the Vickrey auction where winners pay the *(N+1)* price, where N is the number of items (here banners). This is an instance of the generalised Vickrey auction, which has the same auction characteristics as above (see Section 2.3.5 for details).

Note that in such a setup, the same price is charged to the winners of a banner placement. The auction is only theoretically guaranteed to work well if the sold goods (the attention spaces) are assumed to be identical, an assumption that is dependent on the way a customer chooses from a list of alternative offerings. In the simulations, we investigate models of customer behaviour where this assumption is valid, as well as a model where it does not hold. In the latter case, we also investigate another payment scheme, the so-called *next-price* auction. Here, each winner pays the price of the next-highest bidder. Such more complicated auctions are notoriously hard to theoretically demonstrate optimal behaviour for, and we use the ACE methods (as discussed later) to show that in the simulations this auction does work efficiently in the case where the valuation of an attention space depends on the position it has on the list and when the highest position is the most valuable (and the second-highest position is the next most valuable etc..).

5.2.5 Effectiveness and feasibility

Although the typical business information for the supplier agent can contain many variables that relate to those in a consumer profile, these cannot be matched directly. Rather, the supplier must find and improve its actual niche in the market, especially in the fine-grained advertisement mechanism of CASy. Similar observations hold even more for the valuation of a consumer.

The need for accurate valuation and targeting is especially pronounced when consumers are significantly contested by competing suppliers. We illustrate this by the case of a very expensive department store: consumers arriving in a fancy car are *a priori* as likely to buy at the store as consumers arriving in a middle-class car. However, when a cheaper department store exists across the street, this competition changes the behaviour of the latter consumers much more than of the former. Similarly, in CASy the valuation of an advertisement space depends on the selection of and competition between suppliers.

An N+1 auction mechanism is theoretically efficient in case of fully rational agents and independent valuation of the items. However, if consumer purchases are like consumer models 2 and 3 (see also Section 5.3.2), the valuation of advertisement space also depends on the selection and competition between various suppliers. It is then unclear whether an efficient allocation of the attention space will emerge, i.e., a correct match between consumers and suppliers with the largest appearing interests for being displayed together.

In the following, we will show via evolutionary simulation as in the field of agent-based computational economics (ACE, [123]) that the market mechanism is indeed effective and results in an efficient allocation. Furthermore, supplier agents learn to properly evaluate their environment and thereby locate their niche in the market.

5.3 Evolutionary simulation model of CASy

In this section, we model the electronic shopping mall for an evolutionary simulation as in ACE, based on Section 5.2. The goal of the simulation is to assess the feasibility of the market mechanism of CASy (see Section 5.2.5). To this end, we will make some additional assumptions and simplifications, which enables us to study, measure, and visualise the emerging behaviour of CASy (results are given in Section 5.4).

5.3.1 Mall manager agent

The MMA has in the simulation 3 banner advertisements to dispatch (see also figure 5.1), and executes the auction as described before.

5.3.2 Consumer models

We abstract away from any interpretation of the profiles. Profiles are represented by a vector of real values. In the simulations, the consumers are classified by a one or two dimensional vector with entries in a $[0 \ldots 1]$ range. The profile can reflect a consumer's interests such as price segment, taste, or quality, or any combination of characteristics projected on 1 or 2 dimensions. We thus model a class of consumers for some given category of products. In the simulation of CASy, several consumers with different profiles arrive and are contested by the suppliers in CASy.

The "buying" behaviour or feedback of the consumers is also simulated. This enables the supplier agents to learn the proper bidding strategy. We first model the purchasing behaviour of a single consumer for one isolated supplier, and then extend the buyer behaviour to models with several displayed suppliers.

Buying behaviour model for one consumer and one supplier

For each supplier i, the *expected gross monopolistic profits* $E\langle\pi_i(c)\rangle$ is its average gross profits for a possible purchase following the observation of a consumer of its advertisement, while no other supplier is shown. We take

$$E\langle\pi_i(c)\rangle = \mu_i P_i(c),$$

where $P_i(c)$ denotes the monopolistic purchase probability for consumer profile c and μ_i is a constant value related to the supplier's average profit when a purchase is made. Note that both μ_i and $P_i(c)$ are taken as an externally imposed model for interaction and are initially not known or available to the supplier.

In the simulation each supplier is given a *centre of attraction* a_i, where $P_i(c)$ is maximised. We used two types of purchase probability functions P_i in the experiments: (1) linear functions, where the P_i is proportional to the Euclidean distance $d(c, a_i)$ in the following way:

$$P_i(c) = 1 - \delta d(c, a_i),$$

and (2) Gaussian functions with the highest point corresponding to the centre of attraction. The width of the Gaussian curve is then set by parameter σ_i. For simplicity the maximal monopolistic purchase probability is set constant to 1. This value can be chosen lower, but is chosen for maximal discrimination between various advanced behaviour models (see Subsection 5.3.2).

Buying behaviour models for several displayed suppliers

As the consumer is presented with a selection of winning "Consumer Attention Spaces", we assume that with some probability p he or she will buy a product. In effect, this stochastic behaviour can be modelled as meaning that a single presentation of banners results in an amount p of products being sold: how much and at which recommended supplier (the buying behaviour) is formalised in the Customer Buying Behaviour Models. Here, we present several Customer Buying Behaviour Models, as the behaviour of consumers shopping for a specific product may be different for different product areas or different consumer populations.

We modelled three classes of consumer behaviour:

1. *Independent visits with several purchases.* In this model (see figure 5.3), the consumer visits all displayed suppliers, and can buy products at several suppliers (e.g. CDs).

2. *Independent visits with one expected purchase.* In this model (see figure 5.4), a consumer visits all displayed suppliers and then buys on average one product in total (e.g. a computer).

3. *Search-till-found behaviour.* In this model (see figure 5.5), the consumer visits the suppliers in sequential order from top to bottom, until he finds a supplier with the proper product, which he buys (e.g. a raisin bread).

The consumer behaviour in these models is stochastic: whether a product is purchased by consumer c at a certain supplier j depends on a probability value $Q_j(c)$. The monopolistic purchase probabilities $P_i(c)$ are the basic parameters, determining these probability values $Q_j(c)$ as shown in figures 5.3 to 5.5. The expected gross profits $E\langle\rho_j(c)\rangle$ for supplier j is then given by

$$E\langle\rho_j(c)\rangle = \mu_j Q_j(c).$$

Notice that in the models of figure 5.4 and 5.5, the probability that an item is sold at one supplier depends on the monopolistic purchase probabilities of its competitors within the list. Importantly, for the third model, the actual position of a supplier on the list influences the expected average proceeds, meaning that the individual banners are no longer identical. We will address this issue, and a solution, in detail in Section 5.4.5.

5.3.3 Supplier models

We will denote by *gross profit* the profit that a supplier earns on a product, before the cost of advertisement is taken into account (but after accounting for all other costs), and by *net profit* the profit after deduction of all costs, including advertisement cost.

The goal of a supplier is to maximise net profits, and therefore a supplier tries to sell as many items as possible at the lowest possible advertising costs. The net profit of a supplier is also referred to as the supplier's *payoff*. The supplier agents in

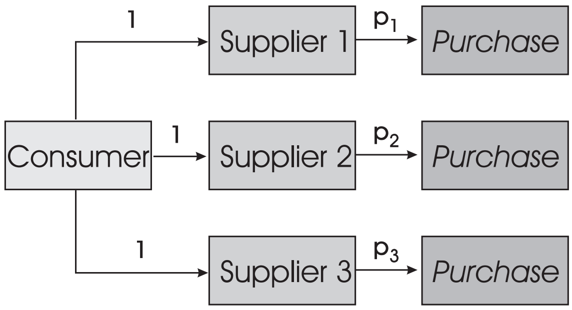

Figure 5.3: Consumer model of independent visits with several purchases, where $P_i = P_i(c)$.

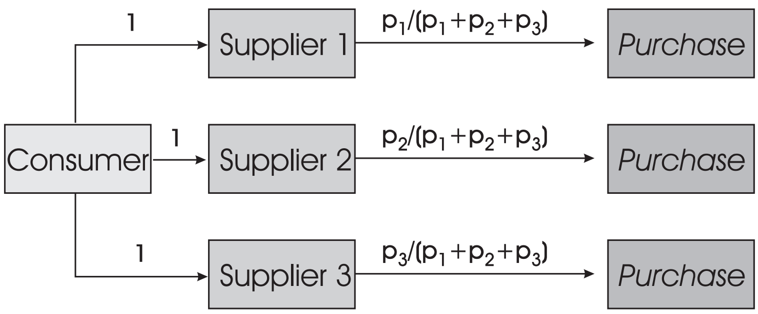

Figure 5.4: Consumer model of independent visits with up to one purchase, where $P_i = P_i(c)$.

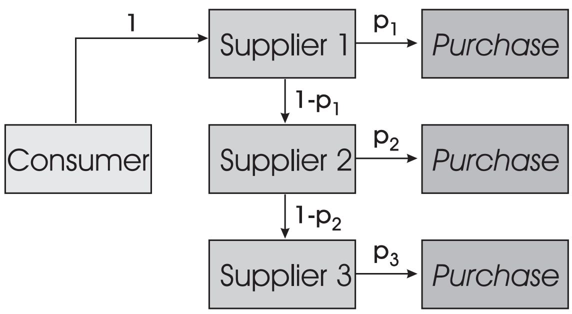

Figure 5.5: Consumer model with search-till-found behaviour, where $P_i = P_i(c)$.

the simulation have no initial knowledge of their own actual niche or payoff function in the market (see Section 5.2.5).

A bidding strategy specifies the monetary bid for each possible consumer profile. Given the feedback in the form of actual payoff from visiting consumers, a supplier agent adapts its bidding strategy and thereby indirectly learns the consumer behaviour and its competitive environment determined by other supplier agents. Note that these two factors are interrelated (see also Section 5.3.2). The strategy of the agent is learned using an evolutionary algorithm (EA). The EA is explained below.

5.3.4 Evolutionary simulation of supplier agents

We simulate the adaptive behaviour of the supplier agents using an evolutionary simulation like in the field of agent-based computational economics (ACE) [1, 42, 123, 127, 130, 138] and similar to the implementation used in previous chapters. Unlike the previous implementations, however, the strategies of each supplier agent evolves independently in a separate population. This is because each supplier agent is of a different *type* (i.e., has a different centre of attraction) and therefore targets different consumers.

We proceed as follows. Each supplier agent is represented by an evolving population of strategies. These strategies are evaluated and evolved according to the amount of profit they earn in a CASy simulation. In such a CASy simulation, a number of consumers arrive, supplier strategies bid for each of these, and the winners get the expected payoffs as described in Section 5.3.2. The strategies that are evolved after repeating this process many times, show the emerging behaviour of the suppliers. Hence, the process of evolution finds effective strategies for a CASy simulation.

An evolutionary algorithm (EA) as described in Section 1.2 is used to adapt the strategies of the supplier agents. The fitness function and the strategy representation are explained below. For further implementation details, see Section 1.2.3.

Fitness evaluation

The fitness of a strategy is equal to the average profits obtained. The actual profit naturally depends on the context, i.e. the profiles of the visiting consumers and the bidding strategies used by the opponents (viz. the competing shops). The populations therefore *co-evolve*. In order to obtain an adequate indication of the performance, the fitness measure is based on several trials with different opponent strategies. The fitness of the opponent strategies is determined concurrently.

We now give a more detailed description of the steps used to determine the fitness of the suppliers' bidding strategies.

1. For each of the suppliers combine the offspring and parent population into a single larger population. We now have m populations, one for each supplier.

2. Reset all previously made profits.

3. Select randomly a single strategy from each population. These bidding strategies are used by the suppliers in the competition. If the competitor is set to *random* (as in Section 5.4.2), however, the strategies are evaluated against random bidding strategies.

4. Let a number of consumers with different profiles visit the shopping mall in a sequential order. We use a fixed set of consumers that are evenly distributed over the profile space (this reduces stochastic variation in measuring the performance of the strategies).

5. For each consumer the supplier obtains feedback on the obtained profits. When a consumer visits the mall the following steps determine the profits:

   (a) Each supplier bids on the consumer using the selected strategy and given the consumer's profile. The strategy is basically a function which maps the consumer profile to a bid. Below, the details on the strategy representation are described.

   (b) The mall manager agent (MMA) selects the winners and determines the advertising costs, as described in Section 5.2.4. Only suppliers who bid higher than zero will participate in the negotiation.

   (c) The MMA shows the list of selected suppliers to the consumer, who decides how much to buy. The purchase amount is determined by the consumer profile and consumer behaviour models described in Sections 5.3.2.

6. The total profits (purchases minus advertising costs) for each strategy are then stored for later reference.

7. If the profits of a strategy have been determined a pre-set, fixed number of trials (and the strategy has thus been tested against different opponent strategies), this strategy is removed from the population.

8. The process is repeated from step 2 until all the populations are empty.

9. The fitness for each strategy then equals the average profit obtained in each of the trials.

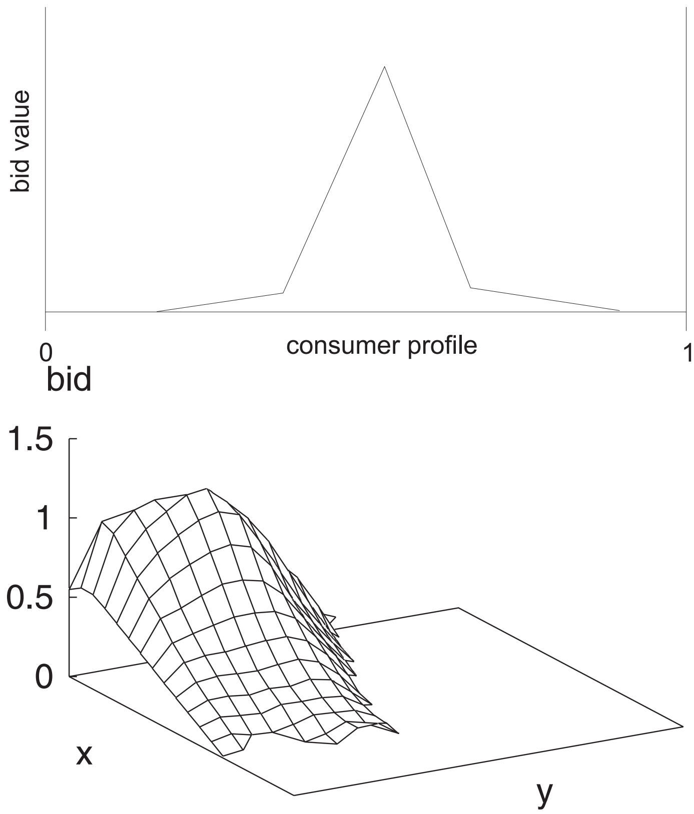

Figure 5.6: Examples of two bidding strategies as learned by co-evolution. The bidding strategy determines the bid value for any consumer profile. The *top* figure shows a strategy for a one-dimensional consumer profile, whereas the *bottom* figure shows a strategy for a two-dimensional consumer profile

Bidding Strategy Representation

In general terms, a supplier's bidding strategy is a function which returns a bid value given the consumer profile. Within the set-up of the simulation the profile has either one or two dimensions. In case of a single dimension, the strategy is represented using a piece-wise linear function that returns the bid given a value along the consumer-profile axis. For a two-dimensional consumer profile, the strategy is represented by triangular planes. Examples of a bidding strategy for a one-dimensional and two-dimensional consumer profile are given in Fig. 5.6 The piece-wise linear bidding strategies are encoded on the chromosome as follows. In case of a one-dimensional profile, the chromosome contains (x, y) coordinates for each of the defining points (the number of defining points is a parameter in the simulation), where x is the consumer profile and y the bidding value. The bidding values for the edges of the consumer profile are always specified within the chromosome. The bidding value for a given consumer profile is then calculated by interpolation between two points neighbouring of the consumer profile on each side.

For a two-dimensional consumer profile, the strategy is represented by triangular planes. The strategy is constructed using Delaunay triangulation of the (three-dimensional) defining points. The bidding value is then determined by interpolation between the three vertexes of the triangle containing the given consumer profile.

5.3.5 Measure for proper selection of suppliers

The selection procedure in an auction should ultimately lead to an appropriate selection of suppliers for consumers. We start from the economic point of view of optimizing the revenue of the collection of shops in the shopping mall as a whole.

Consider the n suppliers with the largest expected payoffs for a given consumer. We measure the *proportion* of properly selected n suppliers as the fraction of these n suppliers that are present in the actual list of 3 displays shown to the consumer.

From the consumer point of view, we can interpret the expenditures of a consumer at a supplier as a measure for his interest in the supplier. In case that the ratio between expenditures and payoff within a certain business sector is similar for the suppliers in that sector, the above measure is related to both the consumer interests as well as the supplier interests.

5.4 Results

We performed a number of experiments in the e-shopping-mall simulation outlined in Section 5.3. The results are given and discussed in this Section.

5.4.1 Simulation settings

Table 5.1 shows the parameters and their values which are varied for different simulation runs. For a description of the mall parameters refer to Section 5.3. For a description EA parameters, see Section 1.2. Two of the parameters are further explained below.

- *Expected gross monopolistic profit functions ($E\langle\pi\rangle$).* The $E\langle\pi\rangle$-functions are explained in Section 5.3.2. The applied settings are specified in table 5.2. Figure 5.7 shows the functions "set2" for 8 different suppliers and a one-dimensional consumer profile. The functions defined in "set3" have different μ_i and δ combinations for each supplier; μ_i varies between 0.5 and 1.0, and δ between 1.0 and 2.0.

- *Number of defining points.* A supplier has to obtain a bidding function on the space of consumer profiles. The function that is learned is an interpolation

	Parameter	Value
EA Parameters	Parental population size (μ)	25
	Offspring population size (λ)	25
	Selection scheme	($\mu + \lambda$)-ES
	Mutation model	self-adaptive
	Initial standard deviations ($\sigma_i(0)$)	0.1
	Minimum standard deviation (ϵ_σ)	0.025
Mall Parameters	Number of suppliers	8
	Number of banner spaces (N)	3
	Maximum bid value	1.5
	Consumer behaviour model	1 / 2 / 3
	Expected gross monopolistic profit ($E\langle\pi\rangle$)	set1 / set2/ set3
	Profile dimensionality	1 or 2
	Number of defining points	8 (1-D,), 16 (2-D)
	Number of consumers	50 (1-D), 100 (2-D)

Table 5.1: Default settings of the simulations.

$E\langle\pi\rangle$ function name	Type	μ_i	δ	σ
Set1	Linear	1.0	2.0	-
Set2	Gaussian	1.0	-	0.2
Set3	Linear	variable	variable	-

Table 5.2: Consumer purchase functions and their general settings.

function, based on a number of defining points. For the one-dimensional case, this results in a piecewise linear function; for the two-dimensional case, we obtain the function values by triangularisation of the profile surface.

5.4.2 Single advertisement model

In this subsection, we illustrate the use and evolution of the bidding function for a supplier for a very simple setting, where the optimal bidding strategy is known from auction theory.

The setting contains a single store competing against a random opponent for the case of one banner. The random player bids any random value between 0 and 1.5. Since a Vickrey (second-price) auction is used, it is a well-known dominant strategy for the supplier to bid its true valuation (i.e. the expected gross profit) [136]; any lower bid risks a missed profit-opportunity, whereas a higher bid might result in direct loss. The dominant strategy maximises the supplier's net profit, regardless of the opponent's behaviour. Thus, the store should learn the profit function as

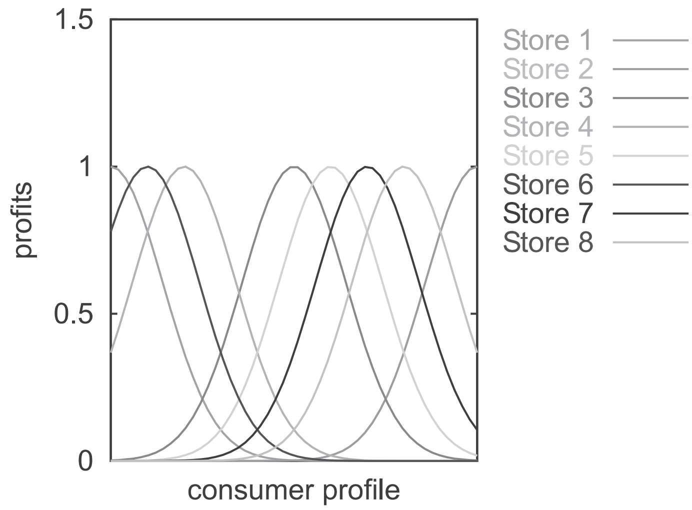

Figure 5.7: Expected gross monopolistic profits at the different stores for "set2" function settings.

the bidding function. The results for experiments on this setting show that this happens indeed. Typical, good results are shown in figure 5.8, where $E\langle\pi\rangle$ is a Gaussian (recall that piecewise linear functions are used).

5.4.3 Consumer model 1: independent visits with several purchases

This consumer model assumes that expected purchases at each supplier can be modelled by the same function as in the single banner case (see Subsection 5.3.2).

The results are shown in figure 5.9. Matching accuracy is measured in several ways. We display the proportion of properly selected n suppliers for 3 banners and $n = 3, 2, 1$ (see Subsection 5.3.5). The reason for including $n = 2, 1$ as well is that the evolutionary system has some degree of stochasticity, and thus small errors occurring frequently can have larger influence on individual outcomes (although relatively little impact on the payoff obtained). Results using these two measures show an almost perfect match. The results after 500 generations of the EA are summarised in table 5.3.

5.4.4 Consumer model 2: one expected purchase

It is more difficult to get a stable system in this situation, since the expected amount purchased at a supplier (and therefore the valuation of a banner space) depends on which other stores are selected as well. Nevertheless, the simulation does stabilise, and the results are comparable to the previous consumer model. See table 5.3.

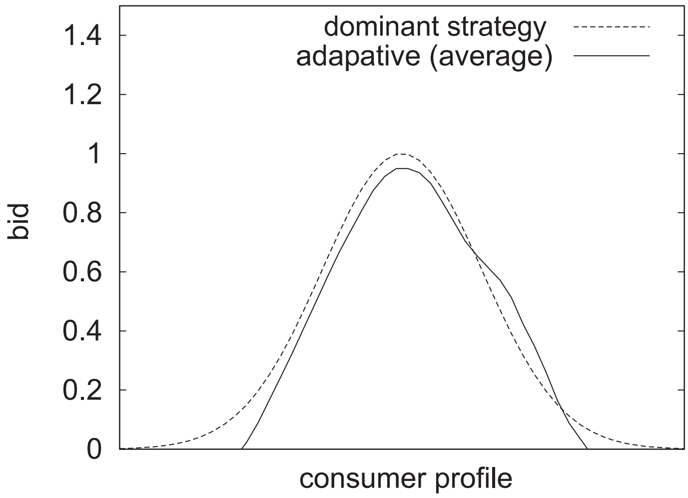

Figure 5.8: Example of a bidding strategy as employed by the supplier after co-evolution no longer increased the profits obtained. Results are shown for a single supplier competing against random supplier. Also shown is the dominant bidding strategy.

Consumer model	$E\langle\pi\rangle$	$n = 3$	$n = 2$	$n = 1$
Regular auction settings				
1	set1	0.95 ± 0.01	0.99 ± 0.00	0.99 ± 0.00
	set2	0.96 ± 0.00	0.99 ± 0.00	1.00 ± 0.00
	set3	0.92 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
2	set1	0.94 ± 0.01	0.99 ± 0.00	0.99 ± 0.00
	set2	0.95 ± 0.00	0.99 ± 0.00	1.00 ± 0.00
	set3	0.90 ± 0.01	0.97 ± 0.01	0.99 ± 0.00
3	set1	0.73 ± 0.03	0.76 ± 0.07	0.79 ± 0.09
	set2	0.83 ± 0.05	0.89 ± 0.06	0.92 ± 0.05
	set3	0.75 ± 0.02	0.89 ± 0.02	0.97 ± 0.01
Next-price auction				
3	set1	0.79 ± 0.03	0.92 ± 0.03	0.97 ± 0.02
	set2	0.75 ± 0.03	0.92 ± 0.02	0.98 ± 0.01
	set3	0.83 ± 0.02	0.95 ± 0.02	0.99 ± 0.00

Table 5.3: Matching results for consumer models 1 through 3. Results denote proportions of properly selected n suppliers for 3 banners and $n = 3, 2, 1$. Averages over 10 runs of the simulation are shown with the standard deviations.

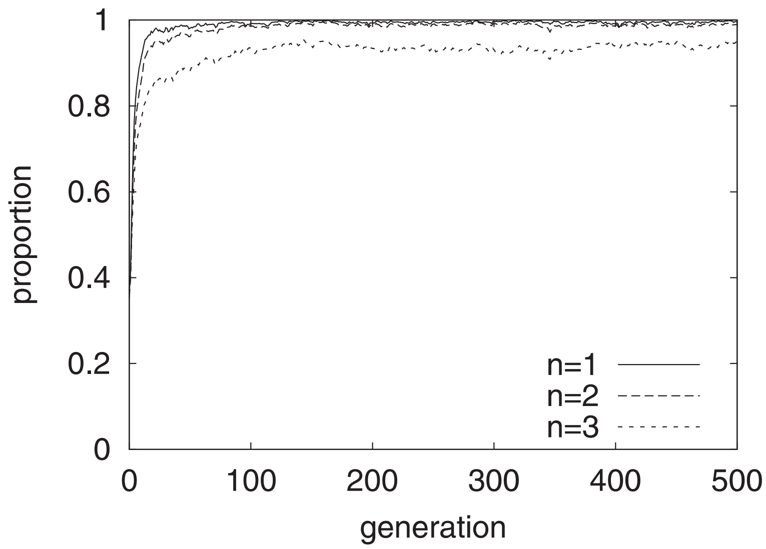

Figure 5.9: Matching results for consumers with independent purchases and $E\langle\pi\rangle$ is set to "set2".

5.4.5 Consumer model 3: search-till-found

In this model, it is not only important for the stores to be in the list, but also to take into account the position on the list (and the other stores above him). Table 5.3 shows that it is indeed more difficult for the stores to find a good matching, in particular when using "set1". This occurs since all relevant suppliers prefer the very top advertisement space and are willing to bid above their valuation (because of the $N+1$-price auction their payment remains relatively low). As a result, the bids reach their limit value (even when this is set to 2.5).

Therefore, we have applied another auction payment procedure as well: each of the winning stores pays the price offered by the next following highest bidder, the so-called next-price auction. This procedure appears to improve the matching, giving comparable results to other consumer models (see table 5.3). Note that a store who obtains the first banner position now pays more than the other stores. This is also reasonable, since the first position is actually more valuable.

We want to remark that we have chosen the maximal purchase probability to 1 (see Subsection 5.3.2) to have maximum difference between this consumer model and the previous ones. When this value is lower, results will become more comparable to the other models also for the regular auction setting.

5.4.6 Two-dimensional profile

We now consider the two-dimensional case, where each consumer profile corresponds to a position within a square. The types of profit functions are similar to the previous

Consumer model	$E\langle\pi\rangle$	$n = 3$	$n = 2$	$n = 1$
1	set1	0.95 ± 0.01	0.99 ± 0.00	1.00 ± 0.00
	set2	0.90 ± 0.02	0.97 ± 0.01	0.99 ± 0.01
	set3	0.93 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
2	set1	0.94 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
	set2	0.92 ± 0.01	0.98 ± 0.00	1.00 ± 0.00
	set3	0.93 ± 0.01	0.98 ± 0.00	0.99 ± 0.00
3	set1	0.85 ± 0.01	0.93 ± 0.01	0.97 ± 0.01
	set2	0.75 ± 0.02	0.89 ± 0.02	0.97 ± 0.01
	set3	0.82 ± 0.02	0.91 ± 0.02	0.94 ± 0.02

Table 5.4: Matching results for consumers with two-dimensional profiles. See also table 5.3 for comparison.

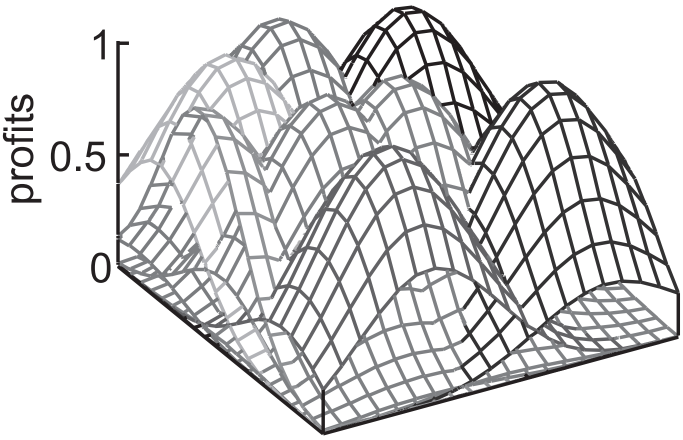

Figure 5.10: Expected gross monopolistic profits $E\langle\pi\rangle$ for "set2" function settings and a 2-dimensional consumer profile.

case, extended for two dimensions. An example is shown in figure 5.10.

The matching results are comparable, but slightly less accurate than for one dimension, see table 5.4. These can be explained through the more difficult learning problem (more defining points are needed for the search function), and thus the settings of the evolutionary algorithms could be further optimised for more accurate learning results in this case.

Specialisation

Interestingly, the suppliers indeed find a niche in the market in case of competition. This becomes clear in figure 5.11, which shows the intersection of a supplier's bidding strategy for two different consumer models, viz. 1 and 2. For consumer model 1,

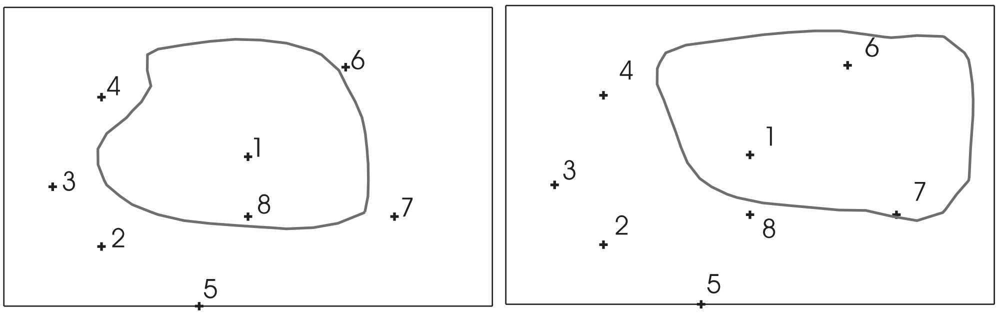

Figure 5.11: Contours of the average evolved strategy at level 0.5 of a supplier 1 at generation 500 for consumer models 1 (left) and 2 (right) using "set2". The points indicate the centres of attraction of the suppliers' Gaussian curves.

a supplier's payoff is independent of the other suppliers displayed. In the second consumer model, however, the payoff is shared amongst the displayed suppliers. In the latter model the payoff thus depends on the competition. We find that this gives supplier an incentive to locate niches in the market, and bid more in places where less competition is present. In figure 5.11, the depicted supplier clearly expands its market to the upper right, and reduces its bids in the lower left region, where competition is relatively greater.

Supplier payoff

The above results mainly focus on the proportion of proper selection. We now briefly discuss the supplier payoffs, i.e. the net profits (see Section 5.3.3). Firstly, we find that in all experiments suppliers obtain positive accumulative payoff in the long run. The strategies emerged are thus individually rational (see Def. 4.2). Secondly, a supplier's payoff depends both on its function settings $E\langle\pi\rangle$ and on the amount of competition. The latter is shown in figure 5.12, which displays the accumulated payoff of the suppliers for consumer model 2 and "set2". The more isolated suppliers, in particular suppliers 4, 6, and 7, obtain a larger payoff than those with much competition (see also figure 5.11). This is due to the difference in advertisement costs. Note that this is in accordance with economics theory: in case of large competition, the net profit of competing suppliers is close to zero.

5.4.7 Conclusions

The experiments show that a proper selection of suppliers emerges with very good to perfect matches. In case consumer model 3 is applicable, a next-price auction mechanism further improves the results. Furthermore, we find that all experiments show positive supplier payoffs. Finally, we observe that shops find their customers

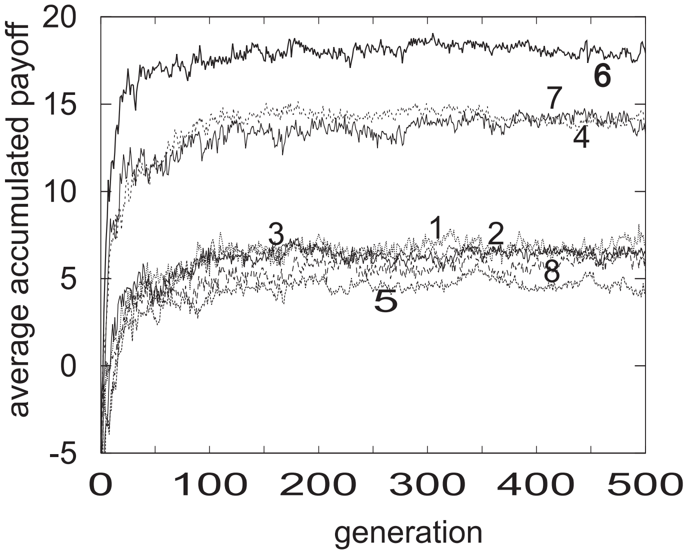

Figure 5.12: The average accumulated payoff for each supplier using "set2" and consumer model 2.

and their niche in the market via CASy.

5.5 Evaluation and further research

5.5.1 Reflections

We can identify a number of commercial and technological advantages of CASy. In CASy, proper matching does not have to be performed or enabled by a third party. This significantly reduces the combinatorial complexity as compared to centrally processing all product ontologies and information about consumers and shops. Furthermore, shops have substantial autonomy and can thus incorporate local domain knowledge and momentary business considerations in their bidding strategies and thus in the ultimate matching process. Especially, they do not have to reveal sensitive business information to a third party, and can take more sales aspects into account: not only product pricing, but also service level, quality, product diversity, or customisation of products. The system also enables them to quickly adapt to market dynamics or their own internal situation (out-of-stock, discount periods, promotion). Note that the relevance of the shop for the consumer is still expressed via the monetary bidding procedure. The mechanism is also a form of dynamic pricing of attention space.

There is much debate about whether or not advances in Information Technology (IT) will increasingly make intermediaries within markets redundant (disintermediation) [37], or whether such advanced IT will help re-establish intermediaries because of new value-added services that become possible (reintermediation) [21]). The results in this chapter can be taken either way: on the one hand, we can conceive the basic auction functionality performed by the MMA to be part of the customer agent, replacing the matching function previously performed by central filtering mechanisms. Alternatively, we noted that there are many possible value-added services regarding user-profile enhancement that could be performed by a central shopping mall intermediary. This conclusion is in line with recent arguments regarding the effects of current agent technology on the disintermediation/reintermedation debate [86].

The proceeds the electronic shopping mall can derive from the matching mechanism (through the auctions) can be used to facilitate additional intermediation services to both customers and shops (e.g. micro payments, 24x7 intermediation). Offering an effective matching mechanism adds considerable value to the customer experience, and can thus be expected to be an important selling point for the electronic shopping mall, and entice suppliers to participate in the mechanism. It will be interesting to investigate the exact economic conditions – such as at which price the suppliers are no longer prepared to follow the customers – for this to be relevant, but we leave that for future research here.

Some points need attention when further implementing CASy. In CASy, information about a consumer is (partially) communicated to suppliers. At the same time however, the consumer's privacy requirements must be respected. We will not extensively address this here, but just mention some approaches: having the consumer decide what information he allows to be communicated, restricting the types of communicated information in general, or conversion of personal information to more sales-related properties. The latter could include restricting the profile to attributes of the desired product (instead of the customer), like "expensive vs. cheap", "ultra trendy vs. conservative" etc.... Such attributes could in principle even be queried from the customer. As argued in [63], no uniform solution for privacy demands exist, rather "privacy will have to be dynamically tailored to each individual user's needs" and requirements.

There remains the issue whether a central entity like the shopping mall would be willing to convey individual user related profile information. Google for instance currently considers its click stream information a business secret. In the setup we introduced here, however, the proceeds that the intermediary obtains from the ongoing auctions, and possibly for additional advanced IT services, will be a strong incentive for the intermediary to consider what parts of the profile information are allowed to be disseminated by its clients (here, the suppliers). Note that when the intermediary charges a (fixed) price for customer profile information services, such information would constitute a sunk cost for each seller, and reduce the available funds for placing advertisements, resulting in lower bids. Since such cost will reduce *all* bids from *all* agents, the relative ordering of the bids remains intact and the market-based selection mechanism itself is not affected by such additional cost.

We remark that once individual shops receive customer (related) profiles, they have the tools developed in information intensive personalised marketing research at their disposal for determining how interested they are in each individual customer: i.e. interactive marketing, database marketing, micromarketing and one-to-one marketing [15, 49, 76, 77, 95]; in [118] these slightly different approaches are considered in more detail. The information filtering mechanism we describe here is then the gate controlling the flood of finely targeted business interests.

Another point concerns the communication between suppliers and shopping mall, which is increased because of the bidding process and the communication of consumer profiles to the suppliers. However, the communication in the mall is linear in the number of customers, and also in the number of participating shops, and the size of the consumer profile. The latter is also typically very small, e.g. up to 100 bytes. In a prototype implementation on a single PC, a single market comprising of 100 *learning* shop-agents was easily able to sustain 100 customers per second, and still continuously update the internal state of the agents (the learning mechanism) [122].

To scale to even larger settings, the market can be divided into a number of segments, with each market handled by different agents. The profile then only needs to be transmitted to agents within a particular market segments, reducing the overall communication. We pursued this approach in a distributed prototype of the electronic shopping mall ([122]). In the extended agent architecture of the prototype different market segments are handled by sub agents (which can run on different machines). In all, we do not perceive the somewhat increased communication as a significant problem, but rather as an issue that can easily be addressed in the process of framework-engineering if necessary.

5.5.2 Open problems and future research

We investigated the concept of CASy for several basic models. The results we describe here show that the market-based approach yields excellent buyer-seller matching given adaptation of the bids made by the sellers. The ACE simulations have been carried out to demonstrate feasibility and learnability of the concept, as these simulations showed effective matching for different auction types and consumer behaviour models. It is also interesting to investigate how software agents can be developed for more advanced settings: one such example would be the extension of the simulations to a dynamical market, with sellers changing their profiles, or sellers entering and leaving the market. For ACE feasibility and learnability studies, methods that can deal with such dynamic environments are only just starting to emerge. As we demonstrated that the steady state version of the problem per-se is both effective and learnable, we would expect that dynamic versions of the problem would also be learnable, but the effectiveness is then rather dependent on the speed and quality of the machine learning techniques employed by the shop agents as well as the actually chosen models for the dynamic environment. E.g., for methods such as neural networks (see [17]), the introduction of dynamics into the market will mean that additional complexity in terms of effective (commercial) exploration/exploitation strategies has to be introduced. At this point we leave the investigation of dynamisation of the system as an interesting problem for future research.

Other points that need to be addressed in future work should be concerned with taking account of the role of (local) ontologies, of marketing and data-mining techniques, and of partial consumer information. Furthermore, in this work, we placed an emphasis on the $N+1$-price auction with single sealed bids. Other types of auctions could be further investigated, for example addressing the possible feedback given on bids of other participants (e.g. multi-round auctions) or to address the revenue of the mall manager.

From the consumer's point of view, we have interpreted the expenditures of a consumer at a shop as a measure for his interest in the shop. CASy gives priority to suppliers with the largest expected payoffs for a given consumer. This thus leads to optimisation of the revenue of the collection of shops in the shopping mall as a whole. In the case that within a certain business sector, the ratio between expenditures and payoff is similar for the suppliers in the sector, this means that CASy completely reacts on the interest of an individual consumer. However, across different sectors, there may be differences or anomalies, leaving the extension of CASy with additional (monetary) correction mechanisms to avoid such anomalies as an interesting open problem. This is part of our future work.

Finally, our system CASy is complementary to existing recommendation systems. It is important to know in what way these together could be used as part of a broader system. Also, which application areas are more suited for the existing recommender systems, and which for the CASy system.

5.6 Concluding remarks

In this chapter, we present a competitive distributed system, CASy, for allocating consumer attention space (Section 5.2). By evolutionary simulation as in agent-based computational economics (ACE), we show the conceptual feasibility of the system (Sections 5.3 and 5.4). We modelled the various parts in the system in a basic and simple way suitable for analysis, visualisation, and comparison, and show that proper matchings emerged while suppliers can learn their niche in the market. Finally, we reflect on the advantages, opportunities, and further open problems concerning the proposed system (Section 5.5).

Chapter 6

Automated bargaining and bundling of information goods

Personalisation of information goods becomes more and more a key component of a successful electronic business strategy [2]. The challenge is to develop systems that can deliver a high level of personalisation combined with, whenever possible, a high adaptability to changing circumstances. In this chapter,[1] we introduce a system which can attain these properties through the manner in which it sells information goods.

We consider a novel approach in this chapter, where bundles of information goods, such as news articles, stock quotes, music, and video clips are sold through automated negotiation. Bundling of information goods has many potential benefits including complementarities among the bundle components, and sorting consumers according to their valuation (see [9] and the references therein). The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interest of buyers of information goods.

The system as developed is also capable of taking into account business related constraints. More specifically, it tries to ensure that customers perceive the bargaining outcomes as being "fair" by having customers end up with equivalent offers whenever that seems fair. This is important for customer satisfaction and acceptance of the system by customers. Partly because of this fairness constraint the actual bargaining process is not really one-to-one bargaining between seller and customer but instead is one-to-many (i.e., between seller and customers).

In the developed system, autonomous "software agents" perform (part of) the

[1]This chapter is based on [120]: K. Somefun, E.H. Gerding, S. Bohte, and J.A. La Poutré. Automated negotiation and bundling of information goods. In *Agent-Mediated Electronic Commerce V*, Springer Lecture Notes in Artificial Intelligence (LNAI). Springer-Verlag, Berlin, to appear.

negotiation on behalf of the users of the system. A seller (or information provider) agent negotiates with several buyer (or customer) agents simultaneously in a bilateral fashion, using an alternating offers protocol like in Chapter 3. The agents are capable of negotiating about several issues simultaneously, such as the price and the quality of the offered goods. Chapter 3 showed that, using such a bilateral negotiation protocol, efficient outcomes can be achieved after a process of learning from several negotiations. In this chapter, we introduce strategies that are capable of finding efficient solutions *within* a single negotiation (i.e., real-time).

To enable efficient and real-time multi-issue bargaining outcomes, we decompose the bargaining strategies into concession strategies and Pareto-search strategies. The concession strategy determines the desired utility level during the bargaining process, whereas the Pareto search strategy looks for *Pareto-efficient* (see Def. 4.3) outcomes that maximise win-win opportunities for a given a desired utility level. Together these strategies produce offers and counter offers for the agents. An important contribution of this chapter lies in the actual development of Pareto search methods that result in efficient solutions while, at the same time, bargainers make concessions using a variety of concession strategies. To that end, we introduce the orthogonal and orthogonal-DF method: two Pareto search methods. We show through computer experiments that the respective use of these two Pareto search methods by the two bargainers, combined with various concession strategies, results in very efficient bargaining outcomes (i.e., these outcomes closely approximate Pareto-efficient bargaining solutions). We obtain these results without assuming any a priori knowledge of other player, nor experience from previous bargaining games.

The remainder of this chapter is organised as follows. First, we introduce a system for selling bundles of news articles through bargaining in Section 6.1. Section 6.2 discusses the buyer and seller agent in more detail and presents bargaining strategies for multi-issue negotiations. In Section 6.3 we investigate the Pareto-efficiency of the introduced bargaining approach through computer experiments. As we only consider the Pareto-efficiency of the deals reached in this chapter, we do not simulate the entire system as developed, but rather restrict attention to bargaining with a single buyer. Experiments using one-to-many bargaining are investigated in the next chapter. Related approaches such as auctions are discussed in Section 6.4. In Section 6.5 we revisit our approach and conclusions follow in Section 6.6.

6.1 A system for selling information goods

The goal is to develop a system for the sales of bundles of news items where buyers bargain over the pricing, quality, and content of the bundles. The negotiated contract applies to a fixed time interval, which is typically a short period of time, e.g. a single day. The bundle *content* defines which news categories the bundle contains. The system distinguishes between k categories. We furthermore distinguish between *low* and *high* quality-of-service categories. If a category with *low* quality of service is selected, a buyer receives only the news headlines for this category. A buyer can, however, after reading the headline, decide to purchase the entire article. In that case, a *variable price* is paid. Alternatively, the buyer can opt for a *high* quality of service category, in which case the buyer obtains all the articles without additional (variable) costs. In the following, we simply use *quality* to refer to the quality of service.

The buyer negotiates about the variable price, the content, and the quality of the categories in the bundle. At the same time, a buyer negotiates a *fixed price* which is an upfront payment for the bundle as selected. Clearly, a high quality category is likely to result in a higher fixed price than a low quality category. Both buyer and seller have private preferences regarding such trade-offs between issues. Differences in preferences allows for the possibility of win-win outcomes (see also Chapter 3). The agents in the system can find these win-win outcomes using Pareto-search strategies, without having to fully disclose their preferences.

The value customers attach to news items may fluctuate heavily due to daily dynamics. Moreover, there may be wide differences in personal interests of customers. The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality and content of the delivered goods may, for example, differ based on daily dynamics and personal interest of customers.

The system as developed is also capable of taking into account business related constraints. More specifically, it tries to ensure that buyers perceive the bargaining outcomes as being "fair" by having buyers end up with equivalent deals whenever that seems fair. Due to the notion of fairness, negotiations are no longer independent and bilateral, but are in fact one-to-many from the perspective of the seller. Fairness and the way in which it affects the seller's bargaining strategy is discussed in more detail in Section 6.1.2. We first continue, however, by describing the bargaining aspect of the system in Section 6.1.1. The bargaining protocol used is explained in Section 6.1.3.

6.1.1 Bargaining using software agents

Within the system, autonomous *software agents* (see Section 1.1.3) perform (part of) the negotiation on behalf of the seller and the buyers. A buyer agent is a software agent owned by the buyer, and a seller agent is owned by the seller. Buyers and seller instruct their agent through a user interface (UI). Figure 6.1 depicts, at a high abstraction level, the bargaining process between a buyer and the seller. There are roughly three possibilities for implementing the starting time of the negotiation process: buyers can negotiate a contract before, after, or during the time that the news becomes available. The system is set up in such a way that all three possibilities

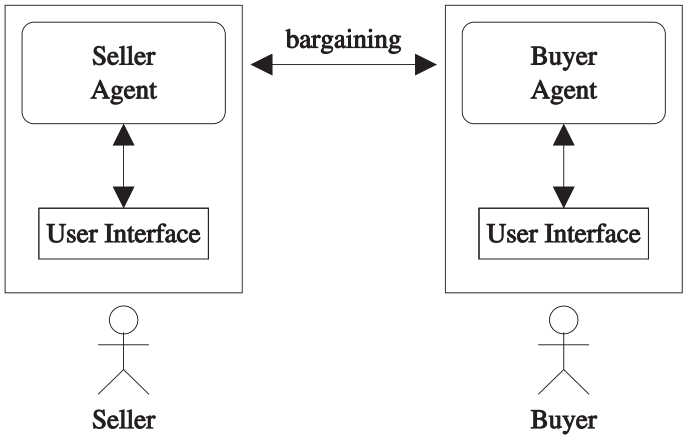

Figure 6.1: The bilateral bargaining process between a seller and a buyer using software agents.

can be implemented.

Given a desired bundle content, a buyer agent can negotiate with the seller agent about the fixed price, variable price, and the quality for each category. The negotiated contracts apply to bundles of news items which become available during a predefined and fixed time interval (e.g., a day). The value buyers attach to news items may fluctuate heavily due to daily dynamics. Moreover, there may be wide differences in personal interests of buyers. The advantage of the developed system is that it allows for a high degree of flexibility. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interest of buyers.

6.1.2 Fairness and one-to-many bargaining

Potentially, bargaining can lead to unsatisfied buyers if buyers perceive the outcomes of the negotiations as unfair. This can occur when, for instance, two buyers obtain similar goods at the same time but end up paying very different amounts. Fairness of negotiation outcomes is important for customer satisfaction, which in turn may be important for a business' long term profitability. The seller agent can prevent unfair outcomes by incorporating a notion of fairness, whereby buyers are treated in a similar fashion. This notion of fairness also implies that any information that is revealed about buyers during negotiation or by using the system in general, is not used to their disadvantage in relation to other buyers. This is also essential in order for buyers to accept the system and delegate responsibilities to software agents.

In the system, the following notion of fairness is incorporated into the bargaining strategy of the seller agent: within a limited time frame, the seller agent maintains an equal expected utility level with buyers who are interested in an identical bundle content. To define fairness more formally, suppose a buyer reaches a deal at time t_d. We say that this deal is fair, relative to a fixed interval $\Delta > 0$, whenever there exist a start time t_s, with $t_d \in [t_s, t_s + \Delta]$, such that the *seller*[2] is indifferent between any other deal reached within the interval $[t_s, t_s + \Delta]$.

Whenever price is the only negotiable issue, the notion of fairness simply implies that all buyers interested in the identical bundle content end up paying the same price for this bundle, given that the deals are reached within a given time frame. This notion of fairness corresponds to the notion of *envy-freeness* in auctions [44], adapted to the more continuous setting of bilateral bargaining. In our case, however, negotiations concern several issues, in which case the expected utility level is used rather than the price. Note that the values for the various issues, such as fixed and variable price, can still vary for different buyers, since buyers can have diverse interests. This is an essential aspect of personalisation which needs to be preserved. Fairness, however, ensures that the seller's expected utility for these different deals is identical.

Because of the fairness imposed on the seller strategy, the bargaining process between the seller and an individual buyer can also affect other negotiations which occur concurrently. Fairness limits the bargaining options of the seller. Therefore, bargaining between a seller and a buyer is not really bilateral, but is in fact one-to-many. Note that this holds only from the perspective of the seller. The buyers can normally not observe the negotiation processes with other buyers, and therefore perceive the negotiations to be bilateral.

We note that besides fairness, also other business side-constraints may be implemented. The actual way in which side-constraints, such as fairness, are implemented may be important because it can alter the strategic behaviour of buyers as well as the seller.

6.1.3 Bargaining protocol

The seller agent negotiates with many buyer agents simultaneously by alternating offers and counter offers. An offer specifies the fixed price, the variable price (uniform for all low quality categories), the bundle content, and the desired quality for each category separately. Formally, an offer is described by the tuple $< p_f, p_v, \vec{b}, \vec{q} >$, where p_f is the fixed price, p_v is the variable price. Furthermore, $\vec{b} \in \{0,1\}^k$ is a binary array describing the bundle content, where $b^i = 1$ if category $1 \leq i \leq k$ is selected, and $b^i = 0$ otherwise, and $\vec{q} \in \{0,1\}^k$ describes the quality settings for each category, where $q^i = 0$ if a selected category is of low quality, and $q^i = 1$ if category i if of high quality.

[2]Note that since the preferences of the buyers are hidden, fairness is defined from the perspective of the seller agent.

Attached to an offer are also preconditions which specify until when the offer is valid. If the offer is accepted within that time, the proposing agent is bound to the conditions specified in the offer. Otherwise, the offer expires. We call the offer combined with the preconditions a proposal. A bargainer can accept, reject, or place a counter proposal. The bargaining process continues until an agreement is reached or one of the bargainers terminates the process. Figure 6.2 depicts the alternating offer bargaining protocol.

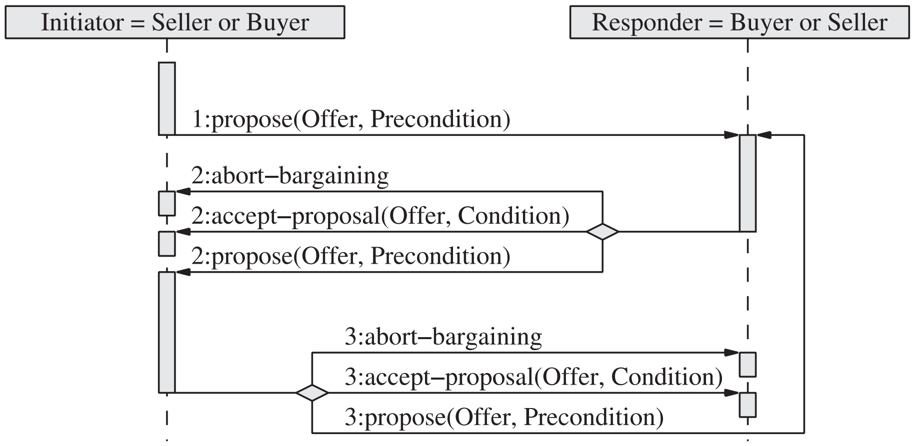

Figure 6.2: The bilateral bargaining protocol.

To accelerate the negotiation process, we allow concurrent negotiation threads for the same bundle content with different quality settings. The buyer can therefore submit several offers at the same time. In order to discern between threads, each thread must have a different combination of quality settings for the selected categories. The seller can only respond by varying the fixed and variable price. The thread in which the agreement is reached first determines the prices and quality settings for the desired categories. Figure 6.3 depicts the one-to-many bargaining process and the possibility of parallel negotiation threads between a buyer and the seller.

Using the above protocol, offers submitted by the buyers could violate the notion of fairness if these offers are immediately accepted by the seller. To provide a seller with the opportunity to ensure fairness (as defined in Section 6.1.2), the bargaining protocol allows for post-agreement negotiation: the bargainer who accepted the offer can propose a post-agreement offer which the other party either accepts or rejects.[3] In case of an acceptance, the original offer is replaced by the post-agreement offer. The process terminates after the post-agreement offer is proposed and is then either accepted or rejected. Post-agreement negotiation can be used by the seller to adjust the offers in favour of the buyers, such that fairness is ensured within the

[3]Post-agreement negotiation is a common approach in the single negotiation text literature [32].

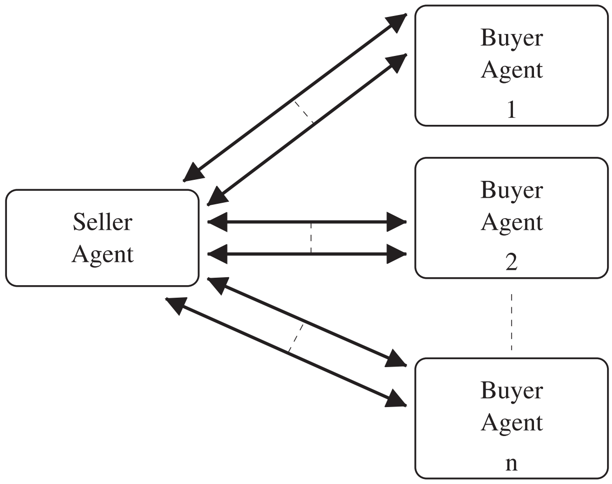

Figure 6.3: One-to-many bargaining with parallel threads.

defined time interval. Note that, in case of multiple issues, the seller can produce a more favourable offer by conceding on one or more issues. Although the buyer's preferences are *private* and unknown to the seller, this approach assumes that a buyer always prefers a lower (fixed or variable) price or a better quality.

6.2 Agents and bargaining strategies

In this Section we discuss the seller agent and buyer agent in greater detail. Additionally, we introduce bargaining strategies that generate good (i.e., closely approximate Pareto-efficient) multi-issue bargaining outcomes.

6.2.1 Seller agent

The seller agent's bargaining behaviour is based on the agent's so-called aspiration level, which we define as follows:

Definition 113.1 ***Aspiration Level*** An *aspiration level* of an agent refers to an agent's desired expected utility level.

Unlike common usage in the literature (where aspiration level is used as a point of reference), the aspiration level is used here as the minimum expected utility the agent is willing to accept at a certain point in time. If the expected utility of an offer received by the buyer exceeds the aspiration level, the offer is accepted, otherwise the offer is rejected. Whenever the seller agent makes a (counter) proposal, the offer's expected utility is set in a way as to match the aspiration level. The aspiration level can change during the process of negotiation. The aspiration is adjusted using the *concession strategy*, whereas the generation of an offer (given a fixed aspiration levels) is achieved by the Pareto-search strategy. These strategies are considered more closely in Section 6.2.3. In this Section, we specify the seller agent's measure of *expected utility*. The seller's expected utility u_s for an offer $< p_f, p_v, \vec{b}, \vec{q} >$ is defined as follows:

$$u_s(< p_f, p_v, \vec{b}, \vec{q} >) = p_f + p_v \sum_{i=1}^{k} e_s^i b^i (1 - q^i), \tag{6.1}$$

where the components e_s^i of vector $\vec{e_s}$ denote the seller agent's expectation about the number of articles an average buyer will read for the duration of the contract, specified for each category $1 \leq i \leq k$. Note that $b^i(1 - q^i)$ equals 1 if category i is selected and is of low quality, and 0 otherwise. Therefore, $\sum_{i=1}^{k} e_s^i b^i (1 - q^i)$ indicates the expected *total* number of articles an average buyer will read in the selected *low*-quality categories (and for which the buyer pays an additional p_v per article). The seller agent can estimate $\vec{e_s}$ based on, for example, aggregate sales data.

Due to the notion of fairness, the seller agent cannot apply different aspiration levels for different buyers in case of identical bundles (within the defined time interval). Consequently, the seller agent must use the same measure of expected utility in different (simultaneous) negotiations. A seller agent is therefore not allowed to use knowledge of individual buyers, such as their past reading behaviour, to directly discriminate between buyer agents in the negotiations. In other words, the seller agent must use the same values for $\vec{e_s}$ in negotiations with different buyers (within the defined time interval). We note, however, that the components e_s^i of $\vec{e_s}$ need not be constants, but can be functions as well. In the experiments described in Section 6.3, for example, the expected number of articles read is a declining function of the variable price p_v. This incorporates the likely assumption that buyers who prefer a high variable price, will purchase less additional articles on average than buyers with a low variable price. This can be used to *indirectly* discriminate between buyers, without violating the notion of fairness. We defer further discussion on the topic of price discrimination until Section 6.5.

6.2.2 Buyer agent

The buyer agent acts on behalf of the buyer. The buyer can indicate her preferences by specifying, for each information category she is interested in, the amount of articles she expects to read. The buyer can furthermore select between several negotiation strategies to be used by the agent and specify a maximum budget b_{max} for the given bundle content and number of articles. The budget provides the agent with a mandate for the negotiation; the total expected costs for the selected bundle should not exceed b_{max}. The value b_{max} can also be interpreted as the buyer's worth for the bundle content and the number of articles specified by her preferences.

Similar to the seller agent, the buyer agent's bargaining behaviour is based on a desired level of expected utility or aspiration level. Given an offer $< p_f, p_v, \vec{b}, \vec{q} >$, the buyer agent's expected utility u_b in case of an agreement is defined as follows:

$$u_b(< p_f, p_v, \vec{b}, \vec{q} >) = b_{max} - [p_f + p_v \sum_{i=1}^{k} e_b^i b^i (1 - q^i)], \tag{6.2}$$

where the components e_b^i of the vector $\vec{e_b}$ describe the *buyer*'s expectations regarding the number of articles she will read, specified for each category. In case of a disagreement, the buyer agent's utility equals zero. Note that the part of Equation 6.2 in squared brackets is identical to seller's expected utility (see Equation 6.1), except that e_s^i is replaced by e_b^i.

As mentioned Section 6.1.3, the negotiation protocol allows for multiple negotiation threads for the same bundle content. Given a bundle content with k categories, in principle 2^k threads are possible (by varying the selected quality of each category). The buyer agent, however, selects only a limited number of combinations based on the buyer's preferences, to reduce the amount of communication. In the current system the buyer agent initiates $k + 1$ threads. In the first thread the quality for all categories is set to low. In the second thread, only the quality for the category with the highest expected articles read is set to high. In the third thread, this is done for the two categories with the first and second highest expected articles read, and so on. Within a thread, a fixed price and a variable price are negotiated.

6.2.3 Decomposing the bargaining strategy

The buyer agents and seller agent are endowed with various bargaining strategies that can bargain over multiple issues. We decompose bargaining strategies into *concession strategies* and *Pareto search* strategies. *Concession strategies* determine what the aspiration level of an offer will be at any decision point. *Pareto search strategies* determine, given the current aspiration level, and given a particular history of offers and counter offers, the actual (multi-dimensional) offer, i.e., the fixed price p_f and the variable price p_v. Note that the quality settings are fixed for a particular negotiation thread. As a result, the Pareto-search strategy in this case is only concerned with continuous issues. In terms of a multi-dimensional utility function, a (counter) offer entails both a movement of the so-called *iso-utility curve* and a movement along the iso-utility curve. Given a specified utility level, an *iso-utility curve* connects all (p_f, p_v) points which generate that utility (see Figure 6.4 for an example). Concession strategies determine the movement of an iso-utility curve; Pareto search strategies determine the movement along an iso-utility curve.

Pareto search strategies aim at reaching agreement as soon as the respective concession strategy allows it. Therefore, it may be good for both parties to use such an approach. The resulting agreements are then also Pareto efficient (see Def. 4.3).

From a system design perspective, Pareto efficiency of the negotiated bundle is desirable since it maximises win-win opportunities.

In Section 6.2.4 we introduce a particular class of Pareto search strategies. The experiments in Section 6.3 show that if the seller agent uses this Pareto search algorithm and buyer agents use a similar Pareto search algorithm, then the bargaining outcome will closely approximate a Pareto-efficient solution given a wide variety of concession strategies.

In the system the seller agent uses an instance of the Pareto search algorithms combined with a concession strategy. Although a buyer is free to select other bargaining strategies, the system is set up such that it is actually in the best interest of buyers to have their agents use Pareto search strategies combined with a concession strategy. We elaborate on this issue in the discussion in Section 6.5.

6.2.4 Orthogonal strategy and DF

Both buyer agent and seller agent may use what we call an *orthogonal strategy* as the Pareto-search strategy. This strategy is probably best explained through an example. Suppose, the buyer (with whom the seller bargains over the combination of p_f and p_v) places the t^{th} offer $< p_f(t), p_v(t) >$ (since the remaining attributes $\vec{b}$ and $\vec{q}$ remain fixed, we omit these attributes in the following). Moreover, the seller's concession strategy dictates an *aspiration level* of $u'_s(t+1)$: i.e., the (counter) offer should have an expected utility of $u'_s(t+1)$. Based on this information, the seller's orthogonal strategy determines a counter offer $< p_f(t+1), p_v(t+1) >$, such that $u_s(< p_f(t+1), p_v(t+1) >) = u'_s(t+1)$ and the point $(p_f(t+1), p_v(t+1))$ lies closest, measured in *Euclidean* distance, to the point $(p_f(t), p_v(t))$. Figure 6.4 provides a graphical example of the orthogonal strategy. In this Figure, function f_s denotes the seller's iso-utility curve at time $t+1$, containing all points (p_f, p_v) such that $u_s(< p_f, p_v >) = u'_s(t+1)$.

The use of the orthogonal strategy by both parties results in a mapping f from a bargainer's aspiration level at t to the aspiration level at $t+2$. Given convex preferences (cf. [72]) and fixed aspiration levels the mapping f can be shown to satisfy the Lipschitz condition $||f(x)-f(y)|| \leq ||x-y||$ for all x and y in the domain of f.[4] Thus, given fixed aspiration levels and convex preferences, the orthogonal strategy does imply that consecutive offers do not diverge. Figure 6.5 illustrates the use of the orthogonal strategy by both parties for the case of tangent iso-utility curves. It draws a sequence of two offers and counter offers with convex preferences and a fixed aspiration level. The figure illustrates, for instance, how the buyer's offer at time $t=1$ is transformed into an offer at time $t=3$ (where the aspiration

[4]The proof is a straightforward application of convex analysis (cf. [140]) given that without loss of generality we can assume that the preferences are bounded. That is, negative and extremely high $< p_f(t), p_v(t) >$ combinations can be discarded, without loss of generality.

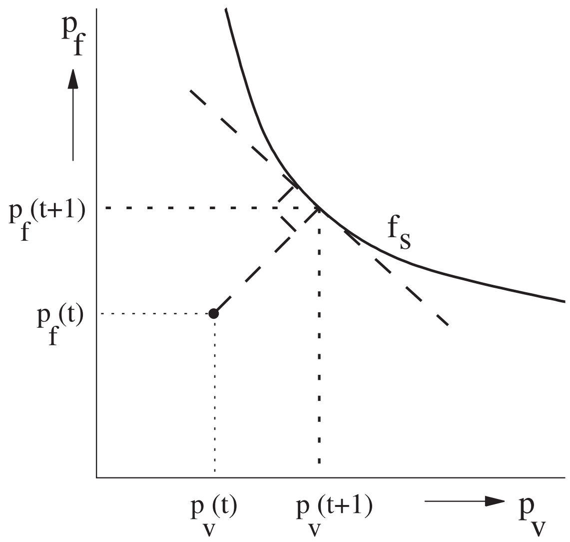

Figure 6.4: Example of the orthogonal strategy, where f_s denotes the seller agent's iso-utility curve.

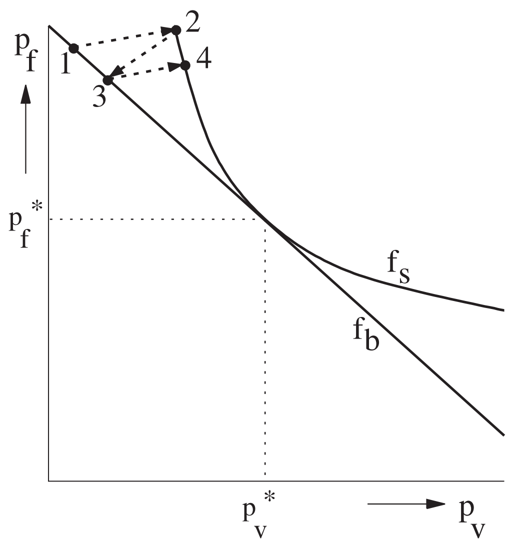

Figure 6.5: Sequence of two offers and counter offers with fixed aspiration levels and convex preferences, where $< p_f^*, p_v^* >$ denotes a Pareto-efficient offer. Here, f_s and f_b denote the iso-utility of the seller and buyer agent respectively.

level remains constant, i.e., $u'_b(1) = u'_b(3)$).

The use of just the orthogonal strategy by both parties may lead to very slow convergence to Pareto-efficient bargaining outcomes. To speed up the convergence process we can add an amplifying mechanism to the orthogonal strategy. As the amplifying mechanism we use the derivative follower with adaptive step-size (ADF).

Henceforth, we will call this strategy the orthogonal-DF.

The derivative follower (DF) is a local search algorithm (cf. [61]). It adjust the variable price p_v returned by the orthogonal strategy by either subtracting or adding δ to it depending on the result of the previous two adjustments, where δ is called the step-size of the DF. Consequently, also the fixed price p_f changes because the adjusted offer still needs to generate the same expected utility level (specified by the concession strategy). The difference between ADF and DF is that the step-size δ becomes adaptive [26, 129]. We use the ADF proposed by [129]. Intuitively, the idea is to increment the step-size relatively far away from a Pareto-efficient solution and decrement it in the vicinity of a Pareto-efficient solution. Consequently, a quicker and more accurate search of the solution space becomes possible. Algorithm 1 (on page 119) specifies the orthogonal-DF in greater detail and figure 6.6 illustrates the use of the orthogonal-DF by the seller (where the buyer uses the orthogonal strategy only).

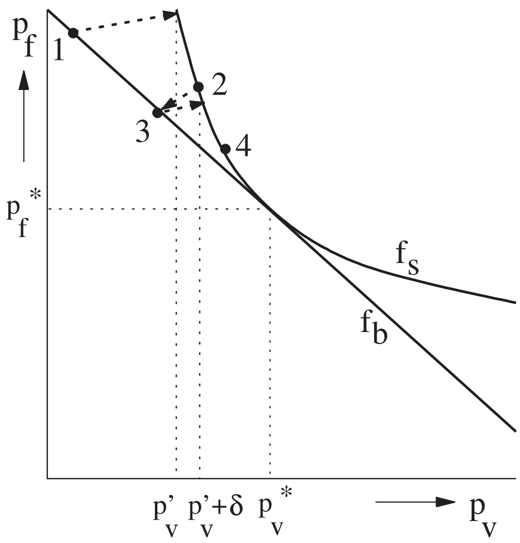

Figure 6.6: Sequence of two offers and counter offers with fixed aspiration levels where the seller uses the orthogonal-DF and the buyer only uses the orthogonal strategy.

6.3 Experimental setup and results

The previous sections outlined the general system for selling bundles of news items to several buyers through negotiation. As discussed in Section 6.2.3, negotiation essentially consists of two strategic aspects: the concession of the agents and the Pareto search method. In this section we focus on the latter aspect of the negotiations. By means of computer experiments we investigate the effectiveness and

Algorithm 1 The orthogonal-DF strategy

The following is given: (a) the opponent's last and before-last offer: $O_1 =< p_f(t), p_v(t) >$ and $O_2 =< p_f(t-2), p_v(t-2) >$ respectively, (b) an agent's utility function $u(< p_f, p_v >)$ and aspiration level $u'(t+1)$ at time $t+1$, (c) the step-size δ, and (d) the search direction $s_{dr} \in \{-1, +1\}$. Based on this information the agent's orthogonal-DF computes the next counter offer $O =< p_f(t+1), p_v(t+1) >$ by executing the following procedure:

1. Use the orthogonal strategy to compute $O'_1 =< p'_f(t), p'_v(t) >$ and $O'_2 =< p'_f(t-2), p'_v(t-2) >$, i.e., the points on the iso-utility curve with expected utility $u'(t+1)$ that lie closest to O_1 and O_2, respectively.

2. Compute d_1 and d_2, the distance of the opponent's last two offers, i.e., $d_1 = ||O_1 - O'_1||$ and $d_2 = ||O_2 - O'_2||$, where $|| \cdot ||$ denotes *Euclidian* distance.

3. Update s_{dr}: whenever $d_1 > d_2$ the orthogonal-DF "turns", i.e., $s_{dr} = -1 \cdot s_{dr}$, otherwise $s_{dr} = s_{dr}$.

4. Update δ: decrease δ whenever the orthogonal-DF turns. For a number of periods directly after a turn δ is not increased, and otherwise δ is increased (cf. [129] for the details).

5. Compute the counter offer $O =< p_f(t+1), p_v(t+1) >$: set $p_v(t+1) = p'_v(t) + \delta \cdot s_{dr}$. Next, calculate $p_f(t+1)$ such that $(p_f(t+1), p_v(t+1))$ lies on the iso-utility curve, i.e., $u(< p_f(t+1), p_v(t+1) >) = u'(t+1)$.

robustness of the orthogonal and orthogonal-DF approach, to find Pareto-efficient solutions for a wide variety of settings. We evaluate the robustness of the search strategy by experimenting with various concession strategies on the buyer side.

Although the system enables buyers to initiate several concurrent negotiation threads, within a thread the Pareto search strategy operates independently from the other threads. For researching the efficiency and robustness of Pareto search strategies it therefore suffices to consider only a single negotiation thread in the experiments. Furthermore, the bundle content in the experimental setup consists of a single category with a low quality of service. The experimental results generalise to negotiations involving multiple categories: only the shape of the iso-utility curves is affected by the number of categories. In the experiments the shape is varied using different parameter settings.

A general specification of the buyer agents and the seller agent was provided in Section 6.2. Sections 6.3.1 and 6.3.2 describe the agent settings which are specifically used within the experimental setup. In particular the agents' preferences and concession strategies are specified in detail in Sections 6.3.1 and 6.3.2 respectively. The experimental results are discussed in Section 6.3.3.

6.3.1 Agent preference settings

We simulate the negotiation with a variety of buyer and seller preferences, expressed by the agents' expected utility functions (see also Sections 6.2.2 and 6.2.1). In the experiments we consider only a single low-quality category. The number of articles e_b (we omit the index for clarity in the following) the buyer expects to read is assumed to be a constant, set randomly between 1 and 20 at the beginning of an experiment. The *buyer agent*'s expected utility therefore reduces to $u_b(< (p_f, p_v) >) = b_{max} - (p_f + p_v \cdot e_b)$. Note that this results in a linear iso-utility curve in the (p_f, p_v) plane (see e.g. Fig. 6.5). Furthermore, since the purpose is to demonstrate the efficiency of the final deals reached, we set the buyer agent's mandate b_{max} for the bundle such that an agreement is always reached.

The expected utility for the *seller agent* is based on e_s, the expected number of articles that the buyers will read *on average* in the (low-quality) category. In contrast to the buyer agent, the expectation is not a constant but a function of the variable price p_v. It is assumed that buyers who are willing to pay a high variable price are expected to read less than buyers with a low variable price (i.e. we assume the law of demand holds cf. [72]). In the experiments we use the linear function $e_s(p_v) = b - a \cdot p_v$ with $b = 20$ and a set randomly between 0.03 and 0.07 at the beginning of an experiment. Note that the seller agent's iso-utility curve is now convex (towards the origin).

6.3.2 Concession strategies

The buyers and the seller can each select their own concession strategies. Although a seller agent's concession in the main system can depend on the behaviour of all buyers (i.e., one-to-many), in the experiments the seller agent's strategy is simply to decrease the desired utility level or *aspiration level* with a fixed amount each round (more advanced strategies are considered in Chapter 7). The initial aspiration level is randomly varied. Note that the number of buyers and their behaviour does not affect the seller's concession when this strategy is used.

On the buyer side, on the other hand, we implemented four classes of concession strategies to investigate the robustness of the Pareto search strategy:

1. *Hardhead.* The buyer agent does not concede when this strategy is used; the aspiration level remains the same during the negotiations.

2. *Fixed.* A fixed amount c in utility is conceded each round.

3. *Fraction.* The buyer agent concedes the fraction γ of the difference between the current desired expected utility level and the expected utility of the opponent's last offer.

4. *Tit-for-tat.* This strategy mimics the concession behaviour of the opponent, based on subjective (expected) utility improvement. If the expected utility of the seller agent's offers increases, the same amount is conceded by the buyer agent. Note that the concession is based on an increment in expected utility perceived by the buyer agent. The seller agent's actual concession is shielded from the buyer agent, as an improvement could also occur when the seller agent moves along his iso-utility curve. Furthermore, note that the perceived expected utility improvement could also be negative. To make the concession behaviour less chaotic, however, no negative concessions are made by the buyer agent.

6.3.3 Results

The seller agent and the buyer agent in the experiments negotiate in an alternating fashion until an agreement is reached. The efficiency of the agreement is then evaluated based on the distance of the final offer from a Pareto-efficient solution. We measure an offer's distance from a Pareto-efficient solution as the maximum possible expected utility improvement for the *buyer* if a Pareto-efficient offer was made, all else remaining equal. This is achieved by moving the buyer's iso-utility curve until the obtained deal is Pareto-efficient.

To evaluate the quality of the results we compare the outcomes using various search strategies and concession strategies of the buyer agent. Table 6.1 provides

	Pareto-search strategy		
Concession strategy	Random	Orthogonal/DF	DF/DF
hardhead	18.92 (±23.56)	**8.03** (±11.44)	18.63 (±32.81)
fixed ($c = 20$)	26.52 (±34.49)	**10.43** (±17.34)	28.82 (±46.71)
fixed ($c = 40$)	38.91 (±49.72)	**16.21** (±23.84)	44.29 (±69.76)
fixed ($c = 80$)	42.12 (±56.88)	**25.61** (±38.72)	48.84 (±72.12)
fraction ($\gamma = 0.025$)	30.26 (±38.37)	**10.07** (±15.03)	32.25 (±52.81)
fraction ($\gamma = 0.05$)	31.53 (±40.00)	**11.52** (±16.16)	28.52 (±52.13)
fraction ($\gamma = 0.1$)	37.81 (±48.82)	**16.91** (±30.80)	26.28 (±42.20)
tit-for-tat	72.78 (±121.35)	**59.60** (±113.27)	56.64 (±116.82)

Table 6.1: Average distance from Pareto-efficient solution for various buyer concession strategies (rows) and buyer/seller search strategies (columns). Results are averaged over 500 experiments with random parameter settings. Standard deviations are indicated between brackets. Best results (see column *Orthogonal/DF*) are obtained if the buyer and seller agents use orthogonal search, and the seller agent's search is amplified with a derivative follower.

an overview of the results. The row labelled *Random* contains the outcomes when both seller and buyer agents use a random search strategy. This strategy selects a random point on the iso-utility curve.[5] The distance of the final offer (from the closest Pareto-efficient solution), when random search is used, lies between 1 and 3 percent of the total costs.

Although the inefficiency with random search is only small compared to the total costs, even better results are obtained when one bargainer (typically the buyer agent) uses orthogonal search and the other (the seller agent) uses orthogonal-DF (i.e., orthogonal search combined with a derivative follower). The results are shown in the column labelled *Orthogonal/DF* of Table 6.1. The improvements are considerable. The distance of the final offer as a percentage of total costs lies then, for almost all concession strategies, between 0 and 1. Only for the tit-for-tat strategy the distance lies around 1.8 percent. Notice that the Orthogonal/Orthogonal-DF strategy combination is also robust, as best results are obtained using this strategy, relatively independent of the concession strategy selected by the buyer agent.

Table 6.1 also shows the results if both buyer and seller agents use orthogonal-DF search (column *DF/DF*). These results are very similar to random, however. The derivative follower relies on a consistent response from the opponent to signal the right direction. If both use a derivative follower, this signal is distorted.

Notice that the average distance depends on the concession strategy used by the buyer. Although in individual cases Pareto-efficient agreements (with zero distance) are reached using the orthogonal/DF search, the average distance consistently shows

[5] Only the downward sloping part of the seller agent's iso-utility curve is used.

some (usually slight) inefficiencies, even when the buyer makes no concessions (i.e., the *hardhead* strategy). The reason for this is twofold. Firstly, the DF accelerates finding the efficient solution by making, at times, large steps on the iso-utility curve. At a certain point the algorithm passed the Pareto-efficient point, and then turns. This way the offers keep oscillating around the optimal point. If the concessions are sufficiently large, an agreement can be reached at a point which is less than optimal.

Secondly, the direction and step-size of the DF are based on changes in the Euclidean distance between the seller and buyer offers through time. The distance can be influenced by both concessions and movements along the iso-utility curve. As the opponent's iso-utility curve is unknown, the agents are unable to distinguish between the two. This can mislead the DF whenever concessions are very large. Two possible solutions are to make either small concessions, or have intervals with no concessions allowing the search algorithm to find the best deal.

Particularly *tit-for-tat* results in a relatively high inefficiency, because of the reasons described above. Recall that tit-for-tat uses a subjective measure of the opponent's concessions. In practice, the perceived utility increments are sometimes quite large, resulting in bursts of very large concessions. If this occurs near the agreement point this can result in inefficient outcomes.

To conclude, the orthogonal/DF strategy clearly outperforms other combinations of search strategies in the experiments. Inefficiencies still occur, especially if the concessions are large. A trade-off therefore exists between reaching an agreement fast (by making large concessions) and reaching an efficient agreement. Since concessions appear to influence the Pareto-efficiency of the outcomes, it is essential that a Pareto-search strategy is evaluated together with a concession strategy.

6.4 Related approaches

In this section some related approaches for multi-issue negotiations are discussed.

6.4.1 Fuzzy similarity criteria

Related to our work, in [33] a heuristic approach for finding win-win trade-offs between issues is introduced. Contracts which are similar to the opponent's offer are selected based on fuzzy similarity criteria, and given a desired utility level. They use fuzzy similarity criteria because most of the considered issues take on very limited discrete values. Based on these similarity criteria, an iterative hill-climbing algorithm is used to find the most similar offer. This hill-climbing algorithm is limited, however, to linearly additive utility functions.

By contrast, we consider negotiation over continuous issues (or issues that can take on many values). For this problem domain, *Euclidean* distance is a more natural choice than the similarity criterion. With Euclidean distance standard mathematical techniques (from fields such as convex analysis) are immediately at our disposal. Moreover, implementing the similarity criterion entails a straightforward application of standard techniques from numerical analysis. The orthogonal search method finds, from the collection of offers that have the desired utility level, the offer closest to the opponents last offer, measured in Euclidean distance. Unlike the heuristics developed in [33], our approach is not restricted to linearly additive utility functions.

A possible limitation of any search method using only a distance method to determine the counter offer, such as the orthogonal search method and the fuzzy similarity criteria described in [33], is that the rate of convergence depends to a large extent on the bargainers' preferences. As we found in our experiments, convergence rate is indeed often very slow (i.e., when both agents use the orthogonal strategy). Therefore, we amplified the search using a derivative follower, which can converge quickly to a Pareto-efficient solution. Slow rate of convergence is especially a problem whenever software agents are not a priori restricted but can search for clever trade-offs and at the same time make concessions (as is the case in our experiments). If the search method is too slow, very little improvement in the efficiency can be realised before a deal is closed. The developed orthogonal-DF, however, is sufficiently fast and consequently can also work very well in conjunction with concessions.

6.4.2 Intermediaries

In the literature the difficulties with bargainers simultaneously making concessions and searching for clever trade offs is generally avoided by assuming an intermediary [32, 62, 68, 101]. The mediator is inspired by the idea of a single negotiation text (SNT). SNT is a mediation device suggested by Roger Fisher [36]. During negotiation, the mediator first devises and proposes a deal (SNT-1) for the two bargainers to consider. The mediator is not trying to promote the first proposal, rather, it is meant to serve as an initial, single negotiation text; a version to be criticised by both parties and then modified in an iterative manner. Modifications to the SNT-1 will be made by the mediator based on the criticisms from the two sides. Thus, both parties need to reveal (aspects of) their preferences to the mediator, hence trust becomes an important issue. Furthermore, additional costs are often involved with a mediator.

The orthogonal-DF method is somewhat related to the work of Ehtamo et al. [32]. They develop the method of improving directions which is a mathematical formalisation of the SNT method (with a mediator). In essence it is a multi-criteria decision making gradient search method. Given a SNT, bargainers give their most preferred direction of the next SNT which is just the gradient. The mediator then uses some relatively straightforward procedure to determine the jointly improving direction which is then used to determine the next SNT. The orthogonal-DF also searches for such a jointly improving direction, but without the use of a mediator, however.

6.4.3 Auctions

Another approach increasingly used to automate one-to-many negotiations is through auctions. Although our system has characteristics similar to those of auctions, bundling and negotiation of information goods have distinct properties which impede the use of current available auction designs. Mainly, information goods have negligible incremental reproduction and distribution costs [18]. The supply of goods can therefore be virtually unlimited. Auctions, however, are more suitable when resources are scarce.

Furthermore, information goods can be easily packaged in a wide variety of configurations, resulting in multi-dimensional products and pricing schemes. Personalisation of information goods then becomes a key component of a successful electronic business strategy [2]. As illustrated in this chapter, and in Chapters 3 and 4, a bilateral approach can be naturally used to perform multi-issue negotiations. Traditionally, auctions have focused on price as the single dimension of the negotiation. Although *multi-attribute* auctions have recently received increasing attention [28, 93], the agents are usually required to reveal their complete preferences. Moreover, the focus is on obtaining Pareto-efficient outcomes and profits are usually not considered. In case of unlimited supply, however, such auctions may fail to provide sufficient profit for the seller. Because of the disentanglement of the concession and Pareto-search strategies, the profits can be regulated by the concession strategy (this issue will be further addressed in the next chapter).

A seller may also have business-related considerations for preferring a bilateral bargaining approach. For example, the bilateral bargaining protocol allows for much flexibility and can be easily applied in case of continuous sales. Using bargaining, new buyers can enter the negotiation at any given time, and buyers can obtain the good at any time by simply accepting a seller's counter offer.

6.5 Discussion

6.5.1 The system revisited

Although the focus of this chapter is the problem of selling bundles of news items, other types of (information) goods can also be sold through the developed system. A key question for extending the use of the system to other application areas is, however, if buyers and (to a lesser degree) sellers are willing to have software agents automate the actual bargaining. A prerequisite would be that the traded goods have a relatively low value and transactions are conducted frequently. Consequently, the risks are low and an agent has many opportunities to learn from past experience and gradually improve performance. Note that the negotiation procedure of the system does not require both seller and buyer to use the same level of automation.

Depending on the particular application of the system, it may be desirable for the buyer to rely more or less on the assistant of the software agent.

An additional important aspect of the relevance to other application areas is the potential benefit of using such a system. The developed system appears particular suitable for selling complex goods with a high degree of personalisation and relatively rapidly changing preferences (as is the case with the news items). More specifically, within the system personalisation entails discriminating between buyers based on the bundle price and the quality of service. Second-degree price discrimination is the economic term for this type of personalisation.

In second-degree price discrimination the price depends on the quantity and/or quality of the purchased good. The distinguishing aspect of second-degree price discrimination is that buyers can *self-select* the best purchase. Traditionally, buyers are offered a menu of price combinations. The work of [18, 59] discusses algorithms which, given a particular pricing scheme, learn the best price combinations on-line. They conclude that (especially in a dynamic environment) complex schemes are generally not the most profitable due to the need of more learning.

The distinguishing aspect of the developed system is that instead of having explicit pricing schemes, buyers can bargain for the most appropriate bundle/price combination. This can result in a similar (or even higher) degree of discrimination between buyers as with explicit complex pricing schemes. In the absence of an explicit structure the seller is, however, more flexible in the degree to which she discriminates. The seller does not have to a priori limit the complexity of the pricing scheme. Whenever bundles of (information) goods are being offered, an additional advantage is that, by initiating the negotiation process, buyers can explicitly express their interest in a particular bundle of goods. This may facilitate the process of offering buyers the appropriate bundles (and consequently it may facilitate the indirect discrimination between buyers).

6.5.2 Bargaining and Pareto efficiency

In the system the seller agent uses the orthogonal-DF as the Pareto search strategy combined with a concession strategy. The concession strategy determines the next concession relatively independently of the ongoing bargaining process with a particular buyer. The idea is that, on the one hand, bargaining with a particular buyer should lead to finding the best possible deal for both parties, given the seller's desired expected utility level. That is, the bargaining outcome should closely approximate a Pareto-efficient solution. On the other hand, the one-to-many aspect of the bargaining process (i.e., bargaining with more than one buyer) should guide the updating of the concession strategy. Thus the seller uses the disentanglement of the bargaining strategy (in a concession and Pareto search strategy) to distinguish explicitly between the one-to-many and one-to-one aspects of the bargaining process.

The experiments in Section 6.3 show that if a buyer agent uses an orthogonal strategy as the Pareto-efficient search strategy then the bargaining outcomes will closely approximate a Pareto-efficient solution. The experiments are conducted for a variety of (buyer) concession strategies, buyer preferences, and seller preferences. Based on the experimental results we can conclude that any other strategy choice of a buyer will probably result in less efficient outcomes. Moreover, such a strategy will not influence the concession strategy of the seller (due to the independence of the concession strategy). Consequently, any alternative bargaining strategy of the buyer is probably at most as good as the orthogonal strategy combined with a concession strategy that mimics the concessions of the alternative strategy. Thus, given the seller's choice of the orthogonal-DF combined with a relatively independent concession strategy, it is in a buyer's best interest to choose the orthogonal search strategy combined with a concession strategy. Moreover, this choice results in (a close approximation of) a Pareto-efficient solution.

6.6 Concluding remarks

We introduce a novel system for selling bundles of news items in this chapter. Through the system, buyers bargain over the price and quality of the delivered goods with the seller. The advantage of the developed system is that it allows for a high degree of flexibility in the price, quality, and content of the offered bundles. The price, quality, and content of the delivered goods may, for example, differ based on daily dynamics and personal interest of buyers.

The system as developed here can take into account business related side-constraints, such as "fairness" of the bargaining outcomes. Fairness ensures that buyers with similar preferences are treated in the same fashion. Because of fairness, the actual bargaining process between seller and buyers is not really bilateral, but is in fact one-to-many since the bargaining process with one buyer can have an impact on a simultaneous bargaining process with another buyer.

Autonomous software agents perform (part of) the negotiation on behalf of the users of the system. To enable efficient negotiation through these agents we decompose the bargaining strategies into concession strategies and Pareto-search strategies. Moreover, we introduce the orthogonal and orthogonal-DF strategy: two Pareto search strategies. We show through computer experiments that the respective use of these two Pareto search strategies by the two bargainers will result in very efficient bargaining outcomes. Furthermore, the system is set up such that it is actually in the best interest of the buyer to have their agent adhere to this approach of decomposing the bargaining strategy into a concession strategy and Pareto search strategy.

Chapter 7

Bargaining strategies for one-to-many bargaining

Through the use of autonomous agents a business can obtain flexibility in prices and goods, and distinguish between different groups of buyers based on their preferences. The previous chapter showed how personalisation of goods in the context of information goods can be achieved using automated negotiation. In this chapter,[1] we focus on the (expected) utility obtained by a seller agent and how different groups of buyers can be targeted having different valuations for obtaining the goods. We consider agent strategies for a one-to-many bargaining setting, where a seller agent negotiates, as before, with many buyer agents simultaneously in a bilateral fashion. We focus on domains where the supply of goods is flexible and new goods can be reproduced quickly, at relatively low costs. Such characteristics apply not only to information goods, but may also apply to other retail markets. As in the previous chapter, the strategies also take into account a notion fairness that is important for maintaining customer satisfaction and acceptance of the system by customers. Fairness ensures that buyers are treated in a similar fashion and is comparable to the notion of envy-freeness in auctions [44] (see Section 6.1.2 for further details).

In many cases, auctions can be used to effectively organise one-to-many bargaining. Depending on the setting, auctions can provide buyers with the incentive to reveal their preferences truthfully, and to allocate the goods efficiently (see also Section 2.3.5). For various situations, however, auctions may not be the preferred protocol for bargainers. In situations of, for example, virtually unlimited supply, multiple issues, and/or continuous sale the appropriate auction protocol becomes,

[1]This chapter is based on [40]: E.H. Gerding, K. Somefun, and J.A. La Poutré. Bilateral bargaining in a one-to-many bargaining setting. Technical Report, CWI, Amsterdam, to appear. A shorter version has been accepted for publication as [39]: E.H. Gerding, K. Somefun, and J.A. La Poutré. Bilateral bargaining in a one-to-many bargaining setting. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS2004), New York City, New York*. IEEE Computer Society Press, 2004.

at best, much more complex. Consequently, businesses may opt for the intuitive and flexible bilateral bargaining protocol, where the seller agent negotiates bilaterally with one or more buyers simultaneously by exchanging offers and counter offers. These motivations are more closely considered in Section 6.4.3 of the previous chapter.

Only little work has been done to study actual strategies for one-to-many bargaining. A few related papers study concurrent bilateral negotiations within a one-to-many setting [85, 98]. In these papers, a framework is described where a buyer negotiates with several sellers simultaneously to find a single best deal. This differs from our setting, however, since the seller in our system can come to an agreement with many buyers as we assume that supply is flexible. The various negotiations in our case are nevertheless related mainly through the notion of fairness.

For the case of virtually unlimited supply, as for information goods, we present a number of one-to-many bargaining strategies for the seller in this chapter, that take into consideration the notion of fairness. In particular, we introduce auction-inspired strategies that achieve good results. We compare the performance of the bargaining strategies using an evolutionary simulation, especially for the case of impatient buyers. These experiments show that the auction-inspired strategies are able to extract almost all the bargaining surplus, given sufficient time pressure of the buyers. The auction-inspired strategies benefit from the fact that the setting is one-to-many, even though bargaining occurs in a bilateral fashion.

This chapter is organised as follows. In Section 7.1 we discusses the bargaining setup and the strategies used by the seller agent. In Section 7.2 we introduce the simulation environment used for testing the performance of the strategies. We present the simulation results of the conducted computer experiments in Section 7.3. Conclusions follow in Section 7.4.

7.1 One-to-many bargaining

Bargaining is performed using the bilateral bargaining protocol described previously in Section 6.1.3. Although the protocol allows for multiple issues to be negotiated simultaneously, we concentrate on single-issue bargaining (e.g. the price) in this chapter and consider the (expected) utility obtained by the agents in the system. The multi-issue aspect is addressed in Chapter 6. We assume here that buyers are impatient and have an incentive to reach agreements early. The buyers' time pressure is further discussed in Section 7.1.1.

An agent representing a business can be endowed with various bargaining strategies. We present a number of strategies for the seller agent in Section 7.1.2. These bargaining strategies take into account a notion of fairness, such that different buyers are treated equally whenever that seems fair. For a detailed description of the fairness concept applied by the seller agent, we refer to Section 6.1.2.

We note that the reader is assumed to be familiar with the contents of Sections 6.1.2 and 6.1.3 of the previous chapter in the following.

7.1.1 Time pressure

An important assumption is that buyers are impatient and prefer an early agreement. Time pressure or time impatience is a common assumption in bargaining, e.g. [110] (see also Section 2.3.2). The seller is simultaneously and continuously negotiating with many buyers and is therefore less concerned with immediately reaching an agreement for a particular bargaining outcome, i.e., he is relatively patient. Furthermore, we assumed earlier that the seller can reproduce the offered goods quickly and at low costs. Therefore, a seller can respond timely to the demand and with little additional costs for matters such as storage of the goods. We model this relative time patience by assuming that the seller, unlike the buyers, has no time pressure.

At least in theory, the seller can benefit from buyers' time-pressure by introducing a delay before submitting a counter offer. An important question is then which bargaining strategies can most effectively utilise these potential benefits. Experimental results discussed in Section 7.3 show that auction-inspired strategies, which we will present in the next Section, are very effective: depending on the time pressure, they are capable of extracting very large shares of the *bargaining surplus* (see Section 1.1.2) for the seller.

7.1.2 Bargaining strategies

The challenge is to develop bargaining strategies for the seller that maximise expected utility by utilising differences in buyers' willingness to pay without violating the fairness constraint. Instead, these strategies make use of differences indirectly through buyers' time pressure. In order to benefit from time pressure all the strategies discussed below introduce a (fixed) delay before the seller agent submits a counter offer.

Fixed and time-dependent threshold strategies

For purpose of comparison we introduce a fixed threshold strategy, where the the seller's desired expected utility level or *aspiration level* (see Def. 113.1) remains constant through time. The seller only accepts offers above the aspiration level and counter offers always have an expected utility level equal to the aspiration level. Whenever the seller agent accepts two different offers within a certain time interval, the bargaining outcome may be unfair. To rule out an unfair outcome, the seller agent immediately engages in post-agreement negotiation with all buyers from which it accepted an offer. During these negotiations the seller agent offers a buyer agent an improved offer where the expected utility corresponds with the seller agent's aspiration level. In case of multiple issues, the seller will concede on one or more issues, and not change the value of the remaining issues.[2]

Clearly, the fixed threshold strategy is not capable of utilising buyers' time pressure. The purpose of the strategy is to provide some insights in the minimal extractable profit, given strategic behaviour of the buyers.

The second strategy we consider is a time-dependent threshold strategy: the current aspiration level depends on time. The aspiration level only changes from one period to the next. Again, the seller only accepts offers above a (time-period dependent) aspiration level and counter offers are always equal to the current aspiration level. As before, the seller agent immediately engages in post-agreement negotiation in case the seller accepts a buyer's offer to ensure fairness.

Unlike the fixed-threshold strategy the time-dependent strategy is capable of utilising buyers' time pressure. Its success, however, depends on how much it knows about buyers' preferences, or how easily buyer preferences can be learned, in relation to time-based pricing strategies.

Auction-inspired strategies

We introduce a bargaining strategy which is inspired by the first-price auction. The auction-inspired strategy operates as follows. The seller agent collects all offers submitted within a certain fixed time interval, after his last offer. Then it sets the aspiration level to the *current highest utility level*, which is equal to the best offer from the collection of offers. It accepts all offers equal to the current aspiration level and counters the unaccepted offers by setting the counter offer's expected utility equal to the current aspiration level. The strategy introduces a fixed time delay before countering the unaccepted offers. Note that because the auction-inspired strategy only accepts offers with the same expected utility within a certain time interval, post-agreement negotiation is not necessary to ensure fairness.

The success of the auction-inspired negotiating strategy does not depend on some (a priori) knowledge of buyer preferences, unlike the fixed and time-dependent strategies. Intuitively, buyers who, due to time pressure, suffer more from delay are inclined to bargain less "hard-headed" than other buyers. Consequently, these buyers may reach a deal sooner and pay more. Thus, at least potentially, the strategy is capable of utilising buyers' time pressure without requiring (a priori) knowledge of buyer preferences. Unlike auctions, actual bargaining occurs in an alternating exchange of offers and counter offers, typically initiated by a buyer. Even though the seller's strategy can be auction-inspired, buyers will be unaware of this fact. They do not know the opponent's bargaining strategy on forehand; they

[2] Note that this approach assumes that the agents have conflicting interests on individual issues and nonsatiation of buyers (i.e., buyers always prefer more than less). If this is not necessarily the case, a weaker form of fairness can be used instead, where the seller tries to improve the offer to the best of his knowledge.

perceive the bargaining process to be bilateral. Buyers may of course suspect some relationship with other ongoing negotiations. The point is that unlike a true auction the relationship with other simultaneously submitted offers is not specified up front, through a set of rules.

Reservation value

A drawback of the auction-inspired strategy is that it becomes vulnerable whenever groups of buyers experience very little time pressure. Without time pressure, buyers have no incentive to buy soon and could independently decide to initially submit very low offers; consequently profits will be very low. To circumvent this we also consider auction-inspired strategies with a *reservation value* (i.e., a lowest acceptable utility level). A seller agent is never willing to sell below the reservation value. This means we alter the earlier definition of the current highest utility level. It now becomes the maximum of the reservation value and the utility of the best offer from the offers collected within a certain time interval. An interesting advantages of introducing a reservation value occurs when some but not all buyers experience very little time-pressure. The auction-inspired strategy can then still utilise the time-pressure of the other buyers.

We consider two approaches for determining the reservation value. Either the reservation value is fixed, like the fixed-threshold strategy, or it is time dependent, like the time-dependent threshold strategy. Thus the auction-inspired strategy with a reservation value is actually a combination of the auction-inspired strategy (without reservation value) and either the fixed or time-dependent threshold strategies.

7.2 Bargaining simulation environment

We apply a simulation environment in order to evaluate the performance and robustness of the above negotiation strategies against many learning buyers. The agents in the simulation are assumed to be boundedly rational: they can learn and adapt their strategies by a process of trial and error, and they do not know the seller's strategy. The bargaining process is repeated many times, enabling buyers and the seller to learn from past interactions. An evolutionary algorithm is used to model the learning aspect of the agents. A similar approach was used in previous chapters (Chapters 3-5).

7.2.1 The bargaining game

The seller agent negotiates with many buyer agents simultaneously by alternating offers and counter offers, where the buyer agents initiate the negotiations. For our simulations we set a maximum number of r rounds, where r is set sufficiently large such that it has no significant impact on the results. At the start of the negotiation, buyer agents submit their offers to the seller agent, which responds by either accepting an offer or sending a counter offer in the next round. Offers consist of a single issue, viz. the *price* of the negotiable good. Negotiation continues after all buyer agents have reached an agreement or the maximum number of rounds is reached, which concludes a so-called *bargaining game*. We note that buyer agents in the simulation do not leave the negotiations or enter later.

We assume that, since buyers are impatient, buyer agents in the simulation will respond to the seller agent's counter offer without delay. This is modelled by having the buyer's counter offer occur in the same round as the seller's counter offer.

7.2.2 Buyers and their agents

Buyers are interested in buying at most one unit of the offered good in each bargaining game. They can have different preferences regarding their time pressure and valuation of the good, which together characterise the buyer *type*. For the analysis we assume buyers can be grouped into a finite number of k types. The number of buyer agents of each type participating in a negotiation game varies randomly and is unknown to the seller agent. The seller agent is also uninformed about the identity or type of a specific buyer agent. The actual number of participants *of each type* is determined independently by a Poisson distribution with average λ.

A buyer agent tries to maximise a given utility function for buyer type i, u_i, which is defined as follows:

$$u_i = (v_i - p)\delta_i^t, \tag{7.1}$$

where v_i is the buyer's valuation of the good, p is the negotiated price, δ_i is the discount factor used to model the time pressure, and t is the negotiation time. In the simulation negotiation occurs at fixed time intervals. Therefore, δ is the discrete representation of time pressure and t therefore also indicates the negotiation round. Note that discount factors are commonly used for modelling time pressure, e.g. in the Rubinstein-Ståhl alternating-offers model (see Section 2.3.2). The agents are furthermore assumed to be *individually rational* (see Def. 4.2): they will not bid nor accept offers with a negative utility.

Within the simulation, buyer agents are endowed with adaptive time-based strategies to produce offers and evaluate the seller's offers. Although this is a relatively simple strategy, the adaptive nature of the strategies provides buyer agents with sufficient flexibility to bid effectively in the long run. A strategy consists of a piece-wise linear function, which determines the price level of new offers and is also used as threshold to accept or reject the seller's offers: if the seller's offered price is above the threshold, the offer is accepted, otherwise the offer is rejected. A post-agreement offer is automatically accepted if this is beneficial for the buyer.

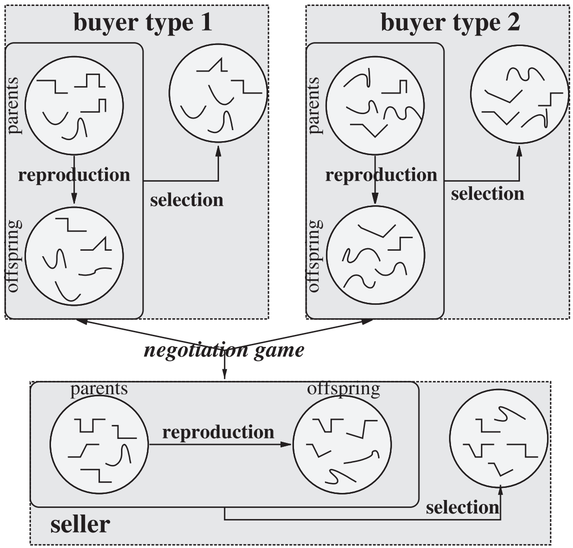

Figure 7.1: The EA cycle for negotiations with two buyer types and an adaptive seller

We also applied an extended strategy in our experiments, where the threshold and offers are determined by separate piece-wise linear functions. The separation of the two functions enhances the bargaining capabilities of the buyer agent. Results using the two representations are very similar. The outcomes presented in this chapter are obtained using the extended strategy.

7.2.3 Seller agent

The seller agent bargains with a number of buyers simultaneously, without knowing the type of these buyers. The seller agent's utility is equal to the total utility or *profit* obtained over all buyers (recall from Section 7.1.2 that the we can assume the seller has no time pressure). Production costs are set to zero.

We consider five strategies for the seller agent: fixed threshold, time-based threshold, auction-inspired strategies and two combined strategies. The first two strategies and the combined strategies are *adaptive*: strategies that maximise total utility are learned using an EA. The time-based threshold strategy is similar to the strategy used by the buyer.

7.2.4 The evolutionary system

Evolutionary algorithms (EAs) are used to produce effective bargaining strategies for the buyer agents and seller agent. The implementation used is described in detail in Section 1.2. The strategies for buyer agents of each *type* are produced by separate EAs, which operate in parallel. Furthermore, a separate EA can also be used to produce strategies for the seller agent, in case the seller uses an adaptive strategy. An example of the evolutionary system with two buyer agents and an adaptive seller agent is depicted in Figure 7.1. Note that, whereas in previous chapters a single EA was used with several evolving populations, the current implementation applies several (independent) EAs. This enables for instance the seller agent to use a different strategy representation than a buyer agent.

The fitness of the strategies is determined by the average utility obtained in a number of bargaining games, which go as follows. At the start of each bargaining game, the number of participating buyer agents of each type is determined randomly using a Poisson distribution as described above. Buyer agents are then assigned a randomly selected strategy from either the parent or offspring population of the corresponding type. Similarly, a strategy is selected randomly for the seller agent (in case of an adaptive seller). The bargaining game is played for a fixed number of times, re-establishing the number of buyer agents and assigning new strategies at the start of each game.

Strategy encoding

As mentioned in Section 7.2.2, the buyer agent's strategy consists of two piece-wise linear functions: an offer and a threshold function. The functions are encoded using real values, where each bending point of a function is encoded by two real values. Additionally, two end points mark the values for the first and last rounds. For example, 8 real values are needed to encode a pair of functions with two line pieces each.

The same representation is used for the seller agent if he uses a time-based threshold strategy. If a fixed threshold is used, only a single real value is needed to encode this. Note that the seller agent uses the same function for both the threshold and for producing offers.

7.3 Experimental results

This section reports on computational experiments using the bargaining simulation environment.

7.3.1 Settings

The following settings are used for the experiments reported in this chapter. Buyers are grouped into three types, each type having adaptive bargaining strategies evolving in separate populations. The time pressure (discount factor) for each type is set as a control parameter. A type's valuation, on the other hand, is randomly selected from a uniform distribution at the beginning of each experiment. In order to make sure that all types have different valuations, the valuation of type 1 is selected between 0 and 1000, type 2 between 1000 and 2000, and type 3 between 2000 and 3000.

The piece-wise linear functions of the buyer agents, and of the seller agent in case of time-based threshold strategy consist of two line pieces. The number of buyers of each type participating in a bargaining game is determined randomly by a Poisson distribution with the average $\lambda = 10$. The length of a bargaining game is set to 40 rounds.

The EA settings are chosen such that results are robust and the EAs are able to find good solutions. All buyer types use equal settings, with 20 strategies in the parent populations and 20 offspring strategies. An exponential decay model is selected to determine the mutation standard deviation (see Section 1.2.3). The mutation standard deviation is initially set to 0.2, and decays with a half-life value of 50 generations. The EA settings for the seller are the same, except that each seller population only contains 10 strategies. Buyers have larger populations because more buyers than sellers participate each game, and because in case of the extended buyer strategy (with two functions) the search space for the buyer is larger (a higher population size is often recommended for larger search spaces). The fitness of the strategies for a single generation is determined by 200 bargaining games. The EAs using these settings are able to find almost optimal solutions for simple test cases.

7.3.2 Results

The reported results are obtained after a process of learning, when the strategies have converged. It is important to note that during learning, the preferences of the buyers remain unchanged, although the number and composition (i.e., number of each type) of buyers can differ in each bargaining game. Experiments are run for 100000 bargaining games (500 generations). Results are averaged over the last 1000 bargaining games of an experiment, and over 30 experiments, accounting for random settings such as the number of participating buyers and the buyer valuations.

The performance of the strategies is evaluated by comparing the fraction of *bargaining surplus* or just surplus (see Section 1.1.2) obtained by the seller agent. Since the seller benefits from any positive agreement (there are no costs for the seller, see Section 6.2.1), the bargaining surplus in this case is equal to the buyer's valuation of the good. Figure 7.2a compares the seller's obtained fraction of surplus for different seller strategies and buyer discount factors, where the buyer types have equal discount factors. The average round an agreement is reached is shown in Fig. 7.2b. The results when buyers have different discount factors are shown in Table 7.1, where δ_i denotes the discount factor for buyer type i and the strategy numbers correspond to

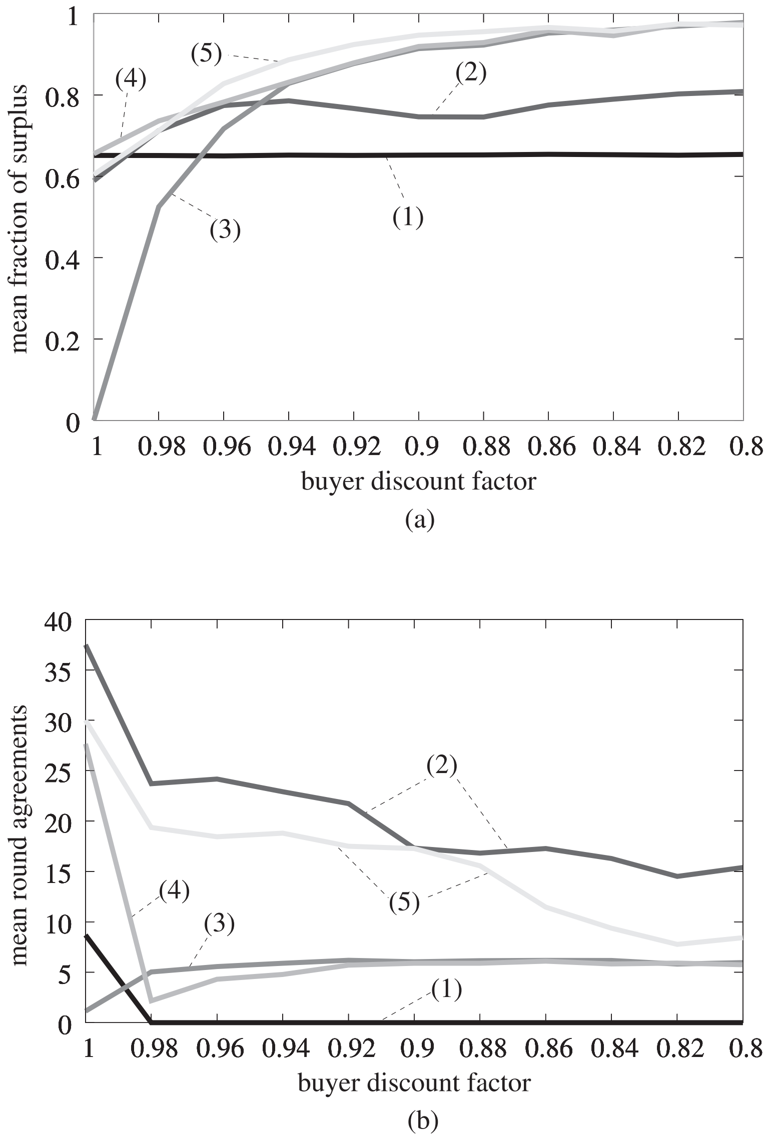

Figure 7.2: Seller's obtained fraction of bargaining surplus (a) and average round of agreement (b) using 5 bargaining strategies: (1) fixed threshold, (2) time-based threshold, (3) auction-inspired, (4) combined (3) and (1), and (5) combined (3) and (2).

Fig. 7.2. As shown in Fig. 7.2a and Table 7.1, a fixed threshold strategy (1) is able to extract around 65% of the surplus. Note that the outcomes are independent of the discount factor. Clearly, the fixed threshold strategy is unable to benefit from the buyers' time pressure.

The time-based threshold strategy (2), on the other hand, shows that higher profits can be obtained if the threshold changes in time, see Fig. 7.2a and Table 7.1. Buyers with a high valuation will settle for an agreement relatively early, since waiting for a better deal does not compensate the loss due to time discounting. Buyers with a low valuation, on the other hand, have the incentive to reach an agreement in a later stage if they can get a better price for it. This way the seller can indirectly discriminate between buyers with different valuations and time pressures. A disadvantage, however, is that this leads to much efficiency loss due to delayed agreements. Figure 7.2b and Table 7.1 show that the average round in which an agreement is reached is relatively high when a time-based strategy is used, resulting in a lower final expected utility for the buyers.

Note that with no time discounting (i.e., when the discount factor is 1) the fixed threshold strategy performs better. This is due to the difference in strategy complexity: only a single value needs to be optimised in case of a fixed threshold, whereas an entire function (encoded by 4 values) needs to be learned in case of the time-based threshold. This is clearly more difficult, especially within a dynamic environment with learning buyers.

Outcomes using the auction-inspired bargaining strategies (see Fig. 7.2a and Table 7.1, strategies (3),(4) and (5)) show an impressive increase in the fraction of surplus when buyers are impatient. If the time pressure becomes sufficiently high, the seller obtains almost the entire surplus. Even for lower time pressure, results are much better for the seller compared to the fixed and time-based threshold strategies. For the case of no or very low time pressure, the results also show that simple auction-like mechanisms are not sufficient in case of unlimited supply. Without competition between buyers, the market price goes to cost level, resulting in a zero profit for the seller. This problem can be resolved in bargaining by combining the auction-inspired strategy with an adaptive reservation value. As shown in Fig. 7.2a, this results in very good outcomes, even if buyers are very patient. This makes the combined strategy very versatile. We note that these outcomes also generalise to settings where different buyer types have different time preferences, assuming that buyers with higher valuation have a higher time pressure, as indicated by Table 7.1.

7.3.3 Bargaining revisited

An important aspect of the bargaining protocol is the ability of the seller to produce counter offers in the next round of bargaining. The auction-inspired strategy only accepts the highest offers in each round. Usually, only a single bid will be accepted

strategy	δ_1	δ_2	δ_3	round	fraction
1	1.0	0.95	0.90	0.88 ± 3.28	0.65 ± 0.09
2	1.0	0.95	0.90	23.94 ± 3.92	0.82 ± 0.08
3	1.0	0.95	0.90	5.75 ± 0.93	0.74 ± 0.08
4	1.0	0.95	0.90	11.88 ± 5.63	0.90 ± 0.06
5	1.0	0.95	0.90	18.05 ± 7.28	0.91 ± 0.03
1	0.95	0.90	0.85	0.00 ± 0.01	0.65 ± 0.08
2	0.95	0.90	0.85	22.44 ± 5.07	0.82 ± 0.05
3	0.95	0.90	0.85	6.29 ± 0.86	0.91 ± 0.04
4	0.95	0.90	0.85	6.75 ± 2.36	0.93 ± 0.03
5	0.95	0.90	0.85	17.02 ± 5.90	0.95 ± 0.03

Table 7.1: Average round an agreement is reached (column "round") and seller's fraction of surplus (column "fraction") when different buyer types have different discount factors and for different seller strategies (the strategy numbers correspond to the strategies of Fig. 7.2).

due to differences in the buyer agents' strategies. Even if buyers are of the same type, small differences remain because of mutations. The outcome where all buyers make the same offer is therefore unstable (this can be compared to e.g. "trembling hand" in game theory, where players are assumed to make small mistakes when executing their strategies). This would result in large inefficiencies because of delays. The counter bid in the next round, however, enables remaining buyers with similar valuations (i.e., of the same type) to accept the seller's bid (albeit with a certain time delay). This way, all remaining buyers of the same type can reach an agreement within a single round. Results (see Fig 7.2b and Table 7.1) show that, in fact, buyers reach agreements on average in the 6th negotiation round when the auction-inspired strategy is used, even though on average 30 buyers participate in each negotiation. This is much more efficient than e.g. the time-based threshold strategy.

A possible disadvantage of producing counter offers by the seller is that buyers could bid very low, and then accept the counter offer of the seller. Such a strategy could be beneficial in case the seller's counter offer is influenced by the buyers' offers, as with the auction-inspired strategies. This could then result in low profits for the seller. To see if indeed buyers profit by using such a strategy, the strategy representation for buyers was extended by separating the functions for producing offers and determining the threshold (see Section 7.2.2). Even with separated function, however, the auction-inspired strategy performs very much in favour of the seller (as shown by the results). This occurs because the counter offer is delayed by the seller, although agreements occur without delay, providing the buyers with an incentive to try and get an agreement immediately.

7.4 Concluding remarks

In this chapter, we consider strategies for a seller agent who negotiates with many buyers simultaneously in a bilateral fashion. These strategies respect a notion of fairness such that buyers are treated similarly. An important assumption is that buyers are impatient and prefer early agreements. Furthermore, buyers can have different valuations and time preferences. A buyer's actual valuation and time preference is only known to himself (i.e., a buyer's type constitutes private information).

We investigate several seller strategies for bilateral bargaining in a one-to-many setting, and introduce several auction-inspired strategies. Five different seller strategies are evaluated and compared: (1) fixed threshold, (2) time-dependent threshold strategies, (3) auction-inspired, (4) auction-inspired with fixed reservation value, and (5) auction-inspired with time-dependent reservation value. The last two strategies are actually a combination of the auction-inspired strategy with the first two strategies.

We use an evolutionary simulation to analyse the performance of the different strategies. The buyers' bargaining strategies adapt and learn through the use of an evolutionary algorithm (EA). The seller's strategies (1) and (2), and the combined strategies (4) and (5) also adapt and learn using an EA. The auction-inspired strategy (3), on the other hand, determines the threshold value based on the offers received by the buyers, and does not require any learning.

The auction-inspired strategies appear to be very successful in utilising the time pressure and consequently extract a very high share of the surplus. For sufficiently high time pressure, the seller obtains approximately the entire surplus, indicating that buyers almost bid their valuations. This is achieved without much delay. Thus buyers self-select to pay their valuation, while the bargaining outcomes respect our notion of fairness. The results also show superior performance of the combined strategies (4 and 5) compared to the auction-inspired strategy (3), in case some or all buyers have very little time pressure. In other words, the combined strategy is very versatile.

Chapter 8

Discussion and conclusion

We investigated both fundamental aspects of bargaining and introduced real-world business applications of bargaining using autonomous agents in this thesis. We applied computational simulations to analyse various situations of bargaining that are difficult to approach mathematically, and demonstrated the feasibility of the suggested applications. The agents in these simulations are not assumed to be completely rational, but rather they learn by doing, and adjust their bargaining policies based on feedback from interactions with other agents. Complete rationality is usually not realistic for actual multi-agent systems, mainly for two reasons. Firstly, agents may not have sufficient time and/or computational power to find optimal or rational outcomes. Secondly, in a multi-agent system with different agents programmed by different parties, one cannot rely on the other agents to act rationally. Nevertheless, game-theoretic or “rational” outcomes serve as a useful benchmark to validate our computational approach.

Evolutionary algorithms (EAs) are used in this thesis to govern the adaptive behaviour of the agents in the computational experiments. EAs are increasingly being used to model societies of learning computational agents and humans, especially within the field of agent-based computational economics. As shown in this thesis, EAs can be used effectively for bargaining both in case of population learning, where several agents select their strategies from a common strategy pool, and individual learning, where genetic material is not exchanged between agents. A possible drawback for using EAs in practice is that off-the-shelf implementations of EAs may require many fitness evaluations before converging to good solutions. If such evaluations are expensive or limited, e.g. when each bargaining game involves large sums of money, a more specialised approach may be required. The applications discussed here, however, mainly involve relatively small-risk transactions that are repeated frequently. Nevertheless, many solutions for learning using limited evaluations already exist in the literature which can be used for high-risk applications. However, it is beyond the scope of this thesis to discuss such approaches in detail.

In order to validate our evolutionary approach, we first compared experimental results to game-theoretic outcomes for relatively simple cases. In Chapters 3 and 4 we considered the game-theoretic subgame-perfect equilibrium as the benchmark for the bilateral bargaining game. Interestingly, in many cases the emerging behaviour of the evolutionary system did coincide with "rational" or game-theoretic behaviour in the long run. In Chapter 5, a validation was carried out for market setting where evolving agents learn to bid in a second-price auction. The second-price auction provides bidders with the incentive to bid truthfully in case of independent valuations. This outcome was indeed found in the evolutionary simulation.

After validating our experimental approach, we applied the evolutionary simulation to analyse situations which are hard or unwieldy to analyse theoretically. In Chapter 3 the agent model was extended with a fairness norm for multi-issue negotiations. The evolutionary outcomes showed that the surplus is more evenly divided when the fairness norm is applied, and that these outcomes are relatively insensitive to the fairness curve if the norms are consequently applied in each negotiation round. If the Pareto-efficient frontier is asymmetric, different types of agreements are reached in the various rounds. Chapter 4 considered a different extension, where agents can have additional opportunities if negotiations fail. Each agent is characterised by her state, denoting the number of opportunities remaining. If the agent's state is common knowledge, the number of opportunities only has a slight impact on the division of the surplus in the simulation. If this information is only privately known, however, the division of surplus reverses if the number of opportunities is sufficiently large and equal for both players. Further extensions, such as the influence search costs and uncertainty about future opportunities were also analysed using the evolutionary framework.

The power of evolutionary algorithms for analysing complex behaviour was also demonstrated in Chapter 5 in a market setting. A framework was presented for selling consumer attention space or "banner space" to the highest bidders (suppliers) in an auction. The value of the attention space is not a-priori known to the bidders and can only be learned with consumer feedback. This value depends on the profile of the consumer, but may also be influenced by other banners which are shown concurrently. Such a setting involving multiple goods, complex interdependencies, and uncertainty in the valuation of the goods is difficult to analyse theoretically. A computational simulation with evolving bidding agents was therefore applied to demonstrate the feasibility of the approach and to compare the effectiveness of various auction designs. With no interdependencies, the adaptive suppliers could accurately learn the profile of the consumers. In case of interdependencies, however, the suppliers also need to take into account the effect of competitive banners. Interestingly, the agents in the framework then learn to target specific niches in the market. The performance of the system, i.e., if a good match between consumers and suppliers is found, appeared to depend on the auction rules in case of interdependencies. Results indicated that the so-called *next-price* auction (where goods are sold at the price of the next-highest bidder) performed well in general.

An important advantage of bargaining is that not only the price, but other product and service related issues can be taken into consideration as well. This can reduce the competitiveness of negotiations if agents have different preferences regarding the relative importance of the issues. The multi-issue aspect has therefore been given much consideration throughout this thesis, especially in Chapters 3, 4 and 6. The main objective of such negotiations is to obtain Pareto-efficiency by finding optimal trade-offs between issues. Chapter 3 showed that the evolutionary agents agree on Pareto-efficient outcomes after a relatively short learning period. Chapter 4 introduced a parameter to tune the competitiveness of two-issue negotiations. Using the evolutionary simulation, the impact of competitiveness in the game with multiple opportunities was investigated. Chapter 6 introduced advanced strategies for autonomous agents that are capable of approximating Pareto-efficiency within a single alternating-offers bargaining game. These strategies require no learning or knowledge of the opponent's preferences. In the example of information goods, we demonstrated that these strategies work well even for non-linear preferences.

In Chapters 6 and 7 bilateral bargaining was applied to the case of virtually unlimited supply as with information goods. Using an alternating-offers protocol as in Chapter 3, a seller negotiates with many buyers simultaneously and aims at reaching as many agreements as possible, and at the same time obtain a large share of the surplus. Chapter 6 focused on the multi-issue problem within the domain of information goods, whereas Chapter 7 considered the profits gained by the agents. For the latter we introduced a number of seller strategies that also take into account a fairness constraint; a negotiation between a buyer and a seller should be fair relative to other concurrent negotiations. One of the introduced strategies is able to extract almost the entire surplus, provided that buyers are sufficiently impatient and prefer to reach agreements early in the negotiation process. This strategy is inspired by the first-price auction, and simply accepts the highest offer received in each period. As before, we carried out evolutionary experiments to compare the performance of the seller strategies. The results showed that, in the absence of time pressure, there is no real competition and prices using the auction-based strategy dropped to cost level. We prevented such low prices by incorporating either a fixed or time-dependent reservation price into the seller's strategy. This combined strategy indeed showed superior performance in the evolutionary experiments.

To conclude, we considered both bilateral bargaining and auction approaches in this thesis, and successfully applied these approaches to several practical settings. An evolutionary framework was developed to investigate various bargaining settings and business applications. The outcomes of the computational experiments resulted in insights that go beyond current game-theoretic findings and demonstrated the effectiveness of automated bargaining for the application to real-world domains.

Appendix

Game-theoretic analysis

In this Appendix we derive subgame perfect equilibrium (see Def. 18.3) strategies for multiple-stage games from Chapter 3 using a backward induction approach. We follow the same approach as in [126], but extend the analysis to multi-issue negotiations and the extended model where the agents perform an additional fairness check. Appendix 1 studies the multi-issue models and is related to Section 3.3.2. In Appendix 2 we analyse the extended model with fairness and is related to Section 3.4.4.

1 Multi-issue bargaining

In Appendix 1.1, we study a model for multi-issue bargaining without a risk of breakdown. The more general model (with a risk of breakdown) is then investigated in Appendix 1.2. Furthermore, Appendix 1.3 presents equations for calculating Pareto-efficient (see Def. 4.3) utility pairs for additive multi-attribute utility functions (see Def. 3.1), given any number of issues and weight settings of the agents.

1.1 Model without a risk of breakdown ($p = 1$)

Because time plays no role in the model without a risk of breakdown, the last agent in turn has the opportunity to reject all proposals from his opponent and demand the entire surplus (for each issue) in the last round. In subgame perfect equilibrium (SPE), the other agent accepts this proposal (see also the discussion in [11, pp. 200-201]). If the maximum number of rounds n is odd, agent 1 will therefore receive the entire surplus, whereas agent 2 receives all in case n is even. Due to the absence of time pressure, multiple subgame perfect equilibria exist in this case. Although these equilibria differ in the timing of the agreements, all result in the same outcome (i.e., the agent in turn at $t = n - 1$ always receives the entire surplus for all issues). It is for instance subgame perfect for the last responder to concede the entire surplus (for all issues) to his opponent before the deadline is actually reached or, alternatively, to accept a take-it-or-leave-it deal from the opponent at any point in time.

1.2 Model with a risk of breakdown ($p < 1$)

We calculate the SPE for the model with a risk of breakdown in this Section. We first show that in Nash equilibrium (see Def. 18.2) the deals are always *weakly* Pareto-efficient (and therefore also in subgame perfect equilibrium). A deal is called *weakly* Pareto-efficient if there exists no other deal that *both* agents prefer. We assume in the following that the agents' decision to accept or reject an offer is determined by a threshold value τ: the offer is accepted if the utility level is above the threshold, and rejected otherwise. Consider a proposing agent i making an offer $\vec{o}_i(t)$ to his opponent, agent j, in round t of the negotiations ($t < n$). Assume that agent i knows that agent j's threshold is equal to $\tau_j(t)$. It is then a best response for agent i to propose a weakly Pareto-efficient deal to agent j.

We show this by contradiction. Suppose agent i proposes an offer $\vec{o}_i(t)$ to agent j which is *not* Pareto-efficient and agent j accepts this offer. Since the offer is Pareto-inefficient, there exists an offer $\vec{o}'_i(t)$ which results in a higher utility for agent i and the same utility or higher for agent j. Since agent j is either indifferent between $\vec{o}_i(t)$ and $\vec{o}'_i(t)$ or prefers $\vec{o}'_i(t)$, agent i would do better by offering $\vec{o}'_i(t)$ instead (which agent j will also accept).

The SPE partitioning can now be calculated as follows. If the maximum number of rounds n is even, agent 2 will be the proposer in the last round (i.e., at $t = n-1$). Agent 2 will then demand the whole surplus for each issue and agent 1 will receive nothing. This division of the surplus would yield agent 2 a payoff (expected utility) of $\pi_2(t = n-1) = p^{n-1}$, where $\pi_i(t)$ denotes agent i's payoff in the bargaining game starting at time t. We now analyse the previous round ($t = n-2$). Suppose agent 1's offer to agent 2 is $\vec{o}_1(t = n-2)$. Agent 2's payoff $\pi_2(t = n-2)$ would then be $p^{n-2}u_2[\vec{o}_1(t = n-2)]$. In equilibrium, at $t = n-2$ agent 1 should propose agent 2 a payoff-equivalent deal [i.e., $\pi_2(t = n-2) = \pi_2(t = n-1)$], . This implies that $u_2[\vec{o}_1(t = n-2)]$ should be equal to p. Agent 1's payoff $\pi_1(t = n-2)$ is then $p^{n-2}f_1(p)$, where $f_1(u_2)$ describes the location of the Pareto-efficient frontier. This function returns the utility of agent 1 when agent 2's utility is equal to u_2 and the agreement is Pareto-efficient.[1] At $t = n-3$, agent 2 can, in a similar fashion, propose an equivalent offer (in terms of payoff) and receive a payoff of $\pi_2(t = n-3) = p^{n-3}f_2[pf_1(p)]$. (The $f_2(u_1)$ function is the inverse of the f_1 function.)

This procedure is then repeated until the beginning of the game is reached (at $t = 0$). The same line of reasoning holds if the number of rounds is odd (simply switch the roles of agent 1 and agent 2). As in the infinite-horizon game [110],

[1] For the bargaining problem studied in this paper (depicted in Fig. 3.3), the Pareto-efficient frontier is described by the equation .3 in Appendix 1.3.

the agents agree immediately on a deal. Table .1 shows the SPE partitionings for different game lengths.

n	Payoff agent 1 $[\pi_1(t=0)]$ (SPE)	Payoff agent 2 $[\pi_2(t=0)]$ (SPE)
1	1	0
2	$f_1(p)$	p
3	$f_1(pf_2(p))$	$pf_2(p)$
4	$f_1(pf_2(pf_1(p)))$	$pf_2(pf_1(p))$
5	$f_1(pf_2(pf_1(pf_2(p))))$	$pf_2(pf_1(pf_2(p))))$
6	$f_1(pf_2(pf_1(pf_2(pf_1(p)))))$	$pf_2(pf_1(pf_2(pf_1(p)))))$
...	...	...

Table .1: Payoffs for agent 1 and agent 2 for different lengths n of the alternating-offers game, assuming that both agents use SPE strategies.

1.3 Calculating the Pareto-efficient frontier

We now show how Pareto-efficient values can be calculated if the agents use an additive multi-attribute utility function. The functions f_1 and f_2 are used to determine the Pareto-efficient utility of agent 1 and 2 respectively, given the utility received by the opponent. We will first give a recursive formula for f_1 and f_2 which can be used for any number of issues, and then present an example for two issues. As was described in section 3.1, each issue i is associated with a weight w_j^i for agent j. We assume that, without loss of generality, the issues are sorted such that

$$\forall i \in \{1, 2, \ldots, m-1\} : \frac{w_1^i}{w_2^i} \geq \frac{w_1^{i+1}}{w_2^{i+1}}.$$

We begin by deriving f_2, the maximum utility which agent 2 can obtain, given that agent 1 receives a utility u. Starting with demanding the full dollar on each issue, agent 2 needs to concede on zero or more issues such that a utility level u for agent 1 is reached. Agent 2 will first concede on issues with a relatively low loss in utility for agent 2 (i.e. with a low weight for agent 2) and a relatively high gain for agent 1. The issues are now sorted in such a way that agent 2 will first concede on issue 1, then on issue 2, etc., until the desired utility level for agent 1 is reached. This is reflected in the following formula, given agent 1's utility u: $f_2(u) = 1 - r_2^1(u)$ where r_2 is a recursively defined function:

$$r_2^i(u) = \begin{cases} r_2^{i+1}(u - w_1^i) + w_2^i & \text{if } u > w_1^i, \\ \frac{w_2^i}{w_1^i} u & \text{otherwise.} \end{cases} \quad (.1)$$

Similarly, agent 1 will start conceding on the *last* issue. The function for agent 1 is defined as $f_1 = 1 - r_1^m(u)$ where m is the number of issues, and r_1 is a recursively defined function:

$$r_1^i(u,i) = \begin{cases} r_1^{i-1}(u - w_2^i) + w_1^i & \text{if } u > w_2^i, \\ \frac{w_1^i}{w_2^i} u & \text{otherwise.} \end{cases} \tag{.2}$$

For any number of issues $m \geq 1$ and any weights, given that $\sum_{i=1}^m w_1^i = \sum_{i=1}^m w_2^i = 1$ and $w_j^i > 0$ for all $i \in \{1, \ldots, m\}, j \in \{1, 2\}$, the following properties hold: $f_1(1) = f_2(1) = 0$, $f_1(0) = f_2(0) = 1$ and $f_1(f_2(u)) = f_2(f_1(u)) = u$ for all $u : 0 \leq u \leq 1$. The next equation is an example of a *two*-issue bargaining situation with weight vectors $\vec{w}_1 = (0.7, 0.3)^T$ and $\vec{w}_2 = (0.3, 0.7)^T$ for agents 1 and 2 respectively:

$$f_1(u) = f_2(u) = \begin{cases} \frac{0.7}{0.3}(1 - u) & \text{if } u > 0.7. \\ 1 - \frac{0.3}{0.7} u & \text{otherwise,} \end{cases} \tag{.3}$$

This function yields the Pareto-efficient frontier depicted in figure 3.3. Note that when the weights are diametrically opposed, the same function applies to both agents, i.e. $f_1(u) = f_2(u)$ for all $u : 0 \leq u \leq 1$.

2 Extended model: Fairness

The fairness models evaluated in Section 3.4.2 (i.e., with a fairness check at the deadline only) and in Section 3.4.3 (i.e., with a fairness check in each round) are analysed in this appendix. As in Appendix 1.2, we apply backward induction to deduce the SPE partitioning.

2.1 General analysis

The fairness function is now formally denoted as $g_r(u)$. This (real-valued) function returns the probability of acceptance of a proposal in round r in case the responding agent's utility is equal to u. If a fairness check is performed only in the last round, $g_r(u) = 1$ for all $r < n$ (where n is the number of rounds). In case the same fairness check is performed each round, $g_r(u)$ is independent of r. We assume that the fairness function is a monotonic non-decreasing function of u and that $g_r(u = 1) = 1$. Let agent j be the agent proposing a deal at round r and agent $-j$ the responder. We then abbreviate $g_r[u_{-j}(\vec{o}_j(r))]$ (the probability of acceptance of agent j's offer $\vec{o}$ in round r) as $p_r^{acc}(\vec{o})$.

If n is even, agent 2 will propose an offer in the last round (at $r = n$). Agent 2 will then propose an offer $\vec{o}_2(r = n)$ which, in SPE, maximises his payoff, i.e., his expected utility. The payoff π_2 received by Agent 2 if his offer is accepted equals $p^n u_2[\vec{o}_2(r = n)]$, where u_2 is agent 2's utility function (see Section 3.1). The acceptance probability is equal to $p_n^{acc}[\vec{o}_2(r=n)]$. Agent 2's payoff in round $r=n$ is therefore:

$$\pi_2(r=n) = \max_{\vec{o}_2(r=n)\in\mathcal{P}} p^n u_2[\vec{o}_2(r=n)] p_n^{acc}[\vec{o}_2(r=n)], \qquad (.4)$$

where $\mathcal{P} \subset [0,1]^m$ is the set containing all Pareto-efficient offers. Analogously, the payoff π_1 for agent 1 in round $r=n$ is equal to:

$$\pi_1(r=n) = p^n u_1[\vec{o}_2(r=n)] p_n^{acc}[\vec{o}_2(r=n)], \qquad (.5)$$

where u_1 is agent 1's utility function.

It is again straightforward to show that it is optimal to propose a Pareto-efficient deal. Assume for instance, that a Pareto-*in*efficient offer is made. The proposer of this offer can then improve his payoff by selecting an offer on the Pareto-frontier which yields his opponent the same payoff. Because the probability of acceptance only depends on the responder's utility of this offer, this will not affect the fairness evaluation.

We now analyse the previous round ($r=n-1$). In SPE, at $r=n-1$ agent 2 only accepts a deal which is at least equal to the payoff $\pi_2(r=n)$ that he receives in the next round (in SPE). Therefore, $\pi_2(r=n-1) \geq \pi_2(r=n)$ in SPE. Effectively, $\pi_2(r=n)$ acts as a threshold used by agent 2 to determine the minimal acceptable offer at $r=n-1$. Some elementary manipulations then show that in SPE agent 1 should make an offer $\vec{o}_1(r=n-1)$ such that

$$p^{n-1} u_2[\vec{o}_1(r=n-1)] \geq \pi_2(r=n), \qquad (.6)$$

otherwise, agent 2 rejects the proposal at $r=n-1$ to earn $\pi_2(r=n)$ in the last round. We now define $\mathcal{R} \subset [0,1]^m$ to be the set of offers for which Eq. .6 is not violated. In SPE, agent 1's payoff in round $r=n-1$ then equals

$$\begin{aligned}\pi_1(r=n-1) &= \max_{\vec{o}_1(r=n-1)\in\mathcal{P}\cap\mathcal{R}} p^{n-1} u_1[\vec{o}_1(r=n-1)] p_{n-1}^{acc}[\vec{o}_1(r=n-1)] \\ &\quad + \{1 - p_{n-1}^{acc}[\vec{o}_1(r=n-1)]\}\pi_1(r=n). \qquad (.7)\end{aligned}$$

In a similar fashion, we can calculate agent 2's payoff at $r=n-1$ in SPE:

$$\begin{aligned}\pi_2(r=n-1) &= p^{n-1} u_2[\vec{o}_1(r=n-1)] p_{n-1}^{acc}[\vec{o}_1(r=n-1)] \\ &\quad + \{1 - p_{n-1}^{acc}[\vec{o}_1(r=n-1)]\}\pi_2(r=n). \qquad (.8)\end{aligned}$$

For $r=n-2$ expressions very similar to Eqs. .7 and .8 can be derived (but the roles of the two agents switch). This procedure is then repeated until the beginning of the game is reached (at $r=1$). The same line of reasoning holds if the number of rounds is odd (simply switch the roles of agent 1 and agent 2).

In the basic model without fairness all agreements occur in the first round in SPE (for $p < 1$). When the agents apply a fairness check in each round, however, even in SPE a significant number of agreements occurs after the first round. In this case, the strategy followed in all rounds comes to play a role in determining the outcome of the game.

We also remark that, although a responder's fairness considerations determines for a large part the offers made by a proposing agent, this does not make the responder's thresholds superfluous in SPE. Recall that the role of the threshold is reflected in Eq. .6.

2.2 Application to a simple case

We will now apply the general approach presented above to a relatively simple example with $m = 1$ (a single issue), $n = 3$ (3 rounds) and $p = 1$ (no time pressure). Because $m = 1$, the offer vector $\vec{o}(t)$ has only a single component. We denote he value of this component as $x(t)$ in the remainder of this appendix. It is obvious (because the agents are assumed to be risk neutral, see Section 3.1) that $u_1[x(t)] = x(t)$, and $u_2[x(t)] = 1-x(t)$ for $0 \leq t \leq n-1$. In this example, the agents evaluate the fairness of the offers (in each round) using fairness function 4 in Figure 3.7 [i.e., $g_t(u) = u$]. Furthermore, we take $n = 3$ and $p = 1$. Notice that, because the number of rounds n is odd in this example, we need to switch the roles of agent 1 and agent 2 when we apply Eqs. .4-.8 in the following.

Agent 1 makes an offer to agent 2 in the final round (at $t = 2$). In SPE, agent 1 applies Eq. .4 to maximise his payoff $\pi_1(t = 2)$. Substituting parameters for this problem, the product term on the RHS of Eq. .4 becomes $u_1[x(t = 2)]g_2[u_2(x(t = 2)]$, which can be simplified further to $x(t = 2)[1 - x(t = 2)]$. This term is maximised for $x(t = 2) = 0.5$, which results in $\pi_1(t = 2) = 0.25$. Using Eq. .5, the payoff of agent 2, $\pi_2(t = 2)$, is then also equal to $[1 - x(t = 2)]x(t = 2) = 0.25$.

Agent 2 makes a move at $t = 1$. We initially assume that the condition stated in Eq. .6 is not violated by agent 2's offer. Agent 2's payoff is then determined by applying Eq. .7. Substituting the parameters of this problem and simplifying, the term that should be maximised in Eq. .7 becomes equal to $[1 - x(t)]x(t) + [1 - x(t)]0.25$. This term is maximised for $x(t = 1) = 0.375$. The condition stated in Eq. .6 is not violated because $u_1(0.375) = 0.375 \geq 0.25$. Our initial assumption therefore turns out to be valid. We can now apply Eqs. .7 and .8 to derive that $\pi_1(t = 1) \approx 0.297$ and $\pi_2(t = 1) \approx 0.391$.

Agent 1 proposes an offer in the first round (at $t = 0$). Again, we initially ignore Eq. .6. Using Eq. .7, agent 1 then maximises his payoff $\pi_1(t = 0)$. This results in $x(t = 0) \approx 0.648$. However, this offer violates the condition in Eq. .6, since $u_2(0.648) = 0.352 < \pi_2(t = 1)$. Agent 1 should therefore propose a payoff-equivalent deal to agent 2 [i.e., $\pi_2(t = 0) = \pi_2(t = 1)$]. For $x \approx 0.609$ this condition is satisfied and agent 2 becomes indifferent between accepting or refusing this deal. Subgame perfection then predicts that agent 2 accepts this proposal, yielding agent 1 a payoff of ≈ 0.419. Tables 3.2 and 3.3 in Chapter 3 summarise these theoretical results. Notice that, in this example, Eq. .6 (i.e., the responder's threshold) indeed plays a role in round 1, whereas in the rounds 2 and 3 the equation does not influence the proposals made in SPE.

Bibliography

[1] F. Alkemade, J.A. La Poutré, and H.M. Amman. Intermediaries in an electronic trade network. In *Proceedings of the fourth ACM Conference on Electronic Commerce (EC'03), June 9–12, San Diego, USA*. The ACM Press, 2003.

[2] R. Amit and C. Zott. Value creation in ebusiness. *Strategic Management Journal*, 22:493–520, 2001.

[3] J. Andreoni and J.H. Miller. Auctions with artificial adaptive agents. *Games and Eonomic Behavior*, 10:39–64, 1995.

[4] J. Arifovic. Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, 18:3–28, 1994.

[5] L.M. Ausubel, P. Cramton, and R.J. Deneckere. Bargaining with incomplete information. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory, with Economic Applications*, volume 3. Elsevier Science Publisher, North-Holland, 2002.

[6] M. Bac and H. Raff. Issue-by-issue negotiations: The role of information and time preference. *Games and Economic Behavior*, 13(1):125–134, 1996.

[7] T. Bäck, U. Hammel, and H.-P. Schwefel. Evolutionary computation: Comments on the history and current state. *IEEE Trans. on Evol. Comp.*, 1(1):3–17, 1997.

[8] Th. Bäck. *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York and Oxford, 1996.

[9] Y. Bakos and E. Brynjolfsson. Bundling and competition on the internet: Aggregation strategies for information goods. *Marketing Science*, 19(1):63–82, 2000.

[10] J. Baron. *Thinking and Deciding*. Cambridge University Press, second edition, 1994.

[11] K. Binmore. *Fun and Games*. D.C. Heath and Company, Lexington, MA, 1992.

[12] K. Binmore, A. Rubinstein, and A. Wolinsky. The Nash bargaining solution in economic modelling. *Rand Journal of Economics*, 17:176–188, 1986.

[13] K. Binmore, A. Shaked, and J. Sutton. Testing noncooperative bargaining theory: A preliminary study. *American Economic Review*, 75:1178–1180, 1985.

[14] K. Binmore and N. Vulkan. Applying game theory to automated negotiation. *Netnomics*, 1(1):1–9, 1999.

[15] R. Blattberg and J. Deighton. Interactive marketing: Exploiting the age of addressability. *Sloan Management Review*, 33:5–14, 1991.

[16] S.M. Bohte, E.H. Gerding, and J.A. La Poutré. Competitive market-based allocation of consumer attention space. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 202–206. The ACM Press, 2001.

[17] S.M. Bohte, E.H. Gerding, and J.A. La Poutré. Market-based recommendation: Agents that compete for consumer attention. *ACM Transactions on Internet Technology*, August 2004. To appear.

[18] C.H. Brooks, S. Fay, R. Das, J.K. MacKie-Mason, J.O. Kephart, and E.H. Durfee. Automated strategy searches in an electronic goods market: Learning complex price schedules. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 31–41. the Association for Computing Machinery (ACM), ACM Press, 1999.

[19] L.-A. Busch and I.J. Horstmann. Signaling via an agenda in multi-issue bargaining with incomplete information. *Economic Theory*, 13(3):561–575, 1999.

[20] J.Q. Cheng and M.P. Wellman. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics*, 12:1–24, 1998.

[21] A.M. Chircu and R.J. Kauffman. Strategies for internet middlemen in the intermediation/disintermediation/reintermediation cycle. *Electronic Markets*, 9(2):109–117, 1999.

[22] S.Y. Choi, D.O. Stahl, and A.B. Whinston. *The Economics of Electronic Commerce*. Macmillan Technical Publishing, 1997.

[23] S.H. Clearwater, editor. *Market based Control of Distributed Systems*. World Scientific Press, Singapore., 1995.

[24] D. Cliff. Evolution of market mechanism through a continuous space of auction-types. Technical Report HPL-2001-326, HP Labs, 2001.

[25] M. Costa-Gomes and K.G. Zauner. Ultimatum bargaining behavior in Israel, Japan, Slovenia, and the United States: A social utility analysis. *Games and Economic Behavior*, 34(2):238–269, 2001.

[26] P. Dasgupta and S. Das. Dynamic pricing with limited competitor information in a multi-agent economy. In Opher Eztion and Peter Scheuermann, editors, *Cooperative Information Systems: 7th International Conference*, volume 1906 of *Lecture Notes in Computer Science*, pages 291–310, Eilat, Israel, September 2000. BoopIs, Springer-Verlag.

[27] P. Dasgupta and E. Maskin. Efficient auctions. *The Quarterly Journal of Economics*, 115:341–388, 2000.

[28] E. David, R. Azoulay-Schwartz, and S. Kraus. Protocols and strategies for automated multi-attribute auctions. In *Proc. 1st Int. Joint Conference on Autonomous Agents and MultiAgent Systems*, 2002.

[29] H. Dawid. *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*, volume 441 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, 2nd edition, 1999.

[30] F. Dignum and U. Cortés, editors. *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Eletronic Commerce Systems*. Springer-Verlag, 2001.

[31] G. Dworman, S.O. Kimbrough, and J.D. Laing. Bargaining by artificial agents in two coalition games: A study in genetic programming for electronic commerce. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 54–63. The MIT Press, 1996.

[32] H. Ehtamo and R.P. Hämäläinen. Interactive multiple-criteria methods for reaching pareto optimal agreements in negotiation. *Group Decision and Negotiation*, 10:475–491, 2001.

[33] P. Faratin, C. Sierra, and N.R. Jennings. Using similarity criteria to make issue trade-offs. *Artificial Intelligence*, 142(2):205–237, 2002.

[34] S. Fatima, M. Wooldridge, and N.R. Jennings. Comparing equilibria for game theoretic and evolutionary bargaining models. In *Proceedings of the 5th International Workshop on Agent-Mediated E-Commerce (AMEC V), Melbourne, Australia*, pages 70–77, 2003.

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

[36] R. Fisher. *International Mediation: A Working Guide.* International Peace Academy, New York, 1978.

[37] R. Gellmann. Disintermediation and the internet. *Government Information Quarterly*, 13(1):1–8, 1996.

[38] E.H. Gerding and J.A. La Poutré. Bargaining with posterior opportunities: An evolutionary social simulation. In M. Gallegati, A. Kirman, and M. Marsili, editors, *The Complex Dynamics of Economic Interaction*, Springer Lecture Notes in Economics and Mathematical Systems (LNEMS), Vol. 531, pages 241–256. Springer-Verlag, 2004.

[39] E.H. Gerding, K. Somefun, and J.A. La Poutré. Bilateral bargaining in a one-to-many bargaining setting. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS2004), New York City, New York.* IEEE Computer Society Press, 2004. To appear.

[40] E.H. Gerding, K. Somefun, and J.A. La Poutré. Bilateral bargaining in a one-to-many bargaining setting. Technical report, CWI, Amsterdam, to appear.

[41] E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Scientific approaches and techniques for negotiation: A game theoretic and artificial intelligence perspective. Technical Report SEN-R0005, CWI, Amsterdam, 2000.

[42] E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. *Computational Economics*, 22:39–63, 2003.

[43] M.A. Gibney, N.R. Jennings, N.J. Vriend, and J.M. Griffiths. Market-based call routing in telecommunications networks using adaptive pricing and real bidding. In *Proceedings of the 3rd International Workshop on Intelligent Agents for Telecommunication Applications (IATA-99), Stockholm, Sweden, August 9-10 1999*, pages 50–65, 1999.

[44] A.V. Goldberg and J.D. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM Conference on E-Commerce (EC'03)*, pages 29–35. ACM press, 2003.

[45] D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, Reading, 1989.

[46] A.R. Greenwald and J.O. Kephart. Shopbots and pricebots. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, volume 1, pages 506–511, 1999.

[47] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review Journal*, pages 147–159, 1998.

[48] R.H. Guttman and P. Maes. Cooperative vs. competitive multi-agent negotiations in retail electronic commerce. In *Proceedings of the 2nd International Workshop on Cooperative Information Agents (CIA'98)*, 1998.

[49] S.H. Haeckel. About the nature and future of interactive marketing. *Journal of Interactive Marketing*, 12:63–71, 1998.

[50] G.W. Harrison and K.A. McCabe. Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory*, 25:303–327, 1996.

[51] J.H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence.* The University of Michigan Press/Longman Canada, Ann Arbor, 1975.

[52] http://www.google.com. Google search engine.

[53] http://www.overture.com. Overture.

[54] http://www.worldmall.com. World Mall, online shopping portal.

[55] P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259, 2001.

[56] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10:199–215, 2001.

[57] E. Kalai and M. Smorodinsky. Other solutions to the Nash bargaining problem. *Econometrica*, 43:513–518, 1975.

[58] J. Kennan and R. Wilson. Bargaining with private information. *Journal of Economic Literature*, 31:45–104, 1993.

[59] J.O. Kephart, C.H. Brooks, and R. Das. Pricing information bundles in a dynamic environment. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 180–190. ACM Press, 2001.

[60] J.O. Kephart, J.E. Hanson, and A.R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.

[61] J.O. Kephart, J.E. Hanson, and A.R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 36(6):731–752, May 2000.

[62] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Negotiating complex contracts. *Group Decision and Negotiation*, 12:111–125, 2003.

[63] A. Kobsa. Tailoring privacy to users' needs. In *Proceedings of the 8th International Conference on User Modeling, Sonthofen, Germany.* Springer Verlag, 2001.

[64] S. Kraus. Beliefs, time and incomplete information in multiple encounter negotiations among autonomous agents. *Annals of Mathematics and Artificial Intelligence*, 20:111–159, 1997.

[65] S. Kraus. *Strategic Negotiation in Multiagent Environments.* Intelligent Robots and Autonomous Agents. The MIT Press, Cambridge, 2001.

[66] B. Krulwich. The BargainFinder agent: Comparison price shopping on the internet. In J. Williams, editor, *Agents, Bots and Other Internet Beasties.* SAMS.NET, 1996.

[67] H. Lin and S. Sunder. Using experimental data to model bargaining behavior in ultimatum games. In R. Zwick and A. Rapoport, editors, *Experimental Business Research*, pages 373–397. Kluwer Academic Publishers, 2002.

[68] R.J. Lin and S.T. Chou. Bilateral multi-issue negotiations in a dynamic environment. In *Proceedings of the 5th Workshop on Agent-Mediated Electronic Commerce V (AMEC V)*, pages 117–123, 2003.

[69] R.D. Luce and H. Raiffa. *Games and Decisions.* Wiley, New York, 1957.

[70] M. Luck, P. McBurney, and C. Preist. *Agent Technology: Enabling Next Generation Computing (A Roadmap for Agent Based Computing).* University of Southampton, 2003.

[71] P. Maes, R.H. Guttman, and A.G. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.

[72] A. Mas-Collel, M.D. Whinston, and J.R. Green. *Mircoeconomic Theory.* Oxford University Press, 1995.

[73] N. Matos, C. Sierra, and N.R. Jennings. Determining successful negotiation strategies: An evolutionary approach. In *Proceedings of the 3rd International Conference on Multi-Agent Systems (ICMAS-98)*, pages 182–189, Paris, France, 1998.

[74] P.J. McBurney. *Rational Interaction.* PhD thesis, University of Liverpool, 2002.

[75] M. Mitchell. *An Introduction to Genetic Algorithms.* The MIT Press, Cambridge, MA, 1996.

[76] A.L. Montgomery. Applying quantitative marketing techniques to the internet. *Interfaces*, 32:90–108, 2001.

[77] A.L. Montgomery and K. Srinivasan. Learning about customers without asking. January 2002.

[78] A. Moukas, R. Guttman, and P. Maes. Agent-mediated electronic commerce: An MIT media laboratory perspective. In *Proceedings of the International Conference on Electronic Commerce, Seoul, Korea*, 1998.

[79] A. Muthoo. *Bargaining and Applications.* Cambridge University Press, Cambridge, MA, 1999.

[80] R.B. Myerson. Utilitarianism, egalitarianism, and the timing effect in social choice problems. *Econometrica*, 49(4):883–897, 1981.

[81] S. Napel. *Bilateral Bargaining: Theory and Applications.* Springer-Verlag, Berlin, 2002.

[82] J.F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

[83] J.F. Nash. Equilibrium points in N-person games. *Proc. Nat. Acad. Sci. U.S.A.*, 36:48–49, 1950.

[84] J.F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[85] T.D. Nguyen and N.R. Jennings. Concurrent bi-lateral negotiation in agent systems. In *Proceedings of the 4th DEXA Workshop on e-Negotiations, Prague, Czech Republic*, 2003.

[86] M. Nissen. Agent-based supply chain disintermediation versus re-intermeditation: Economic and technological perspectives. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 9:237–256, 2000.

[87] E. Oliveira and A.P. Rocha. Agents advanced features for negotiation in electronic commerce and virtual organisations formation process. In C. Sierra and F. Dignum, editors, *Agent Mediated Electronic Commerce: the European AgentLink Perspective*, volume 1991 of *Lecture Notes in Artificial Intelligence*, pages 78–97. Springer-Verlag, Berlin, 2001.

[88] J.R. Oliver. A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of Management Information Systems*, 13(3):83–112, 1996.

[89] M.J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, San Diego, CA, 1990.

[90] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, MA, 1994.

[91] M.J. Owen, L. Lee, S. Steward, and D. Thomas. Multi-agent trading environment. *BT Technology Journal*, 17:33–43, 1999.

[92] J. Padget, O. Shehory, D. Parkes, N. Sadeh, and W.E. Walsh, editors. *Agent-Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.

[93] D.C. Parkes and J. Kalagnanam. Iterative multiattribute Vickrey auctions. Working paper, 2002.

[94] S. Parsons, C. Sierra, and N.R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[95] D. Peppers, M. Rogers, and B. Dorf. Is your company ready for one-to-one marketing? *Harvard Business Review*, 77:151–162, 1999.

[96] S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Co-evolutionary auction mechanism design: A preliminary report. In J. Padget, O. Shehory, D. Parkes, N. Sadeh, and W.E. Walsh, editors, *Agent-Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Computer Science*, pages 123–142. Springer-Verlag, Berlin, 2002.

[97] C. Ponsati and J. Watson. Multiple-issue bargaining and axiomatic solutions. *International Journal of Game Theory*, 26:501–524, 1997.

[98] I. Rahwan, R. Kowalczyk, and H.H. Pham. Intelligent agents for automated one-to-many e-commerce negotiation. In *Proceedings of the 25th Australasian Computer Science Conference (ACSC), Melbourne, Australia*, 2002.

[99] I. Rahwan, S. D. Ramchurn, N.R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, to appear.

[100] I. Rahwan, L. Sonenberg, and F. Dignum. Towards interest-based negotiation. In *Proceedings of the Second International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Melbourne, Australia*, pages 773–780. ACM Press, 2003.

[101] H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, MA, 1982.

[102] S.D. Ramchurn, N.R. Jennings, and C. Sierra. Persuasive negotiation for autonomous agents: A rhetorical approach. In *Proceedings of the IJCAI Workshop on Computational Models of Natural Argument, Acapulco, Mexico*, pages 9–17, 2003.

[103] I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution.* Frommann-Holzboog, Stuttgart, 1973.

[104] T. Riechmann. *Learning in Economics: Analysis and Application of Genetic Algorithms.* Physica-Verlag, Heidelberg, Germany, 2001.

[105] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter.* The MIT Press, Cambridge, MA, 1994.

[106] A.E. Roth. *Axiomatic Models of Bargaining.* Lecture Notes in Economics and Mathematical Systems, No. 170. Springer-Verlag, Berlin, 1979.

[107] A.E. Roth. Bargaining experiments. In J. Kagel and A.E. Roth, editors, *Handbook of Experimental Economics*, pages 253–348, Princeton University Press, Princeton, NJ, 1995.

[108] A.E. Roth, J.K. Murnighan, and F. Schoumaker. The deadline effect in bargaining: Some experimental evidence. *American Economic Review*, 78(4):806–823, 1988.

[109] A.E. Roth, V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir. Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *American Economic Review*, 81:1068–1095, 1991.

[110] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):155–162, 1982.

[111] A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5):1151–1172, 1985.

[112] A. Rubinstein. *Modeling Bounded Rationality.* The MIT Press, Cambridge, MA, 1998.

[113] S. Russell and P. Norvig. *Artificial Intelligence, A Modern Approach.* Prentice-Hall, 1995.

[114] J.B. Schafer, J. Konstan, and J. Riedi. Recommender systems in e-commerce. In *Proceedings of the first ACM Conference on Electronic Commerce (EC'99)*, pages 158–166. ACM Press, 1999.

[115] H.-P. Schwefel. *Evolution and Optimum Seeking.* Sixth-Generation Computer Technology Series. Wiley, New York, 1995.

[116] R. Selten. Spieltheoretische Behandlung eines Oligopolmodels mit Nachfrageträgheit. *Zeitschrift für die Gesamte Staatswissenschaft*, 12:301–324 and 667–689, 1965.

[117] R. Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.

[118] G. Shaffer and Z.J. Zhang. Competetitve one-to-one promotions. *Management Science*, 48(9):1143–1160, 2002.

[119] C. Sierra and F. Dignum, editors. *Agent Mediated Electronic Commerce: the European AgentLink Perspective*, volume 1991 of *Lecture Notes in Artificial Intelligence.* Springer-Verlag, Berlin, 2001.

[120] K. Somefun, E.H. Gerding, S. Bohte, and J.A. La Poutré. Automated negotiation and bundling of information goods. In *Agent-Mediated Electronic Commerce V*, Springer Lecture Notes in Artificial Intelligence (LNAI). Springer-Verlag, Berlin, to appear.

[121] I. Ståhl. *Bargaining Theory.* Stockholm School of Economics, Stockholm, 1972.

[122] P.J. 't Hoen, E.H. Gerding S.M. Bohte, and J.A. La Poutré. Implementation of a competitive market-based allocation of consumer attention space. In J. Padget, O. Shehory, D. Parkes, N. Sadeh, and W.E. Walsh, editors, *Agent-Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, volume 2531 of *Lecture Notes in Computer Science*, pages 273–288. Springer-Verlag, Berlin, 2002.

[123] L. Tesfatsion. Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control*, 25(3/4):281–293, 2001.

[124] L. Tesfatsion. Agent-based computational economics: Modelling economies as complex adaptive systems. In *Proceedings of the 6th Joint Conference on Information Sciences, March 8-13, North Carolina, USA*, pages 40–43, 2002.

[125] L. Tesfatsion. Agent-based computational economics. ISU Economics Working Paper No. 1, February 2003. Available at http://www.econ.iastate.edu/tesfatsi/ace.htm.

[126] D.D.B. van Bragt, E.H. Gerding, and J.A. La Poutré. Equilibrium selection in alternating-offers bargaining models: The evolutionary computing approach. *The Electronic Journal of Evolutionary Modeling and Economic Dynamics (e-JEMED)*, 1, 2002.

[127] D.D.B. van Bragt and J.A. La Poutré. Generating efficient automata for negotiations - an exploration with evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECKO-2001), San Francisco, July 2001*, 2001.

[128] D.D.B. van Bragt and J.A. La Poutré. Why agents for automated negotiations should be adaptive. *Netnomics*, 5:101–118, 2003.

[129] D.D.B. van Bragt, D.J.A. Somefun, E. Kutschinski, and J.A. La Poutré. An algorithm for on-line price discrimination. Technical Report SEN-R0213, CWI, Amsterdam, 2002.

[130] D.D.B. van Bragt, C.H.M. van Kemenade, and J.A. La Poutré. The influence of evolutionary selection schemes on the iterated prisoner's dilemma. *Computational Economics*, 17(2–3):253–263, 2001.

[131] E. van Damme. *Stability and Perfection of Nash Equilibria.* Springer-Verlag, Berlin, 1991.

[132] E. van Damme, R. Selten, and E. Winter. Alternating bid bargaining with a smallest money unit. *Games and Economic Behavior*, 2:188–201, 1990.

[133] H.R. Varian. Economic mechanism design for computerized agents. In *Proceedings of the First USENIX Workshop on Electronic Commerce*, pages 13–21, New York, 1995.

[134] H.R. Varian and J.K. MacKie-Mason. Generalized vickrey auctions. Technical report, University of Michigan, 1995.

[135] B.H.P.J. Vermeer. *Data Quality and Data Alignment in E-Business.* PhD thesis, Eindhoven University of Technology, 2001.

[136] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[137] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, New Jersey, 1944.

[138] N.J. Vriend. Self-organization of markets: An example of a computational approach. *Computational Economics*, 8:205–231, 1995.

[139] N.J. Vriend. An illustration of the essential difference between individual learning and social learning and its consequences for computational analysis. *Journal of Economic Dynamics and Control*, 24:1–19, 2000.

[140] Roger Webster. *Convexity*. Oxford Science Publications, 1994.

[141] E. Weg, A. Rapoport, and D.S. Felsenthal. Two-person bargaining behavior in fixed discounting factors games with infinite horizon. *Games and Economic Behavior*, 2:76–95, 1990.

[142] Y.Z. Wei, L. Moreau, and N.R. Jennings. Market-based recommenders: Design, simulation and evaluation. In *Proceedings of the 5th International Workshop on Agent-Oriented Information Systems, Melbourne, Australia*, pages 22–29, 2003.

[143] Y.Z. Wei, L. Moreau, and N.R. Jennings. Recommender systems: A market-based design. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS2003), Melbourne, Australia*, pages 600–607, 2003.

[144] G. Weiss, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. The MIT Press, Cambridge, MA, 1999.

[145] F. Ygge and H. Akkermans. Decentralized markets versus central control: A comparative study. *Journal of Artificial Intelligence Research*, 11:301–333, 1999.

[146] H.P. Young. *Individual Strategy and Social Structure*. Princeton University Press, 1998.

[147] D. Zeng and K. Sycara. How can an agent learn to negotiate. In *Intelligent Agents III: Agent Theories, Architectures, and Languages*, volume 1193 of *Lecture Notes in Computer Science (LNCS)*, pages 233–244. Springer, 1997.

[148] D. Zeng and K. Sycara. Bayesian learning in negotiation. *International Journal of Human Computer Systems*, 48:125–141, 1998.

[149] R. Zwick and C.C. Lee. Bargaining and search: An experimental study. *Group Decision and Negotiation*, 8(6):463–487, 1999.

Samenvatting (Dutch)

Onderhandelen speelt een steeds grotere rol door de ontwikkelingen binnen de elektronische handel, met name door de ontwikkeling van autonome software agenten. Dit zijn programma's die, geïnstrueerd door een gebruiker, in staat zijn zelfstandig en op een intelligente wijze een gegeven opdracht te verwezenlijken. Door middel van autonome software agenten kan het onderhandelingsproces worden geautomatiseerd, waarmee goederen en diensten met de daarbij horende voorwaarden, zoals garantie en leveringstijd, flexibel kunnen worden afgestemd op de individuele wensen van de betrokkenen. In dit proefschrift wordt aandacht besteed aan zowel fundamentele aspecten van onderhandelen als bedrijfstoepassingen van geautomatiseerd onderhandelen dmv software agenten.

Het fundamentele deel houdt zich bezig met de vraag wat de uitkomst van onderhandelende agenten zal zijn in een gestileerde wereld en hoe deze uitkomst wordt beïnvloed. Hierdoor kunnen inzichten worden verkregen voor het produceren van agenten, strategieën en het opstellen van onderhandelingsregels voor praktijksituaties. Wij bestuderen deze aspecten aan de hand van computer simulaties van onderhandelende agenten. Hierbij wordt gekeken naar adaptieve systemen, dwz waarbij agenten leren hun onderhandelingsstrategie aan te passen aan de hand van ervaringen uit het verleden. Het leergedrag wordt gesimuleerd door evolutionaire algoritmen. Deze algoritmen komen voort uit de kunstmatige intelligentie en zijn geïnspireerd door de evolutie theorie uit de biologie. Oorspronkelijk zijn de evolutionaire algoritmen ontwikkeld om optimalisatieproblemen op te lossen, maar binnen de economie wordt deze methode steeds vaker toegepast om leergedrag van mensen te modelleren. Naast computer simulaties bestuderen wij voor relatief eenvoudige gevallen wiskundige oplossingen uit de zogenaamde spel theorie. De spel theorie houdt zich met name bezig met de "rationele mens", dwz met optimale oplossing in een geabstraheerde situatie (of spel), gegeven dat iedereen zich rationeel gedraagt. De spel-theoretische uitkomsten worden gebruikt om de computer experimenten te valideren. Het voordeel van de computer simulaties is dat minder stricte assumpties nodig zijn en dat complexere interacties die dichter bij de werkelijkheid staan kunnen worden bestudeerd.

Allereerst wordt een onderhandelingssituatie onderzocht waarbij twee spelers boden en tegen boden tegen elkaar uitwisselen, het zogenaamde alternerende boden spel. Dit spel wordt vaak gebruikt als model voor onderhandelen over bijvoorbeeld de prijs van een product of dienst. Het is echter belangrijk om ook andere product of dienst gerelateerde aspecten in beschouwing te nemen zoals de kwaliteit, leveringstijd en garantieperiode. Dit geeft namelijk de mogelijkheid om compromissen te sluiten door toe te geven op minder belangrijke aspecten en meer te vragen voor belangrijke aspecten. Hierdoor zijn onderhandelingen minder competitief en kunnen uitkomsten ontstaan die voor beide partijen aantrekkelijk zijn. Derhalve onderzoeken wij middels computer simulaties een variant op het alternerende boden spel, waarbij meerdere aspecten gelijktijdig worden onderhandeld. Daarnaast gebruiken wij de speltheorie om resultaten van de simulatie te valideren. De simulatie laat zien dat lerende agenten in korte tijd tot optimale compromissen komen, ook wel Pareto efficiënte oplossingen genoemd. Vervolgens bestuderen wij het effect van tijdsdruk die ontstaat als onderhandelingen met een kleine kans worden afgebroken, bijvoorbeeld als gevolg van externe factoren. Bij het ontbreken van tijdsdruk en een maximum aantal rondes, zijn de uitkomsten zeer onevenwichtig: de speler die als laatste de kans krijgt om een bod uit te brengen doet een "alles of niets" bod in de laatste ronde, wat voor de andere speler slechts een fractie beter is dan helemaal geen overeenkomst. Bij een relatief hoge tijdsdruk, is juist het eerste bod het belangrijkste, en worden bijna alle overeenkomsten in de eerste ronde afgesloten. Een andere interessante uitkomst is dat de simulatie resultaten na een lange leerperiode in grote lijnen overeenkomen met oplossingen uit de speltheorie, ondanks dat de lerende agenten niet "rationeel" zijn. In de werkelijkheid is niet alleen de uitkomst belangrijk, maar spelen ook andere factoren mee, zoals de eerlijkheid van de uitkomst. Middels de simulatie wordt gekeken naar de invloed op de onderhandelingsuitkomsten als door de agent met dergelijke normen rekening wordt gehouden. Door deze normen zijn de uitkomsten veel evenwichtiger, ook bij het ontbreken van tijdsdruk, en lijken meer overeen te komen met de werkelijkheid.

Onderhandelingen staan vaak niet op zichzelf, maar worden bepaald door externe factoren zoals additionele onderhandelingsmogelijkheden. Naast het geïsoleerde onderhandelingsspel, bestuderen wij daarom ook onderhandelingen binnen een marktachtige omgeving, waarbij zowel kopers als verkopers meerdere keren kunnen onderhandelen met verschillende spelers om tot een overeenkomst te komen. Deze onderhandelingen worden opeenvolgend uitgevoerd totdat een overeenkomst is gesloten of totdat er geen onderhandelingsmogelijkheden meer zijn. Elk onderhandelingsspel tussen twee spelers wordt hier beperkt tot één ronde, waarbij speler 1 een bod doet, en speler 2 kan dit bod weigeren of accepteren. Met een evolutionaire simulatie onderzoeken wij verscheidene eigenschappen van het markt spel. Het blijkt dat de uitkomsten erg afhangen van de informatie die beschikbaar is binnen het spel. Als de spelers op de hoogte zijn van elkaars onderhandelingsmogelijkheden, dan doet de biedende speler telkens een alles of niets bod en krijgt het grootste voordeel. Dit komt overeen met speltheoretische uitkomsten die wij tevens presenteren in dit proefschrift. Als deze informatie niet bekend is, wordt een theoretische analyse heel moeilijk. De evolutionaire simulatie laat dan echter zien dat de tweede speler, dwz de speler die het bod weigert of accepteert, de beste onderhandelingspositie bezit. Dit komt omdat de eerste speler niet kan inschatten wat de reactie zal zijn van de tweede speler, en hierdoor lager inzet. In het proefschrift wordt verder ook gekeken naar andere factoren die de uitkomsten beïnvloeden, zoals het onderhandelen over meerdere aspecten tegelijkertijd, zoekkosten en afbreekkansen.

Naast de aandacht voor fundamentele vraagstukken worden in dit proefschrift een aantal bedrijfsgerelateerde toepassingen van geautomatiseerd onderhandelen gepresenteerd alsmede generieke onderhandelingsstrategieën voor de agenten die in gerelateerde applicaties kunnen worden ingezet. Als eerste toepassing introduceren wij een raamwerk waarbij onderhandelen wordt gebruikt voor het aanbevelen van winkels aan klanten, bijvoorbeeld op een webpagina van een elektronisch winkelcentrum. Middels de marktwerking van een veiling wordt op een gedistribueerde wijze een relevante selectie van winkels voor de klant bepaald. Hiertoe worden een beperkt aantal advertentieplaatsen in een elektronische veiling aangeboden. Voor elke individuele bezoeker van de pagina kunnen winkels via hun *winkel agent* geautomatiseerd bieden voor deze "aandachtsspanne" van de klant. Het bieden door deze software agent geschiedt op basis van een klantenprofiel, wat persoonlijke gegevens bevat van de bezoeker, zoals zijn/haar interessen, leeftijd, en/of opgegeven zoekwoorden. De winkel agenten zijn adaptief en leren, gegeven terugkoppeling van de klant, op welke profielen ze zich moeten richten en hoe hoog ze moeten bieden. De hoogste bieders worden vervolgens aan de klant getoond. De werking van het op deze wijze gedistribueerd bepalen van relevante winkels is aangetoond middels een evolutionaire simulatie. Wij onderzoeken verschillende modellen van klanten en veilingmechanismen, en laten zien dat het veilingsysteem resulteert in een passende selectie van winkels voor de klant.

Onderhandelen kan vooral van belang zijn als niet alleen de prijs, maar ook andere aspecten een rol spelen. Hierdoor kunnen bijvoorbeeld goederen en diensten beter worden afgestemd op individuele wensen van de gebruiker. Dit aspect wordt benut in een systeem wat wij hebben ontwikkeld voor de verkoop en personalisatie van zogenaamde informatiegoederen, zoals nieuws artikelen, software en muziek. Middels het alternerende boden protocol kan een verkopende software agent met meerdere kopende software agenten tegelijkertijd automatisch onderhandelen over een vaste prijs, een "stukprijs", en de kwaliteit van een bundel informatiegoederen. Het systeem houdt ook rekening met belangrijke bedrijfsgerelateerde voorwaarden zoals de eerlijkheid van de onderhandeling. De agenten gebruiken een combinatie van een concessiestrategie en een zoekstrategie om een onderhandelingsbod te genereren. De concessiestrategie bepaalt hoeveel elke ronde wordt toegegeven, terwijl de zoekstrategie zorgt voor gepersonaliseerde boden. In dit proefschrift introduceren wij een tweetal zoekstrategieen, en wij laten middels computer simulaties zien dat bij gezamenlijk gebruik door een kopende agent en een verkopende agent deze strategieën leiden tot gepersonaliseerde oplossingen, ook in combinatie met verschillende concessiestrategieën. Deze zoekstrategieën kunnen ook gemakkelijk worden toepast bij andere onderhandelingssituaties waarbij personalisatie een rol speelt.

Naast bovenstaande zoekstrategieën hebben wij ook een aantal concessiestrategieën ontwikkeld voor een verkopende agent die met meerdere kopende agenten tegelijkertijd onderhandelt. Ook al is het onderhandelingsproces *op zich* bilateraal (dwz tussen twee partijen), kan de verkopende agent gebruik maken van het feit dat meerdere onderhandelingen tegelijkertijd plaatsvinden. De ontwikkelde onderhandelingsstrategieën zijn gericht op situaties waarbij het aanbod flexibel is en kan worden afgestemd op de vraag, zoals bij informatie goederen. Wij bestuderen hierbij vaste strategieën, tijdsafhankelijke strategieën, en introduceren tevens een aantal strategieën die geïnspireerd zijn door veilingen. Veilingen worden vaak gebruikt in situaties waarbij één partij onderhandelt met meerdere partijen tegelijkertijd. Hoewel deze laatste strategie de voordelen heeft van een veiling, blijft de onderhandeling zelf bilateraal en bestaat uit het uitwisselen van boden en tegen boden. Een evolutionaire simulatie omgeving is ontwikkeld om de strategieën van de verkoper te evalueren. Hierbij wordt voornamelijk gekeken naar de situatie waarbij de kopers tijdsdruk ondervinden en onder druk staat om snel tot een overeenkomst te komen. Uit de simulatie blijkt dat de op veiling geïnspireerde strategieën van de verkopende agent in staat zijn bijna de maximale winst uit de onderhandelingen te halen bij voldoende tijdsdruk van de kopers.

Summary (English)

Bargaining is becoming increasingly important due to developments within the field of electronic commerce, especially the development of autonomous software agents. Software agents are programs which, given instructions from a user, are capable of autonomously and intelligently realise a given task. By means of such agents, the bargaining process can be automated, allowing products and services together with related conditions, such as warranty and delivery time, to be flexible and tuned to the individual preferences of the people concerned. In this theses we concentrate on both fundamental aspects of bargaining as well as business-related applications of automated bargaining using software agents.

The fundamental part investigates bargaining outcomes within a stylised world, and the factors that influence these outcomes. This can provide insights for the production of software agents, strategies, and setting up bargaining rules for practical situations. We study these aspects using computational simulations of bargaining agents. Hereby we consider adaptive systems, i.e., where agents learn to adjust their bargaining strategy given past experience. This learning behaviour is simulated using evolutionary algorithms. These algorithms originate from the field of artificial intelligence, and are inspired by the biological theory of evolution. Originally, evolutionary algorithms were designed for solving optimisation problems, but they are now increasingly being used within economics for modelling human learning behaviour. Besides computational simulations, we also consider mathematical solutions from game theory for relatively simple cases. Game theory is mainly concerned with the "rational man", that is, with optimal outcomes within an stylised setting (or game) where people act rationally. We use the game-theoretic outcomes to validate the computational experiments. The advantage of computer simulations is that less strict assumptions are necessary, and that more complex interactions that are closer to real-world settings can be investigated.

First of all, we study a bargaining setting where two players exchange offers and counter offers, the so-called alternating-offers game. This game is frequently used for modelling bargaining about for instance the price of a product or service. It is also important, however, to allow other product- and service-related aspects to be nego-

tiated, such as quality, delivery time, and warranty. This enables compromises by conceding on less important issues and demanding a higher value for relatively important aspects. This way, bargaining is less competitive and the resulting outcome can be mutually beneficial. Therefore, we investigate using computational simulations an extended version of the alternating-offers game, where multiple aspects are negotiated concurrently. Moreover, we apply game theory to validate the results of the computational experiments. The simulation shows that learning agents are capable of quickly finding optimal compromises, also called Pareto-efficient outcomes. In addition, we study the effects of time pressure that arise if negotiations are broken off with a small probability, for example due to external eventualities. In absence of time pressure and a maximum number of negotiation rounds, outcomes are very unbalanced: the player that has the opportunity to make a final offer proposes a take-it-or-leave-it offer in the last round, which leaves the other player with a deal that is only slightly better than no deal at all. With relatively high time pressure, on the other hand, the first offer is most important and almost all agreements are reached in the first round. Another interesting result is that the simulation outcomes after a long period of learning in general coincide with the results from game theory, in spite of the fact that the learning agents are not "rational". In reality, not only the final outcome is important, but also other factors play a role, such as the fairness of an offer. Using the simulation we study the influence of such fairness norms on the bargaining outcomes. The fairness norms result in much more balanced outcomes, even with no time pressure, and seem to be closer outcomes in the real world.

Negotiations are rarely isolated, but can also be influenced by external factors such as additional bargaining opportunities. We therefore also consider bargaining within a market-like setting, where both buyers and sellers can bargain with several opponents before reaching an agreement. The negotiations are executed consecutively until an agreement is reached or no more opportunities are available. Each bargaining game is reduced to a single round, where player 1 makes an offer and player 2 can only respond by rejecting or accepting this offer. Using an evolutionary simulation we study several properties of this market game. It appears that the outcomes depend on the information that is available to the players. If players are informed about the bargaining opportunities of their opponents, the first player in turn has the advantage and always proposes a take-it-or-leave-it deal that leaves the other player with a relatively poor outcome. This outcome is consistent with a game-theoretic analysis which we also present in this thesis. If this information is not available, a theoretical analysis is very hard. The evolutionary simulation, however, shows that in this case the responder obtains a better deal. This occurs because the first player can no longer anticipate the response of the other player, and therefore bids lower to avoid a disagreement. In this thesis, we additionally consider other factors that influence the outcomes of the market game, such as negotiation over multiple issues simultaneously, search costs, and break off probabilities.

Besides fundamental issues, this thesis presents a number of business-related applications of automated bargaining, as well as generic bargaining strategies for agents that can be employed in related areas. As a first application, we introduce a framework where negotiation is used for recommending shops to customers, for example on a web page of an electronic shopping mall. Through a market-driven auction a relevant selection of shops is determined in a distributed fashion. This is achieved by selling a limited number of banner spaces in an electronic auction. For each arriving customer on the web page, shops can automatically place bids for this "customer attention space" through their *shop agents*. These software agents bid based on a customer profile, containing personal data of the customer, such as age, interests, and/or keywords in a search query. The shop agents are adaptive and learn, given feedback from the customers, which profiles to target and how much to bid in the auction. The highest bidders are then selected and displayed to the customer. The feasibility of this distributed approach for matching shops to customers is demonstrated using an evolutionary simulation. Several customer models and auction mechanisms are studied, and we show that the market-based approach results in a proper selection of shops for the customers.

Bargaining can be especially beneficial if not only the price, but other aspects are considered as well. This allows for example to customise products and services to the personal preferences of a user. We developed a system makes use of these properties for selling and personalising so-called information goods, such as news articles, software, and music. Using the alternating-offers protocol, a seller agent negotiates with several buyers simultaneously about a fixed price, a per-item price, and the quality of a bundle of information goods. The system is capable of taking into account important business-related conditions such as the fairness of the negotiation. The agents combine a search strategy and a concession strategy to generate offers in the negotiations. The concession strategy determines the amount the agent will concede each round, whereas the search strategy takes care of the personalisation of the offer. We introduce two search strategies in this thesis, and show through computer experiments that the use of these strategies by a buyer and seller agent, result in personalised outcomes, also when combined with various concession strategies. The search strategies presented here can be easily applied to other domains where personalisation is important.

In addition, we also developed concession strategies for the seller agent that can be used in settings where a single seller agent bargains with several buyer agents simultaneously. Even if bargaining itself is bilateral (i.e., between two parties), a seller agent can actually benefit from the fact that several such negotiations occur concurrently. The developed strategies are focussed on domains where supply is flexible and can be adjusted to meet demand, like for information goods. We study fixed strategies, time-dependent strategies and introduce several auction-inspired strategies. Auctions are often used when one party negotiates with several opponents simultaneously. Although the latter strategies benefit from the advantages of auctions, the actual negotiation remains bilateral and consists of exchanging offers and counter offers. We developed an evolutionary simulation environment to evaluate the seller agent's strategies. We especially consider the case where buyers are time-impatient and under pressure to reach agreements early. The simulations show that the auction-inspired strategies are able to obtain almost maximum profits from the negotiations, given sufficient time pressure of the buyers.

Curriculum Vitae

Enrico Gerding was born on the 10th of August 1974 in Haarlem, The Netherlands. He studied artificial intelligence at the Vrije Universiteit in Amsterdam. He graduated with honours in August 1999, on the subject of evolutionary models for multi-issue negotiations. The Master's project was carried out at the Centrum voor Wiskunde en Informatica in Amsterdam under supervision of Han La Poutré and David van Bragt.

In November 1999 Enrico started as Ph.D. student at the Centrum voor Wiskunde en Informatica. The research, that resulted in this thesis, was undertaken within a four-year project entitled "Autonomous Systems of Trade Agents", under supervision of Han La Poutré.

Titles in the IPA Dissertation Series

J.O. Blanco. *The State Operator in Process Algebra.* Faculty of Mathematics and Computing Science, TUE. 1996-01

A.M. Geerling. *Transformational Development of Data-Parallel Algorithms.* Faculty of Mathematics and Computer Science, KUN. 1996-02

P.M. Achten. *Interactive Functional Programs: Models, Methods, and Implementation.* Faculty of Mathematics and Computer Science, KUN. 1996-03

M.G.A. Verhoeven. *Parallel Local Search.* Faculty of Mathematics and Computing Science, TUE. 1996-04

M.H.G.K. Kesseler. *The Implementation of Functional Languages on Parallel Machines with Distrib. Memory.* Faculty of Mathematics and Computer Science, KUN. 1996-05

D. Alstein. *Distributed Algorithms for Hard Real-Time Systems.* Faculty of Mathematics and Computing Science, TUE. 1996-06

J.H. Hoepman. *Communication, Synchronization, and Fault-Tolerance.* Faculty of Mathematics and Computer Science, UvA. 1996-07

H. Doornbos. *Reductivity Arguments and Program Construction.* Faculty of Mathematics and Computing Science, TUE. 1996-08

D. Turi. *Functorial Operational Semantics and its Denotational Dual.* Faculty of Mathematics and Computer Science, VUA. 1996-09

A.M.G. Peeters. *Single-Rail Handshake Circuits.* Faculty of Mathematics and Computing Science, TUE. 1996-10

N.W.A. Arends. *A Systems Engineering Specification Formalism.* Faculty of Mechanical Engineering, TUE. 1996-11

P. Severi de Santiago. *Normalisation in Lambda Calculus and its Relation to Type Inference.* Faculty of Mathematics and Computing Science, TUE. 1996-12

D.R. Dams. *Abstract Interpretation and Partition Refinement for Model Checking.* Faculty of Mathematics and Computing Science, TUE. 1996-13

M.M. Bonsangue. *Topological Dualities in Semantics.* Faculty of Mathematics and Computer Science, VUA. 1996-14

B.L.E. de Fluiter. *Algorithms for Graphs of Small Treewidth.* Faculty of Mathematics and Computer Science, UU. 1997-01

W.T.M. Kars. *Process-algebraic Transformations in Context.* Faculty of Computer Science, UT. 1997-02

P.F. Hoogendijk. *A Generic Theory of Data Types.* Faculty of Mathematics and Computing Science, TUE. 1997-03

T.D.L. Laan. *The Evolution of Type Theory in Logic and Mathematics.* Faculty of Mathematics and Computing Science, TUE. 1997-04

C.J. Bloo. *Preservation of Termination for Explicit Substitution.* Faculty of Mathematics and Computing Science, TUE. 1997-05

J.J. Vereijken. *Discrete-Time Process Algebra.* Faculty of Mathematics and Computing Science, TUE. 1997-06

F.A.M. van den Beuken. *A Functional Approach to Syntax and Typing.* Faculty of Mathematics and Informatics, KUN. 1997-07

A.W. Heerink. *Ins and Outs in Refusal Testing.* Faculty of Computer Science, UT. 1998-01

G. Naumoski and W. Alberts. *A Discrete-Event Simulator for Systems Engineering.* Faculty of Mechanical Engineering, TUE. 1998-02

J. Verriet. *Scheduling with Communication for Multiprocessor Computation.* Faculty of Mathematics and Computer Science, UU. 1998-03

J.S.H. van Gageldonk. *An Asynchronous Low-Power 80C51 Microcontroller.* Faculty of Mathematics and Computing Science, TUE. 1998-04

A.A. Basten. *In Terms of Nets: System Design with Petri Nets and Process Algebra.* Faculty of Mathematics and Computing Science, TUE. 1998-05

E. Voermans. *Inductive Datatypes with Laws and Subtyping – A Relational Model.* Faculty of Mathematics and Computing Science, TUE. 1999-01

H. ter Doest. *Towards Probabilistic Unification-based Parsing.* Faculty of Computer Science, UT. 1999-02

J.P.L. Segers. *Algorithms for the Simulation of Surface Processes.* Faculty of Mathematics and Computing Science, TUE. 1999-03

C.H.M. van Kemenade. *Recombinative Evolutionary Search.* Faculty of Mathematics and Natural Sciences, UL. 1999-04

E.I. Barakova. *Learning Reliability: a Study on Indecisiveness in Sample Selection.* Faculty of Mathematics and Natural Sciences, RUG. 1999-05

M.P. Bodlaender. *Schedulere Optimization in Real-Time Distributed Databases.* Faculty of Mathematics and Computing Science, TUE. 1999-06

M.A. Reniers. *Message Sequence Chart: Syntax and Semantics.* Faculty of Mathematics and Computing Science, TUE. 1999-07

J.P. Warners. *Nonlinear approaches to satisfiability problems.* Faculty of Mathematics and Computing Science, TUE. 1999-08

J.M.T. Romijn. *Analysing Industrial Protocols with Formal Methods.* Faculty of Computer Science, UT. 1999-09

P.R. D'Argenio. *Algebras and Automata for Timed and Stochastic Systems.* Faculty of Computer Science, UT. 1999-10

G. Fábián. *A Language and Simulator for Hybrid Systems.* Faculty of Mechanical Engineering, TUE. 1999-11

J. Zwanenburg. *Object-Oriented Concepts and Proof Rules.* Faculty of Mathematics and Computing Science, TUE. 1999-12

R.S. Venema. *Aspects of an Integrated Neural Prediction System.* Faculty of Mathematics and Natural Sciences, RUG. 1999-13

J. Saraiva. *A Purely Functional Implementation of Attribute Grammars.* Faculty of Mathematics and Computer Science, UU. 1999-14

R. Schiefer. *Viper, A Visualisation Tool for Parallel Progam Construction.* Faculty of Mathematics and Computing Science, TUE. 1999-15

K.M.M. de Leeuw. *Cryptology and Statecraft in the Dutch Republic.* Faculty of Mathematics and Computer Science, UvA. 2000-01

T.E.J. Vos. *UNITY in Diversity. A stratified approach to the verification of distributed algorithms.* Faculty of Mathematics and Computer Science, UU. 2000-02

W. Mallon. *Theories and Tools for the Design of Delay-Insensitive Communicating Processes.* Faculty of Mathematics and Natural Sciences, RUG. 2000-03

W.O.D. Griffioen. *Studies in Computer Aided Verification of Protocols.* Faculty of Science, KUN. 2000-04

P.H.F.M. Verhoeven. *The Design of the MathSpad Editor.* Faculty of Mathematics and Computing Science, TUE. 2000-05

J. Fey. *Design of a Fruit Juice Blending and Packaging Plant.* Faculty of Mechanical Engineering, TUE. 2000-06

M. Franssen. *Cocktail: A Tool for Deriving Correct Programs.* Faculty of Mathematics and Computing Science, TUE. 2000-07

P.A. Olivier. *A Framework for Debugging Heterogeneous Applications.* Faculty of Natural Sciences, Mathematics and Computer Science, UvA. 2000-08

E. Saaman. *Another Formal Specification Language.* Faculty of Mathematics and Natural Sciences, RUG. 2000-10

M. Jelasity. *The Shape of Evolutionary Search Discovering and Representing Search Space Structure.* Faculty of Mathematics and Natural Sciences, UL. 2001-01

R. Ahn. *Agents, Objects and Events a computational approach to knowledge, observation and communication.* Faculty of Mathematics and Computing Science, TU/e. 2001-02

M. Huisman. *Reasoning about Java programs in higher order logic using PVS and Isabelle.* Faculty of Science, KUN. 2001-03

I.M.M.J. Reymen. *Improving Design Processes through Structured Reflection.* Faculty of Mathematics and Computing Science, TU/e. 2001-04

S.C.C. Blom. *Term Graph Rewriting: syntax and semantics.* Faculty of Sciences, Division of Mathematics and Computer Science, VUA. 2001-05

R. van Liere. *Studies in Interactive Visualization.* Faculty of Natural Sciences, Mathematics and Computer Science, UvA. 2001-06

A.G. Engels. *Languages for Analysis and Testing of Event Sequences.* Faculty of Mathematics and Computing Science, TU/e. 2001-07

J. Hage. *Structural Aspects of Switching Classes.* Faculty of Mathematics and Natural Sciences, UL. 2001-08

M.H. Lamers. *Neural Networks for Analysis of Data in Environmental Epidemiology: A Case-study into Acute Effects of Air Pollution Episodes.* Faculty of Mathematics and Natural Sciences, UL. 2001-09

T.C. Ruys. *Towards Effective Model Checking.* Faculty of Computer Science, UT. 2001-10

D. Chkliaev. *Mechanical verification of concurrency control and recovery protocols.* Faculty of Mathematics and Computing Science, TU/e. 2001-11

M.D. Oostdijk. *Generation and presentation of formal mathematical documents.* Faculty of Mathematics and Computing Science, TU/e. 2001-12

A.T. Hofkamp. *Reactive machine control: A simulation approach using χ.* Faculty of Mechanical Engineering, TU/e. 2001-13

D. Bošnački. *Enhancing state space reduction techniques for model checking.* Faculty of Mathematics and Computing Science, TU/e. 2001-14

M.C. van Wezel. *Neural Networks for Intelligent Data Analysis: theoretical and experimental aspects.* Faculty of Mathematics and Natural Sciences, UL. 2002-01

V. Bos and J.J.T. Kleijn. *Formal Specification and Analysis of Industrial Systems.* Faculty of Mathematics and Computer Science and Faculty of Mechanical Engineering, TU/e. 2002-02

T. Kuipers. *Techniques for Understanding Legacy Software Systems.* Faculty of Natural Sciences, Mathematics and Computer Science, UvA. 2002-03

S.P. Luttik. *Choice Quantification in Process Algebra.* Faculty of Natural Sciences, Mathematics, and Computer Science, UvA. 2002-04

R.J. Willemen. *School Timetable Construction: Algorithms and Complexity.* Faculty of Mathematics and Computer Science, TU/e. 2002-05

M.I.A. Stoelinga. *Alea Jacta Est: Verification of Probabilistic, Real-time and Parametric Systems.* Faculty of Science, Mathematics and Computer Science, KUN. 2002-06

N. van Vugt. *Models of Molecular Computing.* Faculty of Mathematics and Natural Sciences, UL. 2002-07

A. Fehnker. *Citius, Vilius, Melius: Guiding and Cost-Optimality in Model Checking of Timed and Hybrid Systems.* Faculty of Science, Mathematics and Computer Science, KUN. 2002-08

R. van Stee. *On-line Scheduling and Bin Packing.* Faculty of Mathematics and Natural Sciences, UL. 2002-09

D. Tauritz. *Adaptive Information Filtering: Concepts and Algorithms.* Faculty of Mathematics and Natural Sciences, UL. 2002-10

M.B. van der Zwaag. *Models and Logics for Process Algebra.* Faculty of Natural Sciences, Mathematics, and Computer Science, UvA. 2002-11

J.I. den Hartog. *Probabilistic Extensions of Semantical Models.* Faculty of Sciences, Division of Mathematics and Computer Science, VUA. 2002-12

L. Moonen. *Exploring Software Systems.* Faculty of Natural Sciences, Mathematics, and Computer Science, UvA. 2002-13

J.I. van Hemert. *Applying Evolutionary Computation to Constraint Satisfaction and*

Data Mining. Faculty of Mathematics and Natural Sciences, UL. 2002-14

S. Andova. *Probabilistic Process Algebra.* Faculty of Mathematics and Computer Science, TU/e. 2002-15

Y.S. Usenko. *Linearization in μCRL.* Faculty of Mathematics and Computer Science, TU/e. 2002-16

J.J.D. Aerts. *Random Redundant Storage for Video on Demand.* Faculty of Mathematics and Computer Science, TU/e. 2003-01

M. de Jonge. *To Reuse or To Be Reused: Techniques for component composition and construction.* Faculty of Natural Sciences, Mathematics, and Computer Science, UvA. 2003-02

J.M.W. Visser. *Generic Traversal over Typed Source Code Representations.* Faculty of Natural Sciences, Mathematics, and Computer Science, UvA. 2003-03

S.M. Bohte. *Spiking Neural Networks.* Faculty of Mathematics and Natural Sciences, UL. 2003-04

T.A.C. Willemse. *Semantics and Verification in Process Algebras with Data and Timing.* Faculty of Mathematics and Computer Science, TU/e. 2003-05

S.V. Nedea. *Analysis and Simulations of Catalytic Reactions.* Faculty of Mathematics and Computer Science, TU/e. 2003-06

M.E.M. Lijding. *Real-time Scheduling of Tertiary Storage.* Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2003-07

H.P. Benz. *Casual Multimedia Process Annotation – CoMPAs.* Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2003-08

D. Distefano. *On Modelchecking the Dynamics of Object-based Software: a Foundational Approach.* Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2003-09

M.H. ter Beek. *Team Automata – A Formal Approach to the Modeling of Collaboration Between System Components.* Faculty of Mathematics and Natural Sciences, UL. 2003-10

D.J.P. Leijen. *The λ Abroad – A Functional Approach to Software Components.* Faculty of Mathematics and Computer Science, UU. 2003-11

W.P.A.J. Michiels. *Performance Ratios for the Differencing Method.* Faculty of Mathematics and Computer Science, TU/e. 2004-01

G.I. Jojgov. *Incomplete Proofs and Terms and Their Use in Interactive Theorem Proving.* Faculty of Mathematics and Computer Science, TU/e. 2004-02

P. Frisco. *Theory of Molecular Computing – Splicing and Membrane systems.* Faculty of Mathematics and Natural Sciences, UL. 2004-03

S. Maneth. *Models of Tree Translation.* Faculty of Mathematics and Natural Sciences, UL. 2004-04

Y. Qian. *Data Synchronization and Browsing for Home Environments.* Faculty of Mathematics and Computer Science and Faculty of Industrial Design, TU/e. 2004-05

F. Bartels. *On Generalised Coinduction and Probabilistic Specification Formats.* Faculty of Sciences, Division of Mathematics and Computer Science, VUA. 2004-06

L. Cruz-Filipe. *Constructive Real Analysis: a Type-Theoretical Formalization and Applications.* Faculty of Science, Mathematics and Computer Science, KUN. 2004-07

E.H. Gerding. *Autonomous Agents in Bargaining Games: An Evolutionary Investigation of Fundamentals, Strategies, and Business Applications.* Faculty of Technology Management, TU/e. 2004-08