



**Deutsches
Forschungszentrum
für Künstliche
Intelligenz GmbH**

**Research
Report**
RR-98-02

Decision Theory and Coordination in Multiagent Systems

Klaus Fischer, Christian Ruß, Gero Vierke

October 1998

**Deutsches Forschungszentrum für Künstliche Intelligenz
GmbH**

Postfach 20 80
67608 Kaiserslautern, FRG
Tel.: + 49 (631) 205-3211
Fax: + 49 (631) 205-3210
E-Mail: info@dfki.uni-kl.de

Stuhlsatzenhausweg 3
66123 Saarbrücken, FRG
Tel.: + 49 (681) 302-5252
Fax: + 49 (681) 302-5341
E-Mail: info@dfki.de

WWW: <http://www.dfki.de>

Deutsches Forschungszentrum für Künstliche Intelligenz
DFKI GmbH
German Research Center for Artificial Intelligence

Founded in 1988, DFKI today is one of the largest nonprofit contract research institutes in the field of innovative software technology based on Artificial Intelligence (AI) methods. DFKI is focusing on the complete cycle of innovation — from world-class basic research and technology development through leading-edge demonstrators and prototypes to product functions and commercialization.

Based in Kaiserslautern and Saarbrücken, the German Research Center for Artificial Intelligence ranks among the important “Centers of Excellence” worldwide.

An important element of DFKI's mission is to move innovations as quickly as possible from the lab into the marketplace. Only by maintaining research projects at the forefront of science can DFKI have the strength to meet its technology transfer goals.

DFKI has about 115 full-time employees, including 95 research scientists with advanced degrees. There are also around 120 part-time research assistants.

Revenues for DFKI were about 24 million DM in 1997, half from government contract work and half from commercial clients. The annual increase in contracts from commercial clients was greater than 37% during the last three years.

At DFKI, all work is organized in the form of clearly focused research or development projects with planned deliverables, various milestones, and a duration from several months up to three years.

DFKI benefits from interaction with the faculty of the Universities of Saarbrücken and Kaiserslautern and in turn provides opportunities for research and Ph.D. thesis supervision to students from these universities, which have an outstanding reputation in Computer Science.

The key directors of DFKI are Prof. Wolfgang Wahlster (CEO) and Dr. Walter Olthoff (CFO).

DFKI's six research departments are directed by internationally recognized research scientists:

- ❑ Information Management and Document Analysis (Director: Prof. A. Dengel)
- ❑ Intelligent Visualization and Simulation Systems (Director: Prof. H. Hagen)
- ❑ Deduction and Multiagent Systems (Director: Prof. J. Siekmann)
- ❑ Programming Systems (Director: Prof. G. Smolka)
- ❑ Language Technology (Director: Prof. H. Uszkoreit)
- ❑ Intelligent User Interfaces (Director: Prof. W. Wahlster)

In this series, DFKI publishes research reports, technical memos, documents (eg. workshop proceedings), and final project reports. The aim is to make new results, ideas, and software available as quickly as possible.

Prof. Wolfgang Wahlster
Director

Decision Theory and Coordination in Multiagent Systems

Klaus Fischer, Christian Ruß, Gero Vierke

DFKI-RR-98-02

This work has been supported by a grant from The Federal Ministry of Education, Science, Research, and Technology (FKZ 01 IW 801).

© Deutsches Forschungszentrum für Künstliche Intelligenz 1998

This work may not be copied or reproduced in whole or part for any commercial purpose. Permission to copy in whole or part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of the Deutsche Forschungszentrum für Künstliche Intelligenz, Kaiserslautern, Federal Republic of Germany; an acknowledgement of the authors and individual contributors to the work; all applicable portions of this copyright notice. Copying, reproducing, or republishing for any other purpose shall require a licence with payment of fee to Deutsches Forschungszentrum für Künstliche Intelligenz.

ISSN 0946-008X

Decision Theory and Coordination in Multiagent Systems

Klaus Fischer, Christian Ruß, Gero Vierke

September 1998

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 5 |
| 2 | Decision and Game Theory | 8 |
| 2.1 | Introduction | 8 |
| 2.1.1 | Historical Overview | 9 |
| 2.1.2 | Motivation | 10 |
| 2.1.3 | Epistemological Goals of Decision Theory | 12 |
| 2.1.4 | The Basic Model | 15 |
| 2.2 | Single Agent Decisions | 21 |
| 2.2.1 | Decisions under Certainty | 22 |
| 2.2.2 | Decisions under Risk | 22 |
| 2.2.3 | Decisions under Uncertainty | 23 |
| 2.3 | Two Agents Decisions | 26 |
| 2.3.1 | Game Theoretical Concepts | 26 |
| 2.3.2 | Two-Person Constant-Sum Games | 31 |
| 2.3.3 | Two-Person Non-Constant-Sum Games | 32 |
| 2.3.4 | Iterated Games | 37 |
| 2.4 | Multiple Agents Decisions | 38 |
| 2.4.1 | Coalitions in Cooperative N -Person Games | 38 |
| 2.4.2 | Fundamental Concepts for the Analysis of N -Person Games | 39 |
| 2.5 | Solution Concepts for N -Person Games | 41 |
| 2.5.1 | Von Neumann and Morgenstern's Solution: the Kernel | 41 |
| 2.5.2 | Solutions for Games without Kernel | 45 |
| 3 | Game Theory in Multiagent Systems | 48 |
| 3.1 | Definition of the Domains | 48 |
| 3.1.1 | Task Oriented Domain | 48 |
| 3.1.2 | State Oriented Domain | 50 |
| 3.1.3 | Worth Oriented Domain | 51 |
| 3.2 | Cooperation and Deception | 51 |
| 3.2.1 | Negotiation Set, Protocols and Strategies | 52 |
| 3.2.2 | Nash's Criteria for a Fair Negotiation Mechanism | 52 |
| 3.2.3 | Deception in Task Oriented Domains | 53 |

| | | |
|----------|---|------------|
| 3.2.4 | Cooperation and Deception in SODs and WODs | 54 |
| 4 | Coordination Mechanism Design | 57 |
| 4.1 | Economic Theory of Mechanism Design | 58 |
| 4.1.1 | The Assignment Problem | 58 |
| 4.1.2 | The Revelation Principle | 62 |
| 4.2 | Classifying Auction-Based Mechanisms | 64 |
| 4.3 | Simple Auction Mechanisms | 67 |
| 4.4 | The Vickrey (Second-Price Sealed-Bid) Auction | 68 |
| 4.4.1 | Advantages of the Vickrey Auction | 68 |
| 4.4.2 | Shortcomings of the Vickrey Auction | 69 |
| 4.5 | The Multistaged Extended Vickrey Auction | 72 |
| 4.5.1 | Specification of the MEVA Procedure | 72 |
| 4.5.2 | Analysis of the MEVA | 77 |
| 4.5.3 | Application of the Shapley Value for Profit Division in the MEVA | 78 |
| 4.6 | Matrix Auctions | 80 |
| 4.6.1 | Matrix Auctions for Multiple Identical Items | 81 |
| 4.6.2 | Matrix Auctions for Multiple Heterogeneous Items | 82 |
| 4.7 | Clarke Taxing Auction Procedures | 83 |
| 4.7.1 | The Clarke Tax | 83 |
| 4.7.2 | The Generalised Vickrey Auction | 89 |
| 4.7.3 | The Vickrey-Leonard Auction | 90 |
| 4.7.4 | The VL Chit Mechanism | 93 |
| 4.8 | The DGS Auction | 94 |
| 4.8.1 | The Exact DGS Auction Mechanism for Multiple Items | 94 |
| 4.8.2 | The Approximate DGS Auction Mechanism | 96 |
| 4.9 | The Escalating-Bid Auction | 98 |
| 5 | Application Example: The Transportation Domain | 100 |
| 5.1 | The Domain of Application | 100 |
| 5.1.1 | Modelling the Transportation Domain as a Multiagent Sys- tem | 101 |
| 5.1.2 | Related Work and Previous Results | 102 |
| 5.2 | Analysis of the Transportation Domain | 103 |
| 5.2.1 | Definition and Basic Properties | 103 |
| 5.2.2 | Encounters and Deals in the Transportation Domain | 107 |
| 5.2.3 | Cooperative vs Competitive TODs | 109 |
| 5.3 | Lying in the Transportation Domain | 109 |
| 5.3.1 | Phantom Tasks and Hidden Tasks | 110 |
| 5.4 | Price Determination in Transportation Companies | 112 |
| 5.4.1 | Market Price | 112 |
| 5.4.2 | Modelling the Additional Costs | 113 |

| | | |
|----------|---|------------|
| 5.4.3 | Adequate Modelling of Costs | 115 |
| 5.4.4 | The Statistic Market Model | 116 |
| 5.5 | The Negotiation Situation | 117 |
| 5.5.1 | Evaluation of the Auction Protocols | 118 |
| 5.5.2 | A Negotiation Strategy for Competing Shipping Companies | 120 |
| 5.6 | The Integration of the Auction Protocols and the Market Model . | 123 |
| 5.6.1 | Vickrey Auction | 123 |
| 5.6.2 | Contract Net | 124 |
| 6 | Conclusion | 126 |
| | Bibliography | 127 |

Chapter 1

Introduction

One of the main ideas of multiagent systems (MAS) is to generate approximate solutions to hard problems by distributing them to autonomous rational problem solvers (*agents*) that have local problem solving capabilities and are able to find a solution for the whole problem by cooperating with each other. Therefore, MAS research has great interest in the coordination among autonomous agents. A convenient way of dealing with coordination is by regarding it as *negotiation* among autonomous agents [Chaib-Draa 91]. Hence, coordination and negotiation are two closely related subjects. There is a significant body of publications on coordination in general (e.g. [Corkill & Lesser 83, Durfee 88, Durfee & Montgomery 90, Malone 90, Decker & Lesser 94]), and on negotiation in particular (e.g. [Sycara 87, Decker & Lesser 90, Steiner et al. 90, Conry et al. 91, Lux & Steiner 95]) which document the MAS perspective to these subjects.

When thinking about negotiation, MAS researchers often adopt the ideas of speech act theory which originate from Austin and Searle [Austin 62, Searle 69]. The main idea underlying this approach is to regard communication among autonomous agents as a special type of action aiming at modifying the mental state of the recipient of a message (cf. [Cohen & Perrault 86, Perrault 90, Bussmann & Müller 92, Lux 95]). In this light, making a specific utterance is just one of many options an agent has when it decides which action to take next; negotiation, thus, is a process in which an agent tries to reach a specific goal by means of performing a sequence of communicative acts.

The aim of this report is to provide an introduction to the theoretical foundations of rational decision making and to give an overview of how decision theory and game theory relate to the field of multiagent systems.

The structure of the report is as follows: we start with presenting the basic framework of classic decision and game theory as it was defined by von Neumann and Morgenstern [von Neumann & Morgenstern 44] and continue with explaining the approach of Rosenschein and Zlotkin [Rosenstein & Zlotkin 94a] to apply game theory to multiagent systems. Subsequently, we look at the principles

of coordination mechanism design from an economic point of view and present several auction-based allocation mechanisms. Finally, the transportation domain is introduced as an application example in order to show the importance of the presented concepts in real-world multiagent applications. We conclude this report with a brief review of the presented concepts and how they can be used to enable a social welfare maximising coordination within multiagent systems composed of self-interested agents.

The following paragraphs give a brief summary of the chapters of this report.

Decision and Game Theory: The general objectives of decision theory are to provide formal mathematical models for decision situations, to analyse them according to their special properties, to develop adequate definitions of notions like “solution” or “rational decision” and to provide methods for the computation of solutions and the design decision rules for rational action selection in decision situations. In particular the concepts of game theory (e.g., for the rational behaviour of an agent) are not derived from psychological considerations but solely from the formal structural properties of the abstract decision situation. Game theory can be viewed as a specific branch of decision theory as it concerns decision situations where the consequences of actions depends on the actions of other agents. Therefore, the concepts and methods of decision theory also hold for the decision situations occurring in game theory. Decision theory characterises decision problems and their solutions by the consequences of rules and interaction patterns of the participating agents. The interest in decision theory is not only motivated by the need for constructive advice and decision rules for specific decision problems but also for the design and analysis of organisational structures, allocation mechanisms, voting procedures, conflict management, etc. After a brief history and motivation of decision theory in Section 2.1, the basic concepts of decision and game theory are outlined and decision situations are classified. Section 2.2 discusses classic decision theory by presenting general decision rules for single-agent decision-making. In Section 2.3 constant-sum games, non-constant-sum games, and iterated games are examined using some well-known examples. Section 2.4 concludes the chapter with a discussion of the formation of coalitions and the distribution of profit in the n-player case.

Game Theory in Multiagent Systems: A major contribution of Rosenschein and Zlotkin was to transfer abstract game theoretic concepts to multiagent applications [Zlotkin & Rosenschein 93, Rosenschein & Zlotkin 94a]. We present their classification of multiagent domains into a three-level hierarchy (Section 3.1). Furthermore, we discuss the chances for cooperation and the incentive for deception in these classes of domains and examine whether there exist incentive compatible mechanisms that serve to prevent deception (Section 3.2). We concentrate on the class of task-oriented domains which we need for the charac-

terisation of our example domain in Chapter 5.

Coordination Mechanism Design: The problem of designing auction-based coordination mechanisms for autonomous agents is addressed from an economic viewpoint as presented in [Ruß 97]. The central question is how to solve the allocation problem in a setting where the agents are self-interested, i.e., are willing to betray other agents for their own benefit. The adequate way of dealing with selfish agents is the use of mechanisms that enforce truthful bidding strategies. After a brief introduction to the theory of mechanism design in Section 4.1, we provide a classification of auction protocols (Section 4.2). Subsequently, we shortly discuss the most popular auction mechanisms before we introduce several more sophisticated mechanisms based on the Vickrey principle.

Application Example: The transportation domain is presented as an application example. The domain is characterised as a task-oriented domain following the framework described in section 3.1. However, we come up with the result that it is useful to divide the class of task-oriented domains into the two subclasses of cooperative and competitive task-oriented domains. For both of them, the transportation domain provides examples: The cooperative setting characterises the situation within one shipping company where the main problem consists in a complex scheduling problem, namely to distribute orders to a set of trucks in an optimal manner. In the competitive setting, several transportation companies compete on tasks and profit. Furthermore, we analyse the applicability of some of the auction-based negotiation protocols presented in Chapter 4.

Conclusion: We give a brief overview of how the presented concepts adjust to each other and discuss to what extent they can help to provide a theoretical foundation for coordination which maximises social welfare of self-interested agents.

Chapter 2

Decision and Game Theory

2.1 Introduction

The general objectives of decision theory are to provide formal mathematical models of decision situations, to analyse them according to their special properties, to develop adequate definitions of notions like “solution” or “rational decision” and to design methods for the computation of solutions and decision rules for rational action selection in decision situations. In particular, the concepts developed within decision theory (e.g., for the rational behaviour of an agent) did not originate from psychological considerations but were derived solely from the formal structural properties of the abstract decision situation. Decision theory is able to characterise decision-problems and their solutions by the consequences of rules and interaction patterns of the participating agents. The interest in decision theory is motivated by the need for constructive advice and decision rules for specific decision problems. Additionally, it helps to design and to analyse organisational structures, allocation mechanisms, voting procedures, and conflict management.

While the notion of decision theory is not defined consistently in literature, we will view decision theory in this report as comprising classic decision theory, i.e., decision situations involving one agent, as well as game theory where the situation is characterised by several agents whose decisions are not independent of each others.

The classical decision situation can be sketched as follows: an agent has to select one action from a set of alternative actions. The result of the selected action may depend on factors the decision-maker is not able to control. In deciding to select an action, the agent has to rank the different expected results (which depend on the expected future state of the environment) and then has to determine a preference ordering for the different possible actions based on its goals and preferences.

Game theory considers situations, where several agents simultaneously select one

out of several actions which are called strategies in this context. So far the classic situation is just duplicated, but the new feature is that the outcome for each agent is not independent from the strategies selected by the other agents. Therefore, each agent should consider what the other agents might do which of course crucially depends on their goals, preferences, and assumptions about other agents' decisions. This involves a recursive reasoning about other agents' reasoning which can lead to paradoxical situations in game theory.

The classic single agent decision situation can be looked upon as a game of an agent against nature, where “nature” denotes a probability distribution of future world states which is independent of the decision the agent takes. The “nature-agent” thus represents an ignorant player that bases its decision purely on a probability distribution. On the other hand, from a solipsistic point of view an agent in a game could regard the actions of all other agents as environmentally determined future states and base its decision on classic decision theoretic single agent considerations. This shows the relation between the two domains.

The structure of this chapter is as follows: in the following paragraphs we provide a brief overview of the history of decision theory and make some remarks concerning the motivation for decision theory especially in the context of planning. This leads to the discussion of different epistemological approaches that have been developed in decision theory. Subsequently, we explain the terminology and the basic model of decision making and provide a classification of decision situations. We conclude the introduction section by illuminating the fundamental concepts of game theory.

In Section 2.2, decision situations for one-person decision problems are classified into decisions with certainty, decisions with risk, and decisions with uncertainty. For the later two some standard decision rules are discussed.

In Section 2.3, some classical examples are utilised to illustrate the properties of two-person games. Constant-sum, non-constant-sum, and iterated games are discussed.

Finally, in Section 2.4, the aspects of n-person game theory are discussed: coalition forming and distribution of payoff between the coalition partners.

2.1.1 Historical Overview

The historical roots of decision theory date back to the middle of the 17th century when bored French aristocrats enjoyed to waste their time with gambling. Hoping to increase their winnings, some of them—together with more educated contemporaries—put great effort into the deep analysis of the strategic and systematic rules underlying these games. As a result, elementary principles for dealing with probabilities were discovered and laid the basis for probability theory. From these origins the probability calculus and on that basis modern statistics were developed.

First important contributions to the theory of strategic games came up at

the beginning of our century by [Zermelo 12], but initially failed to meet major attention. Thus, the analysis of strategic games lacked general attention until in 1944 the mathematician John von Neumann and the economist Oskar Morgenstern published their “Theory of Games and Economic Behaviour” [von Neumann & Morgenstern 44]. This was the first time that the analysis of strategic games met with public interest, especially from economists and mathematicians. Therefore, these researchers are considered to be the initiators of modern well-founded game theory.

The book “Games and Decisions” of Luce and Raiffa [Luce & Raiffa 57] later succeeded in introducing decision theory into the areas of social science and psychology. While before 1957, only a few empirical investigations about the behaviour of people in strategic decision situations were published, after the publication of this book, a constantly growing research area of experimental games evolved. Here, more or less abstract variations of the 2-person-game known as prisoner’s dilemma (cf. Section 2.3.3) dominated research until the middle of the 70’s; then the focus was shifted to other 2-person and n-person games.

2.1.2 Motivation

The primary motivation for building a theory that provides mathematical foundations and methods to formally analyse and explain the structure of various decision situations is based on the hope to understand the decision behaviour of people and to find better ways to cope with decision situations when they occur in real life. The classical basic decision situation confronts an agent with several action alternatives. The agent has to react by selecting one of these actions, hopefully that one that maximises its satisfaction in the environmental state achieved by the action. Therefore, the agent bases its decision on its goals, its preferences, and its model about the world, i.e. the relevant objects, their relations to each other, the effects of operations, and the interplay of all these. Additional complexity may be added to this situation in various ways, e.g. by considering incomplete knowledge or uncertainty about the development of the environment or by considering several agents in cooperative or competitive relations.

While all this is treated in decision theory, it is important for the motivation of the theory to ask: “Where does the decision situation come from?” This question puts the decision situation into an application context and depending on the domain and description level of this context several answers may occur. One possible example that will be considered in the following is the context of planning because it is sufficiently abstract to comprise several situations. Planning is reasoning about future states of the world and how to achieve them. This presupposes a representation of the world and a function for appraising different possible world states based on an internal valuation by the agent. A problem initially arises from a difference between the actual state of the world as it is perceived by the agent and a world state that the agent aims to achieve based

on its wishes, aims, and goals. The following steps will be necessary to prepare decision situations:

- Domain model (object system): In most non-trivial cases of real-world problems an appropriate model of the domain initially does not exist but has to be developed using methods from organisational research and system analysis. A problem like “How can we achieve that the accounting department of our company works more efficiently?” can only be solved if we are able to describe how the accounting department is working at all. What are the relevant objects and by which attributes are they characterised? What are the relations between them? Which processes are performed by whom and what are the necessary resources? What are the interfaces to other systems? Which parameters will influence the system into which direction? Design decisions taken during this process of the representation of the domain may have important influences on the decision situations occurring later on in the process of planning.
- Goal system: based on this model the goal has to be specified. What are the criteria that appropriately measure the efficiency of the department’s work? Efficiency has to be described in terms of measurable parameters, i.e., we have to break down the goal into subgoals that allow us to rank the achievement of a goal by access to specific goal criteria (measurability). This problem, also known as operationalisation of the goals, may also involve that subgoals interact with each others and the problem of how the goal criteria can be scaled and accumulated to obtain a uniform measure for the overall goal. Subgoals should cover all aspects of the overall goal (completeness) and on the other hand should not cover certain aspects by several subgoals (no redundancy). Again the decisions taken in this step may provide important constraints for the following steps.
- Decision field: Planning in the narrow sense is necessary to develop the decision field which formally could be described as a mapping of states and actions to results. A set of possible actions presumably predetermined by technical, organisational, legal or any other type of constraints of the system has to be developed and the expected transformations of the system according to these actions must be estimated. Anticipation of future states of the system extends the object model to a prognostic model. The focus will be on those actions that have a crucial effect on the goal criteria.
- Valuing of Alternatives: The result of an action will usually comprise a bundle of effects on different objects, features, and attributes of the system. The goal system determines the relevant criteria for the evaluation. While the overall goal is to obtain a unique utility measure for each alternative action, the problems of the numerical measurement of those criteria

involve questions of collecting data, mapping them onto appropriate scales and combining different scales (e.g. nominal, ordinal, or cardinal scales). Methods from management science like cost benefit analysis or cost effectiveness analysis come into play if a monetary valuation of the results is possible. Subjective preferences will unavoidably influence the utility functions. Thus, if several agents are involved in the decision process¹, it has to be negotiated about their individual preferences. Furthermore, different degrees of determinateness must be treated due to lack of information or due to the fuzziness of attribute values.

- Selection: Finally, the selection of one alternative action is guided by the decision logic of the deciding agent. The decision rules which we will present in Section 2.2 show that different rules may lead to different decisions. Their main differences are due to the parameters they use to model the agent's willingness to take risky decisions.

It should be noted that complex strategic planning problems—e.g. increasing the efficiency of an organisation by a process of several steps of reorganisation²—are obviously far from being solved by the selection of one basic action. The implementation and supervision of a solution to problems of this kind often reveal discrepancies between the reality and the model, which may enforce iterations of the steps presented above.

2.1.3 Epistemological Goals of Decision Theory

The previous section reveals that multifaceted problems occur in the context of decision situations. From an epistemological point of view, it is a matter of discussion what intentions are pursued and which kind of results should be delivered from a science considering decision situations. Should decision theory be treated as a formal tool box for obtaining intrinsic truths in the style of mathematics and logic? Should it result in rules and advice telling how rational decisions are taken? Should it define or examine rationality? Should it describe and explain what people in real decision situations are doing? Several different directions of decision theory grew out of these questions.

The approach of **formal decision theory** is very close to the methodology of mathematics. It tries to discover fundamental structural relationships and causalities based only on the formal definitions of decision situations and some basic axioms defining rational behaviour. This formal approach builds the kernel for several branches in which decision theory split up.

¹Mechanisms for bargaining, voting and electing are another important domain demanding decision theoretic foundations.

²e.g., retraining of employees, investment into computer equipment, etc.

Normative Decision Theory

Normative decision theory is based on fundamentally rational acting agents (**postulate of rationality**). Its goal is to provide instructions for an agent to act in a rational manner when the premises of the decision are given.

The normative approach tries to develop very basic and insightful axiom systems of rational behaviour for different types of decision situations. This is important because the rational behaviour of maximising utility ceases to be a simple concept when several goal criteria, other agents' goals and actions or uncertainty come into play.

Criteria and rules for the appraisal of alternative actions according to those axioms are derived. The normative approach of decision theory claims to be constructive, i.e., to provide advice for an agent in a decision situation how to decide according to some well-defined axioms of rationality.

Its goal does not include to show how decisions indeed are made by human beings in real life. The models used do not consider how to obtain the premises of decisions, how to obtain the necessary informations or how the decision process is influenced by the environment. Therefore, the models of normative decision theory are referred to as **closed models**.

The concept of rationality is central for normative decision theory and the interpretation of rationality is therefore of major importance.

- In a general interpretation the only requirement for the agent's goal system is that it is free of contradictions. Because no further requirements about the substantial contents are demanded, this is called **formal rationality**.
- In addition to this, **substantial rationality** can only be defined in relation to another referential goal system, i.e., substantial rationality can only be evaluated according to another notion of rationality defined by the referential goal system.

A referential goal system might be defined for example as a common goal system of several agents of a society. It can be the case, that an agent acts formally rational according to its own goal system, but substantially irrational according to the common referential goal system. Moreover, it can be distinguished between

- **objective** rationality, if the decisions of an agent base on an objective representation of the world, i.e. if he models a decision situation in the same way as an objective observer would do it.
- **subjective** rationality, if the decisions of an agent base not on an objective representation of the world but only on a restricted subjective representation of the world.

An aspect often neglected in closed models is that information gathering consumes resources and thus increases costs in practice. Therefore, a rational decision in addition to considering the utility of the different results should also take into account these costs of gathering information. This could lead to situations where the expected costs of gathering additional informations to the model is greater than the expected gain of utility that might be achieved in using the more informed model in comparison to the less informed one.

In the sense of the normative approach most human decisions are considered to be not rational. The normative model is a formal deductive approach which bases on a basic set of axioms defining rational behaviour. Therefore, it must be distinguished from an approach to analyse human decision behaviour inductively based on descriptions of real decision situations and their specific contexts.

Descriptive Decision Theory

Descriptive decision theory represents an inductive and pragmatic approach that starts with the description of real decision behaviour of people in decision situations and analyses the consequences of decisions. It takes into account that the rationality of an individual in practice is bounded due to the limited capability of obtaining and processing information. This approach aims at the explanation of human decision behaviour in real situations.

For this purpose, the analysis of the origins of the decision premises and the relations between the deciding agent and its environment are of major interest. Descriptive models are denoted as **open models** because they explicitly consider the decision environment and the bounded rationality of real agents. They do not a-priori exclude non-rational decision behaviour and therefore are sometimes considered as the more comprising and more general approach. On the other hand, the normative theory can often give a more concrete guidance for making rational decisions though it neglects many social and psychological aspects of human decision making.

Synthesis: A Prescriptive Decision Theory

Prescriptive decision theory tries to provide a synthesis of normative and descriptive approaches by working with empirical data of real decision behaviour without neglecting the normative aspects. In contrast to the pure normative approach it also considers the origins of the decision premises as well as empirically observable deviations from the rational behaviour (e.g., caused by limited capabilities of gathering informations). Crucial questions involved in a decision like

- How do decision problems originate?
- How does the deciding agent obtain the premises of its decision?
- How are the agent and its environment influenced by each other?

can only be considered if an open system model is used. On the other hand, descriptions of decision processes and analyses of their consequences do not provide rules or recommendations for an agent in the actual decision situation. The prescriptive approach tries to provide normative instructions for agents in decision situations based on a situation-specific concrete preference behaviour.

Nevertheless, the decision-theoretic concepts that will be presented in the rest of this report originate from formal and normative decision theory.

2.1.4 The Basic Model

In this section, the terminology and the most basic definitions of decision theory are introduced. The basic model of decision theory is presented and its components are discussed in the subsequent paragraphs.

The basic model in decision theory (see Figure 2.1) distinguishes between the decision taker (**subject system**) and the environment (**object system**) and interprets the decision process as an interaction between these systems. The subject system can be further subdivided into the **information system**, the **decision logic**, and the **goal system**.

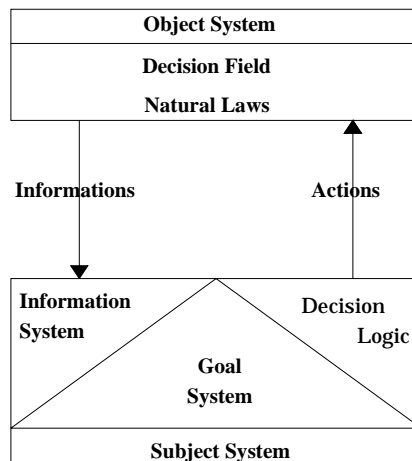


Figure 2.1: Basic Model of Decision Theory

The object system implements the representation of all relevant aspects of the environment. The subject’s information system receives informations about the environment and probably provides—by aggregation or pre-evaluation—the data that might influence the agent’s decision. According to its goal system the agent (subject system) uses its decision logic in order to select one of the possible actions. This action then will transform the object system into the desired state. The decision logic is responsible to take a rational decision on the basis of the given goal and the actual state of the object system (perceived via the information system).

Decision Field

The **decision field** is defined as a triple (A, Z, e) , where:

$A = \{a_1, a_2, \dots, a_n\}$ is the **action space**, i.e. the set of all possible alternative actions.

$Z = \{z_1, z_2, \dots, z_m\}$ is the **state space**, i.e. the set of all possible future states of the environment.

e with $x_{ij} = e(a_i, z_j)$ is the **result function** which maps every action and every state of the environment to a result.

The action space has to obey the **principle of perfect alternatives** which means:

1. all possible actions must be represented entirely (**principle of completeness**) such that the agent is forced to choose one of the actions. (Even doing nothing has to be represented explicitly by a specific nil-action if necessary.)
2. all actions must exclude the other ones (**principle of exclusion**) which means that no action overlaps with another one from the action space. Only one of the alternatives can be selected.

An instance of all relevant factors of the environment is called an environmental state. Analogous to the action space, the state space has to ensure mutual exclusion of the different states and a complete representation of all possible states. Thus the model requires that in a decision situation exactly one action is selected and exactly one state of the environment occurs. The result of an action in a specific environmental state is encoded in the result function e which maps action a_i depending on the state z_j to the result x_{ij} . In general, a result will be a vector of values representing different types of attributes like time, cost, payoff, etc. The result function can be represented in form of a **result matrix**:

| | states | | | |
|----------|----------|----------|----------|----------|
| | z_1 | z_2 | \dots | z_m |
| actions | | | | |
| a_1 | x_{11} | x_{12} | \dots | x_{1m} |
| a_2 | x_{21} | x_{22} | \dots | x_{2m} |
| \vdots | \vdots | \vdots | \ddots | \vdots |
| a_n | x_{n1} | x_{n2} | \dots | x_{nm} |

Knowledge about the probabilities of the occurrence of future environmental states can be obtained via the information system.

Information System

The availability of information about the future environmental states can be classified in the following way:

1. A **complete information system** represents the set of all possible states $Z = \{z_1, z_2, \dots, z_m\}$ and additionally knows in advance which future state will occur, i.e. it knows for sure the environmental state in which the result of the decision has to be evaluated. In this case of a **decision under certainty** the result matrix degenerates to a table with a single column and the result solely depends on the selected action.
2. An **incomplete information system** comprises the possibility of
 - a **decision under risk**, if the set of possible states is known, but for each state z_i the only additional information is the probability of its occurrence $P(z_i)$, such that $\sum_{i=1}^m P(z_i) = 1$. The result of an action is therefore only known statistically depending on the probability distribution of the environmental states.
 - a **decision under uncertainty**, if no information about the probability of the future states is available. The set of states Z just lists possible environmental states without any hint about the frequencies of their occurrences.

Goal System

The goal system contains the goals of the deciding agent ordered according to the preferences he associates with the different goals (e.g. minimisation of cost or maximisation of profit). Without a precise and clear notion of the goals an agent is unable to make rational goal-oriented decisions. Therefore, the design of the goal system is of crucial importance.

The goal system has to meet several requirements as for example the condition of **measurability**. This means, that the degree of achievement of a goal can be exactly measured by the value of a certain goal variable. For example, if the decision problem consists of selecting one of two possible routes connecting city X with city Y and the goal is to minimise the length of the route, then the achievement of the goal is exactly measurable by the amount of kilometres of the driven distance. On the other hand, when the goal is to select a “good” residence, it is not possible to measure the achievement of such an abstract general goal without defining a **goal hierarchy** (cf. [Eisenführ & Weber 93]), i.e. splitting the goal into several measurable subgoals, like for example minimal distance to downtown, minimal distance to the place of work, quietness few through traffic, and so on. Furthermore, the goal system has to specify the preferences given to these different subgoals.

Another requirement to the goal system is that the goals do not overlap with each other (**lack of redundancy**), because if several goals encode the same thing, it might be the case that by the preferences assigned to the goals something implicitly is weighted much more than originally intended. Like e.g. in the above example the goals “few through traffic” and “quietness” may at least partially overlap.

Other preconditions to the goal system are **completeness**, **independence**, and **simplicity**. For further considerations about the design of goal systems and for more examples we refer to [Eisenführ & Weber 93].

Comparability of Results

As mentioned above, the results obtained from the result function in general will be vectors of several values, because an action usually will effect several features of the environment and a decision will take into account several criteria. Problems occur when different attributes of the possible results recommend different actions to select (according to different subgoals).

The following example shows the conflict: Consider a situation where a shipping company has to drive from city X to city Y . There are two possible routes, one is short but often a little bit congested, the other one is longer but rarely used by anyone. The action space consists of driving either on the first or the second route $A = \{r_1, r_2\}$. The state space represents a situation with no congestion z_1 and another one with a congestion on the first route z_2 . The result function considers time and distance as shown in the following result matrix:

| | | |
|-------|-----------------|------------------|
| | z_1 | z_2 |
| r_1 | (55 min, 70 km) | (100 min, 70 km) |
| r_2 | (70 min, 90 km) | (70 min, 90 km) |

Trying to minimise the length of the route driven would recommend to select r_1 in any of the possible states. The goal to minimise time used to execute the transportation task would lead to the selection of r_1 in the case that state z_1 is known to occur. In the case of state z_2 the route r_2 should be selected. If now the goal is to minimise both, time and distance, we will select the better alternative r_1 assuming z_1 to occur, because it is better in both attributes. But in the case of knowing z_2 to occur, it is not clear from the result matrix which action to select because both attributes recommend different actions to choose.

When selecting an action alternative depending on the values of several attributes, we use in the first place the principle of **dominance**³:

Alternative a dominates alternative b if at least for one attribute the value of a is better than that of b , while for all other attributes the values of a are not worse than that of b .

³A formal definition of the game-theoretical term *dominance* will be given in Section 2.3.1.

It is clear that if a dominates b , then a should be selected. In our example this is the case in a certain decision knowing z_1 to occur: r_1 dominates r_2 , because 55 min is better than 70 min and 70 km is not worse than 90 km.

In the case of a certain decision assuming z_2 to occur, the notion of dominance does not help anymore because none of the alternatives dominates the other one. Because of the different measuring units (time in min, distance in km) the values are not comparable. The common approach used in decision theory to overcome this problem is to apply a **utility function** $u(x_{i,j}) = u_{i,j}$ that maps each result $x_{i,j}$ to a utility $u_{i,j}$ encoding the gain that the agent ascribes to the result. Utility is a uniform, comparable and substitutable unit, e.g., measured in the pecuniary units of the profit gained for the result. In our example this approach has to answer the questions “How much does each additional kilometre cost?” and “How much does each additional minute cost?”. The utility function has to encode these questions. The application of the utility function to each result in the result matrix transforms it to the **decision matrix**. It should be noted that the utility function has to encode the preferences the agent associates to the different results. While the results might comprise different incomparable and non-substitutable attributes, the utilities are by definition comparable and substitutable. Therefore, these qualities are kind of artificially imposed by the utility function, which shows why the design of a utility function is the most essential part for the agent to take rational decisions.

In the example above (certain decision in state z_2) it might be, that the shipping company associates costs of 1 unit per each additional kilometre but only 0.10 units for each additional minute. Then the profit of the order is reduced by 80 units in the case of r_1 and by 97 units in the case of r_2 . This will recommend to select r_1 . On the other hand, if a deadline is associated with the order, which specifies a substantial penalty for delivering the goods too late, the utility function has to encode the costs of additional time in a nonlinear way which in our example might recommend to select the faster route r_2 .

Classification of Decision Situations

This section provides an overview of the area of decision and game theory and presents major criteria to classify different decision situations. It is not possible here to illuminate all branches of the area in detail. References to more specific literature for the interested reader will be given. A classification of decision situations can be done according to the following issues:

- The **number of goals** to be used as decision criteria. It should be possible to measure the attainment of a goal by considering the value of a specific attribute of the environment. In the case of several goals several attributes must be considered. A balanced mixture of their values will contribute to the utility. Merging different attribute values into a unique measure of utility is denoted as **value aggregation**. It is often convenient for the model

(but unfortunately rarely adequate for realistic decision problems) to assume that the preferences for values of one attribute are independent from the values of another attribute. Several methods make the even stronger assumption that the units of measurement for the attributes are mutually substitutable. This presupposes a kind of mapping the values of different attributes onto a monetary unit of utility via an aggregation function. Linear aggregation functions are used for example in the “Multi Attribute Utility Theory” (MAUT) (cf. e.g.[French 88]) and in the method of “Analytic Hierarchy Process” (AHP) (cf. [Saaty 80]).

- The **mode of uncertainty** about the decision situation: in a secure decision, a state s_i is known to occur for sure, i.e. $P(s_i) = 1$. Therefore, the decision-matrix degenerates to a single column. A decision under risk is defined by an environment where the probabilities $P(s_i)$ of the occurrences of the different situations s_i are known a priori. From the principle of completeness it follows that $\sum_i P(s_i) = 1$ holds. A decision is made under uncertainty if even such a probability is unknown. Decision problems under uncertainty are sometimes handled like situations under risk by assigning some subjective probabilities to the different possible situations. Encoding the expectations about the results of different possible future states of the world into a unique measure of utility is denoted as **result aggregation**. Several rules for decisions under risk are encountered in literature which mainly differ according to the parameters modelling the agents willingness to take more or less risky decisions. Some popular decision rules for situations involving risk and uncertainty are discussed in Section 2.2.
- The **number of agents** involved in a decision: originally, only single agent decisions were considered within decision theory. Single agent scenarios are sometimes viewed as games against the environment, because the environmental state is besides the agent the only factor influencing the result. Most emphasise is put into the design of the utility function which has to implement reasonable mechanisms for value and result aggregation. Thereby, methods from statistics and probability theory are used.

If several agents are involved we can distinguish between collective decisions and game theoretic settings.

- In collective situations a group of agents has to make a common decision. Different preferences for attribute values, different goal priorities, or different views about the probability of possible states have to be reconciled by negotiation or an appropriate voting mechanism.

Models of negotiation and bargaining are related to the work of Zeuthen ([Zeuthen 30]) and treat bargaining as a process of convergence over time involving a sequence of offers and counter offers. These

models assume that the utility functions of the participants are fixed and known from the outset, and that a compromise zone exists, which can be identified and remains stable over time (cf. [Young 75]).

An interesting result about voting mechanisms (cf. [Schofield 85]) is Arrow's Theorem of Impossibility [Arrow 63], which concludes that any preference-aggregation for more than 2 alternatives cannot obey some rather intuitive and reasonable postulates of rationality.

- In game theory (cf. e.g. [Rauhut et al. 79]) all decisions must explicitly take into consideration the reactions of opponents. Each player tries to maximise its utility (usually called payoff in game theory) which leads to conflicts, especially when the winnings of one player correspond to the loss of another. One can distinguish game theoretic models according to the number of players, the dependency of the payoffs given to the players, the degree of cooperation versus competition, the information available to the players, the kind and number of strategies allowed to play and the influence of randomness.

In two-person games the result is determined by both players and both players are aware of their own and the opponent's possible utilities, which are represented in a payoff matrix. Therefore, strategic considerations and reasoning about the opponent's action come into play. This can lead to paradoxical situations where rational agents are unable to cooperate and are forced to select actions that result in a non optimal payoff.

Scenarios with more than two players are characterised by the possibility of building cooperative coalitions. Viewing a coalition as an agreed company, reduces the number of different parties, such that these scenarios in principle iteratively can be reduced to two-player games. Nevertheless, new problems arise concerning the questions which coalitions should emerge and how the payoff of a coalition should be distributed to its participants.

Iterated games can also be used to model evolutionary aspects (see e.g. [Axelrod 84]). The concept of evolutionary stable strategies (ESS) (cf. [Maynard-Smith 82]) is a useful concept to explain the emergence of agent behaviour as for example in the case of animals the institution of a dominance hierarchy, the territorial and mating behaviour or the behaviour in animal fights. (cf. [Dawkins 76])

2.2 Single Agent Decisions

An agent uses its decision logic to take decisions which gain maximal payoff. The decision logic includes parameters expressing the willingness to run higher risks

or representing its optimism. In this section, we present a number of decision rules that allow to take rational decisions according to the risk preferences of the agent. Thereby, according to the classification of the knowledge about the environment, we will distinguish in the following:

- Decisions under certainty, where the agent knows which environmental state will occur.
- Decisions under risk, where the agent knows the probabilities of the occurrence of the different environmental states.
- Decisions under uncertainty, where the agent only knows that one of the possible environmental states will occur but has no hint about the probabilities of occurrence.

In the following, we will consider these cases and present some decision rules, especially for the case of decisions under uncertainty.

2.2.1 Decisions under Certainty

In this case, the result matrix reduces to one column; that one, that represents the environmental state which is known to occur. The most simple case is the optimisation of one of the attributes of the result (or of a goal that depends on one attribute in a simple way). Then, the selection rule is to select the action whose result optimises (minimises or maximises depending of what this value expresses) that attribute value. If our goal depends on several of the result features, the first thing is to delete all alternatives that are dominated⁴ by others. These alternatives are also called **non-efficient**. For the remaining alternatives, a selection can only be made according to the preferences the agent associates with the different results. These are often introduced by a utility function that assigns a unique substitutable unit (comparable to a currency) to the results so that they can be compared to each other using their utility. It should be noted that the design of a utility function might involve several other problems not elaborated here in detail, e.g., concerning the type, scale and measurement of attributes or the dependencies of the preferences between different attributes. In the following, we assume that multi-attributed results are already mapped onto appropriate utility values.

2.2.2 Decisions under Risk

For each environmental state z_i , the agent knows the probability $P(z_i)$ of its occurrence, such that $\sum_{i=1}^m P(z_i) = 1$ holds. It is of minor concern here how these

⁴confer Section 2.1.4 for the definition of the dominance-relation

probabilities are obtained (cf. remarks about normative and descriptive decision-theory in Section 2.1.3) because the selection rules just assume them as given a priori. Nevertheless, the deciding agent should keep in mind whether they are objectively real (like the probability of throwing a six in a game of dice), based on empirical data or just subjective estimations. In the following, two rules are exemplified.

Bayes' Rule: This rule recommends to select the action with the maximal expected value. The expected value μ is the sum of the utilities of the different states weighted by the probability of these states. The following table shows an example:

| | $P(z_1) = 0.3$ | $P(z_2) = 0.7$ | |
|-------|----------------|----------------|--|
| | z_1 | z_2 | expected value μ |
| a_1 | 20 | 20 | $\mu = 0.3 \cdot 20 + 0.7 \cdot 20 = 20$ |
| a_2 | 30 | 10 | $\mu = 0.3 \cdot 30 + 0.7 \cdot 10 = 16$ |

In this example a_1 will be selected because of its greater expected utility. Bayes' rule is indifferent to the agent's acceptance of risk which becomes clear if we assume $P(z_1) = P(z_2) = 0.5$ in the above example. In this case the expected values of both alternatives are equal to 20 which means that none of the actions is preferred by Bayes' rule. An agent with preference to risk will prefer a_2 in this situation because there is an opportunity to gain 30 units of utility, while an agent which prefers security will omit this opportunity in selecting a_1 but therefore will be sure of gaining 20 units of utility in any of the cases z_1 or z_2 .

μ - σ -Rule: The μ - σ -rule represents the agents preference towards risk by decrementing or incrementing μ by $a \cdot \sigma$, where σ denotes the standard deviation of the probability distribution and a denotes a parameter encoding the agents preference for risk: A positive value for a increases the preference towards risk, a negative value for a encodes a preference for security. The expected value modified in this way $\Phi = \mu + a \cdot \sigma$ is used for the selection of the action.

2.2.3 Decisions under Uncertainty

These situations are characterised by the fact that no probabilities can be associated with the different environmental states. Most of the rules for this case found in literature differ in their assumptions about the agent's preference towards risk. In the examples, we always show the application of the rule to the following decision matrix:

| | z_1 | z_2 | z_3 |
|-------|-------|-------|-------|
| a_1 | 18 | 35 | 5 |
| a_2 | 20 | 14 | 25 |
| a_3 | 12 | 15 | 30 |

The Maximin-Rule: This rule recommends the action that provides maximal utility for the occurrence of the worst case environmental state. Therefore one has to operate as follows: For each action look for the state that provides minimal utility. From these utilities choose the action with the maximal value. In short this rule selects the action that provides the **maximum** of the **minima** of each row of the matrix, which explains the name of the rule that sometimes is also referred to as Wald-rule. In the example this rule will lead to the selection of a_2 :

| | z_1 | z_2 | z_3 | min of row | |
|-------|-------|-------|-------|------------|---|
| a_1 | 18 | 35 | 5 | 5 | |
| a_2 | 20 | 14 | 25 | 14 | ← |
| a_3 | 12 | 15 | 30 | 12 | |

Because this rule always expects the worst case environmental state it has a strong preference for security and against risk.

The Maximax-Rule: As the name reveals this rule is the inverse of the previous one in that it selects the action with maximal gain presupposing the occurrence of the most favourable environmental state. Therefore it is an optimistic selection rule. It recommends the selection of the action that provides the maximum from the maximal values of each row, which in the example is a_1 :

| | z_1 | z_2 | z_3 | max of row | |
|-------|-------|-------|-------|------------|---|
| a_1 | 18 | 35 | 5 | 35 | ← |
| a_2 | 20 | 14 | 25 | 25 | |
| a_3 | 12 | 15 | 30 | 30 | |

The Hurwicz-Rule (Pessimistic-Optimistic-Rule): This rule combines the maximin-rule with the maximax-rule via the parameter $\lambda \in [0, 1]$ which encodes the agent's optimism respectively pessimism. For $\lambda = 1$, the rule degenerates to the maximax-rule while for $\lambda = 0$ it coincidences with the maximin-rule. As in the previous rules, the maxima and minima of each row are considered.

These values are merged together by adding them where the maxima will be weighted with λ while the minima are weighted with $(1 - \lambda)$. The action with the maximal resulting value is selected. For the example, we assume $\lambda = 0.3$:

| | z_1 | z_2 | z_3 | max of row | min of row | weighted sum |
|-------|-------|-------|-------|------------|------------|--|
| a_1 | 18 | 35 | 5 | 35 | 5 | $35 \cdot 0.3 + 5 \cdot 0.7 = 14$ |
| a_2 | 20 | 14 | 25 | 25 | 14 | $25 \cdot 0.3 + 14 \cdot 0.7 = 17.3$ |
| a_3 | 12 | 15 | 30 | 30 | 12 | $30 \cdot 0.3 + 12 \cdot 0.7 = 17.4$ ← |

In this case the maximal value of 17.4 leads to the selection of a_3 . It should be pointed out that the action finally recommended to select crucially depends on the parameter λ . This is why the last three examples all recommend different actions.

The Savage-Niehans-Rule: The Savage-Niehans-rule is also known as the rule of smallest sorrow. It aims to minimise the maximal possible disadvantage resulting from a false estimation of the environment. This disadvantage can be measured by the difference between the gain expected from an action and the maximal possible gain for a state. It is also known as the opportunity cost because it expresses the loss obtained due to a missed opportunity. After computing the disadvantage for each entry of the decision matrix we are able to predict the maximal possible disadvantage for each action by looking for the maximum of each row. The action that minimises this value is selected by Savage-Niehans-rule.

| | | | | opportunity cost | | | max of row |
|---------------|-------|-------|-------|------------------|-------|-------|------------|
| | z_1 | z_2 | z_3 | z_1 | z_2 | z_3 | |
| a_1 | 18 | 35 | 5 | 2 | 0 | 25 | 25 |
| a_2 | 20 | 14 | 25 | 0 | 21 | 5 | 21 |
| a_3 | 12 | 15 | 30 | 8 | 20 | 0 | 20 ← |
| max of column | 20 | 35 | 30 | | | | |

The table above shows how the Savage-Niehans-rule works in the example. The maximum of each column shows the maximal possible utility for each state z_i . Each entry in the matrix is subtracted from its corresponding maximum which results in the opportunity cost. An entry providing the maximum of its column of course has opportunity cost 0 because no opportunity was missed in this case. The maximal opportunity cost for each action should be minimised, therefore action a_3 providing the minimum 20 of all maximal possible disadvantages will be selected.

The Laplace-Rule: The Laplace-rule assumes that there is basically no reason to expect different probabilities for different environmental states, i.e. it assumes all possible states to be equally distributed. Note that this is a kind of brute force method to push a situation of uncertainty into a situation with risk (as considered in the previous section) and that therefore it should only be applied to situations where this basic assumption of the rules seems to be appropriate. In fact, the Laplace-rule suggests to assign a probability of occurrence to each state computed as the reciprocal value of the number of different states and then to proceed like Bayes' rule would, i.e. selecting the action with maximal expected utility. In the example this would result in selecting a_3 : ($P(z_1) = P(z_2) = P(z_3) = \frac{1}{3}$)

| | z_1 | z_2 | z_3 | expected utility | |
|-------|-------|-------|-------|-----------------------------------|---|
| a_1 | 18 | 35 | 5 | $\frac{18+35+5}{3} = 19.\bar{3}$ | |
| a_2 | 20 | 14 | 25 | $\frac{20+14+25}{3} = 19.\bar{6}$ | ← |
| a_3 | 12 | 15 | 30 | $\frac{12+15+30}{3} = 19$ | |

2.3 Two Agents Decisions

2.3.1 Game Theoretical Concepts

Game theory is a formal approach to the representation and formalisation of decision situations that involve several acting agents. From this point of view it can be considered as a specific branch of decision theory and so the decision theoretic concepts and methods also hold for the decision situations occurring in game theory. Additionally the interdependency of the decisions of the different involved agents comes into play. In the following paragraphs we will consider the fundamental premises and assumptions of the theory and formally define the notion of a game in normal form. We discuss the aspect of cooperation and present some concepts of solutions for games.

Premises and Goals

A major goal of game theory is the analysis of strategic decision situations involving several interacting agents or players. These situations will include cooperative scenarios as well as conflicting interests, and, thus, competition. To attain this goal, first of all, the theory has to provide appropriate formal descriptions of games using mathematically valid objects and concepts. Based upon this, concepts like the solutions of a game can be defined. For different classes of games it should be proved that solutions exist and the qualities of these solutions should be worked out and analysed. Last but not least, the theory has to be applied to

existing problems and it has to be shown that the methods provided to obtain solutions are useful and adequate.

As a presupposition for a formal definition of a game, the premises of the situations considered are given:

- Several players interact by concurrently selecting one action from a set of possible alternatives.
- While each player has complete knowledge about the rules of the game, i.e. knows exactly in advance which combination of selections will lead to which payoff for each player, he does not know any of the actual selections of its opponents at the time taking its own decision. All decisions are taken independently of each others and simultaneously.
- The payoff gained by each player after all decisions have been implemented does not solely depend on his own selection but on the combination of the selections of *all* players.
- Each agent is aware of these premises and takes them into consideration while deciding. And he knows that the other agents are doing this also.

Although it is possible to perform several stages or rounds of selections (as one would usually expect by common sense in a game viewed as alternation of moves and counter-moves), the scenario here is restricted to a single round of decisions, i.e. each agent has to decide only once for one action. The reason for this restriction is that game theory has succeeded in showing that any game involving several consecutive decisions (called a game in extensive form) can be reduced to a game in **normal form** without losing its crucial characteristics. The normal form of a game is a single round representation and will be defined formally in the following section.

Definition of a Game in Normal Form

A game $\Gamma = (N, S, u)$ in normal form (strategic form) is defined by:

- the **set of players** $N = \{1, \dots, n\}$
- the **strategy space** which is a family $S = \{S_i\}_{i \in N}$ of sets of strategies. Thereby $S_i = \{s_{i1}, \dots, s_{im}\}$ denotes the set of all strategies available for player i . The conjoint selection of a strategy of all players $s = (s_1, \dots, s_n) \in \prod_{i \in N} S_i$ is called a combination of strategies⁵.

⁵For convenience we use s_i to denote the strategy from S_i selected by player i . Furthermore, $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in \prod_{j \neq i} S_j$ is used to denote the combination of strategies without the strategy s_i of player i .

- a **payoff function** $u : \prod_{i \in N} S_i \rightarrow \mathbb{R}^n$ that assigns a payoff vector to each combination of strategies $u(s) = (p_1, \dots, p_n)$. Here p_i denotes the payoff for the player i . We will also use u_i as a notation for the concatenation of u and the projection to the i -th component, i.e. $u_i(s) = p_i$.

Constant-Sum Games

A game is a **constant-sum game** if there is a constant c such that for all $s \in \prod_{i \in N} S_i$:

$$\sum_{i=1}^n u_i(s) = c$$

holds. In the special case of $c = 0$ we talk of a **zero-sum game**. In general the structure of a game is not changed when a positive linear transformation is applied to the payoff function, therefore the class of zero-sum games captures every situation that can be represented in a constant-sum game. This class is useful to model conflicting interests because the earnings of one player necessarily correspond to the loss of another. Especially in the case of $n = 2$ this constraint leaves no room for cooperation.

Cooperative vs. Non-Cooperative Games

A game is cooperative if it is possible for the players to make obligatory arrangements about their intended behaviour. This condition includes:

- the possibility of **communication** between the players prior to the decision for an action.
- an institution with the capability to enforce the commitments made, or punish the agents that do not hold their promises. (**exogenous execution**)

The second point is much more important than the first one, because without exogenous execution a communicated promise to select a specific action different from that leading to an agent's maximal payoff is worthless as long as the other agents cannot rely on its compliance. There is also no chance for a revenge if one agent does not stick to its promise because we are considering games in normal form, which means that only one round is played. On the other hand, a promise to select that action that leads to the maximal payoff is superfluous because each agent has the complete knowledge about the payoff structure and thus can deduce this information without such an announcement anyway. Therefore games without the feature of an exogenous execution are defined as **non-cooperative** games, while scenarios incorporating this quality and providing channels for communication are considered to be **cooperative** games. (The case of exogenous execution lacking any devices for communication can be neglected obviously.)

Note that this definition formally assigns the attributes “cooperative” and “non-cooperative” to games. This does not mean that in a non-cooperative game cooperative behaviour of the players is impossible. Cooperative behaviour will occur if each agent’s decision logic enables it, based on an incentive to select the appropriate action even in a game that is non-cooperative by definition. Nevertheless, the prisoner’s dilemma introduced later in Section 2.3.3, provides an example of a non-cooperative game, where cooperation for rational behaving agents cannot occur even if they are able to communicate, though the cooperative behaviour would increase the payoff for both players.

Concepts of Solutions

The goal of defining solutions for a game is to provide rules and concepts for the players to guide their decisions. What is considered as a solution may depend on the point of view: A player will consider the combinations of strategies that maximise its payoff. These, however, may not satisfy another player because the payoff may be small for them. A global point of view may consider the set of strategy combinations that maximise the sum of the payoffs of all players. But a single player not satisfied with its payoff may take decisions that impede these solutions to occur. An appropriate concept of solution therefore must mark those combinations that will occur if rational agents play the game. One method to obtain solutions is to exclude those combinations of strategies that will not occur. The following rules and concepts to characterise solutions have been proposed by game theory.

The Principle of Dominance: This is a rule telling a rational player which strategy not to select. It is based on the observation that it is not rational to select a strategy that leads to less payoff than another one independently of the strategies selected by the other players. In short this rule can be postulated as: “Do not select a strategy that is dominated by another strategy”. Formally, dominance is defined as follows:

Definition: Strategy $s_i^* \in S_i$ **dominates** strategy $s_i \in S_i$ if

1. $u_i(t_1, \dots, s_i^*, \dots, t_n) \geq u_i(t)$ holds for all $t \in \prod_{j \in N} S_j$ with $t = (t_1, \dots, t_n)$ and $t_i = s_i$ and
2. there is at least one $k \in \prod_{j \in N} S_j$ with $k = (k_1, \dots, k_n)$ and $k_i = s_i$ such that $u_i(k_1, \dots, s_i^*, \dots, k_n) > u_i(k)$ holds.

In the case that always $>$ holds we say that s_i^* **strongly dominates** s_i .

The Maximin Principle: This rule recommends to select a strategy as follows: For each strategy $s_i \in S_i$ consider the worst-case selection s_{-i}^{wc} of the other players and the payoff gained for that combination, i.e.

$$p_i^{wc} = \min_{s_{-i} \in \prod_{j \neq i} S_j} u_i(s_1, \dots, s_n) = u_i(s_1^{wc}, \dots, s_{i-1}^{wc}, s_i, s_{i+1}^{wc}, \dots, s_n^{wc})$$

Now select that strategy s_k for which the worst-case payoffs obtain their maximum, which means:

$$p_k^{wc} = \max_{1 \leq j \leq |S_i|} p_j^{wc}$$

This principle determines the maximal payoff that a player can obtain on its own strength independent of what other players do, i.e. even if the others behave maximally unfavourable.

The Nash Equilibrium: To explain the idea of a Nash equilibrium we have to introduce the strategy of the best answer or **Bayes' strategy**. It simply says that, if the decisions of the other players are known in advance, it is easy for an agent to determine its best answer, i.e. the strategy that under this presupposition maximises its payoff: Let s_{-i} be the combination of strategies of all players except i , then i 's best answer is to play a strategy s_i^b that maximises its payoff:

$$u_i(s_1, \dots, s_i^b, \dots, s_n) = \max_{s_i \in S_i} u_i(s_1, \dots, s_i, \dots, s_n)$$

In general there may be several best answers for each s_{-i} . The problem is that usually the player does not know the other player's decisions in advance. However, in certain games it could be the case that there exists a combination of strategies s^* that mutually is a best answer strategy for each player:

$$\forall_{i \in N}, \forall_{s_i \in S_i} : u_i(s^*) \geq u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*)$$

In this case s^* is called a *Nash equilibrium point* of the game. A more intuitive characterisation of an equilibrium point is that no player has an incentive to leave the equilibrium as long as he assumes that no other player is going to leave it. If we assume that there is only one such point in a game then by using the knowledge about the payoff function and by reasoning about the other agents' reasoning, each agent is able to deduce what the other agents (rationally) will select, namely the Nash equilibrium point. In this case it is a good idea to define the equilibrium point as the solution of the game because it is the combination of strategies that the agents as a matter of fact will select. Unfortunately, there are several games with no or with several equilibrium points. Furthermore, though the term "mutually best answer" might imply the assumption that the equilibrium point is something desirable for the players, this is definitively not the case: There are games where the equilibrium point in some sense is the worst selection the agents can make (cf. Section 2.3.3 for an example).

2.3.2 Two-Person Constant-Sum Games

Definition 1 (Two-Person Constant-Sum Game) *A two-person constant-sum game is given by $\Gamma = (N, S = (S_1, S_2), u = (u_1, u_2))$ where $u_1(s_1) + u_2(s_2) = c$ with $s_1 \in S_1$ and $s_2 \in S_2$ for a constant c , i.e.*

- *the gain in utility of one player equals the loss in utility of the other player; the game is strictly competitive and it is not possible to play it cooperatively.*
- *both players have diametrical interests.*

Thereby $u(s_{1i}, s_{2j})$ specifies for each pair of actions s_{1i} and s_{2j} the payoff for player 1 and player 2. In the payoff matrix the sum of both elements of each pair is c . A constant-sum game with $c = 6$ is shown in the following payoff matrix:

| | player 2 | |
|----------|----------|----------|
| player 1 | s_{21} | s_{22} |
| s_{11} | (6, 0) | (2, 4) |
| s_{12} | (5, 1) | (3, 3) |

Two-person constant-sum games model decision situations with a total clash in the interests of the two agents. Each constant-sum game can be transformed into a zero-sum game (i.e. $c = 0$) with the same properties.

To obtain this two-person zero-sum game, $u_2 + u_1 = c$ is transformed into $(u_2 - \frac{c}{2}) + (u_1 - \frac{c}{2}) = 0$ to see that we just have to subtract $\frac{c}{2}$ from each value to obtain the payoffs of the equivalent zero-sum game. For the example given above this results in the payoff matrix:

| | player 2 | |
|----------|----------|----------|
| player 1 | s_{21} | s_{22} |
| s_{11} | (3, -3) | (-1, 1) |
| s_{12} | (2, -2) | (0, 0) |

If there is only a finite number of strategies for the players to choose from, the game is completely described by the payoff matrix $U = (u_{ij})$. Therefore the analysis of a game can be done on the basis of this matrix.

Because the profit of one player in a constant-sum game is the loss of the other player, it is assumed to be rational for a player, to maximise his own payoff and therefore to minimise the payoff of his opponent. To reach this goal in a constant-sum game, it is sufficient to use the dominance and the maximin principle as a decision criterion. In the game shown above s_{22} is a dominating strategy for player 2 because he gets more utility when using this strategy than by using strategy s_{21} regardless which strategy is chosen by player 1 ($-3 < 1$ and $-2 < 0$). For player 1 the dominance principle is not applicable ($3 > 2$ but $-1 < 0$). He can

use the minimax principle to choose strategy s_{12} because this will guarantee him a payoff of at least 0 while strategy s_{11} only guarantees a worst-case payoff of -1. The decision criterion *best reply* is also applicable: Strategy s_{12} is the best reply of player 1 (column) to the strategy s_{22} of player 2 (row). That player 2 will choose strategy s_{22} can be deduced by player 1 from the knowledge about the payoff matrix: Assuming that player 2 acts rationally he must select s_{22} because of the dominance.

In a two-person constant-sum game the power of a player — the payoff he is able to guarantee himself — is a good criterion to find out what a player can gain from a game. In the example above strategy s_{12} guarantees a payoff of 0 for player 1, whereas strategy s_{11} only guarantees him a payoff of -1.

2.3.3 Two-Person Non-Constant-Sum Games

The two-person non-constant-sum game is also called *general two-person game* and differs fundamentally from a constant-sum game:

- The profit of one player does not correspond to the loss of his opponent, i.e. the game is not strictly competitive; the sums of the utility-pairs from different matrix positions differ.
- There is not a complete clash in the interests of the two players. Because of this, both of them are able to prefer specific actions conjointly which enables the possibility of cooperative solutions.

In the sequel we present one of the most famous non-cooperative two-person non-constant-sum games: the prisoner's dilemma. In this example we will show that strategies which are based on individual maximisation of utility (dominance, minimax, and Nash equilibrium) will not always result in efficient solutions. These strategies will only provide suboptimal solutions. Optimal solutions can only be obtained by committing to a kind of collective rational behaviour.

Prisoner's Dilemma

This non-cooperative game characterises the situations of two prisoners and was first described by Luce and Raiffa [Luce & Raiffa 57].

Two persons which are suspected to have committed a crime are trapped in solitary confinement. The public prosecutor is convinced that both of them are guilty but there is a lack of evidence to prove them to be guilty. Both of the prisoners have two possibilities to choose: to confess or not to confess the crime. If both of the prisoners do not confess the crime the public prosecutor is only able to sentence them for a minor crime and they will get a mild punishment. If both of them confess the crime, both of them are accused of the crime but they will not get the hardest punishment possible. If only one of the prisoners confesses the crime,

this prisoner will become the chief witness and will be set free shortly after the court procedure is finished. In this case the other prisoner will in this case get the hardest punishment possible. How should the prisoners behave in this situation?

Formal Description of the Game in a Matrix: Both prisoners (players) i , ($i \in \{1, 2\}$), have two strategies s_{ij} ($j \in \{1, 2\}$): not to confess (s_{i1}) or to confess (s_{i2}). Depending on which of the strategies both of them will choose this results in one of $2 * 2$ possible combinations of strategies $s = (s_{1j}, s_{2k})$ which determine the result $e(s)$, i.e. the number of years each of the prisoners will have to stay in prison.

The following matrix displays the situation in the prisoner's dilemma:

| | player 2 | not confess | confess |
|-------------|----------|--|--|
| player 1 | | s_{21} | s_{22} |
| not confess | | 1 year for player 1 1 year for player 2 | 10 years for player 1 3 months for player 2 |
| s_{11} | | | |
| confess | | 3 months for player 1 10 years for player 2 | 7 years for player 1 7 years for player 2 |
| s_{12} | | | |

The assessment of the results is now determined by the preferences of the players. Each player i assigns a utility $u_i(e(s))$ (payoff) to each of the results $e(s)$. In this example we can assume that the prisoners prefer a short punishment. Therefore, using an ordinal *utility function* a large utility will be assigned to a short time in prison.

Derivation of a Payoff Matrix from the Result Matrix: Each combination of strategies s leads to a specific result $e(s)$. Therefore, each player i is able to assign utility to s by using his utility function $u_i(s)$. In doing so the players assign utility $u(s) = (u_1(s), u_2(s))$ to each combination of strategies s and by this the result matrix is transformed into a payoff matrix. The payoff matrix is the basis on which the players decide which strategy they should chose.

Assuming that the prisoners have utility functions which assign more utility to shorter periods in prison than to longer ones, we get the payoff matrix shown below for the Prisoner's Dilemma. The general prisoner's dilemma is characterised by the following structure:

| | player 2 | s_{21} | s_{22} |
|-------------|----------|--------------------|--------------------|
| player 1 | | not confess | confess |
| s_{11} | | (β, β) | (δ, α) |
| not confess | | | |
| s_{12} | | (α, δ) | (γ, γ) |
| confess | | | |

where the variables meet the constraints:

$$\alpha > \beta > \gamma > \delta.$$

A concrete instance of such a game is given by

| | | | |
|-------------|----------|-------------|----------|
| | player 2 | s_{21} | s_{22} |
| player 1 | | not confess | confess |
| | player 1 | s_{21} | s_{22} |
| | | not confess | confess |
| s_{11} | | (6, 6) | (1, 8) |
| not confess | | | |
| s_{12} | | (8, 1) | (2, 2) |
| confess | | | |

Derivation of the Solution from the Analysis of the Payoff Matrix:

If player 2 chooses strategy s_{21} , it is better for player 1 to confess the crime, i.e. chose strategy s_{12} . But also if player 2 confesses the crime, s_{12} is the best strategy for player 1. Analogously, the same is true from the perspective of player 2. Hence, to “confess” is for both players a *dominant strategy* because for both players $i \in \{1, 2\}$ with $j = 3 - i : u_i(s_{i2}, s_{jk}) > u_i(s_{i1}, s_{jk})$ holds no matter which strategy k is chosen by the opponent j . Because of this it is for both players—independent from the action taken by the opponent—individually rational to confess the crime. The combination of strategies (s_{12}, s_{22}) is a Nash equilibrium point because none of the players have an incentive to switch to another strategy. None of the other strategy combinations have this property of being in an equilibrium. Nevertheless, this combination is dominated by the combination (s_{11}, s_{21}) because with this combination both players obtain more utility than in the equilibrium. Even though both players should be interested in changing their strategies and by doing so leaving the equilibrium, the dominance rule will prevent them from choosing another strategy than to confess.

Analysis and Discussion: The equilibrium point of the prisoner’s dilemma is inefficient because the players could get more utility (β, β) out of the situation if they choose the cooperative combination of strategies (s_{11}, s_{21}) instead of choosing the equilibrium combination of strategies (s_{12}, s_{22}) . The problem is that the prisoner’s dilemma is a non-cooperative game. As discussed in section 2.3.1 a non-cooperative game is characterised by the lack of communication between the agents to arrange mutual commitments and—much more important—the lack of an exogenous institution to enforce the execution of the commitments. In the prisoner’s dilemma the possibility of communication would not help the prisoners. Just imagine that both prisoners have communicated and agreed in not

confessing whatever will be. Now they are in the dock and the judge asks the first one of them: “Do you plead guilty?” Without an exogenous institution that enforces the commitment he has no reason to assume that his buddy will stick to the agreement. He himself has promised not to confess, but after answering “No” to the judge his buddy’s decision is simply that between a big or a small punishment. So what should he answer? Note that the assumption of simultaneous independent decisions as it is presupposed in game theory in this case does not matter, at least not for the prisoner that is asked first. The characteristics of his problem is not different from the case of two independent tribunals for the prisoners in separate rooms.

Let us now change the situation and assume the existence of an exogenous institution able to enforce a commitment like in our example to force the prisoners not to confess. This can be modelled as a change of the payoff matrix: The enforcement will rely on the threat of some kind of punishment⁶. The appraisal of that punishment changes the agent’s utility. Then the utility of a strategy combination involving a decision contrary to the agreement no longer solely depends on the time staying in prison but also incorporates the punishment of the enforcing institution. In our example the external punishment should result in a decrease of the utility of confessing so that γ becomes smaller than δ . Then the game is no longer a prisoner’s dilemma.

Results: The Prisoner’s Dilemma shows that

- decision rules which rely on the concept of individual rationality may lead to inefficient solutions. In contrast to constant-sum games, non-constant-sum games have properties which make the usefulness of these decision rules doubtful.
- an agreement, such that a player is able to gain utility in violating the agreement, is meaningless as long as there is no exogenous institution with the capability to enforce the execution of the agreement.

Another Example: The following payoff matrix shows a game in which none of the players’ strategies dominates the other one like it is the case in the prisoner’s dilemma. Furthermore, it has two equilibrium points in (s_{12}, s_{21}) and (s_{11}, s_{22}) , because none of the players has an incentive to leave one of these points.

| | | Player 2 | |
|----------|----------|----------|------------|
| | | s_{21} | s_{22} |
| Player 1 | s_{11} | (1, 1) | (-9, 9) |
| | s_{12} | (9, -9) | (-50, -50) |

⁶E.g. the promise of the Mafia to kill every chief witness.

If both players initially use the minimax-rule as their decision procedure, this results in the combination of strategies (s_{11}, s_{21}) . However, both players have an incentive to leave this point because each could receive the utility of 9 rather than 1 by switching to the other strategy. Note that only a single round of the game is played; if we talk of “switching to the other strategy” we really mean changing the intention from selecting one strategy to selecting the other strategy. Therefore, if player i thinks that the other one will definitively use the minimax-rule for his decision he could increase his expected utility by selecting strategy s_{i2} . Assuming that the other player’s reasoning works the same way this results in the combination (s_{12}, s_{22}) which is rather bad for both players. Therefore using the minimax-rule seems to be a good advice. But now the reasoning starts to loop: If the other player is convinced that the minimax-rule indeed is a good advice, then it is rational to switch etc. We see that rational decisions are not easy to take.

Battle of the Sexes

Another prominent example for two-person non-constant-sum games are games of the type *Battle of the Sexes* which owe their name to the following description of the situation in the game:

A man and a woman fall in love and want to spend their spare-time together. He likes soccer games and she enjoys visiting the theatre but no one of them really enjoys its favourite spare-time activity without the other. They both want to be together. Of course he prefers to take her to a soccer game much more than to accompany her to the theatre. Her preferences are just the other way around.

Games of this type have a payoff matrix of the following kind:

| | woman | |
|---------------------|--------------------|---------------------|
| man | s_{21} soccer | s_{22} theatre |
| s_{11} soccer | (α, β) | (γ, γ) |
| s_{12} theatre | (γ, γ) | (β, α) |

where $\alpha > \beta \gg \gamma$.

It is easy to see that the man prefers the combination of strategies (s_{11}, s_{21}) whereas the woman prefers (s_{12}, s_{22}) . Both of these combinations of strategies are in an equilibrium and therefore the problem to be solved by the players is to agree on one of these equilibrium points.

This setting can be found frequently in decision situations in business. We have one instance of such a situation if two companies try to set up a close cooperation

and for this reason have compiled two cooperation contracts, out of which the first one is favourable for the first company and the second one is favourable for the second company. If each company insists on its own contract, they are not able to find an agreement.

2.3.4 Iterated Games

Up to this point we put the focus of the analysis on *one shot games*, i.e. strategic aspects which arise in repeated games have been omitted. In contrast to one shot games, iterated games consist of a series of *base game*.

Iterated games can be used to model evolutionary aspects (see e.g. [Axelrod 84]). The concept of evolutionary stable strategies (ESS) (cf. [Maynard-Smith 82]) is useful to explain the emergence of agent behaviour as for example in the case of animals the institution of a dominance hierarchy, the territorial and mating behaviour or the behaviour in animal fights (cf. [Dawkins 76]).

Of particular interest are two-person non-constant-sum games. In this case, even if they are non-cooperative, intelligent players will in a series of games agree in silence on a combination of strategies that leads to a result which is for both players beneficial. By choosing specific strategies a player is able to tell the other player which of the combinations of strategies are preferred. The *iterated prisoner's dilemma* is a good example for this.

In the iterated prisoner's dilemma, it is clear that both players gain a substantial strategic option: The players are able to take *revenge* for the actions of the opponent which can be used to enforce cooperative behaviour and suppress non-cooperative behaviour. Axelrod [Axelrod 84] organised a tournament for computer programs, in which computer programs using different strategies played against each other an iterated prisoner's dilemma game. It turned out, that a strategy with the self-explaining name *tit-for-tat*—i.e. start with cooperation in the first game and then behave in the game at present in the same way as your opponent did in the previous game—is the most efficient one. However, the analysis of this situation shows that if there is only a sequence of finitely many games, then there is a large incentive for both players to cheat their opponent in the last game because the opponent will not be able to take a revenge for the behaviour in the last game. Even worse: If the number of iterations is known in advance, cheating in the last game will be deduced by both players and cheating in the game next to the last becomes a matter of consideration. By inductive arguments the cooperative behaviour of the whole series of games may be spoiled.

In iterated games it becomes evident that even in non-cooperative games cooperation can be enforced indirectly via the selected strategies in the series of the game. An exogenous enforcement by the punishment of non-cooperative behaviour via a third party is no longer necessary because each of the players is able to take its own revenge.

2.4 Multiple Agents Decisions

In n -person games besides conflict, cooperation automatically plays a significant role. Because there can be only two preference systems that are completely conflicting. Therefore, if n agents with n preference systems interact with each other, some subset of the set of players \mathcal{N} will have interests in common, which possibly can be supported by cooperation. Therefore n -person games are often called *coalition games*. If an agent has the option to join several coalitions he will always join the one that offers him the highest increase in utility.

2.4.1 Coalitions in Cooperative N -Person Games

In 1944, John von Neumann and Oskar Morgenstern defined in their book “Theory of Games and Economic Behaviour” [von Neumann & Morgenstern 44] the n -person game for the first time. Their approach is based on the assumption that players are able to communicate with each other and are therefore able to make agreements and to form coalitions. That is why their approach is excellently suitable for analysing situations in which it is cooperated within groups.

In a n -person game the players have the option to form coalitions, such that the utility the coalition receives possibly exceeds the sum of the utilities all the members of the coalition could obtain individually based on their own strength.

This is due to that

- the coalition has more resources than each of the individual players and the coordinated application of these resources may lead to a more efficient usage of resources than an application by individual single players guided by their individual rationality.
- the coalition may be able to provide incentives to the members to play strategies that are not individually rational but substantially increase the coalition’s overall profit. This is possible because the distribution of the coalition’s profit to its members is decoupled from the obtainment of the profit.

It is obvious that cooperative players will be better off than non-cooperative players because all results which are possible in the non-cooperative case are still possible in the cooperative case but not vice versa.

Big efforts in theoretical investigations have been made to clarify the question how coalitions should be formed among rational players and how a fair distribution of the profit to the members of the coalition can be computed [Shapley 53b]. The main problem of cooperative game theory is to find a distribution function for the profit of the coalition which reflects the different positions of the players (e.g. their power or strategic importance) in the negotiation process when they build

the coalition.⁷ Closely related with the solution of the distribution problem is the question about the rating of the potential power of coalition members.⁸

2.4.2 Fundamental Concepts for the Analysis of N -Person Games

John von Neumann and Oskar Morgenstern [von Neumann & Morgenstern 44] introduced two basic concepts to be able to analyse the formation of coalitions: The definition of a game in characteristic form using a *characteristic function* to define the payoff structure of the game and the characterisation of *imputations* which are those distributions of the payoff to the members of a coalition that come into question to be solutions.

The Characteristic Function

The formation of coalitions is such an important concept in n -person games that it is straightforward to define the payoff structure of a game according to the profit that the different possible coalitions are able to obtain. Therefore a characteristic function is used to define the payoff instead of the payoff function known from the game in normal form:

Definition 2 (Characteristic Form of a Game) *A n -person game $\Gamma = (\mathcal{N}, v)$ in characteristic form is defined by the set of players $\mathcal{N} = \{1, \dots, n\}$ and the characteristic function $v : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ which assigns a real value $v(\mathcal{K})$ to each $\mathcal{K} \subseteq \mathcal{N}$.*

Thereby $v(\mathcal{K})$ specifies the guaranteed utility the coalition \mathcal{K} is able to obtain, no matter which strategies are chosen by the other coalitions. Thus, the characteristic function is based on the maximin principle and a transformation based on this principle into the 2-person normal form is possible, where one player represents \mathcal{K} while the other player represents the complementary coalition $\mathcal{N} - \mathcal{K}$. The characteristic function has to fulfil the condition $v(\emptyset) = 0$ and the condition of superadditivity, which will be explained in the following:

An important notion for the analysis of game outcomes with respect to social welfare (sum of the agents' payoffs) is the notion of *superadditivity*.

⁷Normally it is assumed that the utility is paid to a player in units of an infinitely divisible, transferable good (e.g. money) so that the profit of the coalition can be divided on its members in any desirable way (this is also called: transferable utility).

⁸Voting games are a typical starting point for n -person game theory. An example for this is the problem of the formation of the government if none of the parties in the parliament has an absolute majority and therefore a governmental coalition has to be formed. In this case a party with 3 seats may have the same power as a party with 20 seats. Important is only that the party joining the coalition will guarantee the absolute majority.

Two disjoint coalitions \mathcal{K}_1 and \mathcal{K}_2 can independently from each other guarantee themselves the utilities $v(\mathcal{K}_1)$ and $v(\mathcal{K}_2)$. If they cooperate by forming the bigger coalition $\mathcal{K}_1 \cup \mathcal{K}_2$, they are almost ⁹ always able to get at least the utility $v(\mathcal{K}_1) + v(\mathcal{K}_2)$, because they can simply simulate the behaviour of the coalitions \mathcal{K}_1 and \mathcal{K}_2 . In some cases through cooperation they may even be able to find better strategies and to get even more utility.

Definition 3 (Superadditivity) *A game is called superadditive if the characteristic function v satisfies the following inequality $\forall \mathcal{K}_1, \mathcal{K}_2 \in 2^{\mathcal{N}}$ with $\mathcal{K}_1 \cap \mathcal{K}_2 = \emptyset$:*

$$v(\mathcal{K}_1 \cup \mathcal{K}_2) \geq v(\mathcal{K}_1) + v(\mathcal{K}_2)$$

This means, in superadditive games the utility several subsets of agents can realise jointly by forming a coalition does never fall short of the sum of the utilities that those agents can realise jointly by forming a coalition.

Therefore in all superadditive games the agents can realise the highest social welfare by forming the *grand coalition* $\mathcal{K}_g = \mathcal{N}$ consisting of all agents.

Some non-superadditive games are *subadditive*:

Definition 4 (Subadditivity) *A game is called subadditive if the characteristic function v satisfies the following inequality $\forall \mathcal{K}_1, \mathcal{K}_2 \in 2^{\mathcal{N}}$ with $\mathcal{K}_1 \cap \mathcal{K}_2 = \emptyset$:*

$$v(\mathcal{K}_1 \cup \mathcal{K}_2) < v(\mathcal{K}_1) + v(\mathcal{K}_2)$$

In subadditive games, agents maximise social welfare by operating alone.

There exist games that are neither superadditive nor subadditive. In those games the characteristic function v fulfils gives us superadditivity for some coalitions and subadditivity for others. In such cases there is no general strategies for maximising social welfare.

Essential and Inessential Games

The only incentive for forming coalitions is to gain more utility than can be obtained without forming the coalition. Therefore, only games in which there are at least two disjoint subsets \mathcal{U} and \mathcal{W} of \mathcal{N} such that

$$v(\mathcal{U} \cup \mathcal{W}) > v(\mathcal{U}) + v(\mathcal{W})$$

holds are interesting. These games are therefore called *essential* games, because at least the players from \mathcal{U} and \mathcal{W} are able to increase their utility by forming a

⁹Except for the case where the coalition formation process involves costs—e.g. coordination costs or payments for securing trust among the coalition members—that exceed the value additionally realisable by coalition formation.

coalition. All games without this condition lack any interesting aspects for cooperative game theory and are called *inessential*.¹⁰ It can be shown that inessential games can be characterised by a rather simple condition:

$$\Gamma = (\mathcal{N}, v) \text{ is inessential} \quad \Leftrightarrow \quad \sum_{i \in \mathcal{N}} v(\{i\}) = v(\mathcal{N})$$

It is easy to see that no constellation of coalitions can gain a higher utility than the maximal coalition \mathcal{N} . Therefore, it is the most rational behaviour for all players to cooperatively form the coalition \mathcal{N} . But the problem that still remains is to distribute the obtained profit $v(\mathcal{N})$ to the players $1, \dots, n$. In the next section we see that a solution for this problem can only be provided by distributions of the profit that satisfy some very intuitive conditions.

2.5 Solution Concepts for N -Person Games

In this section, we present two solution concepts for n -person-games, the *kernel* and the *Shapley value*. Both concepts are applicable for games with a defined characteristic function and propose as solutions payoff vectors that satisfy certain desired properties or specified axioms.

2.5.1 Von Neumann and Morgenstern's Solution: the Kernel

The *kernel* deals with fairly distributing the utility $v(\mathcal{K}) - \sum_{i \in \mathcal{K}} v(i)$ that could be additionally realised by coalition formation.

The kernel has some desired characteristics in that it gives symmetric players equal payoff and more desirable players more payoff than to less desirable ones. Therefore it can be understood as a concept of fairness.

Imputations and Kernel of a Game

Even the simple n -person games are fairly complex and it is usually not possible to derive the best strategies or to predict the final result of the game using a general theory. The problem of distributing the utility $v(\mathcal{N})$ to the players can be reduced by excluding some possible distributions which are rather unacceptable for the players because they contradict some basic assumptions that should be obeyed by all rational solutions. Two of these basic assumptions are the principles of *individual rationality* and *pareto optimality*.

¹⁰Each two-person-constant-sum game is inessential because the minimal payoff a player can guarantee himself cannot be increased by cooperation. An example for an essential game is the prisoner's dilemma because both player can benefit from forming a coalition. Zero-sum games with more than two players can be essential, too.

Informally the individual rationality guarantees for each player an incentive to participate in the coalition implemented as gain in utility compared to the utility he can achieve by his own strength. The pareto optimality guarantees that the utility $v(\mathcal{N})$ gained by the coalition \mathcal{N} is expended entirely without saving any rest. More formally we can define:

Definition 5 (Individual Rationality) *A payoff vector $u = (u_1, u_2, \dots, u_n) \in \mathbb{R}^n$ is individually rational if and only if each of the players gets at least as much utility as he could get on his own:*

$$\forall i \in \mathcal{N} : u_i \geq v(\{i\})$$

Definition 6 (Pareto Optimality) *A payoff vector $u = (u_1, u_2, \dots, u_n) \in \mathbb{R}^n$ is pareto-optimal if and only if there exists no other payoff vector u^* which increases the payoff for a player without decreasing the payoff of any other player:*

$$\nexists u^* = (u_1^*, u_2^*, \dots, u_n^*) : \quad \forall i \in \mathcal{N} : u_i^* \geq u_i \quad \wedge \quad \exists j \in \mathcal{N} : u_j^* > u_j$$

Because $\sum_{i \in \mathcal{N}} u_i \leq v(\mathcal{N})$ must hold for every payoff vector u , the condition of pareto optimality can also be expressed as:

$$\sum_{i \in \mathcal{N}} u_i = v(\mathcal{N})$$

All payoff vectors not obeying the principles of individual rationality and pareto optimality should not be considered as solutions for the problem of distributing $v(\mathcal{N})$ to the players. Von Neumann and Morgenstern defined the set of allowed payoff vectors and called them imputations:

Definition 7 (Imputations) *The set of imputations for an n -person game $\Gamma = (\mathcal{N}, v)$ is defined as:*

$$I(v) := \{u \in \mathbb{R}^n \mid u_i \geq v(\{i\}) \quad \wedge \quad \sum_{i \in \mathcal{N}} u_i = v(\mathcal{N})\}$$

To summarise this, an imputation is a distribution of the overall payoff $v(\mathcal{N})$ gained by the maximal coalition \mathcal{N} to the players which satisfies the criteria of individual rationality and pareto optimality, i.e.

1. Each individual player receives at least as much payoff as he is able to obtain by his own strength.
2. All the profit gained by the coalition is entirely distributed to the players.

Example: A three-person game with players A, B and C is defined by the characteristic function v with:

$$\begin{aligned} v(\{A\}) &= v(\{B\}) = v(\{C\}) = 0 \\ v(\{A, B\}) &= v(\{A, C\}) = 30 \\ v(\{B, C\}) &= 0 \\ v(\{A, B, C\}) &= 30 \end{aligned}$$

We see that no individual player can gain something. All coalitions of several players that have A as a member are able to gain a profit of 30 units of utility.

Rational players will only accept imputations as payoff vectors. But note that not every imputation provides a reasonable solution. In this three-person game $(0,30,0)$ and $(0,5,25)$ are imputations but an objective spectator would be astonished if one of them would be the result of the game. This is due to the fact that it is not evident why in both cases player A does not receive any profit although a coalition without A does not get any utility.

The example shows that though each reasonable solution must be an imputation, not all of the imputations provide a reasonable solution. This should not be too surprising because the principles of individual rationality and pareto optimality only compare the behaviour of individuals (i.e. one-person coalitions) to that of the maximal coalition \mathcal{N} . What is needed is a principle that extends individual rationality to rationality of coalitions with more than one participant. This is provided by the definition of the kernel of a game:

The Kernel of a Game

In natural extension of the principle of individual rationality, it is rational for each subset $\mathcal{K} \subset \mathcal{N}$ of players to demand at least as much profit from the maximal coalition \mathcal{N} as \mathcal{K} could gain as an independent coalition by its own strength. Only imputations that respect this condition will be accepted by all possible coalitions that potentially may be formed. Therefore we define:

Definition 8 (Kernel of a Game) *The kernel of $\Gamma = (\mathcal{N}, v)$ is the set of all imputations $u \in I(v)$ for which the condition of collective rationality holds:*

$$\forall \mathcal{K} \subset \mathcal{N} : \sum_{i \in \mathcal{K}} u_i \geq v(\mathcal{K})$$

If a payoff vector from the kernel of a game Γ is used to distribute $v(\mathcal{N})$ to the players, the game Γ is called stable because neither individual players nor subsets of players have an incentive to form other coalitions than \mathcal{N} . The stability of coalition \mathcal{N} is preserved. Therefore, the kernel of a game can be viewed as the specification of the solution because it contains only payoff vectors which fulfil the

criteria of individual rationality and collective rationality of all possible coalitions of players.

Unfortunately, the kernel is empty for many games.¹¹ For this reason it is obvious that the kernel of a game cannot be viewed as a general solution principle for cooperative games.

The solution of von Neumann and Morgenstern

Payoff vectors from the kernel of a game provide solutions with satisfying properties. However, a large number of games is not solvable with this approach because their kernel is empty. Von Neumann and Morgenstern wanted to be able to provide a reasonable notion of solutions for all games and therefore introduced a new concept for the solution of cooperative games which is based on the definition of a dominance relation on the set of payoff vectors.

Definition 9 (Dominance Relation) *In a game $\Gamma = (\mathcal{N}, v)$ an imputation $u = (u_1, \dots, u_n)$ dominates an imputation $w = (w_1, \dots, w_n)$ if there exists a nonempty coalition $\mathcal{K} \subseteq \mathcal{N}$ such that the following conditions hold:*

1. $\forall i \in \mathcal{K} : u_i \geq w_i \quad \wedge \quad \exists j \in \mathcal{K} : u_j > w_j$
2. $\sum_{i \in \mathcal{K}} u_i \leq v(\mathcal{K})$

The first condition specifies that there is at least one player in \mathcal{K} who is better off with the imputation u than with w while for all other members of \mathcal{K} the imputation u is not worse than w . The second condition guarantees that the players in \mathcal{K} are really able to satisfy the payoffs specified by u_i with the profit $v(\mathcal{K})$ the coalition is able to get.

It is obvious that imputations in the kernel of a game cannot be dominated by another imputation. However, this does not imply that an imputation in the kernel of a game dominates each imputation which is not an element of the kernel.¹²

Von Neumann and Morgenstern now defined \mathcal{L} to be a solution of a cooperative game if \mathcal{L} satisfies the following two stability conditions:

1. No imputation of \mathcal{L} dominates another imputation in the set \mathcal{L} .
2. Each imputation which is not an element of \mathcal{L} is dominated by at least one imputation of \mathcal{L} .

Von Neumann and Morgenstern compare the set \mathcal{L} with an *accepted norm for behaviour* or with a *social order*. The first condition prevents internal contradictions and the second one takes care that behaviours which are not acceptable — an imputation which is not in \mathcal{L} , viz — are eliminated.

¹¹E.g. this is the case for constant-sum games (i.e. $v(\mathcal{K}) + v(\mathcal{N} \setminus \mathcal{K})$ is constant for all \mathcal{K}).

¹²It is also possible to define the kernel of a game as the set of all imputations which are not dominated by any other imputation.

2.5.2 Solutions for Games without Kernel

For games without kernel, several concepts of solutions have been proposed, among others the *Shapley Value* [Shapley 53b]. None of these solutions are stable because if the kernel of a game is empty then there is at least one subset of players which is not satisfied with the presented distribution because they would all together get more utility if they form a coalition. Nevertheless, these solutions may specify an equilibrium in the negotiation process which will be outlined further in the sequel.

The Shapley Value

The Shapley value is a solution concept [Raiffa 82] that bases on an axiomatic approach [Rosenstein & Zlotkin 94c, Rosenstein & Zlotkin 94b]. It is not a model of stability in a bargaining process but more a basis of fairness for joint win distribution.

The Shapley Value specifies a formal solution for cooperative n -person games which is based on the concept of the characteristic function.¹³ Central for Shapley's approach is the specification of an *index for the a priori power of a player* based on the increment of utility that a coalition can gain by including this player as a new member.

Shapley justifies his method by specifying requirements which should be fulfilled by solutions $\sigma \in \mathbb{R}^n$ of a game $\Gamma = (\mathcal{N}, v)$ and proves that the Shapley Value (as defined below) is the only payoff vector which satisfies these requirements¹⁴:

1. $\sigma \in I(v)$
2. if $\forall \mathcal{K} \subset \mathcal{N}, i \notin \mathcal{K} : v(\mathcal{K} \cup \{i\}) = v(\mathcal{K}) + v(\{i\})$ holds, then $\sigma_i = v(\{i\})$
3. if ψ is a solution of (\mathcal{N}, w) then $\psi + \sigma$ is a solution of $(\mathcal{N}, w + v)$
4. if $\pi : \mathcal{N} \rightarrow \mathcal{N}$ is a permutation so that $\forall \mathcal{K} \subset \mathcal{N} : v(\mathcal{K}) = v(\pi(\mathcal{K}))$ holds, then $\forall i \in \mathcal{N} : \sigma_i = \sigma_{\pi(i)}$

These requirements specify an unique payoff vector whose elements represent the agents' Shapley values.

The basic idea of the Shapley Value is based on the assumption that the great coalition \mathcal{N} of the n -person game is formed step by step beginning with a coalition which consists of one player and proceeding by adding one more player to this coalition in every step. For n players there are $n!$ different possible sequences in which the players can join a great coalition. Therefore, all sequences are

¹³Shapley adopts the assumption of super additivity and the transferability of utility from the approach of von Neumann and Morgenstern.

¹⁴Therefore, it is also possible to define the Shapley Value based on these requirements.

enumerated and the increment of utility the coalition gains by the admission of a new player is ascribed to that player. A player's payoff is then determined by taking the average of all these assignments over all $n!$ possible sequences.

Definition 10 (Shapley Value) *In the n -person game $\Gamma = (\mathcal{N}, v)$ the Shapley Value σ_i for each player i is defined by:*

$$\sigma_i := \sum_{\mathcal{K} \subseteq \mathcal{N} \setminus \{i\}} \frac{|\mathcal{K}|!(|\mathcal{N}| - |\mathcal{K}| - 1)!}{|\mathcal{N}|!} \cdot (v(\mathcal{K} \cup \{i\}) - v(\mathcal{K}))$$

In this formula the factor $(v(\mathcal{K} \cup \{i\}) - v(\mathcal{K}))$ expresses the increment of utility the coalition \mathcal{K} gains by the admission of player i . The left factor is explained by the fact that there are $|\mathcal{K}|!$ possible different sequences for the formation of coalition \mathcal{K} and furthermore $(|\mathcal{N}| - |\mathcal{K}| - 1)!$ possible different sequences to complete the generation of \mathcal{N} from $\mathcal{K} \cup \{i\}$. The denominator is the number of all $n!$ possible different sequences that are considered.

The Shapley value apportions the profit of the coalition \mathcal{N} to the n players. Players which on average contribute more to the utility of the coalition get more utility out of the coalition's profit; players which do not contribute to the profit of any coalition do not get any utility.

In the calculation of the Shapley value all possible ways of coalition accumulation are taken into consideration. An agent's contribution to its coalition depends on which other agents have joined before it and which ones will join after it. An agent's Shapley value in a coalition equals the agent's contribution to this coalition averaged over all possible joining orders.

The Shapley value does exist for all coalitions and is unique. Additionally, it is guaranteed to satisfy individual rationality in superadditive games. It satisfies also the property of symmetry among agents and assigns more desirable agents a higher utility than less desirable ones.

Example: Let us illustrate the computation of the Shapley value by reconsidering the example introduced above:

$$\begin{aligned} v(\{A\}) &= v(\{B\}) = v(\{C\}) = 0 \\ v(\{A, B\}) &= v(\{A, C\}) = 30 \\ v(\{B, C\}) &= 0 \\ v(\{A, B, C\}) &= 30 \end{aligned}$$

Obviously A has the power of veto because without him none of the others can enforce the decision. Nonetheless, A is not a dictator because he cannot enforce the decision on his own. The computation of the Shapley value is shown in the following table:

| | A | B | C | C | B | A |
|----|----|----|----|----|----|----|
| | B | C | A | B | A | C |
| | C | A | B | A | C | B |
| A: | 0 | 30 | 30 | 30 | 30 | 0 |
| B: | 30 | 0 | 0 | 0 | 0 | 0 |
| C: | 0 | 0 | 0 | 0 | 0 | 30 |

The rows show the six possible sequences in which the great coalition of 3 players can be formed. The contributions of the players are listed below. In the sequence (A, B, C) , e.g. B contributes 30 to the coalition $\{A, B\}$ because $v(\{A\}) = 0$ and $v(\{A, B\}) = 30$. On the other hand, C contributes nothing to the coalition $\{A, B, C\}$ in this sequence, because $v(\{A, B\}) = 30$ and $v(\{A, B, C\}) = 30$. Therefore, the contributions of A, B and C are given by $(0, 30, 0)$. One can see that the Shapley value for A is 20 and those for B and C are 5. A 's power of veto enables him to get $2/3$ of the money.

Final Remarks

Kernel and Shapley value are only applicable in games with a defined characteristic function but not in games where the value of a coalition may depend on the actions of non-members. For the latter games, that can be modelled as normal form games, the Nash equilibrium is applicable because it is a solution concept in the space of strategies whereas kernel and Shapley value are solution concepts in the space of payoff configurations.

One problem of the Nash equilibrium is that even in a Nash equilibrium subgroups of agents could be motivated to deviate in a coordinated manner. The concept of *strong Nash equilibrium* [Aumann 59] requires that there is no subgroup of players that can increase the utilities of all its members by jointly deviating from their Nash equilibrium strategies when all non-members keep their strategies fixed. Therefore it guarantees more stability but is often too strong because in many games no strong Nash equilibria do exist (nonexistence problem).

Chapter 3

Game Theory in Multiagent Systems

In this chapter we present an approach to utilise the results of game theory in Multiagent Systems. The negotiation between agents can be regarded as a game according to the rules specified by the negotiation protocol.

Rosenschein and Zlotkin [Rosenschein & Zlotkin 94a] provide a good foundation for the analysis of real world multiagent applications. In the following section we present their categorisation of domains. They classify the domains in a three-level hierarchy. Subsequently, we address the possible forms of cooperation and investigate the different possibilities for deception and incentive compatible negotiation mechanisms to prevent it.

3.1 Definition of the Domains

The situation the agents find themselves in is characterised by the general properties of the domain and by the *encounter*, the momentarily specification of the goals they have to reach.

3.1.1 Task Oriented Domain

In a Task Oriented Domain (TOD) a number of agents have non-conflicting jobs (tasks) to do. It can be possible to redistribute the tasks among the agents so that the overall costs for the execution decreases. This provides the opportunity for cooperation by exchanging tasks among the agents.

A TOD is defined by a set of possible tasks, a set of agents and a monotonic cost function which assigns to each subset of tasks the costs of executing them. An encounter is an assignment of tasks to agents.

Definition 11 A Task Oriented Domain (*TOD*) is a tuple $\langle \mathcal{T}, \mathcal{A}, c \rangle$ where:

1. \mathcal{T} is the set of all possible tasks;
2. $\mathcal{A} = \{A_1, \dots, A_n\}$ is an ordered list of agents;
3. c is a function $c : 2^{\mathcal{T}} \rightarrow \mathbb{R}^+$ such that for each finite set of tasks $X \subset \mathcal{T}$, $c(X)$ is the cost of executing all the tasks in X by any single agent. c is monotonic, i.e., for any two finite subsets $X \subset Y \subset \mathcal{T}$, $c(X) \leq c(Y)$.
4. $c(\emptyset) = 0$.

Definition 12 An encounter within a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$ is an ordered list (T_1, T_2, \dots, T_n) such that for all $k \in \{1 \dots n\}$, T_k is a finite set of tasks from \mathcal{T} that A_k needs to perform. T_k will also be called A_k 's goal.

We observe that every agent is able to perform any finite set of tasks with finite costs and that there are no incompatible goals. Therefore the agents have no disadvantages from the existence of the other agents. However they might profit from cooperation by exchanging tasks in a way that reduces the costs for the execution. Besides, no side effects occur in a TOD; an agent cannot fulfil another agents' goal by accident. Hence explicit cooperation is necessary to benefit from the other agents.

Although the definition of the TODs is quite restrictive it includes some classical multiagent domains like the postmen domain, which will serve as a simple example later. Furthermore, the concept of TOD is useful for analysing simplified versions of real-world applications¹. Beyond that, it is reasonable to refine the definition and to introduce a hierarchy of more specific TODs:

Definition 13 A TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$ is called **subadditive** if for all finite $X, Y \subseteq \mathcal{T}$, we have $c(X \cup Y) \leq c(X) + c(Y)$.

In a subadditive domain the costs of performing two sets of tasks together is never higher than the costs of performing them separately. To bunch tasks is often a good idea in subadditive TODs.

Definition 14 A TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$ is called **concave** if for all finite sets of tasks $X \subseteq Y, Z \subseteq \mathcal{T}$, we have $c(Y \cup Z) - c(Y) \leq c(X \cup Z) - c(X)$.

In a concave TOD the costs that an arbitrary set of tasks Z adds to a set Y cannot exceed the costs that Z adds to any subset $X \subset Y$.

Definition 15 A TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$ is called **modular** if for all finite sets of tasks $X, Y \subseteq \mathcal{T}$, we have $c(X \cup Y) = c(X) + c(Y) - c(X \cap Y)$.

¹With some modification of the definition we can even gather the shipping domain (see Section 5)

Modular Domains are the most restricted TODs. The costs of combining two sets of tasks is exactly the sum of the costs of the sets minus the costs of their intersection.

It is easy to see that concave TODs are subadditive and modular TODs are concave.

3.1.2 State Oriented Domain

In State Oriented Domains (SODs), which are more general than TODs, the goals of the agents are not to execute independent tasks but to move the world into certain states. The negotiation objects in SODs are the agreements on a common goal state and the development of joint plans to reach that state.

In the SODs the environment is characterised by a set of states, a set of agents, a set of possible joint plans and a cost function which assigns for every plan the costs of each agents role to that agent. An encounter in a SOD is an assignment of a set of goal states to each agent.

Definition 16 A State Oriented Domain (SOD) is a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{J}, c \rangle$ where:

1. \mathcal{S} is the set of all possible world states;
2. $\mathcal{A} = \{A_1, \dots, A_n\}$ is an ordered list of agents;
3. \mathcal{J} is the set of all possible joint plans. A joint plan $J \in \mathcal{J}$ moves the world from one state in \mathcal{S} to another. The actions taken by agent k are called k 's role in J and will be referred to as J_k .
4. c is a function $c : \mathcal{J} \rightarrow (\mathbb{R}^+)^n$. For each joint plan $J \in \mathcal{J}$, $c(J)$ is a vector of n positive real numbers, the cost of each agent's role in the joint plan. If an agent plays no role in J , his cost is 0.

Definition 17 An encounter within a SOD $\langle \mathcal{S}, \mathcal{A}, \mathcal{J}, c \rangle$ is a tuple $\langle s, (G_1, G_2, \dots, G_n) \rangle$ such that $s \in \mathcal{S}$ is the initial state of the world and for all $k \in \{1 \dots n\}$, G_k is the set of all acceptable final world states from \mathcal{S} for agent A_k . G_k will also be called A_k 's goal.

In a SOD the agents' goals are to move the environment into certain goal states. A goal can either be achieved or not; it cannot be partially achieved. Hence, a conflict between the agents is possible if the intersection of their goal states is empty or very expensive to reach. However, if there is a common goal state the agents might benefit from each other in reaching it by cooperating and thus lowering the costs.

In contrast to the TODs, in a SOD actions have side effects. An agent performing an action might hinder or help another agent without knowing. For example an agent can accidentally achieve the goal of another while working on its own goal. The blocksworld is a well-known example for a SOD.

| | description | cooperation properties |
|-----|---|---|
| TOD | independent goals no conflicts | increase of performance by redistribution of tasks |
| SOD | goal is set of states conflicts possible | cooperation conflict solving |
| WOD | worth instead of goal states | cooperation conflict solving |

Figure 3.1: Summary of the Domains

3.1.3 Worth Oriented Domain

The definition of a Worth Oriented Domain (WOD) equals that of a State Oriented Domain. Yet an encounter is defined in a different way:

Definition 18 *An encounter within a WOD $\langle \mathcal{S}, \mathcal{A}, \mathcal{J}, c \rangle$ is a tuple $\langle s, (W_1, W_2, \dots, W_n) \rangle$ such that $s \in S$ is the initial state of the world and for all $k \in \{1 \dots n\}$, $W_k : \mathcal{S} \rightarrow \mathbb{R}^+$ is the worth function of agent k . W_k assigns some value to each possible final state in the world. W_k will also be called A_k 's goal.*

Like the SODs the WODs are characterised by a set of states some of which are fixed goal states. But in contrast to the goal states in SODs there is assigned worth to the states such that some states are more desirable to reach than others. By assigning worth to the states it becomes possible to achieve goals partially or split them into subgoals which, when achieved, give the agent already some utility. This makes further cooperation possible as it allows the agents to compromise. Like in SODs the objects of negotiation in WODs are joint plans but the agreement on a goal state is more complex because it is possible for the agents to make concessions by relaxing their goals.

Table 3.1 summarises the properties of the three types of domains and describes the differences between them.

3.2 Cooperation and Deception

There are several situations in which an agent might be tempted to take advantage of the other agents' cooperativeness. If every agent tries to do so (what might unfortunately be individually rational), the overall solution will be bad (the tragedy of the commons). Because of the different characteristics of the domains, different forms of deception may be profitable.

Before we turn to deception we should define what we regard as a fair outcome of a negotiation. In the next paragraphs we will define the term negotiation mechanism, list the criteria Nash introduced to describe a fair negotiation mechanism,

give some examples of cooperation in the domains classified above and explain how deception can be prevented by choosing appropriate mechanisms.

3.2.1 Negotiation Set, Protocols and Strategies

A *negotiation mechanism* is defined by a set of possible deals, a negotiation protocol and a negotiation strategy. The negotiation set is the set of possible deals which the negotiation might concern. For example, in a TOD the set of *pure deals* is the set of all possible redistributions of the tasks among the agents. A *mixed deal* is a partitioning of the tasks but instead of explicitly assigning the partitions to the involved agents they agree on performing a lottery with a certain probability to determine who is going to perform which part of the tasks. An *All-or-Nothing deal* is a mixed deal where one of the agents will perform all the tasks. The negotiation over mixed deals allows the agents to agree on solutions which are globally better because the sum of the agents' costs is lower than in any pure deal. However, they are not individually rational as single agents have more work to do than they would if they were alone. Besides "creating more utility" the use of mixed deals and all-or-nothing deals can prevent deception in certain TODs.

The negotiation protocol specifies the rules of interaction which control the process of finding an agreement. Finally, the negotiation strategy of a single agent determines the behaviour of that agent according to the liberties the protocol leaves.

Domain attributes affect the properties of the negotiation mechanism. A protocol that motivates agents to act in an intended way in one type of domain can generate unintended behaviour in a domain of a different type.

The concept of the Nash equilibria from game theory can easily be adopted to negotiation strategies: A pair of negotiation strategies s_1 and s_2 is said to be in Nash equilibrium if s_1 is the best answer to s_2 and vice versa. A strategy is in symmetric Nash equilibrium if it is the best answer to itself.

3.2.2 Nash's Criteria for a Fair Negotiation Mechanism

Nash defined a "fair" negotiation mechanism by claiming a list of properties for the outcome of the mechanism [Nash 53].

- *Individual rationality*: every participant gets at least the utility he would have got without the agreement.
- *Pareto optimality*: it is not possible to change the agreement so that a participant gets more utility without lessening the utility of the others.
- *Symmetry*: if the situation is symmetric the solution should be symmetric as well, that means both agents should obtain the same expected utility.

- *Invariance with respect to linear utility transformations*: If the utility measure of one or more agents is changed by the application of linear functions it should not influence the outcome of the negotiation process.
- *Independence of irrelevant alternatives*: The elimination of deals which are not chosen as solutions does not affect the outcome of the negotiation process.

Nash showed that the only mechanisms that satisfy the above criteria are the *product maximising mechanisms* (PMM). The protocol of a PMM is symmetrically distributed and the strategy is in symmetric Nash equilibrium. Two agents using a PMM will agree on a deal that maximises the product of their utilities. If there are several product-maximising deals, the mechanism will choose the one that maximises the sum of the utilities. If there is more than one deal that maximises the sum and the product of the utilities, the mechanism will choose randomly. These requirements can be used to construct a simple PMM straightforwardly.

Example: Both agents propose a single deal from the negotiation set. From these deals the one with the higher product of utilities will be chosen, if both deals offer the same utility product then the deal with the higher sum of utilities will be chosen, in case that both deals also offer the same sum, a coin toss selects the outcome. The best strategy for an agent using this protocol is to compute the set of product and sum maximising deals and propose the one that gives him the highest utility.

3.2.3 Deception in Task Oriented Domains

The potential for cooperation in TODs is to save costs by exchanging tasks. Therefore, unfair agents will try to manipulate the negotiation to reach a better distribution of the tasks. There are few possibilities to cheat on a product maximising mechanism: any behaviour that violates the rules of the protocol will be detected by the others and can cause them to break off the negotiation; furthermore, no aberration from the supposed strategy can be beneficial because the strategy is in symmetric Nash equilibrium. Hence, the only way to manipulate the outcome of a product maximising mechanism in a TOD is to be untruthful about the initial distribution of the tasks². Actually the possibilities of lying in a TOD are restricted to hiding tasks and inventing new tasks. An invented task is called *decoy task* if it is possible for the liar to produce this task if it is scheduled to another agent. A task which cannot be produced on demand is called *phantom task*. A decoy task is always safe while a phantom task might be detected. In this case a penalty mechanism should serve to prevent lying.

²Which is possible only if the agents have incomplete information about the goals of the others, what seems to be a realistic assumption.

| General | | | |
|---------|--------|---------|-------|
| | Hidden | Phantom | Decoy |
| Pure | L | L | L |
| Mixed | L | T/P | L |

| Subadditive | | | |
|----------------|--------|---------|-------|
| | Hidden | Phantom | Decoy |
| Pure | L | L | L |
| Mixed | L | T/P | L |
| All-or-Nothing | T | T/P | L |

| Concave | | | |
|----------------|--------|---------|-------|
| | Hidden | Phantom | Decoy |
| Pure | L | L | L |
| Mixed | L | T | T |
| All-or-Nothing | T | T | T |

| Modular | | | |
|----------------|--------|---------|-------|
| | Hidden | Phantom | Decoy |
| Pure | L | T | T |
| Mixed | L | T | T |
| All-or-Nothing | T | T | T |

Figure 3.2: Matrix of Incentive Compatibility

In [Rosenschein & Zlotkin 94a] Rosenschein and Zlotkin analyse in detail which types of deception can be beneficial in different TODs. Figure 3.2 summarises the results. The table entry **L** means that there exist negotiation situations such that lying is beneficial, e.g. in a negotiation over pure deals in a general TOD decoy lies may be beneficial. The entry **T/P** means that there is the danger of phantom lies being discovered. So if it is possible to penalise lying agents, it becomes the dominant strategy to tell the truth. **T** means that telling the truth is always the best strategy.

3.2.4 Cooperation and Deception in SODs and WODs

In State and Worth Oriented Domains the agents operate in a common environment. Therefore they have to agree on a goal state they want to reach and on a joint plan to reach that state.

In SODs there is a clear distinction between goal and non-goal states that allows to specify the following four situations:

- The situation is called *symmetric cooperative* if there exists a common goal state and a joint plan to reach it which costs both agents less or equal than

it would cost them to achieve their goal on their own. In this case both agents appreciate the presence of the other.

- The situation is called *symmetric compromise* if there exists a common goal state but the costs of jointly reaching it are for both agents higher than the costs of reaching their goals on their own. In this case both agents would prefer to be alone in the world.
- If there exists a common goal state and the costs of reaching it are for one agent higher and for the other lower than the costs of reaching their goals alone the situation is called *non-symmetric cooperative/compromise*.
- If the intersection of the agents' goal states is empty we speak of a *conflict* situation.

In the cooperative and compromise situations rational agents will agree to perform a plan that changes the world state into a common goal state. The utility of the agents depends on the costs of the role they are playing in the joint plan. A product maximising mechanism can lead to a fair allocation. As in TODs, mixed deals may increase the expected utility for both agents in this case as well.

There are three ways to deal cooperatively with conflict situations of which the simplest is to toss a coin to decide who may reach his goal. Although one of the agents will not get any utility in the end, the expected utility before the tossing of the coin is positive for both agents. The probability of the outcome of the coin may be due to negotiation. There are some domains (e.g. the blocksworld) where two agents can reach a goal at lower joint cost than a single agent. In that case it is possible to increase the expected utility of the agents using *Semi-Cooperative Deals* or *Multi-Plan Deals*. In a semi-cooperative deal the agents agree to save costs by cooperatively reaching an intermediate state from which both agents' goal states are reachable, and then tossing a coin to decide who may carry on to reach his own goal state. A multi-plan deal involves cooperation after the coin is tossed: For each agents' goal a joint plan is generated and a probability for the tossing of the coin is negotiated. Then the coin is tossed and the loser has to help the winner to reach his goal. This conflict solving strategy certainly requires the assuredness that agents keep their commitments after they lost the gamble. Since in a WOD there is worth assigned to the goal states and therefore it is possible to relax goals it is not straightforward anymore to choose that common goal state that is reachable with minimal costs. Instead, the agents have the problem of finding a goal state with maximal utility namely with the highest difference between worth and costs of reaching it. Additionally, the existence of a joint plan that achieves a fair (product maximising) distribution of the utility is necessary. If the agents have the computational power to find such a state and plan, a PMM will lead to an optimal and fair outcome of the negotiation.

Like in TODs, in SODs and WODs the agents have the chance to manipulate the outcome of the negotiation process in their favour by taking advantage of the other agents' incomplete information and lying about their set of goal states and (in WODs) about their worth function. In expanding or altering their actual goal they can pretend that it would be much easier for them to reach their goal alone, or that they obtain lesser utility from the joint plan, and so they can justify to do lesser work in the joint plan.

Like in the cooperative and compromise situation there is some chance for deception in the conflict situation as well. By lying that his own goal is easier to reach a single agent can pretend that he has less utility from the joint plan than his opponent; hence he will have to play a less expensive role in the semi-cooperative joint plan or his chances in the gamble will increase.

Unfortunately there is no general mechanism to prevent lying in a domain with incomplete information. If one agent has more information at his disposal than the others he might be able to invent a beneficial lie. If it is not possible to detect the lie, it is safe. A reasonable way to cut down deception in such domains may be to communicate the relevant information simultaneously such that none of the agents obtains a protrusion of information.

Chapter 4

Coordination Mechanism Design

“Coordination, the process by which an agent reasons about its local actions and the (anticipated) actions of others to try and ensure the community acts in a coherent manner, is perhaps the key problem of the discipline of Distributed Artificial Intelligence.” (N. R. Jennings, in [Jennings 96], p. 187).

Coordination mechanisms for multiagent systems are often chosen without explaining the reasons why precisely these mechanisms were selected. Certainly, in some domains with special types of environments and agent behaviours they work fine. But what happens if some assumptions change that are necessary for the intended performance of the mechanism?

How can we build up a coordination mechanism for a domain with known domain attributes in order to obtain a coordination ¹ of the agents that meets some desired criteria? A promising approach for answering this question provides the *economic theory of mechanism design* that uses the tools of economics and *game theory* to design “*rules of interaction*” for economic transactions that will yield some desired outcome.

This is a normative approach whose goal it is to generate protocols so that when agents use them according to some stability solution concept, e.g. dominant strategies or Nash equilibrium, desirable social outcomes follow.

Outlook on this Chapter

In the first section of this chapter we will provide a brief overview on the research area of coordination mechanism design. Application areas are the *contract theory*, the *theory of collective decision making* and allocation problems in general.

After that in the remaining sections of this chapter a detailed and comprehensive overview of available auction mechanisms as special forms of coordination mechanisms is given. This is done with the pretension to identify *incentive-compatible*

¹In this context by coordination we mean the assignment of items or tasks to agents. The overall goal is to achieve an efficient allocation.

auction-based coordination mechanisms which are truth-revealing. Thus, they guarantee efficient allocation decisions even under incomplete information about the agents private valuations or preferences. Both, a specification of their constituent rules and their impacts on the strategy selections of participating agents is presented in detail.

4.1 Economic Theory of Mechanism Design

An Allocation Problem under Incomplete Information

Assume you adopt the role of a planner who has to determine an allocation of a set of items to a set of individuals. Thereby both the procedure or mechanism you use and the resulting allocation have to meet some desired efficiency criteria. If you are interested in obtaining an economically efficient allocation your objective is to award each item to the individual who values it most. To accomplish this, you need informations about the individuals' private valuations or preferences. But assume you are in a decision situation with *incomplete information*, that is, the individuals are autonomous and you do not have perfect knowledge about their private preferences. The problem you are then faced with is the following: You need to ask the individuals for their preferences but you cannot rely upon that they truthfully reveal their private informations to you when they are *self-interested*².

This would cause no problems for the planner if he³ could use interaction rules under which consumers would truthfully reveal how much an item is worth to them, thus providing the means for the efficient provision to be made. The identification of such interaction rules is well-known as the problem of mechanism design for the efficient assignment of goods or, in short, as the *assignment problem*. It constitutes a central research objective of the economic theory of mechanism design.

4.1.1 The Assignment Problem

The **assignment problem** [Olson & Porter 94] is one of finding a way of efficiently assigning m discrete items to n people when their preferences are unknown to the person or mechanism making the assignment.

Consider the problem of assigning m elements of a set M of items (or slots) to n elements of a set of agents N such that the total system welfare is maximised. An outcome or assignment $x \in \mathcal{A}$ is an $n \times m$ matrix where $x_{ij} = 1$, if agent i is assigned item j , and $x_{ij} = 0$, if he is not.

²I.e., each individual acts only for its own benefit.

³We refer to an artificial agent with "it" and to an agent that could be both, human or artificial, with "he" or "she".

The preferences of each agent depend upon the item allocated, any monetary payment, and the agent's type. The type θ^i of an Agent i parameterises the individual values that i places on the goods being allocated. Each θ^i is an element of the set of all possible types Θ^i that an agent i may adopt. The entirety of all the agents' types $\theta = (\theta^1, \dots, \theta^n)$ is called *profile*. Each agent i calculates his valuation for an assignment $x \in \mathcal{A}$ (whereby \mathcal{A} is the set of all allocations) through an evaluation function $v(x, \theta^i) = \sum_j \theta_j^i x_{ij}$. For simplicity we will use $v_i(x)$ instead of $v(x, \theta^i)$.

e assume that the utility of agent i is *quasilinear*⁴ and is given by $u(x, t, \theta^i) = u_i(x, t) = v_i(x) + t^i$, where t^i is any monetary transfer to (or from) agent i . The assignment problem (A) can be mathematically described as a maximisation problem in which

$$\max_{x \in \mathcal{A}} W = \sum_{i \in N} \sum_{j \in M} \theta_j^i x_{ij}$$

is to be determined with respect to the following set of constraints:

$$\begin{aligned} \sum_{j \in M} x_{ij} &\leq 1 & \forall i \in N \\ \sum_{i \in N} x_{ij} &\leq 1 & \forall j \in M \\ x_{ij} &\in 0, 1 & \forall i \in N, \forall j \in M \end{aligned}$$

If an allocation solves (A), then it is said to be *outcome efficient*. W is called the *total (or social) welfare* of the system.

It is important to point out that the planner's aim is to bring about an allocation, that meets some desired efficiency criteria. The planner's objectives are thereby specified by an *allocation rule* [Groh 96] that defines the planner's desired target allocation(s) for a given θ .

Definition 19 (Allocation Rule) *An allocation rule is a function*

$$f : \Theta^1 \times \dots \times \Theta^n \longrightarrow \mathcal{A},$$

that specifies for each possible profile θ of agent types θ_i a non-empty set of desired allocations $f(\theta) \subset \mathcal{A}$.

In principle, thereby the planner strives after an assignment of allocations to the *true* types of the individual agents. Because incorrect information about the agents' types generally results in an undesired outcome of the planner's allocation procedure. That means a loss of global utility compared to the allocation that would result if the individuals revealed their private preferences or valuations truthfully.

The planner can try to solve this problem through decentralisation by constructing a *mechanism* [Dasgupta et al. 79] [Groh 96], also called *coordination mechanism*.

⁴In *quasilinear* environments the utility an agent gets from any assignment does not depend on the amount of utility he already possesses.

Definition 20 (Mechanism) *A mechanism $mech$ provides each agent i (for $i = 1, \dots, n$) with a finite set S^i of signals or strategies s^i , whereby $s^i \in S^i \forall i$ and $s = (s^1, \dots, s^n)$. The mechanism defines an allocation $\in \mathcal{A}$ by the function*

$$mech : S^1 \times \dots \times S^n \longrightarrow \mathcal{A}, s \longrightarrow mech(s) = x$$

Consequently, a mechanism can be interpreted as a set of rules that determine the permissible interactions of the participating agents. The agents are asked to send free signals or strategies to the planner that imply a certain allocation to them. Thereby the individuals know the mechanism the planner uses. A mechanism may thoroughly consist of a complex dynamic procedure. In this case the strategy sets S^i would consist of plans of actions that inform each agent at each stage of the mechanism what strategies are possible.

With a *direct mechanism* [Groh 96] individuals simultaneously and at once report—instead of any number of signals—only one single information, namely about their type $\hat{\theta}^i \in \Theta^i$, as signal $s^i \in S^i$ to the planner. An individual's type comprises all its private information. However, $\hat{\theta}^i = \theta^i$ must not necessarily hold, i.e. the reported information about the type has not necessarily to be true. Depending on the type of the allocation problem, this may be its private valuation for a single item or a vector of valuations for multiple items.

Definition 21 (Direct Mechanism) *A mechanism $mech_D$ is called a direct mechanism if $S^i = \Theta^i \forall i$.*

Definition 22 (Indirect Mechanism) *Mechanisms with $S^i \neq \Theta^i$ are called indirect mechanisms.*

As mentioned above, the agents know both the mechanism $mech$, which forms the basis for the determination of the resulting allocation, and their individual utility function u_i , which makes it possible for them to assess their utility loss or win $u_i(x, t)$ resulting from each possible allocation $x \in \mathcal{A}$.

On that basis they will deduce their individual strategies by using a game-theoretical solution concept and finally play strategies in the mechanism $mech$ that are *equilibrium strategies* s^{i*} and result in an equilibrium point $s^* = (s^{1*}, \dots, s^{n*})$. The set of all equilibrium strategies in a mechanism $mech$ is denoted by $E_{mech}(\theta)$.

An intuitive characterisation of an equilibrium is that no player has an incentive to deviate from his equilibrium strategy and therefore does not leave the equilibrium as long as he assumes that no other player is going to leave it (see the Nash equilibrium concept for instance).

If a planner strives after coordinating self-interested individuals such that the coordination process results in an allocation promoted by its objective function or allocation rule f the crucial question the planner is faced with is:

Is it at all feasible to design a coordination mechanism $mech$ that induces the participating agents with profile θ to play strategies (or to send signals) s^i such that $mech(s) \in f(\theta)$?

In answering this question the *revelation principle* (see the next section) plays a fundamental role. In the following we want to introduce some concepts needed to be able to understand precisely its motivation and statement.

Definition 23 *A mechanism $mech$ implements the allocation rule f completely if $\forall \theta \in \Theta$ the following holds:*

$$mech(E_{mech_D}(\theta)) = f(\theta)$$

Such a mechanism would generate allocations for the individual agents that coincide with the desired allocations for their true types specified by the allocation rule.

Definition 24 *A direct mechanism $mech_D$ implements the allocation rule f truthfully if $\forall \theta \in \Theta$ the following holds:*

$$\theta \in E_{mech_D}(\theta) \text{ and } mech_D(\theta) \in f(\theta).$$

Definition 25 (Dominant Strategy) *A strategy $s^i \in S^i$ is a dominant strategy for an agent i , characterised by his type θ , if and only if $\forall \hat{s}^i \in S^i$ and $\forall s^{-i} \in S^{-i} = S^1 \times \dots \times S^{i-1} \times S^{i+1} \times \dots \times S^n$ the following holds:*

$$u_i(mech(s^i, s^{-i}), t) \geq u_i(mech(\hat{s}^i, s^{-i}), t).$$

Definition 26 *A mechanism $mech$ implements the allocation rule f in dominant strategies if an equilibrium in dominant strategies s^* does exist and*

$$mech(E_{mech}(\theta)) \subseteq f(\theta) \text{ }^5.$$

Definition 27 *A direct mechanism $mech_D : \Theta \rightarrow \mathcal{A}$ implements the allocation rule f truthfully (in dominant strategies) if it has the truthful report of θ^i as dominant (possible) equilibrium strategy and $\forall \hat{\theta}^{-i} \in \Theta^{-i}$ the following holds:*

$$u_i(mech_D(\theta^i, \hat{\theta}^{-i}), t) \geq u_i(mech_D(\hat{\theta}^i, \hat{\theta}^{-i}), t) \quad \forall \hat{\theta}^{-i} \text{ and } mech_D(\theta) \in f(\theta).$$

The inequality statement above is called *incentive compatibility constraint*.

Definition 28 (Incentive-Compatible Mechanism) *A mechanism that implements an allocation rule f truthfully in dominant strategies is called incentive-compatible.*

⁵($mech(E_{mech}(\theta)) = f(\theta)$ for complete implementation)

An incentive-compatible mechanism makes it a dominant strategy for the agents to report or reveal their information truthfully. Therefore such mechanisms are often called *truth-revealing*.

Definition 29 (Direct Revelation Mechanism) *A direct incentive-compatible mechanism m_D is also called a direct revelation mechanism.*

It is clear that an incentive-compatible coordination mechanism cannot coincide with the interests of all of the individuals (for some of them lying could be beneficial in another mechanism). Now we are prepared to formulate the revelation principle.

4.1.2 The Revelation Principle

The *revelation principle* [Myerson 81, Emons 94, Varian 95, Ma et al. 88] makes a fundamental statement about which allocation rules can be implemented (and therefore which desired outcomes can be achieved) under incomplete information, i.e. when information is decentralised among several agents.

Theorem 1 (Revelation Principle) *If there exists a mechanism $mech$ that implements f in dominant strategies, then there exists an equivalent direct mechanism $mech_D$ that implements $f(\theta)$ truthfully in dominant strategies [Groh 96].*

Consequently, the revelation principle says that for any however complex mechanism resulting in some equilibrium outcome with a corresponding allocation there exists an equivalent direct revelation mechanism that yields the same equilibrium outcome [Emons 94] or allocation⁶.

That is, loosely speaking: “[...]whatever can be done with any mechanism can also be done with a direct revelation mechanism.” (Winand Emons, in [Emons 94], p.482). Negatively formulated, this yields the following *impossibility result*, that answers our crucial question stated above:

If there is no mechanism in the class of direct revelation mechanisms that implements the allocation function f and results in allocations $\in f(\theta)$ then it is not at all feasible to design such a coordination mechanism because there is no such mechanism in the class of all mechanisms.

So the revelation principle tells us that despite the existence of a multitude of conceivable (however complex) coordination mechanisms a planner who strives after implementing a desired allocation function can restrict himself to search in the class of direct revelation mechanisms without loss of generality [Martimort & Stole 97].

⁶For a proof of the revelation principle see [Groh 96].

The Vickrey Principle

An example for such a direct revelation mechanism is the *Vickrey Auction (VA)*, presented later in this chapter. The VA bases on the *Vickrey principle* or *Vickrey pricing rule*.

Definition 30 (Vickrey Principle) *The Vickrey principle lays down that*

- *if an item is auctioned-off, then the agent who places the highest bid for an item is awarded the item. But he is only charged a price equal to the second-highest bid.*
- *if an agent offers his service for performing an order and claims a certain payment for it, then the agent with the lowest claimed payment is awarded the order. But he receives a payment equal to the second-lowest offer.*

In so far, the Vickrey principle couples off the valuation an agent states for an item from the price he is finally charged if he wins it. Therefore, the bid or message sent to the auctioneer is in fact the entire private information of the agent: his utility function.

Certainly, other “indirect” mechanisms like negotiations and bargaining would be applicable. But the revelation principle [Myerson 81] states that “[...]anything that can be achieved by such an indirect mechanism can be achieved by a direct mechanism.”[Varian 95].

Incentive Contracting

Let us assume that individuals act in a non-collaborative environment. They have to perform some tasks and they possess and desire several items that they require for fulfilling their goals. When we consider real-world situations where individuals may benefit from contracting out some of their tasks to other agents or from selling (purchasing) items to (from) other agents we need to implement a monetary system for the provision of rewards into our coordination mechanism that guarantees an efficient allocation of items and tasks to the individuals.

In this case a mechanism is not only faced with the problem of finding an efficient allocation but also of settling prices that gives the individuals an incentive to take part in it and thereby to reveal their true valuations. This problem is well known as the problem of *incentive contracting*.

An approach for solving this problem consists in using incentive-compatible auction-based mechanisms for allocation and price settlement. In the following sections, mechanisms are discussed by which *items* can be auctioned off.

4.2 Classifying Auction-Based Mechanisms

Auction mechanisms can be used in two auction situations:

In the classical auction situation, an auctioneer tries—on behalf of a seller—to sell an item to a set of bidding agents and wants to get the highest possible payment. The bidding agents in contrast, have exactly orthogonal interests. Namely, they strive to acquire the item at the lowest possible price.

Besides this classical situation, also the contracting situation has to be mentioned in which an auctioneer wants to subcontract out tasks at the lowest possible price while the agents competing for the assignment of those tasks strive to acquire the highest possible payment for their task performing activities.

In both situations, the presented mechanisms work analogously and generally lead to a binding contract among two agents, namely the winning agent and the auctioneer.

Single-Item Auctions

A *single-item auction* (or single auction) is an auction where only one single item is sold.

There exist four classical auction mechanisms for auctioning-off a single item that are fundamental for the design and understanding of all other auction mechanisms. Those classical single-item auction mechanisms split up in two *progressive* and two (*one-shot*) *sealed-bid* auction procedures:

- progressive variants:
 - English Auction (EA)
 - Dutch Auction (DA)
- (one-shot) sealed-bid variants:
 - First-Price Sealed-Bid Auction (FPA)
 - Second-Price Sealed-Bid Auction, also called Vickrey Auction (VA)

Further variations on single-item auctions are:

- the Escalating-Bid Auction (EBA) and
- the Demange-Gale-Sotomayor Auction (DGS)

Multiple-Item Auctions

Multiple (i.e. $m \geq 2$) items can be sold by successively carrying out a sequence of m progressive single-item auctions. If the items are to be sold or allocated *simultaneously*, *multiple-item auctions* are used. Multiple-item auction variants are

- the Matrix Auction for Identical Items (MIA),
- the Matrix Auction for Heterogeneous Items (MHA),
- the Demange-Gale-Sotomayor Auction (DGS),
- the Generalised Vickrey Auction (GVA),
- the (monetary) Vickrey-Leonard Auction (VLA) and
- the Vickrey-Leonard Chit Auction (VLC)

Progressive Auctions

In *progressive auctions* bids or prices are lowered or raised incrementally until the most suitable bidder for receiving the item(s) can be selected. Progressive auctions can be *simultaneous* or *successive*. In simultaneous progressive auctions bidders are asked to put up their bids simultaneously.

Sealed-Bid Auctions

These auctions are preferable when an auction is distributed over space and/or time and when communication cost are high. Each bidder submits a sealed bid. The bids are opened and the item is awarded to the most suitable agent (generally the one with the highest bid). Normally the bids are obligatory and therefore cannot be progressively increased or decreased once they have been made.

Overall View

Table 4.1 gives an overall view of how auction mechanisms can be classified according to the number and disparity of the items that have to be allocated and the bidding rules that are used.

Depending on how the agents assign values (monetary equivalent to utility) to an item a further refinement of auction mechanisms is possible. One can distinguish three different value assignment settings in auctions that have a crucial impact on the individual strategies of each agent, because an agent's *strategy* is a function of both the individual value he assigns to an item and prior beliefs about the valuations of other agents.

Private Value Auctions

In private-value auctions an agent's marginal cost/valuation for a task/item is totally determined locally and independent of other agents' marginal cost/valuations.

These auction settings occur whenever a winning agent will not resell the item awarded to him but uses it for his own purposes.

| | Single-Item Auctions | Multiple-Item Auctions | |
|------------------------|----------------------|------------------------|------------------------|
| | | Identical Items | Heterogeneous Items |
| Progressive | EA, DA, DGS | DGS | DGS |
| Mixed | EBA | | |
| One-Shot Sealed-Bid | VA, FPA, MEVA | MIA VLA, VLC, GVA | MHA |

Table 4.1: The classification of auction mechanisms

Public Value Auctions

The opposite to private value auctions are *public-value auctions* in which an agent's valuation⁷ (cost) for an item (task) is entirely determined by other agents' valuations.

Common Value Auctions

In *common value auctions* an agent's valuation depends also entirely on the valuations of other agents but here the value of the item is the same for all bidding agents.

Correlated value Auctions

Here, the value an agent assigns to an item depends partly on its individual local preferences and partly on the values of other agents. As a result, these auction settings occur whenever a winning agent has the possibility to resell the item awarded to him to other agents.

For example, in contracting settings this is the case, when tasks can be subcontracted. An agent can—based on his local task performing capabilities—accept the order to perform a task for a specified payment but recontract out this task if he finds another agent that performs the task at lower cost, for instance. In this case the cost of the first agent depend solely on the cost of the latter.

⁷The valuation of an agent is equivalent to his expected utility.

4.3 Simple Auction Mechanisms

The following auction types are very prominent. Whenever the term “auction” is used you normally think of the English auction (EA). Since it is the best-known one because the EA is generally used when valuablenesses are auctioned off, for instance in the world’s most famous auction houses Sotheby’s and Christie’s. Tulips in Holland are auctioned off according to the rules of the Dutch auction and first-price sealed bid auctions are used in public invitations to bid. We call these three auction mechanisms simple because the way they settle prices for awarded items is very simple: There is only one winner who has made the highest bid, and who has to pay just the amount of his last bid.

The English Auction

Procedure Rules: The price of the item is continuously raised. Bidders who are unwilling to pay the current price drop out until only one bidder remains which has the highest value for the item.

Award: The remaining bidder gets the good for the price of his last bid.

Strategical Analysis: The best strategy for a bidder is to raise his bid until his maximum willingness to pay is reached or no other bidder is willing to overbid him.

Outcome: Consequently, the payed price for the piece of good is the willingness of the bidder with the *second* highest value to pay plus a tiny amount to outbid him. Since the bidder with the highest valuation or willingness to pay always gets the item awarded, the EA provides an economically efficient allocation.

The Dutch Auction

Procedure Rules: The Dutch auction (DA) starts with a price which is very likely to exceed the reservation prices of all bidders. The auctioneer lowers gradually that price until a bidder agrees to buy the item for the lastly announced price.

Award: The first bidder gets the unit at the currently announced price.

Strategical Analysis: Strategical bidding behaviour is inherent to this mechanism. Each bidder is tempted to counterspeculate the other agents’ *reservation prices*⁸ in order to anticipate their bids barely before the auctioneer reaches the corresponding reservation prices.

Outcome: Generally not economically efficient because the outcome is influenced by speculations about other agents’ reservation prices. Therefore the winning bidder is not necessarily the one with the highest willingness to pay.

⁸The maximal price that a bidder is willing to pay for an item or the minimal payment that a competitor in an invitation to bid accepts for his service is called reservation price.

The First-Price Sealed Bid Auction

Procedure Rules: All bidders make one sealed bid.

Award: The highest bidder gets the item and pays the price he bid for.

Strategical Analysis: Just as in the Dutch auction, in the First-Price Auction (FPA) strategical bidding behaviour is inherent to this mechanism. Each bidder is tempted to counterspeculate the other agents' reservation prices in order to give up a sealed bid that is slightly higher than the best bid of the competing agents. In so far, the DA and the FPA are strategically equivalent.

Outcome: In general this procedure will not achieve the desired objective of awarding the item to the bidder with the highest value because the outcome is influenced by "beliefs" about other agents' reservation prices. Therefore the winning bidder is not necessarily the one with the highest willingness to pay.

4.4 The Vickrey (Second-Price Sealed-Bid) Auction

Dominant truth-revealing strategies simplify bid preparation because it is optimal for a bidder to follow the truth-revealing strategy even if he assigns a positive probability to the possibility that his competitors will not reveal the truth but deviate from their equilibrium strategies and lie.

Procedure Rules: All bidders make one sealed bid.

Award: The item is awarded to the bidder with the highest bid, but the winning bidder has to pay only the *second-highest* bid.

Strategical Analysis: The VA is incentive-compatible. That is, placing bids that reflect their true individual valuations for the desired item is a dominant strategy for participating agents. In so far the VA strategically resembles the EA.

Outcome: Thus the Vickrey auction assigns the good to the person with the highest value regardless of the *accuracy of the beliefs* of the participants. In this way it provides a way to achieve the result of the English auction—namely an economically efficient allocation—without iterations.

4.4.1 Advantages of the Vickrey Auction

Vickrey [Vickrey 61] showed that in a sealed-bid second-price auction for symmetric risk-neutral bidding agents it is a dominant strategy to bid their true cost or values (i.e. truth-revealing strategies are not only equilibrium but also dominant strategies). A proof is also given in [Varian 95]. This is due to the fact that the winning agent is charged an amount equal to the second highest instead of his (highest) bid. This splits the correlation between the height of his bid and the price that the agent has to pay for the item. That is, the amount of a bid an agent

places only determines its rank among all bids and nothing else. If an agent bids an amount below his true valuation he has a smaller chance of winning and thus risks losing the award. Overbidding on the other hand involves the risk of losing utility. Since if an agent overbids his true valuation and this makes the difference between winning or not, then he places the highest bid though he has only an individual valuation that ranks behind the second-highest bid. Consequently, he has to pay a price above his valuation and loses utility.

Therefore strategic deliberations do not pay out for an agent in the VA and produce only overhead, which is resource consuming. In both the VA and the EA truth-telling is the dominant equilibrium strategy that leads to complete economic efficiency. That is, the bidder with the highest value for an item or lowest cost in performing a task always wins the item or task. There is no chance that a bidder with a higher value will lose the auction to a bidder with a lower value. Vickrey showed that the VA is “*logically isomorph*” or “*strategically equivalent*” to the English auction procedure in private value auctions. But they are not strategically equivalent in general. In private value auctions they will produce the same allocation at the same prices. But this result does not hold for correlated value auctions because in them the valuation of a bidder depends on the valuations of other bidders. And this information need about other bidders is satisfied in the EA but not in the VA. Therefore, they may lead to different results.

The VA shows the following properties desirable for a coordination mechanism:

- **Efficiency:** It is economically efficient (bidder with highest valuation/lowest cost wins the auction). The fact that the agents reveal their preferences truthfully allows globally efficient decisions to be made.
- **Simplicity:** The VA causes low communication cost (only sealed bid and win/loss message broadcasting) and low computation cost (Strategy deliberations and counterspeculating other agents do not pay out, because the valuations of other agents do not matter in making the bidding decision.).
- **Stability and Incentive-Compatibility:** Agents have no incentive to deviate from truth (truth-telling is the dominant strategy).

4.4.2 Shortcomings of the Vickrey Auction

Even though Vickrey auctions show those desired criteria for an efficient and acceptable coordination mechanism, due to [Rothkopf et al. 90] [Sandholm & Lesser 95] their applications are quite rare.

The Vickrey mechanism was developed to promote truthful bidding and to avoid counterspeculation among *self-interested* agents. Despite its counter-speculation reducing effects **limitations and deficiencies:**

- **Collusion vulnerability:** Bidders could form a bidding coalition in order to coordinate their bids such that the second highest (lowest) bid in the VA stays artificially lower (higher) than the true second lowest (highest) bid of the collusive members of the coalition. That way, the winning bidder has the chance to be charged a lower payment (or to receive a higher payment).

This means that the VA is not *coalition-proof*. But the simple auction mechanisms mentioned above are not coalition-proof, too. If an auctioneer wants to aggravate collusive agent behaviour, he should not use a VA or EA mechanism because collusive agreements are in them more stable than in a DA or FPA mechanism [Rasmusen 89]⁹.

- **Auctioneer reliability:**

The auctioneer has to be truthful. Otherwise he could overstate the second highest (or understate the second lowest) bid to the winning agent and charge him a higher price (respectively give him a lower payment). Additionally, because of the *high grade of information revealing* the auctioneer also learns about the competitiveness, strategies and true cost of participating agents and could pass these discoveries to competing agents. This insincerity about the reliability of the auctioneer has been suggested by [Rothkopf et al. 90] to be the main reason why Vickrey auctions are quite rarely used in auctions among humans.

The problem of Auctioneer reliability is likely to be solved by cryptographic signatures on each submitted bid such that a winning bidder is able to check if the second highest bid announced by the auctioneer really has been placed by another agent.

In the simple auction protocols—presented in the previous section—this problem cannot occur because each bidder pays his own bid amount in the winning case.

- **Recontracting problem:** The VA assures truth-promotion only for *private-value auctions*¹⁰.

In contracting protocols—like the contract net protocol—where a contractee can recontract out a task that he contracted in, the contractee's

⁹You can see this by assuming that some bidder has a value of 10 while the other bidders have only values up to 8. Now, the bidder with lower values agree to understate and to bid at most 5. In the VA the bidder with highest value can bid this value and does not risk losing the item while profiting from the collusion. In the EA, he also does not risk losing the item, because he is always able to overbid other agents that do break the collusive agreement. But in the DA and the FPA, for profiting from the collusive agreement, the bidder with highest value has to underbid even the second highest value among the other bidders and therefore risks losing the item if they break the agreement.

¹⁰In private-value auctions an agent's marginal cost/valuation for a task/item is totally determined locally and independent of other agents' marginal cost/valuations.

marginal cost for a task is, at least to some extent, defined by the potential recontracting prices. Therefore, such settings are not pure private value auctions but correlated value auctions, and truthful bidding is not necessarily the dominant strategy [Sandholm & Lesser 95]¹¹.

- **Future contingencies problem:** Future contingencies could promote strategic behaviour if they can influence marginal cost calculations/valuations when the VA (or another auction) is applied to a series of contracts of task sets. Because the marginal cost determination of a certain task set can effect the marginal cost of other task sets potentially negotiated over in the future. This is also a reason for the general unsuitability of the VA for cases with uncertain or unknown cost.
- **Minimum win claims:** In the contracting setting, dependent on their underlying utility function agents could be lazy in that they are not willing to perform tasks with low win margins. They generally feel not lucky when they have to perform a task and thereby are only compensated for their cost. Instead of this, they strive after realising a certain minimum win with each of their actions depending on the taken efforts. For instance, an autonomous forwarding company generally will not be willing to forward a shipment of oranges from California to Colorado for a win of 5 dollars. Therefore in bidding the company might be tempted to overstate its true cost in order to guarantee itself a minimum win in the winning case.

Whenever only relatively few shipping agents possess the ability to perform a special kind of tasks but the demand for those task performing abilities is comparatively high such strategies may pay out. But whenever there are quite many agents competing for comparatively few orders and when—moreover—those agents have similar cost for performing a task such strategies will fail.

- **Hiding of sensitive information:** If the outcomes of VAs are revealed an agent's competitors can figure out information about the cost or valuation structure and the competitiveness of this agent. First-price auction mechanisms do not as clearly reveal those information because in them bids are generally more strategic than truthful.

In domains where the mentioned problem types are excluded, Vickrey auctions succeed to guarantee truth-revelation and therefore an optimal efficient allocation of tasks.

¹¹The higher the environment dynamics the more the recontracting deliberations can be reduced. Also, the auctioned tasks could be bookmarked to restrict task recontracting to a desired amount.

There are several publications which are dealing with the usability of Vickrey auctions in multiagent systems, e.g. [Huberman & Clearwater 95, Drexler & Miller 88] among others.

4.5 The Multistaged Extended Vickrey Auction

The Multistaged Extended Vickrey Auction (MEVA) is a auction procedure that aims at the promotion of coalition-forming [Weinhardt et al. 96c, Weinhardt et al. 96a, Weinhardt et al. 96b]. As such, the MEVA tries to solve an important problem of DAI, namely how agents can be enabled to make use of synergies by forming coalitions or by starting cooperations. It would be ideal if individual agents could find cooperation partners explicitly through negotiations and bid together with them as a coalition in the auction.

In the process of coalition formation the following inherent problem always arises: An individual agent does not know if he can realise a higher profit by joining a coalition or by making a bid on its own. Thus, a coordination mechanism that encourages coalition forming, has to give each agent an *incentive* to join the coalition. In order to fulfil the *principle of individual rationality* the mechanism must guarantee each single bidder joining a coalition at least that pay-off he would have got bidding on his own—his individually realisable profit or *reference win* (rw). To assure *incentive compatibility* the mechanism has to guarantee it to each participating agent.

For solving this problem the Multistaged Extended Vickrey Auction (MEVA) was developed. It guarantees that each individual agent has an incentive to join a coalition in that it assures each single bidder joining a coalition his reference win.

4.5.1 Specification of the MEVA Procedure

A Rough Overview of the MEVA Procedure

By this coordination mechanism tasks are auctioned off in an *iterative* multistaged bidding process, whereby the number of general stages equals the number of participating agents. After the general stages an additional termination stage follows in which the auctioned task is awarded, its price settled and possible reference wins ascertained.

In each general stage i ($i = 1, \dots, n$) of the MEVA exclusively coalitions consisting of i agents are asked to make a bid.

In each stage, the auctioneer stores the bids together with the coalitions that made them. A feedback about the award of the task and the charged price is given by the auctioneer not in the single stages but only after the n th stage of the MEVA in the termination stage.

The award goes to the coalition with the highest bid of all rounds which is charged according to the Vickrey principle a price equalling the second-highest bid of all rounds. Thereby, in the ascertainment of the second-highest bid, bids that were made in previous rounds by members of the winning coalition are excluded.

After the award, the auctioneer has to investigate a reference win (minimum win a coalition would have got by making a bid on its own) if a subset of the winning coalition has made at any stage of the MEVA as individual or coalition the highest bid of that stage and this bid was higher than the price the winning coalition finally has to pay.

In the following, an algorithmic specification of the stages of the MEVA procedure will be given, for which firstly some notations have to be explained:

Notation for the MEVA:

- hb_i : highest bid of stage i
- hb_i^* : highest bid of all stages $\leq i$
- hb^* : highest bid of all stages
- shb_i : second highest bid of stage i
- $ashb_i^*$: admissible second highest bid of all stages $\leq i$
- $ashb^*$: admissible second highest bid of all stages
- $bidder(bid_i)$: agent or coalition which made bid_i
- $rw(bidder(\cdot))$: reference win of $bidder(\cdot)$

Initial Stage

The initial stage of the MEVA corresponds to the simple VA besides that hb_1 and shb_1 are not announced publicly.

Initial stage of the MEVA:

- Individual agents are asked to make a bid for the item or task.
- The auctioneer stores the highest bid hb_1^* and the second highest bid $ashb_1^*$ with the corresponding bidders.

End of the initial stage of the MEVA

General Stage

If in a stage i (line 01) no agent or coalition of agents has made a bid, the MEVA goes on to the next stage $i + 1$ where coalitions of size $i + 1$ are asked to make




Figure 4.1: Initial stage of the MEVA

bids except if i equals n , i.e. the grand coalition was asked to make a bid (line 03). In that case, the MEVA goes to its termination stage with its highest and second-highest bids being the highest and second-highest bids of all previous stages (lines 06, 07).

General stage i of the MEVA:

```

(01) if there is no placed bid  $hb_i$ 
(02) then
(03)   if  $i < n$ 
(04)   then goto stage  $i + 1$ 
(05)   else
(06)      $hb^* := hb_{i-1}^*$ 
(07)      $ashb^* := ashb_{i-1}^*$ 
(08)     goto termination stage
(09)   end
(10) else
(11)   if  $hb_i < hb_{i-1}^*$ 
(12)   then
(13)      $shb_i^* := hb_{i-1}^*$  (14)      $ashb_i^* := \max(hb_i, shb_{i-1}^*)$ 
(15)     award tentatively  $bidder(hb_i^*)$ 
(16)     goto stage  $i + 1$ 
(17)   end
(18)   if  $hb_i \geq hb_{i-1}^*$ 
(19)   then  $hb_i^* := hb_i$  and
(20)     if  $bidder(hb_{i-1}^*) \notin bidder(hb_i^*)$ 

```

```

(21)         then
(22)              $ashb_i^* := \max(shb_i, hb_{i-1}^*)$ 
(23)             award tentatively  $bidder(hb_i^*)$ 
(24)             goto stage  $i + 1$ 
(25)         end
(26)     else
(27)         if  $bidder(hb_{i-1}^*) \subset bidder(hb_i^*)$  and  $hb_{i-1}^* \geq shb_i$  12
(28)         then
(29)             calculate and store  $rw(bidder(hb_{i-1}^*))$  13
(30)              $ashb_i^* := \max(shb_{i-1}^*, shb_i)$ 
(31)             award tentatively  $bidder(hb_i^*)$ 
(32)             goto stage  $i + 1$ 
(33)         else
(34)             if  $bidder(hb_{i-1}^*) \subset bidder(hb_i^*)$  and  $hb_{i-1}^* < shb_i$ 
(35)             then
(36)                  $ashb_i^* := shb_i$ 
(37)                 award tentatively  $bidder(hb_i^*)$ 
(38)                 goto stage  $i + 1$ 
(39)             end
(40)         end
(41)     end
(42) end
(43) end

```

| |
|---|
| End of the general stage of the MEVA |
|---|

If at least one bid has been placed, the highest bid hb_i of stage i can be determined. If this is lower than the highest bid hb_{i-1}^* of all previous stages, this remains the highest in stage i and hb_i^* is set to hb_{i-1}^* . Then, the bidder who is tentatively awarded the task stays the same, the admissible second-highest bid of all run stages is possibly updated, and the MEVA goes to its next stage (lines 11-16).

Otherwise, if the highest bid in stage i turns out to be the highest bid of all stages, hb_i^* has to be set to hb_i (line 19).

If no subset of the coalition of agents that has made the current highest bid has made the highest bid of all previous stages, nothing special happens. Then the MEVA updates its admissible second-highest bid so far, awards tentatively the coalition with the current highest bid and enters its next stage (lines 19-25).

¹²The auctioneer calculates a reference win for a subset of the coalition with the highest bid hb_i^* at stage i if that subset made the highest bid at a *previous* stage **and** this bid is also higher than the second highest bid of the current stage shb_i^* .

¹³Because $bidder(hb_{i-1}^*)$ would have got the item alone if he had not joined the winning coalition $bidder(hb_i)$

Even if a subset of the coalition of agents that has made the current highest bid has made the highest bid of all previous stages, the MEVA runs through the same steps, if their previous bid is outperformed by the second-highest bid of the current stage (lines 34-38).

An interesting case occurs when a subset of the currently highest-bidding coalition has made the highest bid of all previous stages and this bid outperforms the second-highest bid made in stage i (line 27). Then, in order to make joining the larger coalition individually rational for that subset, its reference win has to be calculated and stored, which may be needed in the later termination stage (line 29). After that the MEVA runs through the habitual steps (lines 30-32).

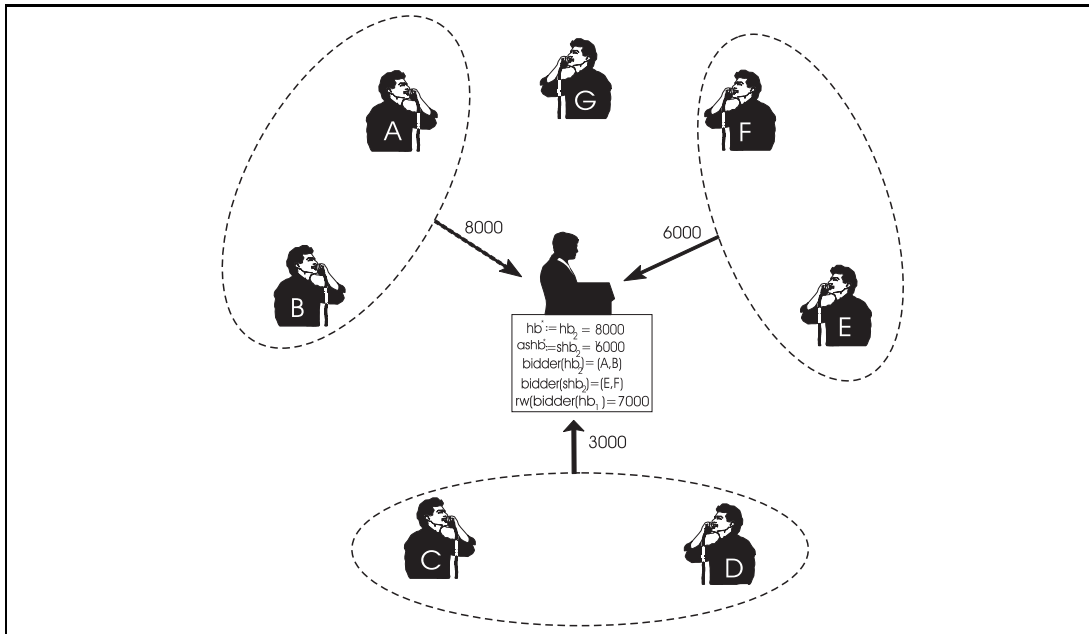


Figure 4.2: Second stage of the MEVA

Termination Stage

Termination of the MEVA:

Awarding and price settlement proceed according to the Vickrey mechanism.

- The current tentatively awarded agent or coalition has made the highest bid hb^* of all stages and hence gets the item at a price equal to the *admissible* second highest bid $ashb^*$ of all stages¹⁴.

¹⁴In determining the charged *admissible* second highest bid the auctioneer does not take into account bids from previous stages that were made by subsets of coalitions that bid at the last stage.

- The coalition win has to be split up among the members of the winning coalition while a possible reference win $rw(sub)$ accrued for a subset sub of the winning coalition has to be taken into account:

do

procedure split_up_win(set_of_agents)

1. if $\exists rw(sub)$ for $sub \in set_of_agents$ then the auctioneer announces sub and $rw(sub)$ to all coalition members
2. sub receives its reference win
3. call split_up_win(sub)
4. the difference between the joint coalition win and $rw(sub)$ has to be split up in further coalition negotiations between all coalition members

End of termination of the MEVA

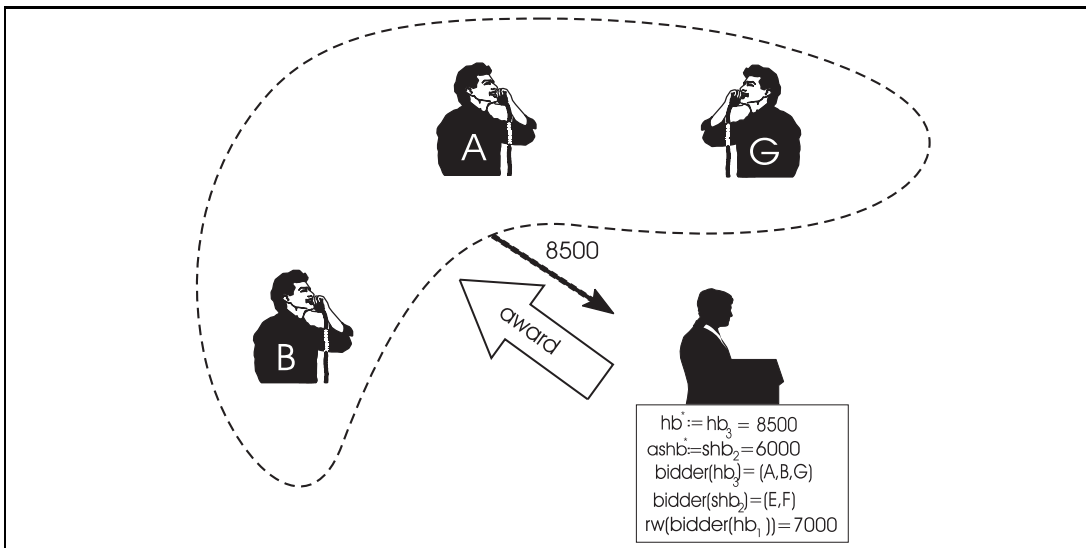


Figure 4.3: Termination of the MEVA after stage 3

4.5.2 Analysis of the MEVA

The MEVA mechanism provides an incentive to join coalitions and therefore seems to guarantee efficient allocations.

Critical Remarks

Regarding the MEVA mechanism three critical remarks have to be made:

1. It is doubtful if the MEVA is really able to rule out strategic bidding behaviour.

Lying could be profitable for agents, because overbidding true valuations in the first stage could ensure higher reference wins if one makes it to join a winning coalition at a later stage. Therefore, agents could be inclined to overbid their true valuations in the first round if they think that they have a good chance to join a competitive coalition in a following stage.

2. Besides this, Weinhardt's explanation [Weinhardt et al. 96c] of the MEVA mechanism lacks elaborateness and clarity—regarding
 - i. the allowed coalition forming processes.

He gives no comment as to whether at stage 3 the forming of coalition (A,B,F) as in Figure 4.4 would be allowed at all after F has joined a coalition with other agents in a previous stage who do not participate in (A,B,F). Assuming it would be allowed, he still leaves it unclear which price $p = ashb^*$ that coalition would have to pay. According to Weinhardt's MEVA rules (A,B,F) would have not to pay $shb_2 = 6000$ because this bid has been made by the coalition (E,F) in which F was member.

Furthermore, according to the rules—because the second and the third highest bid were also made by members of (A,B,F)—the admissible second highest bid would be only 3500 (made by agent G at stage 1) instead of 6000 in the scenario shown in Figure 4.3.

- ii. the division of (coalition win(set_of_agents) - reference win(sub)) among all agents $\in set_of_agents$.

4.5.3 Application of the Shapley Value for Profit Division in the MEVA

There is a lack of fair profit division schemes among agents in the DAI literature. As mentioned above, though the MEVA mechanism introduces a reference win concept it leaves this problem halfway unsolved as far as the division of the residue of the coalition win is concerned. The MEVA mechanism tells us only that the residuing coalition win is split up by a negotiation process but gives us no advice in the choice of a special negotiation mechanism.

In the scenario of Figure 4.3 the joint coalition win of (A,B,G) that has to be split up among A,B and G by negotiation amounts to $8500-6000=2500$. More precisely, because the reference win of (A,B) amounts to $8000-6000=2000$, 500 units of money have to be divided among all coalition members and 2000 currency units among the coalition's subset (A,B).




Figure 4.4: Lack of elaborateness in the MEVA mechanism

| | Sequence of coalition membership and contribution to the joint coalition win | | | | | | | |
|-------|---|------|------|------|------|------|------|----------|
| | 1st | A | A | B | B | G | G | |
| | 2nd | B | G | A | G | B | A | Σ |
| | 3rd | G | B | G | A | A | B | |
| Agent | A | 1000 | 1000 | 2000 | 2500 | 2500 | 1000 | 10000 |
| | B | 1000 | 1500 | 0 | 0 | 0 | 1500 | 4000 |
| | G | 500 | 0 | 500 | 0 | 0 | 0 | 1000 |

Table 4.2: Shapley value calculation for (A,B,G)

For this purpose the Shapley value [Shapley 53a] (see Table 4.2) can be used. According to the Shapley value A's share in the joint coalition win would be $win(A) = 10000/6 = 1666,66$. B's share would amount to $4000/6=666,66$ and F would get $1000/6=166,66$ units of money if the Shapley value is used as profit division scheme.

Now assume the forming of the coalition (A,B,F) (see Figure 4.4) would be allowed. Weinhardt leaves it unclear which price $p = ashb^*$ that coalition would have to pay. If Weinhardt's MEVA rules were followed strictly, (A,B,F) would not have to pay $shb_2 = 6000$ because this bid has been made by the coalition (E,F) in which F was member.

Furthermore, according to the rules—because the second and the third highest

| | Sequence of coalition membership and contribution to the joint coalition win | | | | | | | Σ |
|-------|---|------|------|------|------|------|------|----------|
| | 1st | A | A | B | B | F | F | |
| | 2nd | B | F | A | F | B | A | |
| | 3rd | F | B | F | A | A | B | |
| Agent | A | 1000 | 1000 | 2000 | 5000 | 5000 | 1000 | 15000 |
| | B | 1000 | 4000 | 0 | 0 | 0 | 4000 | 9000 |
| | F | 3000 | 0 | 3000 | 0 | 0 | 0 | 6000 |

Table 4.3: Shapley value calculation for (A,B,F)

bid were also made by members of (A,B,F)—the admissible second highest bid would be only 3500 (made by agent G at stage 1) instead of 6000 in the scenario shown in Figure 4.3.

It is doubtful if the price settlement according to the MEVA pricing rules is accepted by sellers if agents are allowed to leave their old coalition partners and join another coalition from one stage to another.

In the scenario of 4.4 the joint coalition win of (A,B,F) amounts to $8500 - 3500 = 5000$. The reference win of (A,B) amounts to $6000 - 3500 = 2500$. Hence, 2500 currency units have still to be divided among the coalition members by negotiation.

The resulting table for the calculation of the Shapley value is table 4.3. According to the Shapley value A's share in the coalition win joint would be $win(A) = 15000/6 = 2500$. B's share would amount to $9000/6 = 1500$ and F would get $6000/6 = 1000$ units of money.

4.6 Matrix Auctions

Matrix auctions (MA) are applicable for the simultaneous assignment of multiple items or tasks to agents.

We have a standard assignment problem if only one job is to be assigned to a single organisational unit. If we have several tasks to assign, they can be assigned either successively or simultaneously.

The valuation of a set of items can differ significantly from the sum of the valuations of each single item. For instance, this may be the case if the items to be allocated reflect tasks. Performing tasks is resource consuming and therefore may result in a loss of utility which can be reflected by a negative bid. Negative

bids are explicitly allowed in a matrix auction. They are useful if a set of tasks *has* to be completely performed by a group of agents.

4.6.1 Matrix Auctions for Multiple Identical Items

The agents get the information about the item by broadcast or over a blackboard. They calculate their valuations¹⁵ and report them to the auctioneer.

From the transmitted valuations or bids of the agents the auctioneer sets up a matrix whose cells represent the bids of the agents for each number of jobs.

An assignment algorithm can evaluate an efficient allocation from this data. The assignment algorithm has thereby to be adapted to the facts that the maximum of assignments in each row equals one and that the sum of the assignments in each column multiplied with the number of columns adds up to m at most.

Price Settlement According to the Vickrey Principle

For each number of tasks belonging to the optimum allocation the bidder with the highest bid receives the job and has to pay the second-highest price. If there is more than one marked cell the *Vickrey pricing mechanism for m identical items*¹⁶ [Vickrey 61] is applied.

Vickrey points out that this method for price settlement has the “[...] *advantage of reducing the probability that a bidder’s own bid will affect the price he receives, thus inducing bids closer to the full value to the bidder, improving the chances of obtaining or approaching the optimum allocation of resources, and reducing effort and expense devoted to socially superfluous investigation of the general market situation.*” [Vickrey 61]. The proof that it isn’t rational for a bidder to make a bid above or below his true valuation is analogous to the single auction.

But Vickrey also remarks that this “[...] result applies only to cases where each bidder is interested in at most a single unit, and there is no collusion among the bidders.” Problems also arise when the number of bidders is too few to produce a fully competitive market and simultaneously an individual bidder is interested in allocating two or more of the units, because the bid offered for the second unit will be influenced by the possible effect of this bid on the price to be paid for the first.

Example: Assume that there are 4 identical items to assign and that there are 3 agents A, B and C that report the following valuations or bids respectively.

¹⁵Whereby often the valuation of m items or tasks does not equal m times the valuation of one item or task. Often there is neither a linear nor a monotonous correlation between the valuation of one task to another.

¹⁶ m items are sold to the m highest bidders for the uniform price of the $(m + 1)$ st highest bid.

| | Number of identical items | | | |
|-------|---------------------------|-------|----|---|
| Agent | 1 | 2 | 3 | 4 |
| A | *60* | 70 | 80 | — |
| B | *40* | 80 | 90 | — |
| C | 25 | *100* | — | — |

Then the allocation that maximises social welfare respective global utility is the one marked with asterisks. According to the matrix auction rules, A and B have to pay an amount corresponding to the third-highest bid (25). C has to pay the second-highest price (80).

But for this class of allocation problems an efficient allocation does not necessarily mean that the award goes to the agent who made the highest bid. Furthermore, each package size can get assigned to several agents.

When an agent is assigned to an item or item package for which he has not made the highest bid, then he has not to pay the second-highest but the next-lower bid.

Therefore, agent C still has to pay 80 currency units for a package of 2 items, when agent A's valuation for this package is heightened to 110:

| | Number of identical items | | | |
|-------|---------------------------|-------|----|---|
| Agent | 1 | 2 | 3 | 4 |
| A | *60* | 110 | 80 | — |
| B | *40* | 80 | 90 | — |
| C | 25 | *100* | — | — |

4.6.2 Matrix Auctions for Multiple Heterogeneous Items

Here more data has to be communicated to the agents. If the set of items contains m items, the agents are asked to calculate their valuations for $2^m - 1$ potential item or task combinations.

From the transmitted bids or reported valuations of the agents the auctioneer sets up a matrix whose cells represent the bids of the agents for each combination of tasks. Starting from this he identifies the optimal allocation. For this variation of the assignment problem he uses an algorithm that has to consider that the maximum of assignments in each row equals one. Beyond this, assigned task combinations must not have any item or task in common, i.e. they must be disjunctive (form a partition of the item set). How this assignment problem can be formalised and solved within a multiagent system is discussed in [Ruß 97].

Price Settlement According to the Vickrey Principle

The price for each assigned task subset equals the second-highest bid in the matrix column for that task set. This Vickrey pricing assures that the agents reveal their true valuations because the bid of an agent as revelation of his valuation for an item does determine if he gets awarded the item but does not influence the price he has to pay for it.

Example with 3 different items allocated to 3 agents:

| | Task combinations | | | | | | |
|-------|-------------------|-----------|-----|-------|-----------|-------|---------|
| Agent | {1} | {2} | {3} | {1,2} | {1,3} | {2,3} | {1,2,3} |
| A | 10 | 30 | 40 | 60 | -20 | 10 | -60 |
| B | 5 | -10 | 30 | -30 | 70 | 40 | 80 |
| C | -10 | 70 | 60 | 40 | 40 | -20 | — |

According to the Vickrey principle agent C pays 30 currency units for item 2 and agent B is charged 40 currency units for the items 1 and 3 together.

Matrix Auction Analysis

Both Matrix auctions apply the Vickrey mechanism for price settlement and hence are *incentive compatible* (the bid only determines whether the bidder belongs to the m highest bidders, but does not determine the purchase price).

They differ in the used assignment algorithms to calculate an optimal allocation.

4.7 Clarke Taxing Auction Procedures

In the following, we examine variants of the Vickrey and English auction—used as allocation mechanisms for multiple items—in which agents make monetary payments to the planner.

As in the other above presented allocation mechanisms, agents have again no incentive to misrepresent their true preferences (i.e. to lie) in order to get a better outcome when they are not able to form coalitions.

4.7.1 The Clarke Tax

The Clarke Tax Mechanism (CTM) [Clarke 71, Clarke 72, Rosenschein et al. 91, Rosenschein & Ephrati 91a]—also referred to as the Clarke-Groves mechanism [Emons 94]—is a non-manipulable one-shot voting-by-bid mechanism that extracts truth from agents in quasilinear environments. Clarke's Mechanism

makes sure that each voting agent has only one dominant strategy, telling the truth¹⁷, by charging a tax based on how much the agent was able to sway the assignment decision. The problem can be mathematically formalised as follows:

- $\theta = (\theta^1, \dots, \theta^n)$: true types of all agents
- $\mathcal{R} = (r^1, \dots, r^n)$: reported types of all agents
- θ_{-i} : true types of all agents except i
- \mathcal{R}_{-i} : reported types of all agents except i
- $x(\theta)$: social choice function
- $x^*(\theta)$: social choice where the system cannot improve
- t^i : money transfer to (or from) agent i
- $v_i(x)$: valuation of agent i for an allocation x

An agent pays nothing if its announcement does not change the decision, but it does pay if its influence is pivotal. Each agent is fined with a tax that equals the portion of its bid that made a difference to the outcome. By informing agents of the taxing mechanism in advance, the agents have no incentive to lie about their preferences.

$$Clarke_tax_i(\mathcal{R}) = t^i(\mathcal{R}) = \left[\sum_{j \neq i} v_j(x^*(\mathcal{R})) \right] - \left[\sum_{j \neq i} v_j(x^*(\mathcal{R}_{-i})) \right]$$

If an agent does not change the allocation decision, then both sums are equal, so the transfer is 0. In general, the total tax will decrease as the number of agents increases. The more agents make their bids, the greater is the chance of any single agent to actually change the decision approaches zero.

Origin of the CTM

The CTM is a member of the family of Groves mechanisms [Groves 73]. All mechanisms of this class ensure that in quasilinear environments telling the truth is a dominant strategy for agents. Therefore, the agents do not need to waste effort in counterspeculating each others preferences. All Grove mechanisms have the general form

$$t^i(\theta) = \left[\sum_{j \neq i} v_j(all^*(\theta), \theta^j) \right] - \mathbf{Grove_tax}$$

¹⁷A proof that revealing true preferences is the dominant strategy can be found in [Rosenschein & Ephrati 91b].

Thus, the more the agent alters the outcome of the mechanism with his bid, the higher the Grove tax. But unfortunately, the Grove tax cannot equal zero but is a positive or even negative [Raiffa 82] amount that has to be donated to or funded from outside. It has been shown [Laffont & Maskin 80] that among the members of this decision mechanism family, CTM requires the least amount of tax to be paid.

Disadvantages

The Clarke taxing scheme incorporates an effective *incentive mechanism* for telling the truth in a group decision mechanism. Hence, the CTM could be used as an effective “preference revealer” in the domain of automated agents that reduces the need for explicit negotiation. Unfortunately—because it belongs to the family of Grove mechanisms—it has also two important disadvantages. Firstly, it can be manipulated by coalitions of agents—like the Vickrey auction. Currently there exist no techniques to deal with this. Secondly, the Clarke Tax must not be used to benefit those whose voting resulted in its charging. Instead it must either be wasted or spent outside the voting group, which would necessitate having distinct voting groups. The latter disadvantage is called the *tax waste problem*. The following example makes clear how the Clarke Tax is calculated and why the tax must be spent outside a voting group in order to prevent untruthful bidding.

Calculation of the Clarke Tax

Let us assume that the participating agents have the following valuations for a set of items:

| | Items and valuations | | |
|-------|----------------------|------------|------------|
| Agent | 1 | 2 | 3 |
| A | 800 | 700 | 200 |
| B | 700 | 500 | 400 |
| C | 400 | 400 | 400 |

First, each agent i states a report r_i - truthful or not - of his preferences, which defines \mathcal{R} . Then, the total score for $x^*(\mathcal{R})$ and for each schedule $x^*(\mathcal{R}_i)$ (i. e. resulting schedule if agent i had not voted) are calculated. The assignments of the allocation $x^*(\mathcal{R}^{true})$ in the case of unexceptionally truthful reports are highlighted by bold print in the table above.

| Allocation | Assignments | | | Valuations v_j | | | | $\sum_j v_j$ with | | | Tax |
|-------------------------|-------------|---|---|------------------|-----|-----|--------------|-------------------|------|------|-----|
| | A | B | C | A | B | C | $\sum_j v_j$ | j≠A | j≠B | j≠C | |
| $x^*(\mathcal{R})$ | 2 | 1 | 3 | 700 | 700 | 400 | 1800 | 1100 | 1100 | 1400 | |
| $x^*(\mathcal{R}_{-A})$ | - | 1 | 3 | - | 700 | 400 | 1100 | = | | | 0 |
| $x^*(\mathcal{R}_{-B})$ | 1 | - | 3 | 800 | - | 400 | 1200 | | > | | 100 |
| $x^*(\mathcal{R}_{-C})$ | 2 | 1 | - | 700 | 700 | - | 1400 | | | = | 0 |

When all the agents report their valuations truthfully, the group utility maximising allocation $x^*(\mathcal{R})$ is chosen. If A had not voted, B and C would still have been assigned the same items by the planner. If B had not voted, A would have been assigned item 1 instead of item 2. Such an assignment would increase A's utility by an amount of 100 while the utility of C would stay the same. Thus, agent B is fined with its Clarke tax $t^B(\mathcal{R}) = 100$ because he has affected the outcome by his vote by a magnitude of

$$t^B(\mathcal{R}) = \left[\sum_{j \neq B} v_j(x^*(\mathcal{R})) \right] - \left[\sum_{j \neq B} v_j(x^*(\mathcal{R}_{-i})) \right] = 1100 - 1200 = -100.$$

The other agents are not fined because even without their vote, the assignments of items to agents do not change.

Clarke Taxing Prevents Untruthful Bidding

Now let us assume that the valuations of the agents are not known publicly and agent A reports untruthfully a valuation of 1000 for item 1 because he aspires to be assigned to this item. Then the report table looks like

| | Items and valuations | | |
|-------|----------------------|------------|------------|
| Agent | 1 | 2 | 3 |
| A | 1000 | 700 | 200 |
| B | 700 | 500 | 400 |
| C | 400 | 400 | 400 |

and the following taxing table results:

| Allocation | Assignments | | | Valuations v_j | | | | $\sum_j v_j$ with | | | Tax |
|-------------------------|-------------|---|---|------------------|-----|-----|--------------|-------------------|------------|------------|-----|
| | A | B | C | A | B | C | $\sum_j v_j$ | $j \neq A$ | $j \neq B$ | $j \neq C$ | |
| $x^*(\mathcal{R})$ | 1 | 2 | 3 | 1000 | 500 | 400 | 1900 | 900 | 1400 | 1500 | |
| $x^*(\mathcal{R}_{-A})$ | - | 1 | 3 | - | 700 | 400 | 1100 | > | | | 200 |
| $x^*(\mathcal{R}_{-B})$ | 1 | - | 3 | 800 | - | 400 | 1200 | | < | | 0 |
| $x^*(\mathcal{R}_{-C})$ | 1 | 2 | - | 1000 | 500 | - | 1500 | | | = | 0 |

If agent A overbids his true valuations he finally pays with $t^A = 200$ a tax that is larger than his true preferences warrant. Because his true valuation for item 2 is 800 but he has to pay a tax of 200, he finally can only realise a surplus of 600. But this is 100 less than his valuation for item 2 that would have been assigned to him by truthful bidding. Thus, overbidding results in a loss of utility.

Likewise, if an agent underbids to change the outcome and save himself some tax, he will always lose utility. The saved tax will never compensate him for this loss.

Tax Waste Necessity: Similarly, if agent A knew that he would get a portion of the tax (as compensation for the losers or as an equally distributed share), he would also be tempted to overstate his valuation for item 1 (for example by 50) because this would raise B's tax (by 50 up to 150)—and consequently A's share of the tax.

| Allocation | Assignments | | | Valuations v_j | | | | $\sum_j v_j$ with | | | Tax |
|--------------------|-------------|---|---|------------------|-----|-----|--------------|-------------------|------------|------------|-----|
| | A | B | C | A | B | C | $\sum_j v_j$ | $j \neq A$ | $j \neq B$ | $j \neq C$ | |
| $x^*(\mathcal{R})$ | 2 | 1 | 3 | 700 | 700 | 400 | 1800 | 1100 | 1100 | 1400 | |

| | | | | | | | | | | | |
|-------------------------|---|---|---|------------|---|-----|-------------|--|---|--|------------|
| $x^*(\mathcal{R}_{-B})$ | 1 | - | 3 | 850 | - | 400 | 1250 | | > | | 150 |
|-------------------------|---|---|---|------------|---|-----|-------------|--|---|--|------------|

Ephrati and Rosenschein [Rosenstein & Ephrati 91a] suggest a solution to this problem. They propose to use the tax for the benefit of agents *outside* the voting group. They suggest partitioning the entire agent society into disjoint voting groups. When a voting group A_{voting} completes the assignment process, each taxed agent $a_{taxed} \in A_{voting}$ has to distribute his tax equally among $a_j \in A \setminus A_{voting}$ (or a randomly chosen a_j). Agents can be grouped based on common or conflicting interests. Agents who take part in an auction for an indivisible item could form a group because they compete for that special item and thus have conflicting goals.

Ephrati and Rosenschein [Rosenstein & Ephrati 91a] also suggest that the voting group solution to the tax waste problem might impose the need for an extensive bookkeeping mechanism, where each agent's tax debts are recorded. This would allow an agent to pay his debts by giving other agents support in achieving their goals later on. The fact that the tax must leave the group prevents the decision from being pareto optimal because any taxpayer could improve its own utility by the amount of the tax without changing the other agents' utilities.

Conclusion: Usability

The CTM saves each agent from the computational complexity of guessing what the others' preferences and strategies are, what the negotiation set is, and how it can be manipulated. This *simplicity of strategy* is highly desirable in the design of automated agents. Because the agents tell the truth out of their self-interest, the CTM meets both the simplicity and stability criteria (and also individual rationality, etc.).

Ephrati has investigated the usability of the Clarke tax mechanism for reaching a multiagent consensus [Rosenstein & Ephrati 93], e.g. in a multiagent planning problem [Ephrati & Rosenstein 93, Ephrati 93, Rosenstein & Ephrati 91b] and in a meeting-scheduling problem [Rosenstein et al. 91].

Clarke Taxing Results in a Vickrey Auction

When the Clarke Tax mechanism is used to allocate a single item, the solution that extracts truth is exactly the Vickrey auction. The winning agent i is pivotal to the decision because without i someone else would have won the auction (namely the agent with the second highest bid) and the resulting assignment would have been another.

According to the Clarke taxing scheme, agent i pays the value of the outcome to all the other agents (when i participates) minus the value to the other agents if i had not participated.

Suppose agent i bids *first_price*, agent j bids *second_price*, and all other agents bid below *second_price*.

Since there is only a single indivisible item to be allocated, the value to the other agents not assigned to it is 0. And without i , j 's value would have been its bid (the value to the other agents would again be 0). Thus, the payment of agent i is

$$p_i = 0 - \textit{second_price} = \left[\sum_{j \neq i} v_j(x^*(\mathcal{R})) \right] - \left[\sum_{j \neq i} v_j(x^*(\mathcal{R}_{-i})) \right] = t^i(\mathcal{R}).$$

The Generalised Vickrey Auction

1. Each bidding agent i reports a valuation function $r^i(\cdot)$ which may or may not be the truth.
2. The auctioneer calculates the allocation $x^*(\mathcal{R})$ that maximises the sum of the reported valuations.
3. The auctioneer also calculates the allocation $x^*(\mathcal{R}_{-i})$ that maximises the sum of the valuations subject to the constraint that *neither of the valuations nor resources of bidder i are taken into account*.
4. Each participating agent i receives the bundle x_i^* and receives the payment $t^i(\mathcal{R}) = \sum_{j \neq i} [r^j(x^*(\mathcal{R})) - r^j(x^*(\mathcal{R}_{-i}))]$ from the auctioneer.

End of the generalised Vickrey auction

According to this scheme, the final payoff to agent i in the GVA is given by

$$u_i(x^*(\mathcal{R})) = v_i(x^*(\mathcal{R})) + \left[\sum_{j \neq i} r^j(x^*(\mathcal{R})) - \sum_{j \neq i} r^j(x^*(\mathcal{R}_{-i})) \right]$$

4.7.2 The Generalised Vickrey Auction

The Generalised Vickrey Auction (GVA) [Varian 95] is a Clarke taxing procedure that is eligible for resource **reallocation** and not only for resource allocation as the original CTM. Insofar it works for a broader class of resource allocation problems. The GVA has principally the same taxing scheme as the CTM. But it differs from the CTM in that it can not only be used to allocate but also to **reallocate** resources because in the calculation of $x^*(\mathcal{R}_{-i})$ it is not only prohibited to take into account not only i 's valuations but additionally i 's resources.

Just as a bidder does not want to reveal his valuation to the seller, the participants in a **resource allocation problem** will in general not want to reveal their true utility functions. Varian proves that if the GVA mechanism is used it is in the interest of each participating agent i to report his true valuation function $r^i(\cdot) = \theta^i(\cdot)$ to the central resources allocating planner respective auctioneer [Varian 95]. The GVA proceeds as follows:

Suppose there are $i = 1, \dots, n$ agents who each get awarded $j = 0, \dots, m$ goods. Let x_i^j be the award of a piece of good j to agent i . Good 0 will denote "money"

and $x_i = (x_i^1, \dots, x_i^m)$ is the assigned bundle of goods to agent i . Each agent i holds some initial bundle \bar{x}_i and some initial amount of money \bar{x}_i^0 . A reasonable objective in allocating the goods among the consumers is to allocate them in a way that maximises the sum of the individual utilities.

Comparison between the GVA for Multiple Identical Items and the MIA

Assume that there are 4 identical items to assign and that there are 3 agents A, B and C that report the following valuations or bids respectively.

| | Number of identical items | | | |
|-------|---------------------------|-------|----|---|
| Agent | 1 | 2 | 3 | 4 |
| A | *55* | 70 | 80 | — |
| B | *30* | 80 | 90 | — |
| C | 25 | *100* | — | — |

Then the allocation that maximises social welfare respective global utility is the one marked with asterisks. Both the matrix auction and the GVA promote truthful bidding and result in the same allocation. But the settled prizes differ. According to the matrix auction rules, A and B have to pay an amount corresponding to the third-highest bid (25). C has to pay the second-highest price (80).

But applying the GVA, other payments p_i are settled that correspond to

$$p_i = \left| \sum_{j \neq i} r_j(x^*) - \sum_{j \neq i} r_j(x_{-i}^*) \right|, \text{ i.e.}$$

$$p_A = |t^A| = |130 - 180| = 50 > 25$$

$$p_B = |t^B| = |155 - 170| = 15 < 25$$

$$p_C = |t^C| = |85 - 150| = 65 < 80$$

The winning agents' payments in the GVA decrease proportional to their utility contribution to the utility maximising allocation. In contrary to this, in the matrix auction their contribution in maximising the global utility is not taken into account when prices are settled.

4.7.3 The Vickrey-Leonard Auction

The Vickrey-Leonard (VL) auction [Olson & Porter 94, Leonard 83] is a Clarke taxing procedure and as such a multi-item generalisation of the Vickrey auction

to obtain an outcome-efficient allocation. The VL auction has the theoretical property that truthful bidding is a dominant strategy.

Each agent i submits a vector r^i of *sealed* bids r_j^i for all items j being sold to the system. These are interpreted as an indicator of the agent’s welfare, i.e. in the equation system (A) for solving the assignment problem $\forall i, j \theta_j^i$ is replaced by r_j^i (submitted—possibly untruthful—bids in place of the true valuations). Thereafter a (smart market) computer algorithm solves A and thereby determines the assignment of agents to items (i.e. allocation) that maximises “the sum of the agents’ individual welfares”. Then the price—i.e. the Clarke tax—an agent has to pay for his assigned item is determined by solving either a dual program (D) to (A) or by calculating the impact of an additional unit of this item on the optimal assignment (see Price Settlement). Those calculated prices are also called *Vickrey prices*.

Allocation Procedure

The assignment algorithm finds the combination of assignments for which the total of the winning bids for all m units is the greatest (**one buyer can get at most one unit**). If two or more assignments yield the same maximum total, the assignment is chosen randomly.

Example:

| | Items and bids | | |
|-------|-------------------|------------|-------------------|
| Agent | 1 | 2 | 3 |
| A | $\widetilde{800}$ | 700 | 200 |
| B | 700 | 500 | 400 |
| C | 400 | 400 | $\widetilde{400}$ |

The fat table entries mark the resulting allocation. Even though Agent A bid the most on unit 1 he does not receive that unit, because the above determined assignment yields with 1800 the highest possible total of the winning bids.

Price Settlement

1. Price calculation by taking into account an additional unit of an item

1. Calculate the total of the bids for the allocated m units.
2. For each unit i do:
 - (a) Suppose that there is an extra unit i^{extra} available and therefore a total of $m + 1$ units to be sold.
 - (b) Calculate for this problem a new assignment that maximises the total bid amount^a.
3. Calculate the difference between the two totals. This difference is the price charged for unit i .

^aNotice: The new total bid amount is at least as high as the amount in step 1 because there are more combinations available, and all of the combinations previously available are still available.

Example: When an extra unit 1^{extra} is taken into account, a new allocation \tilde{x} results that raises the total bid amount to 1900 and in which agent A receives unit 1, agent B receives unit 1^{extra} and agent C receives unit 3:

| | Items and bids | | | |
|-------|-------------------|------------|-------------------|-------------------|
| Agent | 1 | 2 | 3 | 1^{extra} |
| A | $\widetilde{800}$ | 700 | 200 | 800 |
| B | 700 | 500 | 400 | $\widetilde{700}$ |
| C | 400 | 400 | $\widetilde{400}$ | 400 |

The charged price for unit 1 is 100 because it is determined by the difference between the two bid totals. Extra units of unit 2 or 3 available leave the allocation unchanged. Therefore the price of both units is zero.

Untruthful Bidding: Now let's assume that the valuations of the agents are not known publicly and agent A bids untruthfully 1000 for item 1. According to the notation used above the following allocation results:

| | Items and bids | | | |
|-------|--------------------------------------|------------|-------------------|-------------------|
| Agent | 1 | 2 | 3 | 1^{extra} |
| A | $\widetilde{1000}$ | 700 | 200 | 1000 |
| B | 700 | 500 | 400 | $\widetilde{700}$ |
| C | 400 | 400 | $\widetilde{400}$ | 400 |

The charged price for item 1 is then 200. The total welfare amount of the assignment determined without the item 1^{extra} is 1900. The assignment taking into account the item 1^{extra} yields a total welfare of 2100. Their difference amounts to 200 instead of 100 in the case above where agent A bid truthfully. Thus—because of his untruthfulness—agent A has to pay 100 currency units more for item 1.

2. Price calculation by solving a minimisation problem with constraints

Another possibility to determine prices in the VL auction is to solve the dual program (D) to (A) (see Section 4.1.1) in which we have to minimise the sum $\sum_{j \in M} p_j$ of the prices paid for the assigned items $j \in M$, i.e. to determine

$$\min_{p_j} \sum_{j \in M} p_j$$

with respect to the following set of constraints (the w'_i s are slack variables):

$$\begin{aligned} w_i + p_j &\geq \theta_j^i \quad \forall i \in N, \forall j \in M \\ \sum_{j \in M} p_j + \sum_{i \in N} w_i &= W \\ w_i, p_j &\geq 0, \quad \forall i \in N, \forall j \in M \end{aligned}$$

4.7.4 The VL Chit Mechanism

The VL chit mechanism has also been analysed by [Olson & Porter 94]. This mechanism—unlike the VL auction—does not require the use of monetary transfers to make an assignment. Here, the medium of exchange are called *chits* that have a determined value **only** in the context of the given assignment problem. The VL chit mechanism is implemented as follows:

Each agent has a predetermined chit budget and is allowed to bid a certain number of chits for any of the announced items. The chit bids are used in place of the monetary bids in the VL auction to determine an allocation.

The VL chit auction mechanism without monetary payments does not have a dominant strategy equilibrium but it does have Bayes-Nash equilibria. It does also not provide outcomes that are *pareto optimal* for all profiles θ . Therefore it is worse than the VL auction with regard to the allocative efficiency [Olson & Porter 94].

According to [Olson 92] Bayes-Nash equilibria do not elicit more information about an agent's relative valuations than just ordinal rankings.

4.8 The DGS Auction

The DGS [Roth & Sotomayor 91] auction is a multi-item auction mechanism which assures that truth telling is a dominant strategy for each bidder. It is a generalisation of both the progressive (or English) *and* the Vickrey auction. The DGS auction mechanism is applicable for cases in which a set of items is to be distributed among a set of agents and **no agent is interested in more than one item**. It produces the minimum price equilibrium and therefore is buyer-optimal.

It is assumed that agents have a *demand function*, i.e. that—given a vector of prices for the items—they are able to specify which set of items (possibly empty) they wish to buy.

Demange, Gayle and Sotomayor [Demange et al. 86] proposed two variations, namely the exact DGS auction and the approximate DGS auction.

When there is only one item to allocate, both variations can be implemented by the following process that resembles the EA:

The auctioneer begins with the item's reservation price. Then each bidder announces whether he wants to buy the item at that price. The price is increased by a fixed amount as long as more than one bidder demands the item. If only one bidder demands it that bidder is awarded the item at the announced price and the auction for that item ends.

4.8.1 The Exact DGS Auction Mechanism for Multiple Items

When a set of heterogeneous items is to be allocated, the DGS auction proceeds as follows (see also [Roth & Sotomayor 91], pp. 210-211). The DGS auction algorithm is a version of the Hungarian algorithm for the assignment problem [Kuhn 55].

It is assumed that at each stage of the auction each bidder demands *all* items that maximise his utility (surplus) at the current prices. That is, **a bidder expresses interest in more than one item if and only if they both maximise his surplus at the given prices**.

| |
|------------------------------|
| The exact DGS auction |
|------------------------------|

- **Initial Step ($t = 0$):**

1. The auctioneer begins by announcing an initial price vector p^0 , equal to the vector of the sellers' reservation prices s .
2. Then each bidder announces which item or items j are in his demand set at price p_j^0 .

• **General step** ($t > 0$):

1. Given the demand selections (or bids) at prices p^{t-1} , an algorithm determines if there are *overdemanded*¹⁸ sets of items.
2. The planner locates (with an algorithm [Gale 60]) a *minimal overdemanded set (MOS)*, i. e. an overdemanded set with the property that none of its proper subsets is overdemanded.
3. If there is no MOS the planner can assign each bidder an item he demands and the auction terminates.
4. If he can locate a MOS, he raises the price for each item in it by one unit, what defines p^{t+1} .
5. After that the planner asks the bidders for their demand sets at price p^{t+1} .

• **Termination:**

The auction terminates if there is no more overdemanded set of items. Then, the agents are assigned to their selected items j at current prices p_j^t .

| |
|-------------------------------------|
| End of the exact DGS auction |
|-------------------------------------|

Olson and Porter [Olson & Porter 94] propose the following efficiency increasing extension to this process:

1. They suggest that if an agent selects a *non-overdemanded* item at an iteration, he is committed to select that item still at the next iteration (i.e., an agent cannot withdraw on selections if the price of those selections does not increase). Then he would be automatically assigned to that item and excluded from the auctioning of the remaining items.
2. at the end of the process unassigned items should be randomly allocated among the last unassigned bidders who placed a bid on the items at a previous iteration.

The realisation of the second proposal is outlined in the following:

¹⁸A set of items is called *overdemanded* if the number of bidders demanding only items in this set is greater than the number of items in the set.

| Demanded and assigned items at prices p^t and p^{t+1} : | | | | | | | |
|---|-------|-----|-----|-----------|------------|-----|------------|
| Price | p^t | | | p^{t+1} | | | |
| Item | 1 | 2 | 3 | 1 | 2 | 3 | |
| Agent | A | 200 | 200 | 200 | 300 | 300 | (—) |
| | B | | 200 | 200 | | 300 | 300 |
| | C | 200 | 200 | | (—) | (—) | |

According to the DGS auction rules agents A and B are **assigned** to items 1 and 3 (emphasised by bold printing). Hence they drop out for the assignment of item 2. Agent C has reneged his selection of item 2 after its price raised to p_2^{t+1} . But as unassigned agent he gets the unassigned item 2 at the price p_2^t at which he demanded it in the last cycle.

However, it seems that the exact DGS auction mechanism is difficult to implement in realistic situations. Firstly, the communication cost are quite high. Secondly and more important, in a DGS auction the agents are required to be quite precise in their responses to changing prices. They have to report at each stage t *all* their surplus maximising items at prices p^t . An agent cannot switch in one step from one item to a different one but must demand his surplus maximising items of the previous stage plus perhaps some others.

4.8.2 The Approximate DGS Auction Mechanism

In contrast to the exact DGS auction mechanism the approximate variant does not require the bidders to decide in advance exactly what their bidding behaviour will be. Instead, at each stage they can make use of present and past stages of the auction to decide their next bids. Insofar this variant is more appealing to the bidders. It proceeds as follows:

| |
|-------------------------------------|
| The approximate DGS auction: |
|-------------------------------------|

- **Initial step** ($t = 0$):
 1. An initial price vector p^0 (of the sellers' reservation prices) is announced by the auctioneer.
 2. Each bidder announces his demand for items at price p^0 , thereby *committing* himself to possibly buying the item at the announced price.
- **General step** ($t > 0$):

At each step t , some subset of bidders is committed to some subset of items at some set of prices.

Any *uncommitted* bidder chooses one out of three alternative actions; he may

- (a) bid for some unassigned item j
 \implies he gets committed to that item at its initial price p_j^0 .
- (b) bid for an assigned item j
 \implies he gets committed to that item at price $p_j^t + \delta$ ($\delta =$ fixed amount) and the bidder to whom j was previously assigned to gets uncommitted.
- (c) drop out of the bidding.

- **Termination:**

The approximate DGS auction terminates when there are no more uncommitted bidders in the auction, at which point each committed bidder buys the item assigned to him at its current price.

| |
|---|
| End of the approximate DGS auction |
|---|

The outcome of the approximate DGS variant may depend on the order in which agents bid. [Demange et al. 86] show that the final prices for the items in the approximate DGS auction differ from the minimal equilibrium prices obtained in the exact DGS auction by at most $k \cdot \delta$ currency units, where k is the minimum of the number of items and bidders. [Demange et al. 86] also show that one can come arbitrarily close to the minimum equilibrium prices of the exact auction by making the bid increasing rate δ sufficiently small.

Summary:

The DGS auction generalises the VA to multiple items and produces an assignment that corresponds to a bipartite matching graph between a set of agents and a set of items. That is, each agent receives at most one item but the DGS auction does not assign necessarily all items to agents.

As in the Vickrey second-price auction for a single item, the price a winning bidder pays is not determined by the valuations he states but the valuations of its competitors.

The experiments of Olson and Porter [Olson & Porter 94] showed that the DGS auction outperforms the VL auction concerning the allocative efficiency and the seller surplus.

4.9 The Escalating-Bid Auction

The Escalating-Bid auction (EBA) is a low-overhead auction procedure for allocating resources that are naturally divided into time slices. It uses the Escalator Algorithm [Drexler & Miller 88] and is a *mixed* auction procedure because due to this algorithm it combines properties of the EA and the VA. The EBA was developed as a market-like mechanism for allocating processor time to computational tasks in a decentralised fashion, only guided by local knowledge and decisions concerning task preferences. The auction procedure can be used to schedule any resource allocated on a time-slice by time-slice basis (e.g. communication channels, airport landing rights, etc.).

The auction for a special resource can be viewed as a virtual auction house with multiple escalators in it. An agent enters the auction and begins bidding for the resource required for its task by placing an initial bid on one of the escalators. The size of his bid determines its initial height on the chosen escalator. Each escalator rises at a different *escalation rate* (er) (measured in currency units per time unit), raising its bids at that rate. Fixed bids are held by a special stationary escalator ($er = 0$). An escalator is realised by a priority queue data structure that makes the currently highest bid, i.e. the bid with the highest current priority, readily available. The resource is always awarded to the highest bid.

The height of the initial bid and the escalation rate of the chosen escalator consequently reflect the priority of a task. Given fluctuating prices, faster escalation rates will generally result in higher average cost per time slice. Bids placed on fast escalators will pay a higher than average cost.

Assume a stable price for time-slices, equal to \bar{p} . Then a zero initial bid placed on an escalator with escalation rate er will receive a time fraction of the demanded resource roughly equal to er/\bar{p} .

Price settlement

The Escalating-Bid auction combines characteristics of the English and the second-price, sealed-bid auction. The price settlement is done according to the Vickrey principle:

At the beginning of each time slice, the auctioneer examines the top bid on each escalator in the auction house. The highest bid among them wins the resource for this time slice and is charged with the amount of the second-highest bid. This corresponds to the highest bid among all the other top bids and the follower of the winning bid. After deleting the winning bid its follower gets to the top of the escalator.




Figure 4.5: The escalating-bid auction

Possible Extensions of the Escalating-Bid Procedure

The number of parameters that are adjustable by the agents could be increased. Especially, the initial bid could contain—besides initial height and maximal height of a bid—a kind of bidding strategy function. For example, time horizons for the period of validity of bids could be implemented in this function. It is conceivable to give an agent the option to update his bidding strategy in certain time intervals when its local plan and thus its valuation for the resource has changed.

Analysis

The escalating-bid auction seems well-suited to resource scheduling problems. It avoids the *sleeping-bidder problem* of the classical progressive auctions (EA, DA) because the agents do not need to instantaneously keep track of the auctioning process. Also, it ensures that a resource can be awarded or assigned at any time—crucial, when the commodity to be sold is quite perishable. Hence it avoids the *starvation problem*¹⁹ of the classical sealed-bid auctions (VA, FP). Since the demand for processor time is apt to fluctuate rapidly, proper incentives require rapidly fluctuating prices. Fluctuating prices can provide incentives for delaying less urgent tasks, leveling loads, and so forth.

¹⁹A perishable resource is available but for some reasons there exists no bid for it and the resource decays/keeps unused.

Chapter 5

Application Example: The Transportation Domain

Over the past few years, the scheduling of transportation orders had been established as an important field of application for Distributed Artificial Intelligence both from an academic and from a practical perspective (cf. [Wellman 92, Sandholm 93, Fischer et al. 93b, Fischer et al. 94]). It offers interesting complexity properties, an inherent distribution of knowledge and control, natural possibilities to study coordination and cooperation, and finally, there is a considerable economic interest in obtaining good solutions for these kinds of problems.

In the following sections we give a brief description of the transportation domain and the two systems *COMoves* and MAS-MARS. Subsequently we characterise the transportation domain using the definitions of Rosenschein and Zlotkin as a guideline [Zlotkin & Rosenschein 93, Rosenschein & Zlotkin 94a]. It turns out that we have to modify Definition 11 of Chapter 3 slightly to gather the domain. Furthermore it turns out to be useful to divide the class of task-oriented domains into two classes: cooperative and competitive task-oriented domains. The transportation domain provides examples for both categories of task-oriented domains. In Section 5.4 we introduce some concepts to model market prices, costs, and utility, that provides the foundations for trading transportation tasks between shipping companies.

Section 5.5 discusses the practicability of the auction-based negotiation mechanisms explained in Chapter 4 in the context of the shipping domain.

5.1 The Domain of Application

This section provides an overview of the transportation domain. Important previous results are summarised and related work is discussed.

5.1.1 Modelling the Transportation Domain as a Multi-agent System

The domain of application is the planning and scheduling of transportation orders which is done in everyday life by human dispatchers in transportation companies. Many of the problems which must be solved in this area, such as the Travelling Salesman and related scheduling problems, are known to be NP-hard. Moreover, not only since *just-in-time* production has come up, planning must be performed under a high degree of uncertainty and incompleteness, and it is highly dynamic. In reality these problems are far from being solved.

Cooperation and coordination, both inside a single transportation company and between several companies, seem to be two very important processes that may help to overcome the problems sketched above. Indeed, they are of increasing importance even in the highly competitive transportation business of today. Using the MAS-MARS system, several patterns of cooperation such as the announcement of unbooked legs, order broking, and different strategies for information exchanges have been experimentally evaluated [Kuhn et al. 93, Fischer et al. 93a]. The multiagent approach to solve these problems is to employ two basic types of agents corresponding to the physical entities in the domain: *transportation companies* and *trucks*. Companies can communicate with their trucks and with each other. The user can dynamically dedicate transportation orders to specific companies. Looking upon trucks as agents allows us to delegate problem-solving skills to them (such as route-planning and local plan optimisation). The *shipping company* agent has to allocate orders to its trucks, while trying to satisfy the constraints provided by the user as well as local optimality criteria (costs). A company also may decide to cooperate with another company instead of having an order executed by its own trucks. Each *truck* agent is associated with a particular shipping company from which it receives orders of the form "Load amount a_1 of good g_1 at location l_1 and transport it to location l_2 while satisfying time constraints $\{c_{t_1}, \dots, c_{t_k}\}$ ". More formally the setting can be described as follows:

$\mathcal{S} = \{S_1, \dots, S_l\}, l \in \mathbb{N}$: is the set of shipping company agents.

$\mathcal{L} = \{\mathcal{L}_{S_1}, \dots, \mathcal{L}_{S_l}\} = \{\{L_1^{S_1}, \dots, L_{m_1}^{S_1}\}, \dots, \{L_1^{S_l}, \dots, L_{m_l}^{S_l}\}\}, m_i \in \mathbb{N}$: is the set of truck agents¹, where the set \mathcal{L}_{S_i} specifies the truck agents of shipping company S_i . Each truck has only a limited capacity.

$\mathcal{O} = \{o_1, \dots, o_n\}, n \in \mathbb{N}$: is a set of orders. Each order is specified by a tuple $(a, g, l_1, l_2, c_{t_1}, \dots, c_{t_k})$ where a is the amount, g the type of the good, l_1 the source, l_2 the destination, and c_{t_1}, \dots, c_{t_k} are time constraints. The orders are announced to individual shipping companies at a specific point in time.

¹To avoid name clashes, we use \mathcal{L} for the set of trucks, derived from the British word *lorry*.

Only in a test setting the whole set of orders will be specified in advance. In a system running in the real world the set of orders will dynamically increase.

Within the MAS group at the DFKI two implementations of the system were developed: MAS-MARS and *COMoves*.

MAS-MARS is implemented in DFKI-Oz. The essential goal of MAS-MARS is to produce good results for the transportation scheduling problems.

COMoves is the reimplementaion of MAS-MARS using the multiagent development shell DASEDIS. The main focus of *COMoves* was the development of market and auction models and the analysis of dynamic aspects of the execution of tour plans.

5.1.2 Related Work and Previous Results

A basic research area dealing with these problems is Operations Research (OR). An overview of the work done to solve the scheduling problem for the transportation domain is given in [Bodin et al. 83]. For a more recent approach see [Psaraftis 88]. It turns out that there are problems with these approaches when the number of constraints to deal with grows or when real-time response of the system is required. This is the case if such a system is to be used in order to support a dispatcher who has to tell customers an estimated cost of a delivery. For this class of problems knowledge-based approaches as the one by Bagchi and Nag [Bagchi & Nag 91] are often used. Within their system, global optimisation is reduced to assigning a new shipment to a contract with minimal incremental cost caused by that insertion. This is based on a result of Psaraftis [Psaraftis 88] who shows that in a dynamic scheduling environment, global minimisation over a period of time is best achieved by minimising the incremental cost of each assignment.

A distributed OR-based approach is that of Bachem et al. [Bachem et al. 92]. They present a parallel improvement heuristic for solving vehicle routing problems with side constraints is presented. The solution to this problem is constructed using a procedure called “Simulated Trading” where each tour of a truck can be assigned to a specific processor of a parallel computer. A further example for distributed OR approaches can be found in [Falk et al. 93]. They pursue an approach integrating knowledge-based mechanisms and OR algorithms to model tramp agent companies. Compared to our modelling, this approach represents a specific instance of our domain, namely a single company which is geographically distributed. Thus, the dispatching agents are willing to exchange all the information (in this case, the complete route plans) in a cooperation process.

[Zlotkin & Rosenschein 93, Rosenschein & Zlotkin 94a] gave a decision theoretic classification of domains and presented general results for different classes: task-oriented domains, state-oriented domains, and worth-oriented domains. For the class of task-oriented domains they introduce the *delivery domain* and the *post-men domain* which are both quite similar to the transportation domain. However, on the one hand we will see that the analysis of the transportation domain will lead us to different definitions from the ones given by Rosenschein and Zlotkin. On the other hand, Rosenschein and Zlotkin concentrate on the decision theoretic perspective in these application domains whereas our work tries to join the work done on solving complex scheduling problems and the results developed in decision theory.

To give an overview of the current state of the art of the MAS-MARS system, we will briefly recall major previous results. In [Kuhn et al. 93], the architecture of the basic simulation system was described and a negotiation protocol which allows companies to avoid unbooked legs by offering specific orders to other companies was specified. [Fischer et al. 93b] contains an empirical analysis of how different types of cooperation (namely *vertical*, *horizontal*, and *enhanced cooperation*) influence the performance of the system as a whole. The major result was that by employing (horizontal) cooperation among shipping companies in a straightforward manner, cost savings of about 8% could be achieved. Finally, in [Fischer et al. 94], the results of a benchmark test [Desrochers et al. 92] are described in which we compared our solution for a specific problem, the *static vehicle routing problem with time windows*, to the solutions found by several heuristic OR algorithms. From this experiments we learned that the MAS-MARS system produced solutions close to the optimum of a quality that is generally comparable to heuristic OR approaches in other cases. Moreover, MAS-MARS is able to solve dynamic scheduling problems which are beyond the scope of the OR algorithms used in the benchmark.

In [Pischel et al. 95] the CoMOVES system, which is a reimplementaion of the MAS-MARS system, is described. The results of the project *COMoves-II* are presented in [Siekmann et al. 96].

5.2 Analysis of the Transportation Domain

In this Section, the transportation domain is analysed as a task-oriented domain as proposed by [Zlotkin & Rosenschein 93].

5.2.1 Definition and Basic Properties

Tasks are the main focus of the problem solving and negotiation process in the transportation domain; therefore it is intuitive to classify it as a task-oriented domain (TOD). Unfortunately our domain does not quite match the original

definition of a TOD as given in Definition 11 of Chapter 3: It is easy to see that a single truck with limited loading capacity is not able to execute arbitrary many transportation tasks. Furthermore, there exist trucks with different abilities such that the cost function is not the same for all truck agents.

Additionally, an important feature of the transportation domain is that a shipping company will be paid money for each task it performs. For this reason, we introduce the notion of a price for TODs. At a first glance, this seems to be a technical detail; nonetheless, as we will see below, it allows an important further classification of TODs. One could argue that Rosenschein and Zlotkin introduced the notion of price (worth) in worth-oriented domains (WOD) which is in fact a super-class of the task-oriented domains. However, in WODs, worth is assigned to goal states rather than to single tasks, which is less desirable for the transportation domain.

We continue with a redefinition of the task oriented domains. The main difference to Definition 11 is that the cost function distinguishes between the agents and that it is allowed to assign infinite costs to a set of tasks which means that it is impossible to perform it for that agent.

Definition 31 A Task Oriented Domain (TOD) is a tuple $\Omega = \langle \mathcal{T}, \mathcal{A}, \mathcal{C} \rangle$ where:

1. \mathcal{T} is the set of all possible tasks;
2. $\mathcal{A} = \{A_1, \dots, A_n\}$ is an ordered list of agents;
3. $\mathcal{C} = \{c_1, \dots, c_n\}$ is a set of functions $c_i : 2^{\mathcal{T}} \rightarrow \mathbb{R}^+ \cup \{\infty\}$ such that for each finite set of tasks $X \subset \mathcal{T}$, $c_i(X)$ is the cost for agent A_i to execute all the tasks in X . All c_i are monotonic, i.e., for any two finite subsets $X \subset Y \subset \mathcal{T}$, $c_i(X) \leq c_i(Y)$.
4. $c_i(\emptyset) = F_i$ are the fix costs of agent A_i . For simplicity we assume $F_i = 0 \forall i$.

To be able to assign a price to individual tasks we extend Definition 31 by a price function p :

Definition 32 Let $\Omega = \langle \mathcal{T}, \mathcal{A}, \mathcal{B} \rangle$ be a TOD. A price function for Ω is a function $p : 2^{\mathcal{T}} \rightarrow \mathbb{R}^+$. For each finite subset $\mathcal{X} \subseteq \mathcal{T}$ holds:

$$p(\mathcal{X}) = \begin{cases} \sum_{t \in \mathcal{X}} \hat{p}(t) & \text{if } \mathcal{X} \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

where $\hat{p}(t)$ is the price to be paid to an agent for task t . Throughout this chapter, we will write $\Omega = \langle \mathcal{T}, \mathcal{A}, \mathcal{B}, p \rangle$ for a TOD Ω with price function p .

The utility an agent A_i obtains for performing any given subset of tasks $\mathcal{X} \subseteq \mathcal{T}$ is defined by:

$$tut_i(\mathcal{X}) = p(\mathcal{X}) - c_i(\mathcal{X}).$$

Agents in TODs defined according to Definition 31 will always try to avoid executing tasks because saving costs increases their utility. Agents in the TODs defined according to definition 32 have a desire to execute tasks (in our application domain: transportation orders), because doing so will increase their utility if the price exceeds the costs. The real trick for the agents is to execute tasks at minimal costs.

Rosenschein and Zlotkin characterise TODs specifying several attributes they can have, the weakest of which is *subadditivity*. We refine the definition for our purposes.

Definition 33 (Subadditivity) *Let $\Omega = \langle \mathcal{T}, \mathcal{A}, \mathcal{B}, p \rangle$ be a TOD. Ω is called:*

1. *locally subadditive if for all finite sets of tasks $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{T}$ and every agent $A_i \in \mathcal{A}$, we have*

$$c_i(\mathcal{X} \cup \mathcal{Y}) \leq c_i(\mathcal{X}) + c_i(\mathcal{Y}).$$

2. *globally subadditive if for all finite sets of tasks $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{T}$ and for any agents $A_i, A_j, A_k \in \mathcal{A}$, we have*

$$c_i(\mathcal{X} \cup \mathcal{Y}) \leq c_j(\mathcal{X}) + c_k(\mathcal{Y}).$$

For the price function we require that

$$p(\mathcal{X} \cup \mathcal{Y}) = p(\mathcal{X}) + p(\mathcal{Y}).$$

Theorem 2 *If no time constraints are specified for the transportation orders,*

1. *the transportation domain is a locally subadditive TOD,*
2. *the transportation domain is not a globally subadditive TOD.*

Proof Idea: For 1., the key observation is that an agent can benefit from joining two round trips into a single round trip. Note that this assertion remains true even if capacity constraints (i.e. restrictions of the amount a truck agent can carry at a time) are considered. However, in this case the probability that an agent really benefits from joining two round trips decreases. For 2., it is easy to construct an example in which two trucks located in two geographically distinct locations can execute two sets of orders at less costs than any of them could do on its own. \square

Theorem 3 *If time constraints are specified for the transportation orders, the transportation domain is not locally subadditive.*




Figure 5.1: If time constraints are specified, the transportation domain is not locally subadditive

Proof: Because this result is less obvious than the previous one we give the proof via a more detailed example shown in Figure 5.1. In the example, four orders are specified. The due date (dda) of an order is the latest point in time the truck is allowed to finish the order. Let t_{cur} denote the current time. The four orders are given by:

$$\begin{aligned}
 o_1: & A \rightarrow B, \quad dda_{o_1} = t_{cur} + \text{time-needed}(A \rightarrow B) \\
 o_2: & C \rightarrow D, \quad dda_{o_2} = t_{cur} + \text{time-needed}(A \rightarrow B \rightarrow C \rightarrow D) \\
 o_3: & E \rightarrow F, \quad dda_{o_3} = t_{cur} + \text{time-needed}(A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow F) \\
 o_4: & G \rightarrow H, \quad dda_{o_4} > t_{cur} + \text{time-needed}(A \rightarrow B \rightarrow C \rightarrow E \rightarrow F \rightarrow G \rightarrow H)
 \end{aligned}$$

The orders all together do not exceed the capacity of one truck and no further time constraints are specified. This means that the trucks do not have to worry about being early in a specific city. All they have to care about is not being late. Therefore, the best routes for two trucks starting in city A are:

$$A \rightarrow B \rightarrow E \rightarrow F \rightarrow A \text{ and } A \rightarrow C \rightarrow D \rightarrow G \rightarrow H \rightarrow A.$$

The only possible route satisfying the time constraints for one truck executing all the orders is:

$$A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow F \rightarrow G \rightarrow H \rightarrow A.$$

For these routes we have:

$$\begin{aligned}
 & c(A \rightarrow B \rightarrow E \rightarrow F \rightarrow A) + c(A \rightarrow C \rightarrow D \rightarrow G \rightarrow H \rightarrow A) \\
 & < c(A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow F \rightarrow G \rightarrow H \rightarrow A)
 \end{aligned}$$

This is due to the fact that the following inequations hold for the example:

$$\begin{aligned} c(A \rightarrow C) + c(B \rightarrow E) + c(D \rightarrow G) &< c(B \rightarrow C) \\ c(F \rightarrow A) &< c(D \rightarrow E) + c(F \rightarrow G) \end{aligned}$$

If the distances $B \rightarrow C$, $D \rightarrow E$, and $F \rightarrow G$ are increased while keeping the distances $A \rightarrow C$, $B \rightarrow E$, and $D \rightarrow G$ constant, the difference in costs between the two solutions can grow without any bound. \square

It follows from these two results that our setting is in the range of *general to subadditive* TODs according to the definition in [Rosenschein & Zlotkin 94a], depending on the specification of the orders. We will come back to this point later.

5.2.2 Encounters and Deals in the Transportation Domain

In [Rosenschein & Zlotkin 94a] an *encounter* is an initial distribution of tasks to the agents and a *deal* is a redistribution of the tasks as an negotiation outcome. Because of the dynamic nature of the transportation domain and the necessity of regarding side payings we define as in [Fischer & Müller 96] an encounter as a distribution of tasks and for each agent the side paying he receives from or owes to the other agents. So an encounter is not only the initial distribution of tasks but can also be a deal, a redistribution the agents agreed on.

Definition 34 *An encounter within a TOD $\langle \mathcal{T}, \mathcal{A}, \mathcal{B}, p \rangle$ is a totally ordered list (T_1, T_2, \dots, T_n) such that for all $k \in \{1, \dots, n\}$, $T_k \in 2^{\mathcal{T}} \times \mathbb{R}$. We define two access functions to the elements of an encounter: $tasks(T_k) \subseteq \mathcal{T}$ is a finite set of tasks that A_k needs to achieve and $sp(T_k)$ the side-payment agent A_k will receive. $sp(T_k)$ can be negative. However, $\sum_{i=1}^n sp(T_i) = 0$ must hold.*

The global order scheduling problem of finding good solutions for a given situation can be decomposed into two steps which can be analysed separately. The first question is how orders should be distributed among shipping companies. The second one is how orders can be allocated within one shipping company among its set of trucks. Therefore, in MAS-MARS encounters are defined at two levels. Firstly, at the layer of the shipping companies an encounter specifies which set of tasks is assigned to which shipping company. Secondly, on the layer of the trucks in each shipping company an encounter specifies how the set of tasks which is currently assigned to a specific shipping company is distributed among the trucks of this shipping company. If we assume that there are $n \in \mathbb{N}$ shipping companies, the distribution of the current task is specified by $n + 1$ encounters².

²One within each company plus one among the companies.

In fact, this allocation is the solution to a very complex scheduling problem. In MAS-MARS this solution is initially computed by a contract-net-like negotiation mechanism between the shipping companies and their trucks. This initial solution is further optimised using negotiation mechanisms based on the notion of deals [Fischer & Müller 95].

Definition 35 Let $\Omega = \langle \mathcal{T}, \{A_1, \dots, A_n\}, \mathcal{B}, p \rangle$ be an n -agent TOD; let (T_1, \dots, T_n) be an encounter within Ω . A deal is a redistribution of tasks among agents. It is an encounter (D_1, \dots, D_n) such that $tasks(D_1), \dots, tasks(D_n) \subseteq \mathcal{T}$, and that $\bigcup_{i=1}^n tasks(D_i) = \bigcup_{i=1}^n tasks(T_i)$. The semantics of such a deal is that each agent A_k commits itself to executing all tasks in $tasks(D_k)$.

Given an encounter (T_1, \dots, T_n) within a TOD $\langle \mathcal{T}, \{A_1, \dots, A_n\}, c, p \rangle$, we have the following:

1. For any deal δ , $1 \leq k \leq n$ the deal utility for agent A_k is defined by:

$$dut_k(\delta) = tut(tasks(\delta_k)) + sp(\delta_k) - tut(tasks(T_k)) - sp(T_k).$$

2. The initial encounter $\Theta \equiv (T_1, \dots, T_n)$ is called the conflict deal and remains unchanged if no agreement can be reached.

The variable sp in an encounter is only used in negotiation processes among shipping companies. In the dynamic case, where orders are announced to the system at random points in time, the system starts with an encounter where each of the task sets in the encounter is the empty set and all balances are equal to 0. The side-payments may change within negotiation processes when an agent buys an order from another agent. Their purpose is to express how much the agent has gained from the encounter.

Definition 36 For vectors $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ and $\beta = (\beta_1, \beta_2, \dots, \beta_n)$, we say that α dominates β and write $\alpha \succ \beta$ iff $\forall k (\alpha_k \geq \beta_k)$, and $\exists l (\alpha_l > \beta_l)$. We say that α dominates β weakly and write $\alpha \succeq \beta$ iff $\forall k (\alpha_k \geq \beta_k)$.

We adopt the notion of dominance for deals: for deals δ and δ' , we say that δ dominates δ' , written as

$$\delta \succ \delta' \quad \text{iff} \quad (dut_1(\delta), \dots, dut_n(\delta)) \succ (dut_1(\delta'), \dots, dut_n(\delta')).$$

We say that δ dominates δ' weakly, and write

$$\delta \succeq \delta' \quad \text{iff} \quad (dut_1(\delta), \dots, dut_n(\delta)) \succeq (dut_1(\delta'), \dots, dut_n(\delta')).$$

Deal δ is individually rational for all agents if $\delta \succeq \Theta$ and *pareto optimal* if there is no other deal δ' such that $\delta' \succ \delta$. One possibility to find out if a given encounter

is pareto optimal, is to prove it using a branch and bound algorithm. However, in general this procedure is much too time-consuming. Therefore, we introduce the notion of *weak pareto optimality*: a deal δ is *weakly pareto optimal* if no agent is able to compute a deal δ' which dominates δ within a specified time limit. The set of all deals which are individual rational and weakly pareto optimal is called negotiation set.

Based on the notion of weak pareto optimality, it is possible to define an anytime algorithm [Boddy & Dean 94] for schedule optimisation. If we assume that in the transportation domain the set of tasks \mathcal{T} is given by all tasks which are present at a specific point in time, the set of all possible encounters has an enormous size. Starting with some encounter (T_1, \dots, T_n) which is not weakly pareto optimal in general, we have defined negotiation strategies which will lead to weakly pareto optimal solutions [Fischer & Müller 95].

5.2.3 Cooperative vs Competitive TODs

We are now at the point to refine the definition of a TOD, leading to two subclasses that each describe important cases of TODs with different properties.

Definition 37 A TOD $\Omega = \langle \mathcal{T}, \mathcal{A}, \mathcal{B}, p \rangle$ is called:

1. a competitive TOD if an agent A_k only accepts a deal δ if: $dut_k(\delta) \geq 0$
2. a cooperative TOD if the agents accept a deal δ if: $\sum_{i=1}^n dut(D_i) \geq 0$

According to this definition, task allocation within one shipping company describes a cooperative TOD because trucks of one shipping company switch orders even if one of them obtains less utility by the deal than it had before. In this situation it is possible to define negotiation protocols that result in surprisingly good solutions [Bachem et al. 92]. Results from an analysis of cooperative TODs have been published in [Fischer et al. 94, Fischer & Müller 95]. For the rest of this section, we concentrate on the competitive setting which describes task allocation among different shipping companies.

5.3 Lying in the Transportation Domain

The way the utility is divided among the agents is not the only difference between the cooperative and competitive setting. In a competitive setting it is not very likely that agents always tell the truth. It has been shown in [Zlotkin & Rosenschein 93, Rosenschein & Zlotkin 94a] that lying in subadditive and general TODs may be beneficial.




Figure 5.2: Example: lies may be beneficial

5.3.1 Phantom Tasks and Hidden Tasks

From the previous considerations we learned that the transportation domain is on the verge between general and subadditive TODs. However, lying has different effects in the two settings, as in the TODs originally defined by Rosenschein and Zlotkin agents try to avoid the execution of tasks whereas in our case agents are eager to execute tasks.

Theorem 4 *In the transportation domain phantom lies may be beneficial.*

The proof of Theorem 4 is accomplished by providing an example: Assume that there is a shipping company S_A located in city A and a shipping company S_B located in city B (see Figure 5.2). Let o_1 be an order of $20t$ from C to D offered to S_A and o_3 an order of $20t$ from B to D offered to S_B . The distances the trucks have to go is computed as the Euclidean distance between the points of the cities in the pictures. We assume that the costs are proportional to the distances. The utilities for this initial encounter are

$$tut_{S_A}(o_1) = p(o_1) - c(ACDA)$$

$$tut_{S_B}(o_3) = p(o_3) - c(BDB)$$

One can see that S_B can perform o_1 with less costs than S_A . The cost savings from this redistribution are

$$\Delta = c(ACDA) + c(BDB) - c(BCDB).$$

Hence a product maximising mechanism would exactly halve this utility and choose the deal that specifies a side-paying of

$$\sigma = tut_{S_A}(o_1) + \frac{\Delta}{2}$$

the utility S_A would have reached on its own plus 50 percent of the common savings as a fair outcome.

Now consider the case that S_A invents the phantom order o_2 from A to D . It would still be rational for S_A to sell the order o_1 to S_B . Now the apparent cost savings are

$$\Delta' = c(ACDA) + c(BDB) - c(ADA) - c(BCDB).$$

Therefore, a PMM will choose

$$\sigma' = tut_{S_A}(o_1, o_2) - tut_{S_A}(o_2) + \frac{\Delta'}{2}$$

as sidepaying for the task o_1 . Since

$$\begin{aligned} \sigma' - \sigma &= tut_{S_A}(o_1, o_2) - tut_{S_A}(o_2) + \frac{\Delta'}{2} - (tut_{S_A}(o_1) + \frac{\Delta}{2}) \\ &= p(o_1) + p(o_2) - c(ACDA) - (p(o_2) - c(ADA)) \\ &\quad - (p(o_1) - c(ACDA)) + \frac{\Delta' - \Delta}{2} \\ &= c(ADA) + \frac{1}{2}(c(ACDA) + c(BDB) - c(ADA) - c(BCDB)) \\ &\quad - \frac{1}{2}(c(ACDA) + c(BDB) - c(BCDB)) \\ &= c(ADA) - \frac{c(ADA)}{2} \end{aligned}$$

is always positive we are through. \square

Theorem 5 *In the transportation domain hiding tasks may be beneficial.*

For this proof we refer again to Figure 5.2. We now assume that S_B conceals the order o_3 . Still it would be rational to sell the order o_1 to S_B . So the overall profit of the cooperation seems to be

$$\Delta'' = c(ACDA) - c(BCDB).$$

Hence the side paying S_A receives chosen by a PMM is

$$\sigma'' = tut_{S_A}(o_1) + \frac{\Delta''}{2}.$$

Since $\Delta'' < \Delta$ also $\sigma'' < \sigma$, therefore, the lie turns out to be beneficial for S_B . \square

5.4 Price Determination in Transportation Companies

In the last decades the transportation market in Germany developed a free and atomistic structure with many small and medium sized haulage companies and customers. A single company or customer has no noticeable influence on the market. The ware that is traded is transportation capacity that is offered by haulage companies and that is requested by the customers. The customers can be industrial corporations or other transportation companies. The decisions for the prices and the costs for transportation capacity depend on several parameters such as distance, time, the used transportation device etc.

A customer of a haulage company is anyone who submits a transportation order to the company. We just model the few relevant attributes of an order. For example we assume a single type of good that is diversible in tons or meters-of-loading. The price the client pays for the execution of the order is due to negotiations between him and the shipping company. Usually, it tends to the average market price. The market price depends on various parameters and it is influenced by the actions of every company in the market. It depends on various parameters and is a highly dynamical statistical value.

The actual value of the market price varies with supply and demand. For our scenario it has been sufficient to treat it as a fix value. Of course the concrete prices depend on the actual situation, nevertheless the market price can help to determine a maximal limit for the price and to distinguish between profitable and non-profitable orders.

5.4.1 Market Price

Let a be an order offered to shipping company S . $W(a, S)$ is the *price* that S is payed for executing a . If the customer is not a shipping company itself, we assume that the price equals the *market price* $M(a)$.

Let $Cost$ be the *cost-function* the company uses to compute its costs for executing the task. The costs are independent of the price of the order and can be computed in several ways.

We chose to define the costs of a transportation task as the costs of that truck of the company that can perform the task best. In *COMoves* we used a simple cost model based on the minimal additional way a truck had to go. In *COMoves-II* we presented a more elaborated statistical cost model [Siekman et al. 96].

Definition 38 *Let $a \in A$ be an order and $S \in \mathcal{S}$ a transportation company. $W(a, S)$ is the price S receives for the execution of a . Let $Cost : A \times \mathcal{S} \rightarrow R$ be the cost-function the companies uses to compute the costs for the execution of the orders using their own trucks. The profit $G : A \times \mathcal{S} \rightarrow R$ is the difference between the price and the costs: $G(a, S) := W(a, S) - Cost(a, S)$.*

In general the market price is an upper bound for the price and thus for the profit. Since the costs are independent of the price, the profit can be negative if the costs exceed the price.

5.4.2 Modelling the Additional Costs

This cost model has been used to find the best resource for an order on the basis of additional distance that has to be covered for this order.

Definition 39 Let t be a truck, $o \in O$ an order, $\Pi \in P$ a plan consisting of a set of plan steps $p_i \in \Pi$ and let (Π_o) be the set of extensions of plan Π that include the execution of o . d is the function that assigns to every plan step p_i its distance $d(p_i)$. The function $c : P \rightarrow \mathbb{R}$ computes the total distance a plan Π covers:

$$c(\Pi) := \sum_{p_i \in \Pi} d(p_i)$$

Let v_t be a constant that specifies the variable costs of a truck t in dependence of the distance it covered. Now we can define the function $Cost : O \times P \rightarrow \mathbb{R}$ that computes the costs of adding a new order o to an existing plan Π as follows:

$$Cost(o, \Pi) := (\min_{\pi \in \Pi_o} c(\pi) - c(\Pi)) \cdot v_t$$

The following example serves to illustrate this way of computation of costs.

Example 5.4.1 Consider two orders o_1, o_2 , a truck with capacity 40t and as in 5.3 a map labelled with kilometre-distances. The variable costs are $v = 3 \text{ DM/KM}$. In the initial situation the truck is located at the depot and has no loading and an empty plan. Let o_1 be an order of transporting 10t from A to C and o_2 a order of 5t from C to the depot. It is obvious that the tour $Depot \rightarrow A \rightarrow C \rightarrow Depot$ is the best solution for the truck. According to the specified cost measure its costs are:

$$\begin{aligned} Cost(o_1, o_2) &= (\text{dist}(Depot, A) + \dots + \text{dist}(C, Depot)) \cdot v \\ &= 620 \text{ km} \cdot 3 \text{ DM/km} \\ &= 1860 \text{ DM} \end{aligned}$$

In figure 5.3 we relate the capacity utilisation to the distance.

Using the cost-model we described above we obtain for the orders o_1 and o_2 the following costs:




Figure 5.3: A tour

$$\begin{aligned}
 \text{Cost}(o_1) &= ((\text{dist}(\text{Depot}, A) + \text{dist}(A, C) + \text{dist}(C, \text{Depot})) - 0) \cdot v \\
 &= (620 \text{ km} - 0 \text{ km}) \cdot 3 \text{ DM/km} \\
 &= 1860 \text{ DM} \\
 \text{Cost}(o_2) &= ((\text{dist}(\text{Depot}, A) + \dots)) - (\text{dist}(\text{Depot}, A) + \dots) \cdot v \\
 &= (620 \text{ km} - 620 \text{ km}) \cdot 3 \text{ DM/km} \\
 &= 0 \text{ DM}
 \end{aligned}$$

Evaluation

The example shows the advantages and shortcomings of this approach. The advantages are:

- The cost model is perfect to decide which resource is the best at the actual time point.
- It is easy to compute.
- It is adapted to the current situation.

Nevertheless, the model is not capable for the computation of real prices for transportation tasks:

- There is no fair allocation of costs to transportation orders.
- The same tasks announced in a different sequence may result in different costs.
- It is possible that the costs 0 are assigned to orders.

5.4.3 Adequate Modelling of Costs

Before we introduce the statistical cost model we summarise the requirements for an adequate cost model:

- The cost model should be easy to compute.
- It should be stable and produce good results over a relatively long time.
- It should be pessimistic in that it should rather overestimate than underestimate the costs.
- On the basis of the volume or weight and the distance between the location of loading and unloading it should allow a fair allocation of costs to orders.
- It should allow a fair allocation of unbooked legs to the orders.
- The cost model should offer the possibilities to react to the actual situation, e.g. with special offers.

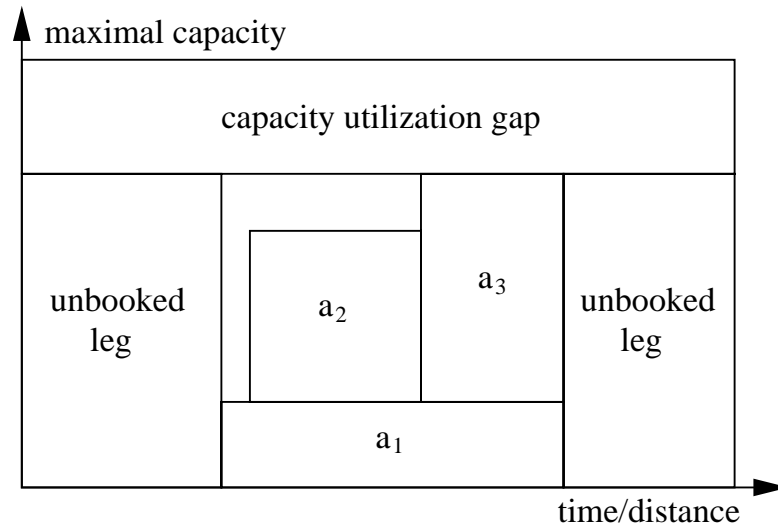


Figure 5.4: A Tour Diagram

Figure 5.4 is an example for a schematic representation of a tour. The maximal performance of transportation that could be reached during that tour is symbolised by the outer rectangle. It is the product of the capacity of the truck and the covered distance.

$$\text{MaxPer} = \text{MaxCap} \cdot \text{MaxWay} = \text{MaxCap} \cdot \sum d(p_i)$$

The actual performance of transportation is the sum of the volume of the rectangles that symbolise the actual performed orders.

$$\text{Per} = \sum (d(p_i) \cdot g(p_i))$$

It is known empirically that trucks are seldom completely utilised. We call that amount of capacity that usually stays unused the *statistical capacity utilisation gap*. Many shipping companies aspire no complete utilisation but one that is reduced according to that gap. Therefore, a truck may be regarded as full although there is some capacity left.

The *proportion of unbooked legs* can as well be computed as statistical value. Hence, we can relate the total driven distance to the distance that is directly connected with orders.

5.4.4 The Statistic Market Model

We are now introducing a cost model that enables us to compute some time in advance an approximation of the costs an order causes when it is executed. This model relies exclusively on statistic information.

Definition 40 *Let t be a truck, $o \in O$ an order, g the function that assigns to the orders its weight, d_m the function that assigns to every order the minimal distance from the location of loading to the location of unloading. The constant v_t represents the variable costs of the truck t .*

c_l is the statistical value that specifies the proportion of unbooked legs and the total distance. c_k specifies the proportion of the average capacity utilisation and the maximal capacity.

Now we define the cost function:

$$\text{Cost}(o) := g(o) \cdot \frac{1}{c_k} \cdot d_m(o) \cdot \frac{1}{c_l} \cdot v_t$$

The following parameters used in this model are statistical or directly related to the order. They are not related to other orders or to the actual situation.

- The weight of the order.

- The minimal distance between the locations of loading and unloading.
- The capacity of the truck.
- The approximated values for the statistical capacity utilisation gap and the proportion of unbooked legs.

Evaluation

This approach offers the following advantages:

- **Simplicity:** The cost function is easy to compute.
- **Adaptive modelling:** It is possible to implement a self-adaptive version of the scenario. It is necessary to protocol the orders and the sum of the costs for executing them to recompute the statistical values.
- **Disturbances:** It is easy to simulate disturbances in the market. Therefore, it is possible to analyse the stability of the approach.
- **Unbooked legs:** The assignment of the costs of unbooked legs to orders is fair.
- **Market price:** It is possible to approximate the market price automatically in the system.

The main shortcoming of the statistical cost model is that the choice of the parameter c_l and c_k is rather arbitrary in the beginning. Using a self-adapting version should overcome this problem soon. Naturally, to reach realistic results it is necessary to use realistic sets of orders. Hence, one should carefully choose the orders if the results are of real-world relevance.

5.5 The Negotiation Situation

In the transportation domain the standard negotiation situation is the following: one agent—either the customer or a shipping company—has a transportation task to perform while there are several companies that are able to perform the task at reasonable costs. Thus the aim of a negotiation mechanism is to find that company which can perform the task at the lowest costs and to negotiate a fair distribution of the profit between the buyer and the seller of the task. Hence we obtain a *one-to-many* negotiation situation with the price for an order as a single topic. Therefore, single-item auction protocols are appropriate mechanisms.

In this section we evaluate whether the single-item auction protocols we presented in Chapter 4 can be used for the negotiation in the transportation domain.

For simplicity we assume that there are no dependencies between the different negotiations. Of course this assumption is not realistic but otherwise we would not be able to do any significant analysis of protocols. In particular we introduced the following limitations:

- The agents cannot use information about the other agents they might have obtained from earlier negotiations.
- The agents do not commit themselves to any actions that do not concern the current negotiation. It is not possible to trade “favours”. Every utility transfer must be explicit using the concept of side paying.
- We do not want to model behaviour that is not rational in short term but that might pay off in the long term, like offering unnecessary concessions to a partner to establish a good relationship or offer dumping prices with the goal to ruin the competitors.
- Furthermore, we treat the statistical values as well as the market price as constant throughout a single negotiation because we assume that the volume of a single order is too small to cause significant changes.

5.5.1 Evaluation of the Auction Protocols

We will now discuss the appropriateness of the following protocols in the transportation domain. We decided to restrict our attention to the three classical auctions and the Vickrey mechanism, which are the basics for the more elaborated single item auctions.

- First-Bid-First-Price (the Contract-Net Protocol)
- First-Bid-Second-Price Protocol (the Vickrey Auction)
- Open English Auction
- Open Dutch Auction

The dominant strategy for the bidder participating in the Vickrey auction is to reveal their true appraisal of the offered object, i.e. the maximal price they would be willing to pay. If every bidder acts rationally the outcome of the Vickrey auction is exactly the appraisal of that bidder with the second highest appraisal of the object. The English auction converges to the same result if we assume an infinitesimal width of steps. The results of the other two auctions may differ from that value if the bidders are not able to assess the appraisal of the others. The more the bidders know about each other the nearer the result will be to the second appraisal. Hence, theoretically if we assume that the bidders have enough

| | costs of computation | costs of communication | revealed information |
|-----------------|-------------------------|---------------------------|-------------------------|
| Contract Net | high | low | medium |
| Vickrey Auction | low | low | high |
| English Auction | low | high | medium |
| Dutch Auction | high | high | low |

Figure 5.5: Evaluation of the auctions

information about each other and that the steps are small all the auctions produce the same result. However in the real world the results can differ quite a lot. We analyse the auctions according to the following criteria:

Costs of computation: For the participants of a Vickrey or an English auction the strategies are straightforward: just reveal the true appraisal, resp. just keep bidding until it is reached. On the other hand for rational bidders in the Dutch auction or the contract net it is necessary to invest effort in the detection or estimation of the appraisals of the other bidders. This can be done by the use of experiences, e.g. based on statistics, or, if possible, even by espionage.

Costs of communication: The contract net and the Vickrey auction get by with a little amount of communication; while the Dutch and English auction need comparatively much communication effort. In the shipping domain time is expensive, and, therefore, the amount of communication is an important factor. Usually a dispatcher decides after a few telephone calls whether to sell an order or not. Of course this factor loses importance if the agents are non-human and dispose of fast communication devices.

Revealed information: In an extremely competitive domain like the transportation domain information can be of essential relevance. Especially if the auctioneer is a shipping company itself, and, therefore, a competitor of the bidders, an auction form that gives him too much insight into the internal state of the bidders is generally not acceptable. Therefore, in such domains the Vickrey auction is only applicable if a totally trusted third party adopts the role of the auctioneer.

Table 5.5 summarises these results.

Assuming a trustworthy auctioneer we can take the first and the second of the above criteria as crucial and so the Vickrey auction can be rated as best and

the Dutch auction as worst. Because of that we decided to integrate the Vickrey mechanism in the *COmoves* system.

Attaching more importance to the criterion of revealed information the Dutch auction gains desirability because the bidders are barely exposed. Especially if the appraisal of the bidders and the auctioneer are not constant but change over time an auction mechanism that can deal with time seems to be beneficial. Fischer and Müller [Fischer & Müller 96] introduced a model for the dynamical change of appraisals of tasks in dependence of the possibility to receive additional tasks and the latest possible time point for the execution. A mechanism similar to the Dutch auction is proposed as a possibility to deal with dynamically changing appraisals.

5.5.2 A Negotiation Strategy for Competing Shipping Companies

Let us consider again the example we had before (5.6). The shipping companies S_A and S_B are located in A resp. B and the task o_1 , an order of 20 from $C \rightarrow D$ is to offered to S_A . Now we can distinguish the following cases:

1. If nothing else is known and both shipping companies do not have further orders to be executed, a truck of shipping company S_B would be the best for executing the order. We define

$$\Delta_{S_A} := tut_{S_A}(o_1) = \hat{p}(o_1) - c(A \rightarrow C \rightarrow D \rightarrow A)$$

and

$$\Delta_{S_B} := tut_{S_B}(o_1) = \hat{p}(o_1) - c(B \rightarrow C \rightarrow D \rightarrow B)$$

where Δ_{S_A} and Δ_{S_B} specify the utility S_A and S_B receive by executing the order. We have $\Delta_{S_A} < \Delta_{S_B}$ because C is closer to B than to A . If $\Delta_{S_A} > 0$, it would not be rational for S_A to give the order to S_B without being paid utility from S_B . On the other hand, in this situation it would not be rational for S_A to execute the order by itself. S_A could offer the deal:

$$((\emptyset, +\Delta_{S_A}), (\{o_1\}, -\Delta_{S_A}))$$

to S_B and it would not be rational of S_B to reject this deal because it obtains a utility of $\Delta_{S_B} - \Delta_{S_A} > 0$ from it.

2. Assume that S_A already has an order o_2 of 20 from $A \rightarrow D$; S_B does not have any orders at all. In this case it is not rational for S_A to give the order to S_B , because




Figure 5.6: Example for the changes of appraisals of tasks

$$\Delta'_{S_A} > \Delta_{S_B}, \text{ where } \Delta'_{S_A} = \hat{p}(o_1) - c(A \rightarrow C \rightarrow D) + c(A \rightarrow D)$$

holds and therefore S_B is not able to pay S_A a price which is higher than the price A could get by executing the order on its own.

Now assume that S_A does not really have an order o_2 of 20 with source A and destination D , but there is a probability of 70% for getting such an order. Is it still rational for S_A to execute the order o_1 on its own? With a probability of 70% S_A receives more utility than it could get by selling the order to S_B . By computing the expected utility [Haddawy & Hanks 90] for both alternatives, we can infer that keeping the order would be individual rational for S_A .

3. The situation again changes if we assume that S_A has an order o_2 of 20 with source A and destination D and S_B has an order o_3 of 20 with source B and destination D . In this situation S_B is the best to execute the order o_1 but it has to pay S_A a high price for the order. Hence, S_B will earn only little money by executing o_1 because S_A was already able to execute the order at low costs.

Now assume that both companies S_A and S_B do not really have the orders o_2 and o_3 but have a 70% probability to get such an order. In this case in the beginning S_B is the best one to execute the order. With a probability of 49% both will get the expected order and S_A would lose utility if it gave away the order without the assumption of getting the additional order o_2 . There is a chance of 9% that both of them will not get the expected additional order and a chance of 42% that only one of them will get the additional order. In both of the latter cases S_B might lose utility if o_1 is transferred for the price where S_A assumed that it would have the additional




Figure 5.7: A time-dependent decision function to determine the price of an order in a negotiation process.

order. Is it rational for S_B to buy the order? At which price should S_B buy the order?

To solve this problem one could try to compute expected utilities for sets of tasks but this is rather awkward because the probability of being offered a specific set of orders is not known beforehand. We therefore suggest another solution to the problem. For this purpose, we return to the situation where we have truck L_A located in city A and truck L_B located in city B . There is only one order o_1 (20 units to transport from C to D) present which was received by the shipping company S_A and is currently scheduled to truck L_A . When L_A executes the order S_A gets a utility of Δ_{S_A} . A threshold utility $\bar{\Delta}$ is specified which defines the amount of utility a truck has to get to stop the shipping company from looking for partners to execute the orders currently scheduled to the truck. The situation for the shipping company agent with respect to truck L_A is now characterised by figure 5.7. S_A invents a phantom task o^* which ideally completes order o_1 — in the example: 20 units to transport from A to D just at the time o_1 has to be executed — and which gives S_A a utility of Δ'_{S_A} . S_A announces o_1 at time t_0 to S_B specifying the expected utility Δ'_{S_A} as a price to be paid if the order is to be transferred. Because Δ_{S_B} is less than Δ'_{S_A} , S_B does not accept this offer.

While time passes and nothing changes significantly, i.e. no new orders arrive, at time t_1 S_A starts to believe that the probability for an order to arrive which can substitute o^* decreases linearly. Therefore, S_A continuously reduces the price Δ'_{S_A} after t_1 . This continues at most until t_3 because this is the time when L_A actually has to start executing the order to meet the specified time constraints. However, at time t_2 we have $\Delta'_{S_A} = \Delta_{S_B}$ which signals that the deal is becoming attractive for S_B . It would not be individually rational for S_B to buy the order exactly at time t_2 because S_B would pass all its utility to S_A . But from the next time on S_A reduces the price of its offer, it will be individually rational for S_B to accept the offer.

For S_B the question is now when is the right time to accept the offer. Even if we assume that S_B knows t_3 , the time L_A has to start executing o_1 , there could be another shipping company waiting for its chance to get o_1 . For the sake of the

example let us assume that

$$\Delta_{S_A} \leq \frac{\Delta_{S_B}}{2}$$

In this case the price of $\frac{\Delta_{S_B}}{2}$ seems to be fair because it exactly splits the utility between S_A and S_B . Therefore, a general strategy a shipping company could choose to negotiate about prices for orders is, to wait until the price of the offer is less or equal to half the utility the shipping company could get out of executing the order.

Although this is a quite robust negotiation strategy in practice, this strategy is not in a Nash-Equilibrium. The problem is that at any time a shipping company who has an advantage for executing a specific order (mostly because of its location) is able to profit from this situation if it is aware of it. More concretely, if a shipping company knows that it is able to execute an order at costs which are significantly lower than the costs of any other shipping company, it can buy this order spending less than half the utility it obtains by executing the order. Still, if the market is quite big, i.e. if there is a large number of shipping companies, the probability that one shipping company gets into such an outstanding position is quite low. Furthermore, even if a shipping company is able to buy an order without splitting the utility fairly, the corresponding partner company does not really lose utility because it is paid at least its expected utility.

5.6 The Integration of the Auction Protocols and the Market Model

In *COMoves-II* the market price is a direct consequence of the statistic cost model that was presented in Section 5.4. We assume that a customer who awards an order to a haulage company pays the market price. So we restricted the possibilities of price negotiation to negotiations between the companies when trading tasks among each other. In this negotiations the market price serves as clue and limit.

We analyse how these defaults influence the strategies of the bidders in the contract net and the Vickrey auction. Let o be an order, $M(o)$ its market price and $cost_S(o)$ the costs that the execution of o causes to company S . We have to distinguish the cases that the costs are lower resp. higher than the market price.

5.6.1 Vickrey Auction

Bidding for a task in the Vickrey auction, it is in general rational to reveal the true execution costs. We have to modify this rule if the costs exceed the market price which is defined as a limit. One solution is not to bid for tasks that are too

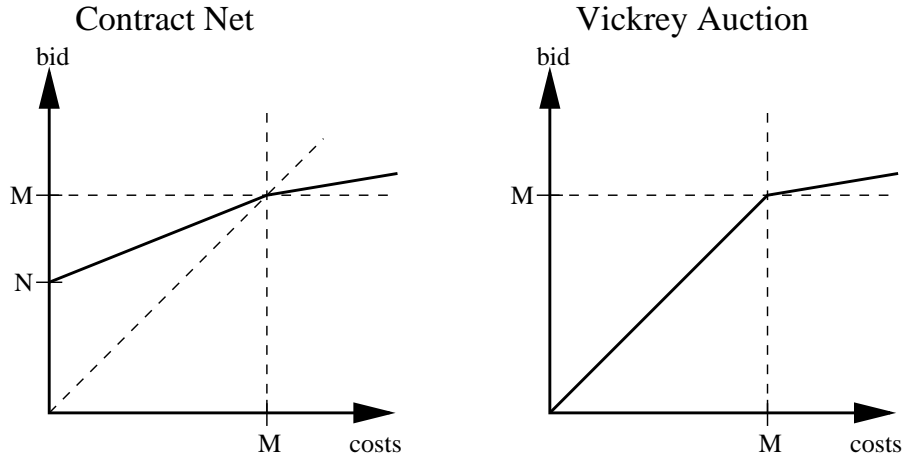


Figure 5.8: The bidding according to the market model

costly to execute. An alternative is to submit a bid at the level of the market price and to risk the difference between the bid and the market price as a loss. This may be rational if there are chances to receive additional tasks that fit well in the planned tour and therefore cause few costs and high profit. In the a-posteriori evaluation the order that seemed to be unattractive in the beginning can turn out to be quite fruitful. Using this strategy the bid that company S submits for order o is computed as follows:

$$G_S^v(o) := \begin{cases} cost_S(o) & \text{if } cost_S(o) \leq M(o) \\ M(o) + \varepsilon(cost_S(o) - M(o)) & \text{if } cost_S(o) > M(o) \end{cases}$$

We require $0 < \varepsilon \ll 1$. Thus, we can guarantee that the function G^v is strictly monotonic increasing. This guarantees that even if the order would cause costs that exceed the market price for all bidders the bidder that could perform the task with the least effort receives it.

5.6.2 Contract Net

The market price is a limit for the bids in the contract net as well. As in the Vickrey auction we choose the strategy of submitting a bid at the level of the market price if the costs exceed it and hoping for additional tasks. If the costs are less than the market price a price between the costs and the market price has to be chosen. The price should be high enough to maximise the profit but low enough to receive the order. To avoid an expensive and speculative reasoning about the bids of the other agents we introduce a factor $0 < \delta_S < 1$ that states how much of the difference between the market price and the costs the bidder

claims as profit for himself.

$$G_S^c(o) := \begin{cases} cost_S(o) + \delta_S(M(o) - cost_S(o)) & \text{if } cost_S(o) \leq M(o) \\ M(o) + \varepsilon(cost_S(o) - M(o)) & \text{if } cost_S(o) > M(o) \end{cases}$$

Figure 5.8 shows the bidding function in the contract net and the Vickrey auction.

Chapter 6

Conclusion

A fundamental problem in multiagent research consists in the question of how the single-agent behaviours within a group of self-interested agents can be coordinated so that the overall group utility is maximised. This research report gives an overview of the decision and game theoretical concepts that are fundamental for approaching this coordination problem. By taking a closer look at the transportation domain in the last chapter, it also shows that there is a great need for these concepts in application-oriented multiagent research.

In order to be able to understand how self-interested agents make decisions, we started with a brief introduction into decision theory as invented by von Neumann and Morgenstern and payed special attention to the foundations of game theory. After that, we took a look at the kinds of domains in which autonomous agents can interact with each other. We presented an approach of Rosenschein and Zlotkin who utilise game theoretical results in order to classify multiagent domains in a three-level hierarchy. Outgoing from that, they identified fair negotiation mechanisms that prevent deception in the defined domain classes. Unfortunately, they did not provide a comprehensive overview of negotiation protocols that are suitable for fulfilling this purpose under incomplete information.

In order to fill this gap, we investigated the theory of coordination mechanism design for theoretical foundations suitable for identifying deception-preventing negotiation mechanisms. We outlined that, according to this theory, the needed negotiation mechanisms are incentive-compatible ones in which rationally acting, self-interested agents are enforced to reveal their preference informations to each other truthfully. In the following, we were able to identify several incentive-compatible auction-based negotiation protocols which can be used to yield outcomes that maximise social welfare in multiagent systems composed of rationally behaving self-interested agents.

In the last chapter, we presented the transportation domain as a fruitful example for the application of multiagent technologies to real-world problems. We extended the theoretical tools of Chapter 3 to enable an intense analysis of the transportation domain and discovered that for self-interested agents it may be

beneficial to deceive other agents by lying about individual preferences. From this observation it should have become clear that the application of incentive-compatible negotiation mechanisms is needed in order to design global utility maximising multiagent systems for domains with competitive, self-interested agents. It also has become obvious that the concepts of game and decision theory serve well in approaching coordination problems within the research area of multiagent systems.

Bibliography

- [Arrow 63] K.J. **Arrow**. *Social Choice and Individual Values*. New York: John Wiley, 2nd edition, 1963.
- [Aumann 59] R. **Aumann**. *Acceptable points in general cooperative n-person games*. In: Contributions to the Theory of Games, volume IV, Princeton University Press, 1959.
- [Austin 62] J. L. **Austin**. *How to do Things with Words*. Clarendon, Oxford, UK, 1962.
- [Axelrod 84] R. **Axelrod**. *The Evolution of Cooperation*. Basic Books, 1984.
- [Bachem et al. 92] A. **Bachem**, W. **Hochstättler**, and M. **Malich**. *Simulated Trading: A New Approach For Solving Vehicle Routing Problems*. Technical Report 92.125, Mathematisches Institut der Universität zu Köln, Dezember 1992.
- [Bagchi & Nag 91] P. **Bagchi** and B. **Nag**. *Dynamic Vehicle Scheduling: An Expert System Approach*. Journal of Physical Distribution and Logistics Management, 21(2), 1991.
- [Boddy & Dean 94] M. **Boddy** and T. L. **Dean**. *Deliberation scheduling for problem solving in time-constrained environments*. Artificial Intelligence, 67:245–285, 1994.
- [Bodin et al. 83] L. **Bodin**, B. **Golden**, A. **Assad**, and M. **Ball**. *Routing and Scheduling of Vehicles and Crews*. Computers and Operations Research, 10(2):63 – 211, 1983.
- [Bussmann & Müller 92] S. **Bussmann** and H. J. **Müller**. *A Negotiation Framework for Cooperating Agents*. In: Proc. of the 2nd Workshop on Cooperating Knowledge Based Systems. Keele, GB, September 1992.
- [Chaib-Draa 91] B. **Chaib-Draa**. *Commitments and Intentions in Distributed AI Systems: Preliminary Report*. Technical Report Report DIUL-RR-9114, Department d’informatique, Université Laval, Québec, 1991.
- [Clarke 71] E. H. **Clarke**. *Multipart Pricing of Public Goods*. In: Public Choice 11:17-33, 1971.

- [Clarke 72] E. H. **Clarke**. *Multipart Pricing of Public Goods: An example*. In: Muskin, S., ed., *Public Prices for Public Products*, Urban Inst., Washington, 1972.
- [Cohen & Perrault 86] P.R. **Cohen** and C.P. **Perrault**. *Elements of a Plan-Based Theory of Speech Acts*. In: B.J. Grosz, K. Sparck Jones, and B.L. Webber (eds.), *Readings in Natural Language Processing*, pp. 423–440. Morgan Kaufmann Publishers, Inc., Calif., 1986.
- [Conry et al. 91] S. E. **Conry**, K. **Kuwabara**, V. R. **Lesser**, and R. A. **Meyer**. *Multistage Negotiation for Distributed Constraint Satisfaction*. *IEEE Transactions on Systems, Man, and Cybernetics (Special Section on DAI)*, 21(6):1462–1477, November/December 1991.
- [Corkill & Lesser 83] D. D. **Corkill** and V. R. **Lesser**. *The Use of Meta-Level Control for Coordination in a Distributed Problem Solving Network*. In: *Proceedings of IJCAI-83*, pp. 767–770, Karlsruhe, Germany, 1983.
- [Dasgupta et al. 79] P. **Dasgupta**, P. **Hammond**, and E. **Maskin**. *The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility*. In: *Review of Economic Studies*, pp. 185–216, 1979.
- [Dawkins 76] R. **Dawkins**. *The Selfish Gene*. Walton Street, Oxford: Oxford University Press, 1976.
- [Decker & Lesser 90] K. **Decker** and V. R. **Lesser**. *A Scenario for Cooperative Distributed Problem Solving*. In: *Proceedings of the 10th International Workshop on Distributed Artificial Intelligence*, Bandera, Texas, 1990.
- [Decker & Lesser 94] K. **Decker** and V. R. **Lesser**. *Designing a Family of Coordination Algorithms*. In: *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, pp. 65–84, Lake Quinalt, Washington, July 1994.
- [Demange et al. 86] D. **Demange**, D. **Gale**, and M. **Sotomayor**. *Multi-item auctions*. In: *Journal of Political Economy* 94, 863–872, 1986.
- [Desrochers et al. 92] M. **Desrochers**, J. **Desrosiers**, and M. **Solomon**. *A New Optimization Algorithm for the Vehicle Routing Problem with Time Windows*. *Operations Research*, 40(2), 1992.
- [Drexler & Miller 88] K.E. **Drexler** and S.M. **Miller**. *Incentive Engineering for computational resource management*. In: *The Ecology of Computation*, ed. B.A. Huberman, pp. 231–266, 1988.
- [Durfee & Montgomery 90] E. H. **Durfee** and T. A. **Montgomery**. *A Hierarchical Protocol for Coordination of Multiagent Behaviour*. In: *Proc. of the 8th National Conference on Artificial Intelligence*, pp. 86–93. Boston, MA, 1990.
- [Durfee 88] Edmund H. **Durfee**. *Coordination of Distributed Problem Solvers*. Boston: Kluwer Academic Publishers, 1988.

- [Eisenführ & Weber 93] F. **Eisenführ** and M. **Weber**. *Rationales Entscheiden*. Berlin, Heidelberg, New York, Tokio: Springer-Verlag, 1993.
- [Emons 94] Winand **Emons**. *The Provision of Environmental Protection Measures under Incomplete Information: An Introduction to the Theory of Mechanism Design*. In: International Review of Law and Economics (1994) 14, pp. 479-491, 1994.
- [Ephrati & Rosenschein 93] Eithan **Ephrati** and Jeffrey S. **Rosenschein**. *A Framework for the Interleaving of Execution and Planning for Dynamic Tasks by Multiple Agents*. In: 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, 1993.
- [Ephrati 93] Eithan **Ephrati**. *Planning and consensus among autonomous agents*. PhD thesis, PhD thesis, Hebrew University, Israel, 1993.
- [Falk et al. 93] J. **Falk**, S. **Spieck**, and P. **Mertens**. *Unterstützung der Lager- und Transportlogistik durch Teilintelligente Agenten*. Information Management, 2, 1993.
- [Fischer & Müller 95] K. **Fischer** and H. J. **Müller**. *Cooperative Problem Solving in the Transportation Domain*. In: Ulrich Derigs (ed.), Proceedings of the International Conference on Operations Research (OR'94). Springer-Verlag, 1995.
- [Fischer & Müller 96] K. **Fischer** and J. P. **Müller**. *A Decision-Theoretic Model for Cooperative Transportation Scheduling*. In: W. van de Velde and J. W. Perram (eds.), Agents Breaking Away — Proceedings of the 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96), volume 1038: LNAI, pp. 177–189. Springer-Verlag, 1996.
- [Fischer et al. 93a] K. **Fischer**, N. **Kuhn**, H. J. **Müller**, and J. P. **Müller**. *Modeling the Transportation Domain*. In: Proceedings of the 3rd Workshop on AI in Economics and Management (AIEM'93), Portland, Oregon, USA, 1993.
- [Fischer et al. 93b] K. **Fischer**, N. **Kuhn**, H. J. **Müller**, J. P. **Müller**, and M. **Pischel**. *Sophisticated and Distributed: The Transportation Domain*. In: Proceedings of MAAMAW-93, Neuchatel, CH, August 1993. Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World.
- [Fischer et al. 94] K. **Fischer**, N. **Kuhn**, and J. P. **Müller**. *Distributed, Knowledge-Based, Reactive Scheduling in the Transportation Domain*. In: Proceedings of the Tenth IEEE Conference on Artificial Intelligence and Applications, pp. 47–53, San Antonio, Texas, March 1994.
- [French 88] S. **French**. *Decision Theory - An Introduction to the Mathematics of Rationality*. New York: John Wiley & Sons, 1988.
- [Gale 60] David **Gale**. *The Theory of Linear Economic Models*. New York: McGraw-Hill, 1960.

- [Groh 96] Christian **Groh**. *Revelation principle*. Seminar workout, 1996.
- [Groves 73] Theodor **Groves**. *Incentives in teams*. In: *Econometrica*, 41:617-631, 1973.
- [Haddawy & Hanks 90] P. **Haddawy** and S. **Hanks**. *Issues in Decision-Theoretic Planning: Symbolic Goals and Numeric Utilities*. In: Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control. Morgan Kaufmann, November 1990.
- [Huberman & Clearwater 95] Bernardo A. **Huberman** and Scott H. **Clearwater**. *A multi-agent system for controlling building environments*. In: Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 171-176, San Francisco, CA, 1995.
- [Jennings 96] N.R. **Jennings**. *Coordination Techniques for Distributed Artificial Intelligence*. In: Foundations of Distributed Artificial Intelligence (eds. O'Hare, G. and Jennings, N.R.):187-210, Wiley, 1996.
- [Kuhn et al. 93] N. **Kuhn**, H. J. **Müller**, and J. P. **Müller**. *Simulating Cooperative Transportation Companies*. In: Proceedings of the European Simulation Multiconference (ESM-93), Lyon, France, June 1993. Society for Computer Simulation.
- [Kuhn 55] H. W. **Kuhn**. *The Hungarian Method for the Assignment Problem*. In: *Naval Research Logistics Q.* 2, no. 1: 83-97, 1955.
- [Laffont & Maskin 80] J. J. **Laffont** and E. **Maskin**. *The theory of incentives: an overview*. In: Hildenbrand, Werner, ed., *Advances in Economic Theory*, chapter 3, pp. 31-94, Cambridge University Press, 1980.
- [Leonard 83] H. B. **Leonard**. *Elicitation of honest preferences for the assignment of individuals to positions*. In: *Journal of Political Economy* 91, 461-479, 1983.
- [Luce & Raiffa 57] R. **Luce** and L. **Raiffa**. *Games and Decisions*. New York: Wiley, 1957.
- [Lux & Steiner 95] A. **Lux** and D. D. **Steiner**. *Understanding Cooperation: an Agent's Perspective*. In: Proceedings of the First International Conference on Multiagent Systems, San Francisco, CA, 1995.
- [Lux 95] A. **Lux**. *Kooperative Mensch-Maschine Arbeit - ein Modellierungsansatz und dessen Umsetzung im Rahmen des Systems MEKKA*. PhD thesis, Universität des Saarlandes, Saarbrücken, 1995.
- [Ma et al. 88] C. **Ma**, J. **Moore**, and S. **Turnbull:88**. *Stopping agents from "cheating"*. In: *Journal of Economic Theory* 46: pp. 355-372, 1988.
- [Malone 90] T. W. **Malone**. *Organizing Information Processing Systems: Parallels Between Human Organizations and Computer Systems*. In: S. P. Robertson, W. Zachary, and J. B. Black (eds.), *Cognition, Computing, and Cooperation*. New Jersey: Ablex Publishing Corporation Norwood, 1990.

- [Martimort & Stole 97] David **Martimort** and Lars **Stole**. *Communication Spaces, Equilibria Sets and the Revelation Principle under Common Agency*. University of Chicago, GSB, 1997.
- [Maynard-Smith 82] J. **Maynard-Smith**. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- [Myerson 81] Roger B. **Myerson**. *Optimal Auction Design*. In: Mathematics of Operations Research 6 (February 1981):58-73, 1981.
- [Nash 53] John F. **Nash**. *Two-person cooperative games*. In: Econometrica, 21, 1953.
- [Olson & Porter 94] Mark **Olson** and David **Porter**. *An experimental examination into the design of decentralized methods to solve the assignment problem with and without money*. In: Economic Theory 4, 1994.
- [Olson 92] Mark **Olson**. *Point voting mechanisms*. In: Mimeo, CREED, University of Amsterdam, Holland, 1992.
- [Perrault 90] C.R. **Perrault**. *An Application of Default Logic to Speech Act Theory*. In: P.R. Cohen, J. Morgan, and M. Pollack (eds.), Intentions in Communication, pp. 161–185. MIT Press, Cambridge, Mass., 1990.
- [Pischel et al. 95] M. **Pischel**, K. **Fischer**, J. P. **Müller**, and R. **Wingenter**. *COMoves: Ein Multiagenten-Ansatz zur Lösung von Dispositionsaufgaben von Speditionen -Realisierung in Dasedis -*. DFKI, 1995.
- [Psaraftis 88] H. **Psaraftis**. *Dynamic Vehicle Routing Problems*. In: Vehicle Routing: Methods and Studies. North-Holland, 1988.
- [Raiffa 82] H. **Raiffa**. *The Art and Science of Negotiation*. Cambridge, Massachusetts: Harvard University Press, 1982.
- [Rasmusen 89] Eric **Rasmusen**. *Games and Information*. : Basil Blackwell, 1989.
- [Rauhut et al. 79] B. **Rauhut**, N. **Schmitz**, and E.-W. **Zachow**. *Spieltheorie*. Stuttgart, 1979.
- [Rosenschein & Ephrati 91a] Jeffrey S. **Rosenschein** and Eithan **Ephrati**. *The Clarke Tax as a Consensus Mechanism among Automated Agents*. In: Proceedings of the Ninth National Conference on Artificial Intelligence, S. 173-178, 1991.
- [Rosenschein & Ephrati 91b] Jeffrey S. **Rosenschein** and Eithan **Ephrati**. *Voting and Multi-agent Consensus*. In: Technical report, Computer Science Department, Hebrew University, Jerusalem, Israel, 1991.
- [Rosenschein & Ephrati 93] Jeffrey S. **Rosenschein** and Eithan **Ephrati**. *Distributed consensus mechanisms for self-interested heterogeneous agents*. In: Proceedings of the First Conference on Intelligent and Cooperative Information Systems, Rotterdam, 1993.

- [Rosenschein & Zlotkin 94a] J. S. **Rosenschein** and G. **Zlotkin**. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- [Rosenschein & Zlotkin 94b] Jeffrey S. **Rosenschein** and Gilad **Zlotkin**. *Coalition, Cryptography, and Stability: Mechanisms for Coalition Formation in Task Oriented Domains*. In: Proceedings of the Twelfth National Conference on Artificial Intelligence, 1994.
- [Rosenschein & Zlotkin 94c] Jeffrey S. **Rosenschein** and Gilad **Zlotkin**. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- [Rosenschein et al. 91] Jeffrey S. **Rosenschein**, Eithan **Ephrati**, and Gilad **Zlotkin**. *A Non-manipulable Meeting Scheduling System*. 1991.
- [Roth & Sotomayor 91] Alvin **Roth** and M. **Sotomayor**. *Two-Sided Matching*. 1991.
- [Rothkopf et al. 90] Michael H. **Rothkopf**, Thomas J. **Teisberg**, and Edward P. **Kahn**. *Why are Vickrey Auctions rare?* In: Journal of Political Economy, Volume 98, Number 1, February 1990, 1990.
- [Ruß 97] Christian **Ruß**. *Economic Mechanism Design for the Auction-Based Coordination of Self-Interested Agents*. Universität des Saarlandes, Saarbrücken: Diploma Thesis, Computer Science Department, 1997.
- [Saaty 80] T.L. **Saaty**. *The Analytical Hierarchy Process*. New York: , 1980.
- [Sandholm & Lesser 95] Tuomas **Sandholm** and Victor **Lesser**. *On Automated Contracting in Multi-enterprise manufacturing*. In: Proc. Improving Manufacturing Performance in a Distributed Enterprise: Advanced Systems and Tools, Edinburgh, pp.33-42, 1995.
- [Sandholm 93] T. W. **Sandholm**. *An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations*. In: Proceedings of the 12th International Workshop on Distributed Artificial Intelligence, pp. 295–308, Hidden Valley, Pennsylvania, May 1993.
- [Schofield 85] N.J. **Schofield**. *Social Choice and Democracy*. Berlin: Springer-Verlag, 1985.
- [Searle 69] J. R. **Searle**. *Speech Acts*. Cambridge University Press, 1969.
- [Shapley 53a] L. S. **Shapley**. *A value for n-person games*. In: Contributions to the Theory of Games. Princeton University Press, 1953.
- [Shapley 53b] L.S. **Shapley**. *A value for n-person games*. In: Contributions to the Theory of Games. Princeton University Press, 1953.
- [Siekmann et al. 96] J. **Siekmann**, K. **Fischer**, M. **Pischel**, and G. **Vierke**. *COMoves-II: Modelle Preisbasierter Verhandlungsführung in der Speditionsdomäne*. DFKI, 1996.

- [Steiner et al. 90] D. **Steiner**, D. **Mahling**, and H. **Haugeneder**. *Human Computer-Supported Cooperative Work*. In: Proceedings of the 10th International Workshop on Distributed Artificial Intelligence. MCC Technical Report Nr. ACT-AI-355-90, 1990.
- [Sycara 87] K. P. **Sycara**. *Resolving Adversarial Conflicts: An approach integrating case-based and analytic methods*. PhD thesis, Georgia Institute of Technology, Atlanta, Georgia, June 1987.
- [Varian 95] Hal R. **Varian**. *Economic Mechanism Design for Computerized Agents*. In: USENIX Workshop on Electronic Commerce, 1995.
- [Vickrey 61] W. **Vickrey**. *Counterspeculation, Auctions, and Competitive Sealed Tenders*. In: Journal of Finance 16: 8-37, 1961.
- [von Neumann & Morgenstern 44] J. **von Neumann** and O. **Morgenstern**. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944.
- [Weinhardt et al. 96a] Christof **Weinhardt**, Peter **Gomber**, and Claudia **Schmidt**. *Efficiency and Incentives in MAS-Coordination*. In: Discussion Paper Nr. 8, 1996.
- [Weinhardt et al. 96b] Christof **Weinhardt**, Peter **Gomber**, and Claudia **Schmidt**. *Ein agentenorientierter Ansatz zur expeditionsinternen Allokation von Transportaufträgen*. In: Discussion Paper Nr. 4a, 1996.
- [Weinhardt et al. 96c] Christof **Weinhardt**, Peter **Gomber**, and Claudia **Schmidt**. *Synergie und Koordination in dezentral planenden Organisationen*. In: Wirtschaftsinformatik 38, pp. 299-307, 1996.
- [Wellman 92] M.P. **Wellman**. *A General-Equilibrium Approach to Distributed Transportation Planning*. In: Proceedings of AAAI-92, pp. 282-290, 1992.
- [Young 75] O. R. **Young**. *Bargaining. Formal Theories of Negotiation*. Urbana: University of Illinois Press, 1975.
- [Zermelo 12] E. **Zermelo**. *Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels*. In: Proceedings of the Fifth International Congress of Mathematicians. Cambridge, 1912.
- [Zeuthen 30] Frederic **Zeuthen**. *Problems of Monopoly and Social Welfare*. London: G. Rotledge and Sons, 1930.
- [Zlotkin & Rosenschein 93] G. **Zlotkin** and J. S. **Rosenschein**. *A Domain Theory for Task-Oriented Negotiation*. In: Proc. of the 13th International Joint Conference on Artificial Intelligence, volume 1, Chambéry, France, 28.8.-3.9. 1993.

**Decision Theory and Coordination
in Multiagent Systems**

Klaus Fischer, Christian Ruf, Gero Vierke