Modelling human behaviour in social dilemmas using attributes and heuristics
Dissertation

Eva Ebenhöh

Institute of Environmental Systems Research
University of Osnabrück, Germany
Eva.Ebenhoeh@usf.uos.de
http://www.usf.uos.de/~eebenhoe

30th June 2006

ii

Abstract

A question concerning not only modellers but also practitioners is: under what circumstances can mutual cooperation be established and maintained by a group of people facing a common pool dilemma? A step before this question of institutional influences there is need for a different way of modelling human behaviour that does not draw on the rational actor paradigm, because this kind of modelling needs to be able to integrate various deviations from this theory shown in economic experiments. We have chosen a new approach based on observations in form of laboratory and field observations of actual human behaviour. We model human decision making as using an adaptive toolbox following the notion of Gigerenzer. Humans draw on a number of simple heuristics that are meaningful in a certain situation but may be useless in another. This is incorporated into our agent-based model by having agents perceive their environment, draw on a pool of heuristics to choose an appropriate one and use that heuristic. Behavioural differences can be incorporated in two ways. First, each agent has a number of attributes that differ in values, for example there are more and less cooperative agents. The second behavioural difference lies in the way, in which heuristics are chosen. With this modelling approach we contribute to a new way of modelling human behaviour, which is simple enough to be included into more complex models while at the same time realistic enough to cover actual decision making processes of humans. Modellers should be able to use this approach without a need to get deep into psychological, sociological or economic theory. Stakeholders in social dilemmas, who may be confronted with such a model should understand, why an agent decides in the way it does.

Acknowledgements

I want to thank Claudia Pahl-Wostl for her support and her competent and commited supervision. I also thank Thomas Berger for his helpful and constructive critique. Thank also goes to all colleagues who have critised this work at one time or another, especially Stephan Jätzold, Georg Holtz, Ole Schröder, and Jens Newig.

For their help and support I want to thank especially my mother and my father, who seem to be impressed with my work more than any other persons. Also, I want to thank all those, who have been taking care of my children while I was at work, especially my parents, Stephan's parents, Peter Gursky, Henrik Dittman, and Sophie Rotter.

Most of all, I thank Stephan for making it possible for me to focus on my career, by insisting on taking his share in child care.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Eva Ebenhöh

1.1 Introduction

> "A [...] basic difficulty is that such aggregation cannot be correctly made without a reasonable model of the same socio-economic system stated in terms of the behavior and interaction of the elemental decision making units. Then, and only then, could ways be found of aggregating relationships without a disastrous loss of accuracy of representation." Orcutt (1957, 116)

50 years ago, Orcutt (1957) argued the need for a new way of modelling decision making of economic agents on the micro level. We are now in a better position to do so, since information technology has evolved to allow a great number of decision making entities to interact in a software environment. This thesis is a contribution to this emerging field.

In the beginning the working title of this thesis was "New approaches to modelling human behaviour". Early on it became clear that the investigation of human behaviour needed to be restricted to certain decision environments, like social dilemmas, and the new approach was to give agents in an agent-based model several simple heuristics to choose from. In an effort to organise different explanations of human behaviour along the lines of several attributes, out of a great number in the literature seven attributes were filtered out as relevant to decision environments implemented within the scope of this thesis. So, the title became "Modelling human behaviour in social dilemmas using attributes and heuristics".

In the next section (Section 1.2) the "objective" of this thesis is briefly stated and the approach motivated. In Section 1.3 the "methods" are explained including a description of some of the methods *not* used, to delineate this thesis from traditional approaches and to justify that it constitutes, indeed, a *new* approach. Section 1.4, the "History", is a synopsis of the seven articles and working papers included in this thesis, including a statement about connections between these papers and their common history. The Appendix A "Implementation" on page 163 deals with some technical details of the implementation. However, no full description of models or the modelling framework is given. The thesis includes a CD with a comprising Java-documentation of the source code, the source code itself, as well as a tutorial. Java documentation, source code, and tutorial are also available on the website[1].

Some of the aspects covered in this introductory chapter are repeated in one or more of the papers. This redundancy has been accepted, in order to be able to present thoughts developed over several years in an aligned way. Here they have been condensed into a meaningful introduction.

1.2 Objective

The objective of this thesis is to contribute to the solution of social dilemmas arising from collective use of common-pool resources or provision of public goods. This contribution is done in form of a new way to model human behaviour in social dilemmas, which can be re-used in other models, easily extended, and altered. This new model of human behaviour claims to be realistic, because it is based on observation.

Social dilemmas are characterised by a discrepancy between individual and group rationality. Traditional approaches tend to fail to explain human behaviour in these models, because they do not take interrelations of humans into account. Numerous field studies deal with very different real situations, where people solve or fail to solve the dilemmas they are facing. Researchers from a number of fields are making an analytical effort to identify reasons for successes and failures, because the importance of a successful solution of environmental and political "social" dilemmas is widely recognised.

Therefore, the aim of this thesis is to further the analysis of social dilemmas by using agent-based models. It is driven by a number of insights, which are further justified later in this Subsection and in Subsection 1.3:

---

[1]http://www.usf.uos.de/~eebenhoe/forschung/adaptivetoolbox/

1. Theories of human decision making in social dilemmas are fragmented and not yet brought into one coherent framework, although a starting point for this has been made by Ostrom (2004).

2. In social dilemmas, rational actor theories tend to fail to predict human behaviour. However, it is observation of human behaviour which should be the basis of a behavioural theory.

3. Extensions of rational actor theories are able to predict the substance of human decision making but not the process. But, in order to investigate the impact of institutions and policy measures on human behaviour in social dilemmas, the procedures of decision making need to be understood.

4. Since the decision makers are the ones being modelled, they are the best experts for their own decision making.

5. A consistent yet flexible way to model human behaviour in social dilemmas is lacking in the modelling community, which usually designs agents specifically for one purpose.

Corresponding to these insights are a number of sub-objectives which are to be fulfilled by the models developed in this thesis:

Theory building The models should contribute to theory building of social dilemmas and their solutions. In order to be able to do this, they need to help to infer regularities. By modelling different scenarios behavioural regularities are integrated into a common framework of behaviour. This theory is referred to in depth in the Behaviour paper (Chapter 7).

Realism Through an analysis of experimental and empirical behaviour in different settings and environments in which these social dilemmas are situated, behavioural regularities of humans are extracted. An approach to bring the different decision environments together is made in the Classification working paper (Chapter 8).

Procedural Rationality Agents' decision making should be procedurally rational and not only substantively rational (Simon (1997), see also Subsection 1.3.1).

Traceability For the new modelling technique to be a benefit in group model building or social learning processes, agents need to be diverse

and use heuristics that are traceable by real humans. Heuristics in the models are based on expert knowledge of the decision making process. For a further discussion of this need, see the Prospects, Section 9.3.

Contribution to modelling community In order for the models to be re-usable by other modellers, they are organised in a framework that allows for easy extensions and alterations. The implementation is introduced in the Appendix A.

To motivate the approach, this section first introduces social dilemmas from a theoretical viewpoint and then proceeds to argue the need for a new modelling tool to represent human behaviour within these social dilemmas. Statements made in the next two subsections are defended more thoroughly in Section 1.3.

1.2.1   Social dilemmas

> The law locks up the man or woman
> Who steals the goose from off the common
> But leaves the greater villain loose
> Who steals the common from off the goose.
> -Anonymous, 18th century

Social dilemmas occur frequently in the provision of public goods and management of common pool resources. The fact that these management processes face social dilemmas can be accountable for environmental problems like overexploitation and pollution of water bodies, land, and the atmosphere.

The modern discussion of social dilemmas and how to solve them goes back to Demsetz (1967) and Hardin (1968). While Demsetz propagates transfer of common goods into individual property, Hardin calls for governmental regulation. Ostrom (1990), on the other hand, adopts a different approach, based on the insight that some common goods can not or should not be made into private goods and that under certain circumstances communities are capable of solving social dilemmas better than governments. She proceeds to research positive and negative influences on the capabilities of groups to govern themselves (Ostrom (2004)).

The dilemma in social dilemmas is a difference between individual and group rationality. In social dilemmas it is *not* the case that if everybody pursues his or her narrowly defined self interest, the result is a benefit for the group as a whole. This belief, which may be true in other situations, is

based on Adam Smith and is the foundation of neoclassical economics. But social dilemmas are precisely defined as situations in which this assumption does not hold. The neoclassical answer to problems arising from social dilemmas is to avoid them in the first place by making everything personal property. This may work for land, like the commons, but it is not feasible or socially unacceptable for some public goods like bio-diversity and the Internet as a low-cost world-wide mass-communication, or for some common-pool resources like clean air and freshwater. There will always remain social dilemmas to be solved, and therefore we need a theory.

At the same time actual human behaviour in social dilemmas like the management of common pool resources is not very well understood. For these management processes, however, it is important to know, what determines and stabilises cooperation and trust within a group of actors. If these situations are analysed within the rational actor theory, some phenomena like trust and cooperation can not be explained. Therefore, rational actor theory needs to be complemented with theories that can explain these phenomena.

1.2.2 The motivation for a new way to model human behaviour

Social dilemmas are usually modelled as prisoner's dilemmas. For the tragedy of the commons the model is the following. Suppose a common is at its capacity limit. A herder has the option of letting another cow graze on this common or not. Doing so would earn this herder an additional profit on the expense of all others, who would receive less profit from this common, because it is overgrazed. If more than one additional cow is put on this common, the additional profit diminishes.

There are, of course, game theoretical alternatives to the prisoner's dilemma (Heckathorn (1996)), the most obvious alternative being the chicken game. The main difference to the prisoner's dilemma is that if both players defect, returns for both are lesser, than the return for a cooperator if the other defects. But this still remains within the framework of game theory, which rests on the foundation of utility maximisation and complete information. If these two assumptions are relaxed, anything can happen even in simple games like the prisoner's dilemma, lowering the analytical power of game theory.

The rigidity of contemporary rational actor theory, which incorporates complete information and utility maximisation, makes it a very useful tool. It allows for mathematical analysis of complex situations like open markets

and other competitive situations, in which learning has taken place (Ostrom (2005, 99f.)). However, not all situations in which humans face decisions are of this kind, for example, rational actor theories give no account on how learning can take place (Vanberg (2002)). The social dilemmas discussed here are not clearly competitive, since cooperation pays. Experience suggests that it may be profitable for decision makers to assume that other decision makers are not following their individual rationality. "Standard economic analysis neglects the identity and past relations of individual transactors, but rational individuals know better, relying on their knowledge of these relations." Granovetter (1985, 491). The point is, that the decision strategy of the other is not known to the decision maker. In this kind of situation, the rational actor theory is still useful as a benchmark and to explain and model some of the decisions made, but not all. But in some cases actors fare better assuming that others do not follow their individual rationality, but behave cooperative and trusting.

A further criticism is that utility maximisation is psychologically implausible. The problem is twofold. The first problem is that one can scientifically doubt that it is a valid way to describe human behaviour, because actual utility maximisation does not happen in one's head. However, game theoretic rational choice theory need not be understood as a tool for *intentional* explanation, that is the explanation of human behaviour on the basis of intrinsic motivation, which is the case if utility maximisation is viewed as a motivation (Lovett (2006, 239)). Game theory can be used to model behaviour *as if* humans were maximising some utility function, regardless of whether they actually do. This is a form of causal explanation, which may or may not be applicable in a certain situation.

The second problem is that stakeholders modelled in this way can not identify with this supposed behaviour, at least not in situations which are not high-cost and competitive. This has in part been alleviated by altering utility functions, for example to include the other's utility or a personal self interest in norm compliance. However, if anything can and has to be incorporated in the utility function, it looses its scientific advantage, its simplicity and generalisability. Like heuristics, the use of utility functions would then have to be dependent on individual differences and altered in a domain specific way. Heuristics, on the other hand, have the advantages of being psychologically plausible, simple, traceable, "fast and frugal", and not based on complete information.

When modelling is involved in field studies either as an additional tool or in form of group model building, agents in these models need to behave in a way that is traceable for non-modellers. If agents in models are to represent

people in real life, they have to behave like people in real life. Those people are the best judges of agent behaviour.

Both these problems come down to the issue of substantive and procedural rationality discussed in Section 1.3. In some cases, including this thesis as we argue, it is not sufficient to model the outcome of a decision, but it is necessary to model the decision making process as well.

Heuristics have disadvantages as well. If decisions are complex, corresponding heuristics tend to get messy and then it is again difficult for humans to trace agents' decision making. Also, the number of possible heuristics for one decision situation may be unknown or infinite. If the range of possible heuristics is not known, it is hard to choose the best fitting heuristic or to validate chosen heuristics. Heuristics are also highly dependent on behavioural analysis. Different modellers come up with different ways to represent human behaviour in agent-based models.

This is a strong reason for models to be based on observation, in our case on behaviour in economic experiments. It is a well and long known fact that the same aggregate behaviour can be modelled with different underlying mechanisms (Orcutt (1957)). The basis, therefore, needs to be individual data. But even on this basis it is still not unambiguous how decisions are made. But at least, this approach is clearly based on observable decision making and thus represents the real world and not only a theory.

However, modelling behaviour in economic experiments neglects the rich social environment, in which social dilemmas like management of common-pool resources take place in the real world. Therefore, the goal is to provide a tool that allows for a comparative analysis of models based on experimental data and models based on empirical observations from case studies. More thoughts on this aspect are presented in Section 9.3 "Prospects" on page 158.

1.3 Methods

Interdisciplinary research uses methods from several disciplines as well as innovative approaches combining traditional theories. This work is no exception. This section outlines methodology from a number of fields, both used and discarded, and then proceeds to describe this new approach and some aspects of its implementation.

1.3.1 Game theory, experimental, and behavioural economics

The standard economic model of decision making is often traced back to Adam Smith's theory of the invisible hand, which states in short, that indi-

vidual pursuit of ones own material good results in a benefit for all (Smith (1776)). Even before Adam Smith, this belief was nicely laid out in Mandeville's fable of the bees (de Mandeville (1714)). Adam Smith's theory is rich and incorporates diverse behaviour. For him self-interest was one of many influences of human behaviour and only a part of the more important "prudence". This theory was shortened on the aspect of self-interest by later economists including John Stuart Mill (Mill (1865)). From one among many influences, self-interest became the only influence. Further, self-interest was used as a synonym for rationality. Sen (1987) shows the absurdity of equating self-interested behaviour with rational behaviour. Still there remains a tendency in the economic discipline to base research on the assumption that humans behave and should behave in a narrowly defined self-interested way. The reason for this tendency can be found in the precision and conciseness of this theory which makes many calculations possible.

There are two mathematical foundations for this concept of self-interested behaviour. (1) The first is for isolated individuals making choices independent of other individuals, but possibly in a probabilistic environment. Those individuals are advised to calculate expected outcome by multiplying the utility of an outcome with its probability. (2) The second foundation concerns humans in social interactions and is necessarily much more complex. It is game theory in the tradition of von Neumann and Morgenstern (1953) who explicitly assume their rational players to maximise expected utility and to assume that all others behave similarly in a utility maximising way. The latter assumption is important in order to resolve mutual interdependence, for instance to calculate Nash equilibria. Game theory is a powerful tool to analyse stylized interactions. These interactions are presented in form of games with payoff matrices or in extensive form. Using stringent assumptions players' behaviour can be predicted. For instance, in order for a minimax analysis to hold (the prediction that players choose the strategy that minimises the losses the other player can inflict), the game has to be a strictly competitive, zero-sum game, the players have to be motivated to do as much harm as possible to the other, and the other's similar motivation has to be known (cf. McClintock (1972b)).

Experimental economics (Kagel (1995)) rests on the tradition of non-cooperative game theory. Experiments are designed to mirror these stylised interactions, the games. Only a few possible choices are presented to players, who also usually receive complete information about the payoff structure and previous decisions of their co-players. Giving complete information on other players' preferences and strategies, however, is not possible. All information belonging to the social environment is reduced by assuring anonymity from

each other and sometimes even from the experimenter. The main research question is whether or not human subjects actually do behave as the rational players von Neumann and Morgenstern assume in their theory. Experimental evidence suggests that sometimes they do, and sometimes they do not, which leads to the more interesting question of when and why discrepancies between actual human behaviour and this theory appear.

There are situations in real life and in the laboratory which seem to be not very well represented by game theory. Additionally, extensions to game theory did not succeed in capturing the essence of these situations and make them accessible to mathematical analysis. As an example, consider the chain store paradox, which is presented in the following in some detail (Selten (1978)):

> "A chain store, also called player A, has branches in 20 towns, numbered from 1 to 20. In each of these towns there is a potential competitor, a small businessman who might raise money at the local bank in order to establish a second shop of the same kind. The potential competitor at town k is called player k. [...]
>
> Just now none of the 20 small business men has enough owned capital to be able to get a sufficient credit from the local bank but as time goes on, one after the other will have saved enough to increase his owned capital to the required amount. This will happen first to player 1, then to player 2, etc. As soon as this time comes for player k, he must decide whether he wants to establish a second shop in his town or whether he wants to use his owned capital in a different way. If he chooses the latter possibility, he stops being a potential competitor of player A.
>
> If a second shop is established in town k, then player A has to choose between two price policies for town k. His response may be 'cooperative' or 'aggressive'. The cooperative response yields higher profits in town k, both for player A and for player k, but the profits of player A in town k are even higher if player k does not establish a second shop. Player k's profits in case of an aggressive response are such that it is better for him not to establish a second shop if player A responds in this way." (Selten (1978, 127f.))

The game theoretical prediction of player A's behaviour is that of choosing cooperation, whenever a second shop is opened. The other players are assumed to predict this behaviour and open their shops if their time comes.

If, on the other hand, player A chooses an aggressive response early in the game, the other players are discouraged from opening their shops, and player A can make much more money later on because of the lack of competition. Selten opposes the game theoretical "induction theory" with a deterrence theory. The latter is based on the assumption that player A *imagines* responses of later players k and thus tries to discourage them from opening a shop by choosing aggressive responses early in the game. This power of imagination is very hard to include in game theoretical models. Further, simpler examples for games in which the rational actor theory does not predict experimental observations very well, are given in the Classification working paper (chapter 8).

Game theory has been extended to be able to explain some of these "paradoxes". One such extension is behavioural game theory (Camerer (2003)). This aims at incorporating norms, emotions and other-regard in the payoff structure of games. That is, if the other player's utility enters the first player's utility function a prisoner's dilemma game can be transformed to have a completely different payoff matrix. Cooperation need no longer be a dominated strategy, and instances of cooperation can be explained. An example for such a transformation is given in the Classification working paper (Chapter 8) in Figure 8.4. The underlying assumption of utility maximisation remains. Self-interest is understood in a broader sense, making it possible to include norm-compliance, utility deriving from negative or positive emotions, and the utility of other players.

Game theory and its extensions serve as benchmarks for alternative models of decision making behaviour. For our purposes, however, it does not provide usable models of human behaviour for two reasons. The first reason is that the focus of our research lies precisely in such decision environments in which the rational actor theory of human behaviour fails to explain the actual behavioural regularities. The second reason is, that it may be substantively rational but it is definitely not procedurally rational (Simon (1997, 293f.)). If a model of human behaviour is used to investigate probable reactions on institutional change, it is necessary to extract and reproduce those regularities which react on this change. Therefore, we can not focus entirely on the *outcome* of a decision making process, but need to model the *process* itself. In his appeal for behavioural economics, Simon calls for an empirical base of processes of decision making:

> "The economic actors' rationality will be defined by the processes they use in making their decisions rather than the substance of the decisions they reach; and the processes they employ cannot

> be deduced from the objective description of the problem situation; they can only be determined inductively from empirical observation or inferred from empirically based theories of behavior" (Simon (1997, 271))

Taken seriously, this calls for more changes than behavioural game theory is willing to make. Experimental economics strives to provide data for this, but as yet, there is no theory able to combine data into one coherent approach.

There are social scientists who strive to build theory on different premises than the rational actor theory. If it is recognised that important human behaviour, as for instance in social dilemmas, can only be modelled as rational choices with additional assumptions or great difficulty, another kind of theory may be useful. "While the rational choice approach is forced, either to adopt the ad hoc solution of treating such behavior *as if* it where based on rational calculation, or to regard it as falling outside of its explanatory domain, the notion of program- or conjecture-based behavior allows one to include the study of conscious rational choices in a broader theoretical framework" (Vanberg (2002, 27)). One kind of program-based behaviour is discussed in the subsection on bounded rationality (Subsection 1.3.3).

1.3.2 Psychological models of human behaviour

I have not delved deeply into psychological theories. The main reason was the fear, not only on my part, that after such an endeavor I would not resurface in time to complete the thesis. There are, however, a number of psychological theory bits which have influenced the attributes and heuristics approach presented here.

A discussion of the attributes is presented in the ESSA paper (Chapter 3) and the WCSS paper (Chapter 6). However, in order for the reader to be able to compare these attributes to theories presented here, a list is given:

1. Cooperativeness

2. Conformity

3. Fairness (concerning others)

4. Fairness concerning me (that is: envy)

5. Positive reciprocity

6. Negative reciprocity

7. Risk aversion

Trait approaches

In psychology there are a number of different approaches to explain human behavioural regularities and differences. One of these approaches is the trait approach, in which differences are explained by differing dispositions of individuals (Liebert and Spiegler (1994)). Which traits are used depends on the theory. An influential theory is the BIG5 or OCEAN theory, in which the main personality traits are identified as openness, conscientiousness, extraversion, agreeableness, and neuroticism. These five big traits include a number of subtraits, which seem to be correlated in individual humans. That is, a person with a high openness tends to score high on the subtraits identified with openness. That does not mean, that all subtraits are equally high for a person. Actually, the argument usually is the other way around: a person scoring high on many subtraits of openness scores high on openness.

Openness A person high on openness tends to be imaginative, curious, creative, adventurous, original, and artistic. A person low on openness is conventional, avoids the unfamiliar, is inartistic, and lacks imagination.

Conscientiousness A person high on conscientiousness is cautious, disciplined, organised, neat, ambitious, goal-oriented, and has a high need for structure. A person low on conscientiousness tends to be unreliable, lazy, careless, negligent, and low on need for achievement.

Extraversion Extraversion is associated with being talkative, friendly, optimistic, sociable, and high on need for stimulation. A person low on extraversion is inartistic, quiet, conventional, less assertive, and aloof.

Agreeableness A person high on agreeableness is compassionate, goodnatured, trusting, and helpful. A person low on agreeableness tends to be irritable, rude, uncooperative, and unsympathetic.

Neuroticism Neuroticism is associated with emotional instability. A person high on neuroticism is anxious, nervous, worrying, insecure, emotional, feels excessive cravings or urges, and entertains unrealistic ideas. Persons low on neuroticism tend to be relaxed, calm, secure, unemotional, and even-tempered.

Apparently, there are interrelations of subtraits. In addition, one can imagine a person being, for example, both optimistic and quiet. With questionnaires on subtraits, psychologists try to ascertain patterns of personalities which can be ascribed to one of the five big traits being more pronounced in a person's personality than others.

The attribute approach presented here has similarities to psychological trait approaches, because it bases agents' differences on differences in their attribute values. The attributes in use, however, are quite distinct from the BIG 5 presented above. Cooperativeness may be associated with agreeableness, but so is fairness concerning others and positive reciprocity. We can, therefore, assume a correlation between these three attributes. In the models, however, this has not been done. Conformity can be associated with low openness, negative reciprocity with low agreeableness, but low negative reciprocity could also indicate low neuroticism. In order to develop a psychologically sound basis for the attributes presented here, it would be necessary to make a thorough investigation in the interrelation between subjects' behaviour and their scoring on different traits. This investigation should contrast the attributes presented here with traits from different psychological trait theories.

An approach to this has been made by Ben-Ner *et al.* (2004) who compared scores on the BIG 5 traits in questionnaires with dictator behaviour. Among other interesting facts, they found that men and women scoring high on agreeableness shared significantly (but not very) more with their recipients, while women scoring high on conscientiousness and neuroticism shared less, and men scoring high on extraversion shared less. Brandstätter (1988) developed an alternative questionnaire in order to determine people's scores on traits. This has also been used in a survey to compare prior dispositions to dictator givings (Brandstätter *et al.* (1999)). They find a correlation between self-reported benevolence and dictator offers, but no significant correlation between tested benevolence and dictator offers.

Social Value Orientation

A comparison of our attribute approach to the social value orientation theory highlights our understanding and use of attributes. McClintock (1972b) introduces social value orientation based on an analysis of behaviour in prisoner's dilemma games (PDG).

> "Rather than attempt to determine what is or should be rational behavior for the players of PDG, most social psychological

	II			**II**	
	C	D		C	D
C	(5, 5)	(-4, 6)	C	(6, 6)	(0, 5)
D	(6, -4)	(-3, -3)	D	(5, 0)	(0, 0)

(Player **I** labels the rows in each matrix.)

Figure 1.1: Payoff matrices in Prisoner's Dilemma (left) and Maximising Difference
Game (right) in the configuration of McClintock [1972].

> research today utilizes the game as a paradigm for investigating
> behavior in situations of social interdependence in which play-
> ers have the option of cooperating or competing, and asks what
> types of variables influence the level of cooperative and com-
> petitive behavior of individuals in such situations." McClintock
> (1972b, 276)

According to McClintock such variables include numerical values in the
game matrix, players' predisposition towards cooperative behaviour, prior
experience, and communication possibilities.

He then proceeds to identify four motivational goals a given player may
pursue: (a) cooperation: maximisation of joint gain, (b) individualism: max-
imisation of own gain, (c) competition: maximisation of one's own gain rel-
ative to the other's gain, and (d) altruism: maximisation of other's gain
(McClintock (1972b, 278)). However, the prisoner's dilemma does not al-
low distinguishing clearly between the four motivations. Leaving the rarely
observed altruism aside, cooperation in prisoner's dilemma games can be as-
sociated with cooperation or individualism, and defection can be motivated
by competition or, again, individualism.

McClintock discusses the Maximising Difference Game (MDG) in or-
der to alleviate this problem (see Figure 1.1). Mutual cooperation yields
6 points to both players, while mutual defection yields zero points. If one
player defects and the other cooperates, the defector gains 5 and the co-
operator nothing. In this game, a clear distinction between individualism
and competition can be made, because maximisation of own gain leads to
a cooperative choice and maximisation of relative gain to defection. In this

game, individualism and cooperation can not be distinguished.

McClintock assumes underlying motivations or goals and designs games to make the different possible motivations discernible. He uses questionnaires in addition to the game results to ascertain underlying motivations. By this process he is able to distinguish two dimensions of motivations. Cooperation and competition are on one axis and mutually exclude one another, and individualism and altruism are on the other axis. It is possible, for one individual to be both individualistic and cooperative, or individualistic and competitive. The social value orientation theory, therefore defines two attributes, cooperativeness and individualism. A low cooperativeness can be associated with competitiveness, and a low individualism with altruism.

The principle of this theory corresponds to our use of attributes, although the attributes do not. We assume, like McClintock, individuals to have a certain cooperativeness which comes into play in situations, where there is a choice between cooperation and competition to make. However, contrary to McClintock, we assume further that independently of this predisposition, the same individuals have a certain inclination to treat others fairly on the one hand and to be annoyed if others do not treat them fairly on the other hand. These two are important in situations, in which the differences in outcomes are more important than the sum. In our model, the same person can act altruistic in one situation and individualistic in another, because the two traits are modelled as two attributes instead of one.

1.3.3   Bounded rationality and the adaptive toolbox

Being unsatisfied with the predominance of rational actor theory, Simon (1997) focused his research on deviations from rationality, namely satisficing and information limits, and thus coined the term bounded rationality. At the same time, it was very important to him, that this behavioural theory was grounded in empirical research (Augier and March (2003)). There are two different kinds of bounds on our rationality, *external*, like information scarcity and limited time, and *internal*, like limited computational and memory capacity of our minds. Focussing on external limitations leads to models of rationality under constraints, while focussing on internal bounds leads to "irrationalities" like cognitive biases, which constitute systematic violations of the rationality paradigm (cf. Todd and Gigerenzer (2003)). If these two sets of bounds are considered to be linked, however, the limits of the human brain can be seen as an advantage with respect to structure of those types of environments within which most of our tasks are situated. As Simon (1990) put it, these two constraints are like a pair of scissor blades

that need to fit together. One scissor blade refers to the environment of a decision, the other to the decision making process. If the decision making process does not fit to the decision environment, decision making does not result in very successful decisions, the scissors do not cut. If, on the other hand, a heuristic fits exactly to some problem, the decision outcome may be very good, even though this decision making process has nothing to do with optimisation. An example is the following heuristic to catch a ball: look at the ball and run so that the angle in which you look at the ball stays the same. Robots can catch balls based on this heuristic better than trying to calculate the point where it will hit the ground on the basis of three points in the path of the thrown ball. The heuristic "keep the same angle" does not work, if the task is to avoid the ball.

On this definition and understanding of bounded rationality rests the research by Gigerenzer *et al.* (1999) on heuristics humans actually use and in what circumstances they are useful. Some very simple heuristics are highly useful in certain environments and not useful in others. The decision environment determines the tools we need to make a decision.

This idea of bounded rationality views human decision making as using an adaptive toolbox (Gigerenzer and Todd (1999); Gigerenzer (2001)). The tools are simple heuristics and their use is adapted to a given decision process. Using simple heuristics instead of maximisation procedures is psychologically plausible. These heuristics work with the cognitive, emotional, social, and behavioural repertoire, humans actually have. Being simple, these heuristics necessarily do not work for any and all decisions. Instead, they are domain specific. Some work in some decision environments and others in other environments. The match between heuristics and environmental structures is precisely what makes them work. In the models in this thesis, agents use different heuristics in different decision environments. Heuristics have different degrees of generality. There are very simple, general heuristics that can be used in a greater variety of decision environments, and there are heuristics that fit to one decision environment only.

Decision making according to the rational actor theory can be seen as a special case, a special tool in the toolbox. Maximisation works very well in a certain kind of decision environment, characterised as "stable, competitive market settings and in competitive electoral and legislative settings where the issue space is constrained" (Ostrom (2005, 100)). But in other contexts maximisation is not appropriate, for instance if strong emotions are involved, if information is scarce, if multiple goals need to be taken into account, without an easy way to weigh them, and if social constraints like morality and accountability limit the decision space.

In their current research program, the ABC research group around Gigerenzer and Todd endeavor to ascertain some of the elements of the adaptive toolbox containing simple heuristics we humans really use, along with the decision environments in which the different heuristics are applicable (Gigerenzer and Todd (1999)). The heuristics researched are defined so precisely that they can be coded as computer programs, which is very useful for our modelling purposes. However, the decision environments investigated in that research do not match the social dilemmas and public goods environments we are interested in. They investigate real-world environments, with a focus on search processes for alternatives. The economic experiments used for our models, in contrast, present all possibilities to subjects, eliminating the need for search. However, the simplicity and precision of their heuristics and building blocks matches that of the heuristics used in our models, while at the same time, their heuristics are more generic than ours.

Selten (1978) developed a three stage theory of decision making. According to this theory, a decision can be made on the routine level, the level of imagination or the level of reasoning. The level of imagination makes use of the routine level, and the level of reasoning employs both, the level of the routine and the level of imagination. A predecision is made in order to decide which level to use in a given decision. This is made on the routine level. A final decision decides between possibly different decisions of the three different levels. This, again, is a routine decision. It is possible, that a final decision decides against the "higher level" decision of reasoning. The theory is concluded by the "decision emergency hypothesis [that] maintains that conscious thinking does not have the task of making decisions. The rational mind is like an adviser to a king. The king is a subconscious hidden mechanism who makes the final decision. The king may or may not listen to the adviser. The adviser does not necessarily have to propose a decision. In some cases he may restrict himself to pointing out advantages and disadvantages of various alternatives. Decisions are not made, they emerge." (Selten (1999, 133))

For search processes Selten developed an aspiration adaptation process on the grounds of Simon's idea of satisficing. According to Simon, humans do not usually maximise a function when making a decision. Rather, they "satisfice", meaning that they choose a satisfactory option if it occurs to them instead of looking for *the* optimal one. Only, if all known options are not satisfactory, they search for new possibilities. Taking this idea, Selten (2001) developed a mathematical description of this process, the Aspiration Adaptation Theory. A decision problem can be modelled as a multi-goal problem without the need to make compromises between the goals or aggre-

gating them into a joint goal. For each goal, it is assumed that a person has an aspiration level, beyond this level an outcome is perceived as satisfactory, below it, it is unsatisfactory, the aspiration is not met. Its mathematical description views decision making as moving in steps on a multi-dimensional grid. The goals make up the dimensions and the grid is determined by the options. Each decision yields an outcome vector on this grid. If the aspiration for one or more goals is not met, a search is triggered to improve the outcome of this goal. The search is guided by rules that can employ different levels of urgency of the goals, and an aspiration adaptation process. The latter can reduce aspiration levels, which prove to be too difficult to fulfill. A simple version of this aspiration adaptation process has been included in the later models of this thesis (see the MABS paper, chapter 4).

1.3.4   Agent-based modelling

Agent-based modelling has developed into a broad field, broad both in its use of techniques and its application. The focus here lies on agent-based social simulation, that is agent-based modelling with an application within the social sciences. These models are simulations that aim to replicate, to some extent, a part of the real world. Using agent-based modelling to simulate social processes has the obvious advantage of modelling micro behaviour at the micro level. Human decision making does not have to be aggregated into average behaviour, but can be diverse and depending on social interaction, which can also be modelled explicitly. The adaptation of a model to a certain real-world or experimental problem can be done by adjusting at the micro or the macro level or both, which is a prerequisite, when the models are to be used in different circumstances, including environmental management (Hare and Deadman (2004)).

   According to Davidsson (2002, 1.4) agent-based social simulation investigates "the use of agent technology for simulating social phenomena on a computer". There is also a focus on the appropriateness of these tools. In case of this thesis, an investigation into the appropriateness of using agent-based models to model human decision making with heuristics encompasses a survey of recent models of agents using heuristics instead of maximisation processes or fixed decisions. We first focus on models using data from experimental economics and then proceed to investigate models using data from field studies.

Data from experimental economics McClintock (1972b) states that game behaviour is behaviour in a social situation in which the players are interdependent, meaning that each can influence but not completely determine the outcome of all players. "For this reason, games provide a convenient paradigm for investigating the motives that underlie and the strategies that evolve in situations of social interaction" McClintock (1972b, 271). This is precisely the reason, why experimental economics provides invaluable insight and a useful database for modelling the motivations and strategies of real humans. The process to use data from economic experiments for agent-based models is labeled by Tesfatsion (2002, 16) as "Parallel experiments with real and computational agents". This approach is relatively new. Studies with this approach include Deadman *et al.* (2000) and Deadman and Schlager (2002), Duffy (2001), Jager and Janssen (2002), and Dal Forno and Merlone (2004).

- Deadman and Schlager (Deadman *et al.* (2000); Deadman and Schlager (2002)) implement agents with bounded rationality. Their agents have complete preferences and information, but limited computational capacity. They can choose from among 16 strategies to make their next decision. The performances of all strategies are measured, and the strategy that would have performed best is chosen in a similar decision next time. This is a model using the experiments of Ostrom *et al.* (1994) and can therefore be compared to the model implemented in this thesis (see Chapters 4 and 5).

  Choosing from a limited set of strategies is similar to the approach presented here, but the way in which strategies are chosen differs. Also, contrary to Deadman *et al.* (2000), our agents exhibit diversity through differences in attribute values.

- The agents in Duffy (2001) maximise their expected utility, that is they choose the best option if it is obvious. The expectation formation for the non-obvious cases uses a learning mechanism that is based on prior experience in order to capture deviations from the predicted rational behaviour exhibited by experimental subjects.

  Similar to the approach here, decision making depends on the case. In an obvious case, a simple decision (take the better option) is used, and in the complicated case, a learning mechanism is employed. A similar learning mechanism to the one implemented by Duffy (2001) is possible to implement within the framework.

- Jager and Janssen (2002) divide their agents' decision making process into four categories: deliberation, social comparison, imitation, and repetition. The first is based on utility maximisation, while the other three categories can be viewed as different forms of heuristics. Agents choose one of these four categories depending on the situation and their satisfaction with a previous choice.

  This process is similar to our choice of heuristics, but their heuristics are more complicated than ours. We have not yet implemented social comparison, because the focus in this thesis is on games with anonymity. There are, however, heuristics based on average previous decisions of other agents, which can be seen as imitation.

- Dal Forno and Merlone (2004) had their human subjects in the experiments indicate heuristics underlying their decisions and coded these heuristics in their agent-based models. This yielded not only agents employing very realistic heuristics as their decision making mechanism, but also a distribution among agents using different heuristics similar to the experiment.

Data from experimental economics   Models using heuristics and simulating real-world situations are more abundant. We focus here on some that explicitly use heuristics for agents' decision making process, and do not use an evolutionary approach to determine the best heuristics. It is an objective of this thesis to provide a basis for models of this kind, so that models of different real-world situations can be modelled in a common framework, or become comparable to some extent.

- Berger (2001) gives his heterogeneous agents different thresholds of adoption of a new technology. When the threshold is reached, the agents calculate possible benefits and adopt the technology if expected benefit is positive.

  Using a threshold is similar to decision making of our agents, although our mechanisms to determine benefits of an action are not as sophisticated.

- Kottonau and Pahl-Wostl (2002) model attitude strength towards one of two parties in an election campaign, by direct communication with other agents, who also have attitudes. Attitudes then change due to the nature of this communication.

- D'Aquino *et al.* (2003) had farmers design a role-playing game and developed an agent-based model on the basis of the game.

  This connection of role-playing games and agent-based modelling is a process that we expect to work very well with our modelling framework. Creating a model can be greatly facilitated by participants of a role-playing game, when they can inform modellers of their decision heuristics used in the game. The model, on the other hand, can be more complex and more often repeated than a game and therefore be used to explore a greater range of possibilities.

- Collentine *et al.* (2004) model decisions as consecutive nodes of either deciding to adopt a new technology, to reject it, or to gather more information. Only the third option leads to the next similar node. By this mechanism, their agents do not have to calculate the cost of gathering more information.

  Although these three modes (adoption, rejection, or gathering more information) of decision making and reflection, are not implemented in such a stringent way in our model, the idea is present. If satisfied with a previous option, no effort is made in determining whether there is a better option or not. Only if not satisfied, an agent enters a meta decision process, possibly leading to a change of heuristic. This can also come up with the same heuristic as before, if the data do not suffice to decide on a new heuristic.

1.3.5 Attributes and heuristics

The attributes in use are cooperativeness, conformity, fairness concerning others, fairness concerning me, positive reciprocity, negative reciprocity, and risk aversion.

In order not to repeat myself, the discussion of the seven attributes is made here with a different focus than in the papers. For a discussion of the attributes from an experimental economists viewpoint, the reader is referred to the ESSA paper (chapter 3), for a topology of the attributes to the WCSS paper (chapter 6). The focus in this subsection lies on the motivation for this approach, on the interplay of attributes and heuristics, as well as on a response on the remark of an anonymous reviewer that attributes are meaningless, when used only in combination with heuristics.

The fact that attributes appeared in the first model was based on the brief discussion of their altruistic punishment experiment of Fehr and Gächter (2002) in *Nature*. The authors state their hypothesis: "Free riding may

cause strong negative emotions among the cooperators and these emotions, in turn, may trigger their willingness to punish the free riders" (Fehr and Gächter (2002, 139)). Post-experimental questionnaires confirm this hypothesis. Subjects were asked to indicate their level of anger or annoyance when they invested much and another player invested nothing or only a little. The result is that generally the higher the deviation, the stronger is the annoyance. But of course, levels of annoyance are different among subjects, as are punishment reactions, as well as investment levels in the first place. Consequently, the first three attributes used to model this particular experiment were "cooperativeness", "inclination to be annoyed", and "willingness to punish" (see chapter 2).

During modelling of the second and third experiment and some further reading of other agent-based models, I came up with a list of twenty-something important attributes. Some defining and refining decreased the number to nine. For example, altruism is not needed as an attribute, because it is, depending on the situation, either cooperativeness, fairness concerning others, or positive reciprocity. Of the remaining nine, two more became obsolete because they could be reduced to one or a combination of the remaining seven. For example, trust is modelled as either expected cooperativeness, expected positive reciprocity, or expected conformity.

In our models, attributes are used in three different ways. The first is that they determine the choice of the initial heuristic. An example is that cooperative players, those with a cooperativeness higher than for instance 0.67, use a cooperative heuristic, and those with a lower cooperativeness use a different heuristic. Of course, this can also be incorporated in the heuristic itself, which is the second way. Most heuristics use one or more attributes and expected attributes to determine the decision value. For example, the `ReciprocalExpectationStrategy` decides on the percentage of a game-defined maximum that corresponds to the expected cooperativeness of a player. The third and most complicated way is that they can be used in the learning mechanism or by a meta search strategy (which is a strategy used in this framework to alter decision making processes of agents). For example, the number of instances an agent can suffer defection from others, before defecting itself, depends on that agent's cooperativeness.

The interplay of attributes and heuristics is an important mechanism in the framework, highlighted in Figure 1.2 and explained here in some depth. The decision environment determines which of the attributes are used for making a decision, in Figure 1.2 attribute 2 is chosen. For example, in a common pool resource game the attribute in question is cooperativeness. At the same time, the decision environment determines the pool of possible heuris-
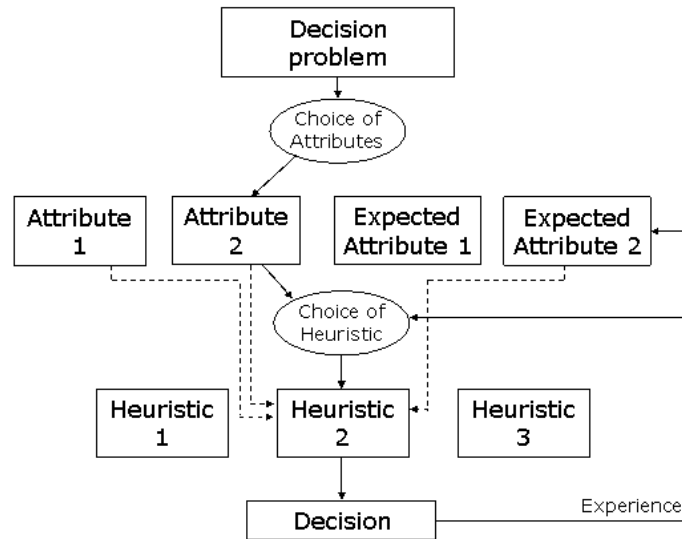
Figure 1.2: Individual attributes play an important role in the decision making process. The choice of attributes is influenced by the nature of the decision problem. Attributes influence the choice of heuristics and may be used within heuristics.

tics, in Figure 1.2 these are heuristics 1 to 3. In a common pool resource game, these could be maximising, reciprocal, or cooperative heuristics.

The three ways of combining attributes and heuristics indicated above are:

1. The attribute value can influence which heuristic or decision strategy is used, in Figure 1.2 the value of attribute 2 results in the choice of heuristic 2. In a common pool resource game, a high cooperativeness results in choosing cooperative heuristic.

2. The employed heuristic can use attribute values and expected attribute values to make the actual choice. In Figure 1.2 heuristic 2 uses attributes 1 and 2 as well as expected attribute 2. The reciprocal heuristic in a common-pool resource game uses the expected cooperativeness in order to decide on the cooperation level, which is the actual decision.

3. The resulting outcome of a decision determines the experiences. This learning process not only alters the expectations and the use of heuristics, but also may be influenced by attribute values. In a common pool resource game, for a very cooperative player to choose a maximising heuristic, it needs more negative or uncooperative experiences, than for a player with a lower cooperativeness.

As has been pointed out to me, a model of a specific experiment could be modelled without any attributes, because these can be incorporated in heuristics and learning mechanisms. However, the meaning of attributes emerge from their use in different models or even in different games within the same model. For example, whether or not punishment acts are correlated with investment decisions can be implemented in this way. These two decisions are part of different games. The agent is the same. If its investment decision and its punishment decision depends on its cooperativeness, the two are modelled as being correlated. As it happens, it was not modelled in this way, because the correlation is unclear. The height of punishment and the height of defection are correlated and therefore there is a link to the original investment decision, but it has been modelled via this link and not directly. It seems to me to be important to model both, the unchanging attitude of an agent, and the volatile use of heuristics. Therefore, these two aspects are modelled in two different ways.

1.3.6 Implementation

One objective of the implementation was a modular design that would allow easy extensions to the framework. This was successful. The later models have been implemented in much shorter time, because it was possible to draw on classes used in earlier models. Model alterations are usually simple to conduct and fast to implement.

In appendix A an overview of packages and classes in the framework is given. The CD provides the program code and Java documentation, as well as models that can be run and altered in the Famoja[2] environment.

For the agent-based models, object-oriented programming was used. This is not only a certain technique but also a specific way of structuring thoughts about a model. In this respect, it is similar to agent-based modelling and suited well. In order to create reusable classes and functionality, processes need to be broken down into little pieces that can be used as abstracted primitives. In case of this thesis one prominent example for this break-down process are the parts of a decision heuristic. In the models a decision process is usually build up from a number of smaller processes:

1. Which attribute is relevant?

2. According to the attribute value(s), which heuristic is used?

3. The heuristic itself may be composed of several building blocks.

4. The learning process may involve several feedback mechanisms.

As one result of this abstraction process, the set of attributes could be derived. Another result is a differentiation into meta decision heuristics (which heuristics are possible in the current situation?), search decision heuristics (according to some feedback, which heuristic is chosen next?), and decision heuristics (that actually make the decision). At first, this structure was useful in altering the models and introducing new decision environments. Later, it was also used for theory building.

A tutorial on the webpage[3] introduces the modelling technique developed in this thesis. It is not necessary to be familiar with the framework Famoja, but the tutorial is based on the assumption that the user knows Java. The tutorial is also included on the CD. It uses the prisoner's dilemma game to illustrate how to implement in successive steps more and more complicated models of prisoner's dilemma games.

---

[2]http://www.famoja.net/

[3]www.usf.uos.de/∼eebenhoe/forschung/adaptivetoolbox/tutorial/

1.4   History

The six articles and one working paper are listed below under the abbreviations used in this thesis, together with a brief statement on their content and publication status.

JASSS 2004 (chapter 2): This article is titled "An adaptive toolbox model. A pluralistic modelling approach for human behaviour based on observation" and published in the Journal of Artificial Societies and Social Simulation (JASSS) in January 2004. It describes the modelling approach and the first model implemented in this approach, the altruistic punishment experiment by Fehr and Gächter (2002).

ESSA 2005 (chapter 3): This paper is titled "Modelling human behaviour with attributes and heuristics based on observation" and was submitted in May 2005 to the third annual conference of the European Social Simulation Association (ESSA05). It is published in the conference proceedings. This paper describes the modelling approach including the seven attributes and exemplary models of ultimatum games.

MABS 2005 (chapter 4): This paper is titled "Modelling non-linear common-pool resource experiments with boundedly rational agents". It was submitted in May 2005 to the workshop on Multi-Agent-Based Simulation within the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS05). An extended version is published in the post-conference publication: Sichman, J.S., Antunes, L. (eds.) 2006. MABS 2005, LNAI 3891, Springer-Verlag, 133-146. It is an extensive description of the appropriation model based on Ostrom *et al.* (1994) for an audience of modellers.

Handbook 2006 (chapter 5): The chapter is titled "Agent-based modelling with boundedly rational agents". It is to be published in Jean-Philippe Rennard's book "Handbook of Research on Nature Inspired Computing for Economy and Management", published in 2006. This chapter contains a description of the modelling framework and two exemplary models, one being an enhanced altruistic punishment model and the other the appropriation model, based on Ostrom *et al.* (1994).

WCSS 2006 (chapter 6): This paper is titled "A Topology Of Agent Attributes For Modelling Microbehaviour in Economic Experiments" and introduces not only the attributes used in the models, but also relates them to one another through defining their different uses. It draws

on generalising thoughts about the model building process, not published but presented here in chapter 8. It has been accepted by the First World Congress on Social Simulation in August 2006 and will be published in the conference proceedings.

Behaviour 2006 (chapter 7): This paper is titled "Agent behaviour between maximisation and cooperation" and introduces the approach to a non-modelling audience, drawing on results from these models, but not presenting an agent-based model itself. The line of argumentation of this paper is a rough classification of agent-behaviour in three categories: maximisers, cooperators, and waverers. The percentages of all three groups varies with the setting and experiences made. The waverers are the easiest to influence in either direction. This paper has been submitted to Rationality and Society in June 2006.

Classification (chapter 8): This working paper is titled "Classification of decision environments". It is about work in progress and is not submitted. This working paper can be read as a conclusion, summarising insights from different models in this thesis. It develops the beginning of a framework to decide in which decision situation, which attributes and heuristics are appropriate.

For these papers to be stand-alone papers it was necessary to repeat certain aspects. In all but the first paper the attributes are described in different length. In the JASSS paper, the attribute approach was not developed. The most extensive discussion of attributes is in the ESSA paper and the WCSS paper. The Classification working paper places the attributes and heuristics in context with decision environments and provides a longer basis for the WCSS paper.

For an illustration of the development of the modelling approach, the JASSS paper can be compared to the ESSA paper, which target the same audience, a social simulation audience. The same audience is addressed by the WCSS paper, although this does not provide a model, but focusses on the attribute approach. The MABS paper also has modellers, but not necessarily social simulators as a target audience, while the Handbook paper and the Behaviour paper try to extend the audience to possible users of agent-based models for other purposes.

The Handbook paper draws on models of the altruistic punishment experiment and the appropriation experiments, which have already been described in other papers. The altruistic punishment model described in the

Handbook paper, however, is a version implemented more thoroughly according to the previously not developed attribute approach and framework. Its implementation is very different from the earlier version, described in the JASSS paper, although the underlying decision making mechanism is not. The MABS paper is concerned with the same model as the Handbook paper, but explains it in more detail.

The ESSA paper contains curt descriptions of dictator and ultimatum models, which have been described in more detail in a working paper accessible on the website[4]. The dictator model also appears in the Behaviour paper in a different version, where it is compared to a model of an investor-trustee experiment. The latter is also described in more detail in a working paper accessible on the website.

Since all submitted papers deal with only one or two out of a number of models, the Classification paper is valuable because it comprises findings from the whole range of models.

---

[4] <www.usf.uos.de/~eebenhoe/forschung/adaptivetoolbox/>

Chapter 2

An Adaptive Toolbox Model: a pluralistic modelling approach for human behaviour based on observation

Claudia Pahl-Wostl and Eva Ebenhöh[1]

Abstract

This article describes a social simulation model based on an economic experiment about altruistic behavior. The experiment by Fehr and Gächter showed that participants made frequent use of costly punishment in order to ensure continuing cooperation in a common pool resource game. The model reproduces not only the aggregated but also the individual data from the experiment. It was based on the data rather than theory. By this approach new insights about human behaviour and decision making may be found. The model was not designed as a stand-alone model, but as a starting point for a comprehensive Adaptive Toolbox Model. This may form a framework for modelling results from different economic experiments, comparing results and underlying assumptions, and exploring whether the insights thus gained also apply to more realistic situations.

Keywords: Social Simulation, Experimental Economics, Common Pool Resource Games, Adaptive Toolbox, Altruistic Punishment

2.1   Introduction

Understanding behaviour of human beings in complex decision making situations is of vital importance for the design of appropriate institutions for sustainable resources management and for managing transition processes towards more sustainable resource management regimes. Most initial efforts have relied on game theoretical approaches and extensions thereof. However, a number of common pool experiments and other empirical evidence showed clearly that the assumptions on rational behaviour are not supported by observation (Ostrom (2000)). Researchers explored for example the importance of trust, reciprocity, and reputation to introduce and stabilize social norms of cooperation in a group (e.g. Hayashi *et al.* (1999)).

Social simulation may play an important role to develop an improved representation of the complex dynamics of human-environment systems.

One may distinguish four different approaches:

1. Start from an established formalized theoretical framework (e.g. rational actor paradigm, game theoretical approach) to test conditions of applications and the consequences of relaxing certain assumptions. The extension can be based on observation or principle considerations about the deficiencies in the framework (cf. Lindenberg (1991)). In general such approaches remain within the boundaries of a given framework.

2. Combine concepts from different social sciences to develop an interdisciplinary approach and build a simulation framework (e.g. Jager *et al.* (2000), Epstein and Axtell (1995), Kottonau (2002)) as experimental laboratory to explore the implications of different assumptions on system dynamics. This is a potentially very rewarding exercise that may support to overcome the fragmentation of different streams of theories in a field.

3. Start from very simplified rule-based representations for social behaviour that are more determined by considerations of complex dynamics rather than explicit social science theoretical considerations. This is the approach of socio-econo-physics or the simple models from Thomas Schelling exploring the importance of spatial interaction for

racial segregation. The recent article by Deffuant *et al.* (2002) and the subsequent comments (Randow (2003); Deffuant *et al.* (2003)) provided an illuminating example for arguments in favour and against the sociophysics approach.

4. Start from observation and extract regularities of behaviour (e.g. Todd and Gigerenzer (2003)). This approach does not claim to achieve an overall synthesis by combining existing theoretical frameworks. But it is guided by some assumptions on human behaviour (in our case the importance of heuristics and different basic dispositions towards cooperativeness in human beings) and relies first of all on observation (also in our case the design of the experiments was not free from theoretical assumptions).

Each approach has its strengths and weaknesses. The work presented in this paper is based on the last approach since we are convinced that starting from observation is required to promote real innovation and integration. The insights derived from a more inductive approach should, however, be confronted with established theories to explore possible contradictions or coincidences. In the end a fruitful exchange among all the approaches outlined above will promote insights and change.

In this paper we present the idea and the first steps towards the implementation of an adaptive toolbox (cf. Gigerenzer and Selten (2001a)) as a multi agent system with diverse agents that can behave differently according to different situations and contexts. In order to capture realistic human behaviour the model is founded in experimental data rather than theory. In contrast to many approaches that use the data to test and extend current disciplinary theories we use the data to extract similar patterns and heuristics that determine human behaviour. By doing so we apply an interdisciplinary approach in the social sciences. We expect that such an approach will promote real innovation in our understanding of human behaviour and its representation in models. It is our goal to contribute to a coherent simulation framework that allows to explore different perspectives on human behaviour and compare their strengths and weaknesses as well as their applicability in different situations. Hence, we make a strong plea for a pluralistic approach.

WHY an adaptive toolbox? Representation of human behaviour in models is still surrounded by huge uncertainties and major controversies. This may be attributed to a lack of an overall accepted interdisciplinary

approach in the social sciences. Many different and partly contradicting approaches for explaining human behaviour coexist within social science disciplines and even more in the different disciplinary approaches. The most formalized theory is the rational actor paradigm (referred to as RAP in the following) in economics. Arguably the success of this approach can be attributed to the fact that formalization provides a base for better communication and unification and to the simplicity of the concept. The supporters of the RAP have refused for years, even decades, to acknowledge criticism and scientific arguments providing evidence for some weaknesses in their assumptions. However, the situation is slowly changing. The rational actor paradigm of economics becomes enriched by insights from psychology and sociology.[2] Some argue even for a more radical approach to abandon the RAP entirely, to move from omniscience and perfect foresight towards more simplified and realistic descriptions that are not due to imperfections of the human brain but evolved to guarantee survival in complex and dynamic environments (Gigerenzer and Selten (2001a)). It must be the goal of social simulation to come up with strong alternatives to the RAP. One should move from the dominance of a single concept to a pluralistic approach with different perspectives on human behaviour that take into account the importance of context and the diversity of human beings. Social simulation should provide the base for an interdisciplinary framework that allows to combine different aspects of human behaviour that are all required to fully understand the complexity of human systems. Without imposing the constraining rigour of analytical mathematics, any simulation approach forces the analyst to be more consistent in his/her assumptions. Development of coherent simulation frameworks will foster the development of more comparative and interdisciplinary approaches. This was also highlighted by Epstein and Axtell (1995) in their pioneering work on artificial social societies, the sugarscape model.

Progress in the representation of human behaviour is crucial for improving the credibility of using social simulation, in particular for real world applications. The choice of behaviour may crucially determine model outcomes (Hare and Pahl-Wostl (2001)). The adaptive toolbox is a step towards representing a range of behavioural types. It will provide a base to explore which assumptions on human behaviour are supported by experimental and empirical evidence. A distinction should be made between experimental and field data. Field data are derived from controlled experiments with human

---

[2]Daniel Kahneman and Vernon Smith were awarded the Nobel price in economics for this, <http://www.nobel.se/economics/laureates/2002/index.html>

beings in repeatable settings. They allow statistical evaluation and have the advantage of comparability. However, as is the case for any experimental approach the settings may be quite arbitrary and aim at reducing complexity by eliminating subjective context as much as possible. The chosen experimental human subjects (often students) may not always be representative for a wider sample of the population. Empirical data are derived from observations in case studies or other real world situations. They have the advantage of portraying the real world and not an artificial setting. However, given the uniqueness of any real world situation and its specific context, interpretation and comparability are difficult. The adaptive toolbox model will support a better use of experimental observations and a testing if they can be applied to real world situations. It will allow to explore what determines which kind of approach is an appropriate representation of human behaviour in a given context.

In the case reported in this paper data were derived from common pool games in experimental economics. An advantage of such data sets is the availability of numerous data which are comparable due to the standardized settings. The games explore mainly the important aspects of fairness, trust, norms, cooperative behaviour versus free-riding which are all crucial for understanding management of common pool resources – our main area of interest. Previous simulation approaches reproduced the aggregated behaviour of experiments with different assumptions on the behaviour of agents. They pointed out that the comparison with aggregated behaviour did not allow them to decide what would be the more appropriate assumptions on agent behaviour. We explore here also the behaviour of individuals and hope to get thus more insights about behavioural processes.

After giving a brief introduction into the concept of an adaptive toolbox in the next section, we describe our modelling approach in detail in section 2.3. This article includes first results in section 2.4 and a discussion of this model in section 2.5. It concludes with section 2.6 on the contribution of the described work to an overall framework and an outlook of this approaches further perspectives.

2.2 Adaptive Toolbox

Human behaviour varies from mentally challenging, deliberate decisions to unconscious behavioural patterns that follow adopted roles or trained routines. In different problem environments different levels of consciousness are employed. This has to be taken into account when these problem en-

vironments are modelled. The adaptive toolbox described below presents a possibility to deal with this diversity by introducing the notion of heuristics that lie in between consciously making decisions and unconsciously following routines.

2.2.1   Concept

Based on Herbert Simon's concept of "Bounded Rationality" (cf. Todd and Gigerenzer (2003)) Gigerenzer developed the notion of an "adaptive toolbox" (Gigerenzer and Selten (2001a)). It captures the idea of decision making as the use of different heuristics under different circumstances. The idea is based on the suppositions of three concepts: psychological plausibility, domain specificity, and ecological rationality (Gigerenzer (2001, 38)).

- *Psychological plausibility* is explicitly opposed to decision making as a maximizing process with unlimited time, memory and computational capacities. It is also different from optimization under constraint (satisfycing), because instead of calculating an optimal stopping point, the stopping rule is also a simple heuristic.

- *Domain specificity* covers the idea that heuristics work in some problem environments, while other heuristics are used in other environments. Heuristics are composed of simple building blocks, that can be re-assembled to form other heuristics.

- Finally, an *ecologically rational* heuristic is one that prevails in an adaptation process (like evolution), rather than being an "optimal" decision making process. This leaves room for the coexistence of different strategies.

The reference to adaptation processes, as described by evolutionary biology, is a strong motivation for the adaptive toolbox. Some heuristics that humans employ, like emotions, can not be justified by the RAP. However, they may have originated and prevailed in an evolutionary adaptation process.

> "The quest for psychological plausibility suggests looking into the mind, that is, taking into account of what we know about cognition and emotion in order to understand decisions and behaviour. Ecological rationality in contrast, suggests looking outside the mind, at the structure of environments, to understand what is inside the mind." Gigerenzer (2001, 39)

For implementing an adaptive toolbox we have to do both.

The adaptive toolbox consists of building blocks that define the actual choice (cf. Gigerenzer (2001, 43)). These building blocks include simple rules for searching for solutions to a given situation (*search rules*), stopping the search, if a satisfying (not necessarily optimal) solution is found (*stopping rules*), and *decision rules* to choose between alternative solutions. The prerequisite of simplicity of these rules assures the outcome of the decision making process to be "fast, frugal, and computationally cheap". For instance, emotions can function as a very simple form of a stopping rule. Decisions are chosen through simple heuristics, that work, because they take the problem environment into account. Information processing can be done with reference of the information providing environment. If, for example, the environment is noisy and information is scarce, then data are not reliable and therefore calculating with it is bound to be less effective than searching for cues. Also, if the environment is a social environment, other agents have to be taken into account, aspects like fairness and accountability have to be considered (Gigerenzer (2001, 46)). Of course, boundedly rational agents often have an incomplete and faulty perception of the environment.

In addition to the rules there is need for learning mechanisms to ensure ecological rationality. These learning mechanisms may be routine-based learning, reinforcement learning, or even cognitive learning (Brenner (1999, 334,338)).

2.2.2 Implementation

The implementation of an adaptive toolbox aims at a reusable agent model for social simulation that is not only complex enough to cover different kinds of human behaviour, while simple enough to be usable for large populations, but also captures *realistic* decision making.

According to the conceptual considerations described above, the implementation has to include at least the following aspects:

- Agents have a representation of their decision environment.

- Agents have a representation of their social environment.

- Agents doubt their perceptions, they know that they have only beliefs and only sometimes "true" factual knowledge. This leads to adaptation through learning processes (see below).

- Agents have a set of possible solutions to a given problem and choose among them by using fast and frugal heuristics. In principle, the search for new solutions should also be possible.

- The solutions themselves are simple and do not usually involve calculation.

Implementing adaptation through learning processes is an important aspect of the model. There are two dimensions in which human behaviour changes due to learning processes, according to different time scales of the underlying processes. According to the concept of an adaptive toolbox outlined above, the first dimension is influenced by the environment of a problem. In some settings the environment changes rather fast. According to changes in the physical environment the strategy choices vary. The social environment (other humans) affects a quick learning process which has also influences strategy choice. The other dimension is the internal disposition the individuals. This includes both the predispositions, for example the individual propensity to behave in a fair way, and the prior assumptions about the others' behaviour. A change in disposition can be seen as a slow learning process. So far we have not modelled change in disposition.

In addition, we strongly believe in modelling heterogeneous agents, so that there is room for different behavioural traits and beliefs. The need for modelling different kinds of behaviour arises from the fact that some experimental evidence can not be explained by neither an "average behaviour" nor the rational model, like non-linearities and self organization. The altruistic punishment experiment (cf. Fehr and Gächter (2002)) that is the basis for our first application is an example that can not be explained by analyzing average behaviour only.

The implementation itself poses a challenge because both the environmental settings and the heuristics/behavioural traits have to be encapsulated. This requires a thorough understanding of the decision making process to be modelled in order to be able to come up with a valid abstraction.

2.3   Building models from data

As has already been indicated in the Introduction there is no unifying theory of human behaviour in the social sciences. One may question if such a state is really desirable given the richness of human behaviour. However, currently the multitude of theoretical approaches characterizes a state of fragmentation and not a vivid multi-perspective approach. Numerous theories coexist,

in different disciplines, but also within single disciplines. Often, these theories contradict each other and, more often, they are heavily disputed within and between disciplines. Example for highly controversial theories are structural behaviouralism by Talcott Parsons and the theory of social systems by Niklas Luhmann. Each of these theories is regarded as conceptional breakthroughs by some colleagues, and as fanciful artifacts by others. This poses a major problem for modellers of social systems and for modellers, who want to include the "human factor" in their interdisciplinary models. If a modeller of climate change wants to include human reactions on perceived or expected weather changes into the model, those reactions should be based on some notion of how people behave in such situations. However, modellers have to face the ambiguity of multiple representations of human behaviour. Matters are complicated further, because usually model outcomes depend crucially on the theoretical assumptions underlying the implementation of human behaviour.

Economics is the only social science discipline with a dominating theory of human behaviour, the rational actor paradigm (RAP). Although originally this was meant to explain only *economic* activity, it has been adapted into other areas, and some economists view the RAP as universal. RAP is very useful in explaining high-cost situations, when actors have a lot of time, knowledge, and can use computational means. Furthermore, involved actors have to have rather clear and consistent preferences, and the important variables have to be quantifiable and comparable. Firms, for instance, can take their time and invest resources in finding out about different, possible alternatives for investing in different markets and then decide on the strategy that yields the highest expected return, all in terms of money. In these cases, utility maximization is a sound tool that can be used in models, because it is easy to formalize. However, even companies may use heuristics when decisions have to be made in very uncertain and complex situations, when investments are made in innovative products.

However, as has been shown multiple times, the RAP fails to explain many day-to-day observations as well as experimental evidence. This is only partly due to the difference of these low-cost situations. People simply do not have a lot of time and computational capabilities for most of their decisions. They use habits, when they face familiar situations, and often they act on emotions, when the situation is new. Additionally, almost always the prerequisites are not met. Preferences are not usually consistent, and quantifiable; knowledge is incomplete; etc. (See Diekmann and Preisendörfer (2001, 68) and Newig (2003) about high-cost and low-cost situations.)

The RAP model starts to be enriched by psychological and social theory

	EXPERIMENTS	CASE STUDIES	STATISTICAL MASS SURVEYS
Comparability / Control of setting	high	low	medium to high
Representativeness	medium (biased sample)	medium (unique situation)	high
Realism / realistic context	low	high	medium
Repeatability	high	low	high
Direct observability of social interactions	high	medium	low

Table 2.1: Comparison of experiments, case studies and statistical mass surveys

to overcome these explanatory shortcomings. (e.g. Kahneman and Tversky (2000)) The basic underlying assumption and ideal is in general retained: humans behave in a selfishly rational and optimizing way. However, it has been shown in many studies, that in order to capture *realistic* human behaviour we have to view selfishness and optimization as possible behavioural traits among others.

By viewing the RAP as only one alternative among different human behaviours it is possible for us to draw on this theory where appropriate and extend and complement even replace it where necessary. This is the path of a pluralistic approach to describing human behaviour. This path is justified by the described shortcomings of the dominant model and the increasing need to model human behaviour in any number of situations. This need is reflected by the emergence of integrated assessment as a discipline (cf. Pahl-Wostl (2002)).

Behaviour depends on the context of the decision environment. The most obvious example is the difference between day-to-day situations, like buying toothpaste, and important, novel, and single decisions, like buying a house. On the other hand, behaviour also varies with the diversity of humans themselves. In order to be able to implement non-linearities and self-organizational processes, we need to be able to implement diverse human behaviour in one model. Agent-based modelling is a suitable tool to do so. But where do we find evidence for this multitude?

Data may be derived from observations drawn from experiments, case studies, or mass surveys. Table 2.1 summarizes advantages and disadvantages of the three approaches. Since we are interested in deviations from

average behaviour, the statistical approach does not provide appropriate information for our problem. We note that the remaining two approaches are complementary in their strengths and weaknesses. Hence a sound strategy should aim at combining both.

Case studies analyze human actions in a given real world context. They are very useful for explaining certain actions in certain situations. They also help to support a theory about the interrelations of a given problem. However, they do not in themselves give generalizable insights into human behaviour, because every case is unique. Inductive reasoning still depends highly on the underlying theory. Therefore empirical evidence derived from case studies can be only one side of our research programme.

Experimental economics is a way of getting experimental rather than empirical data. This has the advantage of repeatability, comparability and statistical evaluation. Recently, a number of such laboratory experiments have been conducted with the objective to prove the limits of the rational model. They focus on different aspects of cooperative behaviour. The experimental setting makes assertions replicable to some extent. By focusing on a few simple games (for example prisoner's dilemma, common pool resource games, ultimatum and dictator games) the experimental evidence becomes comparable between different experiments. Comparability enlarges the data base considerably. Single experiments can only include about 100 subjects, mostly these are undergraduate economics students of a single university, so they form a biased sample. Only some studies deal with comparison between different cultural biases, for instance Henrich *et al.* (2001b). A further constraint is that usually only one or two aspects are covered, for instance the influence of anonymity (Burnham (2003)) or sequencing (Andreoni *et al.* (2002)). Together those different studies constitute a comprehensive data base on the topic of deviations from the RAP under different, simple game settings.

There is a small body of related research labeled "Parallel Experiments with Real and Computational Agents" by Tesfatsion (2002, 16). There are a few economical studies that deal with both, experimental settings with human subjects and parallel experiments with computational agents. However, with the exception of Duffy (2001) these do not try to capture individual human behaviour, but rather have (boundedly) rational computational agents evolve over time to show or explain the observed aggregated behaviour of the human subjects. Learning is usually implemented as a genetic algorithm. (For example see Pingle and Tesfatsion (2001), Andreoni and Miller (1995)).

Duffy (2001) explicitly models individual, heterogeneous behaviour. He uses "hypothetical reinforcement" learning and diverse agents to reproduce

an experiment that is based on the Kiyotaki-Wright trading model. The agent based simulation is then used to design further settings for laboratory experiments. By this the simulation results can be compared with experimental studies that were done only after the simulation runs.

In another interesting study Deadman and Schlager (2002) use experimental learning, but also do not try to reproduce individual decision making of their experimental subjects.

These kinds of models reproduce actual human behaviour better than RAP. However, they focus on only one single economic aspect, that has been covered by the parallel experiment. In contrast, our model aims at reproducing diverse, individual behaviour at an abstract level so that findings from one experiment can be used to explain those of other experiments, of case studies, and eventually also behaviour in everyday situations.

Of course, the explanatory power of simple game settings like those of laboratory experiments has to be mistrusted in respect to day-to-day situations. But this is exactly the gap that our modelling approach may help to bridge. The data base composed of data from multiple controlled experiments contains the inhomogeneous human behaviour in simple environments. This is a fitting starting point for our model. By taking many of these experimental studies into account we build a model of diverse human actions in diverse environments. We plan to test this model against empirical data taken from case studies and reconcile it with this data. By comparing the two different approaches, not only by the results, but also by the information needed to construct a model, we expect to gain valuable information on human behaviour and decision theory.

2.4   First example

2.4.1   Altruistic Punishment Experiment

A first implementation of an adaptive toolbox was based on the data of the altruistic punishment experiment by Fehr and Gächter. The experiment is described in detail in (Fehr and Gächter (2002)). Here only a brief summary is given.

240 participants played an anonymous common pool resource game in groups of four. 12 of these games were played in a row. Participants did not meet each other more than twice. Six of the games are played as simple common pool resource games. The participants received 20 money units of assets and could contribute between 0 and 20 money units to a common project. The common investment of the four participants was increased by

the experimenter by 60% and divided evenly among the four. Hence, free riders who did not invest into the common project received nevertheless an equal share from the common pool including profits and investment made by other players. The other six games were also common pool resource games, but now with a subsequent possibility to punish players for their investment decisions. For every 1 money units (between 0 and 10) invested in the punishment, the punished player had to pay 3 money units. There have been two experimental settings, each with 120 participants, divided into five groups of 24 subjects for each experimental session. One started with six games with the possibility to punish and concluded with six games without punishment. This will be referred to as setting A. The other setting started without the possibility to punish and concluded with punishment. This will be referred to as setting B (see Table 2.2).

2.4.2 Survey of experimental results following Fehr and Gächter

- Common investment increases during games with the opportunity to punish and decreases without.

- With an average investment of about ten in the first games without punishment, the participants' behaviour is far from the prediction of 0 expected for rational behaviour.

- Almost every participant contributed more in games with punishment than in games without.

- In the first games with punishment (game 1 in setting A and game 7 in setting B) the contribution was higher than in the first games without punishment. The punishment threat effectively increases investment.

- Although it is costly, punishment does occur quite frequently and it is correlated to the deviation from the mean investment by the punished player.

- Punished subjects usually increased their contribution in the next game. So, not only the punishment threat but also actual punishment increases investment.

- In games with punishment, the highest return was received by those players who contributed an amount close to the average investment.

Round	1	2	3	4	5	6	7	8	9	10	11	12	
Setting A	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	5 groups of 24 subjects
	P	P	P	P	P	P							
Setting B	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	CPR	5 groups of 24 subjects
							P	P	P	P	P	P	

Table 2.2: Experimental design

2.4.3   Data analysis

Analysis of the aggregated data does not give us clues about the individual decisions over time and reactions to behaviour of other participants by the subjects of the experiment. Thus, we analysed individual data rows to find out about how decisions were altered in reaction to previous experiences. However, we did only analyse data rows of setting B. Our main working hypothesis is that most of the participants tried to invest close to the mean investment, neither defecting nor being the sucker for others to exploit.

The observed aggregated behaviour over time can not be explained by assuming only one average strategy. Expectations decrease as the investment level decreases during games without punishment. Therefore, there have to be participants who constantly contributed less than the average. On the other hand, in games with punishment there have to be participants who contributed more than average and thus lead to an increase in expectations and, consequently, also in investment. This theoretical reasoning is supported by an analysis of individual data.

First of all distribution of investment decisions in the first round has peaks at 0, 10 and 20 with lesser peaks at 5, 8 and 15. Mean investment is 10.52.

At least three classes of "strategies" can be observed in the individual behaviour. By "strategy" we mean the way in which the investment decision is reached, not the investment decision itself. This corresponds to the notion of decision heuristic. (The possibilities are known as contributing 0 to 20 money units, so no search and stopping heuristics are needed). One extreme is permanent defection throughout the games without punishment (maximizing strategy), the other is permanent cooperation (cooperative strategy). In between are participants who change their contribution, presumably according to the recently made experiences (reciprocal strategy). We believe that participants who play reciprocal strategy were trying to contribute close to the expected mean contribution.

Apparently, there are also participants who start out as cooperators or defectors and change to reciprocal behaviour after a number of games and vice versa. Strategy changes can also be seen as based on heuristics, in this case triggered by cues. The following cues for strategy changes have been ascertained from data series of individual participants and their experiences in the game in setting B:

Cues for strategy changes:

- If common investment is much higher than expected, the tendency to switch from maximizing to reciprocal or from reciprocal to cooperative behaviour increases and vice versa with a low investment.

- If a defector is the only defector the likeliness of defecting again decreases. The same is true for cooperators meeting cooperators. Likewise, if a cooperator encounters one, two or three defectors the willingness to cooperate decreases accordingly. The same is true for maximizers meeting cooperators. Also, reciprocalists may imitate behaviour, they encounter often.

- If payoff of previous decisions has been higher than the recent payoff, the current strategy is questioned again. This cue is highly irrational in games where players do not meet each other again, though in other circumstances it might be useful.

- If a decision leads to a lower payoff than the individual contribution, there is a strong force towards maximizing strategy.

- Punishment may lead to reciprocal or cooperative behaviour in the following ways: Higher punishment than expected decreases certainty about maximizing strategies, while lower punishment increases it. A high number of punishers and a higher punishment than total gain in that round also decreases certainty about maximizing strategies

These cues have been retrieved from data analysis by first classifying individual behaviour in the three strategies mentioned above. Then, changes in investment decisions were classified according to events that happened to the deciding person. We tried to find a reason for each drop or rise in the investment decision. Of course, only most of those changes can be explained by the cues listed above. Also, dependence of the height of the change to a corresponding cue could only be guessed. These dependencies would have to be determined in more detail by questionnaires.

There is probably more extensive reasoning involved. On the other hand, there seems to be also less reasoning involved. Some players seem to give 8 or 10 money units for a few rounds and then switch their behaviour to a higher or lower level, which they employ for another few rounds. Additionally, there is probably a good deal of "random" or "irrational" heuristics involved, that is not captured by these cues. An example for this is that some participants drop their investment level without a provocation in the sixth round (which

they assumed to be the last game played). However, the above list of cues is supported by data and was implemented as heuristics.

In order to ascertain the motivation behind the punishment decisions, Fehr and Gächter had questionnaires filled out by the participants after the experiment. Their analysis of the questionnaires led to their deduction that anger is a major driving force for punishment acts and triggers a "willingness to punish" (Fehr and Gächter (2002, 139)). By analysing individual data we could not find out more about why and when punishment occurred than Fehr and Gächter already did (cf. Fehr and Gächter (2002, 139)):

1. Most punishment acts were done by cooperative players and imposed on defecting players.

2. Both the frequency of punishment and the height of punishment seem to depend on the height of the defection of the punished player.

3. Furthermore, punishment acts are expected by defecting players.

2.4.4 Altruistic Punishment Model

Our implementation reproduces not only aggregated but also individual data of the experiment. Data analysis of individual behaviour lead us to the following assumptions as a basis for the model that are summarized in Table 2.3. The assumptions are described in more detail below the table. With the terms Ägent" or "Player agent" we refer to the entities in our multi agent simulation. With "participants" or "humans" we mean the individuals who took part in the experiment.

Agents have individual inclinations to cooperate and punish. The first is implemented as one variable *cooperativeness*, the latter as two independent variables, one indicating the disposition to be annoyed at being cheated (*inclination to be annoyed*), the other defining the likeliness of spending money to punish a defector (*willingness to punish*). These two variables follow the analysis of Fehr and Gächter, who complemented their experiment by questionnaires (Fehr and Gächter (2002, 139)). All three are float values between 0 and 1, 0 indicating no cooperativeness and 1 indicating a high cooperativeness (or respectively inclination to be annoyed, and willingness to punish).

In our model the original distribution of cooperativeness is an equal distribution. This assumption is supported by the fact that the mean contribution in the first game without punishment is close to 10 and there are about as many participants giving 0 as there are giving 20. As mentioned

agents' cooperativeness c	random variable between 0 and 1, normal distribution
agents' "inclination to be annoyed" and "willingness to punish"	independent random variables between 0 and 1, normal distribution
agents expectations about the others	expected cooperativeness expected inclination to be annoyed expected willingness to punish
higher expected common investment (cooperativeness) in games with punishment	offset=0.15 (setting A) offset=0.3 (setting B)
agents' strategies	maximizing strategy ($c < 0.21$) reciprocal strategy ($0.21 < c < 0.68$) cooperative strategy ($c > 0.68$)
strategy changes with cues	high/low common investment no defector/many defectors higher/lower payoff compared to previous games lower payoff than individual investment punishment

Table 2.3: Survey of important assumptions

above, the distribution in the experiment has peaks at 0, 10 and 20 with lesser peaks at 5, 8 and 15. This may be explained by prominence theory, which states that humans are much more likely to choose *prominent* numbers, like 1, 2, 5, 10, 20, ... (Albers (2001)). However, this has not been modelled.

Inclination to be annoyed and willingness to punish are also distributed evenly. This has been decided due to a lack of knowledge of the actual distribution of those attributes in the human participants. In questionnaires participants stated that they would feel angry towards a defecting individual with increasing intensity corresponding to higher defection. They also expect anger when they were the defecting individuals. This is also reflected in the height of the punishment, which increases with the deviation from the group mean. However, the punishment patterns also differ between individuals. Therefore is seems logical to assume an equal distribution. The way, in which punishment decisions are made in the model is described below.

In addition to their own values for cooperativeness, inclination to be annoyed, and willingness to punish, each agent has a representation of the other agents' respective mean values, which indicates general belief about the others. They start out with believing the others to behave similar to themselves. However, their experiences alter the expectations, but not their own values (see the discussion of time scales in section 2.4). By this, agents learn by improving their beliefs about the social environment, but they do not alter their own "character". In pseudo code for every round the learning is:

```
> exp.coop.= (1-learningrate) * exp.coop.+ learningrate * investment
```
With `learningrate = 0.5`

All agents believe the general contribution to be higher in games with punishment. This offset is 3 money units in setting A and 6 money units in setting B. These values have been taken from the aggregated data.

In our model, all agents have three strategies to choose from: maximizing, reciprocal and cooperative. In games without punishment maximizers contribute 0, cooperators contribute 15 to 20, depending on expectations, and reciprocal strategists invest the same amount of money, they expect others to contribute. Only maximizers change their "reasoning" in games with punishment, trying to calculate the lowest contribution that risks no (high) punishment. In fact, in games with punishment, the only difference in contribution between reciprocal strategists and maximizers is that maximizers may risk a slightly lower investment. Contribution close to the mean actually yields the highest return. This was also true in the original experiment (Fehr and Gächter (2002, 138)).

In Table 2.4 the three strategies are described in pseudo code. Decisions are calculated as values between 0 and 1 and are later multiplied by 20 to give the number of money units that the player agent invests.

The initial strategy of each agent depends on its value for cooperativeness. The thresholds were taken from the data. Of the participants in setting B 21% used maximizing strategy, and 32% used cooperative strategy in their first game. Consequently, we used the thresholds of 0.21 and 0.68 as indicators for the starting strategy. That is, a cooperativeness below 0.21 leads to a maximizing strategy, above 0.68 to a cooperative strategy and in between to a reciprocal strategy.

It is important to note the difference between the strategy of an agent and its investment choice in a given game. The strategy is a heuristic and determines in which way the investment decision is made. The same contribution can be made by player agents employing different strategies. The way in which the strategies change according to experiences, can be seen as another form of strategy. In this case it is a heuristic that uses cues. This is the same for every agent.

For example, in our model reciprocalists change their *contribution* according to their experiences because expectations change. However, *strategy* changes induced by experiences also occur and lead the reciprocalist to employ either cooperative or maximizing strategies. The contribution does not even have to change.

For modelling strategy changes, in addition to the actual strategy employed, each agent has as *certainty* for using that strategy. Positive experiences and expected behaviour by other agents increase certainty, while negative experiences and unexpected behaviour decrease it. The implemented cues are described in Table 2.5. In addition, employing a strategy that corresponds to the agent's cooperativeness increases certainty, while noncompliance decreases it. With a low certainty the probability of a strategy change increases.

The cues for strategy changes are checked after every round, those involving punishment are checked for after punishment has been made. The above list is transformed into the following checklist and corresponding cue values, 0 (cue was not encountered), 1 (cue was encountered), or 2 or 3 (`numDefectors`, `numCooperators`, `numPunishers`):

For each employed strategy the influence of the cues is different. For example, if cooperation is higher than expected certainty about cooperative strategy is increased, but about reciprocal and maximizing strategy it is decreased. This is modelled as an array of parameters indicating for each of the three strategies the influence on the certainty about this strategy.

| Maximizing strategy | ```
for (0 .. numSteps) do (
    decision += step
    calculate outcome
    for (all future games) do (
        calculate expected outcome
        outcome += expected outcome
    )
    if (outcome > maxOutcome) (
        maxOutcome = outcome
        remember decision
    )
)
``` | *Note:* Decisions are tried out and outcome is calculated. For this, the loop starts with an initial decision of 0 and increases it by a predefined step in this case 0.05 = 1 money unit. The last decision remembered is the decision that yields the highest expected outcome. The future games are the games that directly depend on this decision, like the "punishment game" after an "investment game". |
|---|---|---|
| Reciprocal strategy | `decision = expected cooperativeness` | |
| Cooperative strategy | ```
if (expected cooperativeness > 0.4)
    decision = 1
else
    decision = 0.75
``` | *Note:* A decision value of 1 is a contribution of 20 money units and a decision of 0.75 is a contribution of 15 money units, which is the minimum that we identified as cooperative behaviour. |

Table 2.4: Strategy implementation

coopIsHigher	`if (investment > (exp.  coop.  + tolerance)`
	`        coopIsHigher = 1`
	`else coopIsHigher = 0`

coopIsLower	`if (investment < expe.  coop.  – tolerance)`
	`        coopIsLower = 1`
	`else coopIsLower = 0`

noDefectors	`if (number of (other) defectors = 0)`
	`        noDefectors = 1`
	`else noDefectors = 0`

numDefectors	`numDefectors = number of (other) defectors` *Note:* defined by `investment <= 1 money unit`

noCooperators	`if (number of (other) cooperators = 0)`
	`        noCooperators = 1`
	`else noCooperators = 0`

numCooperators	`numCooperators = number of (other) cooperators` *Note:* defined by `investment >= 15`

profitIsHigher	`if (profit > last rounds profit + tolerance)`
	`        profitIsHigher = 1`
	`else profitIsHigher = 0`

profitIsLower	`if (profit < last rounds profit – tolerance)`
	`        profitIsLower = 1`
	`else profitIsLower = 0`

profitLtInvestment	`if (profit < my investment)`
	`        profitLtInvestment = 1`
	`else profitLtInvestment = 0`

numPunishers	`numPunishers = number of punishers` *Note:* player agents that punished this agent with any positive number

punishmentIsHigher	`if (punishment > exp.  punishment + tolerance)`
	`        punishmentIsHigher = 1`
	`else punishmentIsHigher = 0`

punishmentIsLower	`if (punishment < exp.  punishment – tolerance)`
	`        punishmentIsLower = 1`
	`else punishmentIsLower = 0`

punishmentGtGain	`if (punishment cost > this rounds gain)`
	`        punishmentGtGain = 1`
	`else punishmentGtGain = 0`

Table 2.5: Cues and cue values

	Maximizing	Reciprocal	Cooperative
coopIsHigher	-1	-1	2
coopIsLower	1	-2	-2
noDefectors	-2	0	0
numDefectors	1	-2	-2
noCooperators	0	0	-2
numCooperators	0	-1	0.5
profitIsHigher	1	2	1
profitIsLower	-1	-1	-2
profitLtInvestment	0	0	-1
numPunishers	-0.5	0	0
punishmentIsHigher	-1	0	0
punishmentIsLower	1	0	0
punishmentGtGain	-1	0	0

Table 2.6: Multipliers for the different cues according to each strategy

As you can see from Table 2.6, only maximizing strategy is influenced by the punishment cues. These parameters are used as multipliers for the cue values. The general procedure is:

```
> certainty = certainty * (1 + enforcement) * matching factor
> for (all cues)
>    certainty = certainty + (cue value * corresp.  parameter * cert.
step)
> if (certainty < certainty tolerance)
>    change strategy
```

Notes:

- `matching factor` is calculated by the rules displayed in Table 2.7.

- `certainty step = 0.2`

- `certainty tolerance = 0.4`

- From maximizing strategy and cooperative strategy any change leads to reciprocal strategy. From reciprocal strategy it depends on whether more cooperative or defecting cues have been encountered.

The multipliers and the step by which certainty is increased or decreased were first subject to logical reasoning and second to fitting the model to the

Maximizing Strategy	`0.5 - cooperativeness`	Cooperativeness below 0.5 increases and above 0.5 decreases the agent's certainty about maximizing strategy.

Reciprocal Strategy	`if (cooperativeness < 0.5)`
	`    (-0.5 + 2 * cooperativeness)`
	`else`
	`    (-0.5 + 2 * (1 - cooperativeness))`
	`)`
		Cooperativeness below 0.25 and above 0.75 decreases and between 0.25 and 0.72 increases the agent's certainty about reciprocal strategy.

Cooperative Strategy	`((-0.5) + cooperativeness) * 2`	Cooperativeness below 0.5 decreases and above 0.5 increases the agent's certainty about cooperative strategy. This is doubled because cooperative strategy seems to depend more on conviction than the other strategies.

Table 2.7: Calculating the matching factor for strategies from agents' cooperativeness

data.

Another decision the agents have to make is the punishment decision. As has been mentioned above, player agents have an attribute for "inclination to be annoyed" and "willingness to punish". The corresponding heuristic involves not only those two attributes but also the height of the defection that is to be punished. We argue for two dependencies. First, the higher the player agents' inclination to be annoyed and the higher the defection, the more likely it is, that punishment occurred. Second, the higher the player agents' willingness to punish and the higher the defection, the higher the punishment decision was. Furthermore, there was also punishment, that did not fall into the pattern, that cooperative players punished defecting players.

The punishment heuristic used in our model is given here in pseudo code:

```
> if (defection > 0.1)
>     angerlevel = annoyance + defection
> else
>     angerlevel = annoyance + defection - 0.8
> if (randomNumber < (2 * angerlevel))
>     punishDecision = (punishment + defection) / 2
> else
>     punishDecision = 0
> punishPoints = 10 * punishDecision
```

Notes:

- `defection` is the difference between the investment decision of the player agent in question and the mean of the other players, this is 0 if no defection occurred. If there was no defection only irrational anger leads to a punishment decision. This can only happen, if the attribute inclination to be annoyed (plus a possible minimal defection) is bigger than 0.8, as indicated in the first `else` clause.

- `angerlevel` is a test level for the random number, the higher the angerlevel, the greater the possibility that punishment actually occurred.

- `annoyance` is the punishing players inclination to be annoyed attribute.

- `punishment` is the punishing players willingness to punish attribute.

- `punishDecision` is the decision, how much to punish the player agent in question. This is still a number between 0 and 1.

- `punishPoints` is the points that the punishing player agent decides to invest in the punishment. (The punishing players pays the number of points in money units and the punished player pays 3 times that amount.)

- Both the tolerance and the value for irrational anger were fitted, rather than taken from the data.

Our model of an adaptive toolbox is programmed as JAVA classes that use the agent based simulation environment *Quicksilver*.[3] For further information about the model, the source code as well as an online version as JAVA applet, please refer to the website.[4]

2.4.5   Results

With this implementation we have been able to reproduce both the aggregated and individual data provided by the altruistic punishment experiment by Fehr and Gächter. We made model runs similar to the experimental setting, model run A starts with games with punishment (see Figure 2.1) and model run B starts with games without punishment (see Figure 2.2). Model runs have been conducted with 1200 player agents. We did not do model runs with only 24 agents because of a strong influence of the random number generator. Even in runs with 120 players the mean investment usually deviates considerably from the mean investment of the experiment, in some cases not even the trend was reproduced. In fact, this effect is interesting and needs to be analyzed in more detail. We believe that the higher variance is due to the lack of prior knowledge of our agents compared to the experiment's participants.

Only individual data of the experimental setting B was used for calibrating the model. As can be seen in Figure 2.2, the data from setting A are not reproduced as well. In setting B the variance of the twelve data points of the model run from those of the experiment is 0.38. In setting A, however, the variance is 2.51. This strong deviation may be explained by two reasons. The first reason is the drop of investment level in the sixth investment game. Some participants defect in that game because they think it is the last one. Whether or not they expect others not to punish in the last game or simply take their chance, would have to be ascertained by questionnaires. The second reason may be that some participants are angry at having been

---

[3]<http://www.usf.uos.de/projects/quicksilver/>
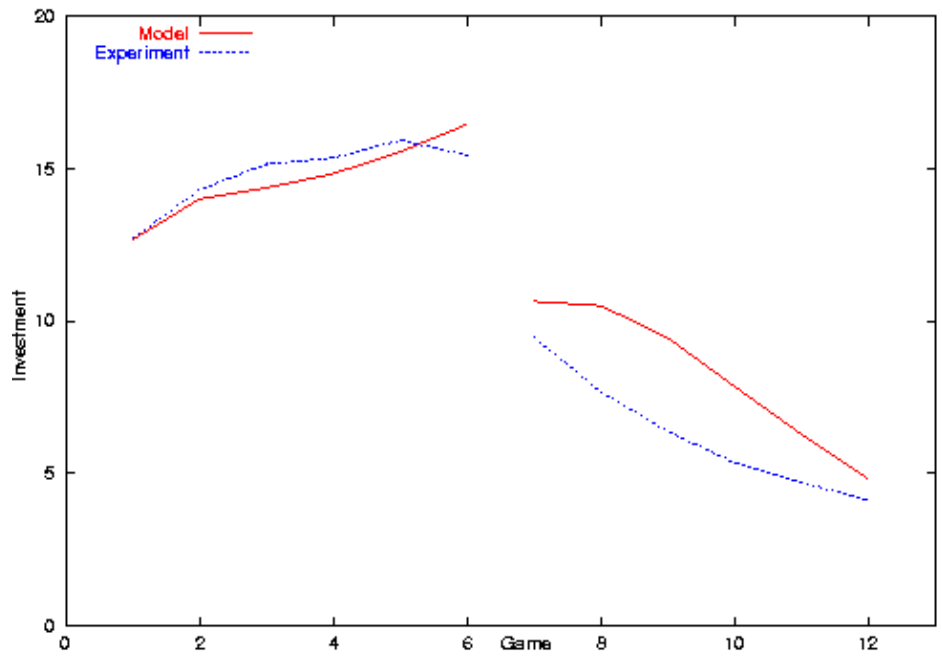[4]<http://www.usf.uos.de/~eebenhoe/forschung/adaptivetoolbox.de.html>

Figure 2.1: Mean investment in the experiment and model setting A
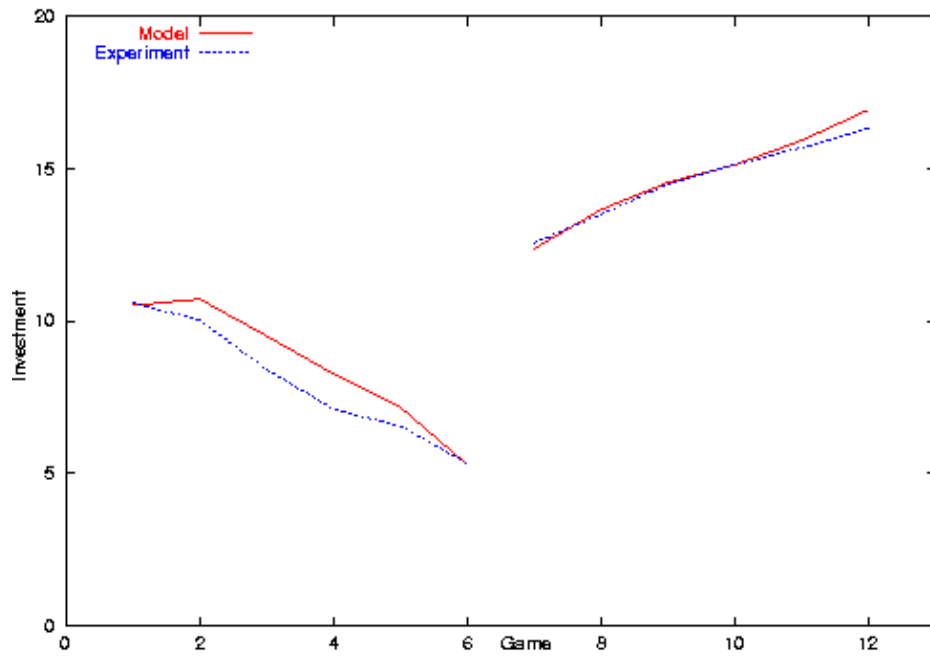
Figure 2.2: Mean investment in the experiment and model setting B

punished in previous games and therefore the aggregated level of investment is lower than in setting B without punishment. Integrating these two aspects in the model leads to Figure 2.3, which shows a better reproduction of the experimental data than Figure 2.1. For this altered setting A the variance is 0.61.

- For the last round effect all agents expected cooperativeness was reduced by `lastGameOffset = 0.1`. This value corresponds roughly to the data. This could have been modelled as an individual trait of agents, because it seemed to be only some participants, who defect without provocation in the last round. However, so far we needed this parameter to be easily accessible and changeable.

- To incorporate increased anger due to previous punishment we increased the percent of player agents employing maximising strategy from 21% to 23%.

Another difference is that mean punishment in the model is higher than in the experiment (1.07 compared to 0.73 mean punishment decision in setting A). The reason for this may be that participants were more risk avers than agents.

Reproduction of individual data is harder to prove. The reason for this is path dependency of individual data. Decisions to increase or decrease investment depend on recent experiences and those are different for every player. However, a few examples of participants' and agents' decisions in setting B are given. As two examples for truly cooperative behaviour see Figures 2.4 and 2.5 for a participant of the experiment and an agent from the simulation respectively.

Examples for reciprocal behaviour alternating with maximizing decisions are given in figures 2.6 and 2.7. Note that the participant's and agent's investment is influenced strongly by the experience made in the prior game.

The participant from Figure 2.8 and the agent from Figure 2.9 start out as maximizers and are forced to invest more in the games with punishment.

In addition to model runs that are similar to the experiment, we also made longer test runs (see Figure 2.10). Results are that it takes about 12 rounds for almost every agent to invest 20 money units in games with punishment and 0 money units in games without punishment. Interestingly, homogeneous investment decisions are possible with different, co-existing strategies. In games without punishment more than 10% of the agents still use reciprocal strategy, but since they expect others to invest nothing, they also do not invest. In games with punishment the setting allows for all

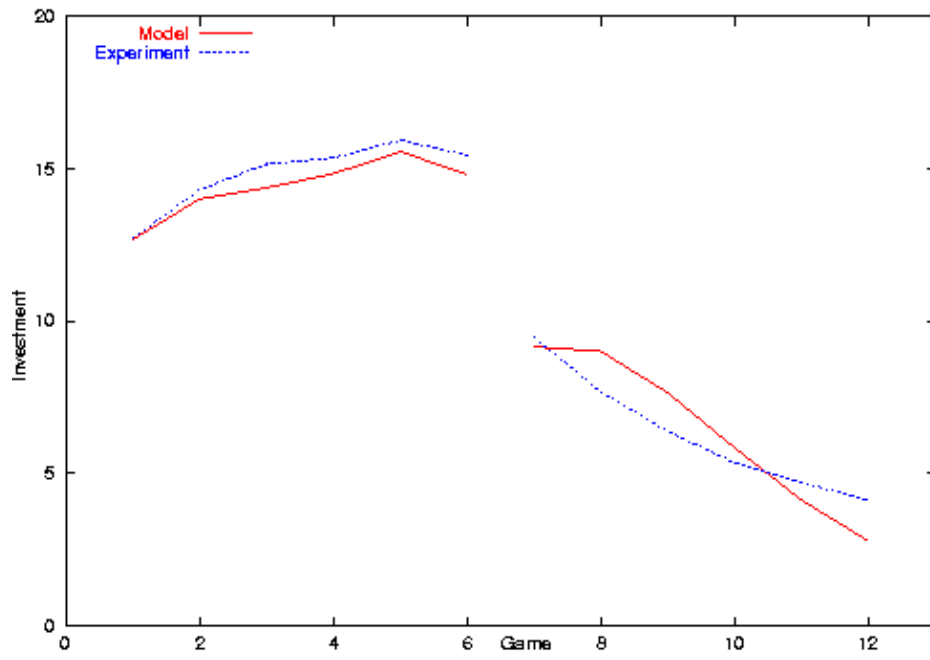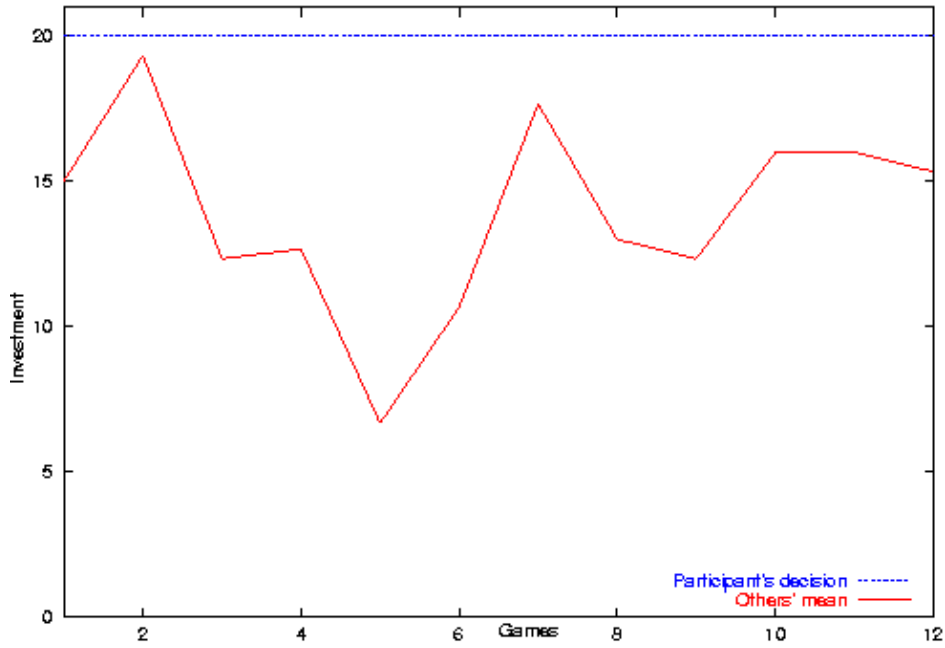*Figure 2.3:* Mean investment in the experiment and modified model setting A

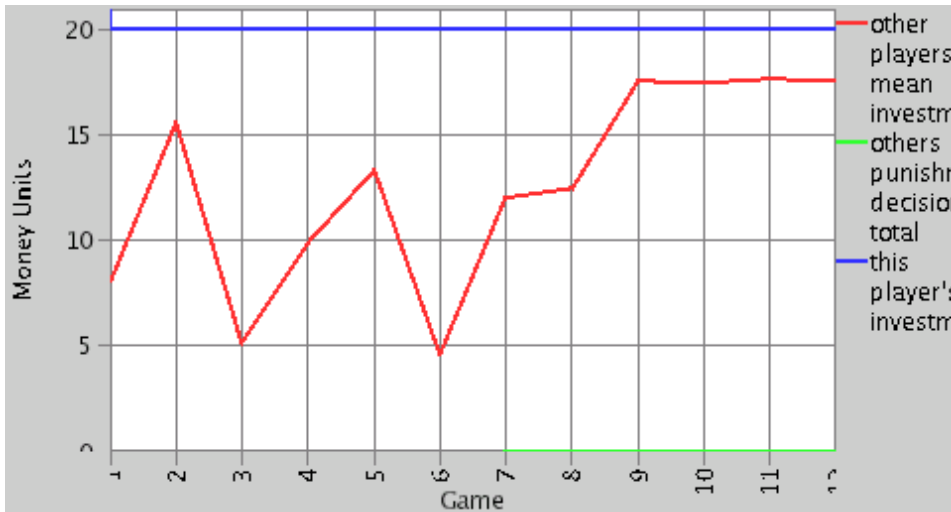*Figure 2.4:* Example of a cooperative participant

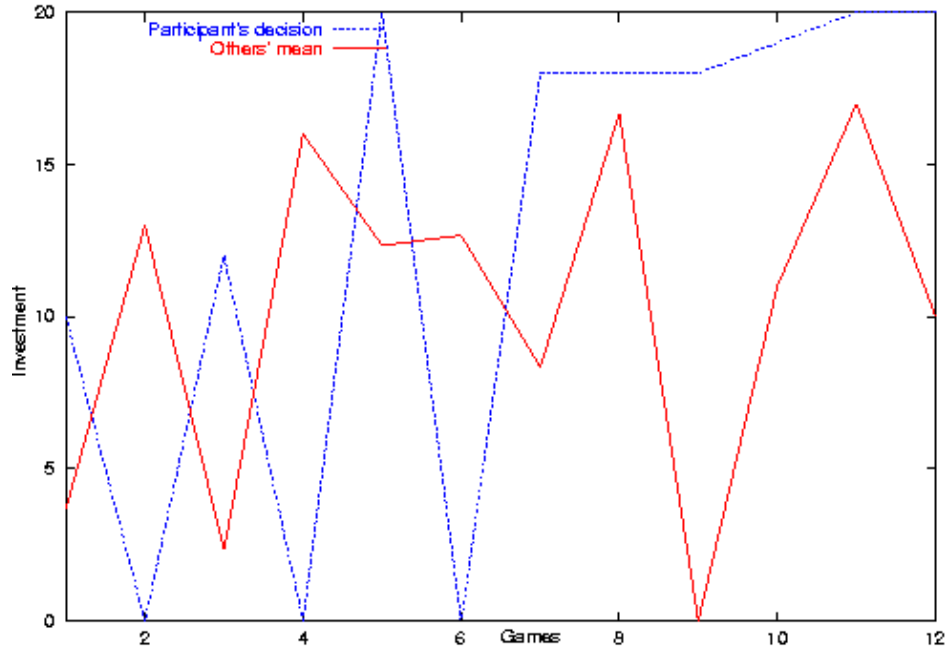*Figure 2.5:* Example of a cooperative agent from the simulation

Figure 2.6: Example of a reciprocal participant

Figure 2.7: Example of a reciprocal agent from the simulation

Figure 2.8: Example of a maximizing participant

Figure 2.9: Example of a maximizing agent from the simulation

Figure 2.10: Mean investment in the long run with and without punishment

three strategies to co-exist. About 60% of the agents are cooperators, about 20% are reciprocalists that invest close to 20 money units because they expect others to contribute that much, and another 20% are maximizers, who contribute about 19 money units, because they want to avoid punishment.

Learning and strategy changes are crucial for model behaviour. Learning rates and, to a lesser extent, the importance of cues are very sensitive parameters. However, from data alone, we could retrieve only limited and unreliable information about these aspects. With questionnaires in addition to the experiment these questions could be addressed more thoroughly.

2.5   Discussion

In the previous section we have shown that the altruistic punishment model reproduces the experiment's data quite well. But we have not yet discussed how it fits into the concept of an adaptive toolbox.

The adaptive toolbox is based on three principles: psychological plausibility, domain specificity, and ecological rationality. By analysing data of

individual participants we captured the actual, individual behaviour, both for the game decision and the reactions to the other participants' actions. For the implementation this behaviour was classified into a set of behavioural types, distinguished by an assumed cooperativeness. From this the actual strategies were derived. The agents' cooperativeness defined the preferred strategy. Strategies themselves were kept very simple. With the exception of the maximizing strategy, no calculation is done by the player agents. Domain specificity was ensured by linking decision strategies to the game setting. This is done by the implementation. That is, only strategies that fit to the game currently played may be used by the players. It is also done by deriving cues for strategy changes from the game setting. For this model, strategies and cues were predefined by the modeller. However, in principle the society of agents could learn in an evolutionary adaptation process, which strategies are possible and which cues are appropriate. Since our objective is to reproduce observed behaviour, we did not choose that path. For this reason, ecological rationality came only from data analysis and not through an evolutionary process.

Decision making within the adaptive toolbox is done by heuristics that are comprised of simple building blocks, so that they can be applied to different kinds of decision environments. As outlined in section 2.2, heuristics are either search rules, stopping rules, or decision rules. In the model player agents chose from a predefined set of strategies. The strategies define not the actual decision, but how the decision is made. That is, they refer to the choice between different solutions. The way in which the choice between strategies is made, is also implemented as heuristics, in this case, checking for cues. Certain cues indicate for player agents that the strategy employed is not appropriate. These cues induce strategy changes. However, new strategies are not (yet) searched for by agents.

Another feature not yet implemented is an evolutionary adaptation process. Adaptation takes place to some extent, but is restricted to learning processes about the social environment.

We have started to derive hypotheses about human behaviour from an economic experiment. By classifying behavioural types it was possible to implement an agent based model to represent data of the experiment. This was done as a first module of an adaptive toolbox that is to be expanded in the near future (see section 2.6). The model is derived from individual rather than aggregated behaviour. By this the idea of an adaptive toolbox has the potential to integrate different coexisting representations of human decision making in agent based models. At the same time it also provides us with a framework for modelling the behaviour and for comparing different

settings.

However, our current approach has some limitations. Experimental economics focuses on a constrained set of behavioural patterns and considers mainly extrinsic motivation for behaviour. Optimization of a utility function is always triggered externally. Furthermore, in economic games all context is removed. However, it is evident from the results that people respond to emotions, so they also have intrinsic motivation. This psychological aspect is very hard to capture in games. In the case of the altruistic punishment experiment, anger was ascertained as major driving force behind punishment decisions. However, this was only possible through corresponding questionnaires. In addition, all people will enter the game with a personality shaped by previous experience. They will change their strategies but not their personalities during a gaming session.

The model can not cover the multitude of behavioural patterns of all the participants. Classifications have the advantage of emphasizing some aspects and general patterns, but always will lessen the variety. Additionally, the interpretation of the individual data, what reasons there were for behaving in that way or another, depends on the modeller.

2.6   Prospects and conclusion

The heuristics explored in this paper focus on what determines the willingness of individuals to cooperate and the development of trust in a group. It is assumed that individuals are characterized by their cooperativeness which may be determined by individual character, individual experience and the cultural context. Nooteboom (2002) suggested quite a useful conceptual framework for sources of cooperation that incorporates many elements and intuitions from literature. Table 2.8 summarizes the main points. One can make a distinction between macro and micro sources and between egotistic and altruistic sources.

This framework is a good base to structure future observations from both empirical and modelling studies. An individual's cooperativeness determines his/her expectations about the behaviour of other players and individual and social learning effects may occur on different time scales. The work reported in this paper provides a start to compile a more comprehensive knowledge base.

We also need to compare insights from the model with theoretical approaches. For instance, cooperativeness was simply modelled as an equally distributed random variable, it might have been modelled as a combination

	MACRO	MICRO
EGOTISTIC	Sanctions from authority or contractual obligation with enforcement	Material advantage or self-interest
ALTRUISTIC	Ethics: values and social norms of proper conduct, moral obligations	Friendship, kinship, routines, empathy

Table 2.8: Sources of intentional reliability (cf. Nooteboom (2002, 9)

of two variables, individualism and altruism, as proposed by Social Value Orientation (as in Jager and Janssen (2002)). Important questions to answer in this comparison between our model and a social value orientation model are: In what way do the results differ? How do investment choices depend on individualism and altruism versus cooperativeness? Does the value orientation of an agent determine, what cues for strategy changes are important to it? How may punishment acts be explained by Social Value Orientation?

The last question is important because Fehr and Gächter found that most punishing acts were done by participants who invested more than average. However, following Social Value Orientation theory this should not be the case. Apparently, there is another aspect, namely anger, involved in that decision.

In order to answer these questions it is necessary to find out about the individual value orientation of the participants by looking at the individual data. One might assume that the classifications of cooperative, reciprocal, and maximizing participants would be the same as a Social Value Orientation classification of cooperative, individualistic, and competitive. This prediction is in line with McClintock and Liebrand, who found that the choices of individualistic players were the most variable ones among those three classes (cf. McClintock and Liebrand (1988, 407)). If the classification was the same, the outcome would likely be the same also, only the representation of cooperativeness in the model would differ. For this experiment we did not need to distinguish between individualistic and altruistic behaviour. For other experiments this might very well be necessary.

The altruistic punishment model is a first step towards an adaptive tool-box that should be comprised of many different modules. For this reason the next logical step is to extend the toolbox by more models of other experimental games. By this it will become possible to compare the validity of

the assumptions made to other findings. This work is currently in progress.

A second step after the extension of the adaptive toolbox is to compare the insights with results from case studies. The question arises whether data from case studies deviate considerably from results in experimental settings. And, if so, what are the main differences?

Major differences, already pointed out, are context dependency and a longer time scale. Short term behaviour may be typically relevant in negotiation processes. However, in case studies we are interested in particular in long-term changes. In our model, these would refer to changes in the individual types, e.g. the attitude to be cooperative. This implies that the incentive for behaviour shifts from extrinsic motivation by sanctions through punishment to intrinsic motivation by the internalization of social norms about socially acceptable behaviour and about a behaviour that leads to an acceptable pay-off in a certain social environment. Hardly any individual is cooperative to an extent to be continuously exploited by others.

Thus, on this longer time scale institutions as another major influence on human behaviour become relevant. Numerous studies have shown evidence for the importance of institutions shaping human nature (Held and Nutzinger (1999)). Hence it will be of interest to explore systematic differences determined by culture and institutional contexts. In these cases more important information can be derived from stakeholder interviews than from experimental settings.

Of course, human behaviour will always remain unpredictable to some extent. Nevertheless for modelling purposes we want to improve our understanding and the representation of how people behave.

Our approach to represent human behaviour by extracting regularities from observations leads to a modelling framework that allows for different approaches to human decision making.

Chapter 3

Modelling human behaviour with attributes and heuristics based on observation

Eva Ebenhöh and Claudia Pahl-Wostl[1]

3.1 Abstract

When modelling humans with agent-based models, modellers face a choice of how to represent human behaviour. Numerous and sometimes contradictory theories exist. We have chosen a pragmatic approach combining the use of attributes and heuristics in an agent-based model. Agents are characterized by a set of attributes and distinguished by their individual attribute values. This way we introduce agent diversity and the possibility to have decisions depend on different aspects of agents' personalities. Their decisions are described by a set of simple decision rules, that are called heuristics. Heuristics are psychologically plausible compared to utility maximization. Due to their simplicity, they are also easy to handle in agent-based models. Both, attributes and heuristics, have been obtained from extensive evaluation of experimental data. Models that reproduce experimental data of various economic experiments are used to test our assumptions and develop heuristics.

This paper presents seven attributes needed in a number of these models and the chosen heuristics for models of dictator and ultimatum games, as an example. These models are part of a framework that allows for quick extensions in the form of models of further experiments.

3.2   Introduction

Human behaviour can be modelled in multiple ways. The range of possible approaches is broad: highly aggregated models where human behaviour is represented by single parameters, indicating for example a general awareness of ecological problems; homo economicus as a rigid model for intentional rationality; several extensive psychological and behavioural theories. Further approaches differ in their taking into account social interactions. Which approach to choose depends, of course, on the application. Our objective is to develop a pragmatic alternative that alleviates some of the shortcomings of conventional approaches.

We aim at a model for human decision making in social dilemmas and negotiation processes that allows for diverse behaviour, is psychologically plausible but simple enough to be incorporated in agent-based simulations and be tested by empirical observation, while at the same time generic in order to be applicable in multiple decision environments. Our approach combines attributes and heuristics. The agents have a number of attributes that simulate traits and preferences. Different attribute values make agents behave substantially different in similar situations. Their decision making is based on heuristics, simple case-based processes. Agents can learn about other agents' attributes and about the usefulness of the heuristics employed.

The attributes implemented in our modelling framework are cooperativeness, conformity, fairness concerning others, fairness concerning me, positive reciprocity, negative reciprocity, and risk aversion. Although all seven attributes are introduced in detail, only CONFORMITY, FAIRNESS CONCERNING OTHERS, FAIRNESS CONCERNING ME, and NEGATIVE RECIPROCITY are use in the model discussed as an example in this paper.

We base the models on experimental data rather than theoretical concepts in order to capture realistic human behaviour. So far, the data considered was taken from economic experiments. From these data sets we obtained attributes as well as heuristics. Several heuristics are generic in the sense that they can be applied to different decision environments, however, some only make sense in one special setting. For our purpose it is important not only to replicate aggregate results but also individual data.

To test the assumptions we implemented models that replicate the data. In this paper, dictator and ultimatum games are used as examples. In dictator games, the first mover allocates an amount of money between him or herself and the other player. In ultimatum games, the first mover still allocates an amount of money between the two players, the second mover, however, has veto power. The second mover can accept the first mover's decision. Then the money is actually paid. He or she can also reject the decision. Then both players get nothing.

Experimental data may be criticised as being artificial and not appropriate to represent human decision making behaviour in real day-to-day circumstances (Loewenstein (1999)). In order to test this approaches usefulness in other settings, we broaden the model to incorporate field data as well.

This paper is organised as follows. The next section *Attributes and heuristics* introduces bounded rationality and our general modelling approach consisting of attributes and heuristics. The attributes are discussed in some detail, while the context dependent heuristics are discussed in the following section *From experimental data to attributes and heuristics*. This deals with data analysis considering as example data from dictator and ultimatum experiments. It describes data and heuristics obtained and used in the model. This section also presents some results of our model. The last section *Prospects* discusses our approach and concludes with possible further applications.

3.3 Attributes and heuristics

3.3.1 Bounded Rationality

Based on the notion of Herbert Simon, bounded rationality has been developed as an explicit alternative to full rationality and the subjectively expected utility maximisation (Selten (1990)). It was proposed "to connect, rather than to oppose, the rational and the psychological" (Gigerenzer and Selten (2001b, 1)). It is not irrational, but has no place for optimisation, or calculations of probabilities and utilities. Instead, the notion is that the cognitive limitations of humans are actually adaptive advantages within the special structure of our environment, which is characterised by uncertainty and incomplete information. Simple heuristics can outperform extensive computations when they take the informational structure of a decision environment into account. However, those heuristics are often useful, only in circumstances that have this special structure. In others they may

not work. Some research has been conducted to understand when and why simple heuristics work (Goldstein (2001); Martignon (2001)).

In bounded rationality decision making is defined as a "search process guided by aspiration levels" (Selten (2001, 13)). Aspiration levels are goals which we try to reach with certain decisions. When the outcome of a decision satisfies the aspiration, there is no need to change behaviour. When the aspiration level is not met, a search for alternatives is initiated. Experimental evidence suggests three levels of decision making, the routine level, the level of imagination, and the level of reasoning (Selten (1990)). If an aspiration level is met by habitual behaviour, nothing needs to change. If it is not met, an individual can change his or her routine slightly, according to simple search rules or social learning processes. The search is stopped, if the goal is achieved. If not, the search for new possibilities may be extended, either by imagination or by reasoning (or both). On the other hand, aspiration levels can change over time in adjustment to the situation (Selten (2001)). So, if an extensive search did not come up with satisfying behaviour, aspiration levels may be adapted downwards, so that a given behaviour becomes satisfactory. This process has been labeled satisficing, to distinguish it sharply from any optimising behaviour.

Habitual behaviour, simple changes to habitual behaviour, simple search rules, and stopping rules are called *heuristics* in this paper. Imagination and reasoning are more elaborate processes. Humans learn a great variety of these simple decision mechanisms together with environmental structures in which they are useful. Gigerenzer has coined the name "adaptive toolbox" for the set of search, stopping, and decision rules (Gigerenzer (2001)).

In this sense, norms and emotions serve as decision mechanisms, that exploit certain environmental structures. For example, the norm of positive reciprocity can be useful in establishing permanent business relationships. In uncertain situations, even when it seems to be only a onetime business, this norm still serves as a useful tool to restrict possible behaviour, even though defection would be feasible and more profitable. There is some research with a focus on norms and emotions as elements of a bounded rationality theory (Mellers (2001); Boyd and Richerson (2001)).

Selten defines a *strategy* as a "system of case distinctions based on simple criteria [that] determines which simple decision rules are employed" (Selten (1990, 653)). The system of case distinctions can be simple or complex depending on a given decision task. However, the application is often very simple.

The concepts of heuristics and strategies are adopted for the development of our agent-based models. Depending on the situations, the agents choose from a number of heuristics which one to apply. A strategy can be used to switch from different heuristics or learn which heuristic worked well in what kind of circumstances.

3.3.2 Attributes

Whenever scientists explain observed human behaviour deviating from what would be predicted by assuming a perfectly rational behaviour they draw on a number of attributes, like altruism, cooperativeness, fairness. Sometimes attributes are defined as parameters in formulas, sometimes they are defined verbally. However, often a precise definition is lacking which leaves the meaning of these attributes ambiguous. What is, for instance, altruism supposed to be? We provide here definitions for those attributes included in our model and arguments for including them. Agents in our model expect other agents to have the same attributes but not necessarily the same value as themselves. These expectations are referred to "expected COOPERATIVE-NESS" etc.

COOPERATIVENESS

This attribute defines the importance of group utility in relation to individual utility. A high COOPERATIVENESS indicates a willingness to invest in group utility even if it is costly. There are a number of reasons, why an agent would invest in group utility. The other agents' welfare may be important, or a possible efficiency increase, or the agent may follow a norm. This attribute is mainly associated with the second reason, an efficiency increase in the sense of Jaffe's "synergistic altruism" (Jaffe (2002)). In addition, this attribute models Andreoni's "impure altruism" (Andreoni (1989)) and "cooperation" in the social value orientation theory (Jager and Janssen (2002)). The agent wants to do its share because it believes the group to be important. Concern for others' monetary welfare is better represented by the attribute FAIRNESS CONCERNING OTHERS, while norm compliance is modelled as CONFORMITY.

COOPERATIVENESS and fairness have not always been separated clearly. For example, Cox' "other-regarding preferences" include cooperation and fairness aspects (Cox (2004)). In our model, fairness is a measure of inequity aversion, while COOPERATIVENESS indicates the importance of group utility. In many cases this may be the same, but

for analytical and modelling purposes, we make a clear distinction. The difference between fairness and efficiency has been investigated in Güth *et al.* (2003). In one-sided gift giving fairness outweighs efficiency considerations. In two-sided gift giving efficiency becomes more important. This goes along with our definition, because in two-sided gift giving the "group" is more apparent than in one-sided treatments.

COOPERATIVENESS is prototypical for social dilemmas in which individuals can free-ride on the group or contribute to the group outcome and profit themselves, if everyone contributes. The basic game for this attribute is the Prisoner's Dilemma. Research on one-shot Prisoner's Dilemma games indicates that a majority defects, but a substantial minority cooperates. Cooperators are willing to initialize cooperation under the risk of being cheated. Defectors are either not willing to cooperate or are not willing to risk being cheated. So, there are three classes of behaviour (Brenner (1999)). Those willing to risk being defected, those willing to cooperate but not willing to risk being defected, and those not willing to cooperate. Depending on the situation, percentages of these three classes vary. In the model, we can represent them by thresholds of the COOPERATIVENESS attribute and the expectations in other agents' COOPERATIVENESS.

CONFORMITY

This attribute defines, how important it is for an agent to appear to be as others expect it to be. A high CONFORMITY indicates norm compliance. If an agent's CONFORMITY is high and its COOPERATIVENESS is low, it may still invest in group utility, although it does not "believe in the group". What norm to comply to, is not defined by this attribute. Rather, it is a tendency to do, what others expect you to do. This may also be used to model commitment into a common project and compliance to agreed upon rules. It is docility, as Simon put it, "the susceptibility to social instruction and influence" (Simon (1997, 285)).

If behaviour is anonymous, this attribute becomes less important. However, it does not have to disappear, since norms may be useful to limit possible behavioural alternatives. The importance of norm compliance depends on the situation. In dictator games, anonymity seems to have a strong effect on first movers' behaviour (Burnham (2003); Hoffman *et al.* (1996)), while in ultimatum games the threat of punishment outweighs this effect (Bolton and Zwick (1995)). In games, in which subjects can verbally agree on a joint strategy, CON-

FORMITY can be used to describe compliance and non-compliance to the agreement. However, there is no simple game to test for this attribute, because it is context and framing dependent.

FAIRNESS CONCERNING OTHERS and FAIRNESS CONCERNING ME

Fairness deals with the difference between an agent's share and the other agents' shares. If an agent's FAIRNESS CONCERNING OTHERS is high, it is willing to give some of its resources to another agent with lesser resources, as in the dictator game. If an agent's FAIRNESS CONCERNING ME is high, it is willing to invest resources into reducing another agent's higher resources, in order to decrease the difference.

FAIRNESS, both CONCERNING ME and CONCERNING OTHERS, is pure inequity aversion. It takes only the outcome into account, not the motives or intentions behind the other players' moves. These are represented by POSITIVE or NEGATIVE RECIPROCITY. A punishing act may be based solely on inequity aversion, or it may be chosen, because of perceived negative intentions, or a combination of both.

The simple game to test for FAIRNESS CONCERNING OTHERS is the dictator game (see below). However, not nearly as much research is conducted to test for FAIRNESS CONCERNING ME. In literature on experimental economics FAIRNESS CONCERNING ME is almost always combined with NEGATIVE RECIPROCITY.

Fairness does not always have an equal split as its focal point. In ultimatum games where the positions are not arbitrary but subject to auctioning or contests, unequal outcomes are more often proposed and accepted (Güth and Tietz (1990); Hoffman *et al.* (1994)).

FAIRNESS is differentiated in CONCERNING ME and CONCERNING OTHERS, as in Fehr and Schmidt (1999), where different and independent parameters are used when a player receives more than his or her opponent and vice versa. Numerous authors, who state that the relative material payoffs effects peoples' utilities, do not necessarily make this difference, as for example Bolton and Ockenfels. In Bolton and Ockenfels (2000) the same parameter is used for unequal outcomes in favor of a player and those that favor the opponent. However, from data analysis of individual decision makers, the difference is quite apparent. They do not have to be correlated. If anything, a high FAIRNESS CONCERNING OTHERS indicates a low FAIRNESS CONCERNING ME. Fair ultimatum first movers would accept significantly less as second movers, than strategic first movers (Straub and Murnighan (1995)).

In some decision environments there has to be a tradeoff between FAIRNESS CONCERNING OTHERS and FAIRNESS CONCERNING ME. Consider for example an ultimatum game with an inactive third player (Güth and van Damme (1998)). First movers have to allocate an amount of money between three subjects. Second movers have to accept or reject the offer and third movers have no say. Virtually no second movers reject on behalf of the third player, even if they would have given more to them themselves. However, since they are not (or think they are not) in the position to act on their FAIRNESS CONCERNING OTHERS, even if it is high, they will act on their FAIRNESS CONCERNING ME instead. On the other hand, rejections in two-player ultimatum games are not necessarily monotone. Hennig-Schmidt et al. found a surprising amount of non-monotone strategies, rejecting too high as well as too low offers (Hennig-Schmidt *et al.* (2002)). This indicates that both FAIRNESS CONCERNING ME and concerning others are taken into account.

In cases, in which no more equitable decision could have been chosen, fairness in this sense is not a good attribute to evaluate an outcome. Sometimes the set of feasible alternatives is incorporated in the fairness model, as in Andreoni *et al.* (2002). We chose to model intentions using POSITIVE and NEGATIVE RECIPROCITY.

POSITIVE RECIPROCITY and NEGATIVE RECIPROCITY

Reciprocity defines how much an agent's behaviour depends on previous behaviour by other agents. If an agent with a high POSITIVE RECIPROCITY encountered an agent that acts nicely, it will reciprocate with a nice act (Cox (2004)).

Agents (and persons) do not know others' intentions. They only perceive their acts and derive possible intentions behind the acts. They can, therefore, only reciprocate *perceived* intentions.

We differentiate in POSITIVE and NEGATIVE RECIPROCITY, as in McCabe and Smith (2001). Experimental evidence speaks in favor of a differentiation, although often analytically the difference is not made (Fehr and Rockenbach (2003)). A subject can reciprocate fair behaviour with high returns, and be unabashed by unfair behaviour, refraining from punishment. The opposite is also possible.

Some authors do not distinguish fairness and reciprocity. Rabin's fairness, for example, is reciprocity as we define it here. "If somebody is being nice to you, fairness dictates that you be nice to him" (Rabin

(1993)). Also, McCabe and others define POSITIVE RECIPROCITY as based on both joint gains and belief about the other player's perceived intentions (McCabe *et al.* (2003)). Reciprocity contrasts from fairness in two ways. As was stated above, fairness takes only outcomes into consideration, while reciprocity is concerned with the intentions of the other players. Secondly, a high FAIRNESS CONCERNING ME (OTHERS), may lead to punishing (rewarding) an opponent in order to equalize the outcome. Reciprocity, on the other hand, can lead to punishment or rewards, even when the outcome can not become more equal (Falk and Fischbacher (2000)). Charness and Haruvy (2002) also make this distinction. Cox (2004) distinguishes between reciprocity and other-regarding preferences, the latter includes COOPERATIVENESS and FAIRNESS CONCERNING OTHERS.

Cox (2004) conducted an experiment with different gift giving games in order to distinguish FAIRNESS CONCERNING OTHERS (other-regarding preferences) and POSITIVE RECIPROCITY. Often, however, we can find FAIRNESS CONCERNING ME and NEGATIVE RECIPROCITY working together. Fehr and Gächter (2002), for example, find evidence, that free riding causes strong negative emotions among the cooperative subjects. However, from their experimental setting it is not clear whether this is due to inequity aversion or perceived intentions. Punishment that equalizes the outcome can result from both causes. It is probably a combination of both, as is indicated by experiments which vary the relation of punishment cost and fining fee (Ostrom *et al.* (1994, 171ff.)). The frequency of punishing acts is inversely related to the cost of punishment.

RISK AVERSION

Agents differ in their willingness to take risks. This may be attributed to a lack of trust (see below), yet we chose to model trust in a different way (see below). A trusting agent may not be sure about its judgement of other agents, and may not risk a trustful move. However, usually RISK AVERSION is used in circumstances where the risk depends on environmental variables and not on other agents.

Kahneman and Tversky conducted a series of experiments to evaluate the RISK AVERSION of their subjects and found that people are usually risk averse with respect to sure gains versus risky gains and risk seeking with respect to sure losses compared to risky losses (Kahneman and Tversky (1979)).

All attributes are modelled as values between 0 and 1. In our model, we assume attribute values to be equally distributed over the whole interval. Attributes do not change over time, because so far only relatively short time episodes have been modelled. They are independent from each other, although some psychological theories suggest correlations between, for example POSITIVE RECIPROCITY, FAIRNESS CONCERNING OTHERS, and COOPERATIVENESS.

In addition to these attributes, agents have a representation of other agents' attributes. These are called *expected attributes* and are subject to quick learning processes. The initial value of expected attributes may either be assumed to reflect past experience or an agent expects other agents to be equal to itself. An agent with a high expected COOPERATIVENESS may cooperate in first moves. If the expectations are not fulfilled, expected COOPERATIVENESS is reduced, which may lead to different behaviour.

Some attributes that have been widely discussed are not included, but can be derived from the list above. Trust, for instance, has been defined as the expectation that a partner will "not be opportunistic even if he has both the opportunity and the incentive to do so" (Nooteboom and Six (2003b, 4)). Depending on the situation, this can be expected COOPERATIVENESS, expected FAIRNESS CONCERNING OTHERS, or expected POSITIVE RECIPROCITY or a combination of those (Cox (2004)). Altruism, defined as the care for others' material well-being (in addition to one's own) (Charness and Haruvy (2002)), is a combination of COOPERATIVENESS and FAIRNESS CONCERNING OTHERS.

By using our pattern of attributes in multi-agent models we can design agents that have all of these attributes and draw on the appropriate attribute(s) when confronted with a certain decision problem. Furthermore, we can take data from single attribute experiments (like dictator games) to define a distribution among our population of agents. The same diverse agents are then used to test data from other experiments with more than one attribute involved (like ultimatum games).

3.3.3   Heuristics

Heuristics are simple decision rules that can be applied to certain decision environments yet fail in others. Our agents are endowed with a number of heuristics to choose from when confronted with a decision task. In our model, the heuristics applicable for a certain decision are defined by the modeller and are not subject to evolutionary processes. The potential danger of working with heuristics is to develop specific models tailored to one decision

Figure 3.1: The use of attributes and heuristics in a decision process. The decision environment defines which attributes are used to decide on the heuristic to use. The heuristic may use attributes and expected attributes in order to make the decision. Experiences alter expectations and the use of heuristics. Choice of attributes, choice of heuristic, the decision, and experiences depend on the decision environment. Attributes, expected attributes, and heuristics are agent properties.

context only. In order to obtain generalisable results we are still developing a classification scheme to characterise decision environments and heuristics used.

Figure 3.1 demonstrates a general scheme for the interaction of heuristics and attributes in our models. The interaction is a threefold process. First, attributes are used to decide on which heuristic an agent uses. Second, the heuristic can use attributes to make the actual decision. And third, the attributes can have an influence on the learning process which may alter heuristics in use.

We have implemented a number of heuristics that work in different environments. However, in order to implement models about certain economic experiments we needed some heuristics that have been designed for one special decision (like responder behaviour in ultimatum games). Instead of presenting them all, we present only a classification. In the next section the heuristics used in dictator and ultimatum games are discussed in detail.

Heuristics depending on attributes

> Most heuristics use one or more agent attributes in order to calculate which decision to chose within a range of possible options. When agents can decide to contribute 0 to 20 money units (MU) to a common project, a simple mechanism of deciding between the 21 alternatives is to choose the proportion represented by the value of an agent's COOPERATIVENESS. A COOPERATIVENESS of 0 corresponds to 0 MU contribution, 0.5 corresponds to 10 MU and 1.0 to 20 MU. The calculation can be extended to include reasonings like, if I contribute anything, it will be at least 3, or rounding to prominent numbers.

Heuristics depending on the range of possibilities

> Simpler heuristics often use only the range of possibilities for a choice. Attribute values are used to decide which heuristic to use, but they are not used by the heuristic itself. An example for such a heuristic would be "give half of what is possible". Also "Yes" or "No" in binary choices are in this class of heuristics.

Learning heuristics

> If heuristics base a decision on past experiences and decisions, they fall into the category of learning heuristics. We could call these heuristics strategies in order to distinguish them from simpler ones. However, this is not done consistently in literature on bounded rationality. Search rules and stopping rules are also considered as heuristics.

Note, that complex formulas that calculate a decision uniformly for all different situations are a special case of "heuristics". Selten (1990) calls these formulas *unified strategies*. We can adopt these cases in our model, by using a formula instead of a heuristic, that is never changed by an agent.

3.4 From experimental data to attributes and heuristics

3.4.1 Experimental economics

Experimental economics investigates human behaviour in artificial, experimental decision environments. This has the advantage of high repeatability and comparability. That is, the experimental setting can be changed in small steps and the effect on decision behaviour can be observed. Similar experimental designs should yield similar results.

However, the complexity in experimental decision environments is usually lower, than in "reality". Still, "the effects of the general bounds of reality are present in the experiment, but they are less than in the historical reality" (Tietz (1990, 659)). Hence, experimental economics has the advantages and disadvantages of most experimental approaches: high ability to infer general conclusions from research results but restricted ability to transfer insights from experimental findings to the settings in the real world that the experimental design intends to approximate.

The experiments discussed and modelled here are about dictator and ultimatum games. In a dictator game, the proposer is given an amount of money and has the task to divide it between him or herself and the other player. Instead of "dividing" the task can also be described as "giving" something to the other player. This framing makes a difference. The second player receives the money but has nothing else to do.

Ultimatum games differ from dictator games insofar that the second player, the receiver, has veto power. He or she can reject the offer. Then, both players get nothing.

3.4.2 Data analysis

The data analysed comes from various dictator (Burnham (2003); Forsythe *et al.* (1994); Güth and Huck (1997); Hoffman *et al.* (1994, 1996)) and ultimatum experiments (Gale *et al.* (1995); Güth and van Damme (1998); Güth and Tietz (1990); Henrich (2000); Straub and Murnighan (1995)), but not
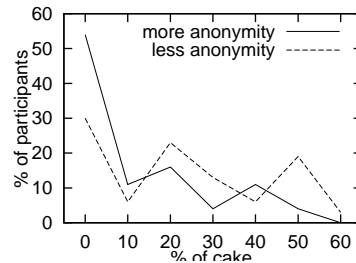
Figure 3.2: Data from two different dictator experiments, with a high degree of anonymity (Burnham 2003) and a medium degree of anonymity (Forsythe *et al.* 1994).

all results are presented. The analysis is concerned with data on aggregated and individual behaviour. This analysis should be seen as an example to illustrate our approach. We present data and some derived behavioural regularities. When more complex decisions are made, the analysis can be guided by questionnaires, as in (Dal Forno and Merlone (2004)).

Dictator Game

Usually the modal offer is 0. There is, however a nontrivial percentage of offers of 50% and some offers in between. The distribution depends on the degree of anonymity. With anonymity from the experimenter assured, the offers are lower, there are more offers of zero and less of 50%. Interestingly, whether a lesser peak of the distribution is at 20 or 30% of the cake depends on the prominence of the actual values. Figure 3.2 presents exemplary data for two different experiments taken from Forsythe *et al.* (1994) and Burnham (2003).

In addition to the anonymity effect, the description of the task and thus the framing of the decision can also have an effect. Different associations are triggered, whether the task is to "allocate an amount of money" or to "make a gift to the other player". When the money to allocate is not yet in the possession of the proposer, he or she is more willing to give money to the other player. If he or she holds the money in hands, giving something away is harder.

Our dictator proposer agents cover the data in the following way. The proposer has to take into account his or her own feelings about a fair distribution. Some may also consider, what impression the decision makes on the experimenter. If we assume anonymity to hold, FAIRNESS CONCERN-

Figure 3.3: Data from dictator model in two different treatments, with and without anonymity effects.

ING OTHERS is the only attribute to apply. If not, the decision may be a combination of FAIRNESS CONCERNING OTHERS and CONFORMITY. A list of possible heuristics includes:

1. Give nothing.
   The reason for choosing this heuristic is probably maximisation of the agent's profit without concern for the other agent's welfare and the experimenter's opinion.

2. Give half the cake.
   More reasons for choosing this heuristic are possible. Either an agent has a high FAIRNESS CONCERNING OTHERS, or it has a high CONFOR-MITY when anonymity from the experimenter is not assured, or does not care and chooses the obvious focal point.

3. Give half the cake times the value of your FAIRNESS CONCERNING OTHERS attribute.
   A reason for this heuristic may be the concern for the fairness of the distribution combined with the assumption that it is appropriate to take more than half, since the proposer possesses more bargaining power.

Which of these heuristics an agent chooses can depend on a learning process. however, in one-shot situations without prior experiences, we chose to assign each agent a start heuristic, depending on its attributes.

Figure 3.4: Data from two different ultimatum experiments, taken from (Hoffman *et al.* 1994) and (Forsythe *et al.* 1994).

1. Without experimenter anonymity and CONFORMITY > 0.9 choose heuristic 2.

2. With experimenter anonymity and FAIRNESS CONCERNING ME > 0.7 choose heuristic 1.

3. With FAIRNESS CONCERNING OTHERS < 0.25 choose heuristic 1.

4. With FAIRNESS CONCERNING OTHERS > 0.8 choose heuristic 2.

5. Else choose heuristic 3.

The thresholds chosen are based on data analysis. Hennig-Schmidt (1999) conducted video experiments in order to ascertain the reasoning behind bargaining behaviour in ultimatum games. In experiments thresholds can also vary with the framing. Anonymity and task descriptions may have an effect on which attribute is used, or on the thresholds. Therefore, the numbers presented here are only approximations. When modelling a specific experiment, the values can be subject to a fitting process.

Figure 3.3 presents the stylised results from a model with the parameters given above.

Ultimatum Game

The modal offer is 50% of the cake. Very few offers are below 20%. Offers of 20% are rejected with a probability of more than 50%. Some offers of 40% are still rejected. Anonymity effects do not have as strong an effect as in dictator games. Figure 3.4 presents exemplary data from two different ultimatum experiments (Forsythe *et al.* (1994); Hoffman *et al.* (1994)).

Figure 3.5: Data from ultimatum model in the first game and after 20 games.

The situation for an ultimatum proposer agent in our model is much more complex than for dictator proposers. But consider first the responder behaviour. Güth ascertained the following reasoning for second movers in ultimatum games: "If player 1 left a fair amount to me, I will accept. If not and if I do not sacrifice too much, I will punish him by choosing conflict" (Güth (1983, 384)). Two questions are involved. What allocation does the agent accept as fair, and what amount would it be willing to sacrifice in order to punish the proposer? The attribute used to decide whether or not an allocation is fair, is FAIRNESS CONCERNING ME. An equal split is obviously fair. In our model, an agent accepts an allocation to be fair if it gets at least FAIRNESS CONCERNING ME times half the cake. If a responder gets less, there is still a possibility that it accepts, when it would have to sacrifice too much. This is modelled as depending on the NEGATIVE RECIPROCITY attribute. If the responder's share is less than the proposer's share times the responder's NEGATIVE RECIPROCITY, it will reject a decision perceived as unfair. This yields the following responder behaviour. Offers of 40% are rejected by 13% of the responders, offers of 30% are rejected by 23%, and offers of 20% are rejected by 45%. The heuristic for responder behaviour is an example of a heuristic derived from questionnaires and expert knowledge. The thresholds are chosen to represent the data.

A proposer does not know the responders' heuristic. However, the proposer's reasoning has to include the possibility of a rejection. Possible proposer heuristics include:

1. Give what you expect to be accepted (half the cake times expected FAIRNESS CONCERNING ME).

2. Give half the cake.

3. Give a little more than what you would give in a dictator game (half the cake times FAIRNESS CONCERNING OTHERS plus 25%), but not more than half the cake.

In our model, the initial distribution of these heuristics is the same as for dictator games without any anonymity effects (consider only cases 3, 4, and 5). We did not implement an anonymity effect, because the punishment threat seems to outweigh anonymity concerns. Figure 3.5 presents the distribution of offers with the above heuristics in the first game and after 20 games with a simple learning process, which alters expected FAIRNESS CONCERNING ME and therefore the decision made with heuristic 1. Note that after 20 games there are no more offers of 20% or lower.

In ultimatum games, learning can take place in several ways. Proposers learn about rejection behaviour. Essentially, their value for expected FAIRNESS CONCERNING ME is altered. Only agents that employ heuristic 1 will adjust their decisions depending on previous experiences about the height of offers that were accepted and rejected. However, agents could also switch heuristics. An agent that used heuristic 3 and gave 50% of the cake, might reconsider, after observing that lesser offers were also accepted. That agent may try heuristic 1. For this kind of learning a more extensive learning process has to be implemented. In our model, different learning processes can be included, for example learning with cues (Pahl-Wostl and Ebenhöh (2004)) or case-based reasoning (Izquierdo *et al.* (2004)). However, in the ultimatum model presented here only the simple learning of others' FAIRNESS CONCERNING ME takes place.

A model validation is possible with respect to certain experimental data. In this case, the objective was to reproduce qualitative aspects of the data with different heuristics underlying the micro processes. The goodness of fit may be measured by the sum of the squared errors of the percentages giving 0, 10, 20, 30, 40, 50 percent. This is, however, no objective measure for model quality. χ^2 for the dictator model with less anonymity compared to the experiment with less anonymity is 4.4, and for dictator model and experiment with more anonymity it is 15.9. For the ultimatum model after 20 games compared to the experiment by Hoffman *et al.* χ^2 is 2.2, compared to the experiment by Forsythe *et al.* it is 24.3. Thus, the anonymity effect discovered by Burnham is not captured in this model. While one ultimatum experiment is reproduced well, the other is not. On the other hand, the qualitative aspects of the data are reproduced in either case.

The heuristics can be further improved by fitting the model to different decision framings (with lesser and higher degrees of anonymity, inex-

perienced and experienced players, and so forth). Our approach makes it possible to validate aggregated behaviour by a goodness of fit to a certain experiment, and at the same time validate the underlying micro processes by matching them to the individual data, results of questionnaires, and expert knowledge (cf. Moss and Edmonds (2005)). A second step in the modelling process would be to integrate expert knowledge to improve the effect of anonymity on dictator behaviour.

For a more extensive discussion of our model of dictator and ultimatum games see the website[2].

3.5 Prospects

The model presented here is a simple example for similar models for other experiments. The implementation would not have to be an agent-based model. However, our aim is to develop agents, that are able to play various games and that can be transferred to other decision environments. Agents need not change, when the model switches from dictator to ultimatum bargaining to voluntary contributions and so on. The reasoning changes and the heuristics to use. Also, in a learning process the expected attributes change. A more complex learning mechanism involves an assessment of a heuristic's usefulness in a certain context. This learning mechanism complies also with the bounded rationality theory, when it is implemented with aspiration levels and simple search rules for other heuristics. An extensive model of this learning process is part of our future work.

The model is part of a framework, that allows for easy extensions to similar games. The attributes and heuristics approach proves to be very generic and expandable within the setting of experimental economics. Different experimental games investigate the same set of attributes from various angles. Questionnaires and analysis of individual data provide heuristics, that are in part interchangeable between games. Thus, for further models, often only the decision environment has to be implemented, agents and part of the heuristics can be reused.

The question remains, whether economic experiments provide useful insight in human day-to-day behaviour. It is doubtful, that a person who gives 50% in a dictator game, can be generally classified as being fair in the sense of sharing all resources with others. The data has to be used with care and analysed with a special focus on the decision environment. However, we expect that data from experimental economics give some insights

---

[2]<http://www.usf.uos.de/∼eebenhoe/forschung/adaptivetoolbox/ultimatum.en.html>

for decision in bargaining situations and for behaviour in institutional settings related to common pool resource problems. Our model is, therefore, currently extended to include data from case studies where the specific context, previous experience, social relationships, and roles in a social network matter and where attributes may change over time.

Chapter 4

Modeling non-linear common-pool resource experiments with boundedly rational agents

Chapter 5

Agent-based modelling with boundedly rational agents

Chapter 6

A Topology Of Agent Attributes For Modelling Microbehaviour In Economic Experiments

Eva Ebenhöh[1]

6.1 Abstract

Manifold evidence for behaviour deviating from the predicted rational norm can be found in economic experiments. Explanations include integration of other-regarding behaviour in utility functions, emotions and norms, as well as irrationality on part of the subjects. These explanations have in common the use of certain attributes that humans seem to have. In modelling a number of games using data from economic experiments, we have developed a topology of attributes that helps to explain experimental micro behaviour. The attributes include cooperativeness and conformity, fairness concerning me and fairness concerning others, positive and negative reciprocity. This paper defines the use of these attributes in the models of experiments and presents a topology of how they relate to one another. Furthermore, it identifies characteristics of decision environments that call for the use of certain attributes.

6.2   Introduction

Agent-based modelling for social simulations opens up opportunities to create decision making units closer to real humans than in models of economic theory. Of course, this comes at a price. The decision making models become messy and less concise than theoretical equilibrium models. When they are not based on a clean theory they need another anchor in order to prevent them from becoming arbitrary. The first anchor we use is to base the models on observations and the second is to use the same underlying assumptions about agents' attributes across different models.

We use data from economic experiments to build agent-based models with heterogeneous agents that use heuristics for their decision making process and draw on a number of attributes. These models are designed as a basis for human behaviour in social dilemma situations as well as situations that call for fairness or reciprocity and should reproduce and help to explain heterogeneous data. These are situations in which the rational actor theory does not provide good explanations and predictions of human behaviour. Therefore, an alternative theory is needed.

Reasons for designing agents in this way are (1) a need for traceability of decision making and (2) that the diversity possible is a means to capture essential characteristics of actual decision making in simple models. A future use will be to model stakeholder behaviour in case studies of common pool resource situations including bargaining and joint decision making. In these more sophisticated models behaviour will be based on simpler models which will be extended with social environments. For this reason, we aim at simple decision making that stakeholders can easily understand and help to validate (Moss and Edmonds (2005)).

This paper takes a look at these models from a theoretical viewpoint with the objective to synthesise results into an integrative theoretical framework. It presents the underlying assumptions about agents' attributes and presents definitions and exemplary use of attributes in different decision environments. Describing different decision environments in which the attributes are used leads to an attribute topology.

6.3   Games in Experimental Economics

Economic experiments are conducted with human subjects, placed in an artificial, laboratory environment in which they have to solve tasks or play games with or against each other. Usually those games are played anony-

mously, in order to prevent the formation of a social environment. In these situations, decision making of human subjects is supposed to be abstracted from a specific context and social interactions. This may be seen as a breach with respect to actual human behaviour in day-to-day situations. However, this method's main advantage is that it produces comparable and reproducible data, which can serve also as a baseline to investigate the influence of a specific decision context and social interactions. In order to test for the effects of a specific rule or rule change in the setting, the decision environment is deliberately kept as simple and free of social aspects as possible.

For the models we used (1) social dilemmas, like prisoner's dilemma, voluntary contribution mechanism, and common pool resource games, (2) fairness games, like dictator games and mutual gift giving, and (3) negotiation games, like ultimatum games. In the following, only games used in this paper to illustrate the approach are described.

6.3.1 Dictator Game

The dictator game is a two player allocation game. The first player, called the dictator, is endowed with an amount of money and has the task to allocate the money between him or herself and the other player. The other player has no task. He or she only receives the amount allocated to him or her by the proposer. (eg. Hoffman *et al.* (1996); Burnham (2003))

The gift should only depend on whether or not or how far the dictator follows a fairness norm.

6.3.2 Ultimatum Game

The ultimatum game is a two player allocation game which represents a stylised version of the last stage of a negotiation process. The first player, called the proposer, has to allocate an amount of money between the two players, similar to dictator games. But now the second player, called the responder, has veto power. If the responder accepts the decision of the proposer, the money is paid as allocated. If he or she rejects the decision, both get nothing. (eg. Güth and Tietz (1990); Hoffman *et al.* (1994); Bolton and Zwick (1995); Hennig-Schmidt *et al.* (2002)).

In ultimatum games, proposers make strategic rather than fair offers. Acceptance, on the other hand, depends on whether or not the allocation is perceived as fair.

II   **II**

	C	D			C	D
I C	(a, a)	(d, b)	**I** C	(3, 3)	(0, 5)
D	(b, d)	(c, c)	D	(5, 0)	(1, 1)

Figure 6.1: The payoff matrix of a Prisoner's Dilemma game has four variables. In order for the game to be a Prisoner's Dilemma the relationship of the variables needs to be $b > a > c > d$ as in the example on the right hand side.

6.3.3   Prisoner's Dilemma

The prisoner's dilemma depicted in Figure 6.1 is a prototypical social dilemma, because individual rationality dictates defection, but mutual defection yields Pareto inferior results to mutual cooperation. Therefore, the group rational strategy would be mutual cooperation. (Colman (2003))

Agents playing prisoner's dilemma not only decide whether they would defect or cooperate. If they would cooperate, they also have to assess the possibility that the other defects, and decide whether or not to take this risk. Important aspects in this decision environment are the possibility to free-ride and the attractiveness of mutual cooperation.

6.4   Attributes

In each of the games discussed above, the experimental setting was designed to investigate deviations from predicted rational behaviour and to create situations that influence behaviour in one way or another. The behavioural regularities observable within and across different experiments are the foundation for the attributes presented here. The process of deriving attributes involves (1) a survey of the literature about a specific experiment; (2) an analysis of the micro data of the experiment; and (3) an implementation of the derived behavioural regularities in an agent based-model in order to test whether or not the experimental data can be reproduced by the assumptions made. The player agent with its set of attributes stays the same across different decision environments. Which attributes are used in a model depends on the specific context.

Figure 6.2: Different attributes used in relationships of the agent to the group, to an anonymous other agent, and to a particular other agent.

Our set of attributes (see Figure 6.2) includes six different characteristic traits. All attributes are represented as real numbers in the interval between 0 and 1, where 0 implies that this trait is not important for the agent and 1 implies that it is very important. The default setting is that the numbers are random numbers, equally distributed, and the attributes are independent from each other. Agents also expect other agents to have these attributes. The values of expectations are changed due to experiences made.

6.4.1 Attribute descriptions

COOPERATIVENESS

In social dilemmas there exists a tension between individual and group rationality. COOPERATIVENESS defines an agent's tendency towards group rationality. A high COOPERATIVENESS indicates the willingness to spend individual resources in order to increase group resources or joint outcomes. COOPERATIVENESS is used in win-win situations or positive sum games. It is used to determine free-riding behaviour in common pool resource games.

COOPERATIVENESS deals with the relationship between an agent and a whole group, but with a group that is poorly organised or anonymous. Cooperative settings are those "in which the individual is aware of his own and other's outcomes, and in which success is defined in terms of a collectivity who share an interdependent relationship" (Mc-

Clintock (1972c)). In these situations a high efficiency of cooperation increases importance of group rationality. On the other hand, the higher an actor's COOPERATIVENESS, the easier he or she responds to the cooperativeness of a decision situation. For instance, the higher the potential profit from cooperation the higher is usually the average cooperation in a group. Likewise, the more participants in the group, the less likely is cooperative behaviour. One advantage of these models is that the intrinsic COOPERATIVENESS can be modelled, as well as the aggregate effects on overall cooperation. With the assumption that agents with a high COOPERATIVENESS react faster on cues for cooperation we model a link between the two. Therefore, we can integrate theoretical approaches, which either explain cooperation based on the structure of the environment or based on the motivational orientations of the persons involved.

In our set of attributes COOPERATIVENESS is distinguished from fairness (see below). While COOPERATIVENESS is defined through importance of total group utility, fairness is concerned with inequality. In some cases this may be the same, while in other cases fairness would dictate a different behaviour than efficiency concerns.

COOPERATIVENESS is essential in the prisoner's dilemma. A highly cooperative player is willing to sacrifice potential profit in order to maximise joint outcomes. The percent of experimental subjects choosing to cooperate in first games varies with the numerical values in the payoff matrix and the framing of the task. The latter can be modelled by triggering norms, using CONFORMITY. The former is in accordance with different levels of COOPERATIVENESS.

CONFORMITY

If a group has decided on joint strategies or has developed shared norms, compliance to the rules is modelled as CONFORMITY. Agents with a high CONFORMITY may play fair because they feel they are expected to, and not because of their own high fairness. If behaviour is anonymous and untraceable, CONFORMITY becomes less important. However, it does not have to disappear, since norms may be useful to limit possible behavioural alternatives. CONFORMITY is used in relationships of an agent to a group of other agents, when the group is organised and has agreed on joint strategies.

The importance of norm compliance depends on the situation. Thus, similar to COOPERATIVENESS, different situations trigger CONFORMITY

in different ways. The tightness of the group, a process of joint decision making, and face-to-face communication can all enhance CONFORMITY. Again, we make the assumption that agents with a high CONFORMITY react easier on these cues.

In prisoner's dilemma games, CONFORMITY can be used in two ways. Consider the initial reference of the prisoner's dilemma, namely two prisoners sitting in separate cells, both being pressed to confess. Not to confess may be seen as a norm among those congenial criminals. In this case, prisoners with a high conformity will not confess no matter what their promised benefits. Further, in repeated prisoner's dilemmas with the same two players, repeated instances of mutual cooperation can turn into an institution where both players expect the other to keep cooperating and view mutual cooperation as their joint strategy.

FAIRNESS CONCERNING OTHERS and FAIRNESS CONCERNING ME

Contrary to COOPERATIVENESS and CONFORMITY, FAIRNESS deals with relationships between one agent and a single other agent. FAIRNESS is concerned with the difference between an agent's share and the other agent's shares (Fehr and Schmidt (1999)). It is divided into fairness with respect to the agent itself, also called envy, and fairness with respect to other actors, that is fairness in the true sense. FAIRNESS CONCERNING OTHERS defines the wish that other actors get roughly as much as the agent itself. Agents with a high FAIRNESS CONCERNING OTHERS are willing to spend money in order to equalize the outcome, if others have less. FAIRNESS CONCERNING ME defines the willingness to spend money in order to equalize the outcome if the agent itself has less than another agent. The latter is also used to model interest in personal utility.

Decision environments can provide different cues for fairness norms. For example, the existence or non-existence of a focal point (like "half") is of importance for fairness.

The basic game for FAIRNESS CONCERNING OTHERS is the dictator game. Ideally, no other reason than inequity aversion should constrain dictator behaviour. A simple example for the use of FAIRNESS CONCERNING ME is the prisoner's dilemma. Agents could be implemented to cooperate, if their COOPERATIVENESS is greater than their FAIRNESS CONCERNING ME and defect, if the opposite is true.

PositiVe reciprocity and negative reciprocity

> RECIPROCITY defines how much an agent's behaviour depends on pre-
> vious behaviour by other agents. It is divided into POSITVE and NEG-
> ATIVE RECIPROCITY, referring to positive or negative motivations per-
> ceived in the acts of others (McCabe and Smith (2001)). If an agent
> has a high POSITIVE RECIPROCITY, it will react in a friendly way on
> an act perceived as nice. If its NEGATIVE RECIPROCITY is high, it will
> tend to retaliate on acts perceived as hostile.
>
> Contrary to COOPERATIVENESS and CONFORMITY, RECIPROCITY deals
> with relationships between one agent and a single other agent, or more
> precisely a previous act of a single other agent.
>
> As for all other attributes, different decision environments also trigger
> reciprocity in a different way. How much one expected a certain act
> of another agent may influence how much the return act is based on
> RECIPROCITY.
>
> RECIPROCITY makes sense in repeated prisoner's dilemmas. If the
> other cooperated, a player with a high POSITIVE RECIPROCITY may
> be induced to cooperate in the next move. For a player with a low
> POSITIVE RECIPROCITY the indication to cooperate may have to be
> repeated. Tit-for-tat is a strategy depending solely on RECIPROCTY.

6.4.2   Expected attributes

In addition to these six attributes, agents hold expectations about the at-
tribute values of other agents. They learn from observed behaviour about
the others' attribute values. These are referred to as expected COOPERA-
TIVENESS, expected FAIRNESS CONCERNING ME etc.

Trust is an important characteristic of the relation between agents that
has a strong influence on the establishment of cooperative relationships.
Trust is modelled as expectations in other agents' behaviour (Cook and
Cooper (2004); Cox (2004); Nooteboom and Six (2003a)). Depending on
the situation this can involve different expected attributes. Trust that other
agents refrain from free-riding is expected COOPERATIVENESS, trust that
they keep promises or stick to joint strategies is expected CONFORMITY, and
trust that gifts are returned is expected POSITIVE RECIPROCITY.

6.4.3 Interaction of the decision environment with the attributes

As indicated in the discussion of the attributes the effect they have are not independent from the decision in which they are used. "Obviously, personalty traits are relevant only in specific situations which provide an opportunity and/or a stimulation to the behavior characteristic for the trait." ((Brandstätter *et al.*, 1999, 10)). The underlying assumption is that the higher the attribute in question the more an agent reacts on cues that indicate the use of this attribute. An indication for using COOPERATIVENESS, for example, is that group utility of a cooperative act is higher than of individually rational acts, like in the prisoner's dilemma. For very cooperative agents, it may be enough that group utility is only slightly higher. For less cooperative agents the efficiency increase needs to be much higher. Likewise, an agent with a low negative reciprocity may not even ask for the motivation of a previous act by another agent unless the situation triggers this kind of reasoning, for example by calling it a "demand". An agent with a high negative reciprocity is more easily looking for the motivation underlying the other's behaviour.

In this way, we explain both individual differences and collective shifts in response to changing rules. An extensive survey of cues Ostrom (2004)) for using attributes and the characteristics of the situation is beyond the scope of this paper, but for the models it constitutes an important part of the situation analysis.

6.4.4 Theoretical discussion of the set of attributes

By using our pattern of attributes in multi-agent models we can design agents that have all of these attributes and draw on the appropriate attribute(s) when confronted with a certain decision problem. Furthermore, we can take data from single attribute experiments (like dictator games) to define a distribution among our population of agents. The same diverse agents are then used to test data from other experiments with more than one attribute involved (like ultimatum games).

The set of attributes is founded in theoretical concepts but derived mainly from modelling practice. Therefore, the modelling technique of multi-agent systems and agent-based modelling had an influence in the choice of our particular attributes.

All attributes listed above deal with the relationship of the agent itself to other agents it perceives. They differ in important aspects of the relationship. One way to characterise them is the other part of the relationship.

Figure 6.3: Different attributes are used in relationships of the agent to an anonymous, unorganised group (cooperativeness) and to a group that has agreed upon joint strategies or norms (conformity).

This other part is either

1. the whole group or the environment, or

2. an anonymous other agent, or

3. a particular other agent and its previous acts.

For each of the three relationships constituted by the above list, there are two attributes to characterise the behaviour of an agent in such a relationship (see Figures 6.2 and 6.3 to 6.5).

If the partner in a relationship is the group as a whole, the attributes in question are COOPERATIVENESS and CONFORMITY. COOPERATIVENESS is the attribute characterising agent behaviour in an anonymous setting without previous agreements or shared norms. CONFORMITY on the other hand replaces COOPERATIVENESS if the group had the chance to agree on joint strategies or shared norms (see Figure 6.3). In a particular situation this differentiation may be fuzzy. For instance, in experiments often the subjects are recruited from undergraduate economics classes. As a subgroup, those subjects can be assumed to share norms.

Figure 6.4: Different attributes are used in relationships of the agent to an anonymous other agent whether the other has more (fairness concerning me) or less (fairness concerning others) than the decision making agent.

If the partner in a relationship is an anonymous other with whom the agent has had no previous experience, the attributes in question are FAIRNESS CONCERNING OTHERS and FAIRNESS CONCERNING ME. Which attribute to use depends on the situation (see Figure 6.4). If the agent feels that it has been treated in an unfair way, the attribute to use is FAIRNESS CONCERNING ME, if it feels, the other has been treated in an unfair way, it is FAIRNESS CONCERNING OTHERS. FAIRNESS may also be used in situations where an agent deals with another agent known to it but not directly in response to a previous action.

If the partner in a relationship is a particular other agent that made an act concerning the decision making agent, FAIRNESS is replaced by (or complemented with) RECIPROCITY. If the other agent's act is perceived as nice, the attribute in question is POSITIVE RECIPROCITY, if it is perceived as hostile, it is NEGATIVE RECIPROCITY (see Figure 6.5). It can also be the case that an agent reacts on the group as a whole. In that case, the agent can reciprocate a joint decision by a group.

Figure 6.5: Different attributes are used in relationships of the agent to a particular other agent when the other's act was perceived as nice (positive reciprocity) or hostile (negative reciprocity).

6.5   Use in the models

6.5.1   Experimental economics

Agents in the model should be able to recognise the characteristics of their decision environment and then choose the attribute(s) on which to act. Using the examples above, a first moving agent in a dictator and ultimatum game has to fulfill essentially the same task. Nevertheless, it is important, whether or not the other agent has veto power. In case the second mover has no veto power, the attribute in question is FAIRNESS CONCERNING OTHERS, else it is expected FAIRNESS CONCERNING ME. For the model, therefore, it is important that the agents can recognise the other agents' influence on joint outcomes. A responder with veto power, on the other hand, needs to know the choices the first mover had, that is the total amount that could be shared, in order to be able to act on its RECIPROCITY or FAIRNESS CONCERNING ME. In a one-shot prisoner's dilemma the attributes to use are COOPERATIVE-NESS and expected COOPERATIVENESS. In a repeated prisoner's dilemma, however, the range of possibly influencing attributes is wider. Therefore, the agent has also to know the context of a certain decision, if it possible to establish a long-term relationship or not. Also, the efficiency of cooperation is important.

In the models implemented so far (Pahl-Wostl and Ebenhöh (2004);

Ebenhöh (2006); Ebenhöh and Pahl-Wostl (2006)), we have given the agents the cues, which attribute to use. They have not decided autonomously. With the topology presented here and the link to decision environments it becomes feasible to let agents learn their environments and decide for themselves according to cues they pick up from the games they play.

An agent that has learnt the prisoner's dilemma can then be confronted with a new environment, like a chicken game. This is like a prisoner's dilemma but with parameter values $b > a > d > c$ (compare to Figure 6.1). If an agent assumes the other agent to defect, it would still be more profitable to cooperate itself. Therefore, it can assess the chance of the other actually choosing to defect as lower than before. Whether or not it reacts on this reasoning with cooperation or defection will depend on the agent's own COOPERATIVENESS and expected COOPERATIVENESS. The game provides the cue: low profitability of defection, assuming the other defects. So, even agents with a medium COOPERATIVENESS should cooperate, because of the reduced fear of being defected.

Another example is a dictator game with a prior beneficial act by the second player. Agents with a high POSITIVE RECIPROCITY will more easily search for the cue of a beneficial motivation by the other player and respond to it accordingly than agents with a low POSITIVE RECIPROCITY.

Even without the possibility or need to learn the environments, however, this topology can be used as a first approach at modelling a new situation. In a validation process of this first model approach through domain experts the use of the attributes can be enhanced and fitted to the characteristics of the situation.

6.6 Conclusion

This paper presents an approach to modelling agent behaviour in games. The presentation does not include whole models, because it omits the actual decision making behaviour and concentrates on differences in decision environments and resulting appropriate attributes for making the decision.

This topology already has been useful in implementing models to reproduce micro data from economic experiments, although the use of attributes alone does not define the model. The decision making mechanism uses heuristics that draw on the attributes. In the simplest case this is only a threshold of the sort: If the COOPERATIVENESS is greater than 0.5 cooperate, else defect. In more complicated heuristics more than one attribute can be used to define the decision. This topology can also help to design

experiments in order to study different, isolated aspects better, in order to distinguish between certain attributes, like Cox (2004) and Güth *et al.* (2003), or in order to investigate how different institutions affect them. A further use for modelling purposes is that the attribute topology helps to characterise decision environments. With these characterisations agents can learn new games. In more complicated settings, the characterisation of decision environments can help to analyse a certain situation and break it down into aspects that are similar to games presented above and other games. This makes it possible for the same agents to tackle also these more complicated settings.

Many aspects are, of course, not included in these six attributes, as for example risk aversion. Nevertheless, they form a core of a behavioural, theoretical approach to agent-based modelling of experimental data.

The approach has limits. The attributes are not as clearly defined in complex situations, where many of them may be important for a single decision. Also, many contradicting cues from the environment may constitute an ambiguity for agents (and modellers) which can not easily be resolved. This has to be subject to validation processes which involve domain experts, which may be experimenters, scientists doing case studies which utilise the model, or decision makers themselves. In these situations understanding of attributes may differ and be subject to the framing of the decision. When a model using this kind of decision making is used to represent real-world dynamics these ambiguities need to be carefully investigated.

Furthermore, the link between cues in the environment which call for certain attributes are not yet well understood. Subjects react differently on cues, which makes it hard to trace and formalise their influence.

Nevertheless, we expect this intuitive but also theoretical framework to be of help in analysing situations and decision making in various problems involving social dilemmas, and questions of reciprocity and fairness.

Chapter 7

Agent behaviour between maximisation and cooperation

Eva Ebenhöh and Claudia Pahl-Wostl[1]

7.1 Abstract

Evidence from economic experiments often shows a pattern of individual choice behaviour that can be classified into three categories. One category is maximisers, another cooperators, and the third is waverers with quickly changing behaviour due to changing circumstances and recent experiences. In this paper the focus lies on the third category consisting of subjects not decided for one extreme but waiting for indications from others' behaviour to make their own choice. This group can be drawn into mutual cooperation or mutual defection. It is this meta-stable behaviour resulting from social embeddedness that causes most social systems to be inherently unpredictable. Theories of utility maximisation and its extensions with fairness or other-regarding behaviour fail to explain this meta-stability, because they omit social embeddedness. In this paper we present a model of agent behaviour in one-sided and two-sided gift giving settings that includes this wavering behaviour drawing on the theory by Ostrom (2004).

7.2   Introduction — The Problem

The dilemma in social dilemmas is a gap between individual rationality and group rationality (cf. Heckathorn (1989)). The dilemma can be defined in the following way: if each actor follows the individual rationality the result is Pareto inferior to another achievable result. That means that each actor is better of, if each actor follows the group rationality rather than the individual rationality. Group rationality is threatened by free-riding behaviour. The question in this and many other studies is under what circumstances people do cooperate.

In experimental economics a strong distinction between one-shot games and repeated interactions is made. In small social communities, one-shot interactions are virtually non-existent, but they serve as a baseline for understanding repeated interactions. Experimental evidence suggests that even in one-shot situations a non-negligible percentage of subjects cooperates. Within rational actor theories it is possible to explain this behaviour by incorporating other regarding preferences into the utility function (cf. Fehr and Schmidt (1999); Bolton and Ockenfels (2000)). However, this fails to explain the varying behaviour found in repeated interactions.

When interactions are repeated, mutual cooperation can be established and maintained or it can be undermined and broken down. An analysis of micro behaviour in repeated interactions clearly indicates that the majority of the subjects change their behaviour over time according to circumstances and experiences they make (Pahl-Wostl and Ebenhöh (2004)). Bicchieri (2002, 221) argues that after observing the group's previous behaviour "subjects quickly conformed to the behavior of the interacting others, regardless of whether it was cooperation or defection."

Yet the reactions to the same kind of experiences can be very different from actor to actor. This can only be investigated by recognising the social embeddedness (Granovetter (1985)) of actors. In accordance with Moss and Edmonds (2005) we define *social embeddedness* "as a state in which an actor is significantly influenced by individual relationships with other actors. That is, the influences on the actor can not be modelled as if all actors were identical". In many rational choice models, social embeddedness is neglected, presumably because it is hard to incorporate into the models. However, it should be considered, if it is, as we believe to be the case here, an important causal explanatory factor. With reference to the recent discussion by Lovett (2006) about where to draw the dividing line between the exogenous and the endogenous, we argue that intrinsic trust can not be explained without social embeddedness which therefore has to be incorporated in the model.

We are interested in reasons and mechanisms for either path (establishment or break-down of cooperation). Analytically, we distinguish three categories of actions: (1) aiming at initiating and sustaining cooperation; (2) refraining to enter a cooperation or exploiting an existing one; and (3) wavering in between.

Many studies of cooperation in social dilemmas (including such fundamental analyses as Heckathorn (1989, 1996) neglect the third category. The analytical confinement to the extremes is predefined when using the Prisoner's Dilemma or other dichotomous situations (Bicchieri (2002); Harvey S. James (2002)) in which the subjects are forced to make an "either-or" choice. In other studies of public goods and common-pool resources (Ostrom *et al.* (1994); Fehr and Gächter (2002)) the choice is more open, allowing the subjects to make a "how much" choice. In these cases the category of waverers usually make up a substantial part of experimental subjects.

We are interested in determinants of the percentages of these three different kinds of behaviour. Since social dilemmas are situations in which outcomes depend on choices of all involved decision making actors, their expectations of the behaviour of others influence their choices. In this case, "the distribution of 'types' in a population of players" (Lovett (2006, 259)) is an important determinant of the situation.

One objective of our research program is to implement agent-based simulation models in which agents[2] represent human decision makers and make choices in clearly defined, simple situations in a way which is as close to human decision making as possible — not only substantively but also procedurally (Simon (1997, 293f.)) close to human decision making.[3]

In our models, we are able to incorporate strategic behaviour as well as intrinsic motivation and social embeddedness in the sense that behaviour is contingent on others' behaviour. The underlying assumption on rationality is one of bounded rationality (Simon (1997); Gigerenzer and Selten (2001a)) or program-based rationality described by Vanberg (2002). However, the short-term interactions we are describing in this article do not offer room for adaptive learning of programs. We therefore treat the programs and

---

[2]In this paper the term "agent" is used for software-agents that are programmed to reproduce behaviour of human subjects in economic experiments. Consequently, agents are referred to as "it". The terms "subjects" or "participants" are used for the humans taking part in these experiments and are referred to as "he or she".

[3]The reason for this desired closeness is that we use this kind of agent behaviour in participatory model building or social learning processes based on agent-based simulation models. The stakeholders involved in these processes should be able to identify with and validate the decision making of software agents.

the differences between agents as given for these interactions. In order for assumptions on programs and individual dispositions not to be arbitrary or ad hoc, we base them on data from economic experiments.

Much has been written about strategic behaviour, but here we are interested in aspects of fairness and reciprocity as determinants of cooperation without strategic elements. In order to be able to abstract from strategic behaviour we look at one-shot games, in which no further interaction is anticipated. This one-shot situation is contrasted with the second and last step in a two-step game, in which also, no further interaction takes place, but the second step is contingent on the other player's decision in the first step. Thus, we examine the link between one-shot and repeated interactions by focusing on the last step of an interaction. This last step shares independence from strategic elements with one-shot games, while it has the influence of reciprocity in common with repeated interactions.

The setting is a two-sided gift giving game, also labelled investor-trustee game. In this game, an initial gift is either reciprocated or not. We first consider the dictator game as an example of one-sided gift giving and as a baseline for the investor-trustee game. Trustee behaviour is then compared to dictator behaviour. We argue that having received a gift before making one has an impact on gift-giving behaviour. This impact, in turn, depends on individual dispositions and prior expectations.

In order to investigate this game, the theory by Ostrom (2004) is adapted for this nucleus interaction. This theory as well as data on this game (Berg *et al.* (1995); Fehr and Rockenbach (2003); Cox (2004)) are drawn upon to develop an agent-based model. The model is developed within a framework for agent-based modelling of micro data of economic experiments. It not only reproduces the data, but offers possible explanations of individual choice behaviour in these experiments. In this model, agents are socially embedded in the sense that the responses to gifts do not only depend on the initial gift but also on the expectations of the responding agent, which in turn depend on its experiences. The model presented here is analytical, meaning that for this particular case, it would not have been necessary to develop an agent-based model, because it presents only one step in a dynamic process, which is not explored further in this paper. We present the model as an agent-based model, because for our explanation of subject behaviour we need agent diversity and the model was developed within an agent-based modelling framework.

Figure 7.1: Adapted and extended from Ostrom (2004, 51). High trust enhances reciprocity (A), cooperativeness, and fairness. High reciprocity enhances cooperation in a group and therefore the benefit in a social dilemma (B), which in turn creates a good reputation (C) and increases trust (D). If a good reputation is established, this also increases trust (E). Dashed arrows indicate the feedback of previous interactions within a group.

7.3 Reciprocity, trust, and cooperation — The framework

7.3.1 The theory by Ostrom

In her theory of collective action Ostrom (2004) lays out the core relationships between reputation, reciprocity, and trust (Figure 7.1). She argues that in social dilemmas the level of cooperation depends on trust within the group and on reciprocal behaviour.

> "Behavior in social dilemmas can be better understood if boundedly rational individuals are assumed to enter situations with an initial probability of using reciprocity based on their own prior training and experience. The more benefits they have received in the past from other reciprocators, the higher their own initial inclinations. The more often they have faced retribution, the less likely will they be to see free riding as an attractive option. Their trust that others will also be reciprocators is highly correlated with their own norms but is affected by the information

they glean about the reputations of other players and their es-
timates of the risk of extending trust, given the structure of the
particular situation." Ostrom (2004, 49)

This is represented in Figure 7.1 in the following relationships among the
variables: Increasing actors' trust in reciprocal behaviour of others increases
not only the level of cooperation, but also their own reciprocity, that is their
willingness to reciprocate nice acts with nice acts and punish acts perceived
as hostile (arrow labeled A in Figure 7.1). By increasing an individual's
reciprocity, the net benefit in a social dilemma increases, because it is more
likely that this individual behaves cooperatively (arrow B). Good experiences
in turn, create a reputation for cooperative behaviour and punishment of free
riding (arrow C) and increase trust in reciprocal behaviour of others (arrow
D). Finally, a good reputation also increases trust (arrow E). Trust and
cooperation can be influenced by other variables depending on the specific
situations in which the actors face their social dilemmas. This feedback
loop, of course, can also work in a "negative" way. That is, low benefits and
unpunished free riding decrease reputation for reciprocity and trust, and
thus decrease reciprocity and cooperation levels.

We have adopted this core feedback loop in our model framework with
some variations. Ostrom takes a view at social dilemmas that integrates
these variables on the group level, while the argumentation is also made
on the individual level. In order to design agent-based models, we mainly
look at the individual level and emphasise differences between individuals'
behaviour. The variables (reciprocity, trust, level of cooperation, and rep-
utation) are more rigidly defined, and fairness is added (see below). Our
approach rests on the methodology of agent-based modelling and uses Os-
trom's theory as undelying rationale.[4]

To illustrate the approach and guide the analysis we use the feedback
loop introduced in Figure 7.1 for the two-person relationship that is inves-
tigated in this paper. We assume two such loops, one in each individual
linked via the joint cooperation level (see Figure 7.2).

In addition to Ostrom's assumption that individuals "enter situations
with an initial probability of using reciprocity", we further assume them to
enter these situations with an initial inclination to behave cooperative and
fair, as defined in the next subsection.

---

[4]Of course, Ostrom (2004) proceeds to investigate the impact of structural variables,
like group size, heterogeneity of interests, and the possibility of face-to-face communica-
tion, on the level of cooperation in a group. These variables are neglected here, because a
very simplistic experimental environment is used to illustrate our approach.

Figure 7.2: Adapted and extended from Ostrom (2004, 51). Two persons are locked in mutual dependency in a social dilemma. Their actions influence the other's behaviour and vice versa.

7.3.2 Agent Attributes

In our modelling framework agents have a number of attributes (Ebenhöh and Pahl-Wostl (2006)) which make up their heterogeneity and determine their decision making, although the influence of structural variables is also recognised. All attributes are modelled as real numbers between 0 and 1. They are independent from each other.

Reciprocity. "All reciprocity norms share the common ingredients that in-
dividuals tend to react to the positive actions of others with positive
responses and to the negative actions of others with negative responses"
Ostrom (2004, 42). Similar to Ostrom, we distinguish between positive
and negative reciprocity. Further, we assume that the two may be dif-
ferently pronounced in a particular agent and therefore model the two
as two different attributes. An agent with a high positive reciprocity
will react nicely on an act perceived as nice. Similarly, an agent with
a high negative reciprocity will react hostile on an act perceived as
hostile. Since negative reciprocity is not used in the model presented
in this paper, in the following we refer to positive reciprocity simply
as reciprocity.

Cooperativeness. Ostrom argues that cooperative behaviour results from norms of reciprocity and is conditional or at least highly dependent on the expectation that others are also reciprocators. She adds, however, that there are unconditional cooperators and individuals who never cooperate (Ostrom (2004, 41)).

> "Some individuals will cooperate in dilemmas only when they have publicly commited themselves to an agreement and have assurances from others that their trust will be returned. Others find it easier to build an external reputation by building their own personal identity as someone who always trusts others until proved wrong. Such an individual will always initiate cooperation in social dilemmas even when there are no explicit agreements." Ostrom (2004, 45)

Along this line, we model cooperativeness explicitly as an agent trait similar to reciprocity. Conditional cooperation is modelled as a combination of reciprocity and cooperativeness. An agent with a high cooperativeness is inclined to initiate cooperation in a group, that is, it will cooperate without assurance that others do so as well. This is not necessarily unconditional cooperation, because in later stages of an interaction, experiences can cause the same agent to refrain from cooperation. It is harder to draw an agent with a low cooperativeness into a joint cooperation than an agent with a medium cooperativeness.

Fairness. In Ostrom's theory, fairness as an individual trait is not included, probably because in social dilemmas, group rationality is more improtant than distributional aspects. In dictator games, however, in which the distributional aspect is dominant, individual fairness norms are used to explain behaviour of experimental subjects (Rabin (1993); Fehr and Schmidt (1999); Bolton and Ockenfels (2000)). We distinguish between fairness[5] and cooperativeness by the aspect of efficiency. Fairness as an agent attribute is concerned with distribution of payoffs only while cooperativeness is assumed to be coupled with an increase in group returns.

Trust. Trust is not modelled as an agent attribute, but as expectations of other agents' attributes. Ostrom's definition of trust quoted above:

---

[5]In other models, we further distinguish between fairness concerning others and fairness concerning me (envy). Fairness in this paper refers to fairness concerning others. The other attribute is not relevant in this model.

"trust that others will [...] be reciprocators" is incorporated in the model as high expected positive reciprocity, if it is trust that nice acts are reciprocated in a nice way, and as high expected negative reciprocity, if it is trust that hostile acts are punished. We also model "trust" that another agent will cooperate as high expected cooperativeness, and "trust" that the others will share resources as high expected fairness. Expected reciprocity, cooperativeness, and fairness can be altered quickly due to experiences an agent makes in a particular group.

7.4 One-sided gift giving — The dictator game

7.4.1 Dictator experiments

A dictator game consists of two players. The first receives an amount of money (the cake) and has the task to share it with or distribute it between the two players. The second player has no task and receives the amount of money assigned to him- or herself by the first player.

Looking at data from experimental economics, it is undisputed that there is evidence for both cooperative *and* maximising behaviour. In all dictator experiments, for example, there is a non-negligible percentage of subjects who give a positive amount to the other player, although they do not have any pecuniary benefit from this act.

Different dictator experiments yield different results. In part these differences can be explained by a great variance in individual behaviour across the human population, which we can also reproduce in these models. Social scientists, however, also search for factors which nudge aggregate behaviour in one way or another even in one-shot games. When there is no dynamic feedback, influencing factors can only be found in the setting as well as the prior experiences and individual differences in character among players (for instance Brandstätter *et al.* (1999); Ben-Ner *et al.* (2004)). Setting effects include anonymity (see below), the existence of a show-up fee[6], actual payment (Forsythe *et al.* (1994)), and the cultural or social group from which

---

[6]One explanation for the exceptionally high dictator gifts in the experiment by Güth and Huck (1997) may be the lack of show-up fees. S. Huck (pers. comm.) pointed out that getting to miss a class may have positive utility, but we will not try to assign a monetary value to this. In the other dictator experiments studied here show-up fees ranged between 30% and 150% of the pie. Consider a show-up fee of 3\$ and a pie of 5\$ (10\$) as in Forsythe *et al.* (1994). The minimum a second player gets from total payoff is 27% (19%), if the dictator keeps everything. It is likely, that in case of side-payments, decision heuristics are different than if there are no side payments. Nevertheless, in this paper the explanation of differences of experimental results focuses on anonymity.

the participants are drawn (Henrich *et al.* (2001a)).

In this paper, we focus on anonymity effects, because these have been investigated in detail and exhibit fairly unambiguous results.

Anonymity   According to Hoffman *et al.* (1994, 1996) gifts are lowest, when anonymity from the experimenter is assured through a double blind procedure. They compare this setting to one in which only cross-subject anonymity is assured conducted by Forsythe *et al.* (1994). In the double blind procedures in Hoffman *et al.* (1994) the average gift is about 10%. In the less anonymous treatment of Forsythe *et al.* (1994) it is 23%. Still, even in the double blind treatments, about 10% of the subjects gave half the cake. In Forsythe *et al.* (1994) there are 21% fair dictators. In the double blind procedures about 60% give nothing, in the less anonymous treatment, only 21% give nothing.

Anonymity can be altered in various ways. Burnham (2003) conducted an experiment in which dictators were shown photos of their opponents, or told that the opponents were shown photos of them. When dictators are shown photos of their opponents, the average gift is 17%; when the recipients are shown photos the average gift is 20%; and in the double-blind control group it is 12%. What is most interesting in this study is that the percentage of zero gifts in all three games varies only between 54% and 58%.

Güth and Huck (1997) conducted dictator games in a school with cross-subject anonymity. The participants did not know the opponent, but knew he or she would be another pupil of the school if not from the same class. 45% of dictators in this experiment gave at least half the cake. The average gift was 32% and only 6% of the subjects gave nothing.

7.4.2   Modelling dictator games

Since the issue of anonymity is investigated most profoundly in the experiments and yields the clearest results, we chose to concentrate on anonymity effects. The experiments described above can be classified into three categories:

	low	high	offset	decided	wavering
Güth and Huck 1995	0.06	0.55	0.1	51%	49%
Forsythe et al., 1994, Hoffman et al., 1994	0.28	0.81	0	47%	53%
Burnham, 2003, Hoffman et al., 1994	0.60	0.92	-0.47	68%	32%

Table 7.1: Parameter settings for models of different dictator experiments. The least anonymous setting is shown in the first row. The second row displays a moderate setting and the third row corresponds to the double blind procedures which ensure the highest degree of anonymity.

1. The least anonymous setting is the school experiment by Güth and Huck (1997). For our analysis we pool data from the two different cake sizes of 16 and 38 DM[7].

2. The setting with moderate anonymity is the experiment by Forsythe *et al.* (1994) and its replication by Hoffman *et al.* (1994, Fig 4a on page 365). Data from experiments with different cake sizes of 5 and 10$ are pooled.

3. The most anonymous setting are the double blind experiments by Hoffman *et al.* (1994, Fig. 4c and 4d on page 366) and Burnham (2003) in the no-photo treatment. Data from these three experiments are pooled.

The data of these settings are displayed in the left column of Figure 7.3, which displays the percentages of subjects giving gifts in categories of 0-5%, 5-15%, 15-25%, and so on. One important finding of the individual data of these experiments is behaviour of subjects who choose neither to keep everything to themselves nor to share 50:50. Between a third and one half of the subjects gave more than 5% and less than 45%. We call this group "waverers". Please note that the highest percentage of wavering people is found in the least distinctive setting (see Table 7.1).

As a first approach to the model, we use the agents' fairness attribute in the following way:

- Dictators with a low fairness keep everything.

---

[7] In 1994, the time of the experiment, 3 DM approximately equaled 2$.

Figure 7.3: Dictator experiments and model results. The top left figure shows results of the experiment by Güth and Huck (1997) (cake size 16 DM, N=22; and cake size 38 DM, N=42). The top right figure shows the model results with parameters set to fit this experiment. The middle row shows results of experiments by Forsythe *et al.* (1994) (cake size 5$, N=21; 5$, N=24; 10$; N=24) and Hoffman *et al.* (1994) (cake size 10$, N=24) and the corresponding model results. The bottom row shows results of experiments by Burnham (2003) (cake size 10$, N=26) and Hoffman *et al.* (1994) (cake size 10$, N=36; and 10$, N=41) with double blind procedures and corresponding model results.

- Dictators with a medium fairness give a positive amount but less than half. In our model the heuristic for this behaviour is

$$gift = (\min(\mathit{fairness} + \mathit{offset}, 1)) * (\frac{cake}{2}) \qquad (7.1)$$

- Dictators with a high fairness give half the cake.

With parameter variations that determine what is *low*, *high*, and *offset*, the model can be fitted to a particular data set.

We do not allow gifts greater than half the cake, and obviously not below 0, in case the offset becomes negative. The *offset* shifts the whole of the wavering agents in one direction or another. By this, we can either increase the average gift and the percentage of fair people, or decrease the average gift and increase the percentage of maximisers. In order to increase or decrease the percentage of wavering, we can alter the parameters *low* and *high*. By this design we model two effects. A change in offset represents an incremental, continuous change, within the same decision making rationale. By altering *low* and *high*, we alter the percentage of decided people.

The three different model results displayed in Figure 7.3 are obtained by using the parameters *low*, *high*, and *offset* according to Table 7.1. The same model structure can be used to approximate all three classes of experiments by varying parameters.

The assumed underlying mechanism is that different degrees of anonymity or sense of group membership trigger fair behaviour in a different way. The parameter values presented here so far, are a result from data analysis and theoretical modelling, and did not result from agent-based modelling. The latter is done by giving the software agents the possibility to perceive the degree of anonymity in their respective environments and react (differently) on it. To obtain the described effects, those of our agents with low and medium fairness react first on increases in anonymity and those with a medium or high fairness react first on decreases in anonymity.

7.5 Two-sided gift giving — The investor-trustee game

7.5.1 Investor-trustee experiments

The dictator game was introduced in the last section as a baseline for the next experiment. In the two-sided gift giving game discussed below, the decision by the second player is similar to the dictator game, but depends on

a previous decision by another player. By comparing these two experimental settings we can focus on how the initial act by the first player influences the second player's decision and thus investigate the effect of positive reciprocity.

The form of mutual gift giving we consider follows Berg *et al.* (1995), Fehr and Rockenbach (2003), and Cox (2004). The game in these experiments is designed to investigate trust. Both players receive an amount of 10 money units (MU)[8]. The first player, called investor, can decide which amount of his or her MU to give to the other player. This amount is tripled by the experimenter, introducing efficiency concerns into the game. The second player, called trustee, then decides which amount between 0 and the tripled gift he or she returns to the investor. The return gift is not tripled.

We are concerned with the second decision, which we interpret as a dictator game. The difference to other dictator experiments is the fact, that a trustee-dictator responds to a previously made decision by an investor-responder and not to an anonymous other who has not made a more or less beneficial move first. The trust placed by the investor is a form of trust that can not be controlled by future interactions (like in Raub (2004)) and it does not have to be either full trust or no trust (as in dyadic trust games in Harvey S. James (2002)) but can be in the range from 0 to 10 MU investment. Based on Ostrom's theory, the process is displayed in Figure 7.4.

Most investors give gifts greater than 0, in average 5.2, 6.0, and 6.5 MU (in Berg *et al.* (1995), Cox (2004), and Fehr and Rockenbach (2003), respectively). Trustees returned 4.7, 4.9, and 7.8 MU on average. The percentages of trustees returning nothing was 30%, 35%, and 18%.

In order to investigate the trustee behaviour as a dictator behaviour the initial gift has to be taken into account. Whether or not an initial gift is seen as a beneficial act may depend on the expectations of a trustee. If a trustee expected a 10 MU initial gift, he or she may be disappointed by a 5 MU gift and return less (in terms of tripled gift) than on a higher initial gift.

In case an investor gives 10 MU, a trustee has no reason to see the gift other than as a beneficial act. In the baseline experiments of Berg *et al.* (1995) and Cox (2004)[9], there have been 18 instances of gifts of 10 MU,

---

[8]In Berg *et al.* (1995) and Cox (2004) there was no additional show-up fee and each MU was worth 1$. In case of Fehr and Rockenbach (2003), on the other hand, there was a show-up fee of about 3 DM and each MU was worth 0.5 DM.

[9]In Fehr and Rockenbach (2003) investors had to indicate how much they would like to get back as a return gift. This indication can itself be seen as beneficial or hostile and therefore this situation is even more complicated to analyse. So, we only considered the

Figure 7.4: As an alteration to Figure 2, this Figure shows the interactions of an investor and a trustee.

enough to compare them to dictator experiments. Assuming the beneficial or trusting act to enhance second players' return gifts, we should find higher gifts than in dictator experiments. The average return gift in these 18 instances is 8.6 MU, which is 29% of the tripled initial gift, but only 21% of the total of 40 MU. This is more than in experiments with high anonymity, but less than in the school experiment. Surprisingly, 44% of trustees in this situation return nothing, 39% return half of the tripled gift or more, and only 17% return something in between. The wavering middle has almost disappeared.

This can be explained on the basis of the model presented in Figure 7.4, because a high initial gift increases the reputation of the investor in the eyes of the trustee. This in turn increases the recipient's willingness to reciprocate and act in a fair way. Those that are wavering because of their initial medium reciprocity and fairness, may react on the cue of a high initial gift. Those that are more certain about their actions, because their fairness and reciprocity are either high or low, do not change their behaviour due to the initial beneficial gift.

In case the investor gave 5 to 8 MU (there was no instance of a 9 MU gift) the returns should be lower, not only in MU but also in percent of tripled gift. In the 24 instances of 5 to 8 MU gifts in the baseline experiments of

---

true baseline experiments.

Figure 7.5: On the left hand side are displayed experimental results (Berg *et al.* (1995); Cox (2004)), in the top row for trustee returns on highest possible gift in percent of tripled gift (N=18) and in the bottom row for trustee returns on medium gifts of 5 to 8 MU in percent of tripled gift (N=24). On the right hand side are percentages resulting from a first model of this experiment.

Berg *et al.* (1995) and Cox (2004), the average return gift was 32% of the tripled gift. Thus, this case does not clearly support a linear relationship between an initial trusting act and reciprocal behaviour of trustees. 29% of trustees returned nothing and also 29% returned half of the tripled gift or more. 42% gave something in between, which makes the percentage of waverers substantially gretaer than in the case of highest gifts.

The experimental results of trustees are displayed in Figure 7.5 in the two categories of initial gifts of 10 MU and 5-8 MU.

7.5.2   Modelling reciprocity and trust

A first approach is to model the trustee behaviour similarly to the dictator setting. This has been done for the two different classes of initial gifts of 10 MU and initial gifts of 5-8 MU. The parameter settings can be seen in Table 7.2 and the model results are compared to the corresponding data from Berg *et al.* (1995) and Cox (2004) in Figure 7.5.

As can be seen from the parameters in Table 7.2, there is a shift to-

	low	high	offset	decided	wavering
10 MU gift	0.44	0.61	0	83%	17%
5-8 MU gift	0.29	0.71	0	58%	42%

Table 7.2: This table presents parameter settings for models of trustee experiments, when the trustees received initial gifts of 10 MU and 5-8 MU.

wards nicer behaviour compared to highly anonymous and moderate dictator games, but not so far as in the school experiment by Güth and Huck (1997). As can be seen from the graphs, the case of trustee returns on medium gifts is not reproduced well. This is because the model introduced above is not capable of reproducing a thinning out of the middle without a movement to either side as in the cases of low or high anonymity.

Furthermore, this first approach conceals the distribution within the category "50+". In case of normal dictator games, there is only occasionally an offer above 50% (in the dictator experiments displayed in Figure 7.3, 2% of gifts were higher than 55%). Giving half is an undisputed focal point for those who want to share in a fair way. In case of trustee experiments, not only does the money originally come from an initial gift of the first player, and therefore can be thought of as belonging to the other. Also, the amount to be split may be considered to be the tripled gift plus the 10 MU the trustee originally obtained. Therefore, there is no single focal point. Among the percentage of participants who obviously want to share in a fair way, perceptions on what is fair, differ in a far greater range than in dictator games. Figure 7.6 shows the distribution of return gifts up to 100% of the tripled gift.

In order to model these differences we chose a two-step heuristic. The first step depends on the fairness attribute as in a dictator game, but is valid only for those with a very high or a very low fairness. Trustees with a very high fairness give not half of the tripled gift but two thirds of it, thus returning the original investment and splitting the surplus. The second step works essentially in the same way as in dictator games before, but using the attribute reciprocity instead of fairness. The process is displayed in Figure 7.7.

Since now there are more parameters, we use *low* and *high fairness* from the intermediary setting of dictator games in Table 7.1. *low* and *high reciprocity* result from fitting the model to a data set. Parameters for this second model are presented in Table 7.3. Results are displayed in the right

Figure 7.6: Displayed are experimental results on the left hand side for trustee returns in percent of tripled gift, for initial gifts of 5 to 10 MU in the baseline experiments of Berg *et al.* (1995); Cox (2004). On the right hand side are percentages resulting from the second model of this experiment.

Figure 7.7: In the second model version, an agent makes a two-step decision. If its fairness is very low or very high, the decision is based on this attribute. Else, the decision is based on positive reciprocity.

low fairn.	high fairn.	decided	low recipr.	high recipr.	decided	wavering
0.29	0.81	47%	0.16	0.74	70%	30%

Table 7.3: Parameter settings for the second model of trustee experiments for gifts of 5 to 10 MU. All offsets are equal to 0.

hand side of Figure 7.6. This process can not produce return gifts of 60, 80, and 100% of the tripled gift, but only two thirds of it, which are in the category 70%.

It would have been possible to spread the fair behaviour from 60 to 100% of the gift, similar to how wavering behaviour is modelled in the dictator game. However, since only very few gifts of those above 50% of the trippled gift were not two thirds, we chose to keep the model simple.

Again, the parameters result from the framework and data analysis. In an agent-based model, agents would have to sense the beneficial motivation underlying another agent's acts and react on this motivation according to their own reciprocity. To get the above results, agents with a high reciprocity react faster on perceived positive motivation by others than agents with a low reciprocity.

To complete the cycle of the core relationship of reciprocity, trust and cooperation introduced in Section 7.3, the return gift should create or reduce trust in the first player. This is the point, where we have to move from purely analytical models to agent-based models. We use a simple learning direction mechanism: The expected positive reciprocity of an investor is altered to become the mean of the previous value and the perceived positive reciprocity of the trustee in terms of per cent of the tripled gift. A full discussion of this model would also have to include an in depth discussion of investor behaviour and exceeds the scope of this paper.

7.6 Conclusion

From data we found that we can loosely order the dictator games discussed above:

1. Games with the lowest average gifts and highest percentage of maximising dictators are those with the highest anonymity achieved through double blind procedures (Hoffman *et al.* (1994); Burnham (2003)).

2. Next in line are experiments with high anonymity between subjects but not with respect to the experimenter (Forsythe *et al.* (1994); Hoffman *et al.* (1994)).

3. Even higher average gifts we find in experiments with prior benign acts by the receivers (Berg *et al.* (1995); Cox (2004)), although the percentage of maximisers is still high.

4. The highest average gift and lowest percentage of maximising dictators we find in an experiment within a school (Güth and Huck (1997)).

We conclude that the social setting matters, which does not come as a surprise. The higher the anonymity the lower the average gift and the higher the percentage of maximising dictators. However, even in the most anonymous one-shot treatment, there are some subjects who share equally. Compared to the school experiment, in the investor-trustee experiment the low percentage of wavering subjects is an interesting new aspect revealed by looking at individual rather than aggregate data.

We base our model on a formalisation of the theoretical framework of Ostrom (2004). The formalisation enforces an explicit and rigid definition of variables and their interdependency. By breaking the feedback loop up into two loops that are connected via the cooperation level, we can analyse the impact of single steps on individuals' decisions and differences in behaviour.

By looking at decision processes in very small steps, we can abstract more complicated decision situations and simplify them step by step. The simple decisions can be building blocks from which more complex decision mechanisms are made up, as in the bounded rationality theory (Gigerenzer and Selten (2001a)).

Investigating individual behaviour allows to focus on the wavering middle, which not only makes up a substantial percentage of subjects, but is also interesting to investigate, because they can tip a group towards mutual cooperation or mutual defection. Investigating especially this behaviour can be a promising strategy to make out which factors enable or hinder the establishment of cooperation.

The experimental data underscore heterogeneity. "Individuals vary substantially in the probability that they will use particular norms, in the way structural variables affect their level of trust and willingness to reciprocate cooperation in a particular situation, and in the way they develop their own reputations" Ostrom (2004, 45). This has implications for the formalisation and resulting models. In order to capture this heterogeneity we need to model the process as an agent-based process and give the agents not only

individual attributes or preferences, but also enable diverse reactions on structural variables or changes in these variables.

The nucleus behaviour introduced in this paper can be used in models with a richer social setting. This can include models adapted for specific case studies and group model building in case studies. Agent behaviour is contingent on the situation. The actual implementation of behaviour and parameterisation of the setting depends on the situation and on individual dispositions. Agent-based models do not determine how to incorporate these effects, but make it possible to implement and simulate them on the level of the individual. At the same time agent behaviour remains simple and traceable. This allows non-modellers to follow up on what the agents are doing in the model.

The models present a feedback of the action of one individual on the actions of another. Here we can only hint at experiments with 10 or more subsequent common-pool resource games (Isaac *et al.* (1984); Ostrom *et al.* (1994); Fehr and Gächter (2002)). The evidence from these experiments suggests that cooperation can be established, and if there is a means to sustain it, can last over several rounds. In addition to effects of sanctions (for instance Heckathorn (1989); Mulder *et al.* (2005)) and communication (for instance Bicchieri (2002)), initial group composition, efficiency of cooperation, and group size also play a role (Ostrom (2004)).

With this kind of modelling using the methodology of agent-based modelling but informed and grounded in sociological theory and based on data, we hope to contribute to building theory, which is neither arbitrary, because it is developed within a framework of contingent actions and based on observation, nor consists of non-testable assumptions about subjective beliefs, because it generates testable models (cf. Vanberg (2002, 14)).

Chapter 8

Classification of decision environments

Eva Ebenhöh

8.1   Introduction

This working paper introduces a classification system for decision environments for modeling purposes. The agents are boundedly rational and draw on a number of attributes as well as heuristics in order to make their decisions. Therefore it is necessary for them to be able to recognize in which kinds of decision environments which attributes have to be used and which heuristics are appropriate. This is no decision theory for all possible decision situations. Rather, this classification system grows with application (modeling practice) and is usable for modeling simple decisions in a realistic way.

Before this paper proceeds to introduce several approaches to decision theory and other classification schemes, requirements of this modeling technique are briefly discussed. Important aspects for a classification system resulting from previous modeling experiences are presented later. But in order to understand the angle from which the literature is discussed, it is necessary to bear in mind the following thoughts.

The agents have a personality, consisting of a number of attributes, which regulate their behaviour in different situations. As an example consider a proposer in a dictator game and in an ultimatum game. Somehow, this agent has to recognize the difference in these two situations, leading to more or less generous gift giving in the first, and to strategic acts in the second.

In dictator games the most important attribute is fairness concerning others: What allocation does this agent judge to be fair? In fact, this game is designed so that ideally no other aspect shall influence the decision. In ultimatum games, however, the situation is completely different. The relevant aspect is expected behaviour of the responder: What fraction is likely to be accepted? Thus the most important attributes become expected fairness concerning me and expected negative reciprocity. The classification system to develop here must facilitate recognition of these differences. Even in this simple example, several key aspects can be recognized: the other player(s) and their positions relative to the decision making agent; (a)symmetry across agents; self-perception; expectations of others' behaviour; the existence of a focal point. This will be discussed in more detail in section 8.3.

Agents do not only have to recognize, which attribute is appropriate, but also which heuristics apply. In the example above the narrowly defined decision task does not change, since the proposer has to decide on a fraction of a certain amount of money to assign to the responder. The decision task alone gives almost no cues for which heuristics to use. The only cue which is appropriate in both situations is using the focal point 50:50. Therefore, we can assume the equal split heuristic to be appropriate in both games. However, for dictators, who do not want to share equally, a decision heuristic has to change between dictator and ultimatum situations. In principle, there can be a decision heuristic which can be applied in both situations, which simply draws on different attributes: In dictator games, give the fraction (of half the cake), which corresponds to your fairness concerning others attribute; in ultimatum games, give the fraction (of half the cake), which corresponds to your expected fairness concerning me (or expected negative reciprocity, or average of these). However, structurally different heuristics are also possible. As before, we state that the classification system has to facilitate the recognition of appropriate heuristics.

These two demands, facilitating recognition of important attributes and of applicable heuristics, are made from a viewpoint of modeling practice instead of theoretical deliberation about differences in decision environments. We bear these in mind during discussion of different decision theories.

8.2 Review of relevant literature

8.2.1 Traditional approaches and counter-examples

Game theory in the tradition of von Neumann and Morgenstern as well as decision theory based on maximization of expected utility have been widely discussed (e.g. Kreps (1988)). The influence of (subjective) expected utility (SEU) maximization theory is undoubtedly great. The formal representation has huge advantages and creates numerous possibilities for further work. For our purposes, however, this theory is of limited use, except as a benchmark and limiting case. We focus explicitly on those decision situations in which SEU maximization does not predict observed behaviour very well. Deviations from predicted rational behaviour are also widely discussed (e.g. Colman (2003)). Some examples are listed below in order to motivate further discussion. Furthermore, expected utility theory "is not always plausible [and] often impractical" (Gilboa and Schmeidler (2001, 27)). This depicts the difference between substantive and procedural rationality, first pointed out by Simon (Rubinstein (1998)). SEU maximization may lead to substantively rational behaviour, but not necessarily to decisions as made by persons who deliberated in an appropriate way about a problem. Humans are capable of planning and utility maximizing decision making. But this is a limiting case and is done only in situations with very high stakes, competitive environment, much time, and abundant information (Ostrom (2004, 39)).

After an overview of some deviations from predicted rational behavior, this section proceeds to introduce several extensions and alternatives to traditional approaches that aim at explaining these deviations.

Prisoner's Dilemma

It would be so reasonable to cooperate in Prisoner's Dilemmas or any other social dilemmas, yet reasoning within the bounds of expected utility maximization dictates defection, leading the game to a Pareto-inferior Nash-equilibrium, see Figure 8.1. In finitely repeated Prisoner's Dilemma games, backward inductions leads to the same impasse. Only in infinitely repeated games, can mutual cooperation be established on the grounds of SEU maximization. However, even in one-shot Prisoner's Dilemma situations there is evidence for a substantial minority choosing to cooperate (cf. Colman (2003))

What feature is responsible for these observed deviations from predicted rational behavior in social dilemmas? Before we can answer this, we have to

		II				**II**	
		C	D			C	D
I	C	(a, a)	(d, b)	**I**	C	(3, 3)	(0, 5)
	D	(b, d)	(c, c)		D	(5, 0)	(1, 1)

Figure 8.1: Payoff matrix in Prisoner's Dilemma games with $b > a > c > d$ as in the example on the right hand side.

ask: Why does game theory predict mutual defection? According to Colman (2003) the answer is in the combination of the two assumptions of common knowledge and rationality:

> "CKR1. The specification of the game, including the players' strategy sets and payoff functions, is *common knowledge* in the game, together with everything that can be deduced logically from it and from CKR2.
>
> CKR2. The players are rational in the sense of expected utility (EU) theory [...] hence they always choose strategies that maximize their individual expected utilities, relative to their knowledge and beliefs at the time of acting. (By CKR1 this too is *common knowledge* in the game.)" Colman (2003, 142, emphasis in original)

So, players have to expect other players to behave selfishly utility maximizing. If this assumption is relaxed, anything becomes possible. And, indeed, anything happens: mutual cooperation, one-sided defection, and mutual defection. The prisoner's dilemma does not incorporate a self-sustaining mechanism for cooperation. Free-riding is possible and attractive.

Agents playing prisoner's dilemma have to know their own preferences, whether they would defect on the other or prefer mutual cooperation. If they prefer mutual cooperation, they have to assess the possibility that the other defects, and decide whether or not to take this risk. Important aspects in this decision environment are possibility to free-ride and attractiveness of mutual cooperation, which is, after all, Pareto superior to mutual defection. In addition, decision maker's self-perception, expectations of other

player's behaviour, and willingness to take the risk are factors that determine decision. The last aspect explains, why stakes can matter. If mutual cooperation is only slightly more profitable than mutual defection, risk may appear greater than if mutual cooperation is much more profitable.

Dictator Game

A proposer is given an amount of money and assigned the task to divide this money between him or herself and another player. The other player has no function in this game. Dictator experiments are usually conducted with anonymity. Nothing should prevent a proposer from keeping the total amount of this money, but a substantial minority gives 50% to the other player, and some give 20, 30 or 40% (cf. Forsythe *et al.* (1994); Hoffman *et al.* (1994, 1996); Güth and Huck (1997); Burnham (2003))

Again the question is, what feature is responsible for the deviations? In this case, there are several possible explanations, and the assumption of maximizing joint *and* individual gains by cooperating, as in social dilemmas, is not one of them. The usual explanations are that players

- care for the other's utility in addition to their own utility because of true other regard.

- have an intrinsic motivation of behaving in a fair way, either because of some fairness norm or because of self-perception.

- want to appear fair to the experimenters because of the existence of a fairness norm.

- choose an obvious focal point.

These alternatives are not easily separable. Another interesting aspect is that other-regard seems to diminish in more strategic interactions, like ultimatum games.

In this case, an important aspect of this decision environment is the existence of a focal point. Other determining factors are social environment (norm or other-regard) and, again, self-perception.

Centipede Game

The Centipede game illustrated in Figure 8.2 is another good illustration for the inappropriateness of backward induction as a method for reasoning. Any outcome from step 3 on is better for player I than the first two outcomes, but

I	II	I	II

(8, 8)

(0, 0)   (−1, 5)   (4, 4)   (3, 9)

Figure 8.2: An example for a Centipede game: if a player moves forward, his or her outcome is reduced by one and the other player's outcome is increased by 5.

	II				**II**	
	High	Low			Heads	Tails
High	(6, 6)	(0, 0)		Heads	(5, 5)	(0, 0)
Low	(0, 0)	(3, 3)		Tails	(0, 0)	(5, 5)

Figure 8.3: Payoff matrices in the High-Low Matching game on the left hand side and the Heads-and-Tails game on the right hand side.

the potential loss in step 2 prevents player I from entering the later stages. (cf. Colman (2003))

Again the assumption that the other player behaves selfishly utility maximizing leads game theoretic analysis to predict that player I does not enter the game. Relaxing this assumption, one can explain that players enter the game and continue until later stages until one of them moves down and ends the game.

As in Prisoner's Dilemmas, from the perspective of agents playing Centipede Games, important aspects of this decision environment are the possibility of defection and attractiveness of mutual cooperation. Determining factors of the decision are expectations of other players' behaviour, self-perception, and willingness to risk being defected.

High-Low-Matching Game

In Figure 8.3 two pure coordination games are presented. Virtually every player chooses *high* in the High-Low-Matching game. Even in the Heads-

and-Tails game 86% of subjects choose *heads* (Colman (2003)). However, game theory does not predict these outcomes. This is a fundamental game to prove the importance of a socially defined focal point for real human decision making. Agents can expect other agents to choose high, and there is absolutely no reason not to choose high oneself.

Framing Effects

Apparently framing of decision problems has an influence on decision making, even when it has no effect on payoffs or outcomes. Anonymity effects in dictator games are an example. Also, the exact words to "divide" an amount of money may lead subjects to believe that the task is to split the money evenly (Hoffman *et al.* (1996)).

Another example is the apparent difference between positive and negative risks. Kahneman and Tversky conducted an experiment in which two sets of choices were identical except for wording of sure gains or sure losses. Subjects choose sure gains over risky choices and risky choices over sure losses (Kahneman and Tversky (1979)).

In light of the definition of rationality of Gilboa and Schmeidler (2001), framing effects can be integrated in "rational choice". They make the following definition: "An act or a sequence of actions is rational for a decision maker if, when the decision maker is confronted with an analysis of the decisions involved, but with no additional information, she does not regret her choices" and conclude: "A decision maker who, despite all our preaching, insists on making frame-dependent choices, will have to be admitted into the hall of rationality" Gilboa and Schmeidler (2001, 17f. and 19). Kleindorfer *et al.* (1993) offer the explanation that some alternatives are eliminated from the possibilities because they are socially or morally unacceptable. The decision to choose a treatment where two thirds of the sick will surely die, instead of a treatment where all may live with probability one third (and die with probability two thirds), is an example for a morally unacceptable alternative (Rubinstein (1998, 13)).

The following subsections deal with different variations of the traditional approach that try to explain deviations discussed above.

8.2.2 Behavioural Game Theory

Camerer (2003) discusses several approaches to include pro-social behaviour in preferences or utility functions of players. Based on experimental evidence, expected utility maximization is extended by making the utility of

	II				**II**	
	C	D			C	D
C	(3, 3)	(−2.5, 2.5)		C	(3, 3)	(−2.5, 4.5)
D	(2.5, −2.5)	(1, 1)		D	(2.5, −5)	(1, 1)

Figure 8.4: Payoff matrices in Prisoner's Dilemma games with altered utilities, according to the theory by Fehr and Schmidt (1999) with $\alpha = \beta = 0.5$ for both players on the left hand side and $\alpha_1 = \beta_1 = 0.5$ for the first player and $\alpha_2 = 1$ and $\beta_2 = 0.1$ for the second player on the right hand side.

one player depend on utility differences and/or the utility total of all players. As an example, consider incorporation of fairness and envy aspects by Fehr and Schmidt (1999). The utility function of each player is altered either by an envy term, if the other players receive more, or by a fairness term, if the other players receive less. In a simple two player game the equation is:

$$U_i(x) = x_i - \alpha_i max(x_j - x_i, 0) - \beta_i max(x_i - x_j, 0) \qquad (8.1)$$

where x_i and x_j are payoffs of the two players and U_i the utility of player i. It is assumed that $\beta_i \leq \alpha_i$ and $0 \leq \beta_i < 1$. Consider α and β to be 0.5 for both players. Then the Prisoner's Dilemma in Figure 8.1 changes to the game illustrated on the left hand side of Figure 8.4. Now assume the second player to have the attributes $\alpha_2 = 1$ and $\beta_2 = 0.1$. Then the symmetric Prisoner's Dilemma becomes the asymmetric game on the right hand side of Figure 8.4.

Behavioural game theory stresses the role of other players' outcomes in utility functions of decision makers, mainly in form of payoff differences and total payoff. It can explain behaviour in social dilemmas, Dictator and Centipede games. Furthermore, it can represent individual differences among players. However, it abides by individual utility maximization as the primary ground on which reasoning takes place. On the other hand, people's attributes have to be considered private information, and so the assumption of common knowledge has to be relaxed.

8.2.3 Psychological Game Theory

Colman (2003) proposes psychological game theory as an alternative approach to behavioural game theory. It focusses on non-standard reasoning processes instead of payoff transformations. Psychological game theory brings the role of a focal point and joint maximum or social optimum into the discussion on decision making. Two of these reasoning processes are discussed here, team reasoning and Stackelberg reasoning.

Team Reasoning

Team reasoning is based on the existence of collective preferences, which differs conceptionally from individual preferences to maximize joint outcome, but may lead to the same increase in group utility. "A team-reasoning player maximizes the objective function of the *set* of players by identifying a profile of strategies that maximizes their *joint* or *collective* payoff, and then, if the maximizing profile is unique, playing the individual strategy that forms a component of it." Colman (2003, 150, emphasis in original). Thus, the existence of a focal point is a prerequisite for team reasoning to come to a successful conclusion. If an agent can assume the other agents to be drawn to the focal point for whatever reasons, it can also choose the appropriate strategy. If this is unclear, team reasoning is replaced by another form of reasoning. This theory can explain cooperation in social dilemmas and centipede games, as well as focal point selection in the high-low matching game.

Stackelberg reasoning

Stackelberg reasoning (Colman (2003)) is strategy choice under the condition that the other player anticipates my choice and reacts on any strategy I choose with his or her best reply. This works only, if the Stackelberg strategies for both players lead to a Nash equilibrium as in the Prisoner's Dilemma and in the High-Low-Matching game. It fails to explain the incidents of cooperation in Prisoner's Dilemma games, but it explains the equilibrium choice of the higher equilibrium in the High-Low-Matching game.

8.2.4 Case-Based Reasoning

Case-Based Decision Theory (Gilboa and Schmeidler (2001)) is an extension to SEU maximization theory, based on the assumption that in many cases people reason by analogies. The main idea is that our memory stores cases,

consisting of a problem, an act, and a result. When we face a problem, our memory comes up with similar problems, the acts chosen, and results obtained. We can then learn from experience by taking an action which led to satisfactory results in the past or refraining from actions leading to unsatisfactory results. This can be done by maximizing a utility function assigned to previous outcomes, but there are simpler methods conceivable.

Case-based reasoning is a learning method as well as a decision method, since we can reason choice of an act on basis of experiences and improve our choices by more experiences. As a learning method it can be combined with different other forms of decision making, for example with the use of heuristics. Satisficing and aspiration adaptation theory (see below) are easily incorporated in case-based reasoning.

In addition to the notion of constrained information processing and memory capacities, this theory sheds light on how a link between actions and outcomes may be represented and modeled.

Case-based reasoning can explain cooperation in social dilemmas, gift giving, and focal point selection. However, as a method it does not define actual decision making, because maximization and simple heuristics can both be easily combined with case-based reasoning.

8.2.5   Integrative Decision Sciences

Kleindorfer *et al.* (1993) integrated multiple aspects from a variety of decision sciences into one conceptual approach. According to their theory four processes are important and interdependent when making a deliberate decision (Kleindorfer *et al.* (1993, 8ff.)).

Problem context: The problem context consists of social context, institutional constraints, and available information. For a decision maker it is important to know who has power and who may be affected by this decision.

Problem finding: The problem first has to be identified and accepted as a problem to be solved. It may also be represented in different ways. "Problem finding involves tracing the perceived source of the problem to the needs, values, and beliefs that the decision maker brings to bear in defining the problem as a decision or choice opportunity." (Kleindorfer *et al.* (1993, 10))

Problem solving: This is making a choice or taking an action to solve a problem according to a decision maker's values and beliefs. This

includes searching for alternatives, evaluating these, and choosing one that seems to resolve the problem.

Legitimation Process: Finally, a decision choice has to be legitimated on the basis of stakeholders' interests and concerns, as well as expected impact this decision has on them.

These processes are interrelated so that for instance the legitimation is of concern during process solving and may lead to rejection of alternatives which can not be legitimated.

This comprising theory brings many new aspects into the discussion, namely social context (relations to other players), institutions, information, a decision maker's needs, values, and beliefs, as well as evaluation and choice of alternatives, and need to legitimate a choice. Within this theory, the counter examples listed above can be explained. However, it is a conceptual approach that does not define decision making mechanisms closely enough to be represented in models. For this abstraction, further work is necessary, depending on the decision task. An important contribution of this integration is a comprising list of relevant features in a decision environment, as well as a combination of these features in a conceptual framework.

8.2.6 Bounded Rationality

For a discussion of the concept of bounded rationality the reader is referred to the Introduction (Chapter 1).

8.2.7 Institutional Analysis

Ostrom (Ostrom *et al.* (1994); Ostrom (2004)) proposes a framework for institutional analysis. In this framework, an action arena, which is the decision environment for actors, consists of participants, positions of participants, possible actions and choices of these actions, information, links of actions to potential outcomes, outcomes, and costs and benefits assigned to outcomes. Actors have preferences, information-processing capabilities, selection criteria to choose among possible actions, and resources. For an institutional analysis an action arena is viewed as embedded in an institutional process. It is influenced by attributes of the physical world, attributes of the community, and rules-in-use. Resulting actions in an action arena form a pattern of interaction which determines outcomes. These in turn influence the action arena, as well as the physical world, the community, and rules-in-use. In

order to judge the performance of an institution, outcomes are evaluated
due to evaluative criteria.

As the integrative approach of Kleindorfer *et al.* (1993), Ostrom's insti-
tutional analysis provides an extensive list of relevant features of a decision
environment. Since the institutional analysis has a focus on social dilemmas,
the provision, maintenance, and use of common-pool resources or common
goods, Ostrom is able to develop her theory on a higher level of abstraction,
rendering it more useful for modelling purposes.

8.2.8   Summary

Table 8.1 summarizes all aspects encountered above and their appearances
and role in various decision theories. These aspects are now discussed with
respect to their potential influence on the choice of relevant attributes and
applicable heuristics.

Agents' objectives

In traditional approaches based on game theory and expected utility max-
imization, an agent's sole objective is maximization of its utility function.
This can be extended to include other agents' utilities, aspects of utility
differences, or a social optimum. In bounded rationality the objective is
expressed in terms of an aspiration level. In broader approaches, the objec-
tives or preferences of decision makers have to be defined with respect to a
decision problem. They can include such diverse aspects like avoidance of
computational effort, ease of choice legitimation, fairness aspects, and norm
or rule compliance. It is not always easy to determine which goals decision
makers follow. But the actual problem to solve depends on this goal. If a
goal is to avoid computational effort, a choice in for example the centipede
game might be very different than if the goal is to maximize joint outcome.
Different agents' objectives, therefore, can lead to different decision problems
even in the same decision environment.

Agents' resources

While resources of agents are either unlimited or incorporated into the choice
problem in traditional approaches, limitations do play a role in real situa-
tions. Limitations can refer to computational capacity, memory, time, and,
of course, money of decision makers. Some decision tasks need to be altered
in order to incorporate these constraints.

	Behavioural Game Theory	Psychological Game Theory	Case-Based Reasoning	Decision Sciences	Bounded Rationality	Institutional Analysis
agent's objectives	SEU maximization	single or group maximization	SEU maximization	need to be defined	aspiration level	as preferences
resources	unlimited	unlimited	possibly constrained memory	limited perception and attention	possible limitations	part of actors
other agents	-	team reasoning	-	social context, stakeholders	for social learning	participants of action arena
positions	-	-	-	social context	-	part of action arena
possible actions	strategy realm	strategy realm	-	search for alternatives	search heuristics	part of action arena
selection criteria	-	-	-	values and legitimation	aspiration adaptation	part of actors
existence of focal point	-	team or Stackelberg reasoning	-	-	part of heuristics	
information	complete	complete	memory based	part of problem context	structure of environment	part of action arena
information processing	no limits	no limits	memory based	possibly limited	part of heuristics	part of actors
links of actions to outcomes	payoff matrix	payoff matrix	memory based	problem representation	learned heuristics	part of action arena
outcomes	payoff matrix	payoff matrix	uncertainty	beliefs	learned heuristics	part of action arena
costs and benefits	same as outcomes	same as outcomes	same as outcomes	legitimation process	-	part of action arena

Table 8.1: Summary of possible aspects of decision environments and their role in different theoretical approaches.

Social context

Relationships to other humans make up the social context of a decision, whether they are stakeholders that are effected by a decision, like dictator receivers, or other actors, whose decisions influence possible actions and outcomes, like ultimatum responders. Even, if they are only observers, as experimenters, instead of participants, they can influence a decision maker. This can be either because their utilities are taken into account, their norms have to be considered, or they are used for social learning processes. Mainly, however, their actions will have to be taken into account, because they influence links of outcomes to actions for a decision maker. Therefore, predicting others' behavior becomes an important aspect of decision making.

However, social context is more, than others being influenced by or influencing outcomes. Social context can influence in itself the appropriateness of heuristics. Within a family calculations of costs and benefits of pro-social behavior is simply inappropriate. In experiments with classmates the social context is different than in experiments with real strangers. In field studies, the social context is both more diverse and more subtle.

Agents' positions

Agents' positions are a concretion of social context with respect to the power relationship of participating agents. Sometimes the power relationship is symmetric, as in the Prisoner's Dilemma. Sometimes it is clearly defined, as in dictator games, and sometimes it is rather vague, as in ultimatum games: Who has more power, the one to propose the ultimatum or the one who can accept or reject?

Possible actions

The range of possible actions is fixed in most decision environments considered in experimental economics. However, in field studies, it is more diverse and open to creativity. In the models considered here, no creativity has been implemented, but in real situations humans can be creative and alter rules of the games they play. Search for alternatives falls into this category.

Selection criteria

From among the pool of possible actions, a decision maker makes choices according to some selection criteria. These need not only be subject to goals

of a decision maker. A selection criterion apart from any other goals is, for example, the possibility to legitimize an act.

Existence of focal point

Decision makers tend to orient on focal points, even in strategic games like the ultimatum game (Hennig-Schmidt (1999)). Therefore, agents should be able to capture this aspect of human decision making. Focal point selection serves multiple purposes, like coordination (heads-tails), reducing computational effort, and facilitating legitimation. In decision realms with one or more clear focal points, decision heuristics can refer to these.

Information

Information of decision makers has, of course, an effect on the representation of a decision problem. A decision maker who does not know other participants' preferences faces a decision under uncertainty, even when the payoff matrix is common knowledge. In case-based reasoning information is stored as cases. In bounded rationality the focus lies on the match of informational structures of an environment with decision heuristics. If information is scarce and variability is great, cue-based heuristics tend to outperform calculative strategies.

Information processing

Not only information can be limited, but also information processing capacity. This is combined with the problem of limited resources. However, it can also be seen as a more procedural problem. What does a decision maker do with an additional piece of information? If it is a cue strongly related to a certain heuristic, it may be considered important enough to trigger a new deliberation process. However, many pieces of information may be ignored, even though they would offer new insight.

Links of actions to outcomes

Links of actions to outcomes may be clear, as in payoff matrices, or rather obscure as in some real world problems. Even in Prisoner's Dilemmas, two possible outcomes are linked to one action, so there is no unambiguous link. This complicates learning. In cognitively rich settings learning of these links happens, but not necessarily in a correct way.

Outcomes

Outcomes are part of the feedback needed in order to learn whether or not an action was appropriate or successful. In most cases, outcomes are monetary outcomes. Additional effects of actions can be captured in the final aspect, the costs and benefits.

Costs and benefits

In addition to pecuniary values, outcomes can be assigned additional costs and benefits. The "warm glow" (Andreoni (1989)) of a nice act may be such an additional benefit. Deterioration of others' willingness to cooperate may be a cost to free-riding.

8.3   Modeling practice

8.3.1   Agent-based modeling

In agent-based modeling there are several well-known techniques for modelling decision making by agents. Before we turn to BDI (Belief-Desire-Intention) as an example, some general aspects of agent-based models are recapitulated.

   Agents have resources, objectives, and knowledge. They perceive their environment, possibly incompletely, they have representations of other agents and objects in their environment, and they communicate with other agents and act to alter objects (Ferber (1999)). This is all there usually is. Thus, in order to model decision environments, the key aspects will have to be reduced to these few categories: other agents (and their relationship structure), representation of objects in the environment, resources, objectives, and knowledge.

Belief Desire Intention

Belief-desire-intention (BDI) architectures base decision making of agents on practical reasoning. First, it is decided what goals to try to achieve (deliberation), then how these goals are to be achieved (means-ends-reasoning). Beliefs about the environment and problem at hand are used to generate options for behaviour. Desires are goals that an agent wants to achieve in general. They are filtered by beliefs about realisability and intentions, which are long-range goals. Only if a desire appears to be achievable and fits to

previous intentions, it is adopted as a new intention. According to a set of intentions, an action is chosen that is believed to further an agent's goal.

This intuitively appealing architecture contains some elements of decision environments identified in the previous section. Objectives are incorporated as desires and intentions. Information processing capacity of an agent exists in form of beliefs about the environment and results of actions. Previous intentions are selection criteria for different possible actions. Other agents exist only as part of the environment. However, since nothing is said about learning methods, a kind of social learning can be implemented. Information is incorporated in a model as generating new beliefs, which includes beliefs about the links of actions to outcomes. Nothing is said about agent resources, positions of agents in a network, costs and benefits of actions, as well as focal point selection.

Several architectures like BDI have been developed for artificial intelligence and collaborative problem solving. Optimization under constraints forms the basis for BDI and other architectures that try to make artificial agents capable of thoughtful and intelligent behaviour. However, since agents we are interested in behave more impulsively and less calculating than a BDI-agent would, the implementation differs. We focus on agents that have characteristic traits and choose from a pool of heuristic one that promises to satisfy their aspirations. How this choice happens depends on the environment.

We now inspect the problem from the side of the applications.

8.3.2 Public Goods Games

In the altruistic punishment model (JASSS paper, Chapter 2, and Handbook paper, Chapter 5), based on Fehr and Gächter (2002), agents' behaviour was based on expected behaviour of others. The overall objective was to earn as much money as possible, at least in the investment decision. Different agents tried to achieve this goal in three different ways. The first is to invest nothing or only just enough to avoid punishment. The second way is to try to achieve group optimum by investing the maximum or close to the maximum. The third is to invest close to the (expected) average investment, in order to avoid both exploitation of others and being exploited. One can identify different secondary goals with these three types, namely a desire not to be exploited, a desire to initialize cooperation, and a desire to behave similar to others.

Several well documented framing effects have not been incorporated in the model, among these are the effect of the return function and group size

(Isaac *et al.* (1984)). This can be done, by changing beliefs about others' behaviour according to these variables. When investment becomes more profitable, expectations in others' investment decisions are raised. Due to a dependency of the third heuristic on expectations, this also raises actual investment. However, maybe the percentage of agents using cooperative heuristic also increases. Similarly, expectations decrease with increasing group size.

The punishment decision involves emotions. Punishing acts can be explained without emotions, based on a belief that the fee pays due to higher future investments by others. In this experiment by Fehr and Gächter it was intended to eliminate this kind of reasoning by using a "strangers" treatment. However, it may be assumed that this has only partly the desired effect. In any case, subjects indicated involvement of emotions in punishment decisions in questionnaires filled in after the experiment. Annoyance about defectors is compensated by retaliatory acts, even if they are costly and do not promise future returns.

A lesson learned from this model is, that strong emotions have to be taken into account in "reasoning" made by decision makers. Except for bounded rationality, none of the approaches described above take emotions into account explicitly. One of the major problems involved is that they are not easy to observe and measure, and another is a great variance among subjects.

Due to the design of this experiment as 12 consecutive games rather than one-shot games, it provided an opportunity to implement learning mechanisms in the corresponding model. In this case, only a simple learning direction mechanism was implemented.

Expected behavior or rather expected attributes are an important aspect for decision making in public goods experiments. Emotions play a role mainly for punishment decisions. Because many experiments with public goods decision environments have been conducted, there are a number of variables (group size, strangers/partners, return function) which have been changed in order to obtain their influence. For a modeling process with heuristics, however, reasons why cooperation depends on these variables, are important.

The implemented heuristics are based on individual data. Data analysis is not a formalized unambiguous process. Therefore, the need to qualitatively validate resulting heuristics by domain experts or questionnaires remains.

As indicated earlier in the discussion on Prisoner's Dilemma games, agents facing a public goods decision situation have to ask themselves what

outcome they would prefer, a mutual cooperation with a risk of free-riders or that they can free-ride on others' attempt to cooperate. This depends on their self-perception. In public goods decision environments with more possible choices than cooperate or defect, like the one implemented here, willingness to cooperate can be indicated by a medium decision. Giving a quarter or half of the possible amount indicates this willingness and a loss would be acceptable. Then, they have to ascertain the risk of free-riding. This depends on the stakes at hand and number of others that would also have to cooperate. A possible heterogeneity among the agents increases risk of free-riding. Communication or group identity lessens this risk. Similarly, framing of a decision task can make cooperation or defection seem more or less appropriate.

8.3.3   Trust Games

The trust game experiment by Fehr and Rockenbach (2003) offered some back up for prominence of emotions in decision making, and a dependency of decisions on expectations of others' behavior (Behaviour paper, Chapter 7). This experiment directed the choice to model trust as expected positive reciprocity. Later this was extended to be expected cooperativeness, expected positive reciprocity, expected conformity, or expected fairness concerning others, depending on the situation (see subsection 8.3.5). Due to a difference in positions, a new feature can be introduced, namely the "sense of control". Although it was made clear to investors that a fine did not enhance the chance of return gifts, on the contrary it seemed to decrease that chance, investors used the fine more often than not. This can be explained by a desire to keep control over events. Although this fine is an inappropriate tool, it was used. Thus, power relationships between subjects enters deliberations.

New features in this model are different agent positions and resulting sense of control.

Again, first movers in trust games first act on their self-perception that can be more or less trusting, that is their expected positive reciprocity has a certain value. Agents have to ask themselves, whether or not they are willing to risk receiving a return gift less than the investment. Risk, in turn, depends again on framing and stakes. Second movers also act on their self-perception. However, since they are second movers, their decisions also depend on first movers decisions. Even a very trustworthy, positively reciprocal, and fair agent may be annoyed at being trusted very little, and then act on its negative reciprocity instead. An agent that would return the equal split

when trusted with the maximum, may decide to return nothing when trusted with only 2 money units. Other agents may not have these resentments. In reactions perceived attitude enters deliberations and complicates matters.

8.3.4   Dictator and Ultimatum Games

The models of dictator and ultimatum games were supposed to shed light on differences between impulsive and strategic behavior (ESSA paper, Chapter 3, and Behaviour paper, Chapter 7). Dictator decisions do not have any consequences, therefore it is assumed that dictator choices depend solely on the fairness (concerning others) attribute of the decision maker. Ultimatum proposers, on the other hand, make a completely different decision based on expectation of responders rejection behavior. Maybe it can be assumed that fair (50:50) dictators make fair ultimatum proposers, but I am not aware of data on this aspect. It is possible that fair dictators are harsh ultimatum proposers.

This is why dictator games are perhaps *not* good baseline games to judge whether ultimatum proposers give fair shares because of fairness aspects or because of strategic thinking. If a situation suggests fairness and an agent's position is superior (as in Dictator games) fairness is the attribute on which this choice depends and a fair choice may or may not be chosen. If a situation demands strategic thinking because its position is not unambiguously superior, a choice will be made on grounds of expectations, and fairness does not necessarily have a part in making this choice.

A second role of these models is integration and comparison of several different modes of decision making. Two theories from behavioral economics have been implemented in addition to the heuristics approach, in order to compare the resulting micro behavior. Different micro behavior can lead to similar macro behavior. Therefore, other validation methods than comparative statistics have to be employed in order to qualitatively validate underlying assumptions. This model demonstrated this aspect clearly, although this was not intended in the beginning.

In addition to back up for influence of agents' positions on attributes and heuristics used, the main new aspect is potential influence of subjects self-perception. If a player has a self-perception of appreciating fairness, he or she will act on behalf of that self-perception. In modelling practice, it may not be necessary to distinguish actual fairness concerning others and perceived self-perception of being fair with respect to others. But it may be necessary to make this distinction theoretically.

8.3.5 Appropriation Games

Importance of a focal point in decision making with heuristics became clear only in the latest of the models, the appropriation games which lack a focal point (MABS paper, Chapter 4, and Handbook paper, Chapter 5). Since the return function is negatively quadratic with a not too easily accessible maximum of 4 *and a half*, players have a hard time finding an optimal solution, even when they can discuss a joint strategy.

Heuristics implemented here, therefore, differ conceptionally from heuristics in previous models. Each possible integer value was modeled as one heuristic, representing the difficulty subjects' had choosing among them. Since this experiment consisted again of a sequence of games, a learning mechanism has been implemented. The choice of heuristics was implemented as a heuristic of the type: invest more, when return per token was higher in market 2 than in market 1. These heuristics changed the actual investment level, but always with respect to the previous value.

Furthermore, in this model, communication was introduced as a two-stage process, the first stage is coordination on a joint strategy, the other is better judgment of other players' cooperativeness and conformity. So far, actual conformity was not altered in a communication process, although there seems to be evidence that commitment is generated or increased by face-to-face communication. This is another model of trust. Trust in this case means expected compliance to the agreed upon rules and therefore expected conformity.

In this decision environment agents do not primarily act on grounds of their self-perception. Rather, they feel their way to an acceptable contribution based on their previous choice and experiences made. However, self-perception is used similarly to public goods environments. An agent may want to invest nothing in market 2 and leave all possible profit to others or it may want to free-ride on others, investing a lot, risking negative returns from market 2. In this environment, free-riding is not risk-free. As before, number of subjects, stakes, and group identity through communication change risk and risk assessment of free-riding.

8.3.6 Summary

Important aspects developed in this discussion of modeling processes are depicted in Figures 8.5 and 8.6 and discussed below.
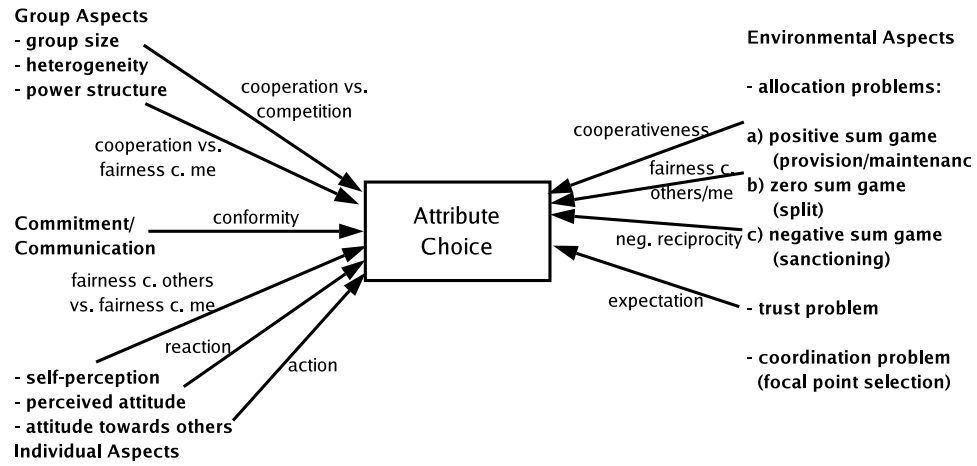
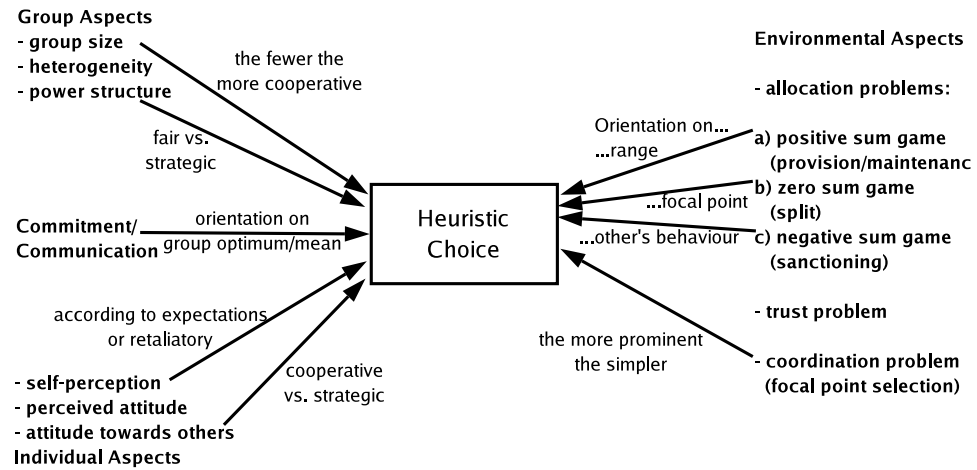*Figure 8.5:* Elements of a decision context influence attribute choice.

Figure 8.6: Elements of a decision context influence choice of heuristics.

Group size

The greater the size of a group, the less important cooperativeness becomes. This is due to several effects. First, contribution of a single individual becomes less important. Second, the probability of free-riding increases. Third, identification with the group is easier when it is small. This effect is reflected in a voluntary contribution experiment varying group size (Isaac *et al.* (1984)), as well as in the theoretical approach by Ostrom (2004).

The same reasoning applies for the choice of heuristics. The greater the group, the higher the probability of choosing a free-riding heuristic, simply because of greater expectations that others will free-ride.

Heterogeneity

A similar effect is achieved by group heterogeneity. The more homogeneous a group, the more individuals will identify with it, and the fewer will expect others to free-ride or free-ride themselves.

Power structure

This is linked to the power structure of a game. However, different power is also reflected in different positions and not (only) in different endowments or goals. Different positions can move a game from a cooperative game to a competitive game. However, it is important to note, that players with greater power can allow themselves to be nicer to players which can not challenge their power, than to players which exert only slightly less power, reminiscent of the pecking order of hens. Choice of attributes can be affected by a players' power, cooperative to players with equal power and with very little power and competitive to players with slightly less power. It can also have a similar effect on a choice of heuristics.

Commitment

Group identification can be improved by communication or a common goal. In these situations choice of the attribute conformity and a choice of a mean oriented heuristic can reflect orientation on a common goal.

Self-perception

Player's self-perception is important in moral situations. This can be influenced by decision framing. A test for fairness will trigger choice of the

fairness concerning others attribute. An indication that a situation is competitive will more likely trigger a player's fairness concerning me.

Perceived attitudes

When a decision is a reaction on another player's move, perceived intentions become important. If intentions are perceived as friendly, friendly attributes and heuristics are more likely to be chosen, as is behaviour according to first mover's expectations. If, on the other hand, intentions are perceived as hostile, behaviour will shift to retaliatory acts.

Attitudes towards others

First movers, who do not respond to a previous act, can act on their attitude towards others. This can make a difference whether cooperation is initialized or not.

Positive sum games

Environmental aspects are first of all distinguished by competitiveness of a payoff structure. Positive sum games induce cooperativeness. Heuristics refer to a range of possible decisions. That is, the more cooperative a player, the higher his or her contribution to a provision or the lesser his or her extraction from a common pool resource.

Zero sum games

If a fixed amount is to be split, the distinction which attribute to choose is most likely between fairness concerning me and fairness concerning others. A heuristic can be oriented on the focal point of an equal split or a split according to the power structure or game positions.

Negative sum games

In a negative sum game, negative reciprocity is used, when a player has a chance to sanction another player's previous move. The heuristics will focus on the other player's behaviour. That is, the greater the perceived hostility of a previous act, the greater the sanction.

Trust problems

Trust problems can employ different expectations of others behaviour. The difference is in what players have to trust. For example, trust can be that others will not free-ride (expected cooperativeness), will stick to agreements (expected conformity), will behave fairly (expected fairness concerning others), or will reciprocate nice acts with nice acts (expected positive reciprocity).

Coordination problems

In coordination problems the existence of a focal point is important. The more prominent a focal point, the more often it is chosen.

8.4 Discussion

This is only a very rough and incomplete classification scheme.

In order to propose a classification system that comprises more decision situations than already implemented, the requirements from the modeling viewpoint will have to be merged with the theoretical approaches. Difficulties expected to arise lie mainly in the huge variability of decision situations in the real world and different applicability of theories.

As modelling is an iterative process, Figures 8.5 and 8.6 can be expanded with more modeling examples. It is to be expected that in some instances, generalizations can be made, and in other instances situations have to be split up even more.

As a next modelling iteration player agents need to be endowed with perception possibilities for different influencing factors and the environments and games need to get the corresponding attributes. Game structure and group size can already be accessed by players, but not a degree of anonymity or focal point. So far, knowledge of these aspects has been provided by the experimenter objects.

Simon (1997, 273) describes the necessary research in behavioural economics as "a long an arduous path, but no longer and more arduous than the paths that scientists are following with great success in such domains as molecular biology or geophysics or any of the other domains where the ratio of validated facts to theoretical generalizations must, by the nature of things, be large." There will be no master equation that explains everything. There will be many facts and generalization will come slowly.

Chapter 9

Conclusion

Eva Ebenhöh

9.1 Results

Starting point for this thesis is the insight that some assumptions on rational behaviour in social dilemmas are not supported by observations. Data from observation are therefore a basis of the models developed here. From this data we extract behavioural regularities that are not only backed up by questionnaires and expert knowledge of economic experimenters, but also by theoretical approaches, like trait approaches and bounded rationality. These regularities form a basis for micro behaviour in our models. At the same time, individual and aggregate data and in some cases even the variance are reproduced in these models. In this sense, our models are successful (see Section 9.2).

Many aspects of this modelling framework were derived from data, like dependency of cooperative behaviour on the existence of sanctioning mechanisms, and dependency of (un)fair behaviour on anonymity. These insights, therefore, can not be claimed as original results of this thesis.

This thesis' result is the abstraction in form of the framework developed from different pieces of data collected from various experimental sources and literature studies. These pieces have been puzzled into a coherent structure, which has holes and fuzzy borders, but is nevertheless a theoretical approach to combine experimental findings and draw meaningful conclusions from them for real-world situations. It shows that a behavioural theory can be developed, although this process has to be done in small steps, with context dependency, and cannot be made without ambiguities.

Data from experimental economics has been useful in establishing this modelling framework. Mostly, these data have been surveyed to attack or defend the traditional economic model of decision making. We use these data to create a new concept.

In this thesis, it has been shown that it is possible to use these data to create agent-based models and draw meaningful conclusions from them. These conclusions support findings from data analyses. By analysing influencing factors across different experiments, findings have been obtained that go beyond those of data analysis. A major result is the attribute topology (WCSS paper, Chapter 6) and its implications on elements of a decision environment and group characteristics. Only by condensing influencing factors to a small number, stringently defined, it became possible to categorise environmental and group aspects and their connections to attributes.

Other, particular findings include:

- In the altruistic punishment experiment (JASSS paper, Chapter 2, and Handbook paper, Chapter 5), the correlation between height of cooperation and height of punishment is not explained by a correlation between the corresponding attributes, but by an underlying process and the dependency of height of punishment on height of deviation. The height of deviations could only be great, if cooperation was high in the first place. This creates a functional dependency, instead of a statistical correlation. It can not be deduced that cooperators are statistically more willing to punish uncooperative acts.

- This model was also used to investigate different degrees of conviction for one heuristic or another. If a cooperativeness is very high or very low, it is harder to make the agent change its behaviour.

- In a comparison of dictator and ultimatum games (ESSA paper, Chapter 3) it was deduced that the only similarity between these two is the focal point of one half. Only truly fair dictators do not change their decision making process. A conclusion is that in strategic environments with little difference in power fairness aspects are neglectable. They are important in relationships with great differences in power.

- Cooperation seems to be easier, when there is a focal point salient for all participants. This became apparent in the appropriation model (MABS paper, Chapter 4, and Handbook paper, Chapter 5) in which a focal point is missing.

- Also fair behaviour is more frequent when a salient focal point is present, as shown in the comparison of dictator games and investor-trustee games (Behaviour paper, Chapter 7).

- In that model, similar to the altruistic punishment model, different degrees of conviction for a chosen heuristic were modelled depending on the value of an attribute. Very fair or very unfair agents are slower in changing their behaviour in response to a previous behaviour of another.

  An assumption underlying these models is that agents (and people) scoring high or low on certain traits are less likely to conform to group norms in decisions depending on this attribute, than agents (and people) having a medium score.

Extracted behavioural regularities are not certainties, but show differences in behaviour both from one person to another and from one setting to another. These differences are reproduced in concept and models. Our approach is a pluralistic approach, because it allows different categories of behaviour. In analysing the potential for cooperation in a group it is important to keep these potential differences in mind and deal with them. This modelling framework allows for this kind of research. Group composition can be explicitly set and changed.

In addition, this framework allows comparison between different situations. This is a prerequisite for investigating influences of institutions on cooperative behaviour in a group. Institutions can be modelled in this framework as having a different impact on stabilisation of cooperation, depending on the acceptance among group members or many other variables. This has consequences for management of natural resources building on cooperation or trust processes in groups of stakeholders. With this concept, we can model path dependencies of the introduction of institutions and interrelations of institutional settings (see Section 9.3).

One very important aspect is still ongoing work. This regards the agents' ability to recognise characteristics of their decision environments. The basis for this has been formed in the Classification working paper (Chapter 8). In the models so far, cues are given to agents by the experimenter and not found by agents themselves.

This pluralistic approach has drawbacks. There will always remain ambiguities in the interpretation of situations and actions. What seems to be a cooperative situation to one person may appear competitive to another. An intentionally nice act may be misinterpreted as being hostile. The modeller

never knows. These models have to be validated in a qualitative way by domain experts.

9.2   Validation

The results of these models are validated in more than one sense. Statistical comparison of model results to experiment results is possible, and has in some cases been made. However, a problem in most models is, that these data have been used to calibrate the model, and therefore a statistical comparison is not very meaningful. Furthermore, data bases with similar settings are rather small, because the research focus lies on explaining differences.

Statistical tests have been reported in the ESSA paper (Chapter 3 and the Handbook paper (Chapter 5). In all models, individual data have been classified in categories, and the distribution of different types of behaviour has been reproduced.

A statistical analysis and comparison of model results and experimental data can be very useful. Especially, if an investigation does not only encompass means and ranges, but also the variance of data sets. If a model reproduces the variability of experimental data, it can be considered as quite good. This has been achieved with the altruistic punishment model, but not so well with the appropriation model (see Chapter 5).

Since these models are not stand-alone models, but part of a framework, a certain kind of cross-validation is already possible. Some processes that work in one environment do not necessarily also work in another environment with many similarities but some differences. Looking at differences in settings and what differences in a model they entail can be seen as a reasonability test. This has been done in a comparison of dictator and ultimatum games (ESSA paper, Chapter 3), and also between one-sided gift giving (dictator) and two-sided gift giving (Behaviour paper, Chapter 7).

However, in order to validate models, micro behaviour of agents would have to be compared to micro behaviour of human subjects not only in a statistical but also in a qualitative way. Moss and Edmonds (2005) propose to incorporate expert knowledge into later models, letting experts criticise earlier models. Experts in this case can be scientists who conducted these experiments. This has not been done with models in this thesis, although papers have been sent to experimenters asking for comments. However, where available, interpretations of subjects' decisions by experimenters as published have been integrated in models. Again, there is the problem that information used to calibrate a model, can not be used to validate it. Also,

experimenters tend to explain aggregate behaviour rather than individual differences.

Of course, subjects of experiments can also be seen as experts for their own decision making. Then, post-experimental questionnaires asking about underlying heuristics give a good indication of micro-processes. Where available, this information has been used to calibrate models. However, the issue of questionnaires is not trivial. "The very act of jointly asking subjects what outcomes they prefer and why seems, not surprisingly, to generate more cooperative orientation." ((McClintock, 1972b, 281)). Therefore, the main source of information was individual data.

Nevertheless, it would be a valuable endeavour to ask subjects of these experiments whether they judge heuristics used in these models to be a good representation of their and others' behaviour in the experiment. This has not been done, because these experiments have not been conducted for this thesis. Experimental data are used, that have been collected in various research programs (see Section 9.3).

When models are used for case studies, participants of case studies can be seen as experts, and their knowledge, which is also used to build a model, can be used for the validation process. In this sense, a model can be built and made better in an iterative process. This validation process is not unusual for modelling processes in general and agent-based modelling in particular (cf. Pahl-Wostl (2002)): (1) a model receives input from domain experts, social entities modelled, and further available data; (2) from this input agents' behaviour and other model behaviour is generated; (3) model results are then fed back into the discussion, so that domain experts are able to test plausibility and stakeholders get data for further discussion that are more transparent and less emotional than before; and (4) feedback is used to iteratively increase these models' realism and validity.

This, of course, is already done in some case studies. An advantage we expect from this approach is a contribution to a conceptual framework of behaviour in social dilemmas and in group processes concerned with regulation of common-pool resources. We expect this conceptual framework to develop from further different sources of data, comprising different scenarios and levels of decision making. The act of breaking down complex behaviour into simple heuristics and determining dependencies on aspects of the social and institutional environment is promising to result in a data base that can be used to model human behaviour. It is not to be expected that this data base will ever be complete and unambiguous.

9.3   Prospects

This work was driven not only by the question of how individual decision making behaviour could be modelled adequately. It was also to be connected to institutions framing a decision. This link has been established only incompletely in the JASSS paper (Chapter 2, punishment as an institution) and the Handbook paper (Chapter 5, communication as an institution).

Our underlying assumption is that institutions based on a theory of self-interest or implicitly assuming humans to act on narrowly defined self-interest would promote self-interested behaviour and discourage other forms of behaviour. It is an interesting and important research question to trace this form of influence. Experiments by Fehr and Gächter (2002) and by Ostrom *et al.* (1994) point in that direction. Punishment and face-to-face communication as institutions promote cooperative behaviour especially when they are combined (Ostrom *et al.* (1994)). Lack of a social context in economic experiments can itself be seen as an institution to discourage cooperative behaviour. Different degrees of anonymity in dictator experiments underscore this effect, as demonstrated in the Behaviour paper (Chapter 7).

What seems to be quite clear in simplistic environments of economic experiments is much harder to trace in real-world situations. Questions regarding relationships of institutions and individual behaviour come from a wide range of topics, including upbringing of children and social welfare. Does punishment and rewards encourage or discourage prosocial behaviour in small children? Do they learn to cooperate *only if* faced with punishment or reward? Does compulsion to work encourage prosocial behaviour in recipients of social welfare? Or does it intensify incentives to cheat?

If this mutual influence exists it has implications on policy. In a common-pool resource situation, for example, an institution that assumes humans to behave as rational egoists may work to promote this kind of behaviour. Benefits from this policy may then be countered by losses due to an decrease in cooperation.

As an example consider a monitoring process conducted by people who are not trusted by the group whose actions are monitored. If this monitoring is not very effective, this may lead to even less compliance to a joint decision compared to no monitoring. Ostrom *et al.* (1994) find that effectiveness of monitoring is higher, when monitoring is done by group members compared to when it is done by "external" officials. Externally imposed sanctioning mechanisms reduce the importance of intrinsic cooperativeness which has an effect on situations in which no sanctions are imposed (Mulder *et al.* (2005)).

Regarding this mutual influence of humans on institutions and institu-

tions on humans (Held and Nutzinger (1999)), rational choice theory needs to be complemented by a theory which makes this link explicit. "The study of institutions has been traditionally a blind spot of rational choice theory, and there is a systematic reason for this. Rational choice accounts face an inherent difficulty when it comes to explaining *rule-following* or *norm-guided* behavior in the sense of a person's general compliance with the rule or norm in question" (Vanberg (2002, 38)). But the assumption of rational actors need not necessarily be dropped. "Standard economic analysis neglects the identity and past relations of individual transactors, but rational individuals know better, relying on their knowledge of these relations" (Granovetter (1985, 491)). What is rational then needs to be defined in a situation specific way. However, as Lovett (2006, 143) states, if we are interested in "sociopsychological effects" of policies in shaping beliefs, attitudes, dispositions, and so on" then these "processes do not primarily involve constrained purposeful action." Rational choice theory can be useful in causally explaining how different institutions directly effect human behavior, but not how they indirectly shape their preferences and dispositions, because this second effect has nothing to do with purposeful action. With these models, we create another tool for investigating this link.

To distinguish between intrinsic and extrinsic motivation of decision making humans, may be a purely analytical effort, because in real situations these two will always be combined. Experimental economics strive to make it possible to separately investigate the intrinsic part. Social context is kept as simple as possible. These experiments are, therefore, a benchmark for investigating behaviour in socially richer environments, encountered in case studies.

In Figure 9.1, which presents an extension of Figure 1.2 on page 23, social context and its possible influences on decision making are introduced. Social context encompasses the social environment of a decision, personal relationships and interactions of individual actors, a longer time scale and repetitions of interactions, and individual and joint framing of a decision context.

This work proposes to use models which allow for a variety of decision making mechanisms in order to support theory building. With this approach it is possible to model and investigate how different mechanisms perform in different institutions and consequently how institutions promote some mechanisms and hinder others.

This question of mutual influences of institutions and actual behaviour will have to be investigated more thoroughly in field studies. Agent-based modelling provides a perfect tool for this question, because the micro-macro
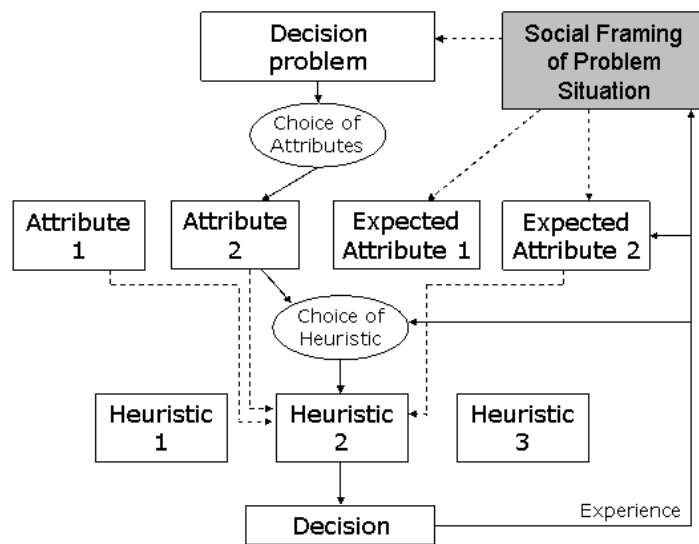
Figure 9.1: Framing of a decision context has an influence on expectations of agents and thus on decision making processes. In contrast to decision situations in experimental economics (cf. Fig. 1.2) where framing is mainly determined by experimental settings, framing in real world situation is influenced by social context such as prior experience or the nature of learning and negotiation processes in a group of actors. [configuration by Pahl-Wostl and Ebenhöh]

link can be explicitly modelled and investigated.

More insights are to be expected from using these models as parts of social learning processes or participatory model building. This usage of this framework is part of ongoing work. This framework has been built on data from experimental economics and will be enhanced by integrating case studies, so that very different data from field situations may also be merged into it. At the same time, this framework may allow a comparison of different field studies on this abstract level.

In the following, three different fields are introduced that can be examples for applications of this framework.

Role-playing games Role-playing games have been conducted in stakeholder groups to illustrate complex interdependencies of different actors (Barreteau *et al.* (2001); D'Aquino *et al.* (2003)). They constitute an intermediary state between laboratory experiments and field studies because the situation is much more controlled compared to field studies but interactions take place in a richer social setting compared with economic experiments. Therefore, we assume that this framework can be used to investigate and analyse behaviour in role-playing games. At the same time, agent-based models can be developed in an iterative process in combination with role-playing games conducted with stakeholders. This can lead to better-fitting heuristics for these models and support development of strategies by stakeholders taking part in these games (Pahl-Wostl and Hare (2004)).

Actors' platform An actors' platform can be used to analyse a common pool resource problem with differing frames and provides a forum for a problem solving process. Success of such platforms depends on the intensity of interactions, stakeholders' cooperativeness, and trust between participants (Panebianco and Pahl-Wostl (2004)). From experimental economics the importance of repeated face-to-face communication for these factors is known. Face-to-face communication not only allows to exchange information about strategies, but also assessing other participants' trustworthiness or their willingness to comply with joint agreements. This in turn allows making promises (Ostrom (2004, 33f.)).

Insights from experimental economics can help to inform decision makers who need to decide whether or not to implement an actors' platform and what to expect from it. This modelling framework can help to compare different settings and their influences on the success of an

actors' platform. It may also help to investigate crucial prerequisites for interactions like length and frequency of meetings, as well as trust in chair persons.

Participatory processes For assessment of effects of a new institution, it is not only important to know how this new institution changes opportunities and potential benefits and costs, it is also necessary to know about its implementation process (Ostrom *et al.* (1994); Ostrom (2005)). This path dependency in the effectiveness of an institution is largely ignored, but shown for effectiveness of sanctioning mechanisms, which are less effective, when they are imposed from outside than if decided upon by stakeholders themselves.

This modelling framework can help to investigate this link by making it possible to model abstracted effects of different institutions under different assumptions about willingness to comply. A comparative analysis of different, more or less participatory processes using this framework may help to find crucial variables for institutional effectiveness or ineffectiveness, level of participation being one important aspect.

Orcutt summarises nicely, what this modelling framework is intended to help to accomplish:

> "Models of the type suggested in this paper could facilitate and improve testing of hypotheses about elemental units by permitting testing of them at any level of aggregation. Such models also would improve the testing of such hypotheses by keeping the interrelated nature of the system in the consciousness of the investigator and by helping him satisfactorily to take it into account." Orcutt (1957, 122)

Appendix A

Implementation

Eva Ebenhöh

The purpose of this appendix is not to repeat the implementation, but to give an overview of the packages and classes in the framework. The CD provides the program code and JAVA documentation, as well as models that can be run and altered in the FAMOJA[1] environment.

There are five packages: `experiments`, `games`, `players`, `strategies`, and `charts`.

A.1   Experiments

The package `experiments` contains the abstract superclass `Experiment-Builder` and the class `Experimenter`. The builder is used to organise the creation and setup process of a certain experiment (cf. Gamma *et al.* (1994, 97ff.)). For new experiments a subclass of `ExperimentBuilder` has to be implemented, because here it is defined which games are played and which strategies players use (see Figure A.1). Also, in this class all parameters have to be set, which are to be changeable at the beginning of a model run. The coordination process during the experiment, however, can often be done by the standard `Experimenter`. That was the reason for choosing this design, decoupling creating from running an experiment. The package also includes the subclasses of `ExperimentBuilder` and `Experimenter` used in this thesis. There is also a class `GamePhases` which is used to synchronise players, games, and experimenter.
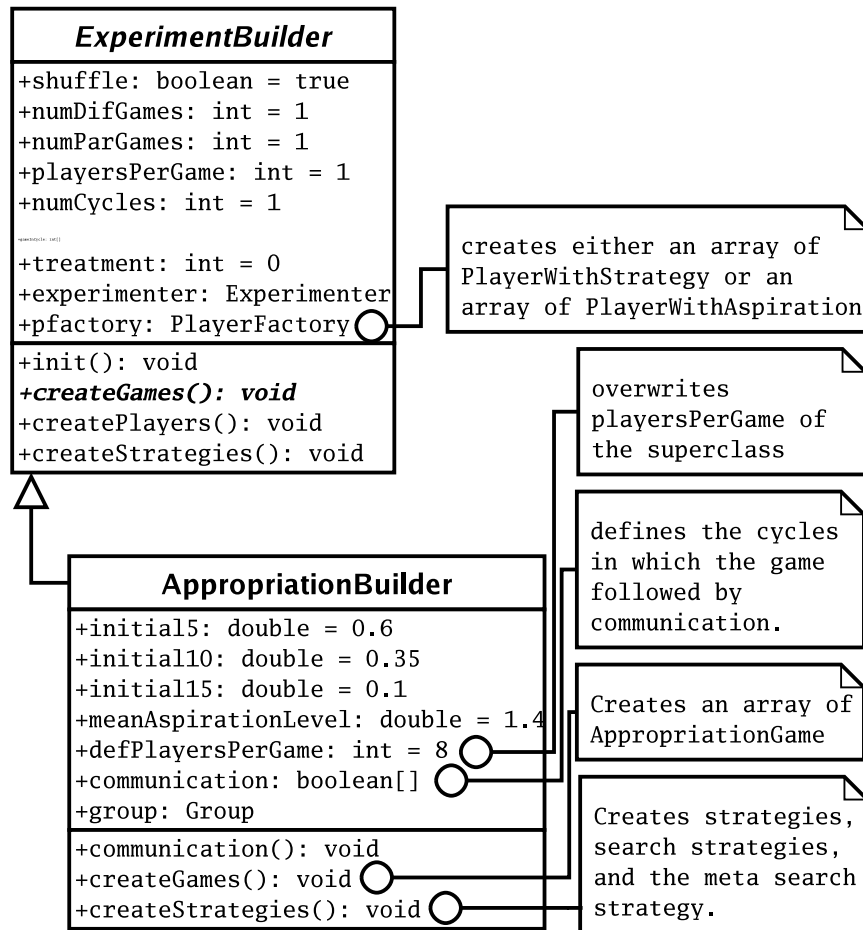
---

[1]http://www.usf.uos.de/projects/famoja/

Figure A.1: UML diagram of `ExperimentBuilder` and one of its subclasses `AppropriationBuilder`, see chapter 5 with selected instance variables and methods.

GamePhases class

`GamePhases` is a mechanism to synchronize the experimenter, games and players. It defines a number of phases which are cycled through, each phase is one step of the model run. The default game phases defined in this class are SETUP, START, PLAYER, GAME, and DATA. Usually, the SETUP phase occurs only once and only the builder acts in it. In repeated games, after the DATA phase, the next phase is START and not SETUP. In START phase, everything that has to happen, before players make their decisions can be done, for example shuffling players. In the PLAYER phase, players make their decisions. In the GAME phase, games evaluate the outcome and inform players about it. In the DATA phase, players learn, data is displayed and everything is triggered that has to happen after a game is played.

ExperimentBuilder class

`ExperimentBuilder` has to set all parameters of games and players. Subclasses can have a large number of instance variables. The superclass provides instance variables for the number of cycles (that is successive games), which game to play in which cycle, the number of different games in this experiment, the number of parallel, similar games, the players per game, the total number of players, the number of charts for displaying data, the chart names, a tolerance value, a way to define different treatments in the same experiment, and a boolean for defining whether or not to shuffle the players. Additionally, it has an instance of `GamePhases`, of `PlayerFactory` to organise the player creation, and of `Experimenter` to instruct it how to organise the experiment.

In the `action` method the default behaviour is implemented: The `init` method is called in the SETUP phase. Nothing else happens. The initialisation process, however, is huge. First, the `Experimenter` is initalised with the number of different games and whether or not to shuffle players between games. Then, games are created. In this abstract superclass this is an abstract method `createGames`, which has to be overwritten by subclasses. After the games, the players are created. This is done by one of several instances of `PlayerFactory`, depending on which kind of players are needed for the experiment (cf. Gamma *et al.* (1994, 107ff.)). After player creation these are assigned to the experimenter. Then strategies are created. This also has to be implemented by subclasses, but is implemented as a default process instead of an abstract method. The method `createStrategies` creates a strategy pool according to the number of different games and de-

cisions of these games, but fills it with the default `Strategy` only. Then the
`Experimenter` sets the players' strategy pool and finishes its setup process
(see below).

Experimenter class

The `Experimenter` class defines a number of instance variables to count the
different games, to know which game to play in which cycle, the current cy-
cle, the games organised in a `Swarm`, the players per game, the total number
of players, player indices (so that not the actual players are shuffled, but
only their indices), the players also organised in a `Swarm`, the `GamePhases`,
the charts also in a `Swarm`, a triple array of strategies, a triple array of search
strategies and a double array of meta search strategies, whether or not to
shuffle, a low limit, a high limit, and a tolerance value.

   The setup of the `Experimenter` is mainly done through the builder. The
games, players and strategies are created before the experimenter enters its
first active phase. The `action` method organises the behaviour during the
model run, referring to the game phases. In the START phase, the method
`start` is called, and in the DATA phase, the method `data` is called. If only
slight changes need to be made for subclasses it suffices to overwrite these
methods and call their super methods and then or before do what else has
to be done. In the default `start` method players are shuffled and assigned
to the games. In the default `data` method, data is collected by the `col-
lectData` method, and the count for the current game and the current cycle
are increased. The method `collectData` calls the players' `collectData`
method and lets them memorise the last game, and then calls the charts'
`collectData` method to let them do the data display.

   The default `Experimenter` provides simple functionality for a communi-
cation process in a group but does not call it itself. Likewise, it can call the
players' aspiration adaptation process, but does not do so per default. The
`boolean aspire` can be set to `true` in order to trigger calling the `aspire`
method.

A.2   Games

The package `games` contains the abstract superclass of all games, called
`Game`. Also, all non-abstract games implemented in the course of this thesis
are included in this package. In the beginning, games were subclasses of ab-
stract classes that defined the attribute or attributes used in the game. The
advantage of this approach was that classes like `CooperationGame` could

implement default learning mechanisms that changed the expected cooperativeness of players, as well as a method that returned the attribute in question. But this mechanism was discarded, when the games became more complex and depended on five or more attributes. In addition to the game classes described in the papers, which I will not discuss here, there are also simple standard games like `PrisonersDilemma` and `ChickenGame`. The `Game` class and some subclasses are illustrated in the Handbook paper, Chapter 5).

Game class

`Game` class provides standard functionality of games. For simple games, the abstract superclass does everything except for the formulation of the actual rules, like for example the payoff matrix of a Prisoner's Dilemma. Games are defined as simple games. Resulting supergames that are a succession of several similar games, are implemented as experiments. The abstract superclass has as instance variables the number of players, an array containing the players, the number of decisions each player has to make, a double array of decisions of each player and each decision, the assets, the decision minimum and maximum, learning rates for altering expected attributes, the attributes needed for the game, an instance of `GamePhases`, and a parameter to define a maximum difference for values to be considered equal.

In the constructor, `Game` initializes the `GamePhases` and sets them to START, because only the builder needs to do something in the init phase. In the `action` method a default behaviour is implemented. This consists of only two calls: `giveAssets` is called in the START phase, and `evaluate` is called in the GAME phase (see below).

All instance variables can be set and obtained. The method `isEqual` decides whether or not two values are considered equal, which tests whether the difference is less than the maximum defined. The method `scaleValue` scales a value assumed to be between 0 and 1 to the defined interval, defined by `scaleMin` and `scaleMax`. Another method `correctDecisionValue` corrects a decision obtained by a player cropping it to the defined interval. This also round it to `int`. It is important to keep track of when decisions are given as values between 0 and 1, and when they represent the actual decision values between the minimum and maximum defined by the game. Usually, strategies call the game's method `scaleValue` to transform decisions between 0 and 1 to the required range (see Section A.4). The method `giveAssets` assigns the players their initial assets. This method is called in the START phase. The method `playersDecisions` fills the array of decisions

with the players' decisions. This method is called in the GAME phase. The method `getAttributesOfPlayer` returns an array with the attribute values of the given players as they are defined in the attributes array of the game. That is, if a game defines its important attributes to be cooperativeness and positive reciprocity, this method returns an array with two double values, the first is the cooperativeness of the player and the second its positive reciprocity.

There are several methods with the names `expectedOutcome` and `expectedOtherOutcome` with different parameters. They may call one another. With these methods players are supposed to be able to form expectations. Currently, only the `MaximizingStrategy` and one meta search strategy use these methods. Subclasses of `Game` do not have to overwrite these methods, but then players strategies, which depend on expectations may not work. Most strategies do not depend on expectations, the most important exception is `MaximizingStrategy`. By letting the methods call one another, a game can define methods for expectation formation which depend on different aspects, like the players decision, the other players' expected decisions, or the representation of an average other player. For the latter, players do not have to assume a decision making mechanism on the part of the other players. Instead, this has to be implemented in the game. This is important for asymmetrical games.

`Game` has a method `learn` which alters the player's cooperativeness, assuming this default attribute to be the attribute in question. Subclasses can overwrite `learn` with their specific learning process. Alternativly, the `Experimenter` can refrain from triggering the player's learning process by setting the players' instance variable `learn` to `false`.

The abstract method `evaluate` has to be overwritten by games with their actual rules.

A.3   Players

The `players` package contains a number of player classes. Their dependency is shown in Figure A.2.

PlayerWithPersonality class

The simplest player, `PlayerWithPersonality`, has a `Personality` instance, which is also a class of this package containing the seven attributes. `Personality` can be altered to include new attributes. The attributes are also implemented in a class `Cooperativeness`, `Conformity` and so forth. In
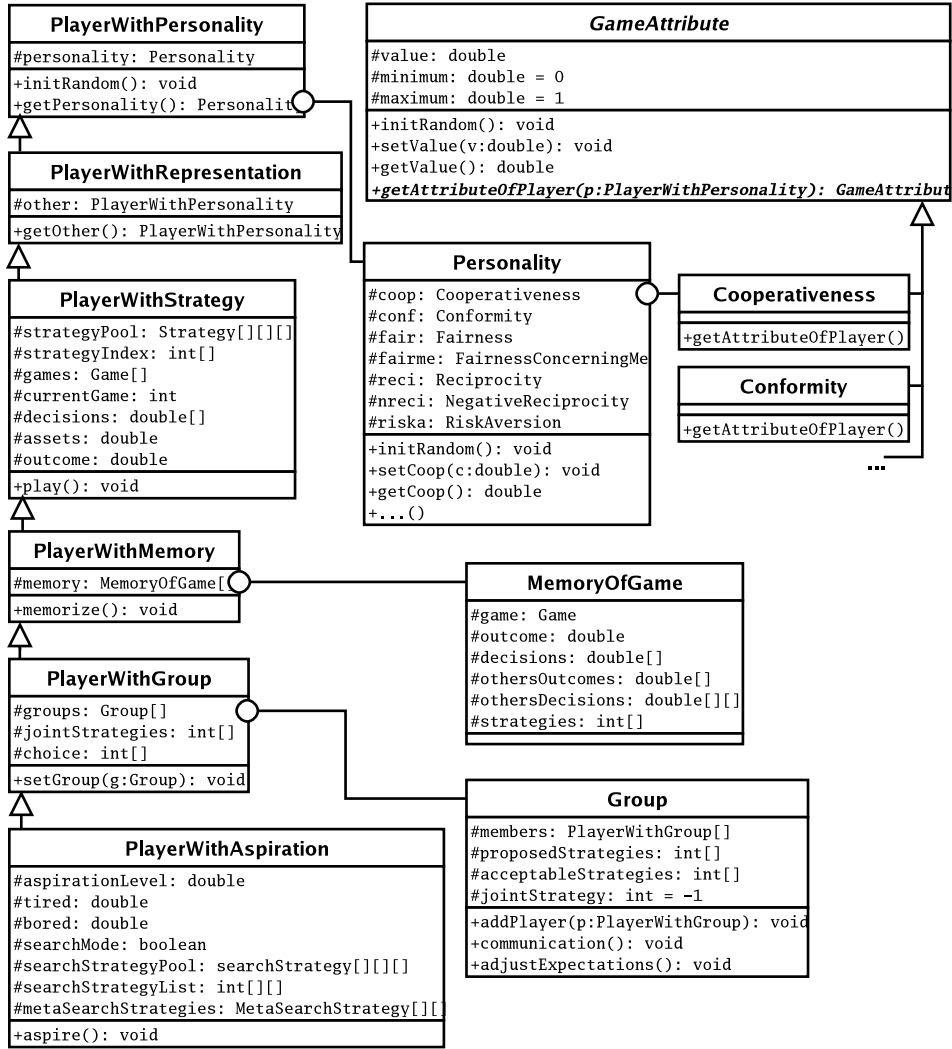
Figure A.2: UML diagram of the tree of player classes with selected instance variables and methods.

these classes new boundaries for attribute values or specific ways to define a distribution of attribute values can be implemented. The default is an equal distribution of values between 0 and 1.

PlayerWithRepresentation class

The second simplest player, `PlayerWithRepresentation`, has an instance of `PlayerWithPersonality` in order to represent the average other player's expected attributes.

PlayerWithStrategy class

`PlayerWithStrategy` is the simplest player class able to play games. It has a strategy pool and a representation of all game types, it has to play. It keeps track of its gains and losses and accumulated assets and has a memory containing one game only as well as expectations concerning games to play in the near future. For example, a player who expects a punishment game after a common pool resource game can behave differently than one who expects another common pool resource game to follow.

Being the first player that actually plays games, this class defined default behaviour in games according to the game phases. In the PLAYER phase, the method `play` is called, which fills the decision array by calling the method `reachDecision` for each decision. In the DATA phase, the method `data` is called, which trigger the learning mechanism, if the `boolean learn` is set to true. The learning process is defined in the game played.

PlayerWithMemory class

`PlayerWithMemory` has extended memory and can store various information about previous games in its memory. The class `MemoryOfGame` has instance variables to represent the game type, the outcome, the assets, the gain and loss, the player's decisions, its strategies, the others' outcomes, and their decisions. `PlayerWithMemory` has an array of `MemoryOfGame` which is filled with all information known to the player. This is important for players learning with case based reasoning mechanisms. This class overwrites the method `memorize` which is defined in `PlayerWithStrategy` but there only remembers the last game played. This functionality is kept by calling the super method.

PlayerWithGroup class

`PlayerWithGroup` is a member of one or more groups of players, and can have representations of the other members of the group. The expectations are no longer mean values, but can be altered to represent other individual players. Furthermore, the class has a choice strategy and a joint strategy, the latter represents the strategy it should follow according to group decisions. Group processes, like communication, are implemented in the class `Group`.

PlayerWithAspiration class

`PlayerWithAspiration` has an aspiration threshold and can be given thresholds for tiredness and boredom that define the aspiration adaptation process. It also has a search strategy pool in addition to the strategy pool and a meta search strategy to decide on a search strategy. The synchronization of strategies, search and meta search strategies is rather lengthy. The class also defines a default aspiration adaptation process.

A.4 Strategies

The `strategies` package contains the default `Strategy` which is the superclass of all strategies, as well as all particular strategies used in the various games in this thesis. Not all of these are discussed here, but a number of generic strategies usable in different situations is introduced. Additionally, this package contains `SearchStrategy` and `MetaSearchStrategy` as well as their subclasses. Examples for subclasses of `Strategy` can be found in the Handbook paper, Chapter 5).

Strategy class

`Strategy` is the superclass of all decision making strategies. It has instance variables to represent the decision value, an array of indices of attributes to use, and a parameters: an offset, lower boundary, upper boundary, two thresholds, and a decision maximum. The parameters can be set and obtained by standard methods. The only other method is `decision` which makes a decision based on the calling player's attributes, the game and an optional index number. The decision is a double value and is returned by this method. In this default class, the decision is equal to the attribute value corresponding to the first in the list of attribute indices. The attribute is obtained by calling the `getAttributesOfPlayer` method of the instance of `Game` given to the method.

Selected subclasses of the Strategy class

Some generic strategies have been implemented that work on either the
attributes of the players or the interval of a decision. The simplest of these
are `NoStrategy` and `YesStrategy` which correspond to returns of 0 and 1.
Note, that some strategies return values between 0 and 1 and the modeller
has to keep track whether or not the decisions have to be scaled to the
actual values by the game or not. In addition to the `YesStrategy` there is a
`MaxStrategy` which returns the maximum possible as scaled by the game, by
calling the game's `scaleValue` method. Equally simple is `ValueStrategy`
which returns a predefined decision value within the boundaries set by the
game. The value has to be set by the `ExperimentBuilder` when the strategy
is created. Slightly more complex is `ExternalValueStrategy`, which uses a
predefined value and calls the game's method `scaleValue` in order to scale
it to the required range.

OneValueStrategy uses the attribute of the calling player, defined by
the first entry in the indeces vector and makes the following calculation: It
sets lower boundary to the minimum possible decision and upper bound-
ary to the maximum possible decision and then calculates $lowerboundary +$
$attributevalue * (upperboundary - lowerboundary)$. Additionally, an offset is
added to the attribute before calculating the decision. `OneExpValueStrate-`
`gy` does the same, interpreting the attribute as an expected attribute of the
player. In order to obtain the expected attribute, `getAttributesOfPlayer`
is called with the player's representation of other players as parameter in-
stead of the player itself. `OneValueModerateStrategy` alters the attribute
value by adding one half and dividing by two. `TwoValueStrategy` takes the
mean of the first two attribute values and calculates as before. `OneValue-`
`OneExpValueStrategy` takes the mean of the first attribute and the second
interpreted as an expected attribute. `MultiplyStrategy` multiplies the first
two attributes before making the calculation.

EqualSplitStrategy returns the maximum possible decision divided by
the number of players.

MaximizingStrategy tries out a number of possible decision values and
calls the game's `expectedOutcome` method to choose the decision with the
highest expected outcome.

SearchStrategy class

`SearchStrategy` is also a subclass of `Strategy` but also defines a step. This
can be set from outside. The step refers to the change within the strategy

pool of a player. If a player uses currently strategy number 4 and moves 2 steps up, it uses strategy number 6 next. The method `searchDecision` returns the new strategy index within the strategy pool. It returns only numbers between zero and the number of strategies in the pool minus 1.

MetaSearchStrategy class

`MetaSearchStrategy` extends `SearchStrategy` and returns the index of the next `SearchStrategy` to use. No basic functionality is provided by this simple class, which only returns the same search strategy index, the player is already using. The meta search of strategies depends strongly on the decision environment.

A.5   Charts

ExperimenterChart class and selected subclasses

Charts are used to display data. The charts are subclasses of `Experimenter-Chart`, which provides a method for initializing the titles and minimum and maximum of the y-axis. It also defines the method `collectData` which is called by the `Experimenter` in the DATA phase. In this method, particular charts can obtain data from the `Experimenter` or players and make calculations to display all possible data. There are a few generic charts implemented, but usually new subclasses will have to be implemented for specific experiments. The generic classes include `DecisionGainExpChart` which displays the average decision and average gain of all players, `DecisionGain-PlayerChart`, which displays a player's decision and gain, `DecisionTotal-ExpChart`, which displays the sum of all players' decisions, `DecMeanOf-StrategyExpChart`, which displays the means of the players' decisions in categories of the strategy employed, and `StrategiesExpChart`, which displays the percentage of players using each strategy.

Bibliography

Albers, W. (2001). Prominence Theory as a Tool to Model Boundedly Rational Decisions. In Gigerenzer and Selten (2001a), chapter 17, pages 297–317.

Andreoni, J. (1989). Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of Political Economy*, **97**(6), 1447–1458.

Andreoni, J. and Miller, J. H. (1995). Auctions with artificial adaptive agents. *Games and Economic Behavior*, **10**, 39–64.

Andreoni, J., Brown, P. M., and Vesterlund, L. (2002). What makes an allocation fair? some experimental evidence. *Games and Economic Behavior*, **40**, 1–24.

Augier, M. and March, J. G. (2003). The economic psychology of Herbert A. Simon: Introduction to a special issue. *Journal of Economic Psychology*, **24**, 135–141.

Barreteau, O., Bousquet, F., and Attonaty, J.-M. (2001). Role-playing games for opening the black box of multi-agent systems: method and lessons of its application to senegal river valley irrigated systems. *Journal of Artificial Societies and Social Simulation*, **4**(2). <http://jasss.soc.surrey.ac.uk/4/2/5.html> (June 10, 2006).

Ben-Ner, A., Kong, F., and Putterman, L. (2004). Share and share alike? gender-pairing, personality, and cognitive ability as determinants of giving. *Journal of Economic Psychology*, **25**, 581–589.

Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, **10**, 122–142.

Berger, T. (2001). Agent-based spatial models applied to agriculture: A simulation tool for techonology diffusion, resource use change and policy analysis. *Agricultural Economics*, **25**(2/3), 245–260.

Bicchieri, C. (2002). Covenants without swords. group identity, norms, and communication in social dilemmas. *Rationality and Society*, **14**(2), 192–228.

Bolton, G. E. and Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, **90**(1), 166–193.

Bolton, G. E. and Zwick, R. (1995). Anonymity and punishment in ultimatum bargaining. *Games and Economic Behavior*, **10**, 95–121.

Boyd, R. and Richerson, P. J. (2001). Norms and Bounded Rationality. In Gigerenzer and Selten (2001a), chapter 16, pages 281–296.

Brandstätter, H. (1988). Sechzehn Persönlichkeits-Adjektivskalen (16 PA) als Forschungsinstrument anstelle des 16 PF. *Zeitschrift für experimentelle und angewandte Psychologie*, **35**(3), 370–391.

Brandstätter, H., Güth, W., Himmelbauer, J., and Kriz, W. (1999). Dispositions and actual behaviour in dictator and ultimatum games. <http://www.proprint-service.de/pp/prvw?idp=oai%3AHUBerlin.de%3A25648> (June 10, 2006).

Brenner, T. (1999). Cognitive learning in prisoner's dilemma situations. In T. Brenner, editor, *Computational techniques for modelling learning in economics*, pages 333–361. Boston, Kluwer.

Burnham, T. C. (2003). Engineering altruism: a theoretical and experimental investigation of anonymity and gift giving. *Journal of Economic Behavior & Organization*, **50**, 133–144.

Camerer, C. F. (2003). *Behavioral Game Theory*. Russel Sage Foundation, New York.

Charness, G. and Haruvy, E. (2002). Altruism, equity, and reciprocity in a gift-exchange experiment: an encompassing approach. *Games and Economic Behavior*, **40**, 203–231.

Collentine, D., Larsson, M., and Hannerz, N. (2004). Exploiting decision heuristics and IT in the design of a DSS for voluntary agri-environmental programs. *Ecological Economics*, **49**, 303–315.

Colman, A. M. (2003). Cooperation, psychological game theory, and limitations of rationality in social interaction. *Behavioural and Brain Sciences*, **26**, 139–198.

Cook, K. S. and Cooper, R. M. (2004). Experimental studies of cooperation, trust, and social exchange. In Ostrom and Walker (2004), chapter 8, pages 209–244.

Cox, J. C. (2004). How to identify trust and reciprocity. *Games and Economic Behavior*, **46**, 260–281.

Dal Forno, A. and Merlone, U. (2004). From classroom experiments to computer code. *Journal of Artificial Societies and Social Simulation*, **7**(3). <http://jasss.soc.surrey.ac.uk/7/3/2.html> (June 10, 2006).

D'Aquino, P., Le Page, C., Bousquet, F., and Bah, A. (2003). Using self-designed role-playing games and a multi-agent system to empower a local decision-making process for land use management: The self-cormas experiment in senegal. *Journal of Artificial Societies and Social Simulation*, **6**(3). <http://jasss.soc.surrey.ac.uk/6/3/5.html> (June 10, 2006).

Davidsson, P. (2002). Agent based social simulation: A computer science view. *Journal of Artificial Societies and Social Simulation*, **5**(1). <http://jasss.soc.surrey.ac.uk/5/1/7.html> (June 10, 2006).

de Mandeville, B. (1990 [1714]). *The fable of the bees or, private vices, publick benefits*. Düsseldorf: Verlag Wirtschaft und Finanzen.

Deadman, P. J. and Schlager, E. (2002). Models of Individual Decision Making in Agent-Based Simulation of Common-Pool-Resource Management Institutions. In H. R. Gimblett, editor, *Integrating Geographic Information Systems and Agent-Based Modeling Techniques for Simulating Social and Ecological Processes*, pages 137–169. Oxford: Oxford Univ. Press.

Deadman, P. J., Schlager, E., and Gimblett, R. (2000). Simulating Common Pool Resource Management Experiments with Adaptive Agents Employing Alternate Communication Routines. *Journal of Artificial Societies and Social Simulation*, **3**(2). <http://jasss.soc.surrey.ac.uk/JASSS/3/2/2.html> (June 10, 2006).

Deffuant, G., Amblard, F., Weisbuch, G., and Faure, T. (2002). How can extremism prevail? a study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulation*, **5**(4). <http://jasss.soc.surrey.ac.uk/JASSS/5/4/1.html> (June 10, 2006).

Deffuant, G., Weisbuch, G., Amblard, F., and Faure, T. (2003). Simple is beautiful ... and necessary. *Journal of Artificial Societies and Social Simulation*, **6**(1). <http://jasss.soc.surrey.ac.uk/6/1/6.html> (March 1st, 2007).

Demsetz, H. (1967). Toward a theory of property rights. *American Economic Review*, **57**, 347.

Diekmann, A. and Preisendörfer, P. (2001). *Umweltsoziologie. Eine Einführung*. Reinbeck: rohwolts enzyklopädie.

Duffy, J. M. (2001). Learning to speculate: Experiments with artificial and real agents. *Journal of Economic Dynamics & Control*, **25**(3/4).

Ebenhöh, E. (2006). Modeling non-linear common-pool resource experiments with boundedly rational agents. In J. Sichman and L. Antunes, editors, *MABS 2005*, number 3891 in LNAI, pages 133–146. Springer-Verlag Berlin Heidelberg.

Ebenhöh, E. and Pahl-Wostl, C. (2006). Agent-based modelling with boundedly rational agents. In J.-P. Rennard, editor, *Handbook of Research on Nature Inspired Computing for Economy and Management*. Idea Group.

Epstein, J. M. and Axtell, R. (1995). *Growing artificial societies: Social Science from the Bottom Up*. Washington: Brokings Institutions Press/MIT Press Cambridge, MA.

Esser, H. and Troitzsch, K. G., editors (1991). *Modellierung sozialer Prozesse*. Informationszentrum Sozialwissenschaften Bonn.

Falk, A. and Fischbacher, U. (2000). A theory of reciprocity. *Institute of Empirical Research in Economics, University of Zurich, Working Paper Series*, **6**.

Fehr, E. and Gächter, S. (2002). Altruistic punishment in humans. *Nature*, **415**, 137–140.

Fehr, E. and Rockenbach, B. (2003). Detrimental effects of sanctions on human altruism. *Nature*, **422**, 137–140.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, pages 817–868.

Ferber, J. (1999). *Multi-Agent Systems*. Addison-Wesley.

Forsythe, R., Horowitz, J. L., Savin, N. E., and Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior*, **6**, 347–369.

Gale, J., Binmore, K. G., and Samuelson, L. (1995). Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, **8**, 56–90.

Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1994). *Design patterns: elements of reusable object-oriented software*. Reading (Massachusetts).

Gigerenzer, G. (2001). The adaptive toolbox. In Gigerenzer and Selten (2001a), chapter 3, pages 37–50.

Gigerenzer, G. and Selten, R., editors (2001a). *Bounded Rationality. The Adaptive Toolbox*. Cambridge, Massachusetts: The MIT Press.

Gigerenzer, G. and Selten, R. (2001b). Rethinking rationality. In Gigerenzer and Selten (2001a), chapter 1, pages 1–12.

Gigerenzer, G. and Todd, P. M. (1999). Fast and frugal heuristics. the adaptive toolbox. In Gigerenzer *et al.* (1999), chapter 1, pages 3–34.

Gigerenzer, G., Todd, P. M., and the ABC Research Group, editors (1999). *Simple Heuristics That Make Us Smart*. New York, Oxford: Oxford University Press.

Gilboa, I. and Schmeidler, D. (2001). *A Theory of Case-Based Decisions*. Cambridge University Press, Cambridge, UK.

Goldstein, D. G. (2001). Group report: Why and when do simple heuristics work? In Gigerenzer and Selten (2001a), chapter 10, pages 173–190.

Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, **91**(3), 481–510.

Güth, W. (1983). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, **3**, 367–388.

Güth, W. and Huck, S. (1997). From ultimatum bargaining to dictatorship - an experimental study of four games varying in veto power. *Metroeconomica*, **48**(3), 262–279.

Güth, W. and Tietz, R. (1990). Ultimatum bargaining behaviour. a survey and comparison of experimental results. *Journal of Economic Psychology*, **11**, 417–449.

Güth, W. and van Damme, E. (1998). Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study. *Journal of Mathematical Psychology*, **42**, 227–247.

Güth, W., Kliemt, H., and Ockenfels, A. (2003). Fairness versus efficiency. an experimental study of (mutual) gift giving. *Journal of Economic Behavior & Organization*, **50**, 465–475.

Hardin, G. (1968). The tragedy of the commons. *Science*, **162**, 1243–1248.

Hare, M. and Deadman, P. (2004). Further towards a taxonomy of agent-based simulation models in environmental management. *Mathematics and Computers in Simulation*, **64**, 25–40.

Hare, M. and Pahl-Wostl, C. (2001). Model uncertainty derived from choice of agent rationality - a lesson for policy assessment modelling. In N. Giambiasi and C. Frydman, editors, *Simulation in Industry: 13th European Simulation Symposium*, pages 854–859. SCS Europe Bvba, Ghent.

Harvey S. James, J. (2002). On the reliability of trusting. *Rationality and Society*, **14**(2), 229–256.

Hayashi, N., Ostrom, E., Walker, J., and Yamagishi, T. (1999). Reciprocity, trust, and the sense of control. *Rationality and Society*, **11**(1), 27–46.

Heckathorn, D. D. (1989). Collective action and the second-order free-rider problem. *Rationality and Society*, **1**(1), 78–100.

Heckathorn, D. D. (1996). The dynamics and dilemmas of collective action. *American Sociological Review*, **61**(April), 250–277.

Held, M. and Nutzinger, H. G., editors (1999). *Institutionen prägen Menschen - Menschen prägen Institutionen*. Campus Verlag, Frankfurt/Main.

Hennig-Schmidt, H., editor (1999). *Bargaining in a Video Experiment. Determinants of Boundedly Rational Behavior*, volume 467 of *Lecture Notes in Economics and Mathematical Systems*. Springer.

Hennig-Schmidt, H., Li, Z.-Y., and Yang, C. (2002). Non-monotone strategies in ultimatum bargaining: First results from a video experiment in the people's republic of china. In *Proceedings of the International Congress of Mathematicians Game Theory and Applications Satellite Conference*, pages 225–231. Qingdao Publishing House. <http://webserver.econ1.uni-bonn.de/people/hschmidt/Papers.html>.

Henrich, J. (2000). Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga of the Peruvian Amazon. *American Economic Review*, **90**, 973–979.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., Alvard, M., Barr, A., Ensminger, J., Hill, K., Gil-White, F., Gurven, M., Marlowe, F., Patton, J. Q., Smith, N., and Tracer, D. (2001a). "Economic Man" in Cross-cultural Perspective: Behavioral Experiments in 15 Small-scale Societies. Working Paper, Santa Fe Institute. <http://www.santafe.edu/research/publications/wplist/2001> (June 26, 2006).

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., and McElreath, R. (2001b). In search of homo economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review*, **91**, 73–78.

Hoffman, E., McCabe, K., Shachat, K., and Smith, V. L. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, **7**, 346–380.

Hoffman, E., McCabe, K., and Smith, V. L. (1996). Social distance and other-regarding behavior in dictator games. *American Economic Review*, **86**, 653–660.

Isaac, R. M., Walker, J. M., and Thomas, S. H. (1984). Divergent evidence on free riding: An experimental examination of possible explanations. *Public Choice*, **43**, 113–149.

Izquierdo, L. R., Gotts, N. M., and Polhill, J. G. (2004). Case-based reasoning, social dilemmas, and a new equilibrium concept. *Journal of Artificial Societies and Social Simulation*, **7**(3). <http://jasss.soc.surrey.ac.uk/7/3/1.html> (June 10, 2006).

Jaffe, K. (2002). An economic analysis of altruism: who benefits from altruistic acts? *Journal of Artificial Societies and Social Simulation*, **5**(3). <http://jasss.soc.surrey.ac.uk/5/3/3.html> (June 10, 2006).

Jager, W. and Janssen, M. A. (2002). Using artificial agents to understand laboratory experiments of common-pool resources with real agents. In M. A. Janssen, editor, *Complexity and Ecosystem Management*, pages 75–102. Bodmin, Cornwall.

Jager, W., Janssen, M. A., De Vries, H. J. M., De Greef, J., and Vlek, C. A. J. (2000). Behaviour in commons dilemmas: Homo economicus and homo psychologicus in an ecological-economic model. *Ecological Economics*, **35**(3), 307–416.

Kagel, J. H., editor (1995). *The handbook of experimental economics*. Princeton Univ. Press.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**(2), 263–291.

Kahneman, D. and Tversky, A., editors (2000). *Choices, Values, and Frames*. Russel Sage Foundation, Cambridge University Press.

Kleindorfer, P. R., Kunreuther, H. C., and Schoemaker, P. J. H. (1993). *Decision Sciences*. Cambridge University Press, Cambridge, UK.

Kottonau, J. (2002). Simulating the formation and change of the strength of political attitudes. Diss. Eth No. 14664, Swiss Federal Institute of Technology Zurich.

Kottonau, J. and Pahl-Wostl, C. (2002). Simulating the formation of political attitude strength during election campaigns. CSC.

Kreps, D. M., editor (1988). *Notes on the theory of choice*. Westview Press.

Liebert, R. M. and Spiegler, M. D., editors (1994). *Personality: srategies and issues*. Wadsworth, Belmont, California.

Lindenberg, S. (1991). Die Methode der abnehmenden Abstraktion: Theoriegesteuerte Analyse mit empirischem Gehalt. In Esser and Troitzsch (1991), pages 29–78.

Loewenstein, G. (1999). Experimental economics from the vantage-point of behavioural economics. *The Economic Journal*, **109**, F25–F34.

Lovett, F. (2006). Rational choice theory and explanation. *Rationality and Society*, **18**(2), 237–272.

Martignon, L. (2001). Comparing fast and frugal heuristics and optimal models. In Gigerenzer and Selten (2001a), chapter 9, pages 147–171.

McCabe, K. A. and Smith, V. L. (2001). Goodwill accounting and the process of exchange. In Gigerenzer and Selten (2001a), chapter 18, pages 319–340.

McCabe, K. A., Rigdon, M. L., and Smith, V. L. (2003). Positive reciprocity and intentions in trust games. *Journal of Economic Behavior & Organization*, **52**, 267–275.

McClintock, C. G., editor (1972a). *Experimental Social Psychology*. Holt, Rinehart and Winston, New York.

McClintock, C. G. (1972b). Game behavior and social motivation in interpersonal settings. In McClintock (1972a), chapter 9, pages 271–297.

McClintock, C. G. (1972c). Social motivation - a set of propositions. *Behavioral Science*, **17**, 438–454.

McClintock, C. G. and Liebrand, W. B. G. (1988). Role of Interdependence Structure, Individual Value Orientation, and Another's Strategy in Social Decision Making: A Transformational Analysis. *Journal of Personality and Social Psychology*, **55**(3), 396–409.

Mellers, B. A. (2001). Group report: Effects of emotions and social processes on bounded rationality. In Gigerenzer and Selten (2001a), chapter 16, pages 263–278.

Mill, J. S. (1865). *Principles of political economy; with some of their applications to social philosophy*. London: Longman.

Moss, S. and Edmonds, B. (2005). Sociology and simulation: Statistical and qualitative cross-validation. *American Journal of Sociology*, **110**(4), 1095–1131.

Mulder, L. B., Van Dijk, E., Wilke, H. A. M., and De Cremer, D. (2005). The effect of feedback on support for a sanctioning system in a social dilemma: The difference between installing and maintaining the sanction. *Journal of Economic Psychology*, **26**, 443–458.

Newig, J., editor (2003). *Symbolische Umweltgesetzgebung. Rechtssoziologische Untersuchungen am Beispiel des Ozongesetzes, des Kreislaufwirtschaft- und Abfallgesetzes sowie der Großfeuerungsanlagenverordnung*. Schriften zur Rechtssoziologie und Rechtstatsachenforschung. Berlin, Duncker & Humblot.

Nooteboom, B., editor (2002). *Trust: Forms, Foundations, Functions, Failures and Figures*. Edward Elgar, Cheltenham UK.

Nooteboom, B. and Six, F. (2003a). Introduction. In Nooteboom and Six (2003b), chapter 1, pages 1–15.

Nooteboom, B. and Six, F., editors (2003b). *The Trust Process in Organizations*. Edward Elgar, Cheltenham UK, Northampton MA USA.

Orcutt, G. H. (1957). A new type of socio-economic system. *Review of Economics and Statistics*, **39**(2), 116–123.

Ostrom, E., editor (1990). *Governing the commons: the evolution of institutions for collective action*. Cambridge, New York.

Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of Economic Perspectives*, **14**, 137–158.

Ostrom, E. (2004). Toward a behavioral theory linking trust, reciprocity, and reputation. In Ostrom and Walker (2004), chapter 2, pages 19–79.

Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.

Ostrom, E. and Walker, J., editors (2004). *Trust and Reciprocity*, volume VI of *Russel Sage Foundation Series on Trust*. Russel Sage Foundation, New York.

Ostrom, E., Gardner, R., and Walker, J., editors (1994). *Rules, games, and common-pool resources*. Ann Arbor, Michigan: The University of Michigan Press.

Pahl-Wostl, C. (2002). Participative and Stakeholder-based policy design, evaluation and modeling processes. *Integrated Assessment*, **3**(1), 3–14.

Pahl-Wostl, C. and Ebenhöh, E. (2004). An adaptive toolbox model: a pluralistic modelling approach for human behaviour based on observation. *Journal of Artificial Societies and Social Simulation*, **7**(1). <http://jasss.soc.surrey.ac.uk/7/1/3.html> (June 10, 2006).

Pahl-Wostl, C. and Hare, M. (2004). Processes of social learning in integrated resources management. *Journal of Community & Applied Social Psychology*, **14**, 193–206.

Panebianco, S. and Pahl-Wostl, C. (2004). Obstacles in launching a participatory modelling process. In C. Pahl-Wostl, S. Schmidt, A. E. Rizzoli,

and A. J. Jakeman, editors, *Complexity and Integrated Resources Management, Transactions of the 2nd Biennial Meeting of the International Environmental Modelling and Software Society*. International Environmental Modelling and Software Society (iEMSs), Manno, Switzerland. <http://www.iemss.org> (June 10, 2006).

Pingle, M. and Tesfatsion, L. (2001). Non-Employment Benefits and the Evolution of Worker-Employer Cooperation: Experiments with Real and Computational Agents. ISU Economic Report No. 55. <http://www.econ.iastate.edu/tesfatsi/sce2001.pdf> (June 10, 2006).

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, **83**(5), 1281–1302.

Randow, G. V. (2003). When the center becomes radical. *Journal of Artificial Societies and Social Simulation*, **6**(1). <http://jasss.soc.surrey.ac.uk/6/1/5.html> (June 10, 2006).

Raub, W. (2004). Hostage posting as a mechanism of trust. *Rationality and Society*, **16**(3), 319–365.

Rubinstein, A. (1998). *Modeling Bounded Rationality*. The MIT Press.

Selten, R. (1978). The chain store paradox. *Theory and Decision*, **9**, 127–159.

Selten, R. (1990). Bounded rationality. *Journal of Institutional and Theoretical Economics*, **146**, 649–658.

Selten, R., editor (1999). *Game Theory and Economic Behaviour*, volume One. Edward Elgar, Cheltenham UK.

Selten, R. (2001). What is bounded rationality? In Gigerenzer and Selten (2001a), chapter 2, pages 13–36.

Sen, A. (1987). *On Ethics and Economics*. Basil Blackwell.

Simon, H. (1990). Invariants of human behavior. *Annual Review of Psychology*, **41**, 1–19.

Simon, H. A. (1997). *Models of Bounded Rationality*, volume 3. The MIT Press, Cambridge, Massachusetts.

Smith, A. (1986 [1776]). *An inquiry into the nature and causes of the wealth of nations*. Düsseldorf, Verlag Wirtschaft und Finanzen.

Straub, P. G. and Murnighan, J. K. (1995). An experimental investigation of ultimatum games: information, fairness, expectations, and lowest acceptable offers. *Journal of Economic Behavior & Organization*, **27**, 345–365.

Tesfatsion, L. (2002). Agent-Based Computational Economics: Growing Economies from the Bottom Up. ISU Economics Working Paper No. 02001. <http://www.econ.iastate.edu/tesfatsi/surveys.htm> (June, 10, 2006).

Tietz, R. (1990). On Bounded Rationality: Experimental Work at the University of Frankfurt/Main. *Journal of Institutional and Theoretical Economics*, **146**, 659–672.

Todd, P. M. and Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology*, **24**, 143–165.

Vanberg, V. J. (2002). Rational choice vs. program-based behavior. alternative theoretical approaches and their relevance for the study of institutions. *Rationality and Society*, **14**(1), 7–54.

von Neumann, J. and Morgenstern, O. (1953). *Theory of games and economic behavior*. Princeton, N.J.: Princeton University Press.