

6. Using artificial agents to understand laboratory experiments of common-pool resources with real agents

Wander Jager and Marco A. Janssen

6.1 INTRODUCTION

Most natural resource systems used by multiple individuals can be classified as common-pool resources. Conventional economic theory predicts that when agents have free access to a common-pool resource they will consume ecosystem services to the point where private costs equal the benefits, whereas externalities are imposed on the rest of the community. This can lead to the well-known tragedy of the commons (Hardin, 1968). Many laboratory experiments have been performed to study this phenomenon. Even in the simplest case of these experiments, without communication between the participants, anomalies were found that are not in line with conventional economic theory, which is non-cooperative game theory (Ostrom *et al.*, 1994).

Conventional theory predicts that players in a non-cooperative game will follow a Nash equilibrium. In none of the reported experiments on common-pool resources was such a Nash equilibrium observed. Furthermore, the total consumption of the common-pool resource fluctuates in time (Ostrom *et al.*, 1994).

Many studies have focused on this phenomenon. Dudley (1993) analysed individual data from Indiana University on resource-use and classified most of the participants according to the strategy they used, such as non-cooperative and cooperative behavior. Deadman (1997) developed a simulation model based on artificial intelligence, which reproduced patterns similar as observed in the laboratory experiments. Casari and Plott (2000) report laboratory experiments that are consistent with their proposed analytical model that includes heterogeneity of social orientation among the agents.

The original aim of this chapter was to explain the phenomenon in question as originating from heterogeneity in Social Value Orientations (SVO's) of the participants, using the *consumat* approach, a multi-agent simulation approach from social psychology (Jager, 2000; Jager *et al.*, 2002) as a research tool. Indeed differences in SVO explained the observed aggregated data. However, when we got hold of individual data, our model was not able to explain individual patterns as observed in the various laboratory experiments. Therefore, we broaden this chapter, and use the common-pool resource experimental data for a discussion on the validation of multi-agent models by using laboratory data.

The chapter is built up as follows. In the next section the original common-pool resource experiments are discussed, as well as relevant other studies that tried to explain the observations. In Section 6.3 we discuss Social Value Orientations, then we provide a brief overview of theories on decision making and introduce the *consumat* approach. The *consumat* implementation of the common-pool resource problem is discussed in Section 6.6, and the experiments with the model in Section 6.7. Section 6.8 discusses the problems which have arisen during the simulation experiments and propose new additional laboratory experiments. The chapter closes with a discussion on the relation between laboratory research and multi-agent modeling.

6.2 COMMON-POOL RESOURCE EXPERIMENTS

Ostrom *et al.* (1994) developed a series of laboratory experiments in an effort to understand the degree to which predictions about individual and group behavior, derived from non-cooperative game theory are supported by empirical evidence. The baseline experiments as described in Ostrom *et al.* (1994) form the basis for the simulation experiments in this chapter. In these baseline experiments, eight subjects were presented with a situation in which they had to choose to invest tokens in two alternatives, or markets. Market one, a safe alternative, provides a constant rate of return on investments. Market two provides returns that vary in relation to the total group investment and the investment of the individual. Market two is the common-pool resource. Ostrom *et al.* (1994) performed experiments for two amounts of tokens, namely ten tokens and 25 tokens. The eight subjects had information about the functional relationship of the two markets, derived the aggregated level of token investments after each round, and were not allowed to communicate with other subjects.

Each individual i has a number of e tokens each round. An amount of x_i is invested in market two, which functions as a collective resource, and an amount of $e - x_i$ is invested in market one. The payoff function is:

$$u_i(\mathbf{x}) = \begin{cases} 0.05 \cdot e & \text{if } x_i = 0 \\ 0.05 \cdot (e - x_i) + (x_i / \sum x_i) \cdot F(\sum x_i) & \text{if } x_i > 0 \end{cases}$$

where

$$F(\sum x_i) = (23 \cdot \sum_{i=1}^8 x_i - 0.25 \cdot (\sum_{i=1}^8 x_i)^2) / 100$$

According to this formula, the payoff of someone investing all his ten tokens in market one ($x_i = 0$) is $0.05 \cdot e$, thus 0.5 tokens. Investing a part or all of the tokens in market two ($x_i > 0$) yields an outcome that depends on the investments of the other players. For example, when player i invests all his ten tokens in market two, whereas all the other players invest in market one, player i will receive 2.05 tokens, considerably more than what market one yielded. However, when all players invest ten tokens in market two, the yields are -0.2 tokens, thus implying a loss. In the case when the players have 25 tokens at their disposal, they may experience even larger losses when they all invest a lot in market two.

If the players behave according to the non-cooperative game theory, they would derive the Nash equilibrium where each player maximizes payoff given the strategies chosen by the other players. The Nash equilibrium for both $e = 10$ and $e = 25$ is equal to eight tokens invested in market two.

If the group act as one agent, they should fully cooperate to maximize their outcomes. In this case their total investment in market two should be 36 tokens, which implies an investment level between four and five tokens per individual. In order for perfectly rational individuals to make themselves better off by achieving the optimal group returns, each individual would have to substantially cut back his or her investment. However, when all other players restrain their investments in market two, it would be very tempting for a player to invest more in market two as this would increase his outcomes significantly. Once individuals give in to this temptation, other individuals may follow, and the group may reach a Nash equilibrium, thereby lowering the group performance.

Looking at the outcomes of different group strategies, we see that a fully cooperative strategy yields the highest outcomes for the group as a whole (Table 6.1). The Nash equilibrium yields higher outcomes than investing exclusively in market one.

Table 6.1: Individual earnings given different group strategies. e can be ten or 25 (when $e=25$ dividing $u(x)$ is divided by 2)

Earnings/subject	$e = 10$	$e = 25$
Group maximum ($\Sigma x_i = 36$)	\$0.91	\$0.83
Nash equilibrium ($x_i = 8$)	\$0.66	\$0.70
No investment in CPR ($x_i = 0$)	\$0.50	\$0.63

The performance of a group can be measured as the extra earnings by investing in market two as a percentage of the maximum extra returns. The Nash equilibrium level of investment leads to a level of 39 $(=(0.66 - 0.50)/(0.91 - 0.50))$ percent of this maximum.

Many laboratory experiments have been conducted on finitely repeated CPR dilemmas. The central question was to what degree the empirical data support the outcomes as hypothesized by non-cooperative game theory for finitely repeated, complete information games (Ostrom *et al.*, 1994). At the aggregate level the results seem to converge to a Nash equilibrium. But instead of monotonic convergence to the predicted equilibrium a different repetitive pattern was observed. The net yield drops toward zero and then rebounds as subjects reduce the level of investment in the common-pool resource. There is a difference in observed aggregate behavior in the two levels of endowment (ten and 25 tokens) (Figure 6.1). In the low-endowment setting, the aggregate results remain close to the predicted Nash equilibrium. In the high endowment setting, however, the aggregate behavior is far from the Nash equilibrium during the first part of the experiment, but begins to approach Nash in later rounds. More interestingly, at the individual decision level, no behavior is found which is consistent with the Nash equilibrium.

When individuals do not use a decision strategy that yields a Nash equilibrium (abbreviated as Nash strategy), the question becomes what strategy do they use in deciding how many tokens to invest in which market, and to what extent is this strategy biased. Dudley (1993) tested how a player should have behaved, given he had perfect foresight, to maximize his/her outcomes given the fixed behavior of seven other players. Dudley derived the fixed behavior from the other seven players from 27 experiments, including 216 subjects. The optimal decision behavior of the artificial players has been calculated for each of the experiments. The strategies behind this decision behavior were categorized as (1) non-cooperative Nash strategy, (2) the cooperative strategy, (3) the average strategy that aims at generating the same returns in both markets, and (4) the remaining non-classified strategies. In the ten-token experiment 24% of the subjects followed a Nash strategy, no player followed the cooperative strategy, 12% played the average strategy and 64%

of the subjects could not be classified. The results for the 25-token experiments were somewhat different. Here, 30% followed a Nash strategy, 3% a cooperative strategy, 66% played the average strategy, and only 1% could not be classified.

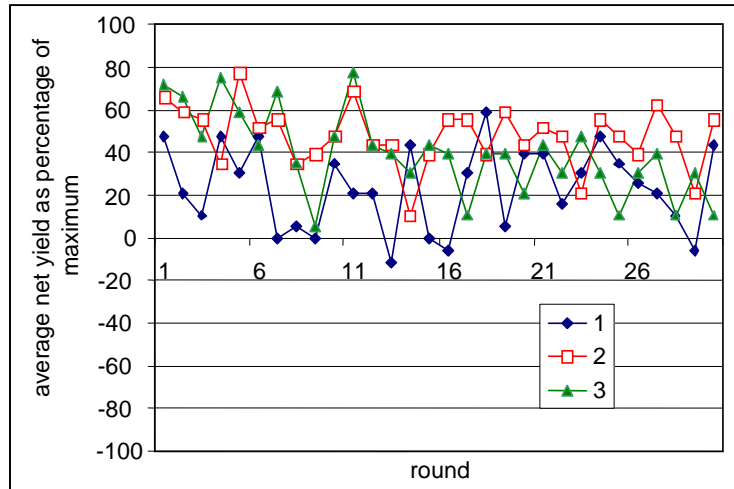


Figure 6.1a: Outcomes of three laboratory experiments with ten tokens. Nash equilibrium generates a yield of 39.5%




Figure 6.1b: Outcomes of three laboratory experiments with 25 tokens. Nash equilibrium generates a yield of 39.5%

The assumption of perfect foresight is of course not very realistic. Therefore Dudley performed experiments equipping the artificial agent with forecasts derived from real people. The forecasts were measured in 128 experiments with real people. Dudley found that in the case of ten-token experiments 29% of the subjects follow a Nash strategy, no subjects follow the cooperative strategy, 11% follow the average strategy, and 60% of the subjects could not be classified. In the case of 25-token experiments 13% followed a Nash strategy, 4% a cooperative strategy, 28% an average strategy and 34% could not be classified. The hypothesis that the reported forecasts of the subjects are unbiased could be rejected on the basis of these results. Dudley argued that there is strong evidence supporting that subjects use adaptive learning in their forecasts.

Deadman (1997, 1999) uses intelligent software agents to simulate the laboratory outcomes as reported by Ostrom *et al.* (1994). From the perspective of bounded rationality a limited set of rules of thumb have been formalized in software agents. These rules are based on questionnaires submitted by individual participants during the baseline common-pool resource experiments run by Ostrom *et al.* (1994). The outcomes of the questionnaires revealed that many participants followed a rule of thumb that stated: 'Invest more in market two whenever the rate of return is greater than \$0.05 per token.' Whenever the per-token rate of return for a market two investment exceeded that of market one, participants increased their market two investment. When the rate of return fell below that of market one, participants invested more tokens in market one. In the ten-token endowment experiments, the authors found a tendency for participants to invest all their tokens in market two whenever the rate of return exceeded that of market one. Many investors followed this strategy, despite the fact that the full information allowed the participants to follow a more optimal (Nash) strategy (Ostrom *et al.*, 1994). Formalizing such rules of thumb in agents, and allowing agents to switch their strategy during the simulation, Deadman was able to replicate the cyclic patterns as shown in Figure 6.1.

However, if we look at the individual data of the players, we may conclude that there exists a considerable diversity in how much people invest in market two. Figure 6.2 shows some individual harvest patterns of subjects in a ten-token experiment (experiment 36 from Ostrom *et al.*, 1994).




Figure 6.2: Some exemplary individual data from a ten token experiment (data by kind permission of J. Walker)

What can be observed here is that the subject represented by series 2 shows large variations in his/her investing. Series 6 on the contrary shows a subject that always invests ten tokens. Series 8 shows a subject that continuously changes his/her investment in market two, but the changes usually involve one token. Apparently the subjects behave quite differently in the same experiment, and hence heterogeneity amongst the subjects may play an important role in understanding the aggregated data.

Casari and Plott (2000) suggest that there is heterogeneity among the agents regarding how they want to interact with other agents. They distinguish altruistic, self-interested and spiteful agents in their analytical model. Spiteful agents derive utility from decreasing the earnings of others. They replicate the Ostrom *et al.* (1994) baseline experiments, and added additional experiments to test different sanctioning regimes. Their experimental data are consistent with their analytical model and they conclude that heterogeneity of social orientation explains why a Nash equilibrium is not reached. However, Casari and Plott (2000) did not test social orientations of the participants directly by surveys.

In sum, previous studies argue that heterogeneity of strategies among the participants is the main cause of the observed aggregate phenomena. Especially differences in the way people want to interact with each other are assumed to explain the cyclic behavior of the investments in the common-pool resource. In the next section we discuss in more depth Social Value Orientation (SVO) as a formalization of this interaction.

6.3 SOCIAL VALUE ORIENTATION (SVO) AND THE VALUATION OF OUTCOMES

In the research on social dilemmas, much attention has been given to how the Social Value Orientation (SVO) of persons affects their harvesting behavior. The SVO of a person is defined as the preferences one has for particular distribution of outcomes for oneself and others. Because in common-pool resources the choice behavior of people may depend on the preferences they have for a certain distribution of outcomes, this perspective may be relevant for understanding differences between people regarding their behavior in a common-pool resource.

The SVO can be measured by using a task where people have to make a choice between two distributions of outcomes for oneself and the others. People are confronted with a series of different choice dilemmas as formulated in, for example, *The Ring Measure of Social Values* (Liebrand, 1984). Graphically depicting the preferences for outcome preferences on a x -axis (own outcomes) and y -axis (other outcomes) results in a circumplex of SVO (Figure 6.3, see also Wiggins, 1980). In Figure 6.3 the most commonly researched prototypical SVOs are being depicted. The SVO of a person can be expressed by the angle of the vector, and the length of the vector can express the coherence of one's SVO. A person who makes choices that are perfectly consistent with his/her SVO is indicated with a vector that touches the outer circle. The less consistent a person is in his choices, the shorter the vector gets. In Figure 6.3 some vectors are denoted with grey arrows.

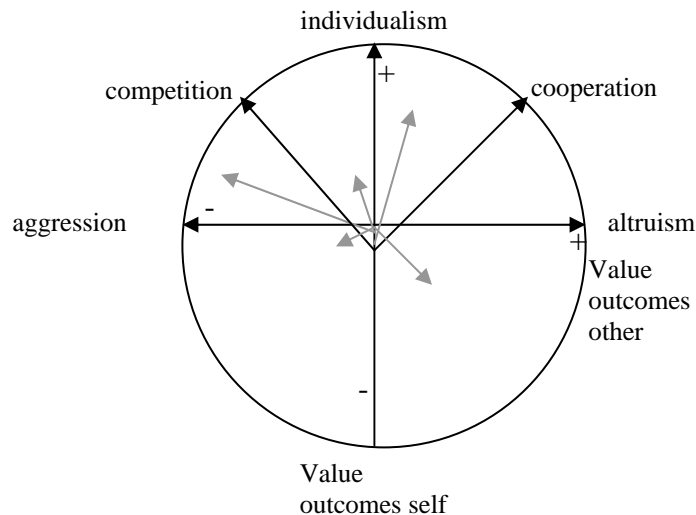


Figure 6.3: Social value orientations

For matters of simplicity we will not formalize different vector lengths in our model, but rather focus on the most common prototypical SVOs. Whereas in principle eight prototypical SVOs can be imagined, the three empirically most frequently occurring orientations have been the topic of theorizing and empirical study. These three orientations are (1) cooperation, aimed at maximizing the outcomes of self and others (2) individualism, aimed at maximizing one's own outcomes at the neglect of the others' outcomes, and (3) competition, aimed at maximizing one's own outcomes in comparison to others' outcomes. In one experiment Van Lange (1999) reports that the prosocial or cooperative is the most frequently observed orientation (57%), the individualistic orientation is observed less frequently (36%), and the competitive orientation the least frequently observed orientation (7%). Two other orientations that have been investigated are (4) altruism, aimed at maximizing the others outcomes at the neglect of own outcomes, and (5) aggression, which is aimed at minimizing the others' outcomes at the neglect of the own outcomes. An important conclusion from research on SVOs is that not all people are *a priori* inclined to value only their own outcomes, or to see the pursuit of self interest as rational (Van Lange *et al*, 1992, p.17). Including the outcomes of others in some way in the outcome matrix leads to a transformed outcome matrix, which may lead to other optimal solutions than choices on the basis of pure self-interest (Kelley and Thibaut, 1978; Kuhlman and Marshello, 1975; McClintock and Liebrand, 1988). The SVO people have is thus an important behavior determining factor in social dilemmas (Messick and McClintock, 1968; McClintock, 1978).

This SVO appears to be an important factor in describing heterogeneity between people in the management of a common-pool resource. However, the SVO of a person does not neatly describe the harvesting behavior of a person. For example, a cooperative person is likely to maximize the joint outcomes, but when the other players are systematically exploiting him, the chances are high that this person abandons a joint maximization strategy, in favor of a punishing strategy. Rather we conceive the SVO of a person as a factor determining his/her satisfaction level with a certain distribution of outcomes. This satisfaction level may affect the decision-making process of the person, which may continually change as a consequence of the own and other players' behavior. People may spend more or less cognitive effort in their decision making, and may use more or less information regarding the behavior of others in this process. First of all this introduces a heterogeneity *between* people, as one person may be inclined more towards extensive elaboration than another. On top of that, heterogeneity can also be observed *within* people, as they usually change their decision-making strategy when repeatedly making the same type of decision, as is the case in a common-pool

resource game. In the next section we therefore provide a psychological perspective on human decision making.

6.4 DECISION MAKING

People make decisions all the time. During a day you decide for example what to eat for breakfast, what alternative route to follow to your work when the usual route unexpectedly happens to be closed, what to wear at that reception tonight and how to set up this laboratory experiment with 120 respondents. Because we make thousands of decisions each day, and our cognitive capacity is limited, people have developed very smart strategies to allocate their limited cognitive capacity over this multitude of decisions. Instead of maximizing the outcomes of behavior, as a prototypical *Homo economicus* would do, humans also optimize the decision-making costs (cognitive effort) that are associated with making a choice. As such, people use an abundance of decision rules or so-called heuristics in their daily lives. For example, people may habitually eat cereals for breakfast, follow the traffic stream (imitating) when confronted with the closed route, think of what other people of the same age will wear at that occasion (norm), and contemplate extensively on the advantages and disadvantages of different experimental set-ups. This allocation of resource works very efficiently, and allows us for example to deliberate about an experiment whilst (almost automatically) driving a car.

The critical question is of course how people decide on which decision strategy to employ in a given situation. The work of Simon (1955, 1959, 1976) on bounded rationality offers a perspective on why habits and complying with a norm may be a rational thing to do. The essential argument is that humans optimize the full process of decision making (*procedural rationality*), not only the outcomes (*substantive rationality*) (Simon, 1976). This holds that consumers may decide that a certain choice problem is not worth investing a lot of cognitive effort in (e.g., deciding on your breakfast), whereas another choice problem requires more cognitive attention (e.g., setting up an experiment). The less important a decision problem is, the less cognitive energy one is willing to invest in the decision, and, hence, the simpler the decision heuristic that will be employed.

Often people use their own previous experiences in a heuristic. For example, when you are satisfied with a certain type of cereal for breakfast, you may not waste any cognitive energy on deciding what to eat, but rather grab for the cereal in an automated way (which may be very convenient early in the morning). However, people may also employ the behavior and experiences of other people in their decision making. For example, when

confronted with an unexpected road obstruction, the car driver may use the behavior of the other drivers as a clue regarding which direction to proceed. Instead of deliberating about all the possible alternative routes, the driver may assume that the predecessors have thought about it, and hence he/she may follow without giving too much thought on the issue. Also when thinking about which clothes to wear on that occasion, you use the experience of other people indirectly. Remembering the negative remarks people made on the too-casual outfit of a colleague on a previous similar occasion, you may decide to wear a suit, despite your personal preference to wear a more casual outfit. It appears that the cognitive strategies that people employ can be organized on two dimensions: (1) the amount of cognitive effort that is involved, and (2), the individual versus social focus of information gathering.

Regarding the first dimension, amount of cognitive effort, the basic idea is that people allocate their limited cognitive capacity over various decision problems they face so as to maximize their utility. When one is frequently being confronted with the same or similar decision tasks and the previous behavior yielded satisfactory outcomes, it is a good strategy to economize on cognitive effort by using simple heuristics or a habitual script in making the decision. This allows for allocating most of the cognitive capacity to decision problems that require more attention in order to find a satisfactory solution, such as non-routine decisions with important consequences. Because cognitive processing takes time, using simple decision heuristics will save time. This explains why people tend to use simpler decision heuristics when under time pressure (e.g., Smith *et al.*, 1982; Wallsten and Barton, 1982; Wright, 1974; Ben Zur and Breznitz, 1981). Also when the decision is less important (in terms of consequences) the decision maker may use a simpler heuristic instead of using all information available (e.g., Tversky, 1969, 1972). The simplest type of behavior in terms of cognitive effort refers to preconscious habits (Fiske and Taylor, 1991), i.e. behavior which bears a reflex-like character.

Regarding the second dimension, the individual versus social focus of information gathering, uncertainty is the key-factor that determines the focus of the information search process. When people are certain of themselves, they usually refer to their own previous experiences when making a deliberate or automated decision. When uncertain, people may use the experiences of other people to come to a decision in a cognitive efficient manner. Especially the behavior of other people with about similar abilities may provide a useful clue in the decision-making process. Simple imitation may be an economical way of allocating cognitive capacity to a decision. The *Social Learning Theory* (Bandura, 1977, 1986) states that seeing someone else's behavior being reinforced may affects one's own behavior. This imitating however requires more cognitive effort than a simple habit, because one should

be *attentive* to the behavior of someone else, *understand* and *remember* that behavior, be able to *reproduce* that behavior, and experience *reinforcement* after performing the behavior yourself (Bandura, 1977). Following simple norms can also be considered as a social focused heuristic that requires relative little cognitive effort. However, according to the *Theory of Planned Behavior* (Ajzen, 1985, 1988, 1991), the *subjective norm*, may require more cognitive effort in making a decision. The subjective norm here refers to a person's perception of the opinion of others about him/her performing the relevant behavior. The subjective norm is proposed as a function of one's beliefs that referents think whether the person should or should not perform the behavior (called the injunctive norm), weighted by the motivation to comply with those referents. Social comparison (Festinger, 1954) is a key process here, involving the fact that people consciously compare their opinions and abilities with those of other people. These comparisons follow dimensions such as the possession of material goods, financial means, status, principles, attitudes and skills. With respect to opinions, people have a drive to roughly conform to others. With respect to abilities, people have a drive to be (somewhat) superior to others. Becoming aware of a subjective (social) norm would involve an assessment of relevant others and an appreciation of their behavioral intentions, which involves considerable more cognitive effort than simple imitation or obedience to a simple norm.

In organizing the various decision strategies that people employ, we find it instructive to use the two dimensions as graphically depicted in Figure 6.4.

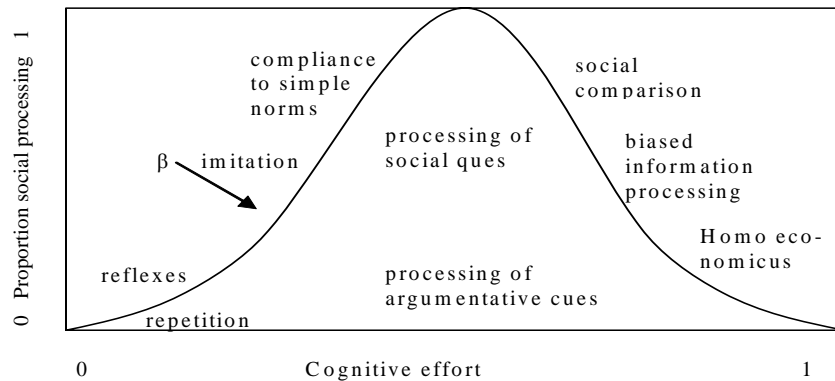


Figure 6.4: Different decision processes organized along the dimensions of cognitive effort and use of social information. β indicates the maximum level of social processing that allows for procedural optimality

Figure 6.4 shows that decision strategies that hardly require any cognitive effort (reflexes), or require very much cognitive effort (the prototypical *Homo economicus*) do not use social information. Strategies that require an intermediate cognitive effort may use both social and non-social information. Here, uncertainty is a key factor that determines the degree to which social information is being used in the decision-making process.

In organizing the decision strategies along these two dimensions, a perspective emerges regarding how people differ regarding their abilities and motivations to invest cognitive effort in a decision, and to what extent they use social information. Hence, this perspective contributes to the understanding of heterogeneity between people as regards their decision making. For example, some people may be more inclined towards using social information, and other people may have a larger cognitive ability, making it easier to invest cognitive effort in the decision-making process. On top of that, understanding how these abilities and motivations may change in a repeated decision-making situation provides a perspective on how people switch between decision strategies over time, and hence contributes to the understanding of heterogeneity within people. For example, when people become more uncertain, they will tend to use more social information in their decision making, and when people are not satisfied, they may be inclined to spend more cognitive effort in their decision-making process as to find a better behavioral opportunity.

To test hypotheses regarding the effects of heterogeneity in the decision-making process on collective outcomes we developed the consumat approach. This approach involves a multi-agent simulation model of decision-making processes. In the next section, we briefly elaborate on the consumat approach.

6.5 THE CONSUMAT APPROACH

The consumat approach is based on a comprehensive conceptual model of choice and decision-making behavior (Jager *et al.*, 1999; Jager, 2000). As such it tries to offer a more psychological based meta-theory of human decision making than the frequently used ‘rational actor’ approach. The consumat approach considers basic human needs and uncertainty as the driving factors behind the human decision-making process.

Based on this conceptual model, a multi-agent simulation model has been developed, in which the agents are called ‘consumats’. The driving forces at the collective (macro-) and the individual (micro-) level determine the environmental setting for consumat behavior. This may be represented by a

collective resource. The individual level refers to the consumats: they are equipped with needs which may be more or less satisfied, they are confronted with opportunities to consume, and they have various abilities to consume opportunities. Furthermore, consumats have a certain degree of uncertainty, depending on the difference between expected and actual outcomes of their behavior.

The various decision-making processes as organized along the two dimensions of Figure 6.4 are reduced to four decision rules: deliberation, social comparison, repetition and imitation. This simplification serves to keep the simulation model simple, and the results transparent for interpretation. Which of these decision rules consumats use at a given moment in time depends on their level of need satisfaction and degree of uncertainty. Consumats having a low level of need satisfaction and a low degree of uncertainty are assumed to *deliberate*, that is, to determine the consequences of all possible decisions given a fixed time-horizon in order to maximize their level of need satisfaction. Consumats having a low level of need satisfaction and a high degree of uncertainty are assumed to engage in *social comparison*. This implies comparison of own previous behavior with the previous behavior of consumats having roughly similar abilities, and selecting that behavior which yields a maximal level of need satisfaction. When consumats have a high level of need satisfaction, but also a high level of uncertainty, they will *imitate* the behavior of other similar consumats. Finally, consumats having a high level of need satisfaction and a low level of uncertainty simply *repeat* their previous behavior. When consumats engage in reasoned behavior (deliberation and social comparison) they will update the information in their mental map, which serves as a memory to store information on abilities, opportunities, and characteristics of other agents.

After the consumption of opportunities, a new level of need satisfaction will be derived, and changes will occur regarding consumats' abilities, opportunities and uncertainty. Moreover, the environment the consumats behave in, for example, a collective resource, will change as a consequence of their behavior, thereby affecting the behavior in subsequent time steps.

6.6 A CONSUMAT MODEL FOR COMMON-POOL RESOURCES

To study how heterogeneity in SVO and decision rules affects the behavior in a common-pool resource in a very controlled setting, we decided to formalize the consumat approach for the common-pool resource paradigm as sketched in the introduction. Three needs are formalized: a personal need, a social need and a need for exploration.

The personal need relates to subsistence, and is assumed to be equal to $u(x)$ the returns from investment. This personal need is equal for all agents ($N_{i,i} = u_i(x)$).

The social need relates to how an agent wants to relate to other people, and is a formalization of the SVO, and hence comprises individualistic, competitive and cooperative preferences for outcome distributions. The agents thus differ regarding their social need.

When an agent is purely individualistic then the social need satisfaction is equal to $u(x)$.

When an agent is competitive the social need satisfaction takes into account the relative returns compared to the average returns. The higher the relative returns, the higher the social need satisfaction.

$$N_{s,i} = 1 - \exp(-c \cdot u_i(\mathbf{x}) / \sum_{j=1}^8 \frac{1}{8} \cdot u_j(\mathbf{x})) \quad (6.1)$$

The social need satisfaction of a cooperative agent is higher the closer the returns come to the cooperative optimum. This cooperative optimum is measured in difference with the cooperative amount of tokens, although the cooperative amount of returns could also have been used.

$$N_{s,i} = \exp(-c_2 * (\sum_{j=1}^8 x_j - 36)^2) \quad (6.2)$$

Besides the personal and social need, we also formalized a need for exploration. The exploration need is a combination of understanding, creation and freedom needs. Whereas exploration is often described as specific, and may relate to, for example, the search for food, Berlyne (1966) also proposed a diversify type of exploration, motivated by the need to know. Large numbers of experiments by Berlyne and others led to further notions to be linked to exploration, such as the novelty of a situation (e.g., Hutt, 1970; see also Gibson, 1988). This type of exploration appears to describe how people learn to understand how a resource system (the resource and the other players) reacts to certain actions. Hence exploration serves to increase the understanding of the system. This exploration need is conceived to be less satisfied the more stable the outcomes are, because in such a situation nothing new is being learned about the system. The exploration need $N_{E,i}$ is being formalized as a standard deviation in outcomes over the last n ($=5$) rounds:

$$N_{E,i} = \sqrt{\frac{\sum_{j=t-n}^t \{u_i(\mathbf{x}(j)) - u_i(\bar{\mathbf{x}})\}^2}{n-1}} \quad (6.3)$$

where $u_i(\bar{\mathbf{x}})$ is the average return for n rounds.

When consumats are dissatisfied, they deliberate about all possible courses of action. When all three needs are formalized, the investment opportunities are evaluated on expected outcomes for the self, the expected outcomes for the other and the contribution to the standard deviation in the outcomes. We assume that dissatisfaction is not absolute: it is more a rule to decide when to employ cognitive energy. Exploration is less important than subsistence, but dissatisfaction may lead to the same level of cognitive effort. Stated differently, if you don't have serious problems to think about, you'll think about less serious problems.

The total need satisfaction is defined as the weighted sum of needs, where $\sum_i \beta_i = 1$:

$$N_i = \beta_1 \cdot N_{S,i} + \beta_2 \cdot N_{I,i} + \beta_3 \cdot N_{E,i} \quad (6.4)$$

Uncertainty, U , is defined as the standard deviation of the individual returns during the last five rounds.

$$U_i = \sqrt{\frac{\sum_{j=t-5}^t \{u_i(\mathbf{x}(j)) - u_i(\bar{\mathbf{x}})\}^2}{4}} \quad (6.5)$$

The agents are equipped with a memory, which is being used in calculating the expected returns from the decisions. In this calculation the agent makes use of an expected aggregate level of tokens. To estimate how many tokens other agents are expected to invest, a neural network is used. A neural network is an algorithm that resembles the way in which the brain works. It is composed of *nodes* representing physiological neurons, and *weights*, which are connections of differing strength between two nodes. Some of the neurons receive their input from the environment and some others give back their output to the environment. We use a single layer neural network, which is a method to describe changes in the weights based on physiological principles, and is described by the following equation:

$$EY_{t,i} = w_{t,i,0} + \sum_{k=1}^{ks} w_{t,i,k} \cdot y_{t,k} \quad (6.6)$$

where $EY_{t,i}$ is the level of tokens of the other seven agents, $y_{t,k}$ are the inputs of the neural network, the observed total investments during the previous ks ($=3$) rounds. Finally, the inputs are weighted by $w_{t,i,k}$.

A neural network is trained when new information about the input values is used to update the weights (w). The widely used Widrow-Hoff delta rule is used to train the neural network during the simulation (Mehrotra *et al.*, 1997). This simple neural network simply weights the observed token investments of the other agents during the last few rounds to estimate the total investments in the next round.

$$\Delta w_{t,i,j} = \eta \cdot \delta_{t,i} \cdot \frac{y_i}{\|y_i\|} \quad (6.7)$$

where

$$\delta_{t,i} = Y_{t,i} - w_{t,i,0} - \sum_{k=1}^{ks} w_{t,i,k} \cdot y_{t,k} \quad (6.8)$$

where $\delta_{t,i}$ is the difference between the observed values and the expected value, and $Y_{t,i}$ is the observed token investments by the other agents. The delta rule updates the expectations according to the observed errors. The rate of updating depends on the value of η , which is suggested to lie between 0.1 and 1 (Gallant, 1993). We will assume η to be 0.5 reflecting relative adaptive agents.

In making an investment decision, the agent employs one of the four cognitive processes as described in the previous section, depending on the (combined) level of need satisfaction and uncertainty and the thresholds U_{\max} determining when an agent is uncertain, and N_{\min} determining when the agent is satisfied. When the agent is dissatisfied and certain, it will engage in deliberation. Deliberation has been formalized as calculating the level of investment that maximizes the agent's expected need satisfaction. We assume that if consumers deliberate they have full information and understanding of the problem, and are able to calculate the Nash equilibrium.

When the agent is satisfied and certain, it will engage in repetition, and hence invest the same quantity as in the previous round. When an agent is dissatisfied and uncertain, it will engage in social comparison. This implies comparing the average investment of the previous round with the own investment of the previous round, and choosing that investment with the highest expected level of need satisfaction. Finally, when the agent is

satisfied and uncertain, it will engage in imitation. This implies copying the average investment of the previous round.

The threshold values for aspiration level (N_{\min}) and uncertainty tolerance (U_{\max}) have an empirical counterpart in the Intellect factor and the Emotional Stability factor of the Big five personality structure (Goldberg, 1990, see also Janssen and Jager, 2001).

In the next section we discuss the results of simulation experiments in which the agents are confronted with the same common-pool resource as used in empirical studies by Ostrom *et al.* (1994).

6.7 RESULTS

In this section we discuss the results for a series of experiments. Tables 6.2 and 6.3 show statistics of six original laboratory experiments as described in Ostrom *et al.* (1994). The average investments are among the Nash equilibrium of eight tokens, although the experiments with 25 tokens are systematically above eight tokens. The variability among individual investments is higher in the 25-token experiments, showing that there is more behavioral change in this condition. Furthermore, variability among rounds is generally larger in the 25-token experiments. Furthermore, the higher the average investment, the higher the interround variability. These statistics are used to compare the simulation experiments with the laboratory experiments.

Table 6.2: Experimental values of three experiments with ten tokens. The first column of numbers denotes the average token investment in the common-pool resource over 30 rounds. For each round the standard deviation of the eight individual investments is calculated. The second column contains the average standard deviation over 30 rounds and is a measure of variability within each round. The last column shows the total absolute changes in the total investments, which is a measure of investment variability among 30 rounds.

Experiment	\bar{x}	$\overline{stdev}(x)$	$cum(\bar{x} - \bar{x}_{t-1})$
1	8.46	1.37	139
2	7.72	0.60	125
3	7.94	0.88	137

Table 6.3: Experimental values of three experiments with 25 tokens. The first column of numbers denotes the average token investment in the common-pool resource over 20 rounds. For each round the standard deviation of the eight individual investments are calculated. The second column contains the average standard deviation over 20 rounds and is a measure of variability within each round. The last column shows the total absolute changes in the total investments, which is a measure of investment variability among 20 rounds.

Experiment	\bar{x}	$\overline{stdev}(x)$	$cum(\bar{x} - \bar{x}_{t-1})$
1	8.64	3.36	283
2	9.22	2.89	316
3	8.54	2.15	108

Experimenting with cognitive processing, needs and social value orientation

A series of simulation experiments has been performed, and by varying the characteristics of the agents we created different conditions. A first characteristic we varied was the cognitive processing the agent could employ. Two conditions were created, namely the *Homo economicus* (HE), which engages exclusively in deliberation, thereby representing the rational agent from standard economic theory, and the *Homo psychologicus* (HP), which could employ all four decision strategies. For the HE conditions the values of N_{\min} and U_{\max} are put on such values that the agents only deliberate. In the HP condition, the values of N_{\min} and U_{\max} are 0.5 and 0.1 respectively, so that all four cognitive processes can be used.

A second characteristic we varied in the experimental design refers to the combination of needs the agents have. Four conditions were created, respectively; (1) agents with only an personal need (noted with P), (2) agents having an personal need and a social need (PS), (3) agents having a personal need and an exploration need (PE), and (4) agents having a personal need, a social need and an exploration need (PES).

The third characteristic we varied was the SVO of the eight agents. In the conditions where the social need is formalized a fixed number of the agents are either cooperative, individualistic or competitive. We analyse all the combinations of these three SVOs.

The initial expectations the agents have regarding the total investments are set in line with the equilibrium outcomes of the SVO of the agent. In case of a Nash equilibrium (competitive or individualistic SVO) this expectation is

set at 64 tokens, and in case of a cooperative equilibrium (cooperative SVO) this expectation is set at 36 tokens.

Table 6.4 summarizes the statistics of the 16 conditions. In ten conditions we observe that the theoretical Nash equilibrium ($x = 8$) is derived, namely in the conditions where only the personal need is taken into account (HE-P, HP-P), or when the social need is weighted but where the agents are all competitive or individualistic.

Inclusion of the need for exploration leads to fluctuations in the token investments, which are more extreme in the case of the *Homo economicus*, since the *Homo psychologicus* can be satisfied with a lower variability.

Table 6.4: Summary of 16 conditions: Simulated values ten tokens, all four cognitive strategies and three social value orientations. The average standard deviation on between agent variation is not depicted since all agents have the same characteristics.

Experiment	\bar{x}	$cum(\bar{x} - \bar{x}_{t-1})$
HE-P	8	0
HP-P	8	0
HE-PS-coop	8.03	448
HE-PS-ind	8	0
HE-PS-comp	8	0
HP-PS-coop	6	0
HP-PS-ind	8	0
HP-PS-comp	8	0
HE-PE	7.93	440
HP-PE	6.3	344
HE-PES-coop	8	424
HE-PES-ind	8	0
HE-PES-comp	8	0
HP-PES-coop	6	216
HP-PES-ind	8	0
HP-PES-comp	8	0

When we assess all combinations of SVO we have to simplify the presentation to keep an overview of how different experimental variations affect the results. Therefore we have developed an indicator I that describes the difference between the simulation results and the empirical results or theoretical expectation.

$$I = \frac{a_1}{3}(x - x_C)^2 + \frac{a_2}{3}(stdev - stdev_C)^2 + \frac{a_3}{3}(cum - cum_C)^2 \quad (6.9)$$

This indicator I is constituted on three relevant outcomes (see Table 6.2 and 6.3). First, the difference between the simulated average investment behavior x and the empirical/theoretical average investment x_C is calculated in $(x - x_C)^2$. The empirical results are the statistics generated by the 24 real subjects in the three reported experiments. The theoretical expectation reflects the Nash equilibrium of eight tokens for these experiments. The second component of the indicator, $(stdev - stdev_C)^2$, reflects the difference between the variability in each round for the simulation and the empirical/theoretical investment. The theoretical (Nash) expectation for $stdev_C$ is 0, as all subjects are expected always to invest eight tokens. For the empirical results we calculate $stdev_C$ using the average indicator values of Tables 6.2 and 6.3. The third component of the indicator $(cum - cum_C)^2$, reflects the difference between the variability over the rounds for the simulation and the empirical/theoretical investment. Again, the theoretical (Nash) expectation of cum_C is zero, as all subjects are expected always to invest eight tokens. For the empirical results we calculate cum_C using the average indicator values of Tables 6.2 and 6.3. In this indicator I we defined the values of a_1 , a_2 and a_3 in such a way that the indicator remained between 0 and 1.

In the following we present the results for a series of experiments. Graphically we present the value of I as a bar. Each bar represents the difference between the statistics of a simulation run and the theoretical expectation (upper, Nash equilibrium) or empirical observations (lower). In the left figures the artificial agents make decisions in line with the *Homo economicus* (HE), while in the figures on the right, the artificial agents make decisions in line with the *Homo psychologicus* (HP). The axes represent the number of individuals and the number of competitors. The number of cooperative agents is equal to eight minus the agents with individualistic or competitive SVO. In Figure 6.5 we present the results for agents having only a need for personal returns and the need for identity.

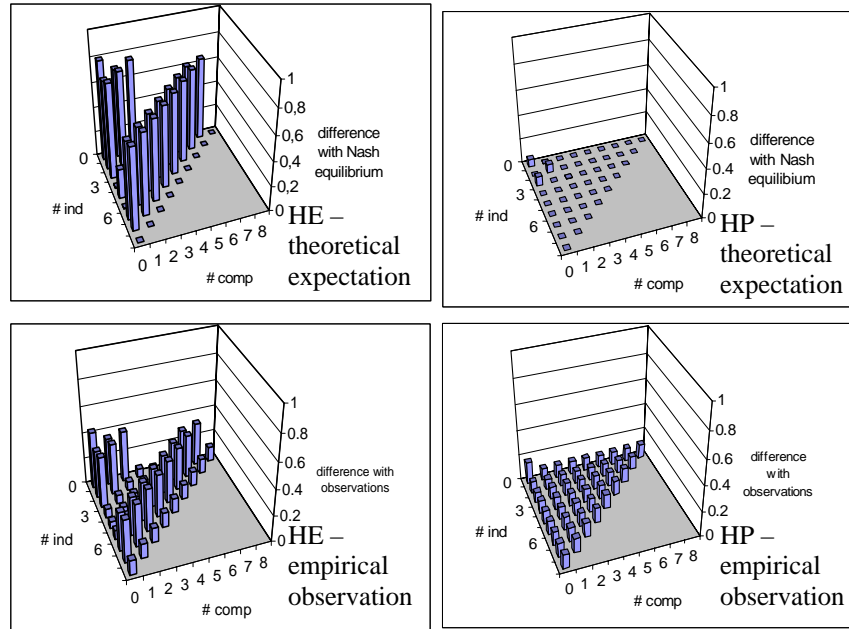


Figure 6.5: Results of all combinations of SVO of experiments with ten tokens with artificial agents without the need for exploration

It can be observed that in the conditions without cooperative agents (on the diagonal line), the theoretical outcomes can be reproduced in both the HE and the HP conditions. When agents act according to the HP, the results in most combinations do not differ much from the theoretical outcome. For the HE we observe that the inclusion of cooperatives causes a large difference with the theoretical outcomes. Only when the number of cooperatives is about four can we see a valley in the results, indicating that the results are closer to the theoretical outcomes. This valley is caused by oscillations that are less extreme than when fewer cooperators take part. Compared with the observed statistics, the HP clearly performs better than HE. It can be seen that more cooperative agents lead towards a better match with the empirical data.

Figure 6.6 shows the results for the same experiments, only now including the need for exploration. The 'valley' in the HP-empirical data figure shows that for the HP the observed statistics can be reproduced closely when about four cooperative agents (close to the proportion of cooperatives as identified by Van Lange, 1999) have been formalized. However, for most combinations of SVO the results are still more close to the theoretical expectations than to the empirical data.

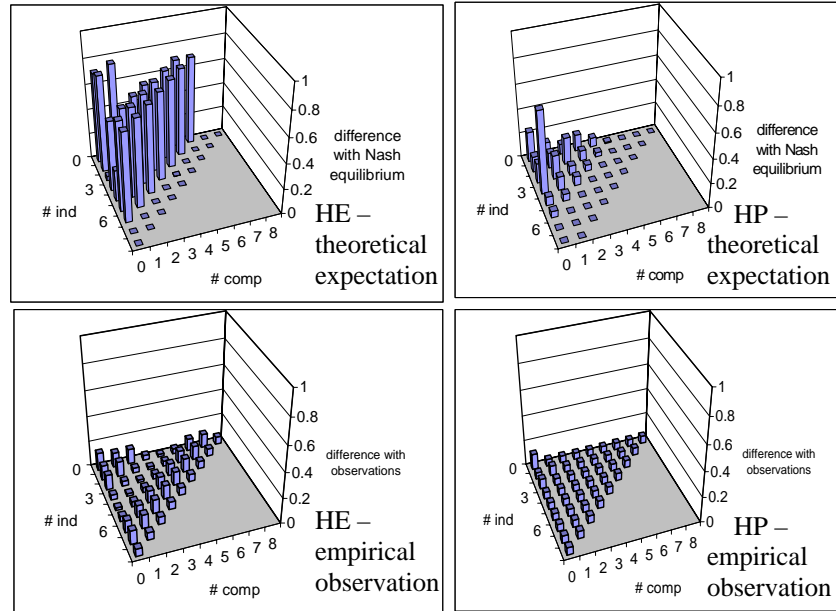


Figure 6.6: Results of all combinations of SVO of experiments with artificial agents with ten tokens, with equal weighting of needs

Experimenting with heterogeneity of needs

In the next series of experiments the agents differ with respect to the weighting of their needs, thereby adding an extra source of heterogeneity between the agents. The values of β_i are drawn from a random distribution. Still the sum of β_i 's is equal to one. The procedure of generating the β 's is as follows. For each agent the value of β is randomly drawn between 0 and 1, but then divided by the sum of β_1 , β_2 and β_3 . For each combination of SVO, we performed 1000 runs and calculated the average value of the indicators.

The results of these experiments are shown in Figure 6.7. Whereas in the previous experiments the agents performed better in matching the theoretical expectations than the empirical data, here we observe the opposite. However, this match on empirical data is less close than for the HP with about four cooperatives in the previous experiment. Remarkable is that the HP now have the best performance when there are only competitive agents. This is mainly due to its good fit on the second component of I that measures the difference between the variability in each round. The theoretical outcomes are now better reproduced when agents have an individualistic SVO, which reduces the variability of the decisions.

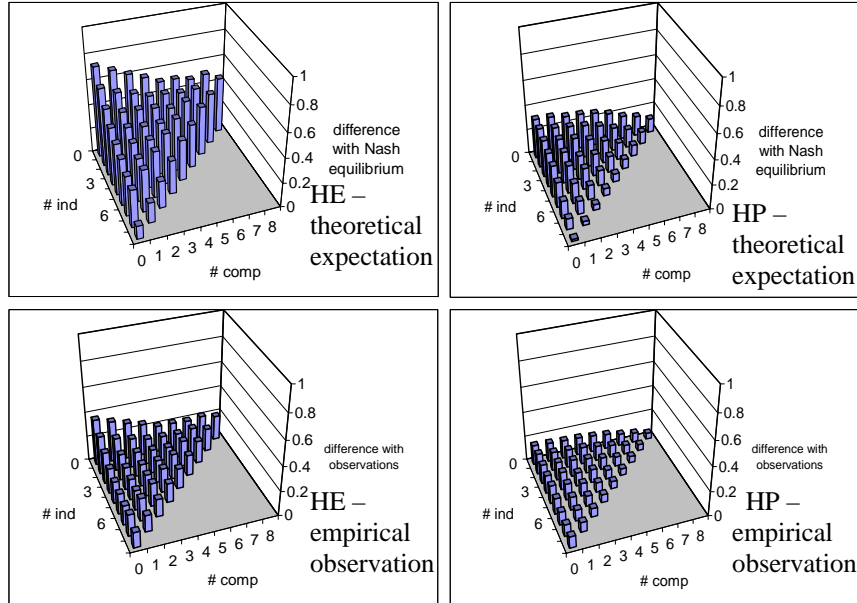


Figure 6.7: Results of all combinations of SVO of experiments with artificial agents with ten tokens, with heterogeneous weighting of needs

The same experiments with heterogeneous weighting of the needs as performed for the ten-token case are repeated for the 25-token case (Figure 6.8). For the HP, the results match the closest to the empirical data when there are only competitors. When there are less competitors, the results match quite well for as long as there is a balance between the number of individualists and cooperatives, as can be seen in the ‘valley’. The theoretical prediction can be replicated the best when agents are individualistic and agents perform like the HE.

These computational experiments show that we are not able to replicate the observations perfectly, but that with different assumptions different types of agent formulations are more suitable to approximate the statistics of the observations. Remarkable is the influence of heterogeneity among the weighting of the needs. If we do not assume heterogeneity, groups that include cooperative agents are better able to approximate the statistics of the observations, while inclusion of heterogeneity leads to the result that no cooperators should be in the group in the attempt to replicate the statistics of the observations. What is clear is that all three needs are of importance to understand the observations, and that agents conforming to the *Homo psychologicus* have a better performance than the *Homo economicus* in approximating the empirical data.

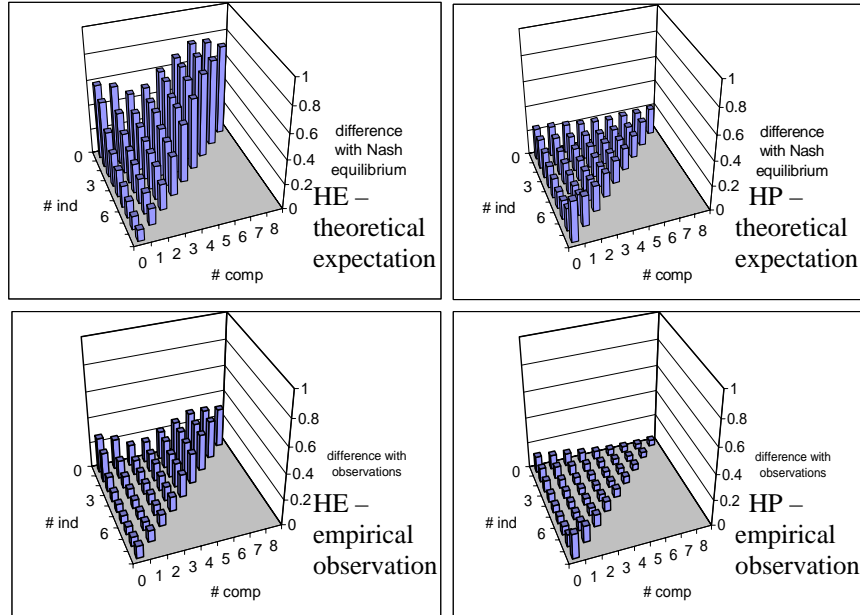


Figure 6.8: Results of all combinations of SVO of experiments with artificial agents with 25 tokens, with heterogeneous weighting of needs

6.8 DISCUSSION AND CONCLUSIONS

In this series of experiments we demonstrated how agents can be equipped with decision rules that are based on psychological theory, and experimented with different settings of these rules as to mimic the individual behavior of real people acting in a resource dilemma. We showed that mimicking the behavior of real people on an individual level requires more psychological realism in the agents than mimicking the aggregate outcomes as in previous experiments. However, we realize that the reproduction of statistics that fit with a limited set of empirical observations does not provide sufficient proof that the simulation model captures the most relevant dynamics that guide the behavior of the subjects in the Ostrom *et al.* (1994) experiments. To do so would require more empirical data. Hence we argue that more experiments are required to unravel the decision-making process of real people. These experiments should address the factors underlying the heterogeneity of decision making. In this chapter, we identified several factors that may contribute to this heterogeneity in decision making. We distinguished the different needs that play a role in the decision-making process, the relative

importance of these needs, the SVO of the decision maker, the cognitive process that people employ when making a decision, the personality characteristics that determine the tendency to use certain cognitive processes more often than others (Intellect, Emotional Stability) and the time-horizon that is taken into account when making a decision. To test all these factors simultaneously in experimental research would yield an enormous task, both for the experimenter and the subjects. This task is especially difficult because we are dealing with complex behavioral dynamics, as different people having different (changing) values on the different factors are interacting for 20 time-steps. Therefore it would be practical if we could test beforehand the relevance of factors and develop hypotheses concerning the effects of varying these factors. We argue that simulation research provides a tool capable of doing so. By performing many experiments and conducting sensitivity analysis one may identify the conditions under which certain behavioral dynamics are more likely to happen, and which factors play a crucial role. Following that, hypothesis and a research design can be formulated for testing these specific effects in empirical experiments.

We claim that it would be most efficient to combine simulation research and empirical research as described above to harvest synergetic benefits. For example, our simulation experiments suggest that the cognitive processes a person is most likely to use may be very important in his/her harvesting behavior. Someone having a low aspiration level is more likely to develop a habit, whereas someone that has a low uncertainty tolerance is more likely to engage in imitation and social comparison. Field experiments could be focused on the question how personality characteristics (aspiration level and uncertainty tolerance) of people are related to their cognitive processing and behavior. The data obtained in this empirical research can subsequently be used to formalize the relation between personality and cognitive processing more validly in agent rules. However, until now experimental and simulation research are rather distinct, despite the fact that they more and more deal with similar research questions. Only three studies are known to the authors that use multi-agent models to formulate hypotheses which are tested with real agents (Duffy, 2001; Pingle and Tesfatsion, 2001; Tobias, in preparation). We think that both multi-agent modeling as well as experimental research can benefit from more interaction between both fields. Whereas both fields appear to be concentrated around the research methodology that is being used, a focus on the research question would benefit this interaction. As a start we discuss five hypotheses based on our computational model, and formulate tests for the experimental research.

Hypotheses:

1. Groups of people having a high aspiration level and a high tolerance for uncertainty are more likely to engage in deliberation. As a consequence they are less sensitive to the imitation effect. This would result in fewer fluctuations in investment decision between agents and between rounds.
2. Groups of people that systematically differ as regards their distribution of SVO also systematically differ as regards their investment levels. We expect that homogeneous groups of cooperatives invest less in market two than relative homogeneous groups of individualists and competitors. Experiments would require a (unobtrusive) pre-measurement of SVO, which will be used to allocate subjects later to experimental settings.
3. Displaying only aggregate outcomes will inhibit the social need in the subjects' decision making, and hence will moderate the SVO effect in comparison to conditions where individual outcome levels are displayed.
4. Giving the subjects a fee for participating in the experiment on the basis of their returns would increase the weighing of the personal need, whereas providing a standard fee for participation would inhibit the importance of the personal need in the decision-making process.
5. Allowing the subjects to explore issues that are not directly relevant for the experiment (e.g., particular information on the other players) would decrease the influence of the exploration need on the harvesting behavior, and hence the results will show less variance.

To be able to compare simulation results involving cognitive processes with experimental data, it would be practical to measure the cognitive process in an experimental setting. The approach of Hine and Gifford (1996) demonstrates that it is possible and necessary to ask people about their decision-making strategies. However, a more structural approach of measuring cognitive processing is proposed by registering the quantity and type of information subjects retrieve from a matrix board (on screen) before making an investment decision. The decision-making process can be tracked by registering the quantity and type of information subjects retrieve from a matrix board (on screen) before making a harvesting decision. The information that can be retrieved relates to the state of the resource and the harvesting behavior of other people (aggregated or individual). Measuring information retrieval in real time allows for discrimination between habitual harvesting and deliberate stable harvesting. A possible avenue for further research on cognitive processes is to use MRI-scan data obtained during the

decision-making task of the subjects. First experiments in which MRI scans are conducted in a cooperation game are conducted by McCabe *et al.* (2001) and show a relation between pre-frontal brain activity and promoting cooperative behavior in the other player.

Several other factors that have been identified as potentially influential in simulation research can be measured in empirical experiments. Moreover, many experiments can be performed in which subjects are being grouped together according to their scores on relevant factors. For example, to test hypothesis 2 we would have to obtain SVO data before assembling subject groups. The more clearly experimental research shows how a certain factor affects the behavior of real people, the better it is to formalize this factor into a valid way in a simulation model. This would give more insight in the dynamics behind the investment behavior and give rise to refining the hypothesis on relevant behavioral processes. We are convinced that the combination of different research tools to address the same research question is a promising way to get a better understanding of the very basic behavioral dynamics that determine our use of collective resources.

ACKNOWLEDGEMENTS

We thank Jim Walker for providing the laboratory data and giving feedback on an earlier manuscript. Furthermore we thank participants of seminars at Indiana University and the University of Amsterdam for providing constructive and helpful feedback.