Title

A Cognitive Model of Agents in a Commons Dilemma

Permalink

https://escholarship.org/uc/item/8v91b76s

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors

Nerb, Josef
Spada, Hans
Ernst, Andreas M.

Publication Date

1997

A Cognitive Model of Agents in a Commons Dilemma

Josef Nerb, Hans Spada and **Andreas M. Ernst**
Department of Psychology, University of Freiburg
D-79085 Freiburg, Germany
{nerb, spada, ernst}@psychologie.uni-freiburg.de

Abstract

KIS (knowledge and intentions in social dilemmas) is a process model of a cognitive-motivational theory of acting in a three person commons dilemma. The model provides an experimental tool to study how ecologically harmful actions evolve in commons problems by having differently parameterized variants of KIS interact with each other and with human subjects. KIS models the application and acquisition of ecological, social, and practical knowledge using a motive-driven decision procedure. To test this model, 42 subjects played a commons dilemma game in an unselfish or greedy social environment. Both environments were realized by pairs of appropriately configured KIS variants. Subjects did not recognize these co-players as being artificial and judged their motives accurately. Subjects' behavior in the unselfish environment was well predicted, however, in the greedy environment subjects based their decisions more on the state of the resource than was expected. To further test the model, we constructed a KIS variant for each subject with respect to the assessed individual motive structure and knowledge. These variants played the game in the same environments. Their actions were compared to the subjects' on both an aggregate and individual level. We obtained good fits in the unselfish environment. Systematic deviations in the greedy environment revealed that under this condition behavior was more determined by ecological aspects than by social comparison.

Introduction

KIS (knowledge and intentions in social dilemmas; Ernst & Spada, 1993; Ernst, 1994) is a model that integrates assumptions about problem solving, motivation, social cognition, and learning into a cognitive-motivational model of an agent involved in a commons dilemma. A commons (or resource) dilemma is a specific type of a social dilemma (Dawes, 1980). It refers to a situation in which a group shares a common resource (e.g., fish, water, forest, or clean air) from which the individual member can harvest. If too many members take too much from the common source, it is exhausted. Thus, the group interest requires moderate harvests, but personal interests may induce the individual members to harvest excessively (Van Lange, Liebrand, Messick, & Wilke, 1992). The classical example of a commons dilemma is where herdsmen involuntarily ruin a shared pasturage by adding animals to their individual herds (Hardin, 1968). Two traps are characteristic of such a dilemma situation. The first is a social trap: Gain of an action for one, losses to all. The second is a temporal trap: Gain of the action now, losses later. As a general rule, the immediate gain for individuals exceeds their share in the damage, which affects everybody in the community on a long term basis. The group as a whole would be better off if everybody restricted his/her resource use.

Commons dilemmas are typical for a variety of critical environmental problems—such as resource depletion or pollution. Agents may be individual persons, organizations, or countries as the quotation illustrates:

> The main reason for over-fishing is one familiar to economists. If all fishermen restrained themselves, each would benefit from enlarged stocks in the future. But no individual boat or fleet sees any gain from holding back unless it knows competitors are doing the same. (The Economist, 18 March 1995, p. 74)

The benefits of having a computational process model of a social agent involved in such a dilemma are twofold. (a) *Artificial social environments* consisting of these agents can be created enabling the controlled variation of social conditions in experimental work. (b) These artificial social environments can also be used to test the theoretical assumptions underlying the model, i.e., the model is able to assist its own evaluation in a non-trivial way, as will be shown later.

Related Work and Desiderata

Much work in the area of social dilemmas is based on a very restricted *minimalistic action scenario*. A minimalistic scenario is defined by two (N) parties, two mutually exclusive actions (defect, cooperate), and four (2^N) possible outcomes (pay-offs) linked to these actions. Usually, outcomes are held constant in iterated games. Subjects' decisions (defect or cooperate) are often analyzed in the framework of classical game theory (Dawes, 1980). Another venue of research uses Monte-Carlo simulations (Axelrod, 1984; Danielson, 1992; Glance & Huberman, 1994; Macy, 1995; Messick & Liebrand, 1995), where two or more artificial players employ simple decision strategies (e.g., tit-for-tat, win-stay–lose-change, win-cooperate–lose-defect) or statistical decision rules in iterated games.

In laying down *cognitive* prerequisites for adversarial problem solving and social problem solving in general, Thagard (1992) complains about the neglect of social inferences (e.g., understanding intentions and traits, anticipating the action of the opponent) in these approaches leading to two severe deficiencies: (a) the theoretical narrowness does not provide a sufficient *explanation* for behavior, nor (b) does the minimalistic scenario realistically account for the broad variety of possible human behavior in social conflict situations. These shortcomings enormously restrict the validity of generalizations based on minimalistic action scenarios.

In the following sections, we introduce a more complex and realistic instantiation of a specific social dilemma (a commons dilemma), the Fishing Conflict Game, and present a cognitive process model of the behavior in such a conflict situation, which allows to predict and explain behavior.

The Fishing Conflict Game

One instance of a commons dilemma is the *Fishing Conflict Game*, the stock of fish being the resource, the harvesting the conflict partners' actions (Spada, Opwis, Donnen, Schwiersch, & Ernst, 1987). In the setting of this game, three players act as fishermen at a simulated lake. They are instructed to pursue the goal of achieving a maximum gain from fishing. A game consists of two phases, but the number of rounds is unknown to the players beforehand. The first phase of the game starts with the simulated resource (i.e., the fish population) in the range of optimum propagation, whereas in the second phase the simulated resource is in the range of suboptimum propagation. The players are not allowed to communicate with each other, but throughout the game each player is informed about the other players' harvesting decisions. The fish propagation rate is a non-linear function, which is not imparted to the players. Overharvesting reduces the fish population, thus its propagation, and subsequently, the possible long term gain. In the extreme, it may lead to the extinction of the resource. From a psychological point of view, the setting that is simulated by the Fishing Conflict Game, can be characterized by: (a) multiple agents; (b) a dynamically changing environment; (c) inter-individual conflicts: The participants compete over the resource; (d) intra-individual conflicts: Given a particular state of knowledge and motives, cognitive-motivational conflicts are likely to result (e.g., maximizing one's gain or protecting the resource?); (e) incomplete knowledge about the robustness of the resource; (f) incomplete knowledge about the behavior of the other participants; and (g) knowledge acquisition through participation in the game.

Empirical Findings

The empirical results point to the importance of several factors that influence behavior in a resource dilemma situation (e.g., Dawes, 1980; Van Lange et al., 1992). (a) A ubiquitous finding is that behavior is tied to the individual motives of the participants. (b) Domain knowledge and experience with the problem are generally considered important determinants of the behavior. Domain knowledge in the present case means ecological knowledge. (c) Findings concerning interpersonal trust highlight the role of social knowledge in forming an estimate of the other participants' intentions and predicting their future behavior. At our laboratory, a series of five experiments was conducted with the Fishing Conflict Game (Spada et al., 1987). The findings can be summarized as follows: Participants of groups with high individual gains can be characterized by less destructive motives. They show a better ecological and social knowledge. A tit-for-tat strategy (i.e., repay overharvesting by overharvesting) of an instructed participant is misunderstood as unpredictable and exploiting. On the other hand, a confederate of the experimenter using a resource adapted equal share strategy brings about positive effects: cooperation among co-players and a stable resource at the state of a maximum sustainable yield.

The Model[1]

KIS (Ernst & Spada, 1993; Ernst, 1994) models the behavior of an agent participating in the Fishing Conflict Game. In the model, we spell out what types of knowledge and which motives come into play in determining people's behavior in this particular multi-agent situation, in which an agent has to balance multiple ecological, economic, and social short-term and long-term goals. Furthermore, following the desiderata outlined above, we have to be specific about how people understand, anticipate and adapt to the behavior of others in these situations.

Motives

In the KIS model, motives define a set of individual preferences, that are relatively constant over time. This interpretation of motives is in line with the conceptualization of *social values orientation* (cf. Liebrand, 1984), according to which subjects' preferred social orientations remain invariant over time.

In the KIS model, three motives are postulated that correspond with psychological findings (e.g., Van Lange, et al., 1992), economic theories (e.g., Etzioni, 1988), and philosophical considerations (e.g., Danielson, 1992): (a) *greed*, a utilitaristic 'get as much as possible', (b) *resource orientation*, the interest or moral obligation to stabilize the resource at the state of maximum sustainable yield, and (c) *social fairness*, which aims to minimize the differences between the own gain and those of the other players. In the KIS model, the strengths of the motives—graded as strong, moderate, and weak—make up the so called *motive structure* of a simulated player; different motive structures lead to different types of players. As three motives are postulated, KIS allows to construct six different patterns of motive structures. The motive structure influences the decisions of a player and determines the evaluation of a situation as being more or less desirable.

Knowledge, Actions, and Decisions

In a commons dilemma three sources of knowledge play a major role: ecological knowledge, social knowledge, and practical knowledge (Spada et al., 1987). Ecological knowledge refers to detecting the regularities about the resource propagation. It improves with experience and is implemented by a simple learning mechanism. Social knowledge permits the ascriptions of others' intentions, motives, and trustworthiness and allows to predict their future actions. Practical knowledge (Ohlsson, 1995) is represented by action schemas and a decision procedure.

Actions An action schema in the KIS model reflects the knowledge of how to generate behavior necessary to achieve a short-term goal. Furthermore, an action schema in the KIS model represents the experiences about its past and/or hypothesized future success with regard to the different motives in the form of motive-specific schema strength parameters. Thus, an action schema comprises a (short-term) goal, a method how to achieve that goal, and an evaluation how this short-term goal might satisfy the different motives (see Table 1).

---

[1] KIS was implemented using KEE and Common LISP on a UNIX workstation.

Table 1: Action schemas within the KIS model.

		schema strength parameters		
short-term goals	actions (methods for specifying the catch quota)	greed	fairness	resource
(a) *equal share*	$\frac{1}{2} \times (\text{catchQuota*[player1]} + \text{catchQuota*[player2]})$	-1	+1	0
(b) *relative gain*	$\frac{1}{2} \times (\text{catchQuota*[player1]} + \text{catchQuota*[player2]}) \times w_1$	+1	-1	0
(c) *res. adapted. equal share*	$\frac{1}{3} \times \text{optimumQuota*}$	-1	+1	+1
(d) *overharvest*	$\frac{1}{3} \times \text{optimumQuota*} \times w_2$	+1	-1	-1

Note. w_1 and w_2 are greater 1. Schema strength parameters are motive-specific and change through learning; higher values mean the action schema is evaluated as more useful to the motive. '*' indicates that this value has to be estimated by the model based on its social and ecological knowledge. All methods are state dependent and are specified for a game with three players.

Four action schemas are implemented in the KIS model: (a) The *equal share* action schema yields a catch quota which is as close as possible to the predicted quotas of the other players. Predicting others' catch quotas is accomplished by bringing social knowledge to bear. (b) Similarly, the *relative gain maximization* action schema is social in nature and uses predictions of the others' quotas as well, but the result will be a quota exceeding the others' quotas. (c) Integrating both fairness and ecological concern, the *resource adapted equal share* action schema uses the equality principle, but takes into account the (estimated) optimal resource propagation at the same time. (d) In contrast, the *ecological-social overharvesting* action schema generates a catch quota markedly above the latter one (see again Table 1).

Based on a player's knowledge about the current state of the resource and/or the catch quotas to be expected by the other players, the selected and instantiated action schema specifies the player's catch quota. Thus, there are three steps taken for goal-directed action: (a) Selecting a schema (build a short-term goal), (b) instantiating it (adapt to the situation), and (c) executing its procedural part (act). The selection procedure is described next.

Decisions A decision in the KIS model consists of selecting the (subjectively) best action schema in a given situation. The decision procedure integrates the action schema parameters and the motive structure into one single value by simply summing up the products of the motive-specific action schema parameters multiplied by the strength of the corresponding motives. For instance, a KIS agent with high resource and moderate fairness orientation using the schema strength parameters presented in Table 1 would decide for the *resource adapted equal share* action schema.

Social Knowledge Implementing the two social action schemas (equal share, relative gain) and predicting future developments of the resource requires anticipating the actions of the fellow players. Prediction implies understanding. But for an observer, the overt behavior in the KIS framework is ambiguous, because it is determined by an action schema *and* by ecological and/or social knowledge (see method part of Table 1). Understanding an actor thus essentially consists of disambiguating possible explanations for the observed actions. Here, social knowledge comes into play. Social knowledge allows an observer to attribute short-term goals and motives to other people's actions.

According to the predominating view in social psychology, an observer acts in this situation like a naive scientist and applies a *primitive theory of mind* to understand other people's actions. This approach, often called 'theory-theory', is at odds with 'simulation theory' which states that we understand others' actions in the absence of any theory of mind, by using the resource of our own minds to simulate the beliefs and intentions of others (Goldman, 1993).

Recently, Barnes and Thagard (in press) have argued that these alternatives are not mutually exclusive. In presenting a computational account for empathy they concluded that empathy always involves simulation, but may simultaneously include theory application. They present a computational model in which the observer tries to construe correspondences between prior own experiences and the action of the other person by analogical mapping. Close analogies appear to involve little or no theoretical work. Long-distance analogies inevitably require the application of a theory to compensate for disparate situations and goals. A similar view is taken in the KIS model.

Because in the Fishing Conflict Game all players are in an identical ecological situation, the KIS model interprets the catch quota of another player by applying all its own action schemas, which get instantiated by the model's own ecological knowledge. The model then compares the outputs with the observed catch quota. The action schema whose result comes closest to the observed behavior is attributed as being the short-term goal of that person. Note that as a consequence of this process, poor ecological knowledge will lead to substantial mis-interpretation of the observed behavior.

In a similar vein, the guiding motive of another player is determined. The situation and the already ascribed action schema are fed into the model and this action schema then is evaluated for each motive. Again, this simulation is done using the own subjective knowledge (in this case, the own action schema strength parameters). The motive that best explains the decision to use this particular action schema is considered the momentarily guiding motive of the fellow player.

Since motives are assumed to be relatively constant over time, the overall motive ascription is conservative. Applying its naive theory about motives and short-term goals, the model ascribes that motive as stable and trait-like for a fellow player which it had determined most often as his or her momentarily guiding motive during the whole game.

Learning

The KIS model behaves adaptively through instantiating action schemas to each given situation. In addition, knowledge acquisition mechanisms are provided. In KIS, all learning of practical knowledge results in decreasing or increasing schema strength parameters.

Learning by Doing This mechanism is based on the evaluation of the consequences of one's own actions. Learning by doing is considered to occur with every action taken. It only affects the dominant action schema, i.e., the one that was chosen for action.

Learning by Mental Simulation Learning by mental simulation is triggered by *impasses* (Rosenbloom, Laird, Newell, & McCarl, 1991; VanLehn, 1988). In KIS, an impasse results from the detection of an inconsistency between the expected and the observed behavior of a co-player, or when a tie between the evaluation of action schemas exists. The outcomes of mental evaluations of possible future states are then integrated into the strength parameters of the action schemas, thus possibly leading to new rankings. Possible future states are simulated using own ecological knowledge and by anticipating actions of the co-players. Mental simulation is restricted to a look-ahead depth of two game rounds.

Performance of the Model

Figure 1 shows as an illustrative example the development of the resource in a Fishing Conflict Game with three artificial players. A resource oriented artificial player with very poor ecological knowledge interacts with a greedy and a fairness oriented artificial player with good ecological knowledge. While the greedy player decides on high catch quotas for several rounds—based on the action schemas relative gain and overharvesting—and then reduces its catch quotas drastically due to a negative evaluation of these action schemas in the light of a rapidly diminishing resource, the resource oriented player takes too high harvests according to its poor ecological knowledge. With the negative experience of the third round, this player decides to select the equal share action schema. Thus, the poor ecological knowledge is overridden by a socially oriented action schema, namely the equal share schema. This leads to a recovery of the simulated resource.

Evaluation

The aim of our study imposes three empirical demands on the model: (a) The different artificial social environments defined by KIS models should be useful as experimental settings. One point is that artificial agents should be believable. (b) The effects of different social environments should be in line with previous findings in this domain. (c) Human and artificial subjects confronted with the same social environments should behave in the same way.

Method

Two experimental environments were built differing only in the motive structure of one of their two artificial players. In the *unselfish environment*, this player had a dominant resource orientation accompanied by a moderate fairness orientation. In the *greedy environment*, this player had a high greed orientation and a moderate resource orientation. The other artificial player had an identical motive structure in both environments being highly fair and moderately greedy. Both artificial players had good ecological knowledge. The third player was either the human or a matched artificial subject, i.e., a modeled KIS counterpart. For each subject, we assessed in each round the estimation of the optimum total catch quota, the actual catch quota, and the ascription of short-term goals of the co-players.

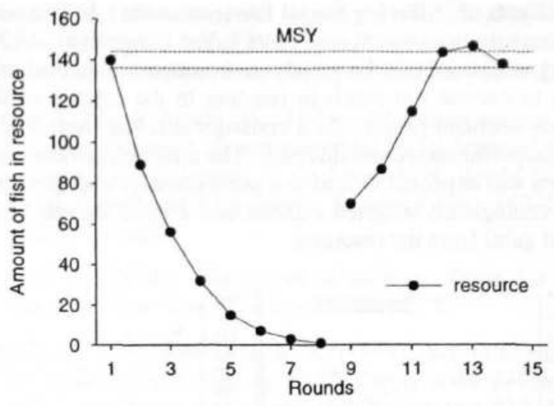

(a) Fish population

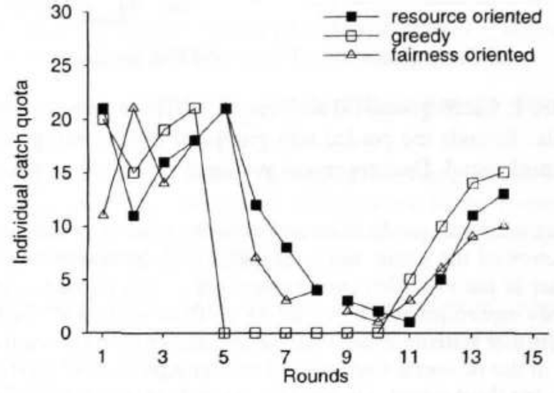

(b) Catch quotas

Figure 1: Development of the resource (a) and individual catch quotas of three artificial players (b). MSY denotes the range of maximum sustainable yield. A resource oriented player with poor ecological knowledge (■) acts in a greedy social environment (□, △).

Forty-two students took part in the experiments in single sessions, distributed equally between both experimental environments. Participants were told that their co-players were seated in adjacent rooms with networked computers.

Results and Discussion

(a) Authenticity of Artificial Social Environments Did human subjects consider the artificial environments realistic? In a short, open debriefing session after the experiment, no subject conjectured that the co-players were artificial. Moreover, in a questionnaire following the game the frequency of correct estimations of the co-players' dominant motives was quite high with 67% and 69% correct estimations for the resource oriented and gain oriented player, respectively. Thus, the model succeeded in a kind of *Social Turing Test* as proposed by Carley and Newell (1994).

(b) Effects of Differing Social Environments In line with previous studies (Spada et al., 1987; Van Lange et al., 1992) we hypothesized that the greedy environment would lead subjects to harvest too much in reaction to the catches of the greedy artificial player. As a consequence, bad overall ecological performance should result. The unselfish environment in turn was expected to lead to a good resource management, i.e., ecologically adapted catches and a good overall yield (joint gain) from the resource.

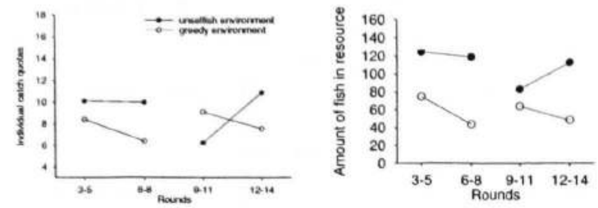

(a) Catch quotas (b) Fish population

Figure 2: Catch quotas (a) and tons of fish (b) in both environments. Rounds are pooled into groups of three starting with the third round. Data represent averaged values ($N = 42$).

According to predictions and reflecting the very dilemma structure of the game, the joint gain of all three players was higher in the unselfish environment ($M = 432.0$) than in the greedy environment ($M = 302.9$), $t(40) = 8.38$, $p < .0001$. For further statistical analysis, catch quotas and the amount of fish in the resource were pooled into groups of three starting with the third round. Using these pooled variables, ANOVAs containing the between-subject variable of social environment (unselfish or greedy) and the within-subject variable of rounds (3-5, 6-8, 9-11, or 12-14) were computed. For both variables (catch quotas and amount of fish in resource), the ANOVAs revealed significant (all $ps < .05$) main effects for environment ($Fs(1,40) = 4.38$ and 70.43) and for rounds ($Fs(1,120) = 3.06$ and 11.18), and a significant interaction between environment and rounds ($Fs(1,120) = 11.23$ and 13.31). Figure 2 shows the respective means.

Whereas in the unselfish environment the resource remained stable during the first phase and recovered rapidly in the second phase, the resource was twice exploited nearly to extinction in the greedy environment (Figure 2b). As predicted, subjects in the unselfish environment behaved sensitively to the size of the resource throughout the entire game—taking more when the fish population was high and vice versa—and exhibited positive transfer of their acquired ecological knowledge from the first phase to the second phase. Subjects in the greedy environment showed a notable deviation from this principle. At the beginning of the second phase, which started for both environments with a new fish population of 70 tons, they exhibited, on the average, catch quotas far above the corresponding quotas in the unselfish environment. But beside this exception and contrary to predictions, subjects in the greedy environment showed relatively low catch quotas, when the resource was on the way to extinction. Because subjects should have acquired almost identical ecological knowledge in the fist phase of the game, we interpret the high quotas in the greedy environment at the beginning of the second phase as retaliation in conjunction with an attempt to compensate a meagre yield during the first phase in which the greed oriented simulated player had overexploited the resource dramatically.

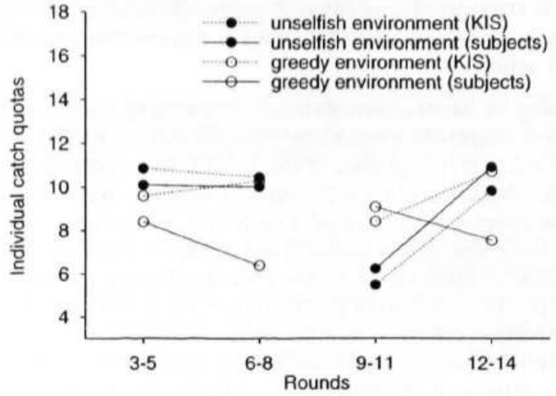

(a) Catch quotas

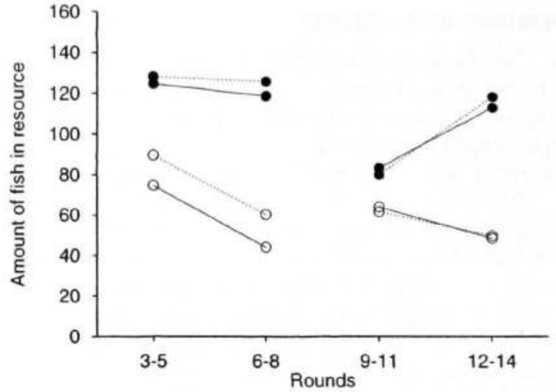

(b) Fish population

Figure 3: Subjects (solid lines) and their artificial counterparts (dashed lines): Catch quotas (a) and tons of fish (b) in both environments. Rounds are pooled into groups of three starting with the third round. Data represent averaged values ($N = 42$).

(c) Comparison between Human and Artificial Subjects On the basis of their behavior in the second round of the game, we categorized our subjects as more resource, fairness, or greed oriented. Then we built the motive structure of the KIS model according to our diagnosis to match each subject by a twin. Similarly, we adjusted the ecological knowledge of the artificial subjects. The aim of this design was to compare the behavior of human and artificial subjects from Round 3 onward.

On an aggregated level, we replicated the effects concerning the fish population for both environments (Figure 3b). Although the environments (i.e., the experimental manipulations) themselves contribute to this finding, it is nevertheless important for an unbiased comparison of catch quotas.

For the mean catch quotas of the artificial subjects, we obtained a good fit for the unselfish environment. In the greedy environment, the artificial subjects showed much higher catch quotas in the first phase compared to our subjects. In fact this is in line with earlier empirical findings (Spada et al., 1987). But we also got different trends in both phases (Figure 3a). Contrary to human subjects in the greedy environment, the artificial subjects took in reaction to the overharvesting of the other players even more although the fish population was decreasing. On an individual level, we could model about one third of our human subjects exactly, mainly due to good fits between pairs within the unselfish environment. For the greedy environment, this analysis revealed major disparities. Most notably, the modeled subjects harvested more than human subjects in the first phase and consequently did not show compensation or retaliation behavior at the beginning of the second phase.

These mismatches suggest an important shortcoming of the KIS model. KIS assumes the existence of stable motives that guide behavior. In the Fishing Conflict Game, motives themselves can be affected by situational factors (e.g., when the state of the resource diminishes dramatically) and by emotional reactions (retaliation).

Conclusions

A major contribution of this research consists in using a process model for defining plausible, realistic, and believable agents in artificial social environments. The validity of the model was tested in a study that was designed by means of these environments. Our prior finding, that a resource adapted equal share strategy has positive influence on the co-players by producing predictable and ecologically sensible behavior, was confirmed and could be explained by the model. In a greedy social environment, the model did not account appropriately for the influence of situational and emotional aspects. But despite this shortcoming, the model allows to create reactive and sufficiently realistic social learning environments based on a broad range of different artificial agents. An advantage of such a procedure is that one can define standardized sequences of learning opportunities, for which the effects can be deduced on the basis of the model.

Acknowledgments

This research was supported by grant no. Sp 251/5-x from the German National Research Foundation (DFG) to the second author. We would like to thank Michael Scheuermann and Hansjoerg Neth for assistance, and Volker Franz, Frank Ritter, and Paul Thagard for comments on an earlier draft.

References

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.

Barnes, A., & Thagard, P. (in press). Empathy and analogy, *Dialogue: Canadian Philosophical Review*.

Carley, K. M. & Newell, A. (1994). The nature of the social agent. *Journal of Mathematical Sociology, 19*, 221–262.

Danielson, P. (1992). *Artificial morality*. London: Routledge.

Dawes, R. M. (1980). Social dilemmas. *Annual Review of Psychology, 31*, 169–193.

Ernst, A. M. (1994). *Soziales Wissen als Grundlage des Handelns in Konfliktsituationen (Social knowledge as a basis for acting in conflict situations)*. Frankfurt/M.: Peter Lang.

Ernst, A. M., & Spada, H. (1993). Modeling agents in a resource dilemma: A computerized social learning environment. In D. Towne, T. de Jong, & H. Spada (Eds.), *Simulation–Based Experiential Learning* (pp. 105–120). Berlin: Springer.

Etzioni, A. (1988). *The moral dimension: Toward a new economy*. New York: The free Press.

Glance, N. S. & Huberman, B. A. (1994). Social dilemmas and fluid organizations. In K. M. Carley & M. J. Prietula (Eds.), *Computational organization theory* (pp. 217–239). Hillsdale, NJ: Lawrence Erlbaum Associates.

Goldman, A. (1993). *Philosophical applications of cognitive science*. Boulder: Westview Press.

Hardin, G. R. (1968). The tragedy of the commons. *Science, 162*, 1243–1248.

Liebrand, W. B. G. (1984). The effect of social motives, communication and group size on behavior in an N–person multi–stage mixed motive game. *European Journal of Social Psychology, 14*, 239–264.

Macy, M. (1995). PAVLOV and the evolution of cooperation: An experimental test. *Social Psychology Quarterly, 58*, 74–87.

Messick, D. M., & Liebrand, W. B. G. (1995). Individual heuristics and the dynamics of cooperation in large groups. *Psychological Review, 102*, 131–145.

Ohlsson, S. (1996). Learning from performance errors. *Psychological Review, 103*, 241–262.

Rosenbloom, P. S., Laird, J. E., Newell, A., & McCarl, R. (1991). A preliminary analysis of the Soar architecture as a basis for general intelligence. *Artificial Intelligence, 47*, 289–325.

Spada, H., Opwis, K., Donnen, J., Schwiersch, M., & Ernst, A. M.(1987). Ecological knowledge: Acquisition and use in problem solving and decision making. *International Journal of Educational Research, 11*, 665–685.

Thagard, P. (1992). Adversarial problem solving: Modeling an opponent using explanatory coherence. *Cognitive Science, 16*, 123–149.

Van Lange, P. A. M., Liebrand, W. B. G., Messick, D. M., & Wilke, H. A. M. (1992). Social dilemmas: The state of the art: Introduction and literature review. In W. B. G. Liebrand, D. M. Messick, & H. A. M. Wilke (Eds.), *Social dilemmas. Theoretical issues and research findings*. (pp. 3–28) Oxford: Pergamon Press.

VanLehn, K. (1988). Toward a theory of impasse–driven learning. In H. Mandl, & A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems*. (pp.19–41). New York: Springer.