

INCA: A Multi-Choice Model of Cooperation Under Restricted Communication

Lutz Prechelt (email: prechelt@ira.uka.de)
Fakultät Informatik
Universität Karlsruhe
76128 Karlsruhe, Germany
++49/721/608-4068, Fax: ++49/721/694092

To appear in "BioSystems" special issue about the Prisoner's Dilemma, 1995

Abstract

The advantage of the iterated prisoner's dilemma as a model of interaction is its simplicity, which allows deep mathematical analysis. On the other hand, it is unrealistic to assume that there are only two possible actions, that one always knows in advance what constitutes cooperative actions, and that successive actions are independent of each other. The present work proposes an extended interaction model with multiple action choices, only vague a priori notion of cooperativeness, action selection constraints from previous choices, and the possibility of gradual cooperation. The model assumes that the partners cannot communicate except for through their action choices. The model was explored experimentally. The results indicate that cooperation may not always be achieved as easily as the iterated prisoner's dilemma model suggests and that extended interaction models should be used as additional tools for interaction analysis.

1 Motivation

For decades, many experiments and analyses have been using the iterated prisoner's dilemma as a model of interaction, for instance to study the conditions under which cooperation can be expected to emerge between egoistic interacting partners. In this context, Axelrod (1984) proved empirically as well as analytically that strategies that were nice, provocable, and forgiving lead to mutual cooperation. A strategy is *nice*, if it never is uncooperative before its partner is uncooperative. A strategy is *provocable* if it will become (at least somewhat) uncooperative after its partner has (too long or too much) been uncooperative. A strategy is *forgiving* if it will stop being uncooperative, provided its partner has stopped being uncooperative. Axelrod showed that strategies having these three properties were both rational and cooperative: Such strategies guarantee high degrees of overall success independent of the kinds of strategies the partners use and they do this by achieving a high degree of cooperation. This result lead to optimistic expectations about the emergence of cooperation in populations of egoistic individuals.

However, the iterated prisoner's dilemma has several important limitations as a model of interaction:

L1: The choice of action for each player in each iteration is only binary in the iterated prisoner's dilemma. This is not representative of real life situations where in most cases many different actions are possible.

L2: The degree of cooperativeness or non-cooperativeness of each available action is exactly known in advance, and is fixed for each action throughout all iterations of the interaction. This is unrealistic too, because the cooperativeness of an action often depends on the expectations of the interacting partner, on the partner's understanding of what constitutes a cooperative action, and on the history of the interaction and the state it has created.

L3: With respect to the iterated prisoner's dilemma model itself, each iteration is completely independent of the previous ones. This neglects the fact that in reality actions often impose restrictions on what can be done in the (near) future.

This paper presents an extended interaction model transcending these limitations. The motivation behind the extended model is the wish to study whether the emergence of cooperation is indeed as easy as the iterated prisoner's dilemma analysis indicates. The following hypothesis will be explored:

In general, cooperation is achieved not as easily as the iterated prisoner's dilemma model suggests, even if both partners try to cooperate. Furthermore, strategies that try to be cooperative can not as easily defend against uncooperative partners as provokable, forgiving strategies can in the iterated prisoner's dilemma.

We will study this hypothesis not in the context of human subjects, but in a context of limited communication as will be explained in section 3 after the introduction of the extended model.

The model proposed takes the form of a game. Hence, we will use game terminology to describe the model: The entirety of the interaction of two partners will be called a *game*; the partners are called *players*; every single iteration of the interaction will be called a *move*; the rewards given for each move are called *points*. Further concepts in this spirit will be introduced below.

The subsequent sections present the model, which is called the INCA model, discuss its properties and connections to related work, and summarize the results of computer simulations of INCA strategies created by many different human authors.

2 The INCA model

I propose a framework for describing bilateral interactions with respect to cooperation, called the *iterated numerous choice action (INCA)* model. Its main ideas are

1. Give the interacting partners a diverse selection of possible action alternatives.
2. Make it difficult to predict the effective degree of cooperation exhibited by a particular action choice.
3. Let the actions taken by a player impose restrictions on allowed future choices.

These ideas are implemented in a two-person game with the following rules.

R1: Both players at the same time choose an integer number in the interval $a \dots b$. The number chosen by player P in a move is called P 's *throw*. The players can watch each throw as it is made, i.e., they know all numbers they and their opponent have thrown before the current move. For the following let us assume that player P chooses number p_1 and player Q chooses q_1 .

R2: If $p_1 = q_1$, neither player wins a point.

R3: Otherwise, the player with the higher number wins, unless the number is more than twice that of its opponent. Let us assume that $p_1 > q_1$, then P wins if $p_1 \leq 2q_1$ and Q wins if $p_1 > 2q_1$.

R4: A player who wins a move with some number n receives $\lfloor \log_2(n) \rfloor$ points in this move, i.e., the rounded-down binary logarithm of the winning number. The other player receives 0 points in this move.

Example: if P wins, it receives $\lfloor \log_2(p_1) \rfloor$ points, e.g. if $p_1 = 6800$, player P receives 12 points, because $\log_2(6800) \approx 12.73$.

R5: A game consists of L moves. L may or may not be known in advance.

R6: Both players must throw series of non-decreasing throws. These series must (for each player individually) have a length of at least k moves; longer series are allowed (k is much smaller than L). This means that if P throws (p_1, p_2, p_3, \dots) then $p_1 \leq p_2 \leq p_3 \leq \dots \leq p_k$ is required. After that, $p_k > p_{k+1}$ is allowed. If $p_k > p_{k+1}$ then $p_{k+1} \leq p_{k+2} \leq \dots \leq p_{k+k}$ is required; otherwise there exists some smallest number j (with $j > k$) for which $p_j > p_{j+1}$ and then $p_{j+1} \leq p_{j+2} \leq \dots \leq p_{j+k}$ is required — and so on through the whole game.

Example: If $k = 3$, then the sequence of throws (of one player)

$$1, 2, 3, \quad 1, 4, 5, 6, 8, \quad 2$$

is allowed, while

$$1, 2, 3, 4, \quad 1, 2, \quad 1$$

is not because the final 1 (which is less than 2) comes too early, namely after only one non-decreasing throw instead of the required three. Each player can choose its sequence lengths individually, i.e., the beginnings and ends of the two players' sequences need not coincide and are not influenced by the sequence lengths of the other; the only restriction is that each sequence of each player must be at least k throws long.

As an example instantiation of the model consider the case $a = 1$, $b = 12288$, $k = 8$, and $L = 1000$; these are also the parameters that are used in the experimental exploration described below. Given these parameters, the maximum number of points to be won in a single throw is 13 and the maximum number of sequences played is 125. Perfect cooperation would result in 6500 points for each player. Note that with these parameters, a nice method of cooperation is the following, called “taking turns”: Begin with a random number in the range 8192 to 11288. Then whenever you won the last move, choose the same number as in that last move again; whenever you lost the last move, choose the number that is one above your partner's last throw for your next move; on a tie, use a random number again.

Just like for the iterated prisoner's dilemma, other parameter choices can significantly change the character of the game.

3 Related work and properties of the model

There is a large body of research around the iterated prisoner's dilemma, including a number of modifications of the original setting. For instance the iterated prisoner's dilemma has been extended to interactions in which the degree of cooperation in a single action can be varied arbitrarily instead of just binary (e.g. Tsebelis, 1990) and to interactions in which noise degrades the observation of the partner's moves (e.g. Bendor, Kramer, and Stout, 1991). Other changes explicitly involve populations of strategies. For instance, (Grim, 1995) describes the behavior of spatial arrangements of variable strategies, and (Ashlock et al., 1995) analyze strategies that can explicitly choose and refuse their interaction partners besides choosing actions.

None of these modifications, however, fully addresses limitations L2 and L3. For instance, even in a noisy prisoner's dilemma it is clear what constitutes the maximally cooperative action, whereas in the INCA model it is not.

Of course, despite its higher complexity compared to the iterated prisoner's dilemma, the INCA model is still quite restricted in its modeling capabilities. For instance it can neither represent multilateral interactions or environmental influences on the players; nor does it reflect the effects of noise and errors in the communication, although it is not realistic that a choice from, say, 12288 possibilities can always be transmitted without errors; nor does it allow for freedom of action, whereas real subjects can always choose actions that are not reflected by the rules of *any* game model.

Models based on games (i.e., models that are exactly formalized and can be explored using either human subjects or computer programs) have often been criticized for being simplistic and thus unrealistic and for producing exaggerated conclusions (Knox and Douglas, 1971; Nemeth, 1970; Nemeth, 1972; Shubik, 1970). Such criticism is justified if one regards games as simulations of real interactions; however, this should not be done. Instead, games can and should be considered vehicles for producing new ideas and for testing *specific* hypotheses about particular kinds of interaction situations (Schlenker and Bonomi, 1978). With a specific hypothesis in mind, interesting results can be obtained from game models.

The INCA model represents a situation where there is only very restricted communication between the interacting partners. There is no a priori notion of what cooperative behavior is and the only way for the players to negotiate a cooperation is through the game interaction itself; there is no meta-game language. This property makes the INCA model inappropriate for learning about human interactions but appropriate for learning about life forms that do not have a general purpose language.

As said above, INCA is meant to be used to test the hypothesis that under such conditions it is difficult (more difficult than the iterated prisoner's dilemma model would suggest) to achieve cooperation, even if both partners are trying to cooperate. Mutual cooperation in INCA requires that the partners develop a common *cooperation procedure*, such as the "take turns" method described above.

The counting rule R4 introduces a tendency that high numbers are more cooperative than lower ones, because more points can be won when both players throw high numbers. But throwing large numbers is also dangerous for a player, because the sequence rule R6 creates an asymmetrical mechanism of commitment: once a cooperative (high number) decision has been made, the player is committed to that decision for a while and can be exploited by an uncooperative opponent, if that opponent threw a number that is less than half as high. The commitment does not occur for uncooperative (low number) decisions; they can always be changed to high number decisions later during the same sequence. Note that the commitment is due to the rules of the model, not due to the expected reaction of the partner. In contrast to the prisoners dilemma, there is no choice of a single constant number that leads to permanent mutual cooperation in INCA.

4 Experimental exploration

Unlike the iterated prisoner's dilemma, the INCA model is not well suited for rigorous mathematical analysis, because of its complexity and because there is no canonical definition for basic notions such as "niceness" or even just "cooperation". To explore the model, I have thus carried out a case study by organizing two international contests (1993 and 1994). In these contests, each participant had to write a computer program implementing a strategy according to the INCA rules with the parameters mentioned above ($a = 1$, $b = 12288$, $k = 8$, and $L = 1000$). The strategies had to strive for high scores not within individual games but over a whole tournament of games, so that for each game not winning the game was the goal, but scoring many points.

The contests consisted of each-pair-once tournaments as well as evolutionary tournament simulations. In an each-pair-once tournament, each strategy played one game against every other strategy. Strategies were ranked according to the sum of points they received over all these games. In addition, strategies also had to play against themselves the 1994 contest (but not in the 1993 contest). For an evolutionary tournament, multiple each-pair-once tournaments are played with multiple instances of each strategy participating. In the first round, each strategy has the same number of instances. The relative success of each strategy S determines the fraction of instances for the next round that are S 's. For example, if in round n the instances of strategy A win 4000 points per game on the average, whereas the instances of strategy B win only 2000, then in round $n + 1$ there will be twice as many A 's as there will be B 's, independent of how many A 's and B 's participated in round n . The total size of the population is held constant from each round to the next. Round after round is computed until no further changes in the population occur. The strategies are then ranked according to the fraction of the population they represent. Evolutionary tournaments were played only during the 1994 contest.

Altogether, 69 strategy author teams from 12 countries participated. Due to this diversity we can expect that the study has a high degree of external validity. A significant amount of data was gathered about the INCA strategies during the contests, so only the most interesting findings can be presented here and only in a qualitative manner. However, the result data, contest software, and participant strategy programs of both contests are available electronically on the Internet via anonymous FTP from the machine `i41s10.ira.uka.de` in directory `/pub/knobeln`. These are the most interesting findings:

1. The winner of the first each-pair-once tournament, let us call it W , was not programmed by hand, but instead had been evolved automatically by a genetic algorithm. This indicates that it is indeed difficult to find good strategies for INCA, even for intelligent beings. The principle structure of the strategies produced by the genetic algorithm is a decision tree. Each node is a condition and each leaf contains one or more numbers or intervals from which the next throw is chosen. The conditions and intervals use fixed numbers and variables, e.g., the last throw of the opponent or the current total points of the strategy or the remaining number of throws in the current sequence — anything which the author of the genetic algorithm could imagine to have some relevance in the game. A crossover of two strategies is produced by simply cutting the two trees anywhere and building a new tree from the parts; simple mutations are also used.
2. This program W was extremely uncooperative; for instance it typically gets only about 40 points (out of 6500) when it is playing against itself. (In the first contest, a strategy did not have to play against itself during the tournaments.) The genetic algorithm had evolved two very successful features in the strategy. The first one is the exploit-the-naïve trick: W often manages to make its own sequences end one throw after the sequences of the opponent. It then throws 12288 in the last throw of its own sequence, misleading naïve opponent strategies into throwing 6143 in the first throw of *their* sequence, enabling W to exploit the opponent for 7 throws by throwing 3071 each time, while the opponent is committed to 6143, resulting in 77 points for W ; afterwards the same cycle begins again. Figures 1 to 3 show the first few throws of a game of W against an (otherwise also very strong) opponent strategy s ; you can see the exploit-the-naïve trick at work, although it does not work perfectly against this particular opponent. The second successful feature is the pretend-cooperation trick: The strategy loses the first few throws on purpose, because this makes many other strategies so extremely cooperative that they are willing to lose many throws later without becoming uncooperative.
3. Programs with a strong bias towards cooperative action were the minority in both contests; most programs had at least some uncooperative aspects. All successful programs, however, were *able* to cooperate, when they could not exploit their opponent.
4. Most of the more successful programs used several different heuristics for computing their

n	throw		points		n	throw		points		n	throw		points	
	W	s	W	s		W	s	W	s		W	s	W	s
1	2	2	0	: 0	18	4162	8324	55	: 61	35	4	6143	130	: 89
2	3	2	1	: 0	19	4162	12288	67	: 61	36	3070	6143	141	: 89
3	4	2	3	: 0	20	12288	12288	67	: 61	37	3070	6143	152	: 89
4	4	8	3	: 3	21	12288	12288	67	: 61	38	3070	6143	163	: 89
5	4	8	3	: 6	22	12288	2	67	: 62	39	3070	6143	174	: 89
6	16	8	7	: 6	23	12288	2	67	: 63	40	3070	12287	185	: 89
7	16	32	7	: 11	24	4	6143	69	: 63	41	12286	12287	185	: 102
8	16	65	11	: 11	25	4	6143	71	: 63	42	12288	12287	198	: 102
9	130	65	18	: 11	26	3070	6143	82	: 63	43	12288	12288	198	: 102
10	130	65	25	: 11	27	3070	6143	93	: 63	44	12288	2	198	: 103
11	130	260	25	: 19	28	3070	6143	104	: 63	45	12288	6143	198	: 115
12	130	260	25	: 27	29	3070	12287	115	: 63	46	4	6143	200	: 115
13	520	260	34	: 27	30	12286	12287	115	: 76	47	3070	6143	211	: 115
14	520	1040	34	: 37	31	12288	12287	128	: 76	48	3070	6143	222	: 115
15	520	2081	43	: 37	32	12288	12288	128	: 76	49	3070	6143	233	: 115
16	2080	2081	43	: 48	33	12288	2	128	: 77	50	3070	6143	244	: 115
17	4162	2081	55	: 48	34	12288	6143	128	: 89					

Figure 1: The first 50 throws and corresponding won points of a game of the 1993 winner strategy W against another strong strategy s .

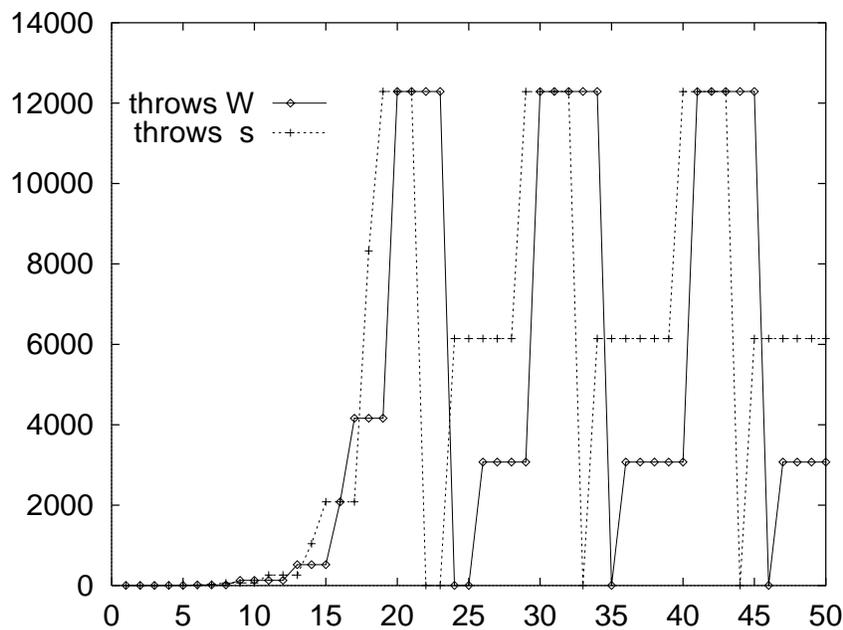


Figure 2: The first 50 throws of the game of W against s .

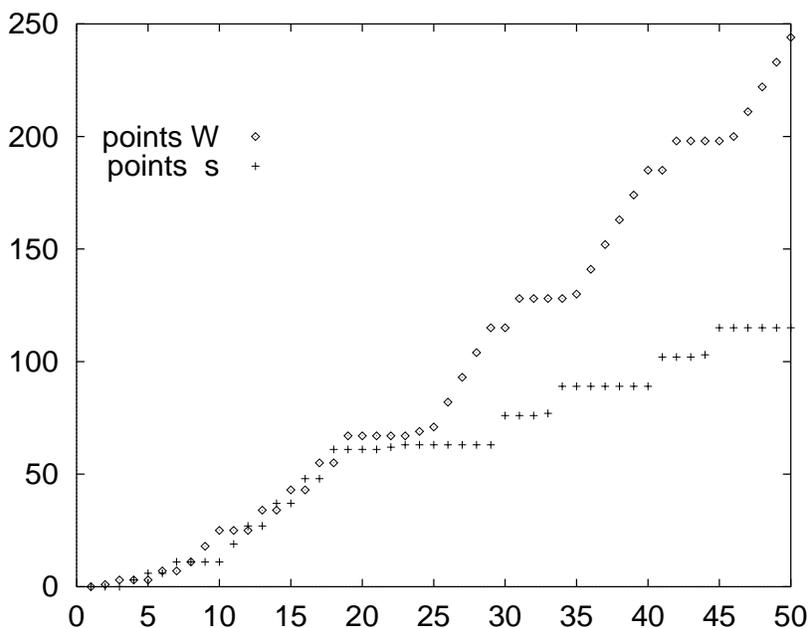


Figure 3: Accumulation of points during the first 50 throws of the game of W against s .

throws. In particular, most programs had a distinction between a cooperative and a non-cooperative mode and used complicated decision algorithms to switch between them. No simple strategy was found that is cooperative, provokable, and forgiving, as is the TIT FOR TAT rule in the iterated prisoner's dilemma.

5. Although the “take turns” cooperation scheme was discovered by many participants in the first contest and hence was well known before the second, not all participants adopted it in the second contest. Even those who did complemented it with non-cooperative parts and a rule for switching between cooperative and non-cooperative modes in order to become provokable.

6. The winner of the 1994 evolution tournament was a very cooperative yet provokable and forgiving strategy from the 1993 contest that had not been changed at all in the meantime. It also ranked third in both the 1994 and the 1993 each-pair-once tournament. This constant success of a cooperative strategy suggests, first, that there *are* strategies that are good in a general sense in the INCA setting and, second, that such strategies are primarily cooperative ones.

7. The winner of the 1994 each-pair-once tournament uses a mix of cooperative and complicated aggressive partial strategies, switching back and forth. Although it should not be called generally aggressive, it was too aggressive to succeed in the evolutionary tournament, where it ranked only 9th. This suggests that the basic property of the iterated prisoner's dilemma, namely that cooperation tends to pay off, is maintained by the INCA model.

8. The less successful strategies are not generally the less cooperative ones. Many of them just lack the provokability or the forgiveness needed to cope successfully with less cooperative partners.

9. The average number of points scored per game in the 1994 contest was only 3338 (out of 6500 possible with perfect cooperation). For the better half of the participants it was 4502, for the winner 5583. Thus, there is considerable room for improvement.

10. In a sense, there was a significant shift towards higher cooperation in the best few strategies from the 1993 contest to the 1994 contest: If a strategy wins less points than it loses to its opponents, the strategy is either cooperative, or weak, or both. In the 1993 contest, each one of the 20 most successful strategies won more points than it lost. In the 1994 contest, there was only one strategy with this property under the best six.

It would be interesting to explore the behavior of the INCA model for other values of the parameters, in particular for smaller and for much larger sequence length k . However, the exploration relies on strategies that are carefully designed and programmed for a particular set of parameters by many different people. This approach makes it very difficult to extend the exploration to many different sets of parameters. Nevertheless, the exploration uncovered some interesting aspects of the INCA model, as we have seen above.

5 Conclusion

The INCA model is a framework of bilateral interactions in which cooperation is useful for the partners, yet the interaction partners cannot readily determine how to be cooperative, let alone how to achieve cooperation. The model serves to help understanding why cooperation may be more difficult than one might expect from the analysis of the iterated prisoner's dilemma. Computer simulations show that within the INCA model, non-cooperative strategies have much better chances of success even in a predominantly cooperative environment than within the iterated prisoner's dilemma model.

In particular, the simulation results suggest the following statements about the properties of the INCA model. First, the more complicated model structure leads to more complicated interaction dynamics, compared to the iterated prisoner's dilemma. Second, simple strategies that are generally successful are difficult to find, despite the availability of the simple "take turns" cooperation rule. This problem arises because it is not clear what is a good method of being provokable and forgiving. Third, an interesting property of the iterated prisoner's dilemma model is maintained by the INCA model: successful strategies tend to be cooperative, provokable, and forgiving.

Advanced models with multiple action choices and non-obvious mechanisms for cooperation do not allow for mathematical analysis as easily as the iterated prisoner's dilemma does. Nevertheless such advanced models should be used as one of the tools for the analysis of co-evolving systems, because they can provide new insights and can point out artifacts introduced by simpler models.

References

- [1] Dan Ashlock, Mark D. Smucker, E. Ann Stanley, and Leigh Tesfatsion. Preferential partner selection in an evolutionary study of Prisoner's Dilemma. *BioSystems*, in this issue.
- [2] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, Harper Collins, 1984.
- [3] Jonathan Bendor, Roderick M. Kramer, and Suzanne Stout. When in doubt...: Cooperation in a noisy Prisoner's Dilemma. *Journal of Conflict Resolution*, 35(4):691–719, December 1991.
- [4] Patrick Grim. The evolution of greater generosity in the Spatialized Prisoner's Dilemma. *BioSystems*, in this issue.

- [5] R. E. Knox and R.L. Douglas. Trivial incentives, marginal comprehension, and dubious generalizations from Prisoner's Dilemma studies. *Journal of Personality and Social Psychology*, 20:160–165, 1971.
- [6] C. Nemeth. Bargaining and reciprocity. *Psychological Bulletin*, 74:297–308, 1970.
- [7] C. Nemeth. A critical analysis of research utilizing the Prisoner's Dilemma paradigm for the study of bargaining. In L. Berkowitz, editor, *Advances in Experimental Social Psychology*, Vol. 6, pages 203–234. Academic Press, New York, 1972.
- [8] Barry R. Schlenker and Thomas V. Bonoma. Fun and games: The validity of games for the study of conflict. *Journal of Conflict Resolution*, 22(1):7–38, March 1978.
- [9] George Tsebelis. Are sanctions effective?: A game-theoretic analysis. *Journal of Conflict Resolution*, 34(1):3–28, March 1990.