Bargaining in a Three-Agent Coalitions Game: An Application of Genetic Programming

Garett Dworman, Steven O. Kimbrough, James D. Laing

The Wharton School, University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104-6366

Introduction[1]

We are conducting a series of investigations whose primary objective is to demonstrate that boundedly rational agents, operating with fairly elementary computational mechanisms, can adapt to achieve approximately optimal strategies for bargaining with other agents in complex and dynamic environments of multilateral negotiations that humans find challenging. In this paper, we present results from an application of genetic programming (Koza, 1992) to model the co-evolution of simple artificial agents negotiating coalition agreements in a three-agent cooperative game.[2] The following sections summarize part of the scientific literature that motivates this research, describe briefly the genetic programming approach we use to model this game, then present results demonstrating that, through a process of co-evolution, these artificial agents adapt to formulate strategies that cope reasonably well under difficult circumstances to negotiate coalition agreements that not only rival those achieved by human subjects but also approximate those prescribed by cooperative game theory as the solution of this game.

Background

Adaptive processes of multilateral negotiations. Results from the rapidly growing collection of laboratory studies indicate various ways in which humans depart in multilateral negotiations from the perfect rationality envisioned in the theory of games (e.g., Aumann, 1989; Camerer, 1990), consonant with a strong research tradition initiated in the seminal work of Allen Newell and Herbert Simon (e.g., Newell and Simon, 1972; Newell, 1990). We interpret these results to indicate the following general scenario (cf. Albers and Laing, 1991). People approach somewhat novel multi-person conflict situations, not *de nova*, but from perspectives of tentative rules of thumb developed from their previous experience. They cope with this new environment by keying on prominent features of the situation and employing simple heuristics to form an initial, but tentative, assessment of the bargaining problem. As suggested by behavioral decision theory (cf. Tversky and Kahneman, 1974; Dawes, 1988), this "first-cut" analysis produces the starting points ("anchors") from which subsequent adjustments are made, whether via cognitive or adaptive responses, to the competitive process of negotiations. This adjustment process is sensitive to vagaries in the co-evolution of the parties as they respond to one another. It is only with considerable experience that players begin to approach outcomes approximating a game-theoretic equilibrium. Even then, people have difficulty articulating a well-defined, comprehensive strategy. Their behavior seems to be based on rules of thumb that represent, at best, mere fragments of the type of complete and optimal strategies identified in the theory of games for negotiating coalition agreements (cf. Selten, 1988, ch. 9; Laing, 1991). Players apply these strategy fragments myopically, and have great difficulty anticipating the other players' moves. Thus, although there is some learning with experience, much of the movement towards the game-theoretic solution appears to result from adaptive, rather than cognitive, processes.

Game-theoretic approaches. In response to these results from laboratory studies, recent progress in game-theoretic reasoning injects into the analysis some elements of nonrational behavior. For one example, Selten (1988, ch. 1) introduced small probabilities that each player might make a mistake ("tremble") when implementing (but not computing) a strategy. For another example, to account for the fact

[1] File: gpconf08.doc, July 18, 1995. We thank David K. Hildebrand and two anonymous referees for helpful comments, and Martin McCormick for conducting some of the simulations.

[2] Technically, this is a cooperative game with sidepayments represented in characteristic function form. Two papers (Dworman, Kimbrough and Laing, 1995 and 1995b) that provide additional documentation on this research and our present implementation are available on the Internet at URL: http://opim.wharton.upenn.edu/users/sok/comprats/ourdocs.html.

that humans can cooperate in the finitely repeated prisoners' dilemma, Kreps, Milgrom, Roberts, and Wilson (1982) introduced the possibility that one player believes that another player is, with some small probability, a "crazy type" who is programmed to always play a particular strategy, such as "tit-for-tat." The rational player takes this into account and may conclude that it pays to emulate the "crazy type," at least for a while. Paradoxically, such injections of a little nonrationality into the analysis greatly complicate the decision problem facing any player who is computing a best response. Such work testifies to the increasingly sophisticated body of rigorous ideas we call game theory that continues to yield significant insights into the structure of mixed-motive situations. Yet such attempts to incorporate elements of nonrationality into game theory just nibble at the edges of the basic problem, because they do not generate a plausible causal explanation of how finite agents, such as humans, cope with the complexities of mixed-motive situations.

A fundamentally more direct approach is to assume that boundedly rational players address complex game situations, not by solving the game *de nova*, but rather by employing fairly simple rules of thumb that help to govern behavior but are too incomplete to constitute a well-formulated strategy. Aumann (1989) calls this "rule rationality," in contrast to "act rationality." From this perspective, the rules may evolve with experience over a sequence of games, and help to determine the starting point of deliberations within any bargaining game. The process of negotiations from this initial stage within the play of the game is guided by these rules and the dynamic processes of competition. Thus, in contrast to previous game theory's almost exclusive focus on equilibrium and scant attention to the processes through which equilibrium is reached in rational play, we think the time is ripe to investigate the adaptive coping processes through which boundedly rational agents adjust to complex, dynamic, and reactive environments.

Some research in animal behavior (e.g., Hammerstein and Reichert, 1988), inspired by the seminal work of the biologist Maynard Smith (1982), is based upon this boundedly-rational process model. This approach assumes that each animal is born genetically endowed with specific pure strategies, and that the population as a whole evolves to a *distribution* of pure strategies that constitutes a *mixed*-strategy (Nash) equilibrium. The biological approach has been used also to understand markets and formal organizations (Cornell and Roll, 1981).

Finite automata and genetic algorithms. A related approach is to represent alternative, complete strategies as a set of finite-state automata that compete with each other in a tournament of, say, repeated two-agent, binary choice (cooperate *vs.* defect), prisoners' dilemma games (e.g., Abreu and Rubenstein, 1988; Miller, 1989; Andreoni and Miller, 1990). These structures play a tournament of repeated prisoners' dilemma games with each other. The strength (cf. fitness) of each structure increments or decrements in accordance with the payoff this strategy wins in each game it plays. At the end of the tournament, the next generation of structures is created by means of a genetic algorithm. Miller's results, for example, indicate that, after an initial adjustment period from the random start, the frequency of cooperation and the average payoffs realized from the repeated game encounters increase across generations until a plateau is reached. Moreover, the top performers that emerge from Miller's computer studies perform credibly against strategies submitted by game theorists for Axelrod's (1984) second tournament of repeated prisoners' dilemma games.

Two problems. These results indicate that adaptive processes can generate relatively simple but reasonably successful and robust strategies for dealing with a difficult environment. Yet studies such as Miller's, however promising and provocative, have two limitations. The first, which we call the *complexity problem*, is that the artificial agents are playing in a fairly simple environment. How will they perform in more complex environments, especially environments with a richer variety of relevant information and in the presence of multiple other, co-evolving agents? The second limitation, which we call the *encoding problem*, arises from the fact that all these experiments employed a specific, usually fairly clever, representation of the agent's problem. The investigator imposed this encoding, rather than permitting the agents to discover a useful formulation of their problem. In sum, recent computational work with genetic algorithms is providing important new insights into bounded rationality. This work, however, lacks generality in at least two important ways, leading us to use a genetic programming approach.

Fitness	Average number of points attained per game in which agent attained the floor.
Parameters	Generations = 4,000 Population size = 50 Creation procedure = Ramped half-and-half (Initial population: random) Selection = Fitness proportionate
Syntactic Structure	strategy :: if-then-else-stmt if-then-else-stmt :: (IF-THEN-ELSE condition-stmt action-stmt action-stmt) condition-stmt :: (condition-fn condition-stmt condition-stmt) | condition action-stmt :: if-then-else-stmt | action condition-fn :: AND | OR | NOT action :: ACCEPT | offer condition :: {player-id amount-lower-bound amount-upper-bound} offer :: {player-id amount}

Figure 1: Tableau for strategies of each agent

Methods

Our application of genetic programming can best be understood in light of the two problems discussed in the last paragraph: complexity and encoding.

Complexity. We have designed and conducted a series of computer simulations to investigate the evolution of alternative strategies for bargaining in three-agent coalitions games (Dworman, 1993; Dworman and Schoenberg, 1993; Dworman, 1994; Dworman, Kimbrough and Laing, 1995). In these games, three agents negotiate competitively to decide which two-agent coalition forms and to agree on a division between the coalition partners of the points constituting this coalition's value. The agent who is excluded from the coalition wins zero. The game is not symmetric, in that each coalition (*ij*) of two players has a unique value: v(AB)=18, v(AC)=24, v(BC)=30. This is a dynamic environment in which three agents are co-evolving throughout while they negotiate a considerably more difficult decision problem than has been addressed in the previous research cited above. We are addressing a more complex decision problem than has been reported in the literature.

Encoding. The finite automata and genetic algorithm experiments, discussed above as well as our earlier work (Dworman, 1993; Dworman and Schoenberg, 1993; Dworman, 1994) have relied upon experimenter encoding of the agents' problem representation. Ultimately one cannot avoid this entirely, but we believe genetic programming can increase substantially the distance between what the experimenter determines and what the agents discover for themselves.

Experimental Design[3]

Strategy Trees. Our simulation is summarized by the tableau in Figure 1. Each agent's alternative strategies for playing the game are represented by a population of 50 strongly-typed GP trees. Each strategy is a nested IF-THEN-ELSE statement. The condition statement applies only to the most recent offer made to the agent. If the condition statement evaluates to TRUE then the first of the two action statements is evaluated. Otherwise, the second action statement is evaluated. The condition statement may be a condition or a Boolean combination of conditions. Action statements are either ACCEPT, an offer (which implies a rejection of the previous proposal), or another IF-THEN-ELSE statement. Therefore, the strategies can be of almost arbitrary complexity by nesting Boolean condition statements and IF-THEN-ELSE action statements.[4] The output of a strategy is either the ACCEPT symbol or an offer. The condition and offer terminals are represented by a triple of values and a pair of values respectively.

[3]This section summarizes the design of our simulations. For greater detail, see (Dworman, Kimbrough, and Laing; 1995b.)

[4]Actually, we limit the complexity of the strategies to a maximum tree depth of 8. See Dworman, Kimbrough, and Laing (1995b) for details.

These complex data structures (CDS) were introduced because a coalition structure consists of two pieces of information: A set of players in the proposed coalition, and a distribution of the coalition's value over those players. Therefore, combinations of player names and payoff amounts are essential building blocks for constructing strategies.

Conditions specify contingencies for the strategy and evaluate to TRUE or FALSE. A condition CDS – written as {Player Lower-Bound Upper-Bound} – evaluates to TRUE if, in the proposed agreement, the specified player would receive a payoff between Lower Bound and Upper Bound (inclusive). For example, suppose B offers 5 points to A. Taking v(AB)=18 into account, this implies the proposed coalition agreement ((B 13)(A 5)). Then the following conditions in A's strategies would evaluate as follows:

{A 5 14} TRUE: A's proposed payoff of 5 points lies in the closed interval [5,14].

{B 2 10} FALSE: B's proposed payoff of 13 would exceed the Upper Bound.

{C 0 12} TRUE: C would win 0 points (implicitly).

An offer signifies a rejection of the current proposal and a subsequent counteroffer. An offer CDS – written as {Player Amount} – represents a proposal to the specified respondent (Player) who would receive the payoff stipulated by Amount. The initiator's own proposed payoff equals the coalition's value minus Amount. Therefore, if Player C makes the offer {A 10}, then, given v(AC) = 24, the proposed coalition agreement is ((C 14)(A 10)).

Figure 2 gives an example of a strategy for each agent. In the first example, agent A will accept any proposal in which it is offered between 2 and 8 points (inclusive). Otherwise, it will reject the offer and propose an ((A 6)(C 18)) split of v(AC)=24.

Genetic Operators. We employ four genetic operators to create new strategies, and constrain these operators to ensure that the syntactic structure is maintained. *Crossover* swaps subtrees from two parents. *Mutation* changes the symbol at a randomly picked function node or a terminal node in the tree. If a terminal node is chosen then a new CDS may be created for that node. We provide two additional operators – *CDS crossover* and *CDS mutation* – to evolve the genetic material inside the CDS terminals. Crossover and mutation operate on entire subtrees or nodes of a program tree; they do not operate on the values inside a CDS. Therefore, without CDS-specific operators, the values inside a CDS would never change, and the agent would be prevented from evolving new offers or conditions.

Both of the CDS operators work on a randomly chosen field within a CDS. CDS mutation toggles a random bit in the chosen field, thereby creating a new value. CDS crossover selects two CDSs, one from each parent, then picks a random point in the chosen field's bit string and swaps the right-hand sides of that field's bit string in the two CDSs. Again, both CDS operations are constrained to ensure legal syntactic structures.

When a child is created without crossover (e.g., reproduction) it is probabilistically subjected to mutation and CDS mutation. When children are created via crossover the post-crossover children are probabilistically subjected to mutation, CDS crossover, and CDS mutation.

Agent A	(IF-THEN-ELSE {A 2 8} ACCEPT {C 18})
Agent B	(IF-THEN-ELSE (NOT {B 10 30}) {C 15} ACCEPT)
Agent C	(IF-THEN-ELSE {B 10 30} (IF-THEN-ELSE {C 0 16} {B 12} ACCEPT) (IF-THEN-ELSE (OR {C 20 30} {A 12 18}) ACCEPT {A 5}))

Figure 2: An example of a strategy for each agent

Strategy Evaluation. Our research simulated negotiations in the three-player cooperative game defined earlier. Each agent is represented by a population of 50 strategies. A tournament is used to evaluate all the strategies. Ideally, every strategy in an agent's population of rules would play each combination of opponent agents' strategies and do so three times, so that each of the players in the combination is given a turn to begin the negotiations. Unfortunately, that would require $3(50^3) = 375{,}000$ games per generation, and thus is not feasible. Instead, selecting randomly without replacement, we match each strategy of a player with exactly 250 combinations of the other two players' strategies. Each selected combination then plays three games, each initiated by a different player. Thus, each strategy plays 750 games in the tournament. A strategy's strength (fitness) is based on its score in the tournament, which we calculate as the average payoff (points) won by the strategy in the outcomes of games in which, sometime during the negotiations, it attained the floor.

Each game is played as follows: To start the game, we make a null offer to the agent selected to be the first initiator. If the agent accepts, the game ends and all agents receive zero points. Otherwise, this agent makes an offer to another agent. If this other agent accepts, the game ends, and the coalition partners win the payoffs as specified in the offer, while the third player wins zero. If the responder does not accept, then it becomes the initiator and makes an offer, tacitly rejecting the previous offer. The game continues in this fashion until an agreement is reached, or until a specified maximum number (here, 6) of offers have been rejected, in which case the game ends by default, and all three agents receive a payoff of zero.[5] Figure 3 presents the bargaining sequences that would be generated in the three games played by the strategies displayed in Figure 2.

First Initiator:	Bargaining Sequence:
A	A: ((A 6)(C 18)) C: ((C19)(A 5)) A: ACCEPT
B	B: ((B 15)(C 15)) C: ((C 18)(B 12)) B: ACCEPT
C	C: ((C 19)(A 5) A: ACCEPT

Figure 3. Three games played by the strategies displayed in Figure 2

Simulation Termination. We terminate the simulations when a specified number of generations (here 4000) has been completed, rather than using a stopping criterion based on the extent to which optimum behavior is approximated. We are interested in the behaviour of this system over time, and the extent to which these simple, heuristically-driven, adaptive agents that do not calculate optimal strategies can realize optimal or near-optimal behavior in dynamic, complex bargaining situations.

Results: Artificial Agents and Human Subjects

This section analyzes data from 36 simulations of this model through 4000 generations. In discussing these results we shall compare where possible the behavior of these artificial agents with the behavior of human subjects who were observed in a laboratory study of a game that is comparable to ours (Kahan and Rapoport, 1974: Game IV, Experiments 1-3 combined, n=48 triads).[6]

The simulations are based on a 12-cell experimental design consisting of a Cartesian product of parameter values for the genetic programming representation.[7] In each of these twelve cells, separate runs were conducted for three different seeds of the pseudo-random number generator. The data result from 12×3 = 36 simulations. We recorded data at the random start (generation 0), and then every five generations in the 4000-generation runs; thus the results for each of the 36 runs are based on variables contained in 801 reports, each of which averages the results from up to 3(players) × 50(rules/player) × 750(games/rule) = 112,500 distinct bargaining sequences played in the tournament for that generation. Due to space limitations, we shall report the results in terms of averages of these variables.

Game length and agreement frequencies. Overall, these simple artificial agents are reasonably successful in reaching coalition agreements. On average, even during the first 1000 generations, they require less than two rounds to agree on an outcome. Typically, the first or second offer from an agent is accepted. Moreover, fully 84.6% of the games result in an agreement, rather than ending by default through failing to agree after 6 rounds of offers. Although this is a high level of agreement, the artificial agents are

[5] The restriction to 6 offers per game is not very limiting. In our simulations, almost all agreements are reached within 3 offers. Initially, we limited games to 10 offers, but discovered that any game that went beyond 5 offers continued until the game ended by default. Therefore, we lowered the limit to 6 to speed up the simulations.

[6] The characteristic function of their Game IV is: v'(AB)=66, v'(AC)=86, v'(BC)=106, else v'(S)=0. To compare payoffs achieved by human subjects in Game IV to those obtained by our artificial agents, we use the following transformation for positive-valued coalitions: v(S) = 0.3[v'(S)-6] if v'(S)>0; hence v(AB)=18, v(BC)=24, v(BC)=30.

[7] For all runs reported in this paper, mutation probability = 0.5. The 12-cell design is based on the full Cartesian product (2×2×3) of the following parameter sets: crossover fractions = {0.2, 0.5}, leaf crossover probabilities = {0.2, 0.5}, and leaf mutation probabilities = {0.2, 0.5, 0.8}. See Dworman, Kimbrough, and Laing (1995b) for details.

somewhat less successful than human subjects in forming agreements: all 48 of Kahan and Rapoport's triads succeeded in forming a coalition.

Equilibrium. We are interested particularly in determining the extent to which these simple, cognitively constrained agents can co-evolve in this rather complex bargaining environment to produce coalition agreements that approximate those prescribed by the theory of cooperative games. Various solution concepts from cooperative game theory prescribe for this game a *quota solution* (e.g., see Shubik, 1982: 177n), akin to a vector of equilibrium prices, which specifies the equilibrium values of the three agents' positions in this situation when each agent seeks to maximize its total payoff. Specifically, the quota solution to this game is $q = (q_A, q_B, q_C) = (6, 12, 18)$. It prescribes that if coalition (*ij*) forms in this game, then the payoffs should be divided between coalition partners in accordance with their quotas, while the third player must, in this game, win zero. For example, the quota solution prescribes that if A and B form a coalition, then they should split $v(AB) = 18$ such that A wins $q_A = 6$ points and B gets $q_B = 12$. Note that $q_A + q_B = v(AB)$.

We now describe the extent to which, at "equilibrium," the artificial agents reach agreements approximating those prescribed by game theory, even though they are being selected in the evolutionary process on the basis of their fitnesses as measured by the average payoffs they win in the tournaments, and *not* by how closely their payoffs approach the quota solution.

In these simulations, "equilibrium" is at best transitory. In particular, crossover and mutation episodically can cause severe disruptions which reverberate throughout the co-evolving system of agents. To measure the extent to which the system is presently in a state that approximates an equilibrium, we use the following operational construct. Let $x(t)$ denote any variable x at a point t in the sequence of generation reports. In particular, we shall be concerned with the case in which $x(t)$ represents the average payoff won as a member of a coalition agreement by an agent's *best rule*, that is, the agent's rule that wins overall the greatest average payoff. First, let us define $x(t)$ to be *settled prima facie* if both of the following conditions obtain: (1) the $m{:=}5$ immediate predecessors of $x(t)$, with mean $x(m,t)$, have a standard deviation no greater than $d{:=}0.5$, and (2) $x(t)$ deviates no more than d from $x(m,t)$. Second, if $x(t)$ is settled prima facie, then we define to be *settled retrospectively* any of its m immediate predecessors that deviates no more than d from $x(m,t)$. Then, we say that $x(t)$ is *settled at t* if it is settled prima facie or retrospectively at t. Finally, we define the *system* of agents to be *settled at t* if, for each of the three agents, the average payoff won by that agent's best rule when it succeeds in forming a coalition agreement is settled at t. We shall report results in terms of weighted averages for those generations (n=3158, out of a total of $801 \times 12 = 9612$ generation reports) during which, in this sense, the systems were settled.

Coalition frequencies. During these periods of relative stability, the artificial agents formed coalitions with the following frequencies: p(AB)= 0.359, p(AC)=0.360, p(BC)=0.162. This contrasts with the oft-observed proclivity of human subjects in characteristic-function games to opt for the "most valuable" coalition: the 48 coalitions observed by Kahan and Rapoport are distributed n(AB)=6, n(AC)=7, n(BC)=35. Yet if payoffs conform to the quota solution then each player should be indifferent as to which coalition that player joins. Thus, it might seem that, at equilibrium, all two-person coalitions should be equally likely to form in this game.[8] Apparently the artificial agents, ignorant of the this game's characteristic function, were not distracted by the fact that (BC) is the "most valuable" coalition.

Coalition payoffs and the quota. According to the quota solution, players A, B, and C, respectively, should win payoffs of 6, 12, and 18 as members of the coalition. During the periods in which the system was settled, the best rules of artificial agents A, B, and C, respectively, averaged in their coalition agreements a payoff of 8.6, 11.9, and 18.5 points. (For these same periods, the payoffs won by each of the agent's median rules averaged only about a half-point less: 8.1, 11.4, and 18.1, respectively.) Thus, A's best rule averages 2.6 points more than the quota in these simulations. In comparison, human subjects in Kahan and Rapoport's experiments average 4.8, 12.7 and 17.7 (transformed — see prior footnote) to players A, B, and C, respectively. Presumably, A is disadvantaged by the proclivity of human subjects, unlike the artificial

---

[8] We have restricted these agents to employ strategies that do not select actions probabilistically (i.e., pure strategies). In contrast, for a game in which negotiations are governed by rules that differ from those used in our games, Selten (1988: ch. 11) has characterized probabilistic strategies for negotiating three-person quota games that, at equilibrium, both implement the quota solution and generate precise coalition probabilities such that, indeed, BC is the most likely coalition.

agents, to form the (BC) coalition. Players B and C approximate their quotas in both the human and artificial data.

Overall, the mean absolute deviation from the quota solution is roughly 2.5 points in the artificial data, rivaling that of 2.0 points in the human data. Moreover, in the artificial data, the distance from the quota decreases in a strictly monotonic progression as the number of agents whose best rule is settled increases in unit steps from 0 to 3: the more settled the system, the more it approaches the game-theoretic solution. In addition, fully 94.2% of the coalition agreements formed by artificial agents, but only 68.7% of those formed by the humans in Kahan and Rapoport's laboratory study, lie closer to the quota solution than to an even split of the coalition's value. In this respect, the artificial agents, perhaps because they are not encumbered by equity considerations or the cognitive and social prominence of even splits, are closer than humans to the game-theoretic solution.

Opportunity costs and best responses. But do these rules approximate optimal responses to the environment? The quota solution can be implemented by a class of (Nash) equilibrium strategies, and thus is consistent with optimal play of the bargaining situation represented as a noncooperative game in extensive (game-tree) form (Selten, 1988, ch. 9; Laing, 1991). At a Nash equilibrium, each agent's strategy is a best response to the strategies being played by the other agents, and thus incurs no opportunity cost.

We operationalize opportunity cost in this situation as follows. Consider any game in the tournament in which a rule obtains the floor in the negotiations. At this stage in the bargaining, the rule can make an offer to either of the other two agents. We compute the rule's *opportunity cost* for this game as the maximum payoff the rule could win by making an offer that would be accepted immediately, minus the actual payoff achieved by the rule in this game's outcome. For example, consider the bargaining sequence displayed in Figure 3 for the game in which B is the first initiator. B achieves a payoff of 12 points in this game's outcome. Yet the strategies shown in Figure 2 reveal that A (respectively, C) would accept an offer from B in which B wins 16 points (respectively, 13). The maximum of these opportunities for B is 16. Thus, in agreeing to an offer of 12 points, B incurs an opportunity cost of 16-12 = 4. Clearly, B's strategy is not a best response to the strategies of the other agents. At a Nash equilibrium, each agent's opportunity cost would be zero.

The opportunity costs incurred in the simulations during the settled periods by the best rule of agents A, B and C averaged approximately 0, 3 and 4 points, respectively. Thus, A employs a strategy that is a best response to the strategies being played by the other agents, but B and C do not. We wonder whether the human subjects in the Kahan and Rapoport study formulated strategies that more closely approximated a Nash equilibrium.

Conclusion

We are encouraged by these results. In a cognitively deprived, information poor environment (e.g., no memory of previous play, lack of full knowledge of the game's conditions) these artificial agents have been able to evolve strategies that reach coalition agreements in the neighborhood of those prescribed by the quota solution and rival those of human subjects. The artificial agents do so with rules that are non-optimal: even the best rules of B and C incur significant opportunity costs. Nonetheless, these agents achieve reasonably effective play in this game. That this has been done in the current context, when a normative solution is known, augers well for contexts in which no unique solution is known to exist. We may hope that computational approaches, GP in particular, will yield insights into games for which solution theory provides no clear answer.

Finally, the performance of these artificial agents should be seen, and must be evaluated, in the context of the large search space they face. Conservatively calculated (Dworman, Kimbrough, Laing, 1995), the search space is at least on the order of 10^{12}. Note further that: (1) These (artificial) players are co-evolving; they are not playing against constant opponents; and (2) In a run of 4000 generations, the maximum number of rules created in the simulation is 6×10^5, yet the search space is at least on the order of 10^{12}. Given this, we can but stand in awe of these simple agents' performance in achieving quite reasonable coalition agreements.

These results encourage us to believe that cognitively limited, adaptive artificial agents can co-evolve to approximate optimal behavior in even more complex and dynamic environments.

References

Abreu, D. and A. Rubinstein (1988). The Structure of Nash Equilibrium in Repeated Games with Finite Automata. *Econometrica* 56: 1259-81.

Albers, W., and J.D. Laing (1991). Prominence, Competition, Learning, and the Generation of Offers in Computer-Aided Experimental Spatial Games. In R. Selten, ed., *Game Equilibrium Models, III: Strategic Bargaining* (Berlin: Springer-Verlag): 141-185.

Andreoni, J. and J.H. Miller (1990). Auctions with Adaptive Artificially Intelligent Agents. University of Wisconsin, Social Systems Research Institute Working Paper 9032.

Aumann, R. (July 25, 1989). Perspectives on Bounded Rationality. Notes of lectures presented at the Workshop on Bounded Rationality, Department of Economics, Stanford University.

Axelrod, R. (1984). *The Evolution of Cooperation* (New York: Basic Books).

Camerer, C. (1990). Behavioral Game Theory. In R.M. Hogarth, ed., *Insights in Decision Making: A Tribute to Hillel J. Einhorn* (Chicago: Univ. of Chicago Press): 311-336.

Cornell, B. and R. Roll. (1981). Strategies for Pairwise Competitions in Markets and Organizations. *Bell Journal of Economics* 12: 210-16.

Dawes, R.M. (1988). *Rational Choice in an Uncertain World* (San Diego, CA: Harcourt Brace Jovanovich).

Dworman, G. (1993). Games Computers Play: Simulating Characteristic Function Game Playing Agents with Classifier Systems. The Wharton School, University of Pennsylvania, OPIM working paper 93-05-05.

Dworman, G. (1994), Games Computers Play: Simulating Characteristic Function Game Playing Agents with Classifier Systems. In *Proceedings of the First IEEE Conference on Evolutionary Computation IEEE World Congress on Computational Intelligence,* Volume 2 (IEEE Press, Los Alamitos, CA): 696-701.

Dworman, G. and E. Schoenberg. (1993). Games Computers Play (Part 2): Applying a Bucket Brigade to a Classifier System Implementation of Characteristic Function Games. The Wharton School, University of Pennsylvania, OPIM working paper 93-01-05.

Dworman, G., S.O. Kimbrough, and J.D. Laing (1995b). Implementation of a Genetic Programming System in a Game-Theoretic Context." The Wharton School, University of Pennsylvania, OPIM working paper 95-01-02.

Dworman, G., S.O. Kimbrough, and J.D. Laing (1995). On Automated Discovery of Models Using Genetic Programming in Game-Theoretic Contexts. In *Proceedings of the 28th Hawaii International Conference on System Sciences, Decision Support and Knowledge-Based Systems*, J.F. Nunamaker, Jr. and R. Sprague, Jr., eds., (IEEE Press, Los Alamitos, CA), pp. 428-438.

Hammerstein, P. and S.E. Reichert (1988). Payoffs and Strategies in Territorial Contests: ESS Analyses of two Ecotypes of the Spider Agelenopsis Aperta. *Evolutionary Ecology* 2: 115-138.

Kollman, K., J. H. Miller, and S.E. Page (1992). Adaptive Behavior and Political Theory. Northwestern University, Department of Political Science, working paper.

Koza, J.R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection* (MIT Press).

Kreps, D., P. Milgrom, J. Roberts, and R. Wilson (1982). Rational Cooperation in the Finitely-Repeated Prisoner's Dilemma. *Journal of Economic Theory* 27: 245-252.

Laing, J.D. (1991). A Noncooperative Model of Bargaining in Simple Spatial Games. In R. Selten, ed., *Game Equilibrium Models, III: Strategic Bargaining.* (Berlin: Springer-Verlag): 80-117.

Marimon, R., E. McGrattan, and T.J. Sargent (1990). Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents. *Journal of Economic Dynamics and Control* 14: 329-373.

Maynard Smith, J. (1982). *Evolution and the Theory of Games* (Cambridge Univ. Press).

Miller, J.H. (1989). The Coevolution of Automata in the Repeated Prisoner's Dilemma. Santa Fe Institute, Economics Research Program working paper 89-003.

Nash, J. (1975). Non-Cooperative Games. *Annals of Mathematics* 54(1951): 286-95.

Newell, A. (1990). *Unified Theories of Cognition* (Cambridge, MA: Harvard University Press, ISBN: 0-674-92099-6).

Newell, Allen and H. A. Simon (1972). *Human Problem Solving* (Englewood Cliffs, NJ: Prentice-Hall).

Selten, R. (1988). *Models of Strategic Rationality* (Dordrecht, Holland: Kluwer).

Shubik, M. (1982). *Game Theory in the Social Sciences: Concepts and Solutions.* (Cambridge, MA: MIT Press).

Tversky, A., and D. Kahneman (1974). Judgment Under Uncertainty: Heuristics and Biases. *Science* 185: 1124-1131.