DOI: http://dx.doi.org/10.7551/978-0-262-31709-2-ch085

Evolving Behaviour-Dependent Strategies in Agent Negotiations

Darius P. Falahat, Enrico H. Gerding and Markus Brede

Electronics and Computer Science,
University of Southampton, SO17 1BJ
dpf1g11@ecs.soton.ac.uk

Abstract

We use genetic algorithms to evolve trading strategies for iterative bilateral negotiations between buyers and sellers. In contrast to previous work we evolve purely reactive strategies that base decisions on memories of behaviour in previous negotiation rounds. We find that simulations lead to three main types of behaviour: (i) cooperative outcomes in which bargaining leads to an agreement and equal sharing of profits, (ii) uncooperative outcomes in which negotiations are not successful and (iii) outcomes in which one party profits at the expense of the other. The frequencies of each type of behaviour vary when the probability for negotiations to terminate is changed, confirming our hypothesis that cooperation should decrease as this break-off probability increases. Comparisons of the results to tit-for-tat (TFT) strategies and previous research on the iterated prisoner's dilemma (IPD) are used to understand simulation results, and we observe the emergence of TFT behaviour during periods of agent cooperation.

Introduction

Trading over the internet and other communication networks continues to grow ever more prevalent in both high-income economies and emerging markets. Research in e-commerce is relevant to high frequency trading, supply chain management and many areas that involve some sort of online transaction. Automated negotiation is central to many of these systems. The application of game theory to automated negotiation is well-established (Binmore & Vulkan, 1999) with practical uses in e-commerce realised early on (Oliver, 1996). The core mechanics of this type of negotiation are retained in a widely used, simple bilateral negotiation model with alternating-offers (Rubenstein, 1994). We use this framework in conjunction with evolutionary computation to make a fresh contribution in a field neglected by the literature in recent years: behaviour-dependent negotiation.

Previous work (Gerding et al., 2004) has focused on using GAs to investigate the emergence of time-dependent negotiation strategies. In the framework of Matos et al. (1998) and Faratin et al. (1997), an agent's strategy can be determined by closeness to a negotiation deadline (time-dependence), the scarcity of a diminishing resource (resource-dependence) and the actions of an agent's opponent (behaviour-dependence). However, regarding behaviour-dependence the authors only consider variations of TFT. The intention of this paper is to address this gap in the literature and develop behaviour-dependent negotiation. Building on previous work in the context of the IPD (Lindgren, 1992) we aim to test for the emergence of cooperative behaviour within a framework of reactive strategies.

Related Research

Automated negotiation has been shown to be vital to e-trading (Sierra et al., 1997) and the use of GAs to identify the most successful negotiation strategies has long been widespread (Matos et al., 1998). GAs use the powerful processes observed in biological evolution: selection of the best-performing (fittest) individuals to replace a population, combined with a small probability of these new individuals undergoing mutations in order to generate diversity and explore the strategy space. GAs have been popular in many fields because they make no assumptions about agent rationality, or the fitness landscape in general. The propagation of agent strategies into new generations is purely based on their fitness.

The complexity and capability of these agent-based models has progressed over time such that, in addition to bargaining over the price of goods (single-issue negotiation), agents can argue over additional properties such as deadlines, cope with multi-issue negotiation (Gerding et al., 2000) and incomplete information (Fatima et al., 2004). This model attempts to take a simple approach that only involves bargaining over a price for the goods; fitness is simply defined in terms of agent utility.

In the context of iterated social dilemma games like the IPD, the situation we consider has been studied extensively by the artificial life community. The prisoner's dilemma is a classic one-shot game, where cooperation rewards higher utility but the rational choice is to defect. In the IPD agents play each other repeatedly, which introduces the potential for more complex behaviour, such as punishment for defection in previous games. Theoretically for a finite number of games, backwards induction implies that the rational choice is to defect each time, although this reasoning does not hold if the number of negotiation rounds played is uncertain and not a priori known to agents. When strategies were tested against each other in an IPD tournament (Axelrod, 1980), the most effective and robust strategy was for an agent to cooperate on the first move and then retaliate with the opponent's previous action. This strategy, TFT, was successful because it rewarded cooperation and punished defection. It should be noted that with the IPD there is a payoff at the end of every negotiation round, whereas in the bargaining model we consider there is only a chance of a payoff at the end of a full negotiation between two agents. Hence results from the IPD can not be translated directly to the simulation results we present in this paper.

Contributions

We have seen that most related research on bilateral negotiation has focused on time-dependent strategies. Our main contribution is to remedy the lack of behaviour-dependent research: the actions of the agents in this model depend entirely on their own past behaviour, and their opponent's past behaviour. The approach we take is similar to Lindgren (1992), where the author models a single population of agents playing the IPD against each other. Lindgren allows the mutations to make the agent strategies more complex, and finds that selection favours cooperative strategies. This exploration of the strategy space via mutations allowed the author to observe extinctions, periods of stasis and other phenomena. Our research also uses a one-population model, but the agents interact via the bilateral negotiation method instead of the IPD. We define cooperation as the sharing of profits equally, meaning exactly equal utility for both agents. Furthermore mutations are only used to randomly select among existing strategies, not to change the size of agent chromosomes. This means the strategy space remains unchanged throughout a simulation, unlike in Lindgren (1992) where the author allows the strategy space to increase.

Despite using a simple negotiation framework, the model produces agent strategies that can be compared to well-known results from the literature (see Results section). The extensibility of the model means the existing framework can be easily built up to further complement the automated negotiation literature from a behaviour-dependent perspective.

Overview

The following section will describe the model in detail, first tackling the negotiation framework and then explaining how the GA works. The section after that reports the results, stating the experimental setup and the parameters used in the simulations. We then present and discuss the observed types of agent behaviour, including analysis of a frequency distribution showing how the different classes of evolved behaviour change over time. The final section summarises the paper including the results, and we suggest several avenues for future research.

Model Description

The model can be understood best by treating it as two distinct components: a negotiation framework and a GA. This section will describe how these components work in detail (see Figure 1 for an overview). The negotiation framework has the bilateral alternating-offers protocol at its core, but the discussion will also include a description of agent strategies and their time-independent nature. The GA is essentially a search algorithm that is applied to the framework; it finds the best performing agents and ensures there is a high probability of them being passed onto the next generation.

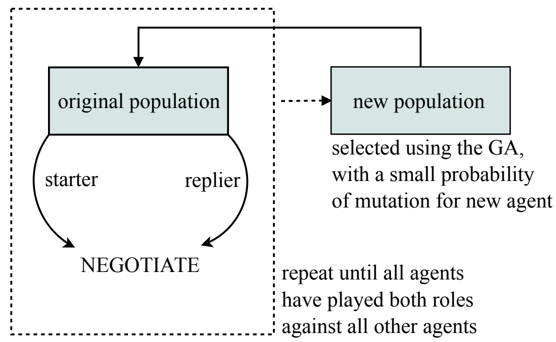

Figure 1: An overview of the interaction between the negotiation framework and GA. One full circuit of the diagram comprises a single *generation*, i.e. the complete interaction and replacement of an agent population.

Negotiation Framework

The bilateral negotiation protocol involves two agents bargaining over goods. The mechanics of how they make offers, counter-offers and contemplate agreement varies considerably over the field of research. In this paper we use a specific protocol with a small strategy space, which keeps analysis relatively simple.

Negotiation between two agents means a Buyer agent and a Seller agent proposing offers and counter-offers to each other, in an attempt to agree on a price for the item. The Buyer initiates proceedings and the Seller replies with a counter-offer. During a negotiation, the Buyer will always begin by offering 0 for the item, while the Seller will make an initial offer of 10. An agent's reserve price is defined as his opponent's initial offer, so the price of the goods stays between the range $[0, 10]$ even though there are no explicit constraints on its value. To simplify matters, the agreed price is converted to utility in the following way. If the Buyer accepts the Seller's offer,

$$U^b = O^t_{s \to b} \tag{1}$$

$$U^s = 10 - O^t_{s \to b} \tag{2}$$

where U is an agent's utility, and $O^t_{s \to b}$ is an offer from the Seller (s) to the Buyer (b), at negotiation round t. Hence utility for both agent types follows the same convention. An inability to come to an agreement results in both agents being punished with zero utility.

There are only two actions, concede and non-concede, available to agents in this model because only two are necessary for bilateral negotiation. Including further actions would be interesting but an area for future work; a spartan philosophy was used for the base model in an attempt to reduce unnecessary complexity. The actions are the same for Buyers and Sellers, but the concede action has different outcomes depending on the agent.

For the Buyer the concede (C) and non-concede (N) actions are as follows,

$$\mathrm{C} : O^{t+1}_{b \to s} = O^t_{b \to s} + 1, \qquad \mathrm{N} : O^{t+1}_{b \to s} = O^t_{b \to s} \tag{3}$$

While for the Seller,

$$\mathrm{C} : O^{t+1}_{s \to b} = O^t_{s \to b} - 1, \qquad \mathrm{N} : O^{t+1}_{s \to b} = O^t_{s \to b} \tag{4}$$

This incremental approach to modifying offers is used instead of other approaches, such as making a complete offer every time, because it restricts the strategy space and limits the complexity of the model.

The alternating-offers protocol is sequential in nature, which means the Buyer will make an offer, followed by a counter-offer from the Seller. The Seller compares the utility it could get by accepting his opponents offer, to his own counter-offer. If the Buyer's offer is more favourable, the Seller will accept. Otherwise, the Seller will decline and the process will continue.The negotiation method M is described in more detail in Eq. (5) from the perspective of a Buyer b who has received an offer $O^t_{s \to b}$ from a Seller s at round t,

$$M^b\left(O^t_{s \to b}\right) = \begin{cases} Quit & \text{if } t > t_{max} \\ Accept & \text{if } U^b\left(O^t_{s \to b}\right) \geq U^b\left(O^{t+1}_{b \to s}\right) \\ O^{t+1}_{s \to b} & \text{otherwise} \end{cases} \tag{5}$$

where t_{max} is the deadline. The notation used to describe the bilateral negotiation protocol is borrowed from (Fatima & Wooldridge, 2002). To lessen the impact of a hard deadline on the negotiation mechanics, the model uses the equivalent method of a break-off probability. This means there is a very small chance of the negotiation ending in each round. Another motivation for using break-off probabilities is that this removes defection via backward induction (see Related Research) as a rational choice, because agents do not know how long a negotiation will last in advance. Using a break-off probability p, we have the expected utility,

$$\langle U \rangle = U\left[1 - (1-p)^T\right] \tag{6}$$

where T is the number of negotiation rounds until agreement. The break-off probability p can be varied: this is likely to have an effect on cooperation if cooperative strategies evolve. Raising the break-off probability limits the time left to negotiate, which would mean fewer strategies can lead to equal sharing, lowering chances for cooperation to develop. We would therefore expect to see less cooperation at higher break-off probabilities, and an increased likelihood of negotiations that end with no agreement at all.

The novel extension of this model involves each agent having a memory consisting of its own and its opponent's previous offers. These memories are central to the strategies that define agent behaviour. An agent's strategy is a mapping of every possible memory to the actions concede (C), or non-concede (N). These actions correspond to those defined in (3) and (4). Every agent's strategy holds the information shown in Table 1 below.

		Initial		Main			
Buyer	Memory	S		C,C	C,N	N,C	N,N
	Action	C or N					
Seller	Memory	C	N	C,C	C,N	N,C	N,N
	Action	C or N					

Table 1: A representation of how agents' strategies are encoded. The order of memories is important: C, N means an agent conceded and his opponent did not. N, C means the opposite.

It should be noted that every agent carries a Buyer and Seller genome, as in Table 1. These genomes are separate: the Buyer genome is used when the agent plays as a Buyer, and similarly for the Seller genome. In a single generation, every agent will play once in both roles. An agent's utility is calculated as the average of the utility from its roles as a Buyer and Seller.

In the first round of a negotiation, the Buyer has no real memories to base an action on. So the special initial S memory is used by the Buyer, only for the first round. In the second round, the Seller does not have a full memory, only a memory of its opponent's move. This means the Seller needs its own initial special case, which depends on what the Buyer did. After the two initial offers, an agent's decision depends on any combination of its own last move, and its opponent's.

Throughout this paper we assume that agents have a one-step memory. This means agents only remember the previous round. A strategy with two-step memories would be much longer: the first two possible memories in the main body would be (CC, CC), (CC, CN), and so on.

Genetic Algorithm

Selection of the fittest agents (those with the highest utility) is achieved using fitness proportionate selection, also called roulette wheel selection in the literature. Unlike some other types of GA, there is still a small probability of choosing less fit agents. This is important because a strategy that is weak against particular strategies may be strong against others, making it slightly more realistic than completely discarding unfit strategies as with truncation selection. See the algorithm below for pseudocode showing where the GA fits in relation to the negotiation framework and the rest of the model.

Algorithm 1 A simplified representation of the one-population model and GA. n_{max} is the agent population and x_{max} is the number of generations for which the simulation continues. Offers, counter-offers and utility calculations are made in the NEGOTIATE subroutine. All agents play each other in both Buyer and Seller roles, but are not allowed to play against themselves.

```
Require: n_max, x_max
  while n < n_max do
    agents ⇐ InitialiseAgents()
    while x < x_max do
      if AgentsPlayed() = False then
        buyer, seller ⇐ PickAgents(agents)
        Negotiate(buyer, seller)
        ResetMemories(buyer, seller)
      end if
      GetUtility(agents)
      agents ⇐ Selection(agents)
    end while
  end while
```

Agents also have a small chance of undergoing single-gene mutation during the selection process, and are given a smaller probability of being completely replaced by a new random agent, to properly explore the strategy space without having to run extremely long simulations. Gene mutations are simply the possibility of an action in an agent's strategy to switch randomly. For example, a Buyer's initial move could change from N to C. If the Seller's two initial moves (see Table 1) are identical, the change will make no difference to the negotiation. On the other hand if they are different, it could change the effective behaviour completely. Since mutations can affect genes that are not used during negotiations, it is possible that mutants could invade populations via neutral drift. This is a possibility because certain types of behaviour use very few genes of an agent's genome (e.g. Table 2 in the Results section), making them potentially vulnerable to drift.

Results

In this section, the different types of agent behaviour are categorised and the change in strategies, and agent utilities, over a simulation are plotted. The stability of the distinct negotiation scenarios are analysed, and some specific strategies are discussed in depth. Finally, the relationship between cooperative behaviour and the break-off probability is plotted and explained.

Experimental Setup

The key parameters for the following simulation results (unless specified otherwise) are given below :

- Population size 100.
- Simulation length of 2000 generations.
- Mutation rate of 10^{-5} for every gene per generation, and 10^{-4} for the mutation of an entire agent. These values were arrived at by slowly decreasing from a large mutation rate, to reach a point where noise from mutations did not dominate the system. For example, there is approximately one gene switch per population every 10 generations, and the introduction of a new randomised agent happens roughly a few times over 100 generations.
- Single time-step memories, i.e. agents only remember the previous negotiation round.
- A break-off probability of $p = 0.005$ (see Eq. (6)) is used for the simulations shown in Figures 2, 3 and 4. This probability gives agents 200 negotiation rounds on average to come to an agreement. Figure 6 uses a larger break-off probability of $p = 0.05$ and the parameter is varied in Figure 5.

Strategy Analysis

Three main negotiation outcomes have been observed in the model. We define these as follows:

- **Domination** is the category of outcome where one agent type has finished a negotiation with a higher utility than the other.
- **Cooperation** is the label for negotiations that finish with both agents walking away with equal utility, if agent utilities are above zero.
- **Zero Utility** outcomes are when both agents finish the negotiation with no utility. This outcome represents a failure to negotiate successfully, so it is treated as distinct to cooperation.

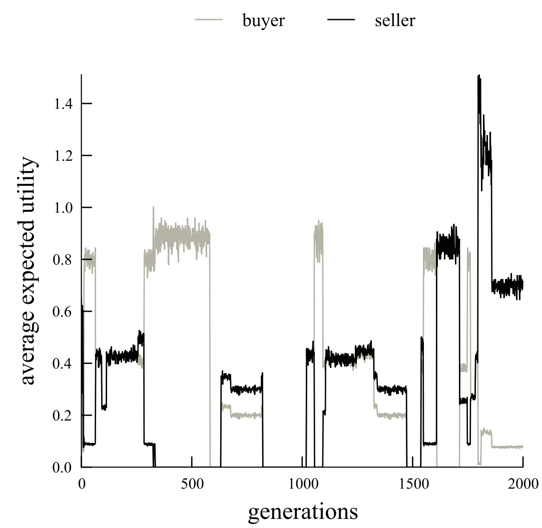

Figure 2: A long simulation run showing the change in agent utilities over time, and the volatility caused by the persistent instability of strategies.

A long simulation using the parameters set out above is shown in Figures 2 and 3. In our analysis of the negotiation outcomes we use two types of plot to clarify the agent dynamics. Figure 2 is a plot showing the average expected utility against the number of generations. The utilities presented in the plot were calculated using Eq. (6) for each agent, with a different utility value for the Buyer and Seller agent types. The expected utilities are then averaged over the population every generation to produce the following plots. Since utilities can range from 0 to 10, the expected utilities will always be less than 10, and often considerably less.

Figure 3 shows the fraction of negotiation outcomes at every generation. Every negotiation is classified as one of the three categories mentioned previously. This type of plot is more sensitive to the strategy dynamics, while the utility plots give a clearer idea of how agents are being selected over the generations.

The most obvious property of Figure 3 is that the simulation never reaches a stationary point, i.e. the strategy fractions never stabilise. In situations where the zero utility (total defection) strategy is in the majority, this lack of stability could be because these agents are scoring no utility, so the population can be invaded by mutants that use any other strategy because no strategy can perform worse. This behaviour highlights a difference to the IPD where defection clearly matters, because agents have the opportunity to punish their opponent straight away.

Invasions of these strategies consist of a build-up of mutants, followed by a swift replacement of the population as soon as utility-scoring mutants start to be selected. Although

		Initial		Main			
Buyer	Memory	S		C, C	C, N	N, C	N, N
	Action	$\boldsymbol{N}$		C	N	C	$\boldsymbol{N}$
Seller	Memory	C	N	C, C	C, N	N, C	N, N
	Action	N	$\boldsymbol{N}$	C	N	C	$\boldsymbol{N}$

Table 2: A strategy table corresponding to the most stable kind of zero utility strategy. Due to becoming locked in a cycle, only the actions in bold are taken by the agents.

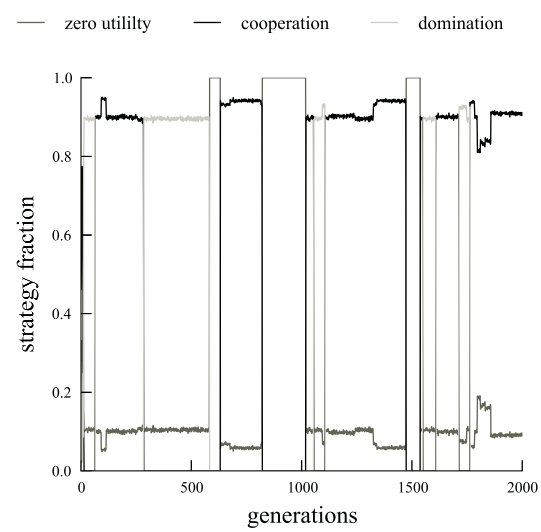

Figure 3: A plot showing the strategies corresponding to the simulation in Figure 2, displaying the changes in agent behaviour.

zero utility strategies are all inherently unstable, some are more resistant to invasion than others: an interesting case is that of a population using the genomes in Table 2.

The most common type of behaviour after the initial stage of the simulation is cooperation. There are many types of strategy that can lead to cooperative behaviour, including all-concede strategies and TFT-like behaviour. Table 3 is an example of the emergent TFT strategies we observed during periods of cooperation.

		Initial		Main			
Buyer	Memory	S		C, C	C, N	N, C	N, N
	Action	C		C	N	C	N
Seller	Memory	C	N	C, C	C, N	N, C	N, N
	Action	C	C	C	N	C	N

Table 3: A strategy table corresponding to cooperative behaviour. Both agents employ TFT-like strategies to arrive at a cooperative outcome.

The Buyer uses a pure TFT strategy, as it initially cooperates and thereafter responds with its opponent's previous move (Axelrod, 1980). The Seller effectively follows the same strategy, resulting in mutual conceding until an equal agreement is reached. For the Seller's strategy to be pure TFT, its initial actions would need to be $C \rightarrow C$ and $N \rightarrow N$. The TFT strategies are not stable, possibly because groups of TFT agents are vulnerable to conceders creeping into the population via neutral drift. The concept is similar to the findings described in Nowak & Sigmund (2005): once non-discriminating conceders invade a population, they are in turn vulnerable to non-conceders (defectors). This has the effect of temporarily replacing a cooperating population with a dominating Buyer or Seller population. A sharp transition of cooperation to domination (and the reverse) can be seen at certain points in Figure 3; the abrupt nature of these transitions may be due to the relatively small size of the agent population.

We have seen from the strategy tables that, during a negotiation, an agent's actions are only determined by a fraction of its entire strategy. By tracking all possible actions available to agents and comparing this to what agents are doing, we can investigate the vulnerability of a majority-conceder population and check for neutral drift. The top plot of Figure 4 shows the fraction of all possible actions in every agent's strategy. The bottom plot shows actual agent behaviour. For the first 300 generations, as cooperative behaviour increases the fraction of concede actions also increases. Shortly after 300 generations have passed, these majority conceders are exploited: this can be seen in the dip of cooperative behaviour in the bottom graph and the corresponding reduction of conceders in the top plot. An indicative example of the neutral drift effect can be seen in the period between 400 and 600 generations, where non-conceding actions increase until they represent over 70% of all actions despite the prevalence of cooperative behaviour.

The analysis so far has focused on simulations that use a low break-off probability. Figure 5 tests our earlier prediction that higher break-off probabilities, corresponding to shorter negotiations, would mean cooperation has a lower chance to evolve. The figure was generated by recording the outcome of each negotiation (cooperation because an agreement was reached, or zero utility because negotiation terminated before agreement) every 100 generations. The fraction of each outcome was calculated at the end of the simulation, and this result was then averaged over 10 simulations to account for variance.

As expected, we find that the fraction of cooperative games with a *small* break-off probability is notably larger than when a larger break-off is used. Furthermore, Figure 5 shows zero utility outcomes are not typical until the break-off probability is increased: the trend for this type of behaviour is roughly inversely proportional to the fraction of cooperative outcomes.

Figure 6 illustrates this point further, showing how the zero utility outcomes take precedence when there is less time

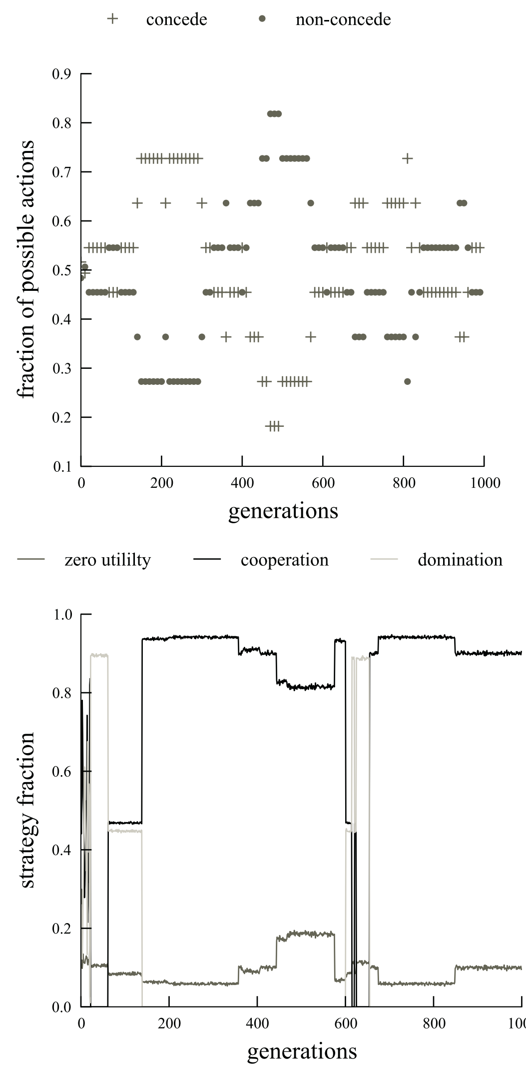

Figure 4: Top: A plot showing the fraction of possible actions for all agents over 1000 generations. Bottom: The behaviour of agents throughout this simulation.

to reach an agreement. The stark difference between Figures 3 and 6 makes sense because cooperative negotiations, which can show transient TFT-like behaviour, are often more complex (since they are discriminatory) and thus need more time to come to an agreement. Non-discriminatory strategies, such as pure defection that results in zero utility outcomes, thrive when there is less time to negotiate.

Concluding Remarks

We built a model using evolutionary algorithms and the bilateral negotiation framework with alternating-offers protocol. Motivated by the literature discussed in the introduction and the relative lack of research on behaviour-dependent negotiation, agents were given behaviour-dependent strategies

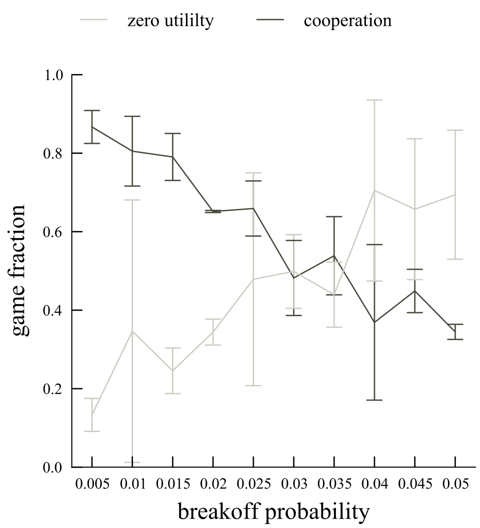

Figure 5: A plot showing the fraction of cooperative and zero utility outcomes as a function of the break-off probability. The error bars give the standard deviation for every averaged point.

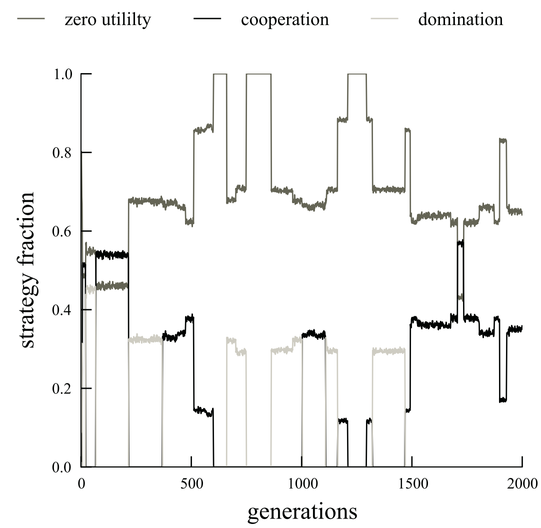

Figure 6: The change of strategy types over time, using a high break-off probability ($p = 0.05$) and otherwise identical parameters to the simulation shown in Figure 3, where the break-off probability was lower ($p = 0.005$).

that base their actions on the behaviour of their opponents. These reactive strategies were evolved and agents selected based on their performance against all other agents in the population.

Three distinct types of negotiation outcomes were observed: cooperation, zero utility and domination outcomes, although the first two types were far more common overall. Simulations proved to be unstable even when left to run for several thousand generations; this is likely due to the neutral drift discussed earlier. Analysis of agent strategies revealed the emergence of TFT-like behaviour during periods of cooperation, although direct comparisons with TFT strategies in the IPD are not possible due to the different way payoffs are handled. In particular, the concept of *cooperation* as discussed in this paper refers to the bilateral agreement on a price for goods; this is not entirely analogous with its meaning in the evolutionary game theory literature. Finally, the assumed relationship of cooperative behaviour with the break-off probability was verified: cooperation is observed more often when negotiations can continue for a longer period of time.

Although this paper was partly motivated by research that used game theoretic concepts like the IPD, our model is not limited to applications within evolutionary game theory. The negotiation framework can essentially be used with any search algorithm in order to select the agents that perform best, not only a GA. The results reported in this paper validate our model as an effective way to investigate the emergence of cooperation in the context of a sequential, bilateral bargaining framework.

Future Work

In this paper we developed a general modelling framework that allows for straightforward extensions in several directions. The first priority in future work would be to verify if it is possible to produce evolutionary stable strategies in our model. One aspect worth more detailed investigation is that of longer agent memories. So far preliminary results from doubling agent memories have mainly shown an extended initial period of noise, before the simulation continues into the familiar patterns of constant strategy invasions discussed above. However, an evolutionary stable solution should not be ruled out and there may be potential in an analytic approach due to the relatively small strategy space.

There are also many ways of increasing the strategy space, such as adding a third action. An interesting possibility is to make it a random choice between the existing two actions. This would allow us to explore if deterministic strategies are more favourable when pitted against unpredictable agents. Explicitly expanding the strategy space can be done by making time a parameter of the agent strategies, making the model both behaviour- and time-dependent. Currently a strategy is a single line of mappings; introducing time would effectively add another dimension, giving a line of possible actions for every round until the deadline.

References

Rubenstein, A. (1994). *A Course in Game Theory*. MIT Press.

Fatima, S. and Wooldridge, M. (2002). Optimal Negotiation Strategies for Agents with Incomplete Information. *Intelligent*

Agents VIII, Springer Verlag.

Binmore, K. and Vulkan, N. (1999). Applying game theory to automated negotiation. *Netnomics*, 1(1):1–9.

Oliver, J. R. (1996). A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of Management Information Systems*, 13(3):83–112.

Sierra, C., Jennings, N. R., Noriega, P. and Parsons, S. (1997). A Framework for Argumentation-Based Negotiation. *Proc. 4th Int. Workshop on Agent Theories, Architectures and Languages (ATAL-97)*, pages 177-192.

Matos, N., Sierra, C., and Jennings, N. R. (1998). Determining successful negotiation strategies: An evolutionary approach. *Proc. 3rd Int. Conference on Multi-Agent Systems (ICMAS-98)*, pages 182–189.

Faratin, P., Sierra, C., and Jennings, N. R (1998). Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems,* 24(3): 159-182.

Gerding, E. H., Somefun, D. J. A. and La Poutré, J. A. (2004). Automated bilateral bargaining about multiple attributes in a one-to-many setting. *Proc 6th Int. Conf. on Electronic Commerce*, pages 105-112.

Fatima, S., Wooldridge, M. and Jennings, N. R. (2004). Optimal negotiation of multiple issues in incomplete information settings. *3rd Int. Conference on Autonomous Agents and Multi-Agent Systems*, pages 1080-1087.

Gerding, E., Bragt, D. van, and Poutre, J. La. (2000). Multi-issue negotiation processes by evolutionary simulation. *Workshop on Complex Behavior in Economics*, pages 1–28.

Axelrod, R. (1980). Effective Choice in the Prisoner's Dilemma. *Journal of Conflict Resolution*, 24(1):3–25.

Lindgren, K. (1992). Evolutionary phenomena in simple dynamics. *Artificial life II*, 10:295–312.

Nowak, M. A., and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature,* 437:1291-1298.